

Freeing the Web from the Browser: Revisiting Open Hypermedia

Joe Savage

Bachelor of Science in Computer Science with Honours
The University of Bath
May 2018

This dissertation may be made available for consultation within the University Library and may be photocopied or lent to other libraries for the purposes of consultation.

Signed:

Freeing the Web from the Browser: Revisiting Open Hypermedia

Submitted by: Joe Savage

COPYRIGHT

Attention is drawn to the fact that copyright of this dissertation rests with its author. The Intellectual Property Rights of the products produced as part of the project belong to the author unless otherwise specified below, in accordance with the University of Bath's policy on intellectual property (see <http://www.bath.ac.uk/ordinances/22.pdf>).

This copy of the dissertation has been supplied on condition that anyone who consults it is understood to recognise that its copyright rests with its author.

Declaration

This dissertation is submitted to the University of Bath in accordance with the requirements of the degree of Bachelor of Science in the Department of Computer Science. No portion of the work in this dissertation has been submitted in support of an application for any other degree or qualification of this or any other university or institution of learning. Except where specifically acknowledged, it is the work of the author.

Signed:

Abstract

The World Wide Web is, without a doubt, the most powerful research tool currently available to man. No longer must researchers comb through endless indices and catalogues to find what they are looking for. With just a few clicks, it is possible to navigate through a vast global space of information comprising documents on virtually any subject. In spite of its tremendous improvements on what came before, however, the Web is somewhat disappointing in comparison to what it could be.

In this dissertation, we explore the potential of so-called *open hypermedia systems* in better supporting contemporary knowledge work. These systems extend the capabilities of users beyond what is possible with the Web in two main areas. Firstly, they empower users to naturally traverse information spaces that span the domains of multiple programs, and secondly, they allow users to explore multiple perspectives on how information is connected.

Identifying a number of potential benefits and issues with systems of this type, we employ an iterative, user-centered approach to propose and evaluate a series of novel open hypermedia design ideas. Specifically, in the areas of backtracking, information organisation, and link display.

In doing so, we find that participants see our designs as coherent and natural to use on some level, representing significant contributions to the state of the art in open hypermedia interface design. Further, our results provide some empirical evidence for the potential value of open hypermedia systems in research activities, suggesting that these systems could dramatically improve intellectual effectiveness in knowledge work.

Contents

Abstract	i
List of Figures	viii
Acknowledgements	ix
1 Introduction	1
2 Literature Review	3
2.1 Motivating Open Hypermedia	3
2.1.1 Origins	3
2.1.2 Adoption	5
2.2 Existing Open Hypermedia Research	6
2.2.1 Microcosm	6
2.2.2 Display of Links & Linkbases	7
2.2.3 Information Organisation	9
2.2.4 Backtracking & History	11
2.3 Knowledge Work	12
2.3.1 Sensemaking	13
2.3.2 Multi-tasking	14
2.4 Summary	15
3 Exploratory Study Pilot	17
3.1 Task Design	17
3.1.1 Task Description	18
3.2 Procedure	19
3.2.1 Participant Recruitment	20
3.2.2 Data Collection	20
3.2.3 Data Analysis	22
3.3 Findings	22
3.4 Summary	23

4	Exploratory Study	25
4.1	Procedure	25
4.2	Participants	26
4.3	Results	26
4.3.1	Getting into the Zone	26
4.3.2	Separating the Wheat from the Chaff	27
4.3.3	Representations of Structure	29
4.3.4	Information Organisation	31
4.3.5	Comparisons to the WWW	34
4.3.6	Discussion of Advanced Hypermedia Features	35
4.4	Critical Analysis	37
4.5	Summary	37
5	Storyboarding	39
5.1	Prototype Design	39
5.1.1	Link & Linkbase Behaviour	40
5.1.2	Spaces & History	46
5.1.3	Window Search	51
5.1.4	Exploration Mode	54
5.2	Summary	57
6	Storyboard Evaluation	59
6.1	Procedure	59
6.1.1	Participant Recruitment	59
6.1.2	Data Collection	60
6.1.3	Data Analysis	60
6.2	Results	61
6.2.1	Link Behaviour: Link Highlighting	61
6.2.2	Link Behaviour: The Marker Bar	62
6.2.3	Link Behaviour: Navigation	63
6.2.4	Link Behaviour: Linkbases	64
6.2.5	Link Behaviour: Doubts	67
6.2.6	Spaces: Fundamentals	68
6.2.7	Spaces: Granularity	68
6.2.8	Spaces: Questions & Suggestions	69
6.2.9	Window Search: Fundamentals	71
6.2.10	Window Search: Suggestions	73
6.2.11	Exploration Mode: Fundamentals	73
6.2.12	Exploration Mode: Resource Suggestion	74
6.2.13	Exploration Mode: Transpointing	75
6.2.14	Exploration Mode: Extensions	76
6.2.15	General Comments	77
6.3	Critical Analysis	78
6.4	Summary	79

7	High Fidelity Prototyping	81
7.1	Prototype Design	81
7.1.1	Form	81
7.1.2	Narrative	82
7.1.3	Details	83
7.2	Evaluation Procedure	86
7.3	Evaluation Results	86
7.3.1	General Sentiment	86
7.3.2	Links	87
7.3.3	The Marker Bar	90
7.3.4	Perspectives	91
7.4	Critical Analysis	93
7.5	Summary	93
8	Conclusion	95
8.1	Contributions	95
8.2	Critical Reflection	96
8.3	Future Work	97
8.3.1	Application Interface Design	97
8.3.2	History	98
8.3.3	Window Management	98
8.3.4	Dynamic Organisation & Exploration	99
8.3.5	Annotation	99
8.3.6	Supporting Collaborative Work	100
8.4	Summary	100
A	Study Data	109
A.1	Exploratory Study Interview Transcripts	109
A.2	Storyboard Interview Transcripts	133
A.3	High Fidelity Prototype Survey Responses	177
A.4	Exploratory Study Informed Consent Form	183
A.5	Storyboard Evaluation Informed Consent Form	184

List of Figures

2.1	A mockup of transpointing from 1972 (Nelson, 1999).	10
2.2	A screenshot from Intermedia (Garrett, Smith and Meyrowitz, 1986). In the lower right, a global resource map is shown.	11
3.1	An excerpt from a hex dump of a Mach-O “Hello, World!” executable.	19
3.2	A screenshot of the resources directory to be given to participants. .	21
5.1	The first page of the link behaviour storyboard.	41
5.2	The second page of the link behaviour storyboard.	42
5.3	A mockup showing source anchor display for links of varying lengths (Weinreich, Obendorf and Lamersdorf, 2001).	43
5.4	A mockup showing destination anchor display via a ‘marker bar’ (Weinreich, Obendorf and Lamersdorf, 2001).	43
5.5	A two pane extract from the link behaviour storyboard showing link display.	44
5.6	A pane from the link behaviour storyboard showing link creation. . .	45
5.7	A three pane extract from the link behaviour storyboard showing how linkbases can be quickly enabled from the menu bar.	46
5.10	A two pane extract from the storyboard showing switching between spaces.	47
5.8	The first page of the spaces storyboard.	48
5.9	The second page of the spaces storyboard.	49
5.11	The window search storyboard.	52
5.12	A two pane extract from the storyboard showing window search behaviour.	53
5.13	The Exploration Mode storyboard.	55
5.14	An extract from the storyboard showing Exploration Mode.	56
5.15	A two pane extract showing transpointing in Exploration Mode. . .	56
7.1	A frame from our final high-fidelity prototype.	82
7.2	A still from the prototype showing the marker bar.	84
7.3	A still from the prototype showing marker bar display with multiple links on the same line.	84

- 7.4 A still from the prototype showing link creation. 85
- 7.5 A still from the prototype showing perspective selection. 85

Acknowledgements

I would like to express my sincere gratitude to my supervisor Dr Leon Watts for the guidance, encouragement, and thoughtful critique he has provided throughout the length of this project. In addition, my grateful thanks are extended to the participants of the studies described in the pages of this dissertation, each of whom kindly gave up several hours of their time in order to make this project possible.

I would also like to thank my friends and family for their continued support, and my parents in particular for their assistance in the arduous task of proofreading. Lastly, my special thanks are extended to both Bret Victor and the late Doug Engelbart, who inspired both this project and my interest in interaction design more generally.

Chapter 1

Introduction

“Any possibility for improving the effective utilization of the intellectual power of society’s problem solvers warrants the most serious consideration. This is because man’s problem-solving capability represents possibly the most important resource possessed by a society.”

– Engelbart (1962, p.131)

Today, technical progress is changing the way we live our lives at an ever increasing pace. Scientists, engineers, designers, and inventors are amongst those that continue to break new scientific and technical ground, the products of which span the corners of the globe. From the fight against cancer to the supercomputer in your pocket, the impact of knowledge workers on today’s society cannot be overstated.

Most critical in supporting this kind of work today is the computer. These complex systems of hardware and software can significantly enhance the capabilities of their users, allowing for the immense amplification of intellectual effectiveness. Steve Jobs famously likened this ability to the remarkable efficiency gains that can be attained in locomotion through the use of a bicycle, declaring the computer as “the equivalent of a bicycle for our minds” (*Memory & Imagination: New Pathways to the Library of Congress*, 1990). Unlike the bicycle, however, the computer systems of today have a long way to go in terms of the efficiency they can extract (Wilson and Papadopoulos, 2004; Victor, 2014).

Particularly central to contemporary computer-aided knowledge work are hypermedia systems such as the World Wide Web (WWW). These systems allow users to create and explore complex interconnected structures of information, facilitating associative interactions that may more closely reflect the nature of the human mind than other forms of media (Borsook and Higginbotham-Wheat, 1992). The incredible value of this kind of functionality in knowledge work is well established (Choo, Detlor and Turnbull, 2000, pp.79–82), and the WWW is, without a doubt, the most powerful research tool currently available to man.

In spite of its tremendous improvements on what came before, though, the WWW is somewhat disappointing in comparison to what it could be. Tim Berners-Lee first proposed the idea for internal use while employed at CERN in 1989 (Berners-Lee, 1989), and proceeded to write the first Web browser with the help of Robert Cailliau in 1990 (Berners-Lee and Cailliau, 1990). Due to these humble origins, the foundations of the system were primarily established on technical grounds, and development was largely focused on technical issues such as the specification of protocols and markup languages (Berners-Lee, 1991). Relatively little attention was put towards how the fundamentals of the system should be tailored to best enhance human capabilities.

More powerful solutions are possible, and may better serve knowledge workers in achieving their goals. The purpose of this dissertation, then, is to explore ways in which the WWW could be different. To propose and evaluate alternatives that may better amplify human intellect. To this end, we re-visit the largely forgotten area of so-called *open hypermedia* research, which pre-dates the Web by some number of years and shows great potential in this regard. In the chapters that follow, we aim to both prove and improve the value of open hypermedia systems in enhancing knowledge work. In more detail, our work is structured as described below.

Chapter 2 opens with a review of the literature central to this project. Through this, we introduce key ideas in open hypermedia and knowledge work, and identify a number of foundational open problems in open hypermedia system design.

Chapter 3, focusing on these open problems, describes the design of a study examining user behaviour in multi-window knowledge work. This aims to ground these problems in practical experiences and inform the design of potential solutions.

Chapter 4 details the results of conducting this study beyond a single pilot run, establishing a greater awareness for how users approach knowledge work using contemporary computing environments.

Chapter 5, reflecting on these insights, describes the design of four potential solutions to open questions in open hypermedia system design.

Chapter 6 details the results of evaluating these design proposals via a user study, discussing the strengths and weaknesses of each at length.

Chapter 7, making use of this feedback, describes the design of a single high-fidelity video prototype, iterating on a few fundamental ideas from previous designs. It then goes on to discuss the results of briefly evaluating this prototype.

Chapter 8 summarises the overall contributions and limitations of our work, and additionally proposes a number of opportunities for future work in this area.

Chapter 2

Literature Review

2.1 Motivating Open Hypermedia

2.1.1 Origins

In 1968, Douglas Engelbart demonstrated a working hypertext system to the public for the first time. To many, this “Mother of All Demos”, as it became known, was revolutionary. Demonstrating a number of the fundamental elements of modern computing for the first time over a mere 90 minute period, Engelbart was described as “dealing lightning with both hands” (Markoff, 2005). The system being demonstrated, NLS, was truly groundbreaking, and would shape the direction of computer science research for many years to come.

Lacking an established software ecosystem within which it could operate, NLS was built from the ground up as a complete, monolithic solution (Grønbæk and Trigg, 1999, p.5). Users were to interact exclusively with system-specific software components, purpose-built to work with NLS. While this was a sensible system design decision at the time, research in the following decades — surrounded by the fast-changing landscape of the digital revolution — identified the potential benefits of structuring hypermedia systems as separate communicating components. In fact, Engelbart himself was among those conducting research on this topic (1990).

Østerbye and Wiil (1996) present a taxonomy to characterise this kind of separation across two dimensions. In the first dimension they consider the separation of data storage from runtime behaviour, and in the second dimension they consider the separation of hypermedia structure (such as links) from resource content. By splitting a hypermedia system into four conceptual components around these axes, systems can be categorised based on how these components interact with one another. In the case that any two components are fused together rather than communicating via some protocol, this indicates a potential loss of flexibility within the system.

Under such a classification, the World Wide Web (WWW) is considered to separate data storage and runtime behaviour via client-server interactions through HTTP, but to make no attempts at separating hypermedia structure from resource content. While the value of the client-server separation within the WWW is clear — underpinning the globally accessible information space of the system today — the loss of flexibility from combining structure and content is often left unconsidered.

Separating Structure and Content

One of the earliest notable systems to make such a separation is Intermedia, a hypermedia system developed a few years prior to the WWW at Brown University (Garrett, Smith and Meyrowitz, 1986). Instead of storing link information inline with resource data, this stores links in ‘web’ databases separate from the content, more frequently referred to in open hypermedia literature as *linkbases*.

By making such a separation, Intermedia allows for multiple *perspectives* on how information is connected. That is, the links displayed on any particular document do not have to come from the document author. Rather, they can originate from one of any number of sources. As such, users are empowered to both craft new relationships between existing documents (without modification!) and to explore relationships created by others.

In his visionary essay *As We May Think* (1945), Bush pondered the benefits of exactly these kinds of capabilities as they related to somewhat similar functionality within his hypothetical “memex”. A relevant excerpt from this follows.

“Wholly new forms of encyclopedias will appear, ready-made with a mesh of associative trails running through them, ready to be dropped into the memex and there amplified. The lawyer has at his touch the associated opinions and decisions of his whole experience, and of the experience of friends and authorities. The patent attorney has on call the millions of issued patents, with familiar trails to every point of his client’s interest. . . .

There is a new profession of trail blazers, those who find delight in the task of establishing useful trails through the enormous mass of the common record. The inheritance from the master becomes, not only his additions to the world’s record, but for his disciples the entire scaffolding by which they were erected.”

(Bush, 1945, p.108)

Although this vision of the future may be somewhat grandiose, it makes clear that the potential of these types of capabilities is immense. Connecting information

together becomes a powerful tool available to all rather than a privilege granted only to content producers. Different interpretations of the relationships between resources can be represented, manipulated, and shared digitally. And with this, users can create and explore varying perspectives on information as they see fit — for instance, to better reflect their goals and motivations in consuming pieces of media or to connect ideas in works of mass media to areas that may have less widespread appeal.

Furthermore, by exposing the interactions between content and structure as an open protocol, any third-party program written to use the protocol can interact with the hypermedia system. While more closed, monolithic hypermedia systems “demand the user disown his or her present computing environment to use the functions of hypertext and hypermedia” (Meyrowitz, 1989, p.107), these so-called *open hypermedia systems* allow hypermedia support to be integrated into the user’s favourite tools.

Additionally, these tools may support a wealth of different experiences and resource types. This is much unlike the WWW today, in which a vast spectrum of desired behaviour — including that to support written works, interactive experiences, social applications, and e-commerce — must be squashed within the confines of only a few first-class resource types and protocols that cannot possibly serve the requirements of all applications perfectly. While the WWW and WWW browsers can in theory support external programs through application-specific URL schemes (W3C, 2015), these are clearly an afterthought. Such schemes have largely un-standardised, application-specific URL formats, and treat non-browser applications as second-class citizens — which cannot, for instance, contribute to a shared chain of history. As such, specialised file formats cannot and do not form a seamless part of the WWW’s global information space today.

By the early ’90s, open hypermedia had become a hot topic in the hypermedia research community. Halasz’s keynote speech at Hypertext ’91, for instance, discussed it as a clear path for future hypermedia systems (1991). This reflected on, amongst other things, Engelbart’s arguments for such systems (1990) and Pearl’s work on Sun’s open hypermedia link service (1989). A seminal paper from Malcolm, Poltrock and Schuler (1991) even justified some aspects of open hypermedia functionality from an industrial perspective. The next generation of hypermedia systems, it seemed, would be open (Davis et al., 1992).

2.1.2 Adoption

Today, it is exceedingly rare to encounter an open hypermedia system. Although the area attracted some continued research throughout the 1990s and early 2000s, the ideas have not achieved any mainstream success and are no longer common topics of research. Atzenbeck et al. (2017) — an exception to this trend — attribute

this general decline in traditional hypertext research to the increasing “omnipresence of the Web”, noting that “the Web never adopted fundamental findings from hypertext research widely” (pp.35–36).

In spite of the WWW’s ubiquity and dominance, though, it seems that open hypermedia systems could provide significant additional value. Particularly, by allowing users to naturally explore seamless information spaces that span the domains of multiple programs, and empowering them to construct and explore these spaces through multiple perspectives. As such, where Atzenbeck et al. use the decline in traditional hypertext research to identify a research gap in advanced hypermedia infrastructure, we wish to highlight the wider research gap surrounding these systems on the whole. Further, we take the view that this research gap may cloud value judgements around open hypermedia systems, and thus may be impeding progress in this area.

Contributing to this research gap are the great many open questions and opportunities surrounding the design of open hypermedia systems (Nürnberg, Leggett and Wiil, 1998). As this project cannot address all such problems, we instead select a handful that seem particularly neglected by previous research or particularly appropriate to the implementation of such systems today. Moreover, in doing so we aim to aid in proving the opportunities and benefits of such systems in supporting knowledge work, and further to encourage future work in this area.

2.2 Existing Open Hypermedia Research

2.2.1 Microcosm

Perhaps the most famous example of an open hypermedia system is Microcosm. Originally developed at the University of Southampton in 1988 (Hall, Davis and Hutchings, 1996), this system was the subject of much development and research, and is one of the few practical, working systems of its type. Its key features include:

- Support for multiple linkbases, much like Intermedia’s ‘webs’
- Fine-grained links to specific locations within resources
- Selection-based link navigation, supporting hidden and overlapping links
- Link navigation to arbitrary closed source applications
- A ‘universal viewer’ able to wrap arbitrary closed source programs to allow them to handle selection-based link invocation
- ‘Generic’ links that link any occurrences of a string in any document to a particular resource
- Link filters that allow for powerful extensions to the system’s core infrastructure (e.g. dynamic links that perform a textual search)

Like a large proportion of work in this area, the research around Microcosm puts great emphasis on the technical implementation of these features, but does not explore their usability in any real depth. As a result, Microcosm — like many other systems of its kind — has a number of fairly serious usability issues, a few of which we will highlight in the proceeding sections.

Further, although many of Microcosm’s capabilities are impressive when compared to the WWW, it is also the case that Microcosm is — to some degree — an artifact of its time in terms of its functionality. Much more can be done with open hypermedia systems today than was demonstrated in these initial forays into the field. A natural evolution of open hypermedia within the hypermedia infrastructure research community, for instance, is the “component-based open hypermedia system” (CB-OHS). These can accommodate multiple types of hypermedia structure in a single system, supporting more than just node-link type connections (Nürnberg, Leggett and Wiil, 1998). More recently, Atzenbeck et al. (2017) put forwards an infrastructure proposal for an extended CB-OHS, allowing for the addition of some degree of machine intelligence to such systems.

Looking further ahead, the ability for file formats to be heterogeneous and specialised within open hypermedia systems may better accommodate semantic structure for machine consumption. In an advanced system, this structure could be leveraged to bring computing environments closer to the visions of Berners-Lee’s “Semantic Web” (Berners-Lee et al., 2001), Apple’s “Knowledge Navigator” (Sculley, 1987), or Victor’s context sensitive software (2006). Prior to any serious exploration of such advanced endeavors, however, there are a number of fundamental usability questions surrounding open hypermedia systems that need to be addressed. Particularly, we have noticed, in the areas of link and linkbase display, information organisation, and backtracking.

2.2.2 Display of Links & Linkbases

In their paper from 1998, Nürnberg, Leggett and Wiil detail and discuss a number of open questions surrounding the design of open hypermedia systems. Although this research is almost 20 years old at the time of writing, many of these questions remain unanswered today due to the research gap in this area. The authors draw particular attention to a lack of consensus surrounding fundamental issues in the “client interface” layer. That is, considering how interfaces to open hypermedia systems can best support users in achieving their goals.

Among the open questions in this area are uncertainties around how links should be best displayed in advanced hypermedia systems, as discussed in more detail by Weinreich, Obendorf and Lamersdorf (2001). One challenge, for instance, concerns the readability of linked text. In any consideration of link display techniques, care must be taken to ensure that linked text is sufficiently readable and does not intro-

duce unnecessary visual noise. In open hypermedia systems in particular, however, the use of linkbases may produce many more links than is typical on the WWW today, amplifying these issues significantly. If a great many links displayed on a single resource are highlighted with bright blue underlined text, for instance, this is likely to reduce general readability and create distracting visual noise.

This topic is discussed in detail by Obendorf and Weinreich (2003). Citing well-established principles of human vision and conducting a study of their own, the authors of this work conclude that “underlined links seem to substantially decrease the reading performance on Web pages” (p.744). In evaluating alternatives, they then provide evidence that translucent coloured overlays may be a significantly better solution, concluding that these “not only avoid the disadvantages of underlined links in our study, they also offer many advantages such as being more flexible”.

Link Discoverability

In *Microcosm*, these readability concerns are addressed to some extent through selection-based links. That is, some links are simply not highlighted inline at all, and instead rely on the user’s explicit selection for navigation. It is clear that such an approach reduces link discoverability tremendously, however, and thus is unlikely to be suitable for all contexts of use. This demonstrates the fundamental conflict between discoverability and subtlety in link display.

This conflict is also central to the second study presented by Obendorf and Weinreich (2003). Following on from their first study, discussed previously, this second study compares the performance of “on-demand” links — which only appear when a dedicated key is pressed — with always-visible, underlined links. In doing so, the results show that the underlined links performed better on the whole than the on-demand links in terms of completion time across a variety of tasks. This indicates that even though always-visible and underlined links can make text more difficult to read, the increased discoverability that they grant may make for a worthwhile trade-off for certain tasks. There is a balance to be struck between discoverability and subtlety, in which the translucent overlays suggested earlier can be considered to occupy some kind of middle ground.

It is worth noting, though, that this study used an average of only 5.9 links per document. As the negative effects of always-visible links on readability may increase with the number of links, higher link densities may cause the balance of what is best to shift. Yet another factor in this calculation is that increased discoverability appears to come at some cost. Obendorf and Weinreich report that users were sometimes too keen to follow always-visible links, navigating away when the information they were looking for was right in front of them. “If a text is to be read completely”, they conclude, “it should not contain always-visible links” (p.743).

Other Issues

If links are to be displayed consistently across many programs, it is also critical that the display of non-textual links is properly thought out. How should link anchors around images, videos, audio, and code be shown to users in a way that is consistent? Furthermore, how should overlapping links and links with multiple endpoints best be handled? One common solution to disambiguate navigation actions around overlapping links is a dialogue box that asks the user to select one endpoint out of a set of potential endpoints to navigate to, but as Weinreich, Obendorf and Lamersdorf (2001) suggest, this may slow down navigation.

In addition, open hypermedia systems present a number of design challenges surrounding linkbases (Weinreich, Obendorf and Lamersdorf, 2001). These are particularly poorly explored, and no standard design solutions exist for linkbase discovery, selection, or creation. In terms of existing systems, Microcosm made use of a configuration dialog to allow users to select linkbases to enable from a given set, but the effectiveness of this solution is unclear. In a similar light, the process for adding new linkbases in Microcosm is described by Weinreich, Obendorf and Lamersdorf as “quite complicated” (p.22). These issues clearly represent serious barriers to the usability of open hypermedia systems.

2.2.3 Information Organisation

Another key question surrounding the design of open hypermedia systems is how such systems can best aid users in organising and finding information. While many Web browsers today employ the lightweight in-program organisation mechanism of browser tabs, these aren't terribly applicable to open hypermedia systems where the set of open resources can span multiple programs. Instead, open hypermedia systems typically structure information by displaying separate resources in separate windows.

A study of tab and window usage among Mozilla Firefox users in 2010 (Dubroy and Balakrishnan) indicates that while multiple windows may be useful for viewing resources at the same time, users may consider them more confusing or less organised than tabs. We also note that in our subjective experiences, navigation between browser tabs using keyboard shortcuts can be far more efficient than navigation between windows in contemporary operating systems. If open hypermedia systems typically display resources in separate windows, then, these concerns clearly present potential issues.

To help alleviate these issues, it may be valuable for open hypermedia systems to aid users in navigating and establishing structure amongst a number of heterogeneous windows. Effectively, extending the operating system's core window management and organisation capabilities to support the efficient utilisation of open hypermedia

systems. One technique for introducing this kind of structure is “transpointing” (Nelson, 1999). This represents the relationships between windows visually, displaying cross-window references as connecting lines that remain anchored to the appropriate sections of the resources even as windows are moved around or scrolled through, as illustrated in Figure 2.1.

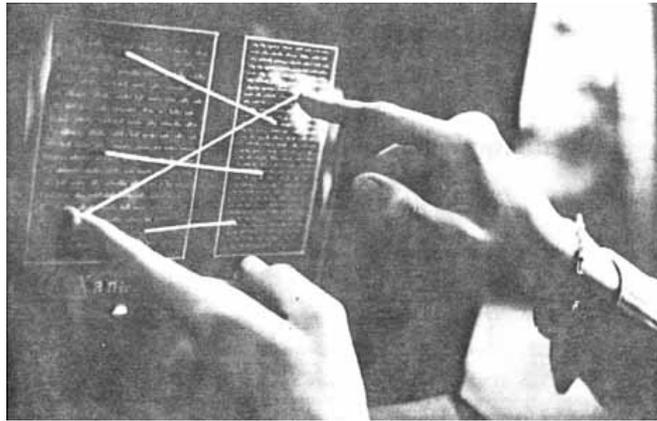


Figure 2.1: A mockup of transpointing from 1972 (Nelson, 1999).

By realising the precise relationships between resources in a visual form, this technique may help users to better ground themselves within an unorganised tangle of open windows. Towards a similar end, another technique to add structure to, and aid navigation through, information spaces are visual resource “maps”. These constructs can be found within Intermedia (Garrett, Smith and Meyrowitz, 1986), for instance, as pictured in Figure 2.2.

In laying out groups of resources and their high level relationships in a compact form, resource maps may help users to navigate through digital resource spaces much like physical maps can aid navigation through physical spaces. While numerous variants of this idea exist, however, these kinds of representations often scale poorly to large or highly connected information spaces (Utting and Yankelovich, 1989). Maps of this type may still prove highly valuable for visualising smaller information spaces, though — for instance, in mapping out individual linkbases, guided tours, or smaller sub-sections of a larger information space.

Yet another area with close ties to these ideas are *spatial hypertext* systems such as VIKI (Marshall, Shipman and Coombs, 1994). These systems allow users to organise resources spatially within virtual environments, for instance on a 2D plane or within a 3D space. This allows relationships between resources to be expressed implicitly through visual cues such as proximity and colour rather than requiring explicit node-link relationships. In doing this, such systems may better allow users

to exploit their spatial memory in information organisation, as has been shown to be highly beneficial (Robertson et al., 1998; Czerwinski et al., 1999).

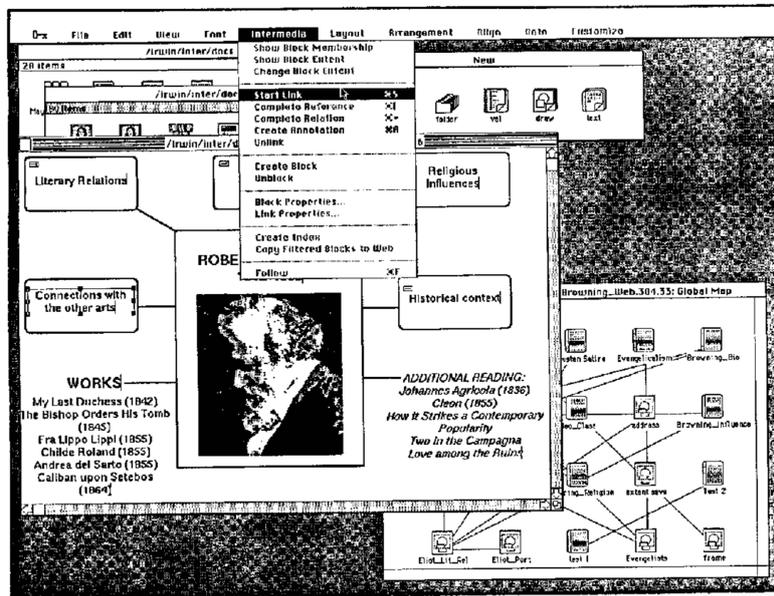


Figure 2.2: A screenshot from Intermedia (Garrett, Smith and Meyrowitz, 1986). In the lower right, a global resource map is shown.

While the core ideas behind many of the kinds of mechanisms described in this section seem to show promise (Chen and Macredie, 2002), many of the specifics of each appear to remain relatively unexplored in terms of both their interface design and technical implementations. There are also, of course, many more ways in which information management and organisation can be enhanced. Perhaps the proposal of new information management and navigation techniques in the context of computing environments today would be a fruitful path for this project to take. Regardless, what is certain is that much more can be done in this area than has been demonstrated in existing open hypermedia systems.

2.2.4 Backtracking & History

Yet another set of open questions surround the behaviour of backtracking and history in open hypermedia systems. Tauscher and Greenberg (1997) identify the “Back” button as the most frequently used history mechanism in traditional Web browsers, accounting for around 30% of all navigation actions in their study. It is important, then, that the revisitation patterns supported by this mechanism are also well-supported by open hypermedia systems.

While this might at first seem to be a trivial piece of functionality to carry over from existing Web browsers, the complexities of multi-program multi-window interactions ensure that this is not the case. In today's Web browsers, browser windows and tabs serve simply as separate spaces for browsing WWW resources. In open hypermedia systems running in similar environments, however, the role of windows is more nuanced. Each window serves as a point of interaction with only one of the many programs participating in the system, typically representing a single resource rather than a general space for hypermedia navigation.

As such, the history mechanisms employed by browsers today — in which users rely on tabs and windows as lightweight constructs for distinguishing between tasks, with separate history records for “Back” behaviour within each — cannot be easily carried over to open hypermedia systems. Instead, a number of practical open systems — such as Microcosm (Hall, Davis and Hutchings, 1996) — simply require users to navigate back through a full, linearized history of all page visits.

This is somewhat comparable to having to work within only a single tab in today's Web browsers, but with the accompanying issues significantly amplified. Not only is it the case that unrelated behaviour originating from the same program may clog up the current chain of history, but *any* unrelated behaviour on the system — including interactions with other programs — has the ability to do so.

This issue is discussed to some degree in the context of open hypermedia systems by Bieber and Wan (1994). In their paper, they introduce the concept of “task-based backtracking” and present a simple algorithm to illustrate how a primitive form of this might be implemented in a multi-window hypermedia system. This research is very much introductory, however, and asks many more questions than it answers. While their approach effectively involves starting a new “task” whenever the user follows a link from a window they switched to manually, a key question is how more sophisticated analysis might improve this.

It is also not clear whether this approach of inferring logically connected groups of windows holds the best solution to this problem. Most notably, inference of the ideal window groupings desired by the user may not be generally possible, and thus such solutions may be more confusing than they are helpful. An alternative solution, more comparable to tabbed browsing today, would be to allow users to explicitly manage and switch between logically related ‘spaces’ of windows. Neither of these approaches appears to be well investigated by existing research.

2.3 Knowledge Work

In its mission to better associate information, hypermedia is inherently well suited to knowledge work. Indeed, much of the foundational work in the field has its roots

in this area. Engelbart's *Augmenting Human Intellect: A Conceptual Framework* (1962) and Bush's *As We May Think* (1945), for instance, are heavily focused around empowering knowledge workers and amplifying human intelligence.

This project is similarly motivated by such goals, and focuses on knowledge work as a primary use case. In the following sections, then, we shall consider literature surrounding knowledge work and cognition that may influence the design of open hypermedia systems, paying particular attention to implications in the areas of link and linkbase display, information organisation, and backtracking, as previously discussed.

2.3.1 Sensemaking

The process of sensemaking is key to knowledge work, and thus to the design of hypermedia systems that aim to support it. While there is a great deal of research in this area — many of the specifics of which are beyond the scope of this write-up — some themes appear to persist throughout the variety of perspectives on the sensemaking process.

Dervin, Foreman-Wernet and Lauterbach's methodology (2003) is a perfect example of this, framing the sensemaking process in terms of a set of "gaps" to be bridged using ideas, thoughts, memories, or emotions. Although this framing is somewhat abstract in nature — and may focus too heavily on the individual rather than the group (Paul and Reddy, 2010) — it provides an excellent basis for thinking about the process. Tackling a difficult problem can be viewed as continually making and un-making sense by bridging "gaps" in understanding.

From this perspective, one of the key goals of hypermedia systems in knowledge work is to aid gap-bridging by making it easier for users to explore the relationships between resources. Through this lens, further direction is presented to some of the open questions previously discussed. Related to the questions surrounding information organisation, for instance, Plumlee and Ware (2006) investigate the advantages of various interfaces to support oceanographers in making sense of detailed maps through visual comparisons. In doing so, they discovered that multiple windows are slower than zooming in and out when comparing a small number of objects in various sets, but faster when comparing more objects.

In a similar vein, Chen and Macredie (2002) review the findings of a number of empirical studies surrounding student learning — and thus, sensemaking — in hypermedia systems. They particularly highlight the issue of disorientation, and the need for supporting navigation patterns with varying levels of guidance for students. Their suggestions to alleviate these issues include visualisations of the information space, resource annotation and highlighting, and adaptive navigation support.

Another model of sensemaking that seems especially relevant to this project is that from Russell et al. (1993), considering sensemaking in the context of understanding a large body of information. This model views the process of sensemaking as the refinement of an internal representation, where the representation continually changes to minimise the amount of information that conflicts with it.

Under this model, hypermedia can be viewed not only as a navigational tool, but also as a means for externalising one's understanding of a body of information into a digital form. Further, by structuring and organising information within a hypermedia system, users may clarify their internal representations and thus improve their understandings. This highlights the potential value of the ability to overlay new link structures over existing content in open hypermedia systems, and suggests that material value may be found in making this process natural to users.

2.3.2 Multi-tasking

Another area pertinent to practical knowledge work — such as that carried out by engineers, scientists, and academics in the real world — is the way in which users manage multiple threads of work.

While traditional HCI research focuses on granular “tasks” — such as writing a document or sending an email — some results suggest that people tend to organise their work in larger thematically connected units. In their empirical study of information workers at an investment management firm, González and Mark (2004) provide evidence for exactly such units, which they call “working spheres”. In doing so, they additionally identify how incredibly fragmented information work can be, noting that “people spend about 12 minutes in a working sphere before they switch to another” (p.1).

In the context of this project, such considerations are clearly important to the design of backtracking and history mechanisms. Specifically, it would appear that “working spheres” roughly correspond to the unit at which we wish to draw backtracking boundaries, and that switching between these backtracking spaces efficiently may be important. Moreover, the idea of “working spheres” may have more general impacts on the project. In the conclusions section of their paper, González and Mark write the following about the implications of their work for information technology design:

“Current information technology is designed to support individual events such as word processing or e-mail use rather than to provide mechanisms to integrate the multiple information objects required by some working spheres. The design of information technology needs to consider how

information workers switch constantly among working spheres. Mechanisms should be flexible to enable people to group particular documents and applications but at the same time there should be recognition that many applications are shared among working spheres (e.g. an e-mail client or scheduling tools).”

(González and Mark, 2004, p.119)

These ideas, surrounding the logical connections between windows, integration of information from multiple sources, and switching between sets of related applications, are all exceptionally relevant to the topic of open hypermedia. In fact, from a certain perspective, this extract can be viewed as specifying a number of important properties that ideal open hypermedia systems should possess to best support knowledge work. It would seem, then, that working spheres may provide a good foundational basis for much of the design of open hypermedia systems.

Interruption

The concept of working spheres, and particularly the act of switching between them, is also closely related to the topic of interruption. While interruptions can be beneficial as well as disruptive (Hudson et al., 2002), they raise numerous questions around efficient task resumption and working sphere restoration that ought to be considered in hypermedia system design.

Borst, Taatgen and van Rijn (2015) provide evidence that the “problem state” of the interrupting and interrupted task may play a large role in this. As such, they suggest it is best to interrupt users at low problem state moments, and may be beneficial to aid state restoration through explicit prompts and cues. These ideas may prove useful in designing usable open hypermedia systems.

2.4 Summary

In considering a number of key ideas and open questions around open hypermedia system design, this chapter has established a foundation in existing literature upon which the rest of this project can stand. Despite the seemingly immense potential of open hypermedia systems in supporting knowledge work and amplifying human intelligence — particularly, by permitting the natural exploration of heterogeneous resources from multiple perspectives — open hypermedia is not a common topic of research today.

One area in which existing open hypermedia research appears to be particularly lacking is interface design. In this area, we have identified specific potential for improvement around backtracking, information organisation, and the display of links and linkbases. These highlight critical usability issues and opportunities, and may

also impact the interest in and adoption of open hypermedia systems.

In order to guide the development of solutions in these areas, we have also highlighted a number of key design considerations from existing research into knowledge work and cognition. Most notably, these encourage further thought around how open hypermedia systems might best support users in switching between logically connected activities and making sense of information.

Chapter 3

Exploratory Study Pilot

Having established a foundation in existing literature upon which this project can operate, this chapter marks the start of our second phase of preliminary research. Focusing on the design issues identified in the previous chapter, this research serves to ground these issues in real experiences and to inform the design of potential solutions. To this end, we describe the design of a primary qualitative study to better understand user behaviour in multi-window knowledge work, and discuss our findings from an initial pilot run.

3.1 Task Design

Central to the design of any study on user behaviour is the task around which behaviour will be observed and analysed. For our purposes, if the results of the study are to be at all applicable to the design of open hypermedia systems, great care must be taken in selecting a task that shares essential characteristics with such systems. Particularly, this task ought to involve the exploration of related, heterogeneous information displayed across multiple windows.

To this end, we initially considered the possibility of having users research a topic of interest using a typical Web browser, and then write an article on this topic using the information they had gathered. This has users freely explore and integrate information in a way that is compelling in many respects, but leaves much to be desired in terms of highlighting issues unique to open hypermedia systems. Most significantly, such a task is likely to encourage participants to work and think within the familiar, entrenched structures of the Web, which may obfuscate more primitive goals and desires that are not easily displayed within this infrastructure.

To avoid this issue, our final study design takes a step back to distance itself from the Web to some degree. Rather than interacting with connected HTML documents, participants are encouraged to use a less sophisticated research medium: normal, local files and programs. The advantages of this approach are multiple.

Firstly, stripping back the research process may allow participants to more naturally desire for changes in behaviour rather than be constrained by the limitations of familiar tools. In avoiding the complexity of, and their familiarity with, the WWW, participants may be more likely to see and discuss areas for possible improvement. Further, the change in environment is likely to have broader impacts on participants' usage patterns. Unlike contemporary Web browsers, for instance, file browsers tend not to have special affinities for particular kinds of files. As a result, navigational biases implicit in existing tools may be reduced, allowing participants to consider the true value of resources without worrying about their lack of integration with existing research tools.

Exploring interactions with different types of resources in this way is, of course, also intimately related to the design challenges that arise in linking together heterogeneous resources within open hypermedia systems. And more broadly, opening resources from a directory rather than browsing through them in a single program is likely to result in many more windows than would otherwise be present. This approximates the environments found in many open hypermedia systems, and thus may yield a better understanding of how users think about these kinds of environments and further aid in addressing the numerous design challenges in this area.

Beyond these changes to the research process itself, having participants write an article to integrate this information was also deemed to be lacking in some regards. While writing a high quality article most certainly requires deep research, writing a low quality article does not. Consequently, there is a risk that participants engaged in this task will take the path of least resistance and will not engage in true knowledge work as is desired. To resolve this, an alternative task has been selected that requires participants to engage in meaningful, intense research to make progress. The details of this task, and how participants will be supported in performing it, are described in the following section.

3.1.1 Task Description

In order to ensure that participants engage in absorbing research that is reflective of real knowledge work, this study assigns a challenging technical task. Participants must understand, to the best of their abilities, the hex dump of a provided executable file. Specifically, a Mach-O binary produced by compiling a simple "Hello, World!" C program with `clang` on macOS 10.13. The first few lines of such a hex dump are displayed in Figure 3.1.

For a full understanding, participants must know a great deal about compilation, linking, loading, and execution. This includes details of Mach-O load commands, object file sections & segments, lazily loaded symbols, position independent code, symbol tables, virtual memory protection, DWARF debug information, and so on.

```

00000000  cf fa ed fe 07 00 00 01  03 00 00 80 02 00 00 00
00000010  0f 00 00 00 00 05 00 00  85 00 20 00 00 00 00 00
00000020  19 00 00 00 48 00 00 00  5f 5f 50 41 47 45 5a 45
00000030  52 4f 00 00 00 00 00 00  00 00 00 00 00 00 00 00
00000040  00 00 00 00 01 00 00 00  00 00 00 00 00 00 00 00
00000050  00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
00000060  00 00 00 00 00 00 00 00  19 00 00 00 d8 01 00 00
00000070  5f 5f 54 45 58 54 00 00  00 00 00 00 00 00 00 00
00000080  00 00 00 00 01 00 00 00  00 10 00 00 00 00 00 00
00000090  00 00 00 00 00 00 00 00  00 10 00 00 00 00 00 00
000000a0  07 00 00 00 05 00 00 00  05 00 00 00 00 00 00 00
000000b0  5f 5f 74 65 78 74 00 00  00 00 00 00 00 00 00 00
000000c0  5f 5f 54 45 58 54 00 00  00 00 00 00 00 00 00 00
000000d0  40 0f 00 00 01 00 00 00  43 00 00 00 00 00 00 00
000000e0  40 0f 00 00 04 00 00 00  00 00 00 00 00 00 00 00
000000f0  00 04 00 80 00 00 00 00  00 00 00 00 00 00 00 00
00000100  5f 5f 73 74 75 62 73 00  00 00 00 00 00 00 00 00

```

Figure 3.1: An excerpt from a hex dump of a Mach-O “Hello, World!” executable.

Participants are not expected to gain such a complete understanding within the time constraints, however, and are instead encouraged to simply understand as much as they can with the time they have.

To aid participants in this task, a directory of resources are provided, as has been motivated in some detail already. This includes PDFs, PowerPoint presentations, spreadsheets, source code, binaries, and text files that may be useful in gaining a better understanding of various aspects of the hex dump. These are the only resources that participants should use in carrying out the task. As more information is provided than could reasonably be read through in the time limit, participants must sift through a great deal of relevant, heterogeneous information to reach a good understanding. It ought to be easy to get lost in all this information, and how participants manage this will be of great interest.

Additionally, participants are to use their own computers to complete the task. This should ensure some level of familiarity with the computing environment, allowing the study to focus on the usage of the machine rather than comprehension of its core mechanics.

3.2 Procedure

Having designed a task for participants to carry out, a procedure was established for a pilot study to validate this approach. This formalisation ensures compliance with the University of Bath’s ethics guidelines and acts as a foundation for the

methodology to be followed in a more complete study of a wider group of participants. A description of this procedure follows.

3.2.1 Participant Recruitment

Participants are to be recruited through informal interactions with students from the University of Bath Computer Science Department. While this sampling process is not particularly rigorous, it ought to be adequate for this work. Crucially, computer science students should have a sufficient foundational knowledge to make some progress with the task. In addition, only students showing interest in the subject area are to be selected, as this indicates a higher likelihood of meaningful engagement with the task.

For this pilot study, only a single participant will be used. Of course, a larger sample size should be used when employing this methodology to a more meaningful version of this study. In general, though, as this is an exploratory qualitative study, fewer participants simply means exploring less of the space of behavioural possibilities. Even a very small number of participants ought to yield results of some value in informing future design decisions and is far from disastrous.

Prior to any data collection, each participant must sign an informed consent form (section A.4). This briefs participants on how much time the session will take and outlines what is expected of them within this time. It also emphasises to participants that the task is intentionally difficult, that their progress is not what is being studied here, and that they are not expected to understand everything within the time constraints. In addition, this informs participants about their rights, of the data that will be collected, and of how this data will be used. Each participant is also informed that they should bring a computer of their choosing to work with in the session.

3.2.2 Data Collection

Each single-participant session is to take around 90 minutes. The first 45 minutes of this should be spent understanding and carrying out the assigned task. Participants are to be briefed in-person of exactly what they are expected to do, after which they are to be given a compressed directory of resources and asked to begin. As indicated previously, this directory contains a great deal of information. Particularly, after carrying out the task ourselves over the course of a few weeks, we collected a total of 31 resources for participants to use. These include, for instance, a dense 71 page Mach-O file format reference PDF, and a text file with the results of running `otool` against the binary to list the Mach-O load commands. A screenshot of the partial contents of the directory follows for illustrative purposes.

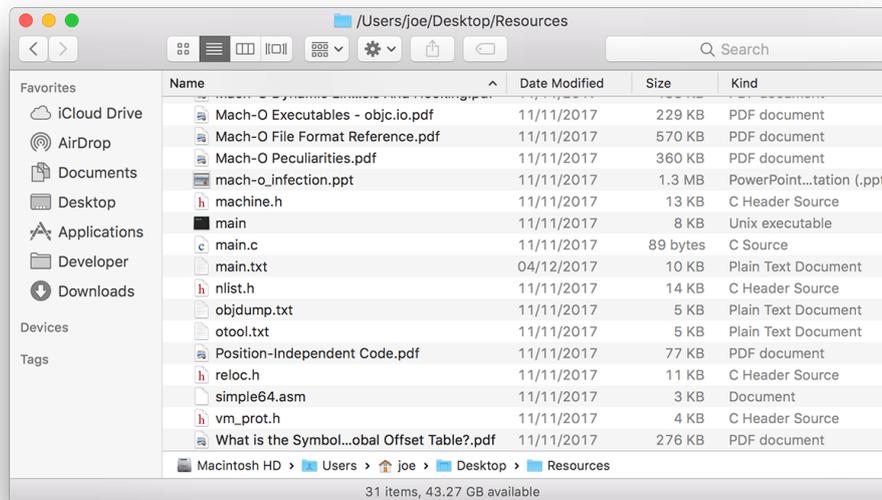


Figure 3.2: A screenshot of the resources directory to be given to participants.

Before participants are to begin the task, a video recording is to be set up to preserve the entire session for future analysis. Then, while the participant is carrying out the task, a researcher should passively observe and take notes regarding the displayed behaviour. Participants are to be informed that they do not need to explain their behaviour to the researcher at this stage, though they are free to speak their thoughts aloud if it feels natural. Additionally, participants should be informed that they can ask questions to the researcher at any stage if they require. This may help to reduce time lost to participants getting stuck in uninteresting quirks such as program malfunction.

Following task performance, the latter 45 minutes of each session is to be spent discussing the experience in a semi-structured interview. This discussion should be guided by the notes taken during task execution and a short pre-prepared script. More generally, the interview should examine both the participant's mental model in completing the task, and any inefficiencies or inadequacies identified in their research process. Particular attention should be paid to the topics of window organisation and backtracking, as suggested by the existing literature.

In conducting the interview, the researcher should not take notes in real time. Rather, it is preferable to rely on the recording for later analysis, as copious note-taking may detract from the quality of conversation. The interviewer should also follow up on any interesting topics raised by participants, even if these fall outside of the expected areas of focus.

Lastly, the researcher should end the interview by asking a number of questions surrounding solutions to problems encountered by participants. Although participants are not expected to have an expert knowledge of the design problems in this space, listening to their thoughts surrounding possible solutions may be valuable to future design work.

3.2.3 Data Analysis

The primary data source of interest from each session is the video recording. For the pilot, only a single participant will be studied, and thus the qualitative data of interest can be extracted from only a single recording. By watching this two or three times and making notes, the general points of interest can be established to be used in assessing the methodology of the study. In more thorough applications of this methodology, multiple participants will be involved, and so multiple recordings must be reviewed. The details of how this analysis should be performed will be described in Chapter 4.

3.3 Findings

Having conducted a pilot run of this study using the procedure described and analysed the results, we can evaluate the efficacy of the study design. The clearest result is that the study seems to be effective at producing data of interest. A number of interesting details were raised in the session, including the following:

- Most of the time was spent by the participant in research, as desired, trying to get a handle on the problem by sifting through resources.
- Particular difficulty was identified by the participant in finding structure amongst the wealth of information, figuring out how different resources were related and what information was most useful at any given time.
- Resources, not programs, were identified as central to the participant's mental model. Programs were regarded merely as a means to access resources.
- Switching between and searching through many resources was highlighted as inefficient. One example of this was in trying to switch to an open resource containing the last encountered occurrence of a specific phrase. With lots of resources open and little recollection of the specific resource's details, finding the relevant information was difficult.
- Internal linking within one of the PDF resources was considered useful, but a lack of history behaviour within this was considered highly frustrating.

One quirk that deviated slightly from the envisioned behaviour was that some file format viewers on the participant's machine displayed multiple documents not in

separate windows, but in a single window with multiple tabs. This occurred with both PDF and plain text files. As such, the number of windows produced was much smaller than imagined. Although the participant did outline this additional indirection as a point of inefficiency, one could also imagine that the additional hierarchy and ordering it provides may make it easier to navigate through resources, and thus reduce the impact of navigational issues that may be present in systems that would produce more windows.

While unexpected, this twist is not seen as an issue warranting changes in study design. Ensuring participants can work within environments familiar to them is considered more important than the number of windows produced. The participant in this pilot still had a great deal to say regarding navigation and getting “lost” in all the resources, and this should be perfectly applicable to the design of systems producing larger numbers of windows. It should be kept in mind when using these results to inform the design of such systems, though, that issues of navigation may be amplified due to the differences between few windows of many tabs and many windows of few tabs.

Another unanticipated detail was that the participant spent a significant amount of time translating bytes from the hex dump into ASCII characters to find strings. As this activity is fairly mechanical in nature and is not particularly interesting to watch, an ASCII column will be added to the hex dump for the main study. The participant also had more trouble finding a starting point than expected. In retrospect, the name of the file format may not have been made clear enough, and so this ought to be further emphasised in the main study.

3.4 Summary

In this chapter, we have described the design of a qualitative study to gain a greater awareness into how users approach computer-aided knowledge work. Trialling this approach through a preliminary pilot run, we have validated the core ideas behind the study and highlighted a few areas that could be improved in future. With a number of minor changes, a more extensive version of this study can be run to collect qualitative data to better guide the direction of this research project.

Chapter 4

Exploratory Study

“It is tempting to describe the essence of hypertext as its ability to perform high-speed, branching transactions on textual chunks. But this is a little like describing the essence of a great meal by listing its ingredients. Perhaps a better description would focus on hypertext as a computer-based medium for thinking and communication.”

– Conklin (1987, p.32)

As with other forms of new media technology such as social media and video games, the utility of hypermedia systems in practical applications can be difficult to fully appreciate without context. In his description of hypertext as a medium for thought rather than a mechanism for navigation, Conklin can be described as looking at such systems through the lens of knowledge work, gaining a better understanding of the utility of these systems in this context. In many ways, such an activity reflects the exact purpose of this chapter: to gain an increased understanding of how hypermedia systems can best support knowledge work.

Building on the study design from the previous chapter, this chapter discusses the results of conducting wider-scale exploratory research into user behaviour in multi-window knowledge work. This aims to establish an increased awareness of how users build understanding and approach research problems in digital environments, and to bring real use cases, desires, and mental models into our design process.

4.1 Procedure

Though this exploratory study will largely follow the same procedure as the pilot from the previous chapter, some details will be modified in accordance with the findings from this. Particularly, an ASCII column will be provided besides the hex dump, and the fact that the hex dump is a Mach-O executable will be further emphasised to participants.

Beyond this, the study will be conducted in much the same way as before. As this study involves a larger number of participants and is being conducted for a different purpose than the pilot, though, the results are to be analysed differently. Rather than extracting findings directly from the video recordings, the interviews are to be transcribed in the context of each full session and the results subject to thematic analysis (Braun and Clarke, 2006).

Thematic analysis as an approach has been selected due to its focus on the discovery of patterns within qualitative data sets. This facilitates the extraction of meaningful themes present throughout the data set, producing general results that can be used both in answering key research questions in this area and informing the design of potential solutions.

4.2 Participants

It was initially planned that around six participants would be recruited to participate in the study. Due to time constraints and unforeseen scheduling issues, however, only four participants were recruited. While this is not ideal, it is far from catastrophic. The study is exploratory and qualitative in nature, and so fewer participants simply results in a less complete picture of potential use cases and perspectives. Though this may bias the results in a less than fully representative direction, they should still prove invaluable in fuelling this project.

4.3 Results

The results of analysing the data follow, with descriptions of each of the themes produced from analysis. As the themes and codes identified are inextricably tied to their interpretations within the surrounding context, discussion around the identified themes is also presented within this section.

4.3.1 Getting into the Zone

A common discussion point across all sessions was the emotional progression in carrying out the task. Participants generally began feeling “overwhelmed” and “disoriented”, but moved into more grounded emotional territory as time went on. Participant 1, for instance, had the following to say about their emotional response:

“At first I was overwhelmed, then I started seeing similar things in files and started to sort of understand a little bit. I started off confused but ended up making something out of it.”

To accompany this emotional progression, data from participant observation suggests a related progression in task progress throughout the 45 minute task performance window. Participants started out making slow progress and seemingly trying to understand the problem and how they might go about solving it. As time moved forwards, however, participants showed an increased familiarity with the environment and began to make improved progress. Participant 4 had the following to say on the subject:

“The first bit is getting used to the problem space and trying to understand what all of this means. But once you’ve got the lingo down, it’s more of trying to find the next piece of the puzzle as opposed to trying to understand what the hell anything is doing.”

In asking participants about the specific barriers that slowed their progress in initially approaching the problem, three areas were consistently identified as contributing factors:

- A lack of familiarity with the subject area
- The volume of information
- The lack of structure in the information

While the first two items in this list are unavoidable in some sense, both are meaningfully related to the final item. Providing some structure to the tangled mess of resources could help to guide users through the vast “sea of information” that is made available and build familiarity more effectively. Discussions around the form that this should take in order to best support knowledge workers in achieving their goals are presented later in this analysis.

4.3.2 Separating the Wheat from the Chaff

Starting Points

Closely connected to this topic of establishing grounding is the subject of finding a starting point. This topic arose in some manner or another in all four sessions. While different participants had varying approaches, the general pattern was that participants looked to gain a high level understanding before moving to finer details, though participants achieved this with varying success.

Participant 1 stated that “finding a starting point wasn’t too bad”, and started analysis from a single PowerPoint presentation, noting that these are “usually pretty easy to get into”. In contrast, Participant 3 chose to look at resources in the order they were listed in the directory as it seemed “as good an approach as any

other to expose myself to the information”. Participants 2 and 4, on the other hand, chose to start with high-level Wikipedia articles that seemed to have relevant titles.

It is not clear whether these wide variety of approaches and varying levels of success indicate fundamental differences between the ways that participants think, or simply reflect the ways in which they chose to approach the problem. It is also likely the case that finding a starting point in a directory of around thirty resources varies significantly from finding a starting point within the many millions of pages on the WWW. It may be the case that the former is easier than the latter due to the hugely constrained number of resources, or indeed that the latter is easier due to the many additional perspectives that can be viewed.

Although participants did not highlight finding a starting point as an area of difficulty, they did seem to call attention to the fact that this could be improved. These details will be discussed further, alongside other potential improvements, in a later section of the analysis.

Internal Structure

From their starting points, all participants aside from Participant 3 — who continued to look at resources sequentially — advanced by skimming resources and trying to determine what would be most useful to look at next. Participant 4 described their experience as initially exploring a wide area to try and “catch a landing place”, adding that “each one sort of told me a new file to open” and that “as my understanding grows, more files become ready to be understood”.

This process of figuring out what documents are currently most valuable and worth looking at next is, of course, closely related to the models of understanding that participants create when carrying out the task. These play a critical role in the sensemaking process that participants follow to gain familiarity with the problem space and integrate information. To explore this further, we asked participants to describe the way in which they went about building understanding.

Participant 1 described their approach as trying to “look at as many [resources] as possible and ... make links between them” and “trying to match each of the different resources to try and form some kind of view of what’s happening”. In a similar vein, Participant 4 described their process as thinking “Well that from there, and that from there, and oh and now I understand this” and “trying near the end to put all of that together”, adding that “where you learn something new, then you have to sort of bring that into what you already know and reinterpret what you already knew”.

Of particular interest here is that participants seem to describe their understanding in terms of a type of internal, constantly changing structure. This is highly related

to the “internal representation” perspective on sensemaking, as described by Russell et al. (1993). In describing this structure, participants tended to use similar language, talking in terms of the relationships between resources. Descriptions of “**connecting** what I read” and making “**links**” between things further support the idea that the associative capabilities of hypermedia systems may closely relate to the way that people think about the relationships between information.

4.3.3 Representations of Structure

Externalisation

While a number of participants described their understandings as a kind of rich informational structure, it is important to remember that these complex structures are contained almost entirely in their heads. This reliance on working memory was identified by many participants as a point of frustration. Participant 4 was particularly vocal about this issue, noting that “if you’re trying to hold it all in your head, it’s a lot harder to remember where everything is”. Regarding external representations of the connections between resources, the participant had the following to say:

“I really had this strong itch of having a whiteboard or a notepad or something to be able to jot things down ... I don’t mean the physical notepad itself, I mean the role that it plays. **Something to coalesce all my thoughts and link things is what I really wanted.** But I had to default to my brain, because at least personally I don’t feel comfortable doing that in a text editor.”

These thoughts are highly related to the Conklin quote at the head of this chapter. The participant desires a medium for thought, but the computer is considered unable to fulfil this role in its current state. After expressing the benefits of a highly flexible medium like paper to fill this need, the participant continues “if I had something on the computer that filled the same sort of niche, that could be useful”.

While a number of participants made digital notes in some form or another, these tended to be fairly simplistic in nature. For the most part, participants used notes for the sole purpose of annotating the hex dump itself — storing details ascertained from their current understandings and associating these with specific byte patterns. Even in the few exceptions to this behaviour where notes were relied on more heavily, they tended to function more as prompts than as methods of offloading any significant portion of the participant’s mental structure. If the distaste for text editors as a medium for thought expressed by Participant 4 is equally felt by other participants, this could go some way to explaining this.

In contrast to this sentiment regarding the text editor, all participants reacted positively to the idea of being able to organise the information, and their thoughts, through cross-program hyperlinks. One participant said this functionality would be “very useful”, while another said it would be “great”.

In discussions around the utility of such functionality, one participant justified its value by saying it “just keeps all the information you need together”. Yet another participant expanded on their thoughts by saying “learning things is creating links”, adding that “in trying to form those cross-references, I’m making my brain work” and that “creating the links is actually part of me solving the problem”. This feedback is highly encouraging, and seems to suggest that separating structure and content in hypermedia systems could be exceptionally valuable to users.

Navigation

Further to discussions of digital structure for the purposes of externalisation, participants also indicated that structure for the purposes of navigation would be desirable. This is hardly surprising given the identification of a lack of structure as obstructing progress, but is nonetheless worth exploring in more detail.

When asked about cross-program hyperlinks for the purposes of navigation rather than externalisation, all participants maintained their positive reactions. Interestingly, however, a number of participants also indicated that while traversing a link structure made by others would be useful, it should not replace the ability to form one’s own links. This emphasises the value placed by participants on being able to link things together themselves, and will be discussed in more detail in a later portion of this analysis.

Participants were much more concise in their descriptions for how digital structure would be useful in navigation. It is assumed that this is likely due to participants being familiar with such functionality from the WWW, and thus seeing it as useful yet familiar and unoriginal. Participant 1 gave the lengthiest description of why such a thing might be valuable, indicating that the connections between documents could prevent them from “faffing around with all the less useful files” and that they “probably would have got further in the task” with such a system.

Beyond this, participants also indicated that some methods of dynamically structuring information — such as by rearranging windows on the fly — could be useful. When asked in broad strokes about the utility of transpointing and resource maps, for instance, all participants indicated that this type of functionality would be useful. In explaining why this was the case, Participant 3 had the following to say:

“A lot of the time I need context. It’s great reading a resource and read-

ing definitions and stuff, but sometimes it's nice to see ... how it all fits together.”

Other participants had similar responses, indicating that such mechanisms could aid the process of making sense of and navigating the information space. In the words of Participant 4, “ways of sort of making it easier to find things are good”. On the whole, then, participants seemed to indicate that some kind of dynamic structuring could certainly be useful — especially for the purposes of organising related windows to aid comparison and navigation.

4.3.4 Information Organisation

Inexplicably tied to discussions around automatic window rearrangement, of course, are the topics of window and information organisation and management. A primary objective in discussions around these topics was to better understand how users think about and manipulate open programs and windows.

One question posed to all participants was whether, in completing the task, they paid more attention to switching between programs or switching between individual resources. Unanimously, participants indicated that resources were the primary considerations in their thought processes, indicating that programs were merely mechanisms to access these. Participant 3 stated that “everything was centred around” resource navigation, while Participant 1 said “[I focused on programs] only when I had problems — other than that, I was just focused on the resources”.

Participant 4 further suggested reaching a state of flow at which programs could become invisible and they could “just focus on what was on the screen as opposed to having to worry about manipulating the programs”. It may be the case that such states are merely transient, however, as all participants indicated some form of inefficiency or frustration in manipulating windows and programs.

Participant 2, for instance, detailed that “I lost my trail of thought sometimes trying to look for windows”. Similarly, Participant 1 mentioned that switching between different programs was “a bit of a pain”. These thoughts support the idea that some form of dynamic structuring of windows to aid navigation may be highly useful to users.

Task Switching

All participants switched between windows using either some form of split view, the Alt-Tab keyboard shortcut, or the task bar. Two participants in particular made use of quick-fire Alt-Tab keypresses to rapidly switch between applications

and compare information. The first of these, Participant 3, explained their use of Alt-Tab in place of split view by saying that it was “just as fast” and that split view could limit the screen space available for each window.

The second in this pair, though, Participant 4, indicated that perhaps split view would have been more efficient, but that there was some perceived switching cost that did not feel worthwhile at the time. They summarised this experience by saying that “it didn’t cause me enough stress or problems or mental work to flick back and forth”, and that while they were “trying to fix the problem rather than think about how I’m doing things” it “didn’t feel worth it” even though this seems “objectively wrong” in retrospect.

This idea of perceived costs factoring in to task switching decisions arose in a few other scenarios too. In discussing the cross-program hyperlinking system, for instance, Participant 2 said that creating links would “would have to be easy enough to do that it wouldn’t be an annoying task in itself”. In the context of switching between resources, Participant 4 also noted that they want navigation “to be as seamless as possible because you want to jump to the other place and continue your train of thought”.

Participant 4 also elaborated on this topic in discussing the value of granular links to specific sections of resources, identifying that such functionality could be significantly helpful in reducing task switching costs. Specifically, they detailed that a fine-grained link to a short section of another document as an “aside” could allow for efficient task switching, indicating that in such a situation they “don’t have to leave what I have in my brain”.

In contrast, the participant said that links to larger sections or to entire documents may result in having to “store what you have” and then “come back”. This clearly relates the topic of problem state restoration, as discussed for instance by Borst, Taatgen and van Rijn (2015).

Tabs

In the findings from our pilot study, the usage of tabbed interfaces by participants was an unexpected result. While undesirable in some respects, however, this quirk provided an excellent opportunity to explore how participants think about tabs, and further to consider how the benefits of tabs might be best replicated in open hypermedia systems.

To this end, we asked participants about their use of tabs in carrying out the task. Participant 3, for instance, found themselves frustrated by the fact that PDF files opened in tabs rather than windows. As this participant relied heavily on using Alt-Tab to switch between resources, they described that this “makes navigation

less efficient” and “meant I had to use the mouse instead of just the keyboard”.

In general, though, most participants seemed not to find tabs especially problematic. The process of navigating to a window and then selecting a specific tab within this may be somewhat indirect, but also seemed to be ingrained into participants’ usage patterns. “That’s probably something I’m just so used to by this point ... that I didn’t notice it as an annoyance”, Participant 1 said.

Participant 4 similarly said they were “conditioned” to this behaviour, but noted that “objectively” “one level of indirection as opposed to two” might be better. Developing these thoughts later in the interview, the participant also gave the insight that “separating things by program and tab is kind of like a two stage directory”, stating that having a single unorganised space to search through would be a “big problem”.

This perspective could go some way to explaining why users may see tabs as more “organised” than windows in certain contexts (Dubroy and Balakrishnan, 2010). The additional level of indirection creates a sort of hierarchy that can make navigation easier. In designing a system that relies on a uniform space of windows, then, care should be taken to ensure that efficient navigation through the space is possible. Curiously, macOS supports this kind of hierarchical access pattern natively through the use of the $\mathfrak{H} \rightarrow$ and $\mathfrak{H} \sim$ keyboard shortcuts for cross-program and intra-program window switching respectively. Whether users know about these shortcuts and how well they solve this problem in practice, though, is not clear.

Novel Organisation Techniques

In discussions around novel information organisation techniques, many participants seemed receptive to new ideas about how windows could be managed and navigated. Though this interest is likely in part due to the demographic of computer science students, it led to some interesting conversations.

In continuing the discussion around tabs with Participant 4, for instance, the participant noted that despite the organisational benefits of tabs, having everything in a uniform space of separate windows could be useful. “I think my connection with them and trying to cross-reference them might be easier [in such a situation]” they said, “splitting screens and putting things side-by-side would be a lot easier”. On the potential organisational issues, they commented that “as long as you can still categorise things and get to things quickly ... it’s not a problem”.

In asking whether search functionality for all open windows could be a useful solution to this organisation problem, the participant responded positively. “As long as I know what I’m looking for then being able to search makes that very easy”, they said. The implication of this is that such a mechanism may in fact be more

efficient for certain types of navigation than the techniques typically used to switch between windows and tabs today.

Participant 2 similarly suggested that some kind of alternative mechanism might be helpful, stating it was “difficult to remember where things were when I wanted them”. Participant 1 had much the same to say, referencing the ease of creating and switching between desktops on macOS as an example of making things “easy”. Such discussions are also highly related to the conversations around dynamic structuring and window rearrangement examined previously.

4.3.5 Comparisons to the WWW

History

In reaching an understanding using ‘normal’ files and programs, participants had no record of history to navigate back through. When asked if this was a constraint in carrying out the task, Participant 1 noted that having closed a file accidentally at some point during the research process, they found themselves unable to find it again when it was desired. “If I was in Chrome when I closed that file I could have just reopened it, which would have been super useful”, they said. Clearly, some kind of history or backtracking functionality could have resolved this issue.

When posed a similar question, however, other participants indicated that history functionality — in the abstract at least — wasn’t especially missed. “The main thing was remembering what files I had read that were relevant,” Participant 2 said, “which is where my notes started coming in handy”. In comparing these thoughts with the functionality of Web browsers today, it would seem that something closer to bookmarks is perhaps what was desired in this case.

Participant 4, meanwhile, indicated that they made up for the lack of history using the ordered positioning of tabs, noting that “the tabs keeping their position kind of made up for the lack of history”. “But the problem with that is that you have this enormous chain of tabs”, they then continued, “as the tab titles become illegible, then I now have to impart effort in trying to find where things are”.

One possible explanation for why participants did not feel especially strongly about the lack of history capabilities is in the specification of the task itself. Though participants were provided with a large directory of resources, the number of files within this was still many millions of times smaller than the number of available documents on the WWW. This difference in scale may result in exploration of a different nature than is common on the WWW, and thus a reduced reliance on backtracking.

Another possible explanation is that, without hyperlinks, participants simply thought

about the problem differently. Without following links, perhaps it does not make natural sense to participants to talk about a “history”. The chain of information being followed is much looser, and thus participants may not expect the computer to track their progress and are happy to substitute for this with their own methods.

Search

The other significant point of comparison with the WWW that took place relates to search. Many participants identified being able to search through information as functionality that was greatly longed for. “Once you learn a term, you want to know as much about it as possible, and rather than using the resources if I’d used a search engine then it would have been a lot quicker”, Participant 1 said.

Participant 2 mirrored this sentiment, indicating that at several points in the task they just wanted to perform a computer-aided search through the available information. This seems to be fuelled in part by a desire for efficient, direct access to information, and in part for exposure to additional perspectives beyond those provided. While the goal of efficient access could in theory be achieved using some kind of local search utility, no participants used such functionality.

Overall, supporting efficient search through both large, remote information spaces and smaller, local ones — such as the set of opened files, as previously discussed — appears to be highly important to participants. Once again, this topic is highly related to the dynamic structuring of information, and fundamentally to how systems for computer-aided knowledge work should allow users to organise information in order to best support them in achieving their goals.

4.3.6 Discussion of Advanced Hypermedia Features

Participant Suggestions

One of the more surprising results from the study was the organic discussion of advanced hypermedia features by participants. From loose hypotheticals surrounding the digital structuring and navigation of information encountered in the task, participants were quick to suggest a wide range of functionality they thought might be useful. Many of these suggestions already have a basis in existing literature and correspond to features discussed in traditional hypermedia research but that are not present with the WWW today.

The natural emergence of these suggestions in the context of the task may go some way to motivating their potential value. In discussing the value of utilising pre-made cross-program hyperlinks for navigation, for instance, Participant 2 commented that “it would be nice to add to something like that, or to even change things if your brain maybe doesn’t work in the same way or you don’t make the

same link”.

Participant 4 also said that “not everyone has the same way of traversing the content”, while Participant 3 noted that one would “have to make sure that the person who did all the gathering and sorting did it right” when making use of such structure. While these viewpoints may have been spurred on by discussing hyperlinks as a tool for externalisation prior to considerations for navigation, they certainly suggest separating structure from content as is done in open hypermedia systems.

In discussions around linking, many participants also indicated that links to specific sections of documents — rather than simply to the start of documents — would be useful. When asked specifically about this topic, all participants responded extremely positively. Participant 4 said that this functionality “speeds up the process” “makes the process a lot more seamless”, and “keeps the mental burden down”. Participant 3 equally said that this would be “much more helpful than linking to the whole thing”, and Participant 2 said it “would help quite a lot”.

Consumption & Creation

The final area to be considered in this analysis concerns participants’ central relationships with interconnected structures. Particularly, the topic of whether participants see themselves more as consumers or producers of informational structure.

As touched upon previously in discussing the use of cross-program hyperlinks for navigation, a number of participants indicated that the ability to follow existing link structures should not replace the ability for them to create links themselves. Participant 2 said consuming pre-made link structures would be “useful to an extent as a starting point, but not as a final point”, adding that “it would be nice to add to something like that, or to even change things”.

Participant 4 mirrored this sentiment, saying that “learning things is creating links between things” and that “creating the links is actually part of me solving the problem”. Expanding on these ideas, the participant noted a separation between tasks where forming relationships is the goal and those where navigating to a particular place to get a direct answer is the goal:

“If I’m trying to break something down, then the act of making those links is actually me doing that decomposition . . . Whereas if I’m just trying to read up on my favourite thing, then I just want it to tell me what’s referring to what. So I think context means a lot there. Because it sort of becomes part of the problem in one area, whereas in the other it becomes part of the solution.”

Context, it would seem, is critical here. In a similar light, participants were also asked what level of hand-holding they would find most useful in the hypermedia structure to support this task — for instance, whether they would prefer a guided tour through resources, like Bush’s “associative trails” (1945), or a more freely explorable space of connections. Unanimously, participants indicated that “it depends”.

In general, guided tours were identified as having great potential value. Participant 2, for instance, suggested that their ability to bring disparate information together in a coherent manner would “definitely” be useful, noting that in performing the task they found that “the information was scattered throughout various articles”. Participants particularly indicated that this ability to follow a tight narrative through the information would be most useful “at the start” in understanding the “fundamentals”.

Many participants, however, expressed a sense of scepticism towards using a guided tour for the entire task. Participant 1 said they would start with a guided tour “with the aim of being able to freely explore later”, while Participant 4 indicated that a guided tour would be helpful as long as they could “jump ship” to perform their own “exploration”. These thoughts are summarised nicely by Participant 1, who commented: “It’s kind of like a city tour, isn’t it? They take you to the main parts, then you can sightsee and explore the little side streets”.

4.4 Critical Analysis

Having presented the results of our thematic analysis, it is important to emphasise the limitations of this work. Recall that the participant pool consisted of four computer science students and that the results are indicative only of their thoughts during uninterrupted 90 minute sessions. As such, the results presented indicate possible perspectives of users, but may not reflect probable perspectives within any wider group.

Due to the qualitative nature of the data and interpretative nature of the analysis, our findings are also subject to some amount of bias. The full transcripts of all four interviews are presented in section A.1, and readers may wish to draw their own conclusions from these rather than to take the presented results at face value.

4.5 Summary

By analysing the results of an exploratory study into multi-window knowledge work, this chapter has established a greater awareness for how users approach such tasks using contemporary computing environments. While our results present an assortment of insights into a number of areas, the high level aggregation of these

provides strong evidence for the utility of open hypermedia systems in this context. Participants in our study considered link structures to be highly powerful tools for both externalisation and navigation, suggesting that open hypermedia functionality could dramatically enhance the role of the computer as a medium for thought.

Far from focusing on operational minutiae, our findings largely concern the fundamental tasks of creating and navigating through large information spaces. In considering the barriers that users may face when working towards these goals within multi-window environments, our results particularly highlight the practical issues of organising and finding information within large spaces of windows. This area sits at the heart of open hypermedia interface design, and seemingly at the core of user considerations, yet as identified in our initial literature review is still very much unsolved. Our hope is that in this — and other — areas, the ideas and insights attained from participants in this chapter can be used to guide the design of better open hypermedia systems. Especially, systems which allow the user to traverse and manipulate information spaces more naturally and efficiently.

Chapter 5

Storyboarding

Reflecting on the insights granted by the last three chapters, the purpose of this chapter is to propose a few potential solutions to some of the significant unsolved problems in open hypermedia system design. Focusing on a handful of issues highlighted by the exploratory study and relevant literature, we outline the construction of four design proposals, communicating these in their envisioned contexts of use through low fidelity storyboards.

5.1 Prototype Design

Prior to any design work being carried out, it was first necessary to establish the problems to which this work would aim to provide potential solutions. Rather than producing some kind of formal specification or framework to encapsulate this, we took a more holistic approach. Thinking about solutions in the context of the entire last three chapters, we have chosen to work on ideas that we believe can maximally improve end user experience across a number of areas. As a result, many of the solutions are highly related to a large swath of issues and topics that have arisen in previous chapters.

Additionally, it is the case that instead of designing a number of alternative solutions to a single problem, we have chosen to work on solutions to largely orthogonal problems. This reflects our desire to stimulate discussion around a wide variety of open questions, and to present the potential solutions that we find most compelling across the space of identified problems.

To communicate the key ideas of each solution along with an imagined use case, rough storyboards have been produced. These provide quick, simple, and malleable bases around which discussion and further design work can take place. The motivation behind this is that by evaluating these primitive representations at an early stage in the design process, little time will be wasted on developing the specific details of immature or flawed ideas. Instead, the fundamental properties of the

solutions can be tuned as seems to be most valuable to users, and these ideas can then be iterated upon.

5.1.1 Link & Linkbase Behaviour

The first of these prototypes concerns an unavoidable area of discussion: the display and behaviour of links and linkbases. In the design of any hypermedia system, the way in which users interact with links is surely of paramount importance. Yet, as highlighted in our literature review, existing research in this area is patchy at best. Even the most established interface issues lack consensus (Nürnberg, Leggett and Wiil, 1998), and a large number of other design issues — including linkbase selection (Weinreich, Obendorf and Lamersdorf, 2001) — remain relatively unexplored.

A full discussion of the design issues in this area could easily fill an entire dissertation. In the interests of sparking discussion around a selection of topics, then, we select only a handful of problems to work on — sketching out just one possible solution to how these could be addressed. Particularly, we propose a solution which describes how links might be displayed, navigated, and created, and how linkbases might be enabled. This draws heavily on existing work, bringing together a number of ideas from the literature in a contemporary context. The full storyboard for this solution can be found in Figure 5.1 and Figure 5.2.

Link Display & Navigation

Central to this storyboard is the way in which open hypermedia systems present links to users. In approaching this rather substantial problem, our work draws most significantly from the work of Weinreich, Obendorf and Lamersdorf (2001). Presenting a whistle-stop tour through some of the key challenges and existing solutions in this area, this research presents a number of questions of direct importance to the design of this prototype. Among these are a number of challenges surrounding the readability of linked text, as touched upon in our literature review.

In discussing readability concerns around anchors “longer than a few words”, the authors draw attention to the potential value in using different strategies to highlight link anchors of varying lengths. Short phrases and long passages are typically marked up differently to one another on paper, so why not follow a similar approach digitally? In their mockup of how such a thing might be realised, shown in Figure 5.3, the authors suggest using coloured overlays to denote short links, and narrow bars besides the document to denote longer links. In theory, this balances readability and discoverability for links of different lengths.

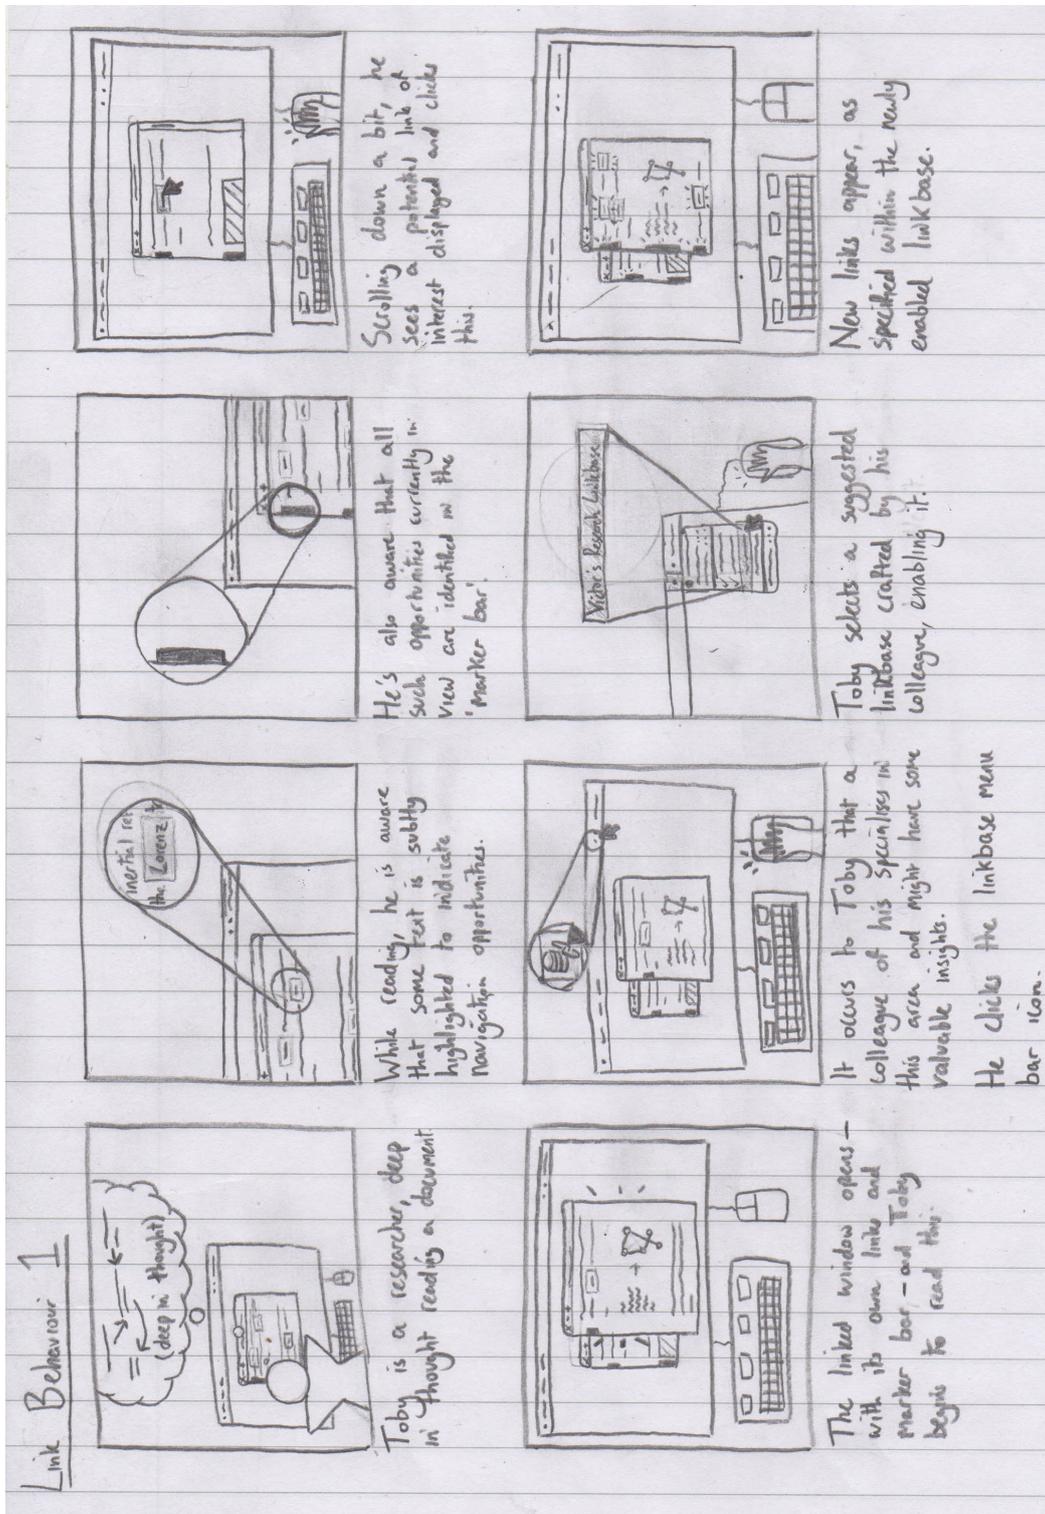


Figure 5.1: The first page of the link behaviour storyboard.

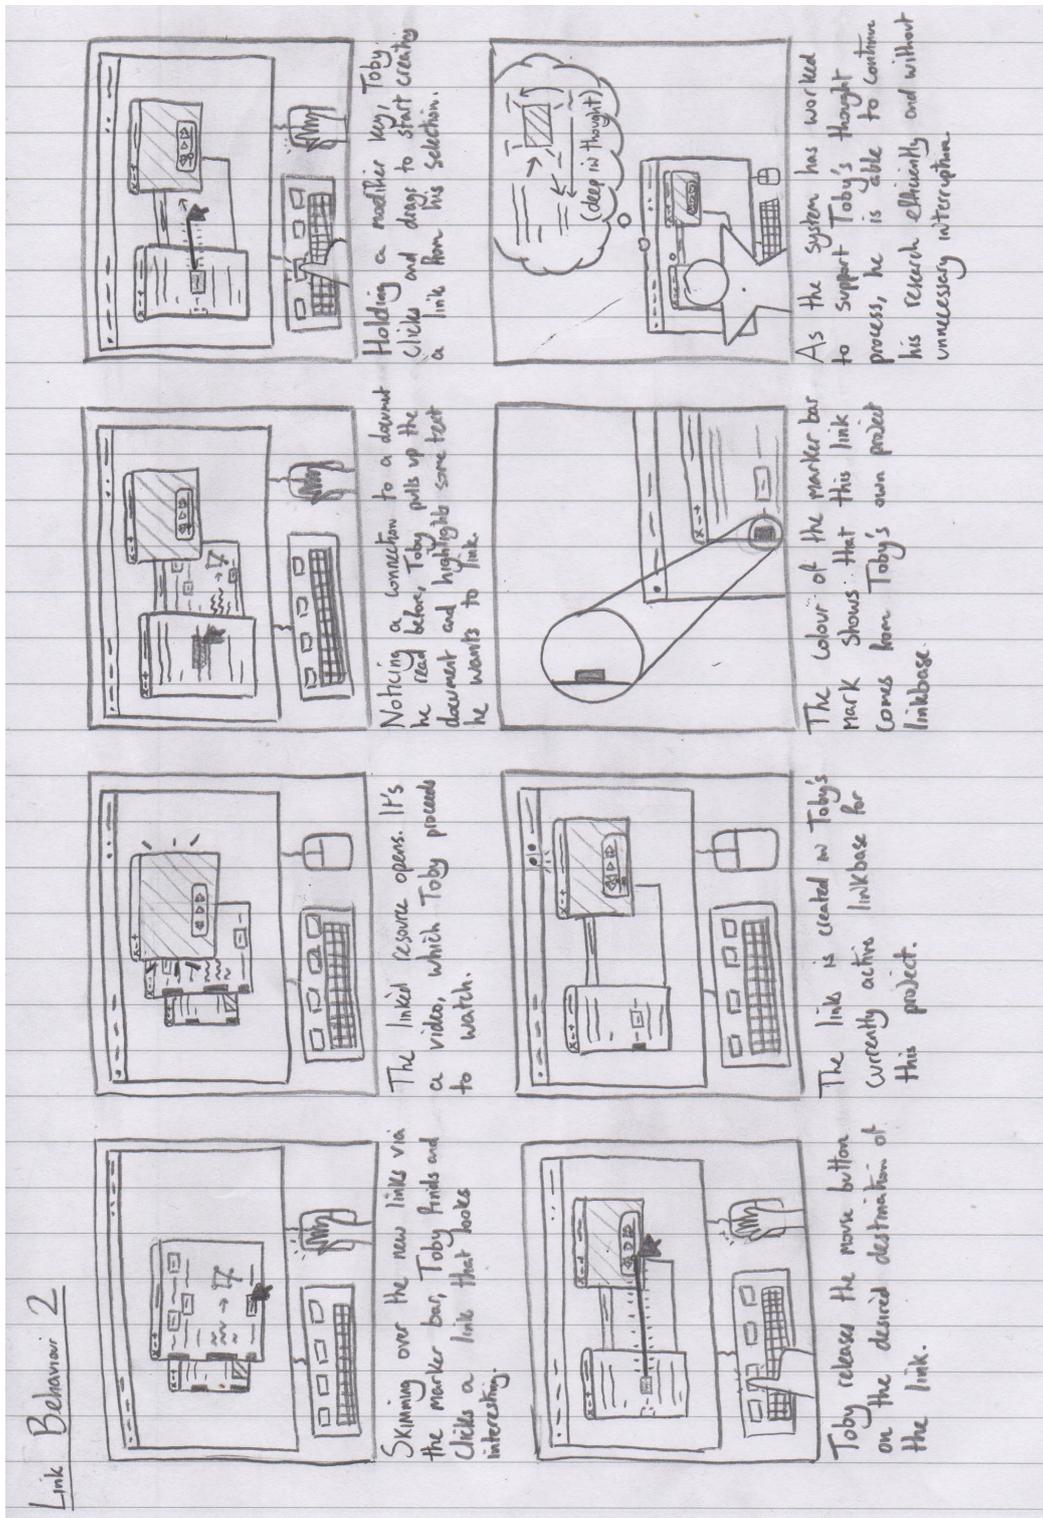


Figure 5.2: The second page of the link behaviour storyboard.

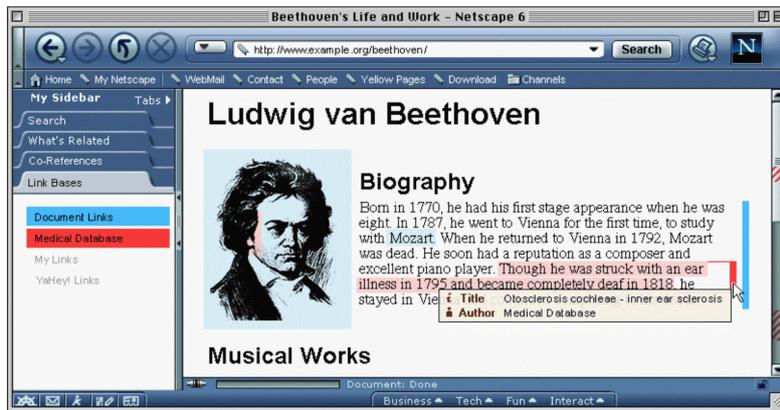


Figure 5.3: A mockup showing source anchor display for links of varying lengths (Weinreich, Obendorf and Lamersdorf, 2001).

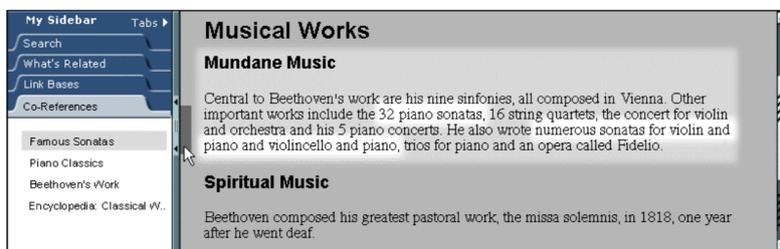


Figure 5.4: A mockup showing destination anchor display via a ‘marker bar’ (Weinreich, Obendorf and Lamersdorf, 2001).

More specifically, their mockups place these ‘long link’ marks on the right hand side of the window. This helps to distinguish them from a “marker bar” interface element they propose for the left of the window to highlight destination anchors, shown in Figure 5.4. Though we consider the central principles underlying these two ideas to be sound, we suspect that as proposed they may introduce too much visual noise — especially when a large number of links are present. As such, our proposal attempts to unify these two ideas.

To complement the scroll bar, we too place a ‘marker bar’ of sorts on the left of the window, operating in the spatial domain of the current view. Just as the scroll bar can be used to highlight all sorts of information across an entire document, our marker bar is intended to be a general purpose tool to highlight information within the current view. Particularly, in this context, to show hypermedia interaction opportunities — including both source and destination anchors. For consistency and predictability, we further propose that marks for both long and short links are displayed in this bar. In this form, the marker bar provides a baseline for the indication of link presence which can then be strengthened through other methods.

To accompany marker bar information, we propose that some subset of short links are subtly highlighted inline, as illustrated in Figure 5.5. This is intended to provide increased link discoverability without seriously impairing readability or producing too much visual noise, even with very many links. In general, the direction of this approach to inline highlighting is strongly inspired by Obendorf and Weinreich (2003), who provide evidence that translucent overlays are significantly more readable than underlined text, as discussed in our literature review. The details of how this set of highlighted links are determined, and other issues such as how the system should display links with overlapping anchors, are not considered within this prototype.

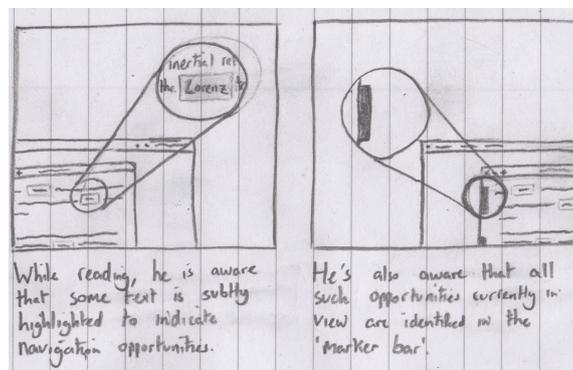


Figure 5.5: A two pane extract from the link behaviour storyboard showing link display.

We also consider the specifics of how information is displayed in the marker bar to be largely out of scope, though we see potential in using the colour, size, and shape of marks for this purpose. Our storyboard indicates, for instance, that the colour of marks might be used to distinguish between links from different linkbases. We also imagine that users will be able to interact with the marker bar in various ways — for instance, being able to hover over marks to view precise visualisations of their anchors, as suggested by Weinreich, Obendorf and Lamersdorf (2001). These ideas have not been explored within the prototype at this stage, however.

In terms of navigation, linked resources in our proposed solution open in a new window by default, as is typical in open hypermedia systems. This seems like a relatively natural way of handling link navigation in this kind of system, and thus we saw no need to deviate from it. While we suspect that a mechanism of this sort may generate too many windows in certain circumstances, we postpone the discussion of such issues to our evaluation of this storyboard in the next chapter.

Linkbases & Link Creation

The storyboard produced for this solution also touches upon methods for fine-grained link creation. As participants in the exploratory study suggested that creating links could play a key role in their process of building and externalising understanding, making this process as natural as possible to users is clearly desirable. To this end, we outline a method of link creation in which users can make connections by simply clicking and dragging from window to window. To make this as concrete and tangible as possible, we visualise the new relationship with a connecting line — an idea inspired by transpointing (Nelson, 1999). A pane within the storyboard illustrating this is shown in Figure 5.6.

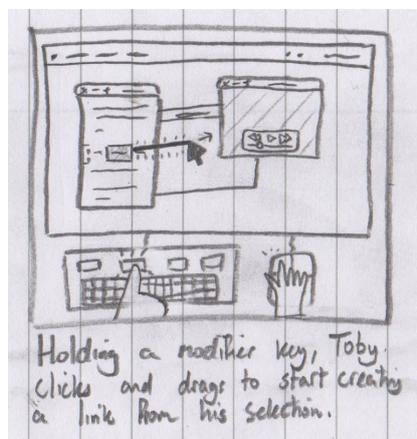


Figure 5.6: A pane from the link behaviour storyboard showing link creation.

In real systems, this procedure for link creation is expected to be accompanied by an alternative that decouples the specification of source and destination anchors, allowing users to create links to resources that are not currently visible. In addition, such systems should have some mechanism by which links can be deleted, though no such mechanisms are described in our storyboard. Further, our proposed solution also neglects to define which linkbase new links are created within. We imagine that there might be a notion of the “current” linkbase that is selected for editing, and that links are created within this, but this is open to discussion.

On the topic of linkbases, the storyboard for this solution additionally presents a simple mechanism by which linkbases can be enabled. As the set of enabled linkbases is a type of global state for the system, it makes intuitive sense for this functionality to be separate from any specific viewer program. At the same time, though, we desire for this process to be intuitive and relatively easy to carry out. As such, we suggest the use of a menu bar application for this purpose. This allows some linkbase operations to be accomplished quickly from the menu bar, while also

providing easy access to a dedicated configuration window through which more complex linkbase operations can be carried out. Some panes from the storyboard that illustrate this are displayed in Figure 5.7.



Figure 5.7: A three pane extract from the link behaviour storyboard showing how linkbases can be quickly enabled from the menu bar.

The critical issues of linkbase discovery and distribution are not addressed here. Additionally, other potential enhancements to this functionality, such as how linkbases are (or are not) related to annotations, and how collaborative linkbase editing might work, are not considered here either. Further, the entirety of this prototype completely excludes transclusion — that is, the embedding of one resource within another. This is not because transclusion is deemed unimportant, but rather that it introduces complexity that would be better dealt with in a more comprehensive consideration of design challenges and solutions in this area.

5.1.2 Spaces & History

Another area clearly identified as inadequate in existing open hypermedia systems is backtracking. As discussed in our literature review, many existing systems constrain users to chronological backtracking through a single, linearized history of all page visits. While backtracking strategies involving the inference of logically connected sets of windows may improve this to some degree, their usability characteristics in practice are not at all clear (Bieber and Wan, 1994).

Taking inspiration from the tabbed interfaces commonplace within Web browsers today, we propose a method by which separate history chains can be explicitly managed by the user. To accomplish this, we lean on the established concept of

multiple desktops. By imbuing each desktop with backtracking state, our solution permits users to switch between distinct sets of related windows with independent history chains. The full storyboard for this solution can be found in Figure 5.8 and Figure 5.9.

As discussed in our literature review, knowledge work can be incredibly fragmented, involving constant switching between multiple working spheres (González and Mark, 2004). As such, it is critical that switching between and manipulating our history-imbued multiple desktops — which we call *spaces* — feels natural and efficient to users. This is additionally emphasised by discussions from our exploratory study, in which many participants noted that the perceived cost of task switching is an important factor in their decision making, and that large switching costs may cause them to lose their place. To best combat these issues, our storyboard suggests the use of various keyboard shortcuts. With these, we believe that spaces can be just as lightweight and efficient as tabs, while also being more flexible and powerful.

Tabs in Web browsers today also seem to possess some desirable spatial characteristics. In our exploratory study, for instance, one participant indicated that “the tabs keeping their position” was highly valuable. As briefly mentioned in our literature review, this kind of utilisation of spatial memory has been shown to be beneficial in information organisation (Czerwinski et al., 1999). By grounding each of our spaces within some virtual location, we believe that many of the spatial advantages of tabs can be replicated. The interface design surrounding multiple desktops in macOS (Apple Inc., 2016) seems to be a good example of this — each desktop is associated with a position in space, and switching between desktops moves through this space. We sketch out something similar in our storyboard, as shown in Figure 5.10.

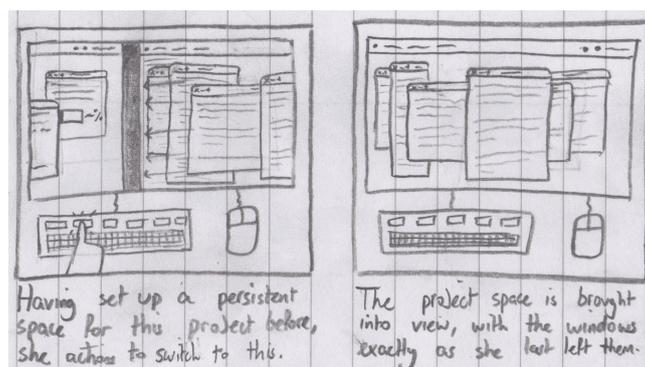


Figure 5.10: A two pane extract from the storyboard showing switching between spaces.

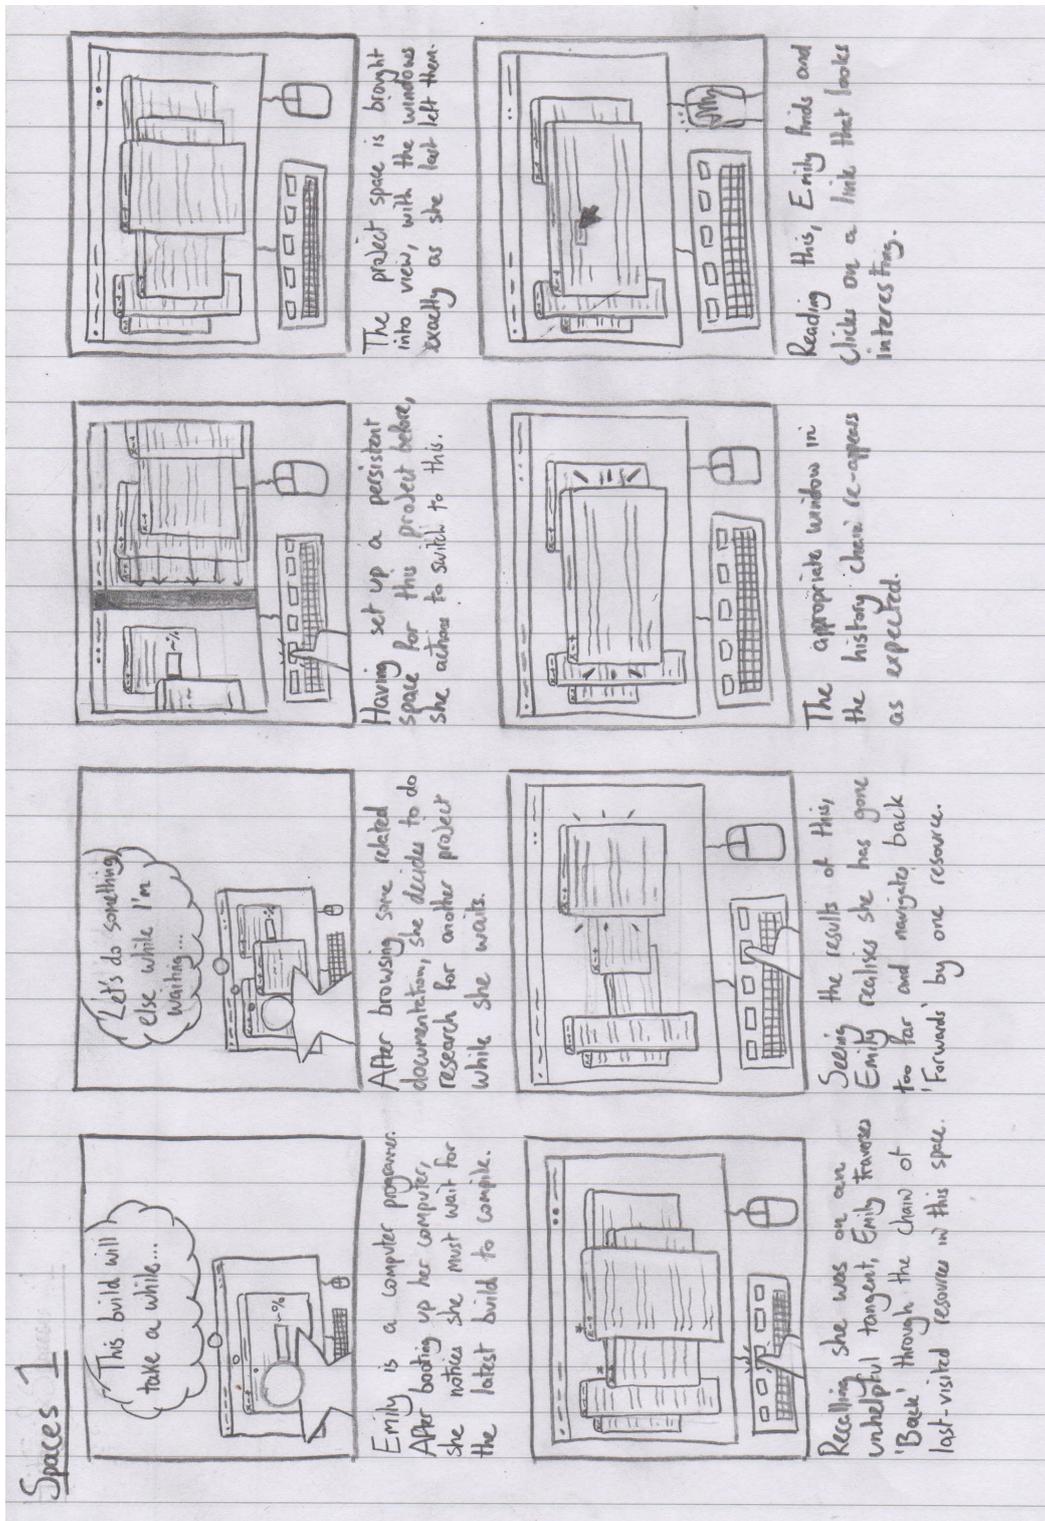


Figure 5.8: The first page of the spaces storyboard.

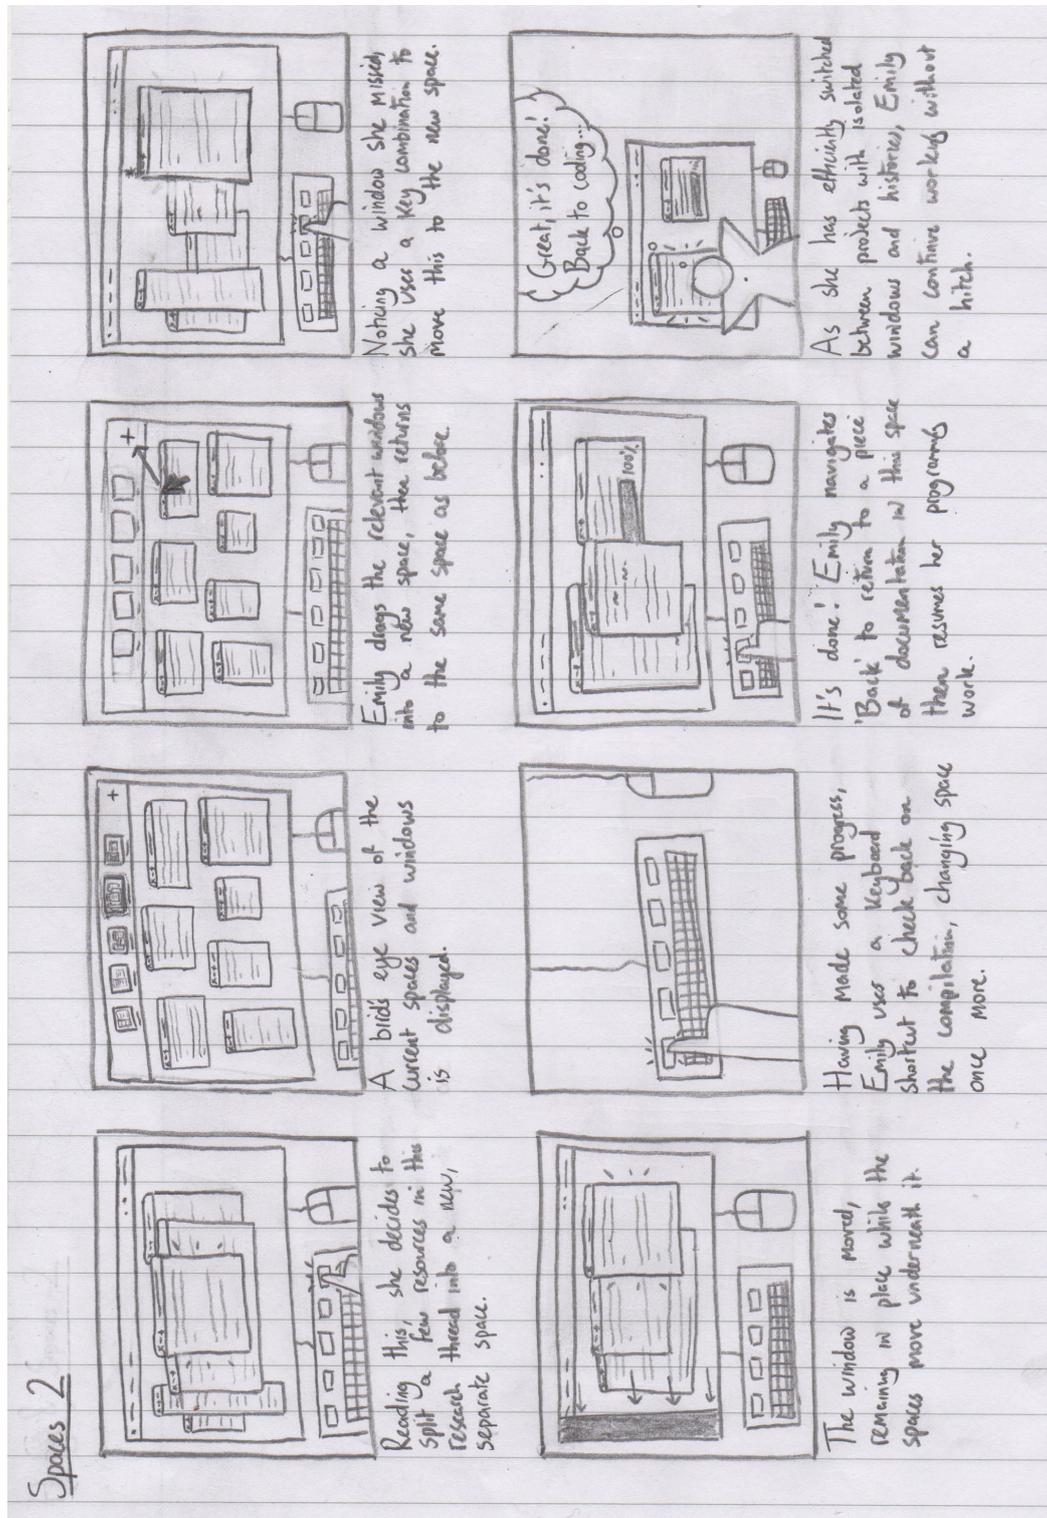


Figure 5.9: The second page of the spaces storyboard.

Multi-space Window Management

In considering the suitability of multiple desktops to history management, a natural question to ask is how people make use of this functionality within existing operating systems. In a study of window management techniques from Hutchings and Stasko (2004), a quarter of participants made use of multiple desktops in some way or another. While this sample appears to be biased towards computer scientists, it seems a reasonable starting point in our context of knowledge work. With this in mind, is it reasonable to expect any more than this proportion of users to utilise spaces in maintaining separate chains of history? Contrary to the figures, we suspect that if space switching and manipulation are sufficiently low cost operations, the increased value from separating unrelated groups of windows and maintaining separate history chains may be enough to push users in this direction. Of course, user testing is required to validate this.

Relatedly, in a study of virtual desktop usage strategies among 20 participants, Ringel (2003) shows some variety in how users make use of multiple desktops. What is most interesting about these results in the context of our prototype is that all but one of the participants partitioned their work between desktops in a way that seems to lend itself well to separated history chains. The majority of participants, for instance, divided their windows among desktops based on tasks or subtasks. Furthermore, the study shows a large degree of consistency within how individual users mapped their windows to desktops over time. The authors suggest this may indicate that users are taking advantage of spatial memory in managing their desktops — a property we have already identified as potentially advantageous.

Returning to the topic of intra-space window management, Hutchings and Stasko (2004) also observe that participants' window organisation strategies fit within one of three broad categories. Curiously, all virtual desktop users were in the “careful coordinators” (CCs) category, where users tended to have multiple windows visible simultaneously. This contrasts most strongly with the “maximizers” category, where users maximise most windows. In some sense, virtual desktops being used exclusively by CCs is unsurprising. These users have the most to gain from multiple desktops, as these allow them to maintain their spatial arrangement of windows across various tasks.

We theorise that tabbed interfaces, with their all-or-nothing approach to switching between resources, may encourage a “maximizer” type approach to resource management. Interfaces that rely more heavily on multiple windows, on the other hand, seem to more naturally suit the “careful coordinator” method of organisation. This is supported to some extent by our exploratory study, in which participants largely relied on “maximizer” type organisation strategies, and where participants indicated that comparing and cross-referencing resources would be a lot easier with multiple windows rather than tabs. In open hypermedia, then, it may also be the

case that purely as a result of the multi-window nature of these systems, users are more naturally encouraged to use multiple desktops.

Other Issues

One final point worth considering in our discussion of spaces is their influence on task and working sphere restoration. As discussed in our literature review, there is some evidence that problem state plays a large role in task resumption (Borst, Taatgen and van Rijn, 2015). Consequently, our prototype takes inspiration from the persistent “projects” of early Smalltalk systems at Xerox PARC (Kay, 2017). In these systems, and our proposed solution, desktops and their contents are maintained across power cycles. We hypothesise that this can significantly reduce problem state restoration when approaching a newly booted machine, as users can find spaces and windows in the exact spatial arrangement in which they were left.

In general, dealing with history in a multi-window multi-space system is a tricky endeavour. As such, we cannot answer all the open questions with regards to how this functionality might work in only 16 panes. If resources have internal links, for example, should these contribute to the same history as cross-program links? And how exactly should backtracking work within spaces anyway? For the latter question, we currently imagine that chronological backtracking through windows is employed, but it may be the case that schemes such as task-based multi-window backtracking (Bieber and Wan, 1994) are more useful to users. Such questions are difficult to answer, and are not considered to be in scope at this stage of prototyping.

5.1.3 Window Search

As discussed previously, a potential issue with many-window hypermedia systems is the difficulty in finding information amongst a large number of open windows. This issue is especially pronounced when comparisons are made with tabs, which are predominantly used for organisation within WWW browsers today. Research shows, for instance, that users may find windows less organised or more confusing than tabs (Dubroy and Balakrishnan, 2010). In addition, participants within our exploratory study indicated that multiple windows could make things more difficult to find, while tabs provided an extra level of hierarchy to help them efficiently navigate through open resources.

Clearly, then, open hypermedia systems should employ some kind of strategy to help alleviate these issues. To this end, one idea that arose naturally during discussions within our exploratory study was some kind of search functionality to quickly filter through and navigate between open resources. Our proposal in this area follows exactly this idea, allowing users to search through the titles and contents of open windows. The full storyboard for this solution can be found in Figure 5.11.

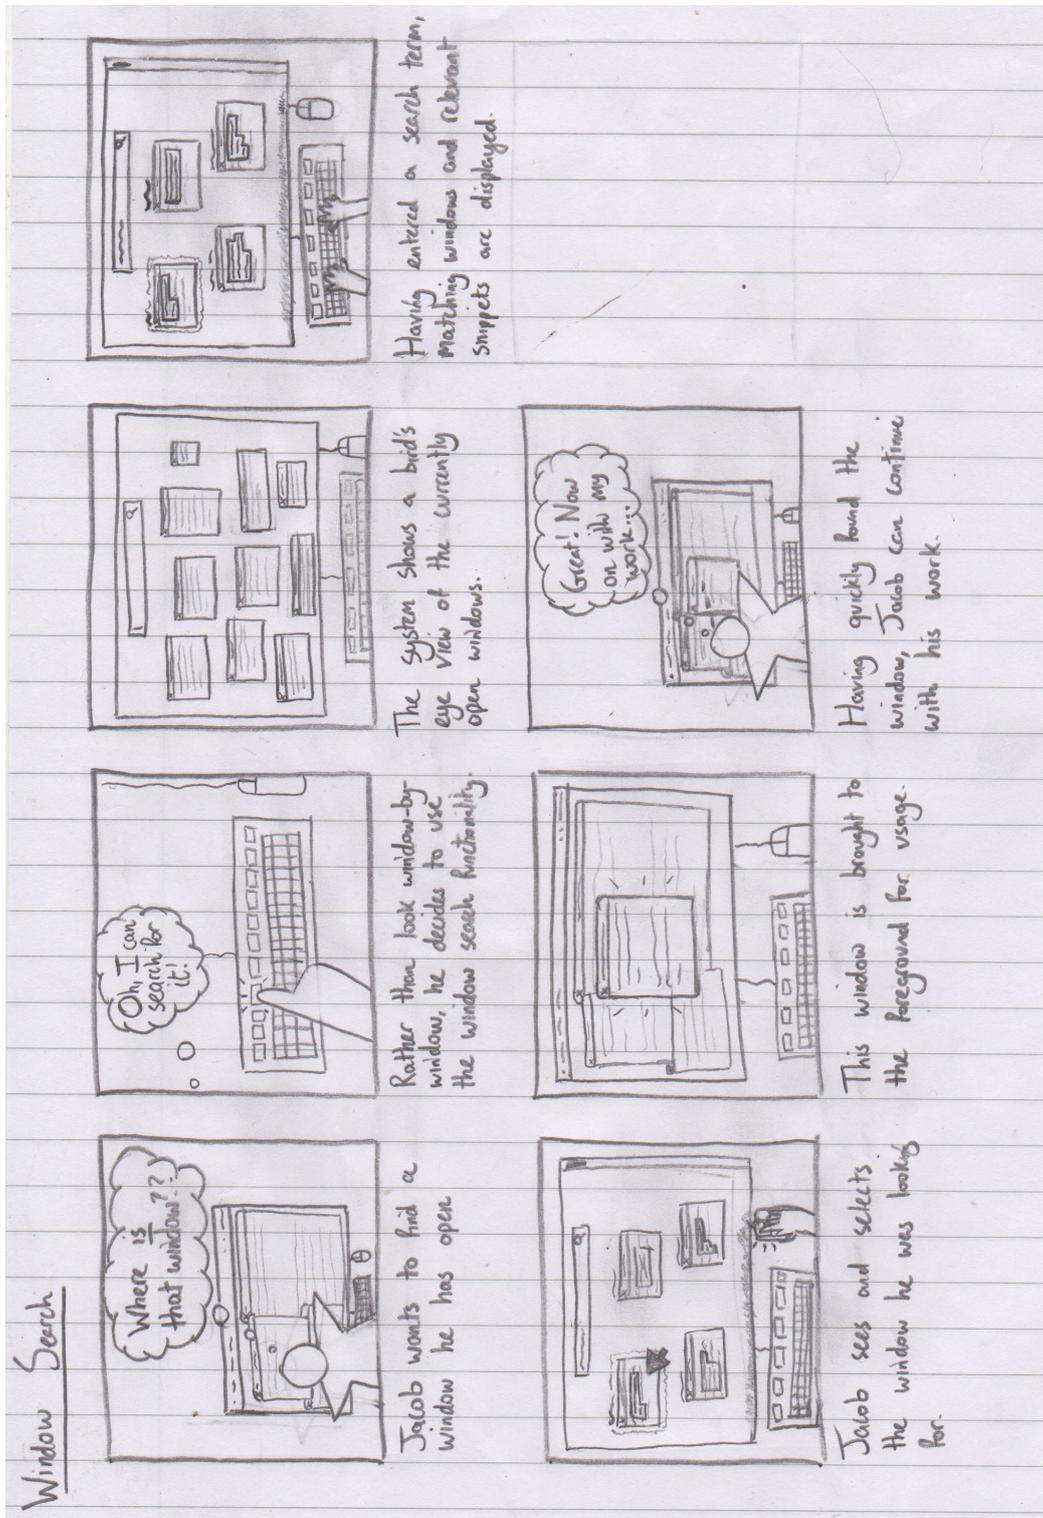


Figure 5.11: The window search storyboard.

Once again leaning on macOS-inspired user interface ideas, our solution is based around a Mission Control style (Apple Inc., 2016) bird's-eye view of the currently open windows, which we additionally extend with search capabilities. As the user types a search term, the set of open windows displayed in this view is filtered through in real time. Results are displayed using both visual previews and matching snippets, as pictured in Figure 5.12, intended to aid the user in efficiently locating a specific window.

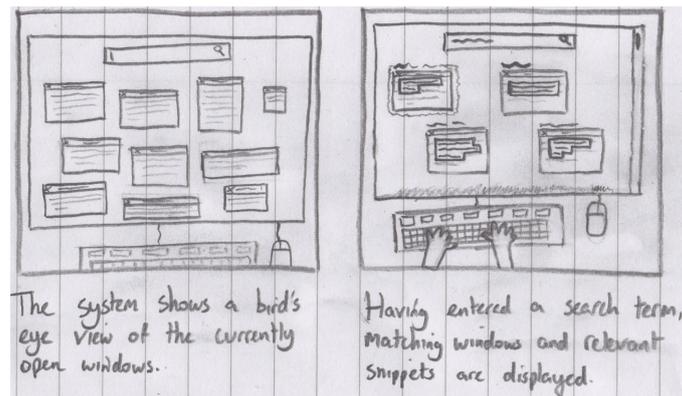


Figure 5.12: A two pane extract from the storyboard showing window search behaviour.

The hope is that by using this fast, direct method of switching between resources, users will be able to navigate between a large number of open windows more efficiently. The potential value in this approach is emphasised by participants' strong desires for the ability to search within the exploratory study. The value of this approach in practice, though, is of course dependent on the extent to which users know what to search for in order to quickly find a desired resource. We suspect that this is not such a big ask in many cases, though, and that users will hone such a skill in using the system over time.

Some readers may question the technical practicality of this functionality. Is it really realistic, for instance, to be able to search through many hundreds of different kinds of resource using a single algorithm? While implementation details are not the primary concern here, we do not consider this functionality to be at all unrealistic. As with all other storyboards outlined in this chapter, a reasonable implementation could have viewer programs implement various pieces of functionality required by the system. In this case, functionality to efficiently search currently open resources. This is fundamentally no different to the many ways in which we expect programs to cooperate with operating systems today, for instance in supporting clipboard behaviour, and can be largely handled by system libraries.

Some readers may also be skeptical that real-time search of a potentially large number of open documents is achievable on typical consumer hardware. Such matters are not of great concern here. The design of the search algorithm (including use of partial evaluation and offloading of work) and capabilities of the hardware can be tuned to achieve the desired level of real-time behaviour within some set of constraints.

5.1.4 Exploration Mode

To complement direct navigation mechanisms like our window search proposal, participants in our exploratory study also suggested that dynamic, associative structuring of information could be incredibly useful in navigating information spaces. While discussions in this area mostly concerned resource maps, the potential issues in scaling such mechanisms to large information spaces (Utting and Yankelovich, 1989) prompted us to pursue a slightly different path. Particularly, our solution is motivated by the insight that the ideal display of connections when primarily focused on the content of a few resources is likely different to the ideal display when exploring the connections between many resources.

This idea is supported by our exploratory study, in which participants exercised intense focus when attempting to understand the technical details of a single resource, but also yearned for better tools to explore the information space when focusing on relationships between resources. To accompany the relatively subtle display of connections in typical viewing, we propose a new *Exploration Mode* to accommodate more exploratory activities. This can roughly be thought to combine the ideas of machine intelligence and recommendation with resource maps and transpointing, yielding an environment designed specifically to support exploration. The full storyboard for this solution can be found in Figure 5.13.

With a resource window in focus, the user can invoke a key combination or click a menu item to activate Exploration Mode. Having done so, the system directs focus exclusively onto a single window and its relationships with other resources. Previously subtle links become prominent and colourful, and machine intelligence is employed to suggest interesting, related resources to the user. These resources may include, for instance, relevant open windows, similar or linked resources in enabled linkbases, popular links that others have made from this document, or pertinent material that link to this resource. This is pictured in Figure 5.14.

To facilitate the rapid inspection of links and associations between resources, this mode additionally provides a form of transpointing, as pictured in Figure 5.15. In effect, Exploration Mode puts the entire power of the computer towards helping the user explore the connections between information. This can be thought to go some way to resolving the tension between focusing on content and connections in hypermedia systems, as is highlighted by Obendorf and Weinreich (2003) in their

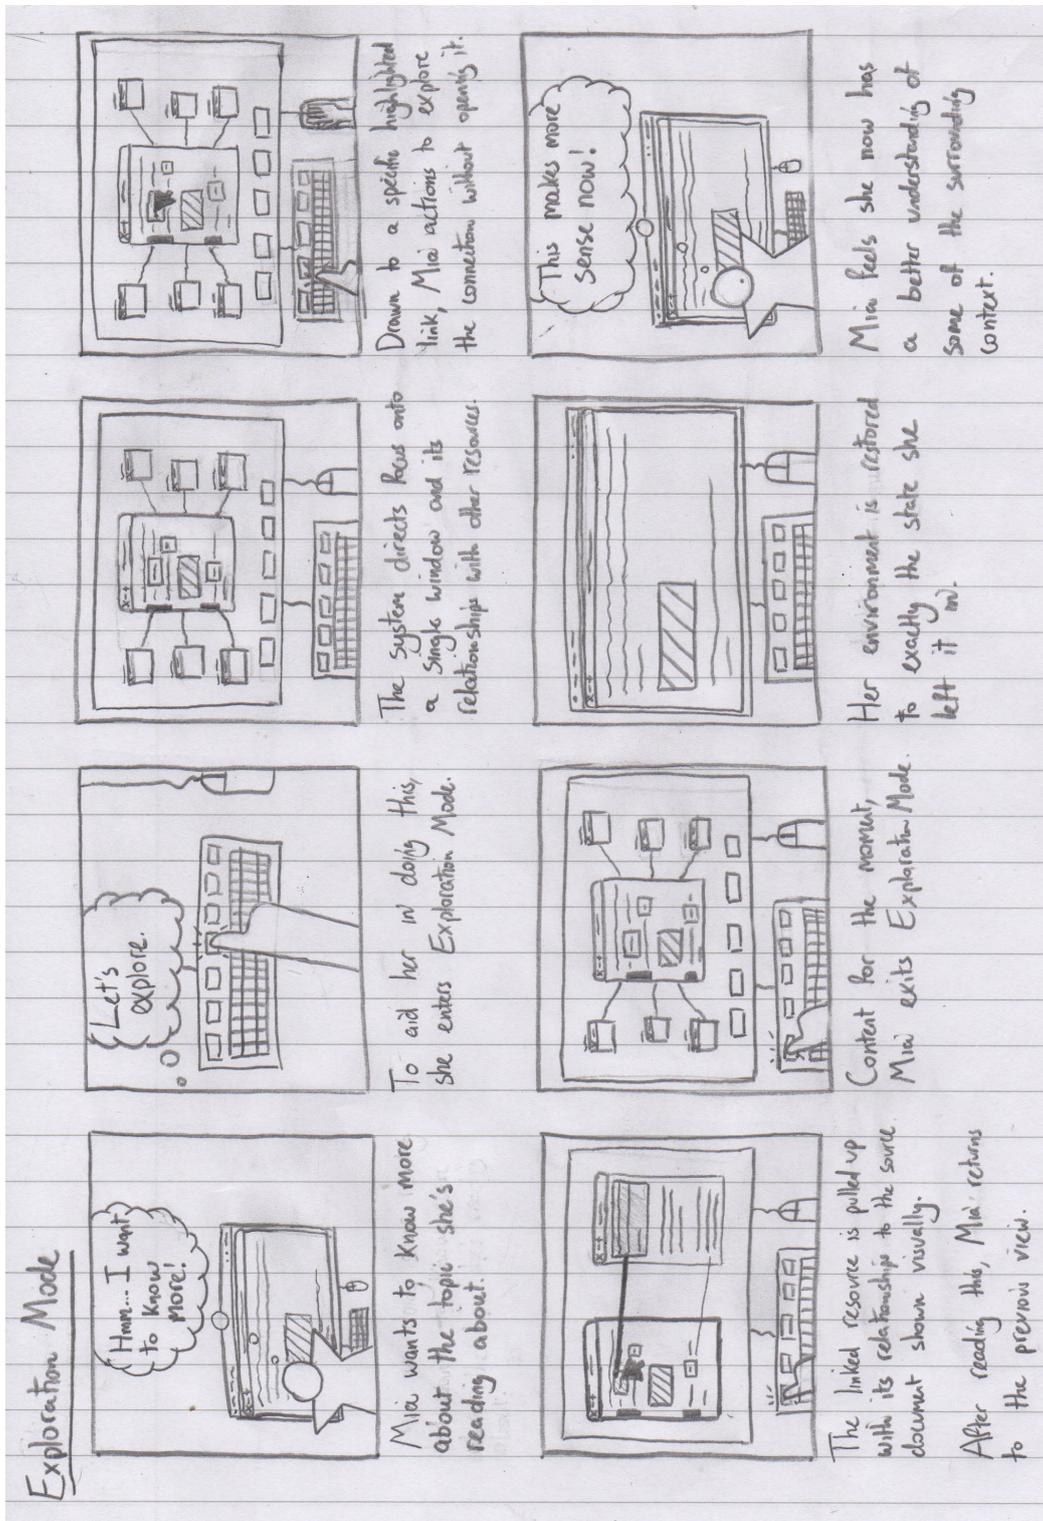


Figure 5.13: The Exploration Mode storyboard.

consideration of always-visible and on-demand links. Subtle link highlighting and the marker bar can indicate links without detracting significantly from readability when focus should be primarily on content, while Exploration Mode can facilitate rich traversal when focus should be primarily on connections.

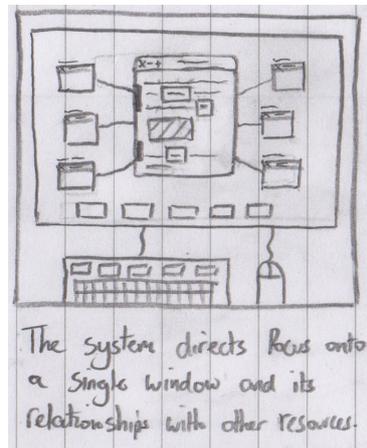


Figure 5.14: An extract from the storyboard showing Exploration Mode.

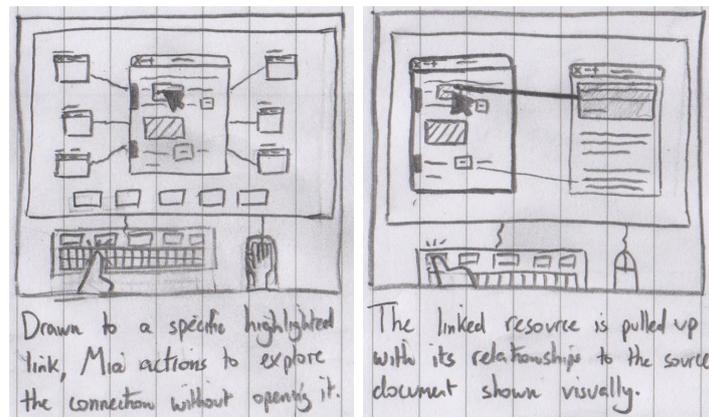


Figure 5.15: A two pane extract showing transpointing in Exploration Mode.

When users are content with the results of their exploration and wish to return to a normal view, they can action to exit Exploration Mode. Having done so, they will find their desktop environment exactly as they left it, though perhaps with a few new resources open or some changes in the marker bar from the exploration. Providing that the transitions in and out of Exploration Mode are sufficiently fast, this should allow users to switch between the two modes freely and naturally as they desire.

5.2 Summary

In this chapter, four potential solutions to significant open hypermedia design problems have been proposed. Taking ideas from existing literature and our exploratory study, these storyboards illustrate exciting possibilities for practical open hypermedia environments of the future. Particularly, surrounding the display and behaviour of links, separation of multiple history chains, efficient filtration of open windows, and exploration of vast information spaces. While each of these solutions aims to support users in achieving their goals, however, what is not yet clear is the degree to which they might be successful in doing this. What is to be done next, then, is to expose potential users to these designs in order to evaluate their potential merits and faults.

Chapter 6

Storyboard Evaluation

Having crafted four storyboards to illustrate various aspects of an imagined hypermedia system of the future, let us reflect on the fact that these design ideas exist purely to support users in achieving their goals. Although each design has been constructed with careful consideration of users' needs and is influenced by practical experiences from our exploratory study, user interactions with software are often complex. As a result, the practical value of our ideas and many of the details of their ideal implementations are not immediately clear. In this chapter, we set out to clarify these uncertainties through a user study to evaluate our design proposals.

6.1 Procedure

As was the case with our exploratory study, we began by establishing a formal procedure for the collection and analysis of our data. This procedure, which we describe in the sections that follow, makes certain that participants are approached with some level of consistency, and further ensures our compliance with the University of Bath's ethics guidelines.

6.1.1 Participant Recruitment

Participants are to be recruited primarily through the university's computer science department. Participants that have taken part in the exploratory study are ideal candidates for this due to their increased familiarity with the surrounding context of our research and with a number of the problems that our design solutions attempt to address. A background in computer science is not required for this study, though, and the inclusion of some non-computer-scientists may be beneficial in obtaining a more diverse set of viewpoints.

As before, each participant must sign an informed consent form prior to any data collection (section A.5). This briefs participants on what they will be asked to

do and informs them of their rights. In this case, a total of six participants were recruited. Five of these had experience with the exploratory study — whether through the main study or the pilot — while one did not. The participant who did not take part in the exploratory study — Participant 5 — is, in fact, a maths student, albeit a maths student with a strong background in computer science.

6.1.2 Data Collection

Each single-participant session is to take around 45 minutes. In the first few minutes, participants are to be briefed on what is expected of them. Included within this should be a reminder of the surrounding context of this project. That is, knowledge work, and supporting such work with a new kind of hypermedia system. Participants without experience in the exploratory study should be especially briefed on this topic as to ensure that the purpose of each design proposal is properly understood.

Immediately following this briefing, an audio recording is to be set up to preserve the session for future analysis. This prevents copious note taking during the session, where the interviewer’s time may be better spent on prompting participants. The interviewer is then to walk through the storyboards pane-by-pane, allowing participants to interject with comments or questions at any point. In doing so, the storyboards are to be presented in the order in which they appear in the previous chapter. This ensures that potential solutions to fundamental issues, such as link and history behaviour, are presented first, so that such issues do not dominate discussions around other prototypes. In walking through the “Window Search” and “Exploration Mode” storyboards, the interviewer may also wish to add additional surrounding context, as the initial panes of these are somewhat general in nature.

After each storyboard has been communicated, a short discussion is then to take place on the participants’ thoughts. This should take the form of first ensuring that participants understand the ideas in the storyboard, and then prompting them for their thoughts on various aspects of the design. Within this, the interviewer should ask participants whether they feel that the core idea within each storyboard actually solves a practical problem, as this may help to evaluate the general utility of each design idea. Besides this general guidance, however, the particulars of each conversation should be tailored to the comments made by participants in order to extract the most value from each session.

6.1.3 Data Analysis

As this study is more focused than our exploratory study, so too is the process of analysing its data. Analysis will first involve listening through the set of audio recordings and selectively transcribing comments of interest. While the default behaviour will be to transcribe comments made by participants, speech from the

interviewer will be transcribed only when it provides context to these comments. Additionally, sections of conversation in which participants are trying to get a handle on how to express their thoughts, or in which participants move to unhelpful or repeated topics, may be excluded from transcription.

Qualitative analysis of these transcripts will then take the form of thematic analysis (Braun and Clarke, 2006), as motivated previously in section 4.1. Unlike our exploratory study, themes will be partitioned into five categories — one for each of the storyboards, and another for general comments. This directly associates feedback from users with the storyboards themselves, clarifying the merits, faults, and potential extensions identified for each, and more easily allowing this feedback to be incorporated into future iterations of our ideas.

6.2 Results

The results of our analysis follows. As before, descriptions of each theme are accompanied by relevant discussions. The raw interview transcripts upon which this analysis was performed can be found in section A.2.

6.2.1 Link Behaviour: Link Highlighting

Perhaps the most visually identifiable characteristic of any hypermedia system is the way in which links are displayed. Breaking the established convention of presenting links using bright blue, underlined text, our storyboard proposes that links are highlighted using subtle translucent overlays. Feedback on this design decision was mixed, splitting participants into three groups.

The first group of two participants stated that they liked the change, emphasising that it could be less distracting than current alternatives. “Sometimes it’s nice just knowing that there’s a link ... [but] if you’re reading a passage[,] you don’t want to be distracted”, Participant 1 stated, praising the subtlety of our proposed approach. “I like the way that links don’t detract your attention too much from the entire article ... it’s just quite subtle which is good”, Participant 2 added.

The second group of two participants, meanwhile, were less certain about this new method of link highlighting. Both participants raised questions around the accessibility of this approach in particular, discussing potential issues such as colour blindness and link display on different monitors. While we consider these issues to be important, we ultimately view them as addressable. High quality colour displays are commonplace in homes and offices around the world today, and accessibility issues can be properly accounted for through careful selection of default colouring in combination with accessibility options where appropriate.

Lastly, the final group of two participants questioned why the system was deviating from established conventions. “I don’t see much point in it being different”, Participant 3 said. In describing the potential benefits of breaking convention in this area, these participants seemed to remain largely ambivalent about this change. Participant 6, whose primary concern was that overlays may be *more* distracting than underlined text, emphasised that “it definitely shouldn’t be too obtrusive”. As our storyboard lacks detailed information about the colour and brightness of our proposed method of link highlighting, higher fidelity prototyping may clarify opinions on this issue.

6.2.2 Link Behaviour: The Marker Bar

To accompany inline link highlighting, our storyboard employs a somewhat novel mechanism for displaying interaction opportunities: the marker bar. Somewhat surprisingly, this rather drastic departure from how things work today was well understood and generally well received by participants. Participant 2 said that it would be “really convenient for noticing where different links are”, and Participant 1 noted that it would be a “really good addition”, identifying its value in helping users “instantly see what’s here”.

There was particular enthusiasm regarding the ability of the marker bar to highlight non-textual links when this topic arose in conversation. “Yeah! That’s even better!”, Participant 1 exclaimed. Not all participants felt that the marker bar would be helpful, though. Participant 5 particularly questioned its motivation. “What would be the point of looking at that instead of just thinking ‘there it is, the link’s there’?”, they questioned. In outlining the imagined advantages to the participant, they seemed to remain unconvinced. More detailed prototyping — or better still, having users actually interact with the proposed ideas — may help to clarify opinions in this area.

Many participants additionally questioned how the marker bar — and links in general — would behave in the case that there was more than one link on a single line or surrounding a single phrase. These are edge cases that were not answered in the design of this storyboard, but that ought to be addressed in future work. Relatedly, Participant 3 considered how the marker bar might handle displaying marks for very many lines simultaneously. “If there are loads of links, that might get really annoying ... [as] the marker bar is just full”, they commented. Once again, this is an edge case requiring further thought. More developed proposals in this area must decide if the marker bar should be designed in some way to help alleviate this problem, or whether users are responsible for avoiding such situations.

Participant 3 — along with a number of other participants — also pointed out that it would be useful if the marker bar facilitated some kind of interaction, for instance through hovering or clicking. This supports our supposition that such functionality

might be useful, as described in the previous chapter. This was not the only suggestion that participants had either. Participant 2, for instance, proposed marking “different kinds of links” in different ways, while Participant 1 suggested that the marker bar could accompany the scroll bar in showing the positions of search results. These varied suggestions go some way to supporting our idea that the marker bar can be a general purpose tool for displaying information, on equal footing with the scroll bar.

Much like the scroll bar, however, the marker bar is not suited to all applications. In contrast to the contrived environment within the storyboard, programs in the real world would have to make their own decisions on how to display links. Ideally there would be some degree of consistency across programs in how this is achieved, but the marker bar is unlikely to play a universal role in this. It is not clear, for instance, how useful the concept of the marker bar might be for displaying links in scientific simulations, spreadsheet programs, or 3D environments. As such, the consistency of link display mechanisms shown in the storyboard may be misleading, and may artificially skew participants towards more positive responses than they would otherwise have regarding links in the proposed system. Once again, higher fidelity and more realistic prototypes may help in better gauging user response on this issue.

6.2.3 Link Behaviour: Navigation

The default navigation behaviour of the proposed system — in which linked resources open in new windows — was not widely commented upon by participants. Perhaps this is because this kind of constructive opening behaviour, in contrast to the destructive “replace the current page” behaviour used in WWW browsers today, is a natural enough way for things to work in a multi-window system supporting heterogeneous programs that it did not draw their attention. Participant 2, one of the few to express their views on this issue, remarked that:

“I like the way it opens a separate window too, so you can still go back to the previous one and there’s kind of like a hierarchy of what you’ve opened.”

Alluding to the potential drawbacks of this approach, Participant 6 commented that “it’s nice to be able to have everything in one place, like sort of a hierarchy, but also I think that could very easily get messy”. This seems like a rational concern, and is related to the fundamental questions of how many windows will practically get spawned in this kind of system and what utility each will have. One thread of enquiry fundamental to answering these questions concerns a topic that arose in discussions with Participant 2 regarding the suitability of a system like this

for interacting with services such as YouTube (and similarly, e-commerce, SaaS, and so on).

Today, Web pages provide both the infrastructure for navigating YouTube and the means of viewing the videos, descriptions, and comments contained within. In this context, Participant 2 seemed concerned that the kind of system being proposed might be overkill or unsuitable for this kind of task. On a practical level, it's easy to imagine that navigation through many separate infrastructural pages and windows may be unintuitive, and would spawn a number of windows that aren't terribly useful to have open at the same time. One fairly natural solution to this would be to simply avoid using the core infrastructure of the information space for navigating between states of an application or service.

Using interconnected WWW pages to interface with applications is a hack that has stuck around (Graham, 2001): the WWW was not originally designed to support these kinds of services (Berners-Lee, 1989), and this is reflected in its architecture today. Similarly, the system we propose is designed for efficient resource navigation, not for changing the state of “Web apps”. Instead, services can exist as normal applications — either as native programs, or as remote “service” type resources to be viewed with some “Web app viewer”. These can then interact with the network as they see fit, just like any other program on the user's computer, and need not unnecessarily open windows or contribute to the history log. In the case of YouTube, you might imagine an interface closer to that of native iTunes or Spotify clients today, though one that may additionally interact with the hypermedia system (e.g. when viewing videos). In outlining this to Participant 2, they indicated that this may resolve their concerns.

6.2.4 Link Behaviour: Linkbases

Fundamentals

Although link display and navigation are important, one of the key features that sets our proposed system apart from existing systems is linkbases. In discussing this behaviour with participants, we were once again somewhat surprised to find that not only did they pick up the concepts relatively quickly, but also that they saw great value in them. Participants expressed particular interest in the way that linkbases allow users to “bring out what's important from different points of view”, and in how they level the playing field such that “anyone could make a link database”.

Participants also responded positively to the proposed interfaces for both creating links and selecting linkbases. Participant 6 commented that “making the link physically drawing a line feels intuitive”, while Participant 3 said that it “just looks really easy to do”. Participant 4 also stated that linkbase selection was “intuitive”

and that it “would be pretty easy for people to understand”, while Participant 1 said that they “don’t know how else you’d make it more intuitive than that”.

With respect to details not shown within the storyboard, Participant 6 identified that the key combination for link creation would need to be taught to users somehow. A few options for how this might be achieved were discussed within the session, including a tutorial at first contact with the system, and having users start by using the mouse exclusively and displaying the keyboard shortcut besides the relevant menu item. More generally, learnability is an important characteristic of any solution in this space. Particularly in the context of the rather unconventional nature of the whole system, which Participant 4 said would be “something you’d need to get used to”.

Regarding linkbase selection, Participant 1 also pointed out that that they would hope that users might also be able to “expand and get more information” about linkbases, such as descriptions. While it’s not entirely clear how this should best be accomplished, the desire for such behaviour certainly seems reasonable. A few participants also questioned how linkbase sharing and discovery might work in the proposed system. These are excellent questions, but are far beyond the scope of this prototype.

A related area raised by participants involved the granularity of linkbases. If linkbases are too small or too large, their value diminishes significantly. Either the management cost of enabling and disabling a few links at a time becomes overwhelming, or the noise in showing a large number irrelevant additional links becomes an inconvenience. With regards to the former case, Participant 5 voiced their opinion that “unless people are taking this seriously . . . you’d just be left with a load of like tumbleweed linkbases”. This is definitely a concern, and is perhaps best dealt with by the linkbase sharing and discovery mechanism. Meanwhile, to tackle the issue of monolithic linkbases, Participant 6 suggested that you might be able to view “a subset of the other dude’s linkbase”. Depending on the ecosystem characteristics, perhaps such behaviour would be desirable.

Terminology

In discussions with Participant 4, the observation arose that the naming of some concepts within our proposals seemed unnecessarily obtuse. “I would question the naming of some of the things and simplify it down a bit”, they said. This is an excellent point, and is not something we put much thought towards in the design of our storyboards. In having to explain the storyboard to others, it quickly became clear that “linkbase” is not an especially helpful term. To communicate the underlying ideas to participants, the interviewer would first expand this to “link database” before clarifying its true purpose as a “perspective on how information is connected”, which participants seemed to understand well.

It seems that carrying over terminology from the existing literature was a poor design choice in this area. In their own words, participants referred to linkbases as “filters on the information to bring out what’s important from different points of view” that allow users to enable someone’s links to “see what they think”. In light of this, and the success in explaining linkbases to participants as perspectives on connectivity, we suggest renaming linkbases to something more indicative of their use cases rather than their implementation details. Something like ‘perspectives’ ought to be more illustrative, and may additionally cause less confusion among knowledge workers unfamiliar with the concept of a database.

Potential Issues

In considering potential issues with linkbases, Participant 5 raised the issue of referential integrity and broken links. “What if the page that you’re linking to disappears, or changes and has different content? . . . How do you distinguish between what’s still valuable and what’s not?”, they asked, adding that “eventually stuff will just get out of sync . . . that’s a load of work just gone”. This is an important point: by decoupling link structure and content, open hypermedia systems are vulnerable not only to changes in URLs, but also to changes in content, invalidating links. This issue has been the subject of much existing research (Davis, 1998). We believe that satisfactory solutions exist, though this topic is beyond the scope of our work.

Participant 4 also touched upon the potential security issues surrounding the system. On the subject of linkbase discovery and link display, they noted that people could make “troll links” to undermine the value of the system, or link to propaganda and viruses. Undoubtedly this system changes the relationship of trust that users have with resources from the Web to some degree. Whereas users today can have some level of trust that links from a BBC news article, for instance, will be safe to follow, linkbases distort this picture.

Though this is somewhat of a concern, we take the viewpoint that linkbase discovery mechanisms can deal with this much in the same way that search engines deal with spam today. This is not to say that search engines today cannot be gamed, of course, but just that we can roughly achieve parity with what exists today. A more pressing issue, perhaps, is how this system might avoid phishing attacks. As proposed, it’s not clear whether the system displays URLs besides resources or how users can verify that a given resource came from the source it was expected to. These issues of trust and authentication are critically important, but are beyond the scope of this project.

6.2.5 Link Behaviour: Doubts

While most participants were generally positive about the ideas underlying this storyboard, and particularly regarding linkbases, one participant in particular expressed significant scepticism. A large proportion of the conversation with Participant 5 — the only participant without exposure to the exploratory study — concerned the central motivation behind why one might use this system. “I can’t really see when I’d use something like this”, they stated, referring largely to link creation and linkbases.

Examining the context in which these remarks were made, we suspect that the participant felt this way largely due to two factors. Firstly, although the participant was briefed and prompted to think in the context of research, they often strayed into considering the wider applicability of the system beyond what was described in the storyboard. “It seems almost too academic for something that would be considered useful for the general population”, they commented, noting also that they “don’t read many papers”. Although the applicability of the system outside of knowledge work is not of particular interest here, this viewpoint emphasises that the benefits of open hypermedia systems are not trivial to appreciate. While we believe that linkbases can be highly useful in a wide variety of circumstances, the ideas that underpin such applications are not obvious and must be carefully and concretely explained such that they can be properly understood.

In addition, the participant seemed very much concerned with implementation details of the system rather than the design details. Instead of supposing that the proposed system could solve some of the potential issues that could arise, the participant’s vision of the system seemed very much clouded by implementation challenges. In discussing referential integrity, for instance, they said that “there’s just too many variables on how stuff’s going to change”, indicating a belief that the problem may be insurmountable. In considering the many hurdles that such a system would face, the participant seemed to react negatively to the drastic nature of the proposed changes, desiring something closer to what exists today.

On the topic of linkbases, for instance, the participant asked whether this behaviour could not “just be replicated by people keeping a list” of links. While it’s true that the functional behaviour of linkbases could be approximated by users creating and uploading many lists of links on the WWW today, the usability characteristics of the two approaches differ wildly. As other participants seemed to appreciate, making link creation and linkbase display seamless parts of the system can lead to significant efficiencies and changes in behaviour. In the words of Engelbart (1962): “The important thing to appreciate here is that a direct new innovation in one particular capability can have far-reaching effects throughout the rest of your capability hierarchy” (p.14).

Ultimately, we do not consider the doubts regarding the fundamental value of the system raised by this participant as especially worrisome. Instead, we view the discussion as a worthwhile examination of some of the potential challenges in constructing such a system. The change-resistant viewpoint provided valuable contrast, and continued to do so through the discussions of other storyboards.

6.2.6 Spaces: Fundamentals

The second storyboard shown to participants concerned history. Participants made clear that history was an important feature to them, and the fundamental idea of splitting history chains across virtual desktops was well received. Participant 1 noted that history “doesn’t seem like something that’s been forced on top of workspaces”, concluding that they “like this idea of history through workspaces” and that “having the history makes this paradigm underlying all the other storyboards manageable”. Participant 2 similarly said that the system was “straight-forward and intuitive”, while Participant 3 said that it would be “awesome” and “super useful” for research.

In comparing the system to existing multiple desktop systems, participants indicated that the essential components were already well-established and that the addition of history would increase the value of using spaces. “I think what adds to this is the history feature”, Participant 1 said, “now I have a reason to keep work separate . . . it gives more of a reason to start breaking things down”. Participant 3 also expressed satisfaction that our proposal was largely based on how multiple desktops work within macOS, saying that multiple desktops are “quite hard to do” on Windows but are “really easy” on macOS.

With regards to the persistence of spaces, few participants commented on this. Perhaps as there were more pressing matters to discuss. Participant 1, the only subject to discuss this matter, indicated that persistence could be useful, but that it was best to “not force it on people”. They remarked that some people like to boot up to a clean environment, and so that this functionality ought to be configurable.

6.2.7 Spaces: Granularity

The most substantial questions that arose regarding spaces surrounded their granularity. While some participants viewed spaces as a natural unit to draw history boundaries, others saw them as only the coarsest-grain unit in history management. Participant 1 said that they “don’t think it’s a hard thing” to use spaces instead of tabs, and that “people would just need to get used to bringing it up a level”. In contrast, Participant 4 questioned how granular spaces should practically be, saying that “you might not want hundreds of thousands of spaces” and suggested that “you might want another way to sort of sub-group within a space”.

Participant 5 expressed their concerns around spaces being the only unit of history management particularly clearly. They observed that our proposal uses spaces “for both different work and different trails of thought on the same piece of work”, which they consider as “different things”. Elaborating on this, the participant noted that spaces felt too detached and separate for slight tangents, stating that they think of switching spaces “almost like context switching”. “When I switch space, I want to be thinking about something completely different”, they said.

This seems to speak to how people think about multiple desktops. In theory, switching between tabs and switching between desktops are very similar operations. Transient spaces ought to be able to deal with tangents just as well as transient tabs. In practice, however, some participants seem to think about spaces as far more heavyweight and independent than tabs. As such, perhaps this solution requires changing the way that people think about multiple desktops. One way of accomplishing this might be to support more efficient space switching. By adopting a keyboard shortcut such as command-click to quickly open links in a new space, for instance, maybe users can begin to view spaces are more lightweight.

Alternatively, if changing the way that people think about spaces is not desirable, perhaps additional in-space history manipulation mechanisms could empower users to better manage history chains within a single space. Task-based backtracking (Bieber and Wan, 1994) could be considered one approach to this. Other approaches might include the ability to preview links or to explore links within a transient history chain, both of which are discussed alongside the Exploration Mode storyboard shortly. In fact, this exact topic is discussed with Participant 5 in subsection 6.2.11, where the participant indicates that Exploration Mode may go some way to resolving their concerns around the unsuitability of spaces for exploring tangents.

One potential solution that participants generally responded poorly to was the ability to group windows within a space, creating sub-spaces of sorts within a single space of windows. While Participant 4 indicated that this “could also be useful to have”, several participants indicated that this would be “more confusing”. “I like having them logically separated”, Participant 3 commented, “if I hit the ‘Back’ button, it will go back to where I expect it to”. Participant 5 added that “if you have too many options, you’re just never going to utilise them in the right way” and that having to mark groups of windows within a space puts “more onus on the user”.

6.2.8 Spaces: Questions & Suggestions

Participants additionally had a number of questions around the details left out of this storyboard. Participants 2 and 4, for instance, brought up the issues of learning and avoiding key combinations in history and spaces. “Would there be any GUI-

based approach for going back?”, Participant 4 queried, continuing that “learning new key combinations ... is always a thing I’m going to be anxious about”.

This is a good question. There ought to be some way of navigating back without using a keyboard shortcut, but as history operations are necessarily cross-window and cross-program in nature, it is not clear how this should work. Having reflected on this question for some time, it seems that history management is very similar to linkbase management in this regard. As such, one possible solution is to have a history menu bar application to accompany the linkbase menu bar application. ‘Back’ and ‘Forwards’ actions could be displayed within this, alongside a list of recently visited pages, allowing users to quickly jump back through multiple resources (as is also supported by WWW browsers today).

Participant 6 also questioned whether for a large number of spaces, there would be a way to view them all at once for direct navigation through the GUI. “Would there be a way to see them all at once, or would you have to go along five at a time?”, they asked. Again, this seems like a reasonable thing to ask for, though it’s not clear how it would best be accommodated.

Additionally, a number of participants touched upon awkward edge cases and ambiguities within the history mechanisms as described. When moving a single window to a new space, for instance, we had imagined that this resource would be removed from the history chain of the source space and appended to the history chain in the destination space. If the moved window was not the latest visited resource in the source space, though, this creates a hole in the history chain, as described by Participant 5. In such a case, it’s not clear whether going ‘Back’ in the source space should re-open this resource in the space when it encounters it, or simply skip over it.

Equally, Participant 4 questioned the behaviour of the system when moving a group of windows from one space to another. In moving a group of windows, is any effort made to preserve the existing history relationships between these windows? If so, what happens if a group of windows are moved separately to a single window that — if moved as a part of the group — would have had its history with the others preserved? The participant also queried what would happen when moving a window into an entirely new space and then pressing ‘Back’. Is it better to do nothing, or to “close the space and go back”?

One last edge case, also raised by Participant 4, regarded chronological and task-based (Bieber and Wan, 1994) backtracking. “If I have like five windows open, and they each went from one to the other ... if I switch to document three and then go back, do I go to document two?”, the participant asked. In explaining that it was imagined that this action would instead go back to document five, the participant

responded that they could “see both pros and cons of that”, but that they “would have said it makes more sense for three to go back to two” but that “it’s kind of hard to know”. Having users actually use each method may help in deciding which backtracking mechanism people practically prefer.

6.2.9 Window Search: Fundamentals

The third storyboard shown to participants concerned a mechanism for quickly searching through open windows. This was perhaps the best received of all the storyboards, with five out of six participants highly praising its value. Participant 1 described the idea as “taking something that exists and powering it up with more usability and more functionality”, adding that it is “not something that’s completely alien to people” and “would just sort of become natural to people’s workflow”. Other participants said that it would be a “really nice feature to have”, would be “very useful”, and could “really help” to “speed up searching”.

In probing participants as to whether the proposal actually solves a real problem, the responses strongly indicated that it did. Participant 4, describing a real world situation in which such functionality would be useful, noted that this provides a “direct way to go between” open resources, and to “search between them for what I wanted”. “I’m sure Jacob could find what he wanted in the end, but it’s definitely going to save him time”, Participant 6 remarked.

Besides directly navigating to known resources, participants further identified that the proposed behaviour could be useful in finding resources without remembering their exact details. “There are a lot of times where you read something, but shortly after you forget where you read it and how you found it”, Participant 2 commented. Participant 3 expressed a similar sentiment, commenting that they often “remember a certain phrase when I’m working and I’m like “Where did that come from?”, so this would help ... this would speed up searching so much”.

Participants also found value in the general ability to search through heterogeneous resources easily, for instance in being able to quickly search through video, PDFs, and text documents simultaneously. Participant 4 particularly highlighted the value of being able to search through video, indicating that it “would be great” if users could see information like “that guy said that at 3 minutes into his video” in search results and “jump straight to that point”. The participant did question how search through images and video might work, though, noting that you could “fall back to searching the title or file name” but that this would not be nearly as useful. Employing optical character recognition and forms of machine learning to aid this process does not seem entirely unreasonable, particularly when results can be pre-computed.

Many participants also questioned how this functionality might interact with spaces.

While it was generally agreed that the search functionality should be scoped to the contents of spaces, some participants noted that the use of spaces may reduce the value of this kind of functionality to some degree. Participant 1 said that when “splitting things into workspaces”, “the need to search may be reduced because you have less to look at”. Participant 4 similarly said that “if you’re having to do that . . . you probably have too much in one space”. To contrast these viewpoints, however, Participant 6 noted that “this would be useful with the spaces as well” as “they might get cluttered after a while” and “could very easily get messy”. “You need some shortcut way of quickly seeing what you’ve got in a space”, they said.

Participant 5 also noted that “if this is within a space, then all your resources might be pretty similar, so I don’t know how specific your search criteria would have to be”. This is a good point: if there are lots of resources on closely related topics within a space, formulating a search query to disambiguate your intent might be difficult. We suspect that users may get better at doing this over time, but this may be an area worth examining more closely in future work.

Doubts

Although most participants saw significant value in the proposed solution, Participant 5 once again proved to be the exception. “If I’m in a space and I need to find a document”, they said, “I generally know, like, what program I’m looking for and what the document looked like”. “I might not remember what I actually searched for to get it”, they followed, “I just remember what the distinguishing features look like”. This is an interesting perspective. In some cases, perhaps the title or contents of the resource are simply not what users will have in mind when they wish to locate a particular window. In such cases, our proposed search functionality is not likely to be useful.

Touching on the motivations behind the storyboard, the participant also emphasised their belief that tabs better facilitate fast switching between resources than windows. “I use tabs a lot, so I can remember where stuff is without other cues”, they said, adding “I don’t have to figure anything out”. This mirrors a sentiment the participant expressed earlier in the session, reasoning that “if I’ve got tabs open, I can sort of organise things in terms of tabs . . . I don’t really see how that’s going to happen here”. Window search was only intended to provide part of the solution to such problems, and so such a line of enquiry cuts to the core of how operating systems allow users to interact with programs.

In discussing potential solutions, the participant even went as far to suggest that “one window, with each program running in a tab” could resolve this issue. This speaks to the immense value that the participant places on tabbed interfaces. Although operating systems today provide a myriad of ways to switch between windows — for instance, Mission Control, →, and ` on macOS — perhaps

additional mechanisms could be provided such that windows can be considered a more adequate alternative to tabs in more situations. A manner by which users can map certain windows to ordered key combinations — such as $\wedge 1$ and $\wedge 2$ — for example, could fill this role.

6.2.10 Window Search: Suggestions

Participants also had some suggestions for how the proposed solution might be improved or extended. Amongst these was the suggestion that some mechanism should be provided to navigate between search results. “If I search for something and five things come up, I want to be able to jump between them all”, Participant 4 commented, “having some way to tab between the results would be super super great”. While this was not envisioned within the original design, it makes complete sense. Such functionality allows search to continue beyond the selection of a single result, empowering users to browse results as they see fit to find what they’re looking for.

“It’s basically like having links without having someone actually make the links for you”, Participant 4 described. Continuing along this line of thought, the participant suggested that “it would be good if there was a way to actually build links from within the search UI . . . I’d like a way to instantly be able to have a link so I can remember that I can go between these things”. This is an interesting suggestion, and may be worth considering in future iterations of this design idea.

Many participants also asked questions around how much control the system would grant users in terms of advanced queries. “Is it just words, or can I use regular expressions?”, asked one participant, while another asked whether users could “put in multiple terms” or “use logic to filter your queries even further”. Participants also suggested “filtering by file type” or recency, and being able to “self-tag things while you’re working”. Some kind of advanced, expressive search functionality is clearly desirable, though is beyond the scope of our work.

6.2.11 Exploration Mode: Fundamentals

The final storyboard shown to participants concerned Exploration Mode — an environment specifically designed for exploratory activities. On the whole, participants were fond of many of the general ideas underlying the storyboard. Participant 2 described the mode as “an efficient way to spend your time looking for stuff”, while Participant 5 declared it as almost a “more natural way of browsing”. Many participants, however, kept an air of healthy scepticism in viewing the somewhat exotic system being presented. While a number of properties of the solution were identified as highly appealing, participants were also quick to raise potential issues and highlight less desirable features.

Participant 1 particularly emphasised value in the transient nature of exploration within the mode. “We’ve all been on search tangents before where we . . . just want to top up our understanding”, they noted. “Being able to go through that, feel comfortable understanding that, then pressing a button and immediately returning to where you were before . . . seems like just a natural contribution to people’s workflows”, they added. That is, the participant views the way in which the mode allows users to browse through a large interconnected information space to learn or find something, and then immediately return to their starting point — bringing back whatever information they want to keep around — as highly valuable.

Curiously, the kind of “tangents” that this participant describes Exploration Mode as particularly fit for seem highly related to the “tangents” and “trails of thought” that Participant 5 described as potentially problematic in polluting and muddling up history. Although we have not defined how Exploration Mode interacts with spaces or history at all, it seems that this kind of ephemeral browsing may be an excellent fit in resolving the history issues previously described. Participant 5 hinted at this in suggesting that it “would be kind of cool if you could just like jump straight into a new space [from Exploration Mode]”, and that being able to say “let’s take this into a new space” could “make it more natural to create spaces and trails of thought”.

6.2.12 Exploration Mode: Resource Suggestion

One of the core features of Exploration Mode is its resource suggestion behaviour. Upon entering the mode, the system directs focus onto a single resource, and machine intelligence is employed to suggest related resources to the user. The great potential of such functionality did not appear to be lost on participants. Participant 3, for instance, emphasised the value in not being “confined to stuff that you’ve sort of already seen before”. Participant 1, meanwhile, imagined that “maybe for maths or something, as a really simple example, you could have a concept, and then all of its applications”, noting that this could be “really powerful” and “really help”.

Many participants had concerns about the volume and quality of these suggestions, however. “In a modern document or webpage or something, realistically you could link from that to a lot of places”, Participant 1 outlined. “I’m mindful of being bombarded with information that I don’t care about if I enter this mode”. Participant 5, meanwhile, remarked that “it’s rare, like, that I look at suggested stuff . . . it’s generally not intelligent enough to be useful or specific enough”. More dramatically, Participant 4 stated that they “would not find this screen that useful”, citing concerns of “having a scary, huge mind map just explode on my screen”. This sentiment was mirrored somewhat by Participant 3, who noted that “there’s a bit too much going on”.

Relatedly, Participant 2 cautioned that “you don’t want your vision to be too im-

paired with just, like, a lot of information at once”. “You just want to kind of look and see what you might need or want”, they concluded. In the context of this caution, the participant also questioned the proposed interface for the functionality, warning that “everything might appear too small”. Participant 4 similarly asked whether the suggested resources would “hold content or just images or icons?”. Participant 3, meanwhile, called into question the new way in which way the mode displays these connections. “There are like three different things highlighting links, which might be a bit overboard”, they said.

Following this topic further, Participant 4 questioned the entire motivation of using a mode to show related resources. “You don’t really need another system to deal with that”, they said, suggesting that the typical links displayed on resources could encapsulate this kind of information. “I just think that with how links are incorporated into the main system, it might not even be necessary to have this kind of system”, they summarised. Perhaps the introduction of ‘dynamic linkbases’ within the system could accommodate such behaviour.

In raising the question of how links *to* a resource might be displayed within such a model, the participant commented that they “see some use” in Exploration Mode for this purpose, but that “the negative of having another screen” might outweigh the benefits. As an alternative, the participant suggested a “little pop-up box at the side” indicating “by the way, these articles link to here”. This suggestion of integrating the behaviour into a single model is quite compelling, though removes the transient exploration behaviour seen as so desirable by some participants.

One possibility to work around this issue is to reshape Exploration Mode into a more lightweight construct. Instead of having a transformative effect on the entire desktop environment, the mode — or some series of ephemeral actions that replace it — could exist purely for the purposes of transient exploration and transpointing. One could imagine, for instance, that some windows are shown to look somewhat ghostly in nature to indicate their transience. Meanwhile, resource suggestion could be delegated to the normal mechanisms of link display, with such links only being shown when demanded by the user.

6.2.13 Exploration Mode: Transpointing

An area that participants universally praised when it arose in conversation was link previewing and transpointing. “I really like the idea of previewing a link”, Participant 2 remarked, “there are a lot of times you might waste time by going onto another resource that you think is relevant, but it’s not”. Touching upon transpointing, the participant added that “seeing the context in which it’s linked to whatever you’re reading — I think that’s really useful”. Participant 5 similarly said that this behaviour seems “really useful”, while Participant 4 said they “definitely want that”.

Following from the discussions of a ‘lightweight’ Exploration Mode in subsection 6.2.12, we also considered the possibility that transpointing might better serve users if it was available outside of any mode. In asking Participant 5 about this, they responded that allowing users link preview and transpointing functionality without a mode switch “would completely make more intuitive sense” to them. Elaborating on this, the participant described how this way of doing things could be materially less confusing, citing the reduced mental burden from making the behaviour unconditional of the mode of operation.

Another point of discussion in this area concerned the potential issue of visual noise in transpointing, raised by Participant 6. “Would you make sure there aren’t too many lines between the two documents?”, they asked, commenting that “you don’t want like a spiderweb of lines”. This is an excellent point, and is certainly something that ought to be considered in the design of any practical transpointing system.

6.2.14 Exploration Mode: Extensions

As Exploration Mode is the most open-ended of the proposed solutions, it should be no surprise that participants had plenty of suggestions for how it might be improved. Chief amongst these were suggestions surrounding how the system should choose resources to show to the user, attempting to address the issues of the quality and quantity of information being displayed. Most all of the suggestions in this category concern trying to direct the exploration in some way. “It might be nice to enter this mode, but for a specific reason”, Participant 1 reasoned.

Throwing out possible solutions, the participant suggested that “maybe if it was before you enter Exploration Mode, you highlight a section of text”, or that “maybe you do it after the fact”, entering Exploration Mode and then filtering the suggestions. Another idea, suggested by participants 1 and 3, proposed tuning the suggestions by discarding those that weren’t relevant. “Maybe it would be helpful to be able to like, remove them? ... If one of the resources isn’t useful, you can just throw it away”, Participant 3 said.

Participant 6, meanwhile, imagined that users might “configure it you just like an ordered list of what would be prioritised above what”. This idea was also brought up by Participant 1, who put forwards that maybe the resources “would have some form of priority”. Continuing this thought, Participant 1 presented the idea that related resources might be shown under some kind of categories, noting that “maybe there’s an ‘Other’ section” as a part of this. Having reflected upon this idea in the context of the other feedback regarding this feature, it seems clear that categorised, compact lists are likely to be more useful to users than space-consuming visual previews. The interface probably ought to be closer to Google

search results — with a title, snippet, and perhaps an estimated reading time — than annotated thumbnails.

Another suggestion made by participants concerned the ability to redirect the focus of Exploration Mode onto other resources. “Would you be able to like, change focus onto one of those [recommended resources]? Like, change the middle window and then there’s more links?”, Participant 3 queried. On this topic, Participant 6 commented that they “think it would be useful to do both ... right click it and say ‘Open out of Exploration Mode’ or ‘Carry on exploring through this one’”. Providing the option to perform either action here certainly seems like a good choice. Sometimes users may wish to bring a resource out of Exploration Mode for normal usage, but other times they may wish to continue their transient adventure.

Related to this topic, Participant 6 posed the question of whether Mia — the protagonist of the storyboard — would have “some way to save resources she particularly liked for later viewing”, as “there’s no guarantee that it’s going to find the resources she found last time if she tries again”. This is a reasonable suggestion, and relates somewhat to previous conversations around being able to retain relationships from the window search functionality. Perhaps a ‘save for later’ action within Exploration Mode could save links with specially highlighted marks in the marker bar, or maybe add resources to a list of bookmarks to be reviewed later.

In a related discussion, Participant 6 proposed that “it could be cool” to be able to make a connection from Exploration Mode into a “permanent link”. “If you like that connection to that document, you could, like, click on that and it turns into like a chain or something”, they described. Further miscellaneous suggestions from participants included being able to “see who else has put emphasis on a particular link to some resource” (somewhat related to read wear, Hill et al., 1992), being able to have “separate workspace[s] just for exploring stuff”, and having a “search bar that searches the whole Internet” within Exploration Mode.

6.2.15 General Comments

Amongst the discussion of storyboard-specific issues, conversations with participants occasionally moved to the vision as a whole. At the end of each session, participants were asked explicitly for their thoughts in this area, concerning a system that combines the ideas within all four storyboards. From this, it became clear that while participants were generally positive about the direction of each individual storyboard, they were far more excited about the potential of the system as a whole. “Individually they’re all quite nice concepts”, Participant 1 commented, “but it gains its power when you really start mixing them together”. “It just seems like it would be a much easier way to do research”, Participant 3 summarised.

In considering how the system might be practically useful in the real world, Participant 1 identified links and linkbases as the “backbone” of the entire system, noting that “this seems like a system prime for having different levels of publishing”. In “departments, businesses, the wider world”, they imagined, “linkbases are your contribution”. Other participants, meanwhile, identified potential applications in other areas. Participant 2, for instance, emphasised the potential of the system in education and learning, while Participant 3 highlighted the system’s applicability to research in general.

In drawing parallels between the proposed system and what exists today, Participant 1 also made a curious comparison between our project and Chrome OS — Google’s operating system built around the Google Chrome Web browser. The participant remarked that both projects aim to bring “browsing habits and concepts to a higher level”. While Chrome OS does this by “making the browser the OS”, our proposals try to “pull what the browser does into the OS . . . making that manageable and adding functionality to it”. This is an insightful comparison. The Chrome OS announcement remarks that “the operating systems that browsers run on were designed in an era where there was no web”, and that — rectifying this — Chrome OS attempts to “re-think what operating systems should be” (Pichai and Upson, 2009). In many ways, our work shares these same goals.

With regards to potential issues, some participants did note that the drastic nature of the changes we propose may be problematic. “It’s just about how easy it is to learn”, Participant 6 commented, “not everyone is going to be able to sit down for half an hour and have someone explain it to them”. On a similar thread, Participant 4 questioned how the system might gain adoption, noting that users “get mad over very minor changes, like moving the start bar” and so that “revolutionise the whole design of everything” might be a tough sell. These are, of course, reasonable points. Practically implementing a system of this kind would be difficult, but gaining mind share and establishing an ecosystem such that the true benefits of the system can be realised are likely to present even greater challenges.

6.3 Critical Analysis

Having analysed and discussed the feedback from participants in some depth, it is important to recall that the entire analysis is seeded by only about four and a half hours of conversation total, with a sample of only six people. As such, responses reflect possible perspectives of users, but may not represent prevailing opinions within the wider population of knowledge workers.

It is also worth keeping in mind that the potential solutions evaluated here are largely orthogonal in nature, and so that participants have only been shown a single solution in each problem domain. In order to reach conclusions about what

solution might be ‘best’ for a particular problem domain, a number of alternatives ought to be produced and evaluated.

Further, it is likely that in some areas, participants’ descriptions of their preferences do not reflect how they would actually behave when faced with real implementations of our proposals. As a result, we believe that studies examining practical user behaviour are necessary to properly evaluate some design decisions. Consider our ‘spaces’ prototype, for instance. If we are to assume that spaces correspond roughly with “working spheres” (as described in our literature review), empirical research in this area suggests that “there should be recognition that many applications are shared among working spheres” (González and Mark, 2004, p.119). The success with which our backtracking proposals handle such a situation was not considered by participants in examining our storyboards, but could likely be clarified with relative ease in more practical circumstances.

6.4 Summary

In the previous chapter, a number of potential solutions to fundamental open hypermedia design problems were devised. By discussing these proposals with users and analysing the results in some significant depth, this chapter has established a greater appreciation for the strengths, weaknesses, and open questions that surround these. On the whole, our findings validate the underlying motivations behind our design ideas, and additionally suggest a number of opportunities for possible improvements. The prototype evoking the greatest number of comments in this regard was our Exploration Mode proposal, which seemed to excite participants, but also to attract generally mixed feedback.

With regards to the broader research goals of this project, our analysis shows that many features of open hypermedia systems were perceived by participants as having great potential value. This provides further evidence for the unrealised utility of such systems, and may help to clarify their practical worth. Further, the generally positive nature of the feedback indicates that in spite of the large number of open questions in this space, many of the fundamental problems discussed are not insurmountable. We believe that further work in this area could yield open hypermedia systems that are significantly more usable than those of the past, and thus can be seen as more serious alternatives to the WWW and WWW browsers.

Chapter 7

High Fidelity Prototyping

This chapter centres around the design of a single high-fidelity prototype. Selecting a few, fundamental proposals from our storyboarding process, we further develop and refine foundational elements of these in accordance with the feedback from the previous chapter. Focusing in particular on the areas of link behaviour and history navigation, we produce a more polished representation of these ideas, better communicating their details and potential advantages. We then proceed to use this representation to briefly evaluate our proposals.

7.1 Prototype Design

7.1.1 Form

Prior to any consideration of specific design details, it was first necessary to determine what form our prototype would take. From discussions of participant feedback in the previous chapter, it is clear that an interactive prototype would be ideal for this purpose, allowing users to get a feel for the dynamics of the system in practice and to explore at their own pace. In considering several possibilities to this end, however, it was decided that such an endeavour was not a suitable choice at this stage. Our proposals allow for a large space of possible interactions, and do so across program boundaries in a way that would be time consuming to replicate. At this early stage of the design process, we believe that such an undertaking would not be an effective use of our time.

Instead, our prototype takes the form of a short video, guiding viewers through a realistic-looking mockup of the proposed system. This still allows for many of the dynamic properties of the system to be communicated, but is considerably less time consuming to produce than a fully interactive prototype. Further, to present our solution within a familiar, contemporary environment, we chose to ground it within the macOS operating system. Using assets that ship with macOS and a handful

of resources from the Web, we utilised Apple’s Keynote presentation software to produce the video. A frame of this is shown in Figure 7.1, while the full video can be found at: <https://www.youtube.com/watch?v=SAYL9YzYOWU>.

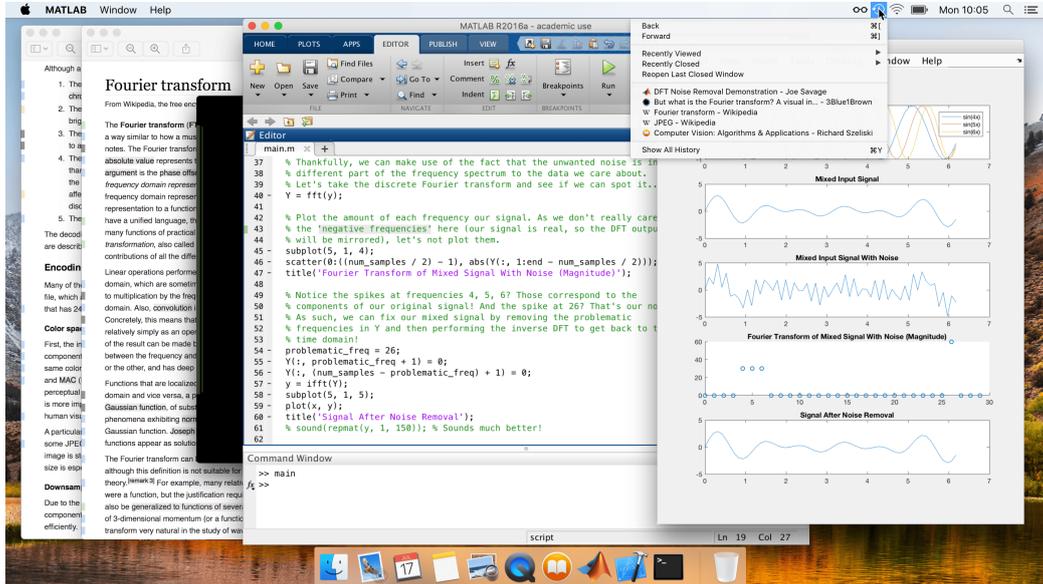


Figure 7.1: A frame from our final high-fidelity prototype.

7.1.2 Narrative

To scope the feature set of the prototype, we chose to concentrate on the ‘link behaviour’ storyboard from previous prototyping with the addition of basic in-space history navigation mechanisms. The prototype thus includes the inline display of links, the marker bar, link navigation, link creation, linkbase selection, and primitive backtracking behaviour. This functionality is highly foundational in nature, yet is sufficient to illustrate some of the key value propositions of open hypermedia systems.

With this set of features established, a narrative was constructed to guide users through these. As this narrative underlies the entire prototype, careful consideration was put towards its design. Ultimately, the scenario of a student making use of the proposed system to better understand JPEG compression was selected. This is a research-based task that has significant potential depth, yet is relatively accessible and comfortably permits heterogeneous resources and multiple perspectives.

Our prototype is then structured such that it introduces the novel properties of the proposed system through this narrative step by step. To begin, a desktop environment is shown containing a single window. This window displays a largely textual

document with links overlaid from a single source, mirroring the kind of behaviour that is typical on the WWW today. As we expect this behaviour to be relatively familiar to viewers, this serves to introduce our proposals around inline link highlighting and the marker bar. From this point, link navigation is demonstrated. As navigation through multiple windows may surprise some viewers, this behaviour is gently introduced and normalised over the course of a few examples.

Linkbases are then introduced at around the 1:12 mark. Specifically, a new linkbase is enabled for the first time in the video, showing a different perspective on the information in view. This progression then continues throughout the remainder of the video, building up to first-class support for heterogeneous resources — including link navigation both to and from a video file — and the interaction of all these components with the history mechanism, first shown at around 3:40. Critically, all of these components are presented in the context of the overarching narrative, and are shown as natural, useful steps to take given the goals of this particular user.

7.1.3 Details

Unlike storyboards, high fidelity prototypes present a fairly detailed picture of a proposed design idea. As such, the production of such prototypes require the designer to describe not only the highest level features of a design idea, but also a number of finer details. As this prototype is only the second iteration of our ideas in this area, many such details — such as the exact shade of grey selected for link highlighting and the precise positioning of marks in the marker bar — have been determined by gut feeling rather than careful consideration. In some areas, however, more deliberative design decisions were made.

Our marker bar design element, for instance, was only loosely defined within its original storyboard. Though its general form factor was made clear — a vertical bar on the left hand side of the window containing coloured marks to show interaction opportunities with links in various linkbases — many of its specific details were left open to interpretation. Developing the idea further, we have chosen to display marks at approximately the same width as the macOS scroll bar, and with relatively subtle, unassertive colours. This is intended to provide a reasonable balance between readability and discoverability — a trade-off discussed in length in Chapter 5 — and is shown in Figure 7.2. While the manner by which mark colours are determined for each enabled linkbase is left unanswered within the prototype, this mapping is assumed to be consistent across all open applications. In addition, we reserve one specific colour — a pale blue — to mark links distributed with a resource, as opposed to those that are overlaid by any other linkbase.

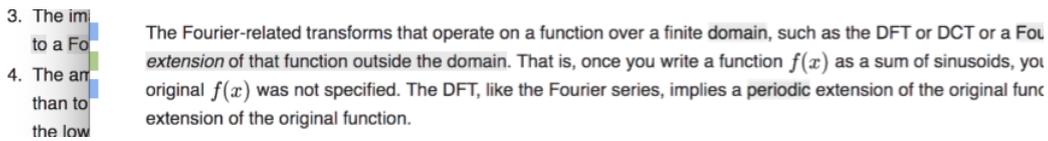


Figure 7.2: A still from the prototype showing the marker bar.

Note, however, that our prototype makes no effort to distinguish between ‘internal’ links *within* a resource, and ‘external’ links *between* resources. In retrospect, not distinguishing between these is likely an error. The prototype shows internal links to have different behaviour to external links — jumping to the linked section in the same window rather than opening a new window — and this difference in behaviour should almost certainly be reflected visually. Similarly, navigation through internal links should likely have a different effect on the current chain of history than that through external links, but such details are not explored here.

We also leave the issues of link display and navigation surrounding overlapping links unaddressed. We do, however, propose a solution for the display of the marker bar when multiple links are present on the same line. If these links are from the same linkbase, we display a single mark on that line, coloured as appropriate for that linkbase. If the links are from different linkbases, however, then we propose that this mark is displayed in black — as shown in Figure 7.3. Black marks stand out in the bar, are easy to notice at a glance, and visually are intended to show the mixing of multiple linkbase colours in a subtractive colour system, like with paint. This design decision is not explicitly discussed within the video, however, and so may confuse viewers.

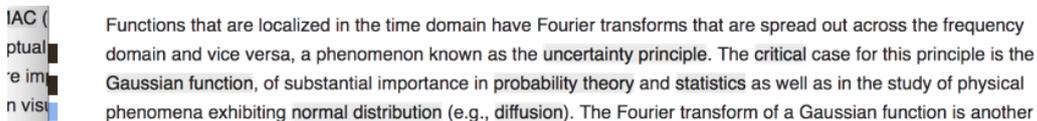


Figure 7.3: A still from the prototype showing marker bar display with multiple links on the same line.

Our prototype also makes no effort to indicate which marker bar marks are associated with hidden links. This too is likely an oversight, as the marker bar is the only way that the presence of such links can be shown. The prototype does define a mechanism by which the specific locations of such links can be viewed, and the links themselves navigated, though. Hovering over a mark within the marker bar highlights all corresponding links inline in their linkbase colour, while clicking a mark pins all such links in place. Hidden links that have been pinned in place can then be followed just like any other visible links.

The way in which our prototype highlights links inline, meanwhile, remains much the same as within the storyboard. The only significant change is the addition of a hover effect, in which the neutral overlay changes colour to the linkbase colour associated with the link. Link creation has also remained much as originally envisaged, as can be seen in Figure 7.4, with the process of selecting a destination for a link remaining largely ambiguous. The prototype shows a hollow circle as a possible destination anchor, and this is relied upon without explanation of how it was specified or whether it is possible to create links in other ways.

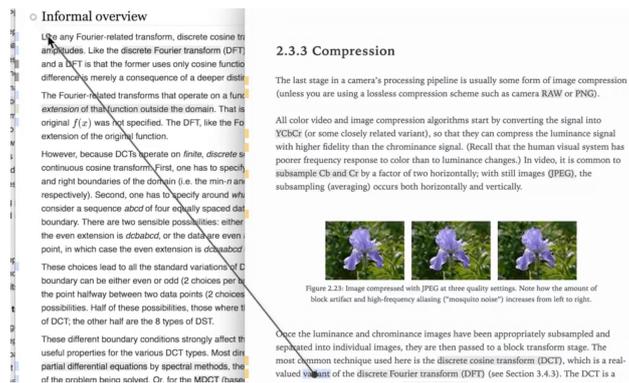


Figure 7.4: A still from the prototype showing link creation.

In accordance with the feedback from the previous chapter, our prototype also rebrands linkbases as ‘perspectives’. This change in terminology is supported by our choice of the perspectives menu bar icon: a pair of glasses. Accompanying this in the menu bar is a history menu bar application, as suggested in the previous chapter. Unlike our original storyboards, the menus spawned by each of these have been more fully developed and are displayed in detail within our video. The design of the perspectives menu, for instance, can be seen in Figure 7.5.

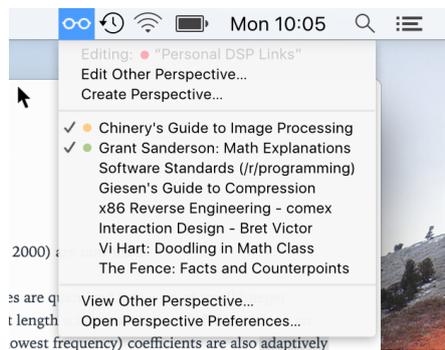


Figure 7.5: A still from the prototype showing perspective selection.

7.2 Evaluation Procedure

To evaluate the prototype, we recruited the same set of six participants that took part in our storyboard evaluation. Although these participants are not likely to provide a fresh perspective on the core ideas underlying the prototype, they come equipped with a level of familiarity that is difficult to find elsewhere. As a result, these participants are likely to get up to speed with the ideas within the prototype more quickly, and may provide deeper insights into its merits and drawbacks.

Unlike the previous two studies within this project, brief comments were obtained from participants using an online survey. Participants were first briefed on what they were expected to do, along with their rights, and reminded of the surrounding context. Following this, they were asked to watch the six minute video prototype, and were then presented with five questions. These asked participants about their general thoughts on the prototype, things they liked, things they would like to be different, and any questions or suggestions they had, in addition to providing a space for other, miscellaneous comments. This procedure is compliant with the University of Bath's ethics guidelines.

Qualitative analysis of this data was then performed using thematic analysis (Braun and Clarke, 2006), as motivated previously in section 4.1. As the body of data under analysis in this study is far smaller than the previous studies undertaken within this project, many fewer themes are to be extracted than before. As participants in this study are already familiar with the general ideas and surrounding context of the project, though, the feedback from participants still ought to be highly informative in evaluating our prototype.

7.3 Evaluation Results

The results of data analysis follow. As in previous analyses, this consists of descriptions of each theme accompanied by relevant discussion. The raw responses from participants can be found in section A.3.

7.3.1 General Sentiment

As a whole, the prototype was well received by participants. Describing it as “very well executed”, “coherent”, and “useful”, participants made clear that they admired many aspects of the prototype. Participant 2 said that they could “see a clear use-case for this form of system”, noting that they “feel it does a much better job than any current system for linking documents” and that they “wholeheartedly believe this to be a very beneficial system”. Participant 4 similarly expressed that they “really like this idea for learning and seeing what other people thought was useful”. These responses are highly encouraging, and suggest that the prototype

was successful in describing a coherent and usable interface to an open hypermedia system.

In describing what they like most about the prototype, participants clearly identified value in the unique capabilities of the system. Even those participants that showed a large degree of scepticism in storyboard evaluation seemed to see a great deal of potential value in this more detailed representation of the system. Participant 2 admired the “simplicity” and “versatility” of the system in handling various types of resources, for instance, while Participant 1 highlighted the value of perspectives in showing “different pathways depending on the requirement and knowledge of the user”. This feedback mirrors the general sentiment from storyboard evaluation, indicating that the fundamental ideas behind open hypermedia — and perspectives, in particular — have great potential.

7.3.2 Links

Inline Highlighting

In evaluating our proposals for link display in the previous chapter, feedback was mixed. The storyboards only vaguely described the envisaged approach to inline highlighting, and a number of participants questioned the utility of highlighting links in this way. This prototype described how links are displayed in much more detail, and did so to greater praise. Participant 5 noted that they liked the “unobtrusive interface” on the whole, particularly highlighting that they were a “big fan of the subtle highlighting to prevent distraction while reading”. Participant 3 similarly admired the “very subtle” link highlighting.

Using a grey overlay is unlikely to be suitable in all situations, however, as pointed out by Participant 2. “The generic grey highlighting . . . could be problematic depending on the document”, they said, “what happens if it has a grey background?”. This is, of course, a valid point. This problem is not isolated to inline link highlighting, though. In fact, this is a problem that all cross-content and cross-program UI elements must deal with in one way or another. The scroll bar and marker bar are equally subject to these issues, for instance. In the case of link highlighting, we suggest resolving this issue by using a different shade to highlight links depending on the background. This is similar to the way in which the scroll bar changes colour based on the content it is scrolling over from macOS 10.7.

Navigation

In depicting the proposed hypermedia system in a realistic setting and being used for a concrete task, the prototype also clarifies how resource traversal is intended to work in a system of this sort. In this area, a few participants reiterated potential issues flagged up in the previous chapter. Two participants, for instance, noted

that the number of open windows might become a problem in certain situations, while another pondered exactly how history navigation should interact with typical window manipulations. These remain important issues, but as discussed in the previous chapter are best addressed through other studies involving user interaction with systems of this sort.

One observation regarding behaviour not shown within the original storyboards came from Participant 5, who commented that it “might also be useful to be able to open internal links as a new resource/window”. This refers to the asymmetry between the way that the prototype handles ‘internal’ links *within* a resource, and ‘external’ links *between* resources. As discussed in subsection 7.1.3, internal links are shown to jump to the linked section in a single window, while external links open a new window. Though natural from one perspective, this asymmetry does feel somewhat unusual. Certainly, some way to open internal links in a new window seems desirable.

In addition to remarks surrounding the destination window in link navigation, many participants also made comments around the destination program. Participant 1 said that it would “be useful to know what the sort of link is that I’m clicking”, adding that “I want to know what I’m about to open as I may be somewhere where I can’t listen to sound for example”. “The type of content you’re navigating to shouldn’t be a surprise you have to react to”, Participant 6 declared. Suggestions for how this might be achieved included additional use of colour, expression of media type within the marker bar, and use of tooltips. It may also be the case that link preview functionality — like that discussed in the previous chapter, or the three-finger tap trackpad action to preview links in Safari on macOS — may help in resolving these concerns.

Participants also raised the question of what should happen when the user has no viewer program installed with which to open a resource. “What happens if I click a matlab link and I don’t have matlab? Is it worth even showing the link?”, Participant 1 asked. Participant 4 voiced similar concerns, noting that “there should be some coping mechanism or an alternative” in dealing with this. While these issues are important in isolation — and boil down to much the same question as “What should the OS do when I try to open a file with an unknown extension?” — they also speak to the larger ecosystem questions around heterogeneous information spaces.

In the same way that heterogeneous files are distributed today (e.g. via email), we expect that communities using a system such as this would settle on a handful of common formats to fill various roles. Though it is to be expected that there will be some variation in these formats over time — allowing for competition and improvement — the overall level of fragmentation (along with other properties such

as format openness) ought to be largely dependent on the cultural norms within these communities. For this reason, the initial viewer programs and formats that integrate with the system are particularly important, setting expectations for all other formats and viewer programs to come.

Similarly, the authoring tools that are made available for these formats are critically important in setting the cultural norms in terms of content (Victor, 2012). With the right tooling, the specialised nature of heterogeneous resource types could allow for far richer content than is available on the WWW today, making better use of the computational medium (Victor, 2013b).

Creation

Participants also had a number of questions surrounding link creation, an area not particularly well defined within the prototype. Participant 1, for instance, asked how links are created to resources that are obscured by other windows, how links to specific sections of videos are created, and when the possible link destination indicators shown in the prototype are. Participant 2 similarly questioned where exactly the link created in the prototype leads to in the target document, and how links could be removed. These are reasonable questions, and ought to be addressed in more thorough design work in this area.

Participant 2 also queried how link creation interacts with perspective sharing. “I create that link, I send X person my perspective, how do they then get access to that document? Does it then have to be uploaded somewhere by me? ... Would all resources have to be available by networks?”, they asked. This is a good question. As the prototype makes no effort to distinguish between local and remote resources, it’s not entirely clear whether the link that is being made is between remote resources, local resources, or some mix of the two.

In a more practical system, it may be desirable to provide some way to distinguish between these cases. The title bar of a window, for instance, could indicate whether a given resource originates from a local or remote source. Following this thread further, more advanced systems may also provide some way to examine the URL of a resource, or for the source of a given document to be verified for authenticity or security reasons. Returning to the issue of how links to local resources should be handled in sharing perspectives, however, one could imagine a number of reasonable solutions. One possibility is to notify the user at the point of sharing that the perspective contains purely local documents, and provide an option to automatically upload such documents automatically to a pre-configured server for the user.

7.3.3 The Marker Bar

Fundamentals

Through the addition of realistic context, colour, and dynamic behaviour, the marker bar is very much brought to life by the prototype. As in the storyboard evaluation, this UI element was largely well received. Participant 6, for instance, commented that the bar “works well to bring emphasis to links without obstructing the actual content of the particular window”. Naturally, however, participants had a number of questions around details that were not made clear within the video.

“What happens in multicolumned text wrt links in the marker bar?”, Participant 1 asked. “Maybe a marker bar for every column but then that would be ugly, not sure how you’d address this”, they continued. As it happens, the prototype does actually show a multi-columned textbook as an example resource, and uses only a single marker bar for this. The reason for this design decision was exactly as the participant described — two marker bars didn’t quite look right. It is easy to imagine why a dual marker bar layout might be preferable, though, and this may be worth exploring further in future design work.

More fundamentally, many participants commented on the use of colour within the marker bar. Participant 2 noted that they “like the simplicity of using colours to label what perspectives lines have links on”, but expressed some concern around recalling which colours map to which perspectives. Participant 1 and Participant 5 expressed similar sentiments, with the former writing that “the number of different colours could become slightly confusing”. This is a reasonable concern — how well does this approach of using colours alone to distinguish categories of links scale? Participant 5 proposed that “an easily accessible key” could go some way to improving this situation, while Participant 2 suggested using tooltips to resolve any ambiguities around link properties.

These were not the only queries raised by participants around marker bar colouring either. Participant 2 also questioned how the mapping from colours to perspectives was to be performed at all, noting that transient mappings could cause additional confusion. Further, a number of participants brought up the recurring question of how the marker bar should best show lines containing links from multiple perspectives. While the prototype uses black marks for this purpose, this was not discussed in any detail, and so appears to have been missed by — and in some cases, to have confused — participants. Taking these questions in aggregate, it is clear that while our vision of the marker bar has potential, it is not yet ready for practical applications. Further work is necessary to ensure that this component can best support users in achieving their goals.

Non-textual Media & Hidden Links

The prototype also shows the use of the marker bar in highlighting less conventional links, including hidden links and links from non-textual sources. While the general ability to have numerous links across heterogeneous resources was well admired, a number of participants raised questions in this area. Participant 1, for instance, noted that it was “neat to have the links appear at certain times in the video”, but was curious as to how this might extend to other forms of media, asking “could this potentially work with audio and other time contextual sources?”. This is certainly an interesting question, and may be worth exploring in future work.

Following a different thread, Participant 6 noted that the marker bar helps to “provide a constant view of all links on the current window”, but that this “seems to get lost when applied to video content as marks only appear whilst there is a link”. This is a reasonable point. While the current frame of a video seems analogous to the current view of a text document in some respects, something about the serial, transitory nature of video makes it different in others. The participant then went on to suggest that “something could be added to the seek/progress bar to show the same kind of information i.e. highlighting of the progress bar on sections where a link is present in the video”, a seemingly logical solution.

In order to follow links within videos, the prototype relies on interaction with the marker bar. This occurs in exactly the same way as the prototype allows for the exploration of ‘hidden’ links that have not been highlighted inline within textual documents: links can be highlighted inline by hovering over their associated marks, and then pinned in place for navigation by clicking. Some participants expressed confusion around this behaviour, however. Participant 1 noted that it “seems slightly confusing” to have to follow such a procedure, while Participant 6 said they they were unsure if they had “grasped the concept or use” of this. This may speak to a lack of clarity around when links are hidden and when they are not, or may indicate that the prototype communicated these ideas poorly or simply that this behaviour is confusing.

7.3.4 Perspectives

The final area of the prototype that attracted significant comment from participants concerned perspectives. As was the case in evaluating the original storyboards, this functionality and the user interface proposals surrounding it were largely well received. “Toggling between perspectives appeared quick and simple allowing attention to remain on the workflow”, Participant 6 commented. “I like that perspectives can be individually toggled with multiple being active at a time using different marker colours”, they added, indicating that this would be highly valuable in avoiding “arduous” switching back and forth between perspectives.

In terms of suggested improvements, Participant 3 noted that some effort should be made to ensure that the list of perspectives shown in the menu bar application does not become too unwieldy or confusing, noting that “descriptive names for them would probably help”. Participant 5, meanwhile, made a feature suggestion around the display and filtration of links from various perspectives. “Could there also be a way some way to quickly display all links in a document, perhaps filterable by type?”, they asked, describing that this could aid the user in quickly flicking through the useful links from certain viewpoints. Both of these appear to be reasonable suggestions, which may be worth exploring further in future work.

Sharing & Discovery

Perhaps unsurprisingly, a large proportion of the feedback from participants in this area concerned sharing and discovery, an extremely important topic that is left almost entirely open by the prototype. As mentioned in the previous chapter, such considerations are largely out of scope within this project. Not wishing to ignore these issues entirely, though, we will briefly discuss a few of the key issues raised by participants here.

Participants 2 and 6 brought up the issues of the scope and permissions as they relate to perspectives. These issues include and relate to: sharing perspectives with a restricted group of people (e.g. within an organisation, perhaps with password protection), the default storage of perspectives (e.g. local, network-local, remote), how new links are managed in shared perspectives (e.g. through local modifications or submission of suggestions to a central source), and how perspectives are to be properly managed and vetted before being shared with others. These are important questions that ought to be answered in future work.

Participant 5 also raised a number of questions in this area. These included how the perspective ecosystem should be bootstrapped, how perspective discovery and ranking should work, and how “bad links” will be policed. With respect to the central discovery mechanism, this is fundamentally a search problem to which one can imagine a number of solutions. In many ways, searching through perspectives is much the same as searching through any other kind of informational content. As such, some form of link analysis — similar to, for instance, PageRank (Page et al., 1999) — may yield highly valuable search functionality.

From another viewpoint, however, the value of perspectives can be viewed to be highly dependent on community and social connections. An objective, worldwide ranking of how a search term relates to a particular perspective is unlikely to be useful for perspectives that have niche appeal. As such, social platforms could also play a key role in the distribution of perspectives. One could even imagine a social platform built especially for sharing and ranking perspectives among communities. Such a platform would present a number of interesting risks and challenges, of

course, including those surrounding echo chambers and groupthink (Wu et al., 2011; Flaxman, Goel and Rao, 2016).

7.4 Critical Analysis

In contrast to the other studies within this project, it is important to recall that this evaluative study was not carried out in person. Instead, participants independently completed a short, digital survey, yielding less detailed results and thoughts than could have otherwise have been captured. As such, the results of this study are intended to be taken in combination with the evaluative results from the previous chapters — which, recall, share much the same pool of participants. Even within this larger context, however, it is worth emphasising that these results exist to present possibilities rather than probabilities, examining the thoughts of small sample of participants which may not be representative of the larger population of knowledge workers.

In spite of this, the study was highly effective in highlighting a number of perceived benefits and potential issues with our prototype. Curiously, though, the results do not contain a great deal of discussion around the highest-level design ideas. Rather, participants seemed to understand and appreciate the fundamental elements of the design with relative ease, and instead looked through these towards a number of specific details. Although this is likely to be in part due to familiarity from the previous study, it does demonstrate the rather encouraging possibility that users are able to consider the proposed system as natural in some sense. Taken in combination with descriptions of the system as being “logical” and “coherent”, this represents an extremely positive development in open hypermedia interface design.

7.5 Summary

Building particularly on the contents of the last two chapters, this chapter has described the design and evaluation of a single high fidelity video prototype. Demonstrating our latest design ideas in the fundamental areas of link behaviour and history navigation, this shows the imagined foundations of a modern, useful, usable open hypermedia system for knowledge work. While there is still much work to be done, evaluation results indicate that our prototype was successful to some degree in achieving its goals. It presents a coherent and somewhat natural interface to an open hypermedia system, through which users can see themselves extracting immense value. As far as we are aware, this video mockup is the first of its kind, illustrating the possibility and potential value of usable open hypermedia systems in the 21st century.

Chapter 8

Conclusion

8.1 Contributions

The central objective of this dissertation has been to both prove and improve the value of open hypermedia systems in enhancing knowledge work. To this end, we have made a number of contributions. We began this work by calling attention to the seemingly untapped potential of open hypermedia systems and the remarkable research gap that surrounds them. Having explored this topic throughout the last six chapters, the results are highly encouraging. Our user studies provide empirical evidence for the potential utility of open hypermedia systems in knowledge work, and further prove that more natural, usable interfaces to such systems are possible.

In our first chapter, we reviewed the existing literature in this area. This was critical in guiding the overall direction of the project, motivating our focus towards open hypermedia interface design and allowing us to identify a number of open questions and potential issues with open hypermedia systems. Mostly, these issues concerned the fundamentals of how information should be organised in such systems and in what ways these systems should allow relationships between resources to be explored. As such, this prompted the first material contribution of our work, an exploratory study of user behaviour in multi-window knowledge work. From this, we gathered insights into how students use — and desire to use — their computers to explore and integrate information from multiple, heterogeneous sources. These results have been instrumental in informing the design work within this project, and may prove valuable in guiding the design of future hypermedia systems more generally.

Building on the insights granted by both this exploratory work and the existing literature, we then presented four open hypermedia design proposals. Taking the form of storyboards, these propose design solutions in the areas of link display and behaviour, information organisation, and backtracking. Although these solutions are heavily inspired by existing work, each is a unique construction that we have

not seen proposed elsewhere and forms a significant contribution to the state of the art in this area. Further, the evaluation of these largely orthogonal design ideas indicates that each idea possesses some degree of merit, illustrating promising possibilities in the design of open hypermedia systems.

Selecting a handful of ideas in the areas of link behaviour and backtracking to develop further, we also produced a single high fidelity video prototype. This shows the imagined foundations of a modern, useful, usable open hypermedia system in a realistic setting. As alluded to previously, our evaluation of this prototype suggests that participants found the proposed system to be coherent and natural to use on some level. This is not to imply, of course, that our proposed system is perfect, but rather that it describes an improvement to some number of existing systems. Moreover, all three studies carried out within this project indicate that open hypermedia may, in many respects, describe considerable improvements to the current status quo in hypermedia: the WWW. In our exploratory study in particular, participants found immense value in the possibility of representing and sharing cross-program connections, indicating that these concepts may allow for dramatic improvements in digital knowledge work.

8.2 Critical Reflection

As contemporary research into open hypermedia interface design is exceptionally rare (as discussed in subsection 2.1.2), this dissertation has opted to broadly survey an assortment of topics rather than to focus on any one topic in particular. This has allowed us to take a wider look at some of the usability issues that can arise in these sorts of systems, though at the cost of failing to reach any firm conclusions about specific design decisions. The prototypes presented within this project explore only a tiny fraction of the possible design space, and are intended to stimulate further discussion, not to depict final solutions.

It is also the case that the prototypes presented within this dissertation are not interactive in any way. As a result, all evaluation is based on reasoning from what the prototypes look like to use, and not from what they feel like to use. To combat this, much effort was put towards ensuring that our prototypes communicate the dynamics of our proposals as best as is possible within their respective mediums. Our evaluation results indicate that we achieved some success in doing this, as participants understood what was being proposed fairly well. Despite this success, however, we believe that these prototypes must ultimately be accompanied by interactive systems in order to assess how users feel about and interact with different design elements in practice.

Additionally, the exploratory and evaluative studies within this project make use of relatively small samples of participants, largely consisting of computer science

students. While a number of our studies made use of a familiarity with computer science in order to accurately stimulate and represent knowledge work, a consequence of this is that the observed behaviours may not be representative of groups in other knowledge domains. Further, although this has been stated numerous times throughout previous chapters, it is important to re-emphasise that all studies within this project were purely qualitative in nature, existing only to show possibilities of use. As such, small sample sizes are not viewed to be a significant issue.

8.3 Future Work

As this dissertation is highly exploratory in nature, there is an extensive space of future work that can follow naturally from it. Many such opportunities have been identified as they arose in discussions throughout the body of the dissertation. Here, though, we will highlight just a few of the opportunities that we consider to be most closely related to our work and potentially valuable in improving the design and usability of open hypermedia systems.

8.3.1 Application Interface Design

An open hypermedia system is only as good as the applications that comprise it. For this reason, the design of open hypermedia viewer programs is critical to the success of these systems. While there are many examples of existing research in this area, the field as a whole still contains many open questions and lacks consensus even around some of the most fundamental design issues (Nürnberg, Leggett and Wiil, 1998).

This dissertation can be seen to make some contributions to this area, but largely chooses to focus on exploring an assortment system-wide design issues rather than considering application specifics. In contrast, what appears to be lacking most in this area is directed research to help achieve consensus around specific design issues. One such issue worth drawing attention to is the display of links in non-textual media such as video, as briefly discussed alongside our high fidelity prototyping. This is examined to some extent by existing literature (Davis, Hall and Heath, 1993), but remains largely unresolved and is likely to make for fruitful future work.

The Marker Bar

Following this topic further, in considering link display within audio and video in Microcosm, Hall, Davis and Hutchings (1996, p.85) suggest the use of a bar besides the content to display link anchors. In some respects, this is similar to the marker bar design element that we have proposed. In our work, however, the marker bar is not uniquely applicable to temporal content. Rather, it is intended to be a more

general interface element, comparable in some ways to the scroll bar. Particularly, our prototypes use this construct to mark the presence of links from various linkbases, allowing for links to be more subtly highlighted inline and enabling users to quickly distinguish between links from different linkbases across various forms of media. Functionality that has been viewed as valuable by participants in our evaluative studies.

In spite of its identified benefits, though, there are still a large number of open questions surrounding the marker bar. These include how the mapping of colours to linkbases should best be performed, the scalability of using colours to denote different linkbases, how the marker bar should best interact with various forms of non-textual media, and what interactions the marker bar should allow.

More generally, the performance of the marker bar in relation to possible alternatives remains unclear. Although this interface element was based on ideas from existing literature, it is highly unlikely to represent a perfect solution in its current form. A study to compare a number of alternatives in this area may help to alleviate these issues in future, establishing a set of foundational guidelines around link display in open hypermedia systems.

8.3.2 History

As identified in our literature review, history is an area in which many existing open hypermedia systems are severely lacking. While the separation of and navigation through various chains of history within typical Web browsers are associated with well established design solutions, the way in which these behaviours should be best supported within open hypermedia systems is much less clear. As a result, a number of existing systems — such as Microcosm — fail to expose any mechanisms at all by which users can efficiently switch between chains of history (Hall, Davis and Hutchings, 1996). Given that knowledge work can be highly fragmented (González and Mark, 2004), this clearly represents a major design issue.

In this dissertation, we have explored one possible solution to these issues through our ‘Spaces & History’ prototype (subsection 5.1.2). The evaluation in section 6.2 suggests that this approach has much promise, but is by no means a complete solution in its current form. We consider developing these ideas further and contrasting them with possible alternatives as excellent opportunities for future work in this area.

8.3.3 Window Management

In creating connections between multiple programs, open hypermedia systems typically display information through multiple windows. As discussed in Chapter 2,

this can be beneficial in some respects, but may be problematic in others. Particularly, interactions with these systems may spawn a large number of windows, cluttering the user's work environment and making it more difficult to find information. Though we have touched upon various solutions to these issues throughout this dissertation — including in our 'Window Search' prototype (subsection 5.1.3) and the discussions in subsection 6.2.3 — the practical impacts of and possible solutions to these problems are yet to be fully explored.

8.3.4 Dynamic Organisation & Exploration

By liberating hypermedia interactions from the constraints of a single window, open hypermedia systems present new opportunities for the dynamic rearrangement, organisation, and display of information. While we have already touched upon the possibilities surrounding the application of these techniques to window management, responses from our exploratory study suggest that there may also be value in dynamically rearranging information for other purposes. We have principally examined this idea in this work through our 'Exploration Mode' prototype (subsection 5.1.4). Although the feedback relating to this prototype was largely mixed, participants widely appreciated the underlying ideas as being highly valuable. As we have focused on only a tiny slice of the possible design space around such functionality, we consider this to be a highly exciting area for further work.

8.3.5 Annotation

In separating hypermedia structure from resource content, our work has naturally considered linkbases as a mechanism by which the externalisation and communication of connections can occur. In thinking more widely about the goals and applications of such functionality, however, linkbases can be considered to make up only single piece of a larger narrative surrounding annotation. Looking back on the results from the exploratory study, it is clear that participants desired for their digital environments to serve as natural mediums for thought and communication. While the ability to create connections between information was deemed to play a large part in this, annotation was also highlighted by participants as highly valuable.

In fact, the value of annotation in knowledge work is widely appreciated in existing research (Marshall, 1997; Sellen and Harper, 2003). While the design of annotation solutions within hypermedia systems (Denoue and Vignollet, 2000) and for knowledge work (Steimle, Brdiczka and Muhlhauser, 2009) have been well explored in existing work, there are many open questions around how annotation might best be integrated into *open* hypermedia systems. How are annotations best created and displayed across a range of media types? (Are annotations fully freeform? Does their display tie in with the marker bar?) To what degree annotations related to linkbases? Can users create new links from annotations? If so, what happens if

they try to create a link from an annotation that other viewers of that linkbase might not see? We view these questions as excellent grounds for future work in this area.

8.3.6 Supporting Collaborative Work

In many areas of knowledge work, collaboration is essential. As such, systems that set out to empower knowledge workers in these contexts ought to support and amplify collaborative activities as best as is possible. While this dissertation has largely ignored collaboration, focusing instead on elementary design issues for single-user knowledge work, there is one area around which collaborative possibilities have been repeatedly discussed: linkbases. These appear to be an excellent fit for collaborative work, but pose a great many design questions.

Key amongst these are the issues of collaborative linkbase editing, linkbase sharing, and linkbase discovery. These have been touched upon throughout this dissertation — particularly in section 7.3.4 — but remain generally unanswered. Perhaps the most interesting aspect of exploring these problems in this context is that open hypermedia systems do not constrain the user to a single program. As such, design solutions in this space need not be constrained to the contents of a single window, and can draw upon the wealth of existing systems level CSCW research. This can allow, for instance, for the use multiple mouse pointers for collaborative purposes as was done in NLS — creating a “shared intellectual space” rather than supporting collaboration through a series of “bolted-on hack[s]” (Victor, 2013a).

8.4 Summary

In this dissertation we have explored the potential of open hypermedia systems in supporting contemporary knowledge work. While these systems can be viewed as ‘only’ extending the basic capabilities of the Web in a few areas, our work suggests that the way in which such systems can better capitalise on fundamental human capabilities for knowledge work and spawn new kinds of rich, information ecosystems could significantly improve intellectual effectiveness. While much work is still to be done in order to realise these efficiency gains in a practical setting, we believe that our design work makes new contributions in this area, and further that open hypermedia systems could dramatically improve the role of the computer as a medium for thought.

Bibliography

- Apple Inc., 2016. *Use Mission Control on your Mac* [Online]. Available from: <https://support.apple.com/en-gb/HT204100> [Accessed 10th March 2018].
- Atzenbeck, C., Schedel, T., Tzagarakis, M., Roßner, D. and Mages, L., 2017. Re-visiting hypertext infrastructure. *Proceedings of the 28th ACM Conference on Hypertext and Social Media*, 4–7 July 2017 Prague, Czech Republic. New York, NY, USA: ACM, pp.35–44.
- Berners-Lee, T., 1989. *Information management: A proposal*. Switzerland: CERN, (CERN-DD-89-001-OC).
- Berners-Lee, T., 1991. *Re: status. re: X11 browser for WWW* [Online]. Available from: <http://lists.w3.org/Archives/Public/www-talk/1991SepOct/0003.html> [Accessed 23rd April 2018].
- Berners-Lee, T. and Cailliau, R., 1990. *WorldWideWeb: Proposal for a hypertext project* [Online]. Available from: <https://www.w3.org/Proposal.html> [Accessed 23rd April 2018].
- Berners-Lee, T., Hendler, J., Lassila, O. et al., 2001. The semantic web. *Scientific American*, 284, pp.28–37.
- Bieber, M. and Wan, J., 1994. Backtracking in a multiple-window hypertext environment. *Proceedings of the 1994 ACM European Conference on Hypermedia Technology*, 19–23 September 1994 Edinburgh, Scotland. New York, NY, USA: ACM, pp.158–166.
- Borsook, T.K. and Higginbotham-Wheat, N., 1992. A psychology of hypermedia: A conceptual framework for R&D. Presented at the 1992 National Convention of the Association for Educational Communications and Technology, Washington DC, USA.
- Borst, J.P., Taatgen, N.A. and Rijn, H. van, 2015. What makes interruptions disruptive? A process-model account of the effects of the problem state bottleneck on task interruption and resumption. *Proceedings of the 33rd annual ACM Conference on Human Factors in Computing Systems*, 18–23 April 2015 Seoul, Republic of Korea. New York, NY, USA: ACM, pp.2971–2980.

- Braun, V. and Clarke, V., 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3, pp.77–101.
- Bush, V., 1945. As we may think. *Atlantic Monthly*, 176, pp.101–108.
- Chen, S.Y. and Macredie, R.D., 2002. Cognitive styles and hypermedia navigation: Development of a learning model. *Journal of the American Society for Information Science and Technology*, 53, pp.3–15.
- Choo, C.W., Detlor, B. and Turnbull, D., 2000. *Web work: Information seeking and knowledge work on the World Wide Web*. Norwell, MA, USA: Kluwer Academic Publishers.
- Conklin, J., 1987. Hypertext: An introduction and survey. *IEEE computer*, 20, pp.17–41.
- Czerwinski, M., Dantzich, M.v., Robertson, G. and Hoffman, H., 1999. The contribution of thumbnail image, mouse-over text and spatial location memory to web page retrieval in 3D. In: M.A. Sasse and C. Johnson, eds. *Human-computer interaction INTERACT '99: IFIP TC.13*, 30 August–3 September 1999 Edinburgh, Scotland. Amsterdam, The Netherlands, The Netherlands: IOS Press, pp.163–170.
- Davis, H., Hall, W. and Heath, I., 1993. Media integration issues within open hypermedia systems. *Proceedings of IEEE International Symposium on Multimedia Technologies and Future Applications* Southampton, England. London, England: Pentech Press, pp.21–23.
- Davis, H., Hall, W., Heath, I., Hill, G. and Wilkins, R., 1992. Towards an integrated information environment with open hypermedia systems. *Proceedings of the ACM Conference on Hypertext*, 30 November–4 December 1992 Milan, Italy. New York, NY, USA: ACM, pp.181–190.
- Davis, H.C., 1998. Referential integrity of links in open hypermedia systems. *Proceedings of the ninth ACM Conference on Hypertext and Hypermedia*, 20–24 June 1998 Pittsburgh, PA, USA. New York, NY, USA: ACM, pp.207–216.
- Denoue, L. and Vignollet, L., 2000. An annotation tool for web browsers and its applications to information retrieval. *Content-Based Multimedia Information Access*, 12–14 April 2000 Paris, France. Paris, France: Le Centre de Hautes Études Internationale D'Informatique Documentaire, pp.180–195.
- Dervin, B., Foreman-Wernet, L. and Lauterbach, E., 2003. *Sense-making methodology reader: Selected writings of Brenda Dervin*. Cresskill, NJ, USA: Hampton Press.

- Dubroy, P. and Balakrishnan, R., 2010. A study of tabbed browsing among mozilla firefox users. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 10–15 April 2010 Atlanta, GA, USA. New York, NY, USA: ACM, pp.673–682.
- Engelbart, D.C., 1962. *Augmenting human intellect: A conceptual framework*. Menlo Park, CA, USA: Air Force Office of Scientific Research, Stanford Research Institute, (Summary report AFOSR-3223 - Contract AF49(638)-1024, SRI Project 3578).
- Engelbart, D.C., 1990. Knowledge-domain interoperability and an open hyperdocument system. *Proceedings of the 1990 ACM Conference on Computer-supported Cooperative Work*, 7–10 October 1990 Los Angeles, CA, USA. New York, NY, USA: ACM, pp.143–156.
- Flaxman, S., Goel, S. and Rao, J.M., 2016. Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly*, 80, pp.298–320.
- Garrett, L.N., Smith, K.E. and Meyrowitz, N., 1986. Intermedia: Issues, strategies, and tactics in the design of a hypermedia document system. *Proceedings of the 1986 ACM Conference on Computer-supported Cooperative Work*, 3–5 December 1986 Austin, TX, USA. New York, NY, USA: ACM, pp.163–174.
- González, V.M. and Mark, G., 2004. "constant, constant, multi-tasking craziness": Managing multiple working spheres. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 24–29 April 2004 Vienna, Austria. New York, NY, USA: ACM, pp.113–120.
- Graham, P., 2001. *Beating the averages* [Online]. Available from: <http://www.paulgraham.com/avg.html> [Accessed 17th March 2018].
- Grønbaek, K. and Trigg, R.H., 1999. *From web to workplace: Designing open hypermedia systems*. Cambridge, MA, USA: MIT Press.
- Halasz, F.G., 1991. "Seven Issues": Revisited, *Hypertext '91 closing plenary* [Online]. Available from: <http://www.eastgate.com/hypertext/Halasz/Halasz7Rev.pdf> [Accessed 10th April 2018].
- Hall, W., Davis, H. and Hutchings, G., 1996. *Rethinking hypermedia: The Microcosm approach*. Norwell, MA, USA: Kluwer Academic Publishers.
- Hill, W.C., Hollan, J.D., Wroblewski, D. and McCandless, T., 1992. Edit wear and read wear. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 3–7 May 1992 Monterey, CA, USA. New York, NY, USA: ACM, pp.3–9.

- Hudson, J.M., Christensen, J., Kellogg, W.A. and Erickson, T., 2002. "I'd be overwhelmed, but it's just one more thing to do": Availability and interruption in research management. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 20–25 April 2002 Minneapolis, MN, USA. New York, NY, USA: ACM, pp.97–104.
- Hutchings, D.R. and Stasko, J., 2004. Revisiting display space management: Understanding current practice to inform next-generation design. *Proceedings of Graphics Interface 2004*, 17–19 May 2004 London, Ontario, Canada. Waterloo, Ontario, Canada: Canadian Human-Computer Communications Society, pp.127–134.
- Kay, A., 2017. *Who was the first to introduce multiple desktops, Linux or macOS?* [Online]. Available from: <https://www.quora.com/Who-was-the-first-to-introduce-multiple-desktops-Linux-or-macOS> [Accessed 4th March 2018].
- Malcolm, K.C., Poltrock, S.E. and Schuler, D., 1991. Industrial strength hypermedia: Requirements for a large engineering enterprise. *Proceedings of the third annual ACM Conference on Hypertext*, 15–18 December 1991 San Antonio, TX, USA. New York, NY, USA: ACM, pp.13–24.
- Markoff, J., 2005. *What the dormouse said: How the sixties counterculture shaped the personal computer industry*. London, England: Penguin Books.
- Marshall, C.C., 1997. Annotation: From paper books to the digital library. *Proceedings of the second ACM International Conference on Digital Libraries*, 23–26 July 1997 Philadelphia, PA, USA. New York, NY, USA: ACM, pp.131–140.
- Marshall, C.C., Shipman, III, F.M. and Coombs, J.H., 1994. VIKI: Spatial hypertext supporting emergent structure. *Proceedings of the 1994 ACM European Conference on Hypermedia Technology*, 19–23 September 1994 Edinburgh, Scotland. New York, NY, USA: ACM, pp.13–23.
- Memory & Imagination: New pathways to the library of congress*, 1990. Film. Directed by Julian Krainin and Michael R. Lawrence. USA: Krainin Productions.
- Meyrowitz, N., 1989. The missing link: Why we're all doing hypertext wrong. In: E. Barrett, ed. *The society of text: Hypertext, hypermedia, and the social construction of information*. Cambridge, MA, USA: MIT Press, pp.107–114.
- Nelson, T.H., 1999. Xanalogical structure, needed now more than ever: Parallel documents, deep links to content, deep versioning, and deep re-use. *ACM Comput. Surv.*, 31.

- Nürnberg, P.J., Leggett, J.J. and Wiil, U.K., 1998. An agenda for open hypermedia research. *Proceedings of the ninth ACM Conference on Hypertext and Hypermedia*, 20–24 June 1998 Pittsburgh, PA, USA. New York, NY, USA: ACM, pp.198–206.
- Obendorf, H. and Weinreich, H., 2003. Comparing link marker visualization techniques: Changes in reading behavior. *Proceedings of the 12th International Conference on World Wide Web*, 20–24 May 2003 Budapest, Hungary. New York, NY, USA: ACM, pp.736–745.
- Page, L., Brin, S., Motwani, R. and Winograd, T., 1999. *The PageRank citation ranking: Bringing order to the web*. Stanford, CA, USA: Stanford InfoLab, (1999-66).
- Paul, S.A. and Reddy, M.C., 2010. Understanding together: Sensemaking in collaborative information seeking. *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work*, 6–10 February 2010 Savannah, GA, USA. New York, NY, USA: ACM, pp.321–330.
- Pearl, A., 1989. Sun’s link service: A protocol for open linking. *Proceedings of the second annual ACM Conference on Hypertext*, 5–8 November 1989 Pittsburgh, PA, USA. New York, NY, USA: ACM, pp.137–146.
- Pichai, S. and Upson, L., 2009. *Introducing the Google Chrome OS* [Online]. Available from: <https://googleblog.blogspot.co.uk/2009/07/introducing-google-chrome-os.html> [Accessed 21st March 2018].
- Plumlee, M.D. and Ware, C., 2006. Zooming versus multiple window interfaces: Cognitive costs of visual comparisons. *ACM Trans. Comput.-Hum. Interact.*, 13, pp.179–209.
- Ringel, M., 2003. When one isn’t enough: An analysis of virtual desktop usage strategies and their implications for design. *CHI ’03 Extended Abstracts on Human Factors in Computing Systems*, 5–10 April 2003 Ft. Lauderdale, FL, USA. New York, NY, USA: ACM, pp.762–763.
- Robertson, G., Czerwinski, M., Larson, K., Robbins, D.C., Thiel, D. and Dantzich, M. van, 1998. Data mountain: Using spatial memory for document management. *Proceedings of the 11th annual ACM Symposium on User Interface Software and Technology*, 1–4 November 1998 San Francisco, CA, USA. New York, NY, USA: ACM, pp.153–162.
- Russell, D.M., Stefik, M.J., Pirolli, P. and Card, S.K., 1993. The cost structure of sensemaking. *Proceedings of the INTERACT ’93 and CHI ’93 Conference on Human Factors in Computing Systems*, 24–29 April 1993 Amsterdam, The Netherlands. New York, NY, USA: ACM, pp.269–276.

- Sculley, J., 1987. *The knowledge navigator*. Educom Keynote Address.
- Sellen, A.J. and Harper, R.H., 2003. *The myth of the paperless office*. Cambridge, MA, USA: MIT Press.
- Steimle, J., Brdiczka, O. and Muhlhauser, M., 2009. CoScribe: Integrating paper and digital documents for collaborative knowledge work. *IEEE Transactions on Learning Technologies*, 2, pp.174–188.
- Tauscher, L. and Greenberg, S., 1997. How people revisit web pages: empirical findings and implications for the design of history systems. *International Journal of Human-Computer Studies*, 47, pp.97–137.
- Utting, K. and Yankelovich, N., 1989. Context and orientation in hypermedia networks. *ACM Trans. Inf. Syst.*, 7, pp.58–84.
- Victor, B., 2006. *Magic ink: Information software and the graphical interface* [Online]. Available from: <http://worrydream.com/MagicInk/> [Accessed 10th April 2018].
- Victor, B., 2012. *Stop drawing dead fish* [Online]. Available from: <https://vimeo.com/64895205> [Accessed 26th April 2018].
- Victor, B., 2013a. *A few words on Doug Engelbart* [Online]. Available from: <http://worrydream.com/Engelbart/> [Accessed 19th April 2018].
- Victor, B., 2013b. *Media for thinking the unthinkable* [Online]. Available from: <http://worrydream.com/MediaForThinkingTheUnthinkable/> [Accessed 26th April 2018].
- Victor, B., 2014. Humane representation of thought: A trail map for the 21st century. *Proceedings of the 27th annual ACM Symposium on User Interface Software and Technology*, 5–8 October 2014 Honolulu, HI, USA. New York, NY, USA: ACM, pp.699–699.
- W3C, 2015. *The app: URL scheme* [Online]. Available from: <https://www.w3.org/TR/2015/NOTE-app-uri-20150723/> [Accessed 24th April 2018].
- Weinreich, H., Obendorf, H. and Lamersdorf, W., 2001. The look of the link - concepts for the user interface of extended hyperlinks. *Proceedings of the 12th ACM Conference on Hypertext and Hypermedia*, 14–18 August 2001 Århus, Denmark. New York, NY, USA: ACM, pp.19–28.
- Wilson, D.G. and Papadopoulos, J., 2004. *Bicycling science*. Cambridge, MA, USA: MIT Press.

Wu, S., Hofman, J.M., Mason, W.A. and Watts, D.J., 2011. Who says what to whom on Twitter. *Proceedings of the 20th International Conference on World Wide Web*, 28 March–1 April 2011 Hyderabad, India. New York, NY, USA: ACM, pp.705–714.

Østerbye, K. and Wiil, U.K., 1996. The flag taxonomy of open hypermedia systems. *Proceedings of the the seventh ACM Conference on Hypertext*, 16–20 March 1996 Bethesda, MD, USA. New York, NY, USA: ACM, pp.129–139.

Appendix A

Study Data

A.1 Exploratory Study Interview Transcripts

Participant 1

[Discussion of specific details of the solution]

Interviewer: So, how did you feel when you were trying to carry out the task?

Participant: Umm... not great. [Laughter] There's just so much. Lots of information to take in. Trying to match each of the different resources to try and form some kind of view of what's happening.

Interviewer: So would you say you were 'in the zone' doing that, or did you feel confused, or what was your emotional response?

Participant: At first I was overwhelmed, then I started seeing similar things in files and started to sort of understand a little bit. I started off confused but ended up making something out of it.

Interviewer: Okay, that makes sense. Do you feel generally that you got into the flow of things then, or do you feel it was always a bit of an uphill battle?

Participant: Uhh yeah, always an uphill battle. [Laughter] Still not really there.

Interviewer: Do you think you made effective use of the resources provided?

Participant: I tried to, yeah. I tried to look at as many as possible and then tried to make links between them. I didn't get through everything though with the time limit. I'm not very good at skim reading is the problem.

Interviewer: Okay. How did you deal with sorting through and searching the great mass of resources that was presented to you?

Participant: Umm... I probably could have done things a bit better, but I just went by memory. I don't know if you saw but I accidentally closed a file at one point and it had something in there --- something to do with GOT --- and I knew I saw that somewhere. But I forgot which file it was so I had to scan through all these files again. So making notes probably would have been the better approach

- to figure out what was going on. So, yeah, probably not well.
- Interviewer: Could you try and describe for me your general process of trying to integrate the information from the different resources?
- Participant: I was trying to figure out what all these ‘__’ things meant, since they seemed kind of important. I guess was kind of focused on the header bit. So I was looking mainly for what these meant.
[Points at various labels within the hex dump]
- Interviewer: Okay, so would you say you focused on one of those at a time and tried to look through the documents, or what was your method for doing that?
- Participant: Yeah it was one at a time. I started at first trying to go from top to bottom, but yeah one at a time trying to figure them out and then obviously referencing back when I forgot.
- Interviewer: Okay. So looking through multiple documents for each, or...?
- Participant: Yeah, then I found this document [Opens up Mach-0 file format reference] which was really useful.
- Interviewer: Okay. So you said you felt kind of overwhelmed --- at least at first --- with how much information there was. What do you think the source of that was? Literally just the amount of information?
- Participant: Probably the amount of information combined with a complete lack of understanding of the topic. So that’s probably what did it. It was nice to learn something new though.
- Interviewer: Did you find it difficult to find a starting point?
- Participant: I don’t think I did actually. I started with this PowerPoint, just went for something that’s usually pretty easy to get into. It was useful finding that [Points to section of PowerPoint presentation] straight away --- I didn’t really understand what it was doing but I sort of had something to go of. And then some of the words were in the names of other files. So yeah, finding a starting point wasn’t too bad, it was just understanding everything. Maybe I just got lucky with that though.
- Interviewer: Okay. To alleviate some of your disorientation and I guess provide some structure to this big mess of resources, do you think it would have been helpful to have some kind of visual relationships between the resources? So, for instance if you could see visual link lines between the related parts of windows, or if you could temporarily view the resources in some kind of map or something like that. Do you think that would have been helpful just as a form of structuring?
- Participant: Oh yeah, definitely. You saw I had quite a few files open at the same time, so yeah that would have been helpful, to map things.
- Interviewer: Do you have any insights as to what structure might be most helpful there?
- Participant: I don’t know. I can see how it would be very useful, I just don’t know how it would work.
- Interviewer. Okay. So you think that some structure would be useful, you’re just not necessarily sure what form that should take?
- Participant: Yeah, it would definitely be useful. I thought that this [Puts windows into split view] was like the epitome of multitasking

[Laughter], but yeah a bit more than that would be awesome.

Interviewer: Okay. And then, occasionally I saw you would skim a resource and then decide to close it because it wasn't useful, but then go back to it --- I think by accident --- and be like "Hang on, what's this thing again?".

Participant: Yeah, that was when I was trying to find that GOT thing. I opened one, forgot that I opened it a second ago, and then reopened it.

Interviewer: Do you think some kind of additional structure could help with that?

Participant: Oh yeah. It probably would have been helpful if I had, say, 'last opened' or something like that. Just so I could know I just opened it a second ago. That might not be hugely useful for most cases, but in this case it would have been.

Interviewer: I think that brings up an interesting point as well. Often for research-type tasks like this, rather than having a directory of resources you would just be browsing the Web. When you just have a directory of resources, though, you don't have a history or any way of going back, you're just on your own. Did you feel constrained at all by the fact that there isn't a history and that you can't go back to where you were and see where you visited?

Participant: Yeah. If I was in Chrome when I closed that file I could have just reopened it, which would have been super useful. In my Final Year Project as well, doing the lit review, I just have millions of tabs open and I'm scared to close them because I don't know what's useful. That's probably something that's not really avoidable. Definitely though having a history would be super useful.

Interviewer: Did the process of sifting and searching through the information feel efficient or inefficient to you would you say?

Participant: I would say inefficient, because of all the little hiccups that I had. Flicking through trying to find where I found that thing previously. Also trying to constantly switch between PowerPoint, Excel, the text editor and all of that. That was a bit of a pain. That's why the side-by-side view was kind of useful.

Interviewer: So that's another question I have actually: did you find the switching between different programs and different windows to be efficient?

Participant: No, it was just more of just an annoyance than anything. I guess depending on the device it might be different as well, because I remember when I had a Mac you can just do that three finger thing and open another desktop, it's just easy. It made it a lot easier. Whereas Windows is just a bit of a bitch. [Laughter]

Interviewer: Did you put any thought at all into how you sort of managed all the different windows?

Participant: No not really. I just sort of opened them and if I needed it I would just go through and find the one I was looking for. I didn't really think about it.

Interviewer: Did you find it more of a blessing or a curse to have the PDFs display in a separate tabbed interface? I guess there's a level

- of inconsistency if sometimes you're having to switch between windows to go to the different resources, but for PDFs you're having to not only switch windows but also find the right tab. What are your thoughts on that?
- Participant: I don't know. That's probably something I'm just so used to by this point --- just flicking between different tabs --- that I didn't notice it as an annoyance.
- Interviewer: Okay, so it's not really a problem for you?
- Participant: No, but I could see why it might be for some. I think it's only because I'm used to it from all the research I've been doing recently.
- Interviewer: Okay. While carrying out the task, do you think you were paying more attention to the specific program you were in --- so for instance, Edge or whatever --- or the specific resources that you were trying to get access to?
- Participant: Umm... only when I had problems --- other than that, I was just focused on the resources.
- Interviewer: I noticed that a lot of the time you had the 'main.txt' file side-by-side with something else, but you also seemed to find that since it was a relatively wide text file that sometimes the lines would wrap. Did you feel constrained by your screen space at all?
- Participant: Yeah. That was a big of a pain. Having to manually adjust the windows so they both fit was a bit of an annoyance, as Windows automatically goes to splitting the screen between the two windows straight down the middle. I have dual monitors at home, and it's just easier to just have two things. It's kind of a pain going to a single screen.
- Interviewer: Alright. Let's move into hypothetical land for these last few questions. If it were possible, how helpful do you think it would have been to organise all the information in some way digitally as you were working through it? So, for instance if you could create cross-program hyperlinks or do some kind of digital structuring to the information as you were looking through it.
- Participant: Yeah, most likely. Knowing me, even if it didn't exist I probably wouldn't have used it [Laughter], but it would definitely help. I could just highlight this [Selects label within the hex dump] and have a little hyperlink, and it could just throw me into that file [Opens Mach-0 file format reference]. That would be great.
- Interviewer: And do you think it would be useful to have specific sections of documents rather than only being able to link to documents as a whole?
- Participant: Oh yeah definitely. That would help quite a lot.
- Interviewer: And on the flip side of this situation --- rather than you having to structure the information yourself --- if somebody had already gone through the process of making sense of the information and created links between the appropriate parts of the different resources, do you think making use of that kind of relational structure would have helped you?
- Participant: Yeah most likely. I wouldn't be faffing around with all the

less useful files, 'cause I would have found one thing that could have told me what I needed to know based on that link or whatever it was. Probably would have got further in the task.

Interviewer: And what level of structure do you think would have been most useful in that? So do you think something that's closer to a sequential guided tour through the resources or something more freely explorable would have been more useful?

Participant: Umm... a guided tour would probably be a bit too far, because you're a bit restrained in that way, you don't have much room--- well I guess... I guess you just could go to a different tab or whatever. I suppose a guided tour would actually be useful thinking about it, yeah. It could give you everything you need, then you could research on the side as well.

Interviewer: Okay. So I guess what you're saying there is that it seems like maybe both would be useful, right? You could go through a guided tour but then you also want to search on your own where you're freely exploring stuff

Participant: Yeah. It's kind of like a city tour, isn't it? They take you to the main parts, then you can sightsee and explore the little side streets.

Interviewer: Right, yeah!

Interviewer: Umm, okay so I think that's everything I have. Is there anything you want to add?

Participant: Umm, no I don't think so. I think I see what you're trying to do with your project now, and that's super cool. It would make life so much easier. That was cool, I enjoyed that.

Interviewer: Thanks for taking part, I'm glad you found it enjoyable.

Participant: That's alright, I hope it was useful!

Participant 2

Interviewer: First of all, tell me about your experience and progress in completing the task.

Participant: So, progress feels limited. There are a lot of resources, the task was very unfamiliar, and I wasn't overly sure where to begin. And I feel there was a lot of stuff I maybe didn't need as well or wasn't sure exactly where to start with. I started by looking at something that sort of looked vaguely familiar. When I found something that actually related to the mess of hexadecimal, that felt good, that was the highlight of my experience... then I realised I didn't know what it was talking about. [Laughter]

Participant: I have a vague understanding of what Mach-0 actually is now, and also a lot of places talked about the linking and loading, but I found it hard to see that reflected in the hex dump. I probably missed something.

Interviewer: Okay, yeah it's hard to make that association I guess. The jump from the conceptual stuff to actual nuts-and-bolts bytes. That makes sense.

- Interviewer: I guess before I go on with my other questions, do you want to give me a brief overview of what you think you now understand about the format?
- [Discussion of specific details of the solution]
- Interviewer: How would you say you felt when you were trying to figure out what was going on?
- Participant: Inadequate. [Laughter] Umm... I guess it's a bit weird being watched while you work, and I felt like I wasn't getting anywhere for a lot of it. I felt like I was scrolling and looking for things that jumped out to me as familiar, but it was kind of hard because none of it was familiar.
- Interviewer: Okay. How would you describe that on an emotional level? Maybe confused or disoriented?
- Participant: Disoriented's a good word, I like that. I guess maybe if I was on my desktop at home it would be a bit nicer, because I could spread everything out a little bit more, but I'm stuck on my slow little laptop.
- Interviewer: Did you find screen real estate was a limiting factor for you?
- Participant: Yeah. You may have seen me play around with trying to get the PDF to display on the left. I couldn't hide the bar [Points to part of the PDF viewer] so I just gave up. Don't want to waste time.
- Interviewer: Yeah. Do you think you made effective use of the resources provided?
- Participant: I think so. I mean, I didn't use all of them. I looked for files that looked interesting and scrolled through a little bit... maybe if they were sorted it might have been easier to navigate.
- Interviewer: Sorted in what sense?
- Participant: Some of them sort of cover similar topics, so maybe if there was some kind of folder hierarchy there. I could use that to find what I wanted more easily. Sorry that's a very vague answer.
- Interviewer: No no, that makes complete sense. I guess just adding some kind of structure to the unstructured mess of information.
- Interviewer: On that topic, how did you deal with sorting through the wealth of information that was available to you?
- Participant: Uhh, did I have a process? [Laughter] I went for the Mach-0 Wikipedia page first, I just wanted to get a high-level understanding of it. And then I wanted to work out--- I opened "Mach-0 Peculiarities", I guess because I was really confused and that sounded like it was talking about things that were weird, so I thought it might help. I think I maybe spent more time in that file than maybe I should have done.
- Participant: As for process, when I felt confused enough by a single file I decided just to move on, cut my losses, and check something else. See if there was something else that might help.
- Interviewer: Okay, so it sounds like from a high-level starting point you went through a process of "What currently seems relevant? I'll try to go to that file, and if I get too confused I'll go elsewhere"
- Participant: Yes. Pretty much.

- Interviewer: Okay. And I guess, following on from that, what would you say your process of then trying to integrate information from different places was?
- Participant: Well I started making some naive notes, but they didn't really get very far. I started using it because so many of these I wasn't finding as relevant at the time, so I wanted to keep track of which ones I'd opened that I liked. Something to go back to later. I guess that was my approach, but I didn't really follow through with it. I ended up sort of branching out rather than digging deep.
- Interviewer: You mentioned that you felt kind of disoriented in all the information. Well, I used that word, but yeah. What do you think the source of that was? Just how much information there was, or?
- Participant: I think it's the volume and the unfamiliarity. And I think, like I was saying, about the structure as well definitely didn't help. It's like a sea of information, it's kind of hard to... uh, yeah.
- Interviewer: Do you think that having some kind of visual representation of the relationships between resources would have been helpful in alleviating some of your disorientation? So, for instance, if you could see visual lines that linked related portions of different windows or if you could temporarily view all the resources in some kind of map, to add some kind of structure to them. Do you think that would have been helpful?
- Participant: I think so. Some kind of mind map showing "If you want to read more about this go here" sort of thing? Yeah definitely. It would help me get to places that are more relevant more quickly rather than me going back to this directory thinking "What looks good?". Basically it would do that job for me, and help me get to the information a lot faster.
- Interviewer: Related to the topic of going back to things that you found relevant, did you find it limiting at all that --- unlike research on the Web --- you didn't have any kind of history?
- Participant: Yeah. There were a few times that I wanted just to search a term, or like right click something and 'open in new tab'. But yeah, definitely.
- Interviewer: Okay, so searching is something you thought would be useful. But also was the lack of history specifically a problem for you do you think, or not?
- Participant: I don't think especially. The main thing was remembering what files I had read that were relevant, which is where my notes started coming in handy.
- Interviewer: I guess if we're comparing to the Web, then, some kind of bookmarks is maybe what you missed?
- Participant: Yeah.
- Interviewer: Okay. So did the process of searching through the information feel efficient or inefficient to you?
- Participant: Inefficient, I'd say.
- Interviewer: Why would you say that is?
- Participant: Uhh, I had to do a lot of work. There was a lot of stuff to

- sift through, and some of it either didn't seem relevant or seemed like it would require more time to understand than I felt I had in the time available. So it was kind of hard to find exactly what I needed to complete the task, knowing that time over my head.
- Especially when I realised that one file was like 70 pages long.
- Interviewer: Yeah. And I guess the other angle to this is did you feel the actual mechanics of switching between different windows to be efficient or inefficient?
- Participant: It's manageable, but again with my small screen it's not particularly efficient I'd say. If I had my two monitors at home, I'd be able to keep this [Points to the hex dump] open all the time, for example, and have a couple of other things open too.
- Interviewer: Okay. So how did you deal with managing the windows in a constrained space like that?
- Participant: Uhh, I used the space as best I could I guess. I had to keep switching between a lot, which I guess isn't ideal. I guess maybe I lost my trail of thought sometimes trying to look for windows I thought I wanted again. If that makes sense?
- Interviewer: Yeah. Related to windows: while you were carrying out the task, did you consciously pay very much attention to the program you were in --- so for instance Chrome or your PDF viewer --- or were you more focused on the actual resources that you were trying to look at?
- Participant: I was more focused on the resources, but I was a little bit annoyed at having to open things in different places. It would have been more convenient to have it all in one place I think. I was a bit confused as to why the PDFs opened in different windows instead of different tabs, I guess I have the older version on here or something, which was a shame. I might update that later.
- Interviewer: So it seems like you're kind of saying that your primary attention was on the resources, since they're what you're using to accomplish the task, but that maybe you noticed the programs when they got in your way? Would you say that's accurate?
- Participants: Switching between programs got in my way, yeah. I'd say so.
- Interviewer: Related to tabs, I guess there's a bit of an inconsistency there if to access most of the resources you have to change windows, but to access one of the tabbed resources you have to change window and then change tab. Did you find that indirection to be sort of...
- Participant: Definitely. It was also difficult to remember where things were when I wanted them. So I might have gone here first, then "That's not it. Is it this one? No. Maybe it's this one?" [Attempts to find a specific resource among a number of windows]. And with the tabs, sometimes they weren't particularly well titled, so if I wanted to find something I would have to scroll quickly just to get my bearings again. So labelling is definitely an issue as well.
- Interviewer: Interesting. So I guess the mechanics of actually navigating the resources did sort of impede your progress as well as just the informational content of all the different resources.
- Participant: Yes. Definitely not to the extent that the information did though, but to an extent yeah.

Interviewer: In the file format reference, there are some links internal to the PDF that help you to get to different sections. Did you make use of that at all?

Participant: Not in the file format reference, but I think there was another file where I did.

[Participant navigates between various files for a few moments trying to find a specific file]

Participant: See, now I have a problem looking for it again! [Laughter]

Participant: I can't find it, let's just say it was the file format reference.

Interviewer: Did you find the internal linking useful?

Participant: Yeah, it saved scrolling. Especially when I realised how long it was, it meant I could more quickly parse what was in the document.

Interviewer: So going back to this question of history again, did you find it annoying at all that when you click a link inside there you have no way of going back? So you can't navigate straight back to the contents, for instance.

Participant: It wasn't a big issue because it's available on the left, but if that wasn't available it would have been an issue. I would have had to scroll back or change the page number of whatever.

Interviewer: Okay. Now we're going to change direction slightly to some more hypothetical questions.

Interviewer: First of all, if it were possible, do you think it would have been helpful to build your understanding if you could organise the information by structuring it digitally in some way by---

[Recording temporarily cuts out]

Participant: Yeah, I think it would. I think it would have to be easy enough to do that it wouldn't be an annoying task in itself. If that makes sense? Because I don't want to spend time playing around with linking things up if it's going to take away from time I could have spent reading --- especially with the 45 minute time limit. But if it was as easy as just "add, click, link": yeah, for sure.

Interviewer: Okay. So if we were to say that cross-program hyperlinks is the way you were doing that, do you think that it would be helpful to be able to link to particular sections, or maybe specific paragraphs or lines of documents, rather than only being able to link to the start of documents?

Participant: Most definitely, yeah. Easy answer. I don't want to scroll through 70 pages, yeah. Could you even link to multiple disparate places in the same file as well? That would be cool. Because there could be both something on page 10 and something on page 60 that are relevant.

Interviewer: Yeah, for sure. So that was you structuring information digitally. I guess the other side of that is if someone had already gone through the process of making sense of the information and creating links perhaps through the resources. To what extent do you think that would have been useful to you?

Participant: I think it definitely would have been useful as a starting

- point. Just because like I was saying earlier, it was hard to know where to start with everything in quite a flat structure. But also it would be nice to add to something like that, or to even change things if your brain maybe doesn't work in the same way or you don't make the same link. So yeah, it would be useful to an extent as a starting point, but not as a final point.
- Interviewer: Okay, yeah that makes complete sense. And in receiving this kind of structure from somebody else, what level of structure do you think would be most useful? So, do you think you'd want something that's very user-driven and that you could freely explore as you wish, or do you think that maybe something that's more of a sequential kind of guided tour through resources would be more helpful?
- Participant: Hmm... I mean, if the person who set up the guided tour knew exactly what you needed to know and the best way to know it, then sure I'll go along for the ride. But if it's just to understand a concept on the whole, probably more of a Web would be more useful. So you think "Oh I want to know more about that, so I'll go there", and the link has been made for you so you know where to go. It's kind of like a tube map --- you want to go there, you see the links there, so you follow it. Does that make sense? I guess it depends on how much knowledge the person who made the links had.
- Interviewer: If someone had made these very kind of guided paths through resources, on the topic of maybe fundamentals that would be useful for you in understanding the bigger picture --- so, for instance, on the header of a Mach-0 file --- do you think it would have been useful if you had maybe guided tours through some of the fundamentals and then you could more freely explore to appreciate some of the details that are less likely to be present in a more strict tour?
- Participant: Yeah, definitely. I found it hard just to find the fundamentals, as the information was scattered throughout various articles, so definitely.
- Interviewer: Yeah. So I think that's all the questions I have. Is there anything else you'd like to add?
- Participant: No, I don't think so. I hope that was helpful! Sorry I didn't get far with the task.
- Interviewer: No no, thank you very much for being a part of this. It was really useful.

Participant 3

- Interviewer: First of all, tell me a little bit about your experience.
- Participant: At the start, I had no idea what anything was. It was quite difficult just going into a new topic like that, but it was interesting. I worked my way top down through the resources that I was provided with. Some stuff made sense after reading multiple resources, but it's still not all there.
- Interviewer: Okay. And how did you feel when you were carrying out the task?

- Participant: At the start I was very confused as to what everything was. I guess because of my lack of knowledge in this area. After that, once I started reading and getting more knowledge, I started to feel a little bit better. I guess at the start you could say it wasn't a very good state because you feel silly for not knowing, and then once you start learning then you begin to understand more, which feels better. And then I felt it was quite interesting towards the end.
- Interviewer: Okay, cool. Do you think you made effective use of the resources provided?
- Participant: No. The resources had varying length, and for some of the resources I would scroll a lot longer or read longer than some of the others. And I didn't even end up using like the last quarter of the resources. So I think I managed my time poorly.
- Interviewer: Could you tell me a little bit about what you actually found out? What did you learn about?
- [Discussion of specific details of the solution]
- Interviewer: You said that towards the end you were starting to feel much better about things. Do you feel like you really got into the research process by that point, or was it still a bit of an uphill battle?
- Participant: It got easier once I started learning what the terms were. Once you find the core definitions, it begins to make sense in the bigger picture of how they're used in different resources. So I thought that was quite helpful, but it was still a bit of an uphill battle because I didn't know anything about the subject.
- Interviewer: Why do you think you went for the approach of looking sequentially through the list of resources from the top to the bottom rather than using some other approach to guide your navigation?
- Participant: That was just my thoughts of "When in doubt, go for the start". I didn't think the files were grouped too much, and I didn't understand what anything was at the start, so it seemed like as good an approach as any other to expose myself to the information. I don't know if that was the right approach, because once there were terms that I understood then I probably could have looked in different files for the information, but I thought I would get more knowledge just by going sequentially.
- Interviewer: Okay. So you mentioned there that there wasn't a real order, so do you think you were getting kind of a random sample of the documents by doing it that way?
- Participant: Yeah. If there was an order or I better understood what I was doing then I would maybe change what I was doing and not go sequentially, but it just depends.
- Interviewer: There was quite a lot of information provided. How did you deal with sorting through it all?
- Participant: I looked for the key phrases that were used a lot, like "dynamic linking", "loading", "libraries", "executables" and all that, and it was just a process of reading through the resources and connecting what I read I guess.
- Interviewer: Did you feel disoriented in the information at all? How

- grounded did you feel with all the information that was available?
- Participant: I feel that given enough time it would have been fine --- 45 minutes is definitely way too short to even begin to grasp most of the concepts here. If you just say "Here take this, look at it and try to figure it out", it's quite overbearing because there's so much information, but it was interesting so I felt okay with it.
- Interviewer: If you had some kind of visual representation of the relationships between the documents --- for instance, if you could see visual lines that link related portions of different documents, or if you could view the documents within some kind of map --- do you think that kind of structure would have been helpful?
- Participant: Yeah, absolutely. I think most people would find that a lot easier, because people learn visually quite easily. For me, a lot of the time I need context. It's great reading a resource and reading definitions and stuff, but sometimes it's nice to see practical examples or actual working demonstrations of what it is you're looking at and how it all fits together.
- Interviewer: Did the process of searching through the information in general feel efficient or inefficient would you say?
- Participant: It's hard to say because I'm the one doing it. I would say inefficient, but that's because I didn't really understand. If I were to do something like this again, it could be more efficient, so it's a process of learning.
- Interviewer: And what about all the switching between programs and windows? Did that feel efficient or inefficient to you?
- [Pause]
- Interviewer: Or was it something you even noticed as an issue? Was it just completely normal?
- Participant: I didn't really think about it too much. I usually just use Alt+Tab to switch between windows, and sometimes I use my mouse. I don't really think about it too much.
- Interviewer: Okay. In a similar vein, did you consciously pay very much attention to what program you were in, or when you were switching were you concentrating more on the particular resource that you wanted to navigate to?
- Participant: I think it was about resource navigation really. Everything was centred around that.
- Interviewer: I noticed that your approach for integrating information mostly relied on having multiple programs open in full screen, and switching rapidly between these to compare the information. Why do you think you went for this approach rather than, say, using a 'split view'?
- Participant: If I had it in split windows I probably would have to scroll around anyway, and Alt+Tab'ing is just as fast if I needed to refer back. I can understand why a lot of people would use dual windows, but I just decided not to.
- Interviewer: If you were doing this research on the Web, things would be a little bit different. For one, you would have history capabilities --- you could go and see where you've been and navigate back and do

- whatever else. What's your feeling about not having these things? Did you miss some of those things?
- Participant: Yeah, just having a directory of resources was harder I think than using a search engine like Google. Once you learn a term, you want to know as much about it as possible, and rather than using the resources if I'd used a search engine then it would have been a lot quicker.
- Interviewer: Okay. So do you think if you had some kind of search capabilities within the directory that would have been equally valuable?
- Participant: Yeah, I think that would have helped a lot.
- Interviewer: Cool. So now we move into more hypothetical questions.
- Interviewer: If it were possible, do you think it would have been helpful for you to organise the information by structurally it digitally as you were figuring things out? So, for instance, if you could create cross-program hyperlinks between resources or something, do you think that would have been helpful?
- Participant: Yeah, definitely. It just keeps all the information you need together, and keeps it cleaner I guess.
- Interviewer: And in doing so, do you think in such a system it would be useful to be able to link to specific sections, paragraphs, or sentences in a document rather than just linking to the document as a whole?
- Participant: Oh yeah, definitely. Because a lot of the time the articles could be about various different things and maybe only a chapter or specific section is relevant. That would be much more helpful than linking to the whole thing.
- Interviewer: And I guess on this flip side of this, rather than having to create this kind of structure yourself: if somebody had already gone through the process of making sense of the information and had created links between the appropriate parts of the resources, do you think making use of that kind of structure would have been useful?
- Participant: Yeah, as long as the implementation was done right then that would be fine.
- Interviewer: What do you mean by 'as long as the implementation is done right'?
- Participant: When it comes to learning and information you have to make sure that the person who did all the gathering and sorting did it right, and that would mean having to read for yourself to make sure that their understanding was correct. I guess it's a question of trust.
- Interviewer: And in a system like that, what level of structure do you think would be most useful? So would you want something that's very freely explorable, much like regular hyperlinks are, or do you think it would be more helpful to have something that's closer to a guided tour through particular resources?
- Participant: For this example, probably a guided tour. Usually I like free learning because you get a lot of context, but in this case it's a very technical, specialised area. People who've written the

articles know what they're doing, and as someone that's new to it you aren't so... I guess a guided tour would be good to start of with the aim of being able to freely explore later.

Interviewer: By the way, I noticed that your PDF viewer uses tabs for open documents. This might not even be something that you've really thought about, but did you feel that the extra indirection of having to navigate through tabs as well as windows effect you at all?

Participant: Yeah, it meant I had to use the mouse instead of just the keyboard. Unless you know the exact hotkey or whatever to switch tabs inside the PDF viewer, then it makes navigation less efficient.

Interviewer: Okay, I think that's all I have to ask. I don't know if there's anything else you have to add?

Participant: Nope.

Interviewer: Okay, well thank you very much, I hope you learned something.

Participant: No problem, my pleasure.

Participant 4

Interviewer: So first of all, tell me a bit about your general experience.

Participant: Uhh... in which regard?

Interviewer: Just your general progress, how you felt about it --- that kind of thing.

Participant: I felt I've got a lot to learn in a lot of different places. So I had to sort of first try and land somewhere useful and then use that to branch out. Not sure about the progress... I think I started making progress at the end in terms of mapping things, but it seemed to be everything needed information from another thing, so it was like "So I'll go to that one but let me understand this first" and then that never ended.

Interviewer: Did you find the task difficult?

Participant: Yes. I'd say difficult within the time constraints. If I had enough time --- well, maybe this counts as difficult --- but if I had more time I would have gotten it eventually. I don't think it was impossible.

Interviewer: Okay. And how would you describe--- what was your emotional response through different stages of the task?

Participant: At the start it was "Oh god. Oh god. Oh god. Oh god." [Laughter], then sort of the second quarter or so was "Oh okay, I've just got to read some documents", then it was "Oh god, I have to read some documents" [Laughter]. Towards the end, there were lots of sort of stages of "epiphany, oh no, epiphany, oh no". Where you learn something new, then you have to sort of bring that into what you already know and reinterpret what you already knew and see if it hasn't sort of changed some stuff. Like sort of near the end, I thought these [Points at top of hex dump] were the actual sections and stuff, but actually I think they point to data later in the file with the data contained in those sections.

Interviewer: That's correct.

- Participant: So my understanding sort of changed at the end. I had read something new, and I had to reinterpret everything.
- Interviewer: I guess do you want to just give me a brief overview of what you now understand in general?
- Participant: Oh god [Laughter], okay.
[Discussion of specific details of the solution]
- Interviewer: Do you think you made effective use of the resources that were provided?
- Participant: No, because there were tonnes of files that I didn't open. Some which I kind of knew, like a PDF on endianness I thought "I don't need to open this", but some of them I was opening just to see what they did. Like this [Points to specific PDF], I thought this was quite nice. For a lot of them, though, I needed to understand the big bit first before I got into the smaller parts.
- Interviewer: Given that that was the case, how did you go about finding a good starting point?
- Participant: So I think I started on executables. Because I was like "This is an executable", so I opened the PDF about executables. I think I read that and thought "Well okay, it hasn't really informed me much, it's just... well, yeah, it's an executable", so I think I opened like a few of them just to sort of catch a landing place. I think I settled on starting with documents about Mach-0. And each one sort of told me a new file to open as well, and when I felt comfortable I could refer to 'reloc.h' and stuff like that. So it was sort of like "Go wide and see what catches" and then once you're there, see what that references. I think that was my approach.
- Interviewer: Do you feel that generally you got into the flow of the research process by the end, or was it more of an uphill battle all the way?
- Participant: Yeah, I think after a while I got used to the information. The first bit is getting used to the problem space and trying to understand what all of this means. But once you've got the lingo down, it's more of trying to find the next piece of the puzzle as opposed to trying to understand what the hell anything is doing. So I think by the end it was just more of the same really, just trying to fill up what I knew and reinterpret what I knew. So I think by the end I was in the flow a bit more.
- Interviewer: And, I guess we've addressed this to an extent around finding a starting point, but more generally how did you deal with sorting through the wealth of information that was made available to you?
- Participant: Hmm... the names really. The titles of the files and sorting just by what I thought would be most relevant when. So not knowing anything I thought it would be more relevant to look up the file structure before I look up, you know, peculiarities about it, if that makes sense? I'm not ready to read that at the start, whereas I think I'm ready to read the more general stuff at the start. And as my understanding grows, more files become ready to be understood, if that makes sense?

- Interviewer: Yeah that makes sense. And what about within documents? A lot of the documents are quite long, so how did you figure out what was important and what wasn't?
- Participant: Uhh, so... a lot of skim reading really. And sort of reading the first parts. So this one was really helpful [Opens a PDF resource], and it was reading the first section, which took quite a lot of time, and things like--- [Looks at a diagram within the PDF] oh that would have helped, that's quite stupid, I skipped that and I was like "I'll come back to that", but yeah that would have helped. Uhh, but yeah, it's kind of just like the reading of it, then certain things caught my eye. It was just a lot of skim reading and seeing what that points to, or hoping that the first section has a sort of contents-y bit. And then trying to find something that matches that, or knowing what I need to find and trying to find something that matches that.
- Interviewer: Okay. Could you try and describe your general process of integrating information from different resources?
- Participant: So... in my head I was thinking about it more like cross-referencing, if that's what you mean? So I would open up a file, read about it, then when it refers to another file have that. And I guess I was integrating that all in my head, if that makes sense? So going "Well that from there, and that from there, oh and now I understand this", and then I was trying near the end to put all of that together --- because you don't want to keep holding all that in your head --- so that's when I started to write all these little notes next to each part. And then reading a new thing, I would re-read these [Looks at notes] to make sure that they still made sense with my new understanding. So I guess it was sort of "read, bring into what you know, and then sort of cross-reference for future reference".
- Interviewer: That makes sense. I think from observing your process, it seems like --- correct me if I'm wrong here --- but maybe that's reflected slightly in the way that you managed the windows for your resources too. Because mostly at the start you were using lots of full screen windows, and you would go through and read, then flick back and try and check something, then go back and continue reading. But then also you had other times when you did use a split view to look at two files at the same time, which is maybe more directly comparing and integrating from the two, versus the more cross-referencing-based approach that you just described.
- Participant: Yeah, yes. That was when I was like "Oh yeah I can split screen. Let's do that!" [Laughter]. Because before I did this, the information was a bit more abstract in terms of what I could cross-reference, whereas with this flicking back and forth with individual numbers is not fun. So being able to have them directly side-by-side is a lot easier. So I think when I switched to this, it was because of cross-referencing being a lot harder with individual hex bytes.
- Interviewer: Okay, yeah. Something I wrote down is, sort of, ask why you

had switched to a split screen perspective for this thing versus earlier when, say, you were comparing the sections in the file format reference to the sections in the ASCII in the file. Would you say it's just because of that extra load from all the numeric stuff, or?

Participant: The extra load, and also being stupid and forgetting that I can split screen. Because, I don't know, I guess it takes like an extra millisecond to like put two screens together, and for some reason my brain was like... I think the problem was that this wasn't, umm, like it didn't cause me enough stress or problems or mental work to flick back and forth. So I guess my brain considered it less effort to do that, to flick back and forth... which I think is objectively wrong from afar, but in the moment when my brain is trying to fix the problem rather than think about how I'm doing things, I think flicking back and forth is less strenuous than the millisecond --- [Laughter] all I'm saying is, right --- I can't believe I'm saying this, but I think this is what was going on: like, the millisecond it would take to use split view didn't feel worth it. But as soon as I get to something a lot more strenuous --- back and forth of numbers --- my brain goes "Hang on a sec, there's an easier way: let's split the screen". I think that's what was going on.

Interviewer: Okay, interesting. Did you feel disoriented in all the information at all?

Participant: Right at the start, yes. Because it's like "Here you go". Once I sort of managed to, like, mark things and say "Well this helps this information and this helps this information", my like fear of the unknown went down progressively. So I'd sort of say it was like a--- it got much less stressful quite quickly once I understood the problem. Then at the end, as I began to understand each file individually, sort of the trickling of--- because still now there's some that I still don't know. But I'm sure that if I had more time, I would begin to piece puzzles and begin to understand why that file was given to me. But yeah, so I guess it's sort of very high, went down quite a lot --- because a lot of the files just from the name you can understand the sort of context --- and I felt able to approach a lot of the files quite quickly. Then at the end it was sort of just "read a file and hopefully that makes me able to approach another file", but it sort of slows down in terms of trickling at the end. Though some of them I still don't know how they're going to help me.

Interviewer: Okay. So in building all of that kind of implicit structure in your head around this unstructured mess of files, did you ever feel like--- did you ever forget what a particular file was about or lose any part of that structure?

Participant: It's a lot harder to--- because like I say if you're trying to hold it all in your head, it's a lot harder to remember where everything is, which is why I started writing things down. So yes is my answer. Hmm... what else can I say on that?

Interviewer: Okay, so this leads into the next thing I have written here, which is: do you think that having some kind of visual representation

- of the relationships between resources --- I guess an explicit representation of the structure in your head --- would have been helpful? So maybe if you could display the resources on some kind of map, or something like that.
- Participant: Umm, yeah. One thing I will say, I'm not sure if this counts or makes it worse, but for the entire thing I really felt like I wanted a notepad to do that, if that makes sense?
- Interviewer: Okay, yeah.
- Participant: So I had this really strong itch --- and I don't know if it's because I've been doing loads of revision and work --- but I really had this strong itch of having a whiteboard or a notepad or something to be able to jot things down and draw --- like being able to link things. And that was one thing that from the get go I felt. I feel like I can't --- it's probably because I don't have the right tools here --- like if I had a tablet or then maybe something electronic could satiate that, but I didn't feel like opening a tab in Sublime, it isn't the same, if that makes sense? So that's one thing I was really missing throughout the process. Being able to put what I've got --- because that's when I had to default to doing everything in my head --- but really what I would have done is put it on something. And for me that's paper. Just because it's a lot easier and dynamic. But I feel like if I had something on the computer that filled the same sort of niche, that could be useful.
- Interviewer: Okay, so on that topic... I'll re-order the questions I was going to ask slightly. If it were possible, do you think it would have helped build your understanding to be able to organise and structure the information digitally in some way? So for instance if you could create cross-program hyperlinks as you were understanding the information.
- Participants: Yeah. Yeah, yeah. So at the start, you've just got that [Opens the directory of resources]. So if I was browsing Wikipedia on my own, things would link to other things, so I'm assuming that's kind of how what you're proposing would work. And being able to go back in and say "Well this references that", and be able to actually map to the section in another file exactly so that when I come back and go "What did that mean again?" I can just go "Ooh" rather than have to re-do that process again, I think that would be very useful, yes.
- Interviewer: Okay. And I guess that kind of also fits in with your desire for a notepad --- potentially that need could be filled by such a system.
- Participant: Yeah. I don't mean the physical notepad itself, I mean the role that it plays. Something to coalesce all my thoughts and link things is what I really wanted. But I had to default to my brain, because at least personally I don't feel comfortable doing that in a text editor. I'm a lot more visual or something. So yeah, it was that sort of role that I was really wanting.
- Interviewer: Okay. Also on this topic, if we're going back to this cross-program hyperlink digital thing, if you link to stuff on the

Web today usually you'll just link to the entirety of a resource rather than a particular section of a particular resource. I guess, do you think that would be problematic? Would you want more granular behaviour in a system like this, or do you think that would be unnecessary?

Participant: I think linking to a specific part of a document just speeds up the process. For example, if something referenced the `__text` section, then yes I know how to read and go down to the text section, but if it just went "Bam. Here's the `__text` section." then it just makes the process a lot more seamless. It keeps the mental burden down and just makes it smooth. It's not impossible to do it the other way, but it's one of those quality of life things --- you've just made my life a little bit easier because there's one less thing to worry about, so I can keep going. As opposed to "What was I doing? Okay, go to the `__text` section, right now I can carry on what I was doing". So I think it's not essential, but it's one of those quality of life things where it's a big help.

Interviewer: Okay. I guess that's kind of similar to what you were saying about the notepad, where you could have done everything through a text editor but it feels more natural to be able to link things in a much easier way, and actually that makes a big difference to the overall process even though it's just a small change.

Participant: Yeah, exactly. It's less 'start-stoppy' and more just "get to the destination".

Interviewer: Yeah. I think had another question around this topic too... oh yeah. So there was this point where you were reading through the file format reference, and I think you saw in the text somewhere it had a reference to another document, it said something like "See the position independent code document". You then went and looked at the position independent code document, but then promptly just ignored it and closed it. What was the process there?

Participant: Because it was when I was saying--- like, this whole--- I'm just going to keep saying what I said before. [Laughter] I didn't feel ready to open that file, so I opened it to quickly see if it was helpful, but then my brain sort of cached it for later and said "I know this will be useful, as I've seen it referenced a lot, but I think it will be more useful---", so it's a sort of priority thing. I thought it would be more useful to completely understand what I was currently reading as opposed to "Stop, and start reading an entirely new section". Because I felt it was different enough of a topic, where I would need to like--- so that's why I opened it quickly to see if it was just something to skim and add to my understanding, but when it seemed like it needed to be --- at least in the moment --- that it wouldn't immediately add to what I need to know for this part, I thought I'll queue it up and let's carry on reading what I know and then come back to this other stuff later.

Interviewer: Okay. Do you think if that particular reference had actually been a link to a particular section that was relevant to the specific thing the referencing document was talking about, do you think that

maybe that could have meant you would delve into the details just to have a look at the specifics there?

Participant: Uhh... probably. I can't say for sure, because I don't even remember what I saw on that page. But uh, yeah. If you think about it like that, it's kind of as if it was an aside at the bottom of the--- it depends on how--- if it points to a section but that section was really big then I think I'd still default to "Okay, come back to it" --- but at least I would have a handy reference to it! But if it was really small then it's kind of like an aside, so I could just click on it, see it, go "Oh okay cool" and then come back, then yeah that probably would have made me read it fully. Even a paragraph I would have been like "Okay cool, read, go back". Because I don't have to leave what I have in my brain --- I don't have to leave that behind, I can add something new to it. Whereas if you have an entirely new section you kind of need to store what you have, approach this, then once you go back, come back. Then you've got the mental task of "So what did you know? How do you bring that in?" or you've got the problem that they might be too different, and you've got the problem of trying to remember where you were and why you were doing it and things like that, which is a much higher burden. So being able to queue it up in that case is a lot easier.

Interviewer: Okay, that makes sense. So you mentioned there that you might be able to just have a look and then go back to where you were. On that topic of history, if you were usually researching on the Web you would have this ability to keep track of the exact places you've been, and easily go back and forwards through this chain. But because we were dealing with heterogeneous programs and just a directory of resources, you don't have any of that. Did you find that an annoyance at any time?

Participant: Kind of, yes. So sometimes Wikipedia might link to something and you can just click on it, whereas here I had to remember where it was and try to find it. I think I kind of maintained that by having tabs, and the tabs keeping their position kind of made up for the lack of history. Assuming that the back and forward buttons took you to exactly the right places --- because sometimes they don't, they just take you to the start of a page --- but at least with the tabs it's exactly where you left it positioned on the screen. So you could just jump back to it. But the problem with that is that you have this enormous chain of tabs. So the problem moves from "Where in the page was I?" to "Which tab was I on?". I was getting there, the list was growing, but at that point I could still see the names of the tabs so yeah I think I was using separate tabs to fill that space.

Interviewer: Okay. Do you think if you'd have kept going that this would have become more of a problem then?

Participant: Yes. So if I had to keep growing it--- I don't know if I would have felt comfortable with closing a tab eventually, once I felt like I fully sort of exhausted it and felt comfortable with its contents that I wouldn't need to come back to it. But I don't know if the speed of that would match the speed of needing a new tab. So at

some point it might have been helpful to do something else.
Interviewer: Okay. Also on this topic of tabs, I guess there's sort of an inconsistency in that the resources in the directory were of several different files types. And to navigate between many of the resources, you needed to change windows. But if you wanted to go from, say, the objdump text file to a particular PDF resource, you not only needed to change window but you also needed to change tab as well. Did that indirection feel like it got in the way at all, or not?

[Pause]

Participant: Probably. I feel like I've been so conditioned to having to switch program for file type that it's less of an issue, because that's what I'm used to. But if the dump had have just been another tab, then it becomes one level of indirection as opposed to two. So objectively yes. In the moment, probably in terms of a minuscule--- like I said before you want it to be as seamless as possible because you want to jump to the other place and continue your train of thought. But when you have to sort of split that and in the middle find the next thing, then that also kind of sometimes takes brain power. Especially if you're stupid like me. [Laughter] So yeah, I'm sure it impacted my process, but I don't think it's as big of a problem as trying to find a place. If this, for example, [Points to tabs] had been really long, then if it wasn't the last thing I had open then yes I think it would have been more of a problem. But because everything was still visible, and if I wasn't on the page I could just take a quick glimpse and read. Or if for example I had three tabs open that all started with Mach-0, which thankfully I didn't, then yes I'd have to try to figure it out. So in this specific scenario it wasn't too much of a problem, as the number of tabs was manageable, but if the list were to have grown, then it would become more of a problem.

Interviewer: Related to that, did you find the switching between different programs and windows to be efficient or inefficient in general?

Participant: Like I said, I'd probably say I'm conditioned to it. There were a few situations where, you might have seen, when I was trying to get a file to open with Sublime, that was taking me out of being able to solve the problem, because I had to fight with the program. So if it was all just a seamless workflow then it would have been a lot nicer. It was the bit of, yes, having to switch between programs --- but thankfully in this case I only really had two, each with a manageable number of tabs --- but if, for example, Chrome couldn't open PDFs then I would also have a PDF viewer, HTML viewer, and the text editor. And then it becomes a bit more of a juggle. If it's kind of different then I don't mind it being a different program, but it's when they're kind of more linked. Like, the Excel file is sort of just an aside to help you, so having it separate is actually even better because it's easier to find. But when you've got all of these sort of bits of data that refer to each other, then it's a lot easier to --- I don't know, find stuff? I forgot what I was trying to get at

- here.
- Interviewer: Okay, no that's fine. While carrying out the task, did you consciously pay very much attention to what program you were in, or were you more focused on the resources?
- Participant: The resources. Because I felt kind of comfortable enough with the programs that they weren't really a problem, apart from fighting with them occasionally to get things working. But once it was there, I felt I could just focus on what was on the screen as opposed to having to worry about manipulating the programs. So like, I happened to know the shortcut to split screen in Sublime. Otherwise I would have to probably thought a bit, or like I said before if I didn't know the shortcut then who knows if the effort of trying to find the split screen would have been less than trying to like cross-reference the numbers. So yeah, I think a lot more of my time was spent more on the resources than on the programs, other than when they misbehaved.
- Interviewer: Okay. So in saying that, I guess you're kind of saying that the tabs are just another mechanism through which you could get to individual resources. Would you say that's accurate?
- Participant: Yes. I think so? So, what exactly do you mean?
- Interviewer: So, I guess I'm saying that if you felt like you were more focused on individual resources than programs, then switching windows and switching tabs are kind of just ways of getting at resources.
- Participant: Yeah. But then if there were more tabs and I couldn't easily switch between them, then more time would be spent in Chrome trying to find where that is, and the same applies to Sublime. As the tab titles become illegible, then I now have to impart effort in trying to find where things are. But at least for now, thankfully it wasn't too bad.
- Interviewer: Okay. You also said a minute ago that you thought it was useful that Excel was totally separate. If it were the case that instead of having all these tabs, everything was actually in a separate window --- like a uniform space of stuff --- what do you think the impact of that would be for you?
- Participant: Umm, so you mean like the Chrome OS approach where everything is just a tab?
- Interviewer: So, I guess like that. I was kind of more thinking about everything being a window than everything being a tab though.
- Participant: Okay, so you could still manipulate and put things side-by-side, yeah yeah. It depends on how... I think it would be useful, because then resources like this and this [Points to two different resources]... it's like a middle ground, you have to be careful, because on the one hand sometimes files are different and mentally being able to know that "For this type of file go to this program", but I'm not sure if that's just because I've been conditioned to that approach. If the text files were here as well [Points to Chrome], I guess it would help? But at the same time it adds to the difficulty of finding things. Separating things by program and tab is kind of like a two stage directory of "Go to here,

and then you find what you're got", whereas if instead you've just got "Try and find it" I think that's quite a big problem. As long as you've got a way of categorising. I think what you lose with that is the categorisation of information, but if you can maintain that then I don't think that's a problem. I guess I've been conditioned to look for something based on what kind of file it is, and then from there continue looking. And I think that's still easier than trying to find it amongst everything? So as long as you can still categorise things... yeah. As long as you can still categorise things and get to things quickly, I think it's not a problem. It might help in terms of the mental, so like I said before because things are all grouped as tabs in one window, if they were all separate windows then I think my connection with them and trying to cross-reference them might be easier, because I'd probably find that my desire to start splitting screens and putting things side-by-side would be a lot easier. Because I don't have to worry about pulling the tab out and then splitting it --- these are very small tasks, but in the moment your brain just doesn't think about them. If it's individual pieces of work, then I'm probably going to be a lot more willing to manipulate it. So I think it would be good for that, but you still have to maintain the ability to find everything.

Interviewer: Okay. So, for instance, if there was a single space of windows and there are no tabs, and you could, for instance, have some kind of search bar you could bring up to search through all current open windows. Do you think that might be, just as an example, a useful way of doing that kind of organisation?

Participant: Uhh... yeah. So as long as you know the title of the web page or the title of the text file that's open, then that's probably a good substitute. Yeah, I was trying to think of a downside for you, but yeah! I think that--- other than you now have to type, but I don't think that's the end of the world. There are ways of--- [Activates Windows 10's Task View] even though I had things arranged differently, now it's brought all of my information so that they're all similar so I can quickly look through stuff like that. So ways of sort of making it easier to find things are good.

Interviewer: I guess this is sort of related to some of the stuff we were talking about earlier, like displaying resources on a map or something like that. It's like a different spatial arrangement of the resources.

Participant: Yeah. So if I just had loads of windows and I just wanted to try and find something, as long as I know what I'm looking for then being able to search makes that very easy. The problem comes when I don't know what I'm looking for, but I guess that's a different problem entirely.

Interviewer: So we've talked about the process of you structuring digital information yourself using cross-program hyperlinks. I guess, the other side of that is that if someone had already gone through the process of making sense of the information and had created links between appropriate parts, do you think making use of someone's

- already-created structure would have been helpful?
- Participant: Yes and no. Yes, because it makes my life easier. But making my life easier might not make the problem easier. Because in trying to form those cross-references, I'm making my brain work, and the more it consciously is working the more likely I am to remember something. So, you know, I don't know how much validity this argument has, but learning things is creating links between things, so the act of you making those references is helpful. Whereas if they're given to you, you don't naturally have the link between the two. But your life is made easier, so you do have more time to solve the problem. I guess that would be something you'd have to compare, what's better in the long run. But if the problem's still hard enough, you're probably still engaging your brain, and that's just a quality of life improvement. But yeah, I can't tell you if that quality of life improvement is something that takes away from your learning experience.
- Interviewer: Okay. I guess those two things aren't mutually exclusive either. You could have something that's already linked and create your own structure on top of it, or maybe it's about different kinds of problems are better suited to different approaches. So maybe for more exploratory tasks like this you might want to create things yourself, whereas for other things you might just want to explore.
- Participant: Yeah! So if I'm trying to look up a topic, then I don't want to create these links probably. If I'm trying to break something down, then the act of making those links is actually me doing that decomposition. Creating the links is actually part of me solving the problem. Whereas if I'm just trying to read up on my favourite thing, then I just want it to tell me what's referring to what. So I think context means a lot there. Because it sort of becomes part of the problem in one area, whereas in the other it becomes part of the solution.
- Interviewer: Yeah, okay. Do you think if someone had already created a bunch of links, do you think that would have reduced or increased any sense of disorientation you have? Kind of having lots of additional information, and potentially noise, in the resources.
- Participant: Umm... I think...
- [Pause]
- Participant: I don't think it would have created noise. It might have reduced my general understanding of the files themselves, which might hamper my progress later on, but yeah I guess if I needed anything, that would also refer to the right section, so maybe I don't need a more general understanding. But no I don't think it would add to the mess.
- Interviewer: And also, if someone had already made this kind of thing, what level of structure in that do you think would have been most useful? So, do you think something that you can explore at your own rate and go where you want or something that's closer to a sequential guided tour through the resources would have been more helpful?
- Participant: Both?

[Pause]

Participant: I think a guided tour helps as long as I'm allowed to say "No" and jump ship. Because not everyone has the same way of traversing the content, if that makes sense? So being told "Look at this" can be problematic, and being able to do your own exploration --- like I said before, like your enqueueing --- would be helpful. But knowing that there is a structure there I think would help at the start. When I started I just got given this [Points to resource directory], and I had to make that first queue and prioritise based purely on the names. But if someone had a guided tour, then I already have a queue. And yes I can change it, as long as I can in the middle go "Actually...". So for example, when it came to endianness, maybe that's part of the guided tour and I can go "I know what endianness is, I want to skip this". Or for example I can say "Well, I think I actually want to read this other thing first". I think it really helps at the start, but you have to be able to take control, is what I'm going for.

Interviewer: Okay, that makes sense. And you also seemed to emphasise there that there will be different view points on different resources, so it makes sense to have multiple relational structures --- including perhaps your own and other people's...

Participant: Yeah. It's really good for example when you're learning at school, there is a structure. They've objectively said that "To understand this you need to understand this", but some people want to understand that before they understand this, and not everyone mentally follows the same structure, that's just life. So having the basis is really good, because you don't have to start from nothing. You can start from what they've given you, and you'll probably read the first few things to introduce yourself to the problem space, but from there if you can maintain that list but add to it, take away, and move about, then I think that's really useful.

Interviewer: Cool. I think that's all I have, unless there's anything else you want to add?

Participant: I think that's it.

Interviewer: Okay, cool. Well thank you very much for your time.

A.2 Storyboard Interview Transcripts

Participant 1

Link Behaviour

Interviewer: So that's the link behaviour storyboard. First of all, do you understand the ideas there, and does that all make sense?

Participant: I seem to understand it. So, as far as my understanding... the linkbase is separate to the documents, which lets you load in a linkbase of your choosing. Or at least one. It seems to me that there was more than one going on at a time. And people can connect their

own documents for different reasons. Like it sounds like his colleague was highlighting things from a certain topic, whereas he was working on another area.

...

Interviewer: Do you have any general thoughts and feelings about the design of this system?

Participant: I liked it, really. It seemed intuitive. And, like, being able to specialise these linkbases to certain needs... I like the idea of splitting up the content of a page. There's a lot there related to different topics, so being able to essentially put filters on the information to bring out what's important from different points of view --- I quite like that. I like the sort of sharing aspect, like the sort of idea of the 'linkbase of linkbases', so I can publish my own for say colleagues or friends or whatever, you know, who can use it or add to it or however permissions work. I liked that it's sort of straight to the content --- as opposed to, you know, sometimes websites have little things where you can mouseover to get just a little bit more information, but I liked that this was, like, you click it and directly "here is the bit I'm talking about" or "this is the video I'm talking about", potentially with time markers or however you want to make it work. And it seemed intuitive to use, like just a modifier key and just clicking and dragging. A lot of researchers aren't very... they're smart, but they're not very computer competent, if that makes sense? So having something that's very simple to use in this area gives you all the benefits without many of the downsides of them not being able to understand what the hell is going on.

...

Interviewer: And what are your thoughts on the process for selecting a linkbase, did that seem intuitive to you?

Participant: Yeah. ... And being able to select whichever linkbase with whatever information from the list. I'd hope from there you'd be able to maybe expand and get more information. Descriptions or whatever. But I don't know how else you'd make it more intuitive than that.

...

Interviewer: What about the way links are highlighted and this idea of the marker bar? What are your thoughts on those things?

Participant: I like the marker bar.

...

Participant: It's like, you kind of want links to be subtle and non-subtle at the same time. If you're reading a passage you don't

want to be distracted by "Ooh, I could click that, and I could click that", but at the same time... sometimes it's nice just knowing that there's a link. Knowing that I could get more information if I did want to. And being able to come back to it later --- so, go somewhere else on the same page and then scroll back up and you can instantly see what's here. So you have the backup of the marker bar. So I thought that was a really good addition to having a link somewhere in a paragraph.

Interviewer: Something that's also illustrated here quite subtly, is that the marker bar can also show you linked images, and other things that maybe aren't totally clear.

Participant: Okay, yeah. Yeah! That's even better! And also it looks like it's something you could integrate into a sort of search system if you used things like... I think most browsers now do that when you Ctrl+F to find stuff: on the scroll they'll highlight things. I'm used to things like that, so it seems like something that's just an expansion of that sort of idea.

...

Interviewer: Is there anything you wish was different, or that you thought was maybe a bit weird or unintuitive?

Participant: I'll be honest it all just seemed simple enough. ... It seemed pretty intuitive.

Spaces

Interviewer: Do you understand what this is proposing, and does it make sense?

Participant: It makes sense to me. Other than... so would moving a window into a new space count as part of the history, or would it bring along its history? How does history interact with moving?

Interviewer: So, that's a good question. It's something I haven't fully defined. I think I was imagining that when you bring it over, it sort of adds a single new entry into the history of the space you're moving it to. So if you go back, it just goes back through what the history chain of that space was before, rather than bringing a chain with it.

Participant: Yeah that's what I was thinking. So just moving it along counts as an instance of history. Yeah, okay. Other than that, I seem to understand it... I hope.

...

Participant: The idea of multiple workspaces is a concept that's kind of well established, so that's nothing new. But I like this idea of history through workspaces. I haven't seen that before. But I don't really use these multiple workspaces. I think, it sounds like it's a

natural progression. It doesn't seem like something that's been forced on top of workspaces. And it seems like something that has real uses. In terms of, like, we've all been there where we're like "Where was I? What was I doing?", so I like this concept of being able to sort of traverse backwards to "What's the first non-stupid window I had open?". Yeah, it just seems like a natural progression to what's already well established and well used.

...

Interviewer: Okay, and do you think the persistence would be useful to you?

Participant: I think the persistence is... I guess it's one of those things where, having the option there... I mean, maybe not force it on people. Some people like to boot up to a clean "This is my new work day". And some people --- the one-hundred tabbers --- like to have things there all the time. I think having the option there doesn't hurt, really, and I think that it will help more than it hurts.

...

Interviewer: Tabs are a very lightweight mechanism --- you don't have to change the whole view of everything just to change your history chain. Does it worry you that maybe this is potentially a bit more heavyweight, even though I'm saying that it's relatively efficient to switch between these spaces? If you do want to switch between history chains, you do have to sort of move into a totally separate space --- do you think that would be a problem at all?

Participant: Seeing as you're sort of defining moving and changing as an efficient process, I don't think so. You sort of have this concept of, on a browser... you still have to sort of manage history on different tabs and different windows. So I think people are used to having to have different concepts of history in different areas. And it sounds like getting from one to the other isn't... just like you click another tab to go there, you can do some quick swipe or key combination to go to the other workspace. I don't think it's a hard thing, I think people would just need to get used to bringing it up a level. They're very used to history being a very one-window/application concept, whereas this is sort of bringing it to sort of the OS level. So it's just, you know, I think it could confuse people at first but it's one of those things where you just have to bring your understanding up a level, if that makes sense, and apply the same concepts.

...

Interviewer: Do you think this is a good way of solving the problem, as opposed to --- for instance --- being able to differentiate between

groups of windows in a single space?

Participant: I see what you mean... I think personally --- idealistically at least --- I prefer the separation. Realistically I'm not sure, it depends again on how efficient managing these workspaces is. I think if you're trying to get this kind of paradigm of workspaces through, there are people who like their 'work' and 'play' computing time to be separate, and I think this sort of helps. And in a work environment, I've been there doing one thing and trying to manage another, and having everything in one space sort of makes it a lot more difficult. And this concept of her checking on the compilation, sometimes it's difficult to try to find that, whereas if it's in a workspace that is the main focus. I think the concept, if you mean like "in this workspace these three windows are for this so bring them into focus", I feel like that sort of sits in-between. Personally I'd lean more towards the discrete spaces. But I could imagine people wanting the middle one. This is one of those things where people have different sort of opinions. I'd push more towards keeping things separate.

...

Interviewer: And you said you don't really use multiple desktops yourself, right? So do you think this would change your usage, or do you think you just wouldn't use this maybe?

Participant: I'd want to try this. ... I think what adds to this is the history feature. That makes it more useful. I've used multiple desktops occasionally, trying to force myself to keep things separate. So umm... I think that what this would contribute is not so much a concept in my mind of keeping work and leisure separate... now I have a reason to keep work separate, because now there's a reason to want to split them up. So I think it gives more of a reason to start breaking things down into... more... compact... things.

...

Interviewer: Is there anything you maybe wish was different?

Participant: [Pause] No, I can't think of any. Again, just letting the user choose what they want. Be configurable. But it seems like it's own sort of contribution. I can't think of anything different.

Window Search

Participant: So this isn't a search for window names, this is a search for content within the windows themselves?

Interviewer: Yeah, titles and contents.

Participant: Okay, so it's even more powerful.

...

Participant: I think, like the other two really, that this just sounds like taking something that exists and powering it up with more usability and more functionality. I mean things like being able to see all your windows already exists, but now being able to search all of them... I'm imagining that this is a reasonably potent computer, so it isn't going to take two years to search all these windows or whatever, this is something that's quite responsive. It just seems like one of those things that's just another tool in the toolbox in terms of managing your workflow.

...

Interviewer: Do you think this solves a problem that actually exists? Like, do you think this would actually be useful, or is it just a cool idea in theory?

Participant: Hmm... [Pause] In isolation --- just this ignoring other things we've looked at --- I think it's another way of trying to solve the problem of trying to find where I was, or trying to find something specific. I think when you start to combine all the elements that we mentioned before, I think it's usability is overshadowed a little, because maybe through splitting things into workspaces, the need to search may be reduced because you have less to look at. So it seems like this is sort of a different approach to that problem. So I'm not sure how much redundancy there is there, but again it's one of those sort of things where it doesn't hurt. And if you've got three giant Wikipedia articles on something, this is sort of a more powerful method of search than just searching each individual page for some word.

[Interviewer highlights the inspiration for this design idea in the exploratory study]

Participant: Ah, okay, yeah! I'm pretty sure I do remember, in that study, searching through pages to find something that had been mentioned in another page. So this helps. And that's also something where I'd also see the linkbases helping. If I'm Ctrl+F'ing a word, ideally the connection is already made by some linkbase. But maybe that hasn't been set up yet, and I'm wanting to set it up, so I would use my window search.

...

Participant: Just like the other two, I like that it's something that exists in some form today, but this improves its functionality and adds to it. So it's not something that's completely alien to people. It's something that with an hour of use or less, it would just sort of become natural to people's workflow. I quite like that it's not an enormous leap for people to move to this style of workflow.

...

Interviewer: Any questions or 'what if's? Maybe this could be more powerful in some way you're imagining?

Participant: Integration into web searches would be interesting. But where does it stop? How much time does the user have to wait for it to scan everything? Also, how would one navigate their searches? If you search in a web page, a lot of modern browsers highlight on their scrollbar. For example if you have a lot of windows open, you now have lots of windows open with pieces of text, so how do you navigate between them? That would be one question to answer, I guess. And maybe refining searches, so how powerful is the search functionality? Is it just words, or can I use regular expressions? You know, things that the average user might not want to use, but things that would really help power users to get even more out of the system.

Exploration Mode

Interviewer: Do you understand this storyboard? Does that make sense?

Participant: Yeah. So, the content that it pulls up... is this following links and requesting new content, or is this only for things she's already opened?

Interviewer: Umm, so, it's not entirely clear within the storyboard, but I think I was imagining it being perhaps some mix of those things. Maybe if you have something open that's very related, that's displayed, but also you could have stuff that's not open at the moment that's just related that you might want to visit.

Participant: And is this sort of separate from your concept of a linkbase, or is this sort of plugging into it? So... what I'm understanding is that this is like the linkbase, but that rather than doing it yourself, in this mode it's going to give you a little preview of everything that's being directed away from that page?

Interviewer: Yeah, so I guess there are different ways you could do this. It could be that it draws on the linkbases you currently have enabled, or it could be that it does some kind of global search --- so maybe it's also displaying stuff from other linkbases that you don't have available but seem relevant --- and this kind of thing.

Participant: Yeah, okay, so there's some form of intelligence of it scanning the content that you're viewing now, and through that it's pulling up other information?

Interviewer: Yeah.

Participant: Okay. I guess my only worry would be, how much is being shown and how much visual noise there is. I think, in a concept like this, it looks quite nice. In a modern document or webpage or something, realistically you could link from that to a lot of places. So, where does it choose to, like... it seems these rectangles at the bottom seems to be like overflow, or maybe this could be other spaces that are linked to, or... Yeah, I guess you just have to be mindful of the way it chooses to display it. Maybe it would have some form of

priority, so if something strongly links it could be shown right there, but if it's very minor maybe there's an 'Other' section or something, to make sure... I think it might be very discouraging or scary to press this button and then get this giant graph thing, and you're thinking "Where do I start?". But I think if you can manage that, I think that it would be a good thing. Because we've all been on search tangents before where we don't want this to be the prime focus, we just want to top up our understanding or come across a term and say 'What does that mean?'. So being able to go through that, feel comfortable understanding that, then pressing a button and immediately returning to where you were before, rather than having to navigate all the way back. It seems like just a natural contribution to people's workflows.

...

Interviewer: Do you think that the fact that it's kind of a separate mode that you need to switch into and out of is a problem? Maybe if you wanted to quickly see a connection to something, it's not the ideal solution because you have to switch in to the mode, check it, and then switch out of the mode.

Participant: I mean in my head I'm just imagining that you sort of just press a button and it transitions into this. I don't think that's such a horrible thing to have to do. For example if you want to search in Windows, you have to press a button and something comes up to search it. I think people are used to having to enter a new state to do something. And it's not like it's a completely distant state, like if something completely new flashes up. Your mind can sort of track what's going on. Another thing to sort of just be aware of I guess is to make sure there's a connection between the two modes, so that the user --- and like, non-power-users especially... I think there's an understanding that it feels like the same mode, just more powerful, like I've entered this super-state. And then when I'm done, return back to where I was. So it doesn't seem like I've lost track of what was going on before, or that I've had to like really put in any effort to enter something new. It's just like "Give me more information on what's there".

...

Interviewer: Also, do you think this solves a real problem, or is it just a neat piece of functionality?

Participant: I think it has the potential to. Again, I'm mindful of being bombarded with information that I don't care about if I enter this mode. So as you said before, I think you said it's highlighting lots of things but different things could be highlighted different ways. So it might be nice to enter this mode, but for a specific reason. So "I want to go into Exploration Mode, but for this type of term". And then everywhere it's mentioned on the page --- that term --- you can

- maybe bring up things in that context. So if you were able to refine it, I think that would be really powerful, and it would really help. So maybe for maths or something, as a really simple example, you could have a concept, and then all of its applications... if I could press this button, 'Exploration Mode', and then see those examples fully developed in a sort of preview, where I might not even have to click into the preview because it's just there, I think that sort of thing is really powerful. But if I'm pressing this button and sort of getting just random information, because it's decided that this is relevant to me, I think that's where it starts to lose its value.
- Interviewer: Okay, that's interesting. So I guess this kind of relates to what you were saying before about there potentially being too many connections, and there being a potential issue with visual noise, or recommendations that you're just not really interested in. So if you had some way to direct what connections you want it to show you, you're saying that might be really useful?
- Participant: Even if it's just showing you that, or maybe like blurring out or fading away, so that you can see other things are still there, but currently we're focusing attention on the things you wanted. Or maybe you bring in other stuff like "Now you can search these previews", or... yeah, letting the user direct--- because it's Exploration Mode, but I think you want to be able to direct your exploration, and more than just clicking would be useful.
- Interviewer: Okay. What exactly do you mean by "more than just clicking would be useful"?
- Participant: So, you can enter exploration and hover and click around, but more than just that. So some form of deciding what kind of thing you want to explore. So maybe if it was before you enter Exploration Mode, you highlight a section of text, and say "Explore this!". Or maybe you do it after the fact, so you press Exploration Mode and then within the mode you can say "Remove all the ones that don't talk about this". Or maybe if there are specific instances where the algorithm seems a bit off, you can say "not interested" or whatever and it would go and replace it with something else. And maybe I could tune it with things I'm more interested in. Kind of just customising it and refining it, rather than just an explosion of information which sounds interesting, but in reality I feel like my eyes would just glaze over.

Whole System

- Participant: Individually they're all quite nice concepts, but it gains its power when you really start mixing them together. The link and the linkbase and things, they have the potential to be the backbone to a lot of these other ideas. In terms of you could have a linkbase which is prepared for a space, or a linkbase which is prepared for an exploration. So you could tie the two together so that the exploration is of the links in a linkbase. And I think the history

one really helps if you're trying to move to this 'bespoke window' paradigm, where content is in its own sort of view. Being able to manage those becomes a lot more difficult without these extra sort of features. So having the history makes this paradigm underlying all the other storyboards manageable really.

...

Participant: To me, this seems like a system prime for having different levels of publishing, if that makes sense? Like, departments, businesses, the wider world. It seems like something where linkbases are your contribution.

...

Participant: In some ways, this almost seems like a take on Chrome OS, if that makes sense? Chrome OS tries to make the browser the OS. So that's at the top level, and those concepts are at the top level. So this seems like you're trying to take these browsing concepts, and move them up a level. So the way that Chrome OS did it is by making the browser the OS, whereas you're trying to pull what the browser does into the OS almost, which is a different take on it I think. That's I think the closest thing I can come up to what you're sort of trying to do here. It just seems like a different way of bringing browsing habits and concepts to a higher level whilst also making that manageable and adding functionality to it.

...

Participant: It's definitely something that I would try. ... Like actually, honestly, those are some pretty cool concepts that you had there.

Participant 2

Link Behaviour

Participant: I think it's really cool. I like the idea of just being able to research a subject, and having all the links relevant to what someone thinks. And I like the way that anyone could make a link database. You can also see what your friends thought as well. Or researchers who know more about the field than you --- you can follow their links to see what they think is relevant to the subject that you're researching.

...

Participant: Would this solely be for a desktop? Because I don't think it would work well on a mobile device like a phone or whatever.

...

Participant: I like how you create links as well. And especially the marker bar --- I think that's just really convenient for noticing where different links are. I think it's quite important to get the user's attention.

...

Participant: If you're going to have different kinds of links, like maybe picture links or video links, maybe you could show that on the marker bar. I think as long as it's consistent, then that would be good.

...

Participant: I like the way that links don't detract your attention too much from the entire article, so you can continue reading if you want. It doesn't say "click me now!" or anything, it's just quite subtle which is good. I like the way it opens a separate window too, so you can still go back to the previous one and there's kind of like a hierarchy of what you've opened. So you can see how far along you are --- deep into whatever you're researching. I think it's quite a cool concept.

...

Interviewer: What do you think about the way linkbase selection works, does that seem intuitive to use?

Participant: Yeah --- I think as long as you're familiar with how linkbases work, then I think it would be a fine system to use.

Spaces

Participant: I like the concept again. It's straightforward and intuitive, because you know that you're going to have different projects so you can put the relevant stuff in whatever space. And I like the idea of the back and forwards navigation --- that's something which I think is quite important to have in a system like that. And yeah, if it was implemented like you've shown it then it would work well, as long as the user's aware of the key combinations and whatever. I like the way your can do the bird's-eye view as well to see what's in your different spaces, so you can put relevant stuff in whatever if you need to manually. Yeah, it's quite cool.

...

Interviewer: Can I ask, do you use multiple desktops much today?

Participant: I did when I was working last year. Not so much now that I'm back at my Windows machine.

Interviewer: So even though you can see how this approach might be valuable, do you think you would actually practically use it? Is the pull of having the separate histories enough value to make you use this, or do you think you would still maybe not use it because multiple desktops is just not something you really use?

Participant: I could see myself using it. There's a lot of value in having different spaces to work on different things. Keep it all relevant to whatever you're doing. I'd definitely use it.

...

Interviewer: Also, do you think this is a good way of solving this problem? So, like, I suppose a different solution could be that instead of having different spaces, you could differentiate between different groups of related windows in a single space. Do you think using separate spaces is a good approach in general?

Participant: I think this approach is quite suitable for what people want to do. Straightforward, intuitive. I think it would work well.

...

Interviewer: Is there anything you wish maybe was done differently, or think might be painful?

Participant: Not really. Once again, it's just user experience --- so as long as people are aware that they can do all these things then it's fine.

Window Search

Participant: Interesting. I feel like this is kind of similar to what Google Chrome does, although that only searches for the name of your tab --- it doesn't look in the content. I think searching through text in individual documents or windows to find snippets is a really nice feature to have.

...

Interviewer: And do you think this addresses an actual problem? Like, I guess going from your experiences in the exploratory study, do you think that something like this would have been useful to have?

Participant: Oh yeah, definitely. There are a lot of times where you read something, but shortly after you forget where you read it and how you found it. So if you've got everything open, you can search and say "Oh yeah, it was in this thing. But it's also in this thing I read as well!". So it's just a good way of keeping tabs on the information that you've seen.

...

Interviewer: It's easy to imagine a number of ways of extending how this kind of functionality works. Do you have any thoughts or questions or ideas about what you'd maybe look for in a more advanced version of this sort of idea?

Participant: So, maybe just some basic filters. Maybe you only want to see pictures or videos --- filtering by file type. And maybe how recent they were looked at too. You might remember that you've seen something in the last five minutes, but you might not remember what it was. So time filters might be good too. Aside from that, I can't really think of much.

Exploration Mode

Participant: In general, Exploration Mode seems quite straightforward. I really like the idea of previewing a link. Hovering over something and maybe seeing the context in which it's linked to whatever you're reading --- I think that's really useful. There are a lot of times you might waste time by going onto another resource that you think is relevant, but it's not. So it's good to see kind of a quick summary of whatever it is. And then you can decide if "Oh I want to spend my time reading this or looking at this" or whatever. Yeah, I think it's a pretty cool idea. I like it.

...

Interviewer: Do you think this solves an actual problem?

Participant: Yeah. It's nice how in Exploration Mode you can see how all these different links are all attached to what you're looking at --- you can see the relevance in whatever you want to do. And maybe you can say "Oh I see that link over here, but maybe I might read that later because it looks like I need to read this resource over here first", and you can preview and figure that out. So it's kind of an efficient way to spend your time looking for stuff.

...

Interviewer: I guess there's also potentially an issue here that if you have lots of related information, how do you decide which ones we should actually show? What do you think about that?

Participant: So, I did think of something you could do to really improve it. Using multiple link databases, you could have, like, the most popular links from a selection of link databases and choose from that. Or you could do user created content. So Amazon, like, does the whispers thing --- whenever you highlight, people can see that you've highlighted. So it might be good for users to see who else has put

emphasis on a particular link to some resource.

...

Participant: It seems quite efficient for learning, which is obviously great. You don't want to be wasting your time. I think it's good.

...

Interviewer: Anything that you would maybe like to be different, or potential problems that we haven't talked about?

Participant: Maybe the links... in the context of a desktop view, everything might appear too small. So, like, taking a look at some of the stuff might be fine, but for like reading I'm not sure how feasible that would be. Maybe if it was, like, I don't know... a different way of viewing it, making it slightly larger or something, I don't know.

Interviewer: Mmm. So which bit in particular do you think might have readability problems here?

Participant: So, the resources definitely. If there's a number of them, you're going to be kind of squinting maybe.

Interviewer: Do you think maybe if these were just visual previews, but then you combine that with this link preview functionality, do you think maybe that could help alleviate some of these issues?

Participant: I think so. Because at the end you don't want your vision to be too impaired with just, like, a lot of information at once. You just want to kind of look and see what you might need or want.

...

Participant: I assume you could do each of these things in separate workspaces if those were a thing. Is that, like, the kind of idea? So for whatever you want to do, maybe you want to have a separate workspace just for exploring stuff on a topic, and then maybe you have a separate workspace for just kind of looking at an article or whatever. That kind of thing.

Interviewer: Interesting, yeah maybe. So, I guess, is what you're saying there if you have different spaces and you go into Exploration Mode into one of them, you would want it so that when you switch back you're still in Exploration Mode and can continue to explore?

Participant: Yeah, exactly, that's what I'm thinking.

Interviewer: Yeah, that's a good idea.

...

Interviewer: Although I've proposed it for the use case of digging in and doing deep exploration, do you think maybe if you just wanted a quick look at the connections between something, do you think that you having to enter the mode, do it, and then exit the mode is a bit of a

barrier? So you can't just go straight to it, you kind of have to go into the thing and then do it and then come out. Or do you not really see that as a problem?

Participant: Personally, I don't see it as too much of a problem, but I can see a lot of people seeing that it is.

Whole System

Interviewer: What do you think about the vision as a whole, if we were to combine all these things? What are your thoughts on that?

Participant: I think it would be a really interesting system. I can see it being used a lot for kind of academia, learning, that kind of thing. Getting interested in the stuff that you want to look at --- I think that would be really useful. Aside from that, I don't know how much I would really use it. I know whenever I go on the Internet that I only look at specific things, I don't generally go into too much detail unless I have to. So if it is literally for uni work then I'll go checking all these different links and everything, but day-to-day I just kind of... yeah. Do you know what I mean? I don't think that's a bad thing, though, I think this system has a really good use.

Interviewer: Okay. So what's an example of somewhere you might go in a day-to-day sort of situation?

Participant: I'll probably open some specific programs, Eclipse or whatever, and a Web browser with just a few tabs. Facebook, YouTube, whatever. Having Spotify open, that kind of thing. So maybe that tells you about how I use my laptop or whatever, other people might have different workflows.

Interviewer: So, do you think this could work as a potential replacement in terms of... so rather than being a set of connected pages, if you think of Facebook maybe as a program almost. So you have Facebook as a program, perhaps, and YouTube as sort of another 'service', which is just kind of one window rather than having this kind of multi-window navigation like we're talking about here. And then maybe for your documentation for your code or whatever you can have the kind of interrelated windows kind of thing that we've been talking about. Do you think that could work as a solution for your kind of day-to-day case, or...?

Participant: Yeah. [Pause] Yeah, I think so.

Participant 3

Link Behaviour

Participant: I like the idea of the linkbase, that's quite cool. The highlighting... is there any point it being different to what it is already? I don't see much point in it being different, especially because the blue underline is like convention now.

Interviewer: I think part of the reasoning was that if you enable lots of linkbases, you potentially have many more links on the page. So if links are too harsh visually, it could be too distracting.

Participant: Hmm, okay, that makes sense. And would it be possible to tell if there are two different links? Say if they were right next to each other or there was just one word.

...

Interviewer: What do you think about the method of link creation?

Participant: Yeah, that's pretty cool. I like that. It just looks really easy to do --- I could see myself using that.

...

Participant: I like the marker bar idea. Is there any sort of information if you hover over the marker bar? Because that could be useful, sort of like an overview of what it is.

...

Interviewer: What do you think about the process of linkbase selection?

Participant: I'd say it's relatively easy, it makes sense.

...

Participant: I guess if there are loads of links, that might get really annoying --- literally like the marker bar is just full. But if a page doesn't have too many links, I think that'll be fine. For modest amounts of information, yeah, I reckon it's good. I like it.

Spaces

Participant: So the 'Back' button is like... instead of just being on one program is for the entire space?

Interviewer: Yeah, exactly.

...

Interviewer: How do you think this compares with something like tabs as a solution to managing chains of history?

Participant: It sounds kind of similar, fundamentally. I feel like you need to be a certain type of worker to be able to utilise something like this though. Because I'm so bad at keeping track of different desktops, and I just have everything everywhere.

Interviewer: Okay. Is it really that different to how you would manage tabs today, though?

Participant: Hmm... no, I guess not. [Pause] No, I guess it isn't.

Interviewer: Okay, so it sounds like it's something that maybe you would have to get used to and see how it felt?

Participant: Yeah. I mean... I'm not very good at like separating out things... having desktops for specific things, like having a work desktop and a game desktop and whatever --- I'm not very good at that.

Interviewer: I see, so you don't use multiple desktops that much today?

Participant: No.

Interviewer: Do you think maybe you would use them more if they had this extra value of being required to separate chains of history?

Participant: Yeah, probably. It'd be awesome for like, at work, doing research and all that. That would be super useful, so I'd probably use it then. And it's just because, like, it's quite hard to do multiple desktops on Windows --- you have to press like four buttons or something. But on Mac it was really easy. I use it a lot on Mac.

...

Interviewer: Do you think this is a good way of solving the problem? So you could imagine that maybe another solution would, rather than splitting things into separate spaces to maintain separate history chains, maybe you could allow users to differentiate between groups of windows in a single space, and using that to do history, or something like that.

Participant: So it's kind of like multiple spaces within a space? That would definitely just confuse me. I like having them logically separated. That makes sense to me. And if I hit the 'Back' button, it will go back to where I expect it to go. But with the groups thing I definitely would get confused.

...

Participant: I like the idea. The prototype looks really similar to multiple desktops on the Mac, I guess it is inspired by the Mac design. And I liked it on Mac, it was easy.

Window Search

Participant: I like that. I would use that a lot. Especially with work --- I always like remember a certain phrase when I'm working and I'm like "Where did that come from?", so this would help.

...

Participant: Search is always nice, just having any level of search anywhere. I just--- I like the look of this. I imagine this would speed up searching so much rather than just Alt+Tabbing through everything. It would really help.

...

Participant: So if we have spaces, is that just searching one space, or is that searching all the spaces?

Interviewer: So, yeah, I haven't really defined here whether that's the case. I think in my head it's to search one space. Because I feel like that's what you want most of the time.

Participant: Okay, yeah, that makes sense. Yeah, because everything's separated logically. Okay.

Exploration Mode

Participant: I personally probably wouldn't use this, but I can see how it's useful.

Interviewer: Why do you think that is? Why do you think you wouldn't use it?

Participant: Umm... it just looks like there's a bit too much going on, you know? Especially if there are like six windows going around. Would you be able to like, change focus onto one of those? Like, change the middle window and then there's more links?

Interviewer: Very possibly, yeah, that's an interesting idea. It's not very well defined at the moment.

Participant: I guess it's... I can see how it's useful if you want to learn about something, but I don't want to learn about much very often.

Interviewer: Right, okay. I guess, if you think back to your experience with the exploratory study, do you think it would be useful for that kind of work?

Participant: Umm... actually, yeah, possibly. Yeah I guess it would, thinking back. It would make a lot of sense actually. Okay, I might use it in situations like that, yeah.

[Laughter]

Interviewer: I see what you were saying, though. You don't see yourself using this day-to-day very much, this is kind of quite a specialised piece of functionality. Right?

Participant: Yeah, yeah. This is like for a very specific purpose. It's not very often that I'm like trawling through loads of information trying to figure something out. But some people might.

...

Participant: Also, it seems to me like there are a lot of different things highlighting links. There are like three different things highlighting links, which might be a bit overboard.

Interviewer: How do you mean exactly?

Participant: Well you've got like the link being highlighted, then you've got the marker bar showing you where that link is, then you've got a big line coming off of the link showing you the actual resource.

There's loads of stuff just about one link. I guess it's not really a problem, but, yeah...

Interviewer: Okay, I see what you mean. I think also, part of the idea here was that in addition to just the links that you'd seen before, you might also see other stuff. So it could almost do like a search, right, for what popular, interesting, related stuff there is. So there might be some degree of intelligence for displaying new, related resources. So you might see different links and different relationships than you would see outside of Exploration Mode. You're right though that we could just display this information using link highlighting and the marker bar. But do you think that generally this intelligence-based approach would be valuable?

Participant: Yeah, that would be quite cool for learning about a specific thing. So it's not like linkbases anymore, it's like an AI guessing what you want to know about?

Interviewer: It might interact to a degree with the linkbases you have enabled, but yeah, it could also maybe pull on wider stuff.

Participant: Yeah, that makes sense. That would be cool. So you're not confined to stuff that you've sort of already seen before.

...

Participant: I like that it's just like a mode. You can just turn it on and start exploring. That's quite cool.

Interviewer: So what do you think you like about that? Just that you can always turn it on from anywhere?

Participant: Yeah, just that it's so easy to jump in and jump out. Things like that are just nice. You can just hit a key combination and bam --- you're in this new mode.

...

Interviewer: I guess, there's also a problem here with what resources we choose to display. We only get to select a small number to display. Do you have any thoughts on that side of it?

Participant: Maybe it would be helpful to be able to like, remove them? So like if one of the resources isn't useful, you can just throw it away.

Whole System

Interviewer: Finally, if we were to combine these things into some sort of single system, what are your thoughts on the whole that consists of these parts?

Participant: It just seems like it would be a much easier way to do research. For like general browsing and stuff I probably wouldn't use a lot of features, but definitely for researchers and stuff it would be a really useful thing to have I reckon. For just browsing YouTube

and Facebook and stuff, you don't need all of this stuff, do you? But I like the idea. I think I would use it a lot if it was a thing when doing like research work.

Participant 4

Link Behaviour

Participant: You said about highlighting links subtly... would it deal with, like, different monitors, different backlight colourings, and different things like that without becoming an issue? Because I can imagine very easily, if you have just say a very subtle grey and let's say someone hasn't got their monitor configured properly, they're not going to notice that issue.

Interviewer: Yeah, I think that's a good point. You don't have to highlight in grey, so for instance the text selection colour in some browsers today is a fairly inoffensive blue, right? You could use something like that potentially.

Participant: Because one of the things with hyperlinks is obviously that they're underlined as well, so you can see them visibly even if you have your monitor in black and white mode for example.

Interviewer: Yeah, for sure. I think part of the justification as well here is that because we have linkbases, you maybe have many more links on a page. And if you have lots of bright blue underlined things, I think that's potentially quite noisy.

Participant: Yeah, okay, that's a fair point.

...

Participant: Let's say I've just started off with this. I've just come to the system for the first time, and I haven't got friends who have used this. Would there be no links at all, or would there be links given by some sort of base...?

Interviewer: So, documents still come with their own links, like a default linkbase for the document. But then, apart from that, yeah --- the discoverability mechanism hasn't really been fleshed out yet.

Participant: Yeah I was about to say, discoverability is the big thing I was going to question. Because one of the issues that I know even search engines have is that you only find things that you're looking for. Which sounds a bit backwards in a sense, because obviously you're searching for something so you want to know about it. But then you're not going to expose yourself to something that you maybe didn't know existed, to look for that sort of thing.

Interviewer: Yeah, I think there are certainly lots of questions around discoverability. And also, if you're made something, how you share that with the world as well.

...

Participant: So you said, like, someone creates these linkbases. Two questions. One, what stops people just making big troll links and destroying what this is. So people can be like "Oh, we're going to link this to...", I don't know, some bad propaganda or something, and make people never trust any of the link sources. Or they could link it to a site that's a virus site.

Interviewer: Yeah, people could definitely do that. So I think this is another issue that sort of relates to this discoverability stuff, right? So, in this case, Toby trusts his colleague and so this isn't a problem, right?

Participant: Okay, yeah, so it involves a bit of trust?

Interviewer: Yeah. Maybe it's the equivalent of like following someone on Twitter, right? "Oh, I follow this guys linkbases because I think he knows his stuff."

Participant: But the difference is that on Twitter, you have trust in Twitter that Twitter isn't going to send me a virus by me looking at... so if I follow you on Twitter you can post whatever you want, but if I suddenly---

Interviewer: Well, people might post some offensive images or something like that. You have some trust in the people you follow.

Participant: There is, yeah. But there's still an amount of... so okay, fine, I might see an offensive image, but it won't affect my computer for example. You can't put a virus on Twitter, whereas you could link straight to a virus, or to something that just grab your IP and DDoS you or something, I don't know. Something evil. There are malicious people on the Internet.

Interviewer: Yeah, I agree there are definitely some security concerns here.

...

Participant: How do you share your linkbases with people?

Interviewer: That's also not really defined here. Toby has created his own links here, and if he wants to share those with someone it's not totally clear how that would happen. Maybe he has to share some specific file that is exported from the system, or maybe it happens more sort of automatically through some kind of infrastructure.

Participant: Yeah, I was going to say, how are linkbases maintained? Are they digitally maintained on a server somewhere, or are they locally maintained?

Interviewer: So, when Toby is creating this linkbase I think at the moment it's probably local-only. But then I think if he was to share that, then yeah there are questions. Maybe you push it to some server or maybe it's some peer-to-peer thing, there are lots of different ways that could work.

Participant: Yeah, that makes sense.

...

Participant: The marker bar... does it do anything other than just being something to see? Like, does it expand out or pop out or anything like that, or is it just...?

Interviewer: Yeah, potentially. Again, this is something that's not really defined here, but interacting with the marks in the bar is definitely something I've thought about.

Participant: Then you also have to deal with, what if you have multiple links in one line? Would it highlight all of them with one mark, or would it become a stacked thing?

Interviewer: Yeah, that's another good question.

...

Participant: Let's say I'm doing my own research on something. Right now, for example, I'd go to a computer and load up my own things, and I'd form mental links between them. In some sense it's useful to then also have physical ways to go between them. But I think the more useful aspect would be when you have someone else's linkbase, to go between things in a more logical way. I think that's probably the bigger help than making my own links, although I can definitely see the benefit in some cases for it.

...

Interviewer: Do the processes of displaying and enabling links and working with linkbases seem intuitive to you?

Participant: It's intuitive, but it would definitely be something you'd need to get used to. Because it's just not standard right now in any way, but I think it's something that would be pretty easy for people to understand. Although I would question the naming of some of the things and simplify it down a bit, just because people might get scared of... what did you call it before? The 'linkbase menu', people might get a bit scared of certain things like that. But these are very minor, trivial things.

...

Interviewer: Are there any things that you maybe wish were different or think could be painful?

Participant: [Points to the list of linkbases in the menu bar menu.] I'm a bit worried about this, in terms of that list becoming a bit gnarly and awkward to run through.

Interviewer: Yeah. So in my head this shows you some that it thinks might be useful, but then you can also open up a more complex dialog to explore and manipulate linkbases.

Spaces

Participant: I like the history aspect. I'm not sure how that works when you, say, drag a group of programs over to one space and then bring over another one. I don't know how that keeps the history --- would that then become a new item, or would you try and put it back in its position? And what if the person wants to view global history, all of it. I know you have the different workspaces, but is there a way that I can say "Oh, I actually want to go back through everything"?

Interviewer: To that last question, I think my answer at the moment is probably not. Spaces are kind of supposed to be kind of link a replacement for tabs in that you just have completely separate history chains in each space.

Participant: Okay. Yeah, that makes sense.

...

Participant: I'm kind of one of the people who also hate having to remember key combinations. ... For me, learning new key combinations and new systems is always a thing I'm going to be anxious about.

...

Interviewer: Do you think this is a good way of just solving this problem of "how do we deal with history when resources are linked in windows"? For instance, you might imagine a different solution in which you provide some way for the user to mark different groups of windows as having separate history chains, even within a single space.

Participant: I think that could also be useful to have. I think of it very much like a tree system. And here you're breaking out the branches to separate spaces. I think that works, and it makes more sense for a lot of environments. For some people there's this downside of learning all these new things, but definitely being able to go back is useful.

...

Participant: If I have like five windows open, and they each went from one to the other --- so I went from document one to document two, document two to document three, and so on --- if I switch to document three and then go back, do I go to document two?

Interviewer: So you mean if you Alt+Tab to document three, for instance?

Participant: Yeah.

Interviewer: So, that's actually a very good question. I think the way the system works at the moment in my head is to just do the very simple thing of, if you clicked to go to document three, the click action itself isn't a new item in the history chain, but 'Back' will take you back to document five. So it stays chronologically based, rather than skipping around depending on your focus.

Participant: Okay. I can see both pros and cons of that. Let's say I want to go back to document two or something, I might not want to have to

go "back, back, back, back, back, back, back --- oh, forward!", but I can sort of see that making sense. ... I think personally right now sitting here, I would have said it makes more sense for three to go back to two, and going back forward it would go to four. But it's kind of hard to know, maybe if I tried that it would be completely unintuitive.

...

Participant: Another question, let's say I have those five documents again. I take item five and put it into a new space, then hit 'Back'. What happens?

Interviewer: Again, this is kind of an open question. But I think the way it works in my head at the moment is that if you've taken it into a space that already has some history, then this resource just becomes one new entry in the history that already exists there, and so you'd go back through the history that was already there before you.

Participant: And what if you move it into a new space that has nothing else in it, an empty space?

Interviewer: Then I guess back probably does nothing.

Participant: Okay, so you wouldn't do anything. You wouldn't close the space and go back to the other thing, or something like that?

Interviewer: Possibly, yeah, maybe that makes sense.

...

Participant: Would there be any GUI-based approach for going back? I know you said there was a key combination, but could you have some sort of icon or something that shows you "oh if you click that you'll end up back on this video"?

Interviewer: So, I think that's a good question. I suppose my counter question is, where do you imagine such a thing might be, or what might it look like? Because the reason I haven't pictured such a thing is that it doesn't really fit within the operating system environment today. I don't know where you'd put a button like that, it's like a cross-window action.

Participant: Hmm, yeah. ... I definitely understand that from a UI point of view, it's very difficult to conceptually build something there.

...

Participant: I also think it would be a good idea to group things. To be able to group applications, and just jump to a group.

Interviewer: So are you talking about something separate from spaces?

Participant: Yeah, separate from spaces. But then it would make sense I guess in some sense to make it a space, but sometimes you might not want hundreds of thousands of spaces for different things --- like, how granular can you make your spaces? So like you have a space for Computer Vision, but then Harris detection: do I make a separate

space just to research Harris, but then it's also a part of Vision?
So you might want another way to sort of sub-group within a space.
Almost like a sub-space within a space, but not.

Window Search

Participant: I could definitely see a use for that. ... It's a very nice thing to have. I slightly feel like at some point, though, if you're having to do that, you've probably got too much open and probably need to organise yourself more. Like, if you had the space system for example, you probably have too much in one space.

...

Participant: Searching window titles is useful, searching contents... I'm a bit dubious of that concept, just mostly for speed reasons.

...

Participant: But yeah, I can see why having search for all your content would be super useful. Especially if it searched across applications... and I guess it would also be very useful for building links between things. So ideally, it would be good if there was a way to actually build links from within the search UI. That might sound a bit weird at first, but if there's a key term I'm looking for and it pops up on this Web page I was looking at, a document I had open at some point, and a coursework spec. I'd like a way to instantly be able to have a link so I can remember that I can go between these things.

Interviewer: I guess this kind of also goes back to variants on how this could work. I was kind of imagining that you enter this search, see the results, click one, and then the results are gone. But you might want it so that even after you've selected to view a result you can say "Right, what's the next result?", and then actually go to a different window with the next result. So maybe that could accommodate this kind of behaviour as well.

Participant: Yeah, that kind of concept, yeah exactly. I definitely think that would be useful. If I search for something and five things come up, I want to be able to jump between them all. I don't want to go back to the search bar again, re-type my search and then go "Okay I think I'm on the second item this time, so it's this one". Definitely having some way to tab between the results would be super super great.

...

Interviewer: Do you think this solves an actual problem, or is it just a neat piece of functionality in theory?

Participant: No, it solves a problem. So, let's go back to coursework as

- an example. I might have the spec open, then I might have a PDF with some article I've read, and then I might also have some other stuff. This gives me a direct way to go between them and search between them for what I wanted.
- Interviewer: So, actually, because of the way that spaces work, you seem to be saying that you shouldn't really have too many windows open such that this is the main way of finding windows in big spaces, but actually that this is much more useful for just being able to do that search between different types of resources?
- Participant: Yeah. It's basically like having links without having someone actually make the links for you. Actually one question I have about this: how do you deal with photos and videos and stuff, that don't have just a text search?
- Interviewer: That's a good question. You can search transcripts if they exist, but yeah it's a little more difficult.
- Participant: I guess you can fall back to searching the title or file name. But if somehow you could also search the content --- hand wavey, let's pretend you could --- that would be great, like "Oh, that guy said that at 3 minutes into his video", being able to jump straight to that point to keep track of what I'm doing.

Exploration Mode

- Participant: So you have all these things showing you what the resource relates to... might get super super cluttered and become overwhelming at some point. If you click on them, does it do anything?
- Interviewer: Yeah, so it's not entirely clear what happens but certainly I think it should do something. So either it exits Exploration Mode and just opens the resource---
- Participant: Yeah, that's what I would expect it to do.
- Interviewer: But also, if you wanted to, you could have it so you stay in Exploration Mode but the selected resource becomes the new main resource for you to explore further. It's an option. But with regards to what you're saying about the number of these resources, I think that's definitely a key question here. You would need to have some kind of intelligence to choose only a select set that you want to display. Do you have any thoughts on that process of how it should decide which ones to show you?
- Participant: That's a really difficult question to answer, because there are many ways that people would approach this. Some people would go down the route of "Oh, let's do the one that most people clicked on". But then there could be situations where, for example in your first storyboard, when someone --- maybe one of my friend researchers who I like and trust --- has given me a linkbase, and I want to view more of their connections. I would imagine it's something that you would have to configure somewhere. So maybe to configure it you just like an ordered list of what would be prioritised above what, like "lecturer friend first, then most popular results next" then further

down I might have other stuff. Then it could show the top things according to what I wanted to see. That's how I'd probably go about solving that.

...

Interviewer: Do you think this actually solves a problem? Maybe reflecting on your experience with the exploratory study.

Participant: Hmm... I personally would not find this screen that useful. [Points at main Exploration Mode screen.] Instead, I think I prefer this kind of approach. [Points at transpointing.] For example I can hold down the Ctrl key and then see a transient window based on what the link would go to. That's more useful to me than having a scary, huge mindmap just explode on my screen.

Interviewer: Okay. Why do you think you wouldn't really make use of that part of it?

Participant: I kind of... I can see what it would do, because it would point out things and where they're going to go to. I don't know how you would recommend the boxes --- would they hold content or just images or icons? I'm not too sure. But umm... if I'm already able to tell that there's links on the screen, I would probably just look at the links to find out what I wanted to go to anyway. Rather than having some thing well me "Oh by the way there are all these possible things that you can go to", when I can already see on the screen the links that might take me somewhere else.

Interviewer: Right. Part of it as well though is that you might see things that you wouldn't see outside of Exploration Mode. So you might see popular, related stuff for example that wasn't visible outside of Exploration Mode.

Participant: Hmm... I guess it depends what your approach to links is --- if your links are strictly done by linkbases, or if your links also come from "a lot of people also look at this" and that kind of thing. You don't really need another system to deal with that, it might just confuse people more. You have to be careful with how many menus you start throwing at people. If you think of a current modern OS, they probably have two main systems --- you have the main screen, and you have a screen that shows all your apps. I think when you start throwing too many UI designs at people, like "Oh well if you actually go to this then you get this UI, but if you go to THIS then you actually end up with this UI", and people are just going to get confused about where they're at and what they're doing. I see why you might have that, don't get me wrong, but I just think that with how links are incorporated into the main system, it might not even be necessary to have this kind of system.

Interviewer: Okay, yeah that's true. Let me just throw another element into the mix of your consideration there. What if this could also display popular resources that link to this, but are not linked from this? So it's kind of awkward to display such things as links on the actual page.

Participant: Hmm... yeah, I could see the benefit there then. Yeah, I could see some sense of that side being more useful then, just showing where links are going to, that could be useful. In a sense it's a bit backwards though, because if I'm already there I probably want to find out more things about what's on there rather than what things can lead to there. As I said, I think I can see some use of that, but I think the negative of having another screen might be more than what that benefit would be. You could just also have a little pop-up box at the side rather than having this big mind map thing. Just be like "Oh, by the way, these articles link to here", if you wanted to do that rather than a big mind map. I like mind maps, don't get me wrong, but I imagine they might get quite scary for people.

...

Participant: But I definitely want that. [Points at transpointing.]

Interviewer: Maybe you could have this without a mode switch, then? Like you could just hold something and it just happens.

Participant: Yeah. That would completely make more intuitive sense to me. Rather than be like "Oh, I've got to be in this mode, and in the other mode I can't do this..." --- that's a bit weird. But if it's just a simple modifier key where you can hover over the link, that would make a lot more sense.

Whole System

Participant: I definitely see it being very useful. I can see the strengths in being able to jump between resources, especially when they're in different applications. Linkbases are super useful, being able to search for things is great. Being able to mentally lay out different topics in different areas is also a good logical idea. Overall I see it being useful, but I do have some concerns about it being such a drastic change from what we already have --- the adoption of it being a key question. People get mad over very minor changes, like moving the start bar, and now you're saying "Oh we're not going to revolutionise the whole design of everything". But no, I think it's really good, and it makes a lot of sense. Especially if it's integrated directly into the OS rather than just being a little application layer on top, it would be great. It would be really useful, especially in what we've studied and what we do.

Participant 5

Link Behaviour

[Interviewer explains link display using subtle highlighting and the marker bar.]

Participant: What happens if there's two links on a line?

Interviewer: Yeah, so there are different ways you can deal with that. Maybe you can show that by showing overlapped marks in the bar, or maybe you use a single mark but vary the colour or shape of the mark, or something like that.

Participant: Also, another point here: what if you're colour blind? Or like, accessibility stuff...

Interviewer: Yeah, I think that's a good point. We should still be able to have enough contrast so that colour blind people can view links just fine, and we could even have special accessibility settings to improve things in this case. And we don't have to use a grey colour as pictured here. So, like, if you highlight text in Google Chrome today, it's relatively inoffensive and readable, yet has enough contrast to be seen in those kinds of situations.

Participant: But how does this sort of tie in if you're blind? In terms of, the normal kinds of accessibility options that you'd usually get with a Web browser. I'm guessing that's outside the scope of your project?

Interviewer: Well, yes. If you're blind, though, I would expect you to have some kind of screen reader program to help you out, and you can tab through links, and this kind of thing. Normal software accessibility guidelines apply.

...

Participant: What happens if you click on the marker bar?

Interviewer: So, that's not really defined within this prototype, but I think we do imagine that there's some behaviour there. So, for instance, if you have a link in a table maybe that isn't highlighted inline, perhaps hovering over the mark in the marker bar shows you the specific location in the link. Or, yeah, you could imagine clicking, right clicking... doing various things with the marker bar.

Participant: Other than pointing out where stuff is, what's the point of it?

Interviewer: Well at this point in the storyboard, pointing out where stuff is is the only point. And that its primary purpose.

Participant: But surely then... the point is you're not highlighting a word enough to attract attention, but then you've got a big marker bar to attract attention?

Interviewer: Well, it's not huge. I would imagine a relatively slim bar --- kind of the size of a scroll bar --- on the right side of the window.

Participant: What would be the point of looking at that instead of just thinking 'there it is, the link's there'?

Interviewer: Well, the idea is that often when you're reading text you don't want to be distracted too much by interaction opportunities. You want to know that they exist, but actually you don't want them to impede readability. So we have some highlighting to see links inline, but especially for links that maybe are entire paragraphs --- where actually you don't want to highlight them inline because that's going

to hurt readability --- or things like images where you can't highlight them, or just to affirm opportunities for things that were subtly highlighted that you maybe want to skim through: you can just look at the marker bar. There are also some other uses, one of which we'll talk about in just a minute. Does that answer your question?
Participant: Yeah.

...

[Interviewer explains link navigation, and how clicking links opens a new window.]

Participant: So would the previous window have closed?

Interviewer: No, the previous resource remains open.

Participant: Oh okay.

...

[Interviewer explains linkbase selection through the linkbase menu bar application.]

Participant: If I clicked on that, how many things would there be? Would there be tonnes?

Interviewer: Potentially, but I think what I was imagining here is that there are a handful of suggested linkbases that are useful in the current context. Or otherwise, you can click to open a dialog that has more in terms of discoverability and configuring what linkbases you want to see.

Participant: Right.

Interviewer: But just for quick access, there are some that you can enable and disable straight from the menu bar.

...

[Interviewer explains linkbase activation.]

Participant: So it hasn't altered the underlying document, it's just what sits on top of it?

Interviewer: Exactly right. The documents are exactly the same, but there are new links, with accompanying new marker bar marks and inline highlights.

Participant: Would someone actually do that?

Interviewer: I think so --- we'll get to some part of how you can create linkbases on the next page.

...

Participant: Can you have more than one linkbase being shown at a time? Because if it's just his, it'll just be one link.

Interviewer: Yes, you can. Yeah.

Participant: What if there are conflicts? Like there are two linkbases and they both highlight the same thing?

Interviewer: So you can certainly have overlapping links, yeah---

Participant: So what happens when you click it?

Interviewer: It's not very clearly defined within this prototype, but there are a few different things that you could do. One of the things you could do is, maybe when you click they visually separate out into options around your cursor, and then you can choose the one you want, right? Or something like this. Or if they're overlapping but not exactly the same, then maybe a rollover effect helps you distinguish between the two.

Participant: Yeah.

...

Interviewer: What are your general thoughts and feelings about this storyboard? In the context, remember, of research.

Participant: Yeah, I mean it seems useful. I'm not sure whether or not... given how many documents there are out there, how many linkbases there are going to be. Is it... does a person have a single linkbase per document, or do they just have one that's then applied to separate documents?

Interviewer: So, I think that's a good question. And it's... I think I was certainly imagining that you cover several documents within a linkbase, but I think there is a question as to well what's the granularity then? Do you have massive linkbases or do you have tiny linkbases or how do you manage that?

Participant: Because I think if they were just small, it wouldn't encourage people to...

Interviewer: Yeah, I agree that small linkbases probably wouldn't be that useful. But monolithic linkbases will probably create too many links for lots of people as well, so there's some middle ground I guess. "Here are my links for this topic."

Participant: Well, it could be in a big one and then you see that someone has got links in their linkbase corresponding to the document you're viewing, and then when you load it, it just shows that part... you don't need to load the whole thing.

Interviewer: That's true, you could partially load linkbases and think about it that way. But I think maybe in lots of cases, so for instance Toby here and his colleague, he's interested in lots of the kinds of connections that his colleague made between all kinds of information. So I think in lots of situations it is useful to see links in various places.

Participant: But how long are you realistically going to be looking at the same document for?

Interviewer: Well, exactly. That's why it's useful to cover multiple documents with the linkbase.

Participant: No, I mean, is someone going to be using a document so often that it's worth them creating links and stuff?

Interviewer: Well, I think in a research area you commonly gravitate around a number of related documents. Like maybe 10 or 20 of the most

- popular documents, which can then link to more obscure stuff too.
- Participant: Yeah, I could see that. This could be useful for something like an intranet I would say. ... But it's a question of whether people would keep stuff updated.
- Interviewer: Okay. So what do you mean exactly by that?
- Participant: What if the page that you're linking to disappears, or changes and has different content? ... How do you distinguish between what's still valuable and what's not?
- Interviewer: Yeah, that's a good question. This has been addressed to some extent in existing research, and there are some relatively reasonable things you can do. But yeah, it's a good point to raise. But let's say we handled the situation so that broken links are dealt with in a sensible way. What utility do you think this has in its applications to the wider Web rather than just an intranet kind of situation?
- Participant: What I'm thinking about when I see something like this is, like, Wikipedia. So maybe just a more visual way of seeing stuff... it would be kind of cool if you could just hover over a link and get a preview or something, that would be pretty good. But, yeah, the only sort of stuff I can think this would be useful for is very information heavy stuff, which is why I just think of Wikipedia. And I would say that's already pretty well linked. Information is kept up to date and maintained and all of that sort of stuff.
- Interviewer: Mmm. I mean, I think I would argue that Wikipedia is a very small portion of the Web, and actually tends not to be massively useful in research work. You tend to use Wikipedia as a starting point, but then don't spend much time there. Usually looking more at research papers and stuff like that, in my experience.
- Participant: Maybe. I don't read many papers to be honest. I read a couple, but if the document itself doesn't have enough information in then it's probably not considered a very good paper. Like, the paper should be all encompassing. I don't think you should have to go and read extra stuff to try and understand the paper.
- Interviewer: Mmm, I'm not sure I necessarily agree, but go on.
- Participant: Well unless it relies heavily on other work. But in terms of just linking stuff, then... I don't know... [Pause.] I can't really see when I'd use something like this. Or, like, want to create my own links unless it was between my own documents? Yeah, maybe as a way for organising my own documents. But I think when you start introducing other people's stuff, linking that all together... there's just too many variables on how stuff's going to change, and eventually stuff will just get out of sync... there's a high level of maintenance. Is it worth it? I don't know. That's what you'd find out: whether people actually maintain links and add links, and I feel like people would start adding links to a linkbase... like, unless people are taking this seriously in like an academic place, I think people would just create one or two and then just forget about it. And then you'd just be left with a load of like tumbleweed linkbases. I don't know how you're sort of mitigate that. And once you change

- like a document, then anything that references something that's just changed is just gone --- that's a load of work just gone.
- Interviewer: Well, potentially, yeah. There are certain ways you can keep the structure, depending on how the original document was organised. But yeah.
- Participant: I don't know. It seems almost too academic for something that would be considered useful for the general population.
- Interviewer: Okay. I mean, we are focusing here on knowledge work in particular. On information-heavy research kind of activities. I'm not too worried about the general population at the moment.
- Participant: But can this kind of behaviour not just be replicated by people keeping a list of like... at the bottom of an article say, the author just references a bunch of stuff that he used? An article containing references and footnotes and whatever.
- Interviewer: Well I think part of the advantage of linkbases here is that you don't have to be the author to link to something. So that allows you to associate pieces of mass media with more niche topics. It allows you to take someone else's research paper and make connections with work that maybe they didn't imagine. To do all these kinds of things. And make connections between things that people haven't related before, or that come from very different time periods.
- Participant: Mhm. But would an author always want that?
- Interviewer: Well, I don't think what the author wants really matters, does it? It's what you want as someone that's exploring and making connections between information. Umm... but yeah, I think you've raised some interesting points. We have a limited amount of time, so let's move on to the next storyboard.

Spaces

- Participant: What happens if you press 'Back' in that space you've just dragged windows out from?
- Interviewer: So, that's a very good question. If you're moving the windows as a group, maybe there's some kind of intelligence which takes the salient portions of history in the source space and let's you keep those. But if you're moving them one-by-one, it seems there's not much you can do there, because at first there wasn't much history and now you're kind of re-creating it.
- Participant: Yeah. So if you move three into a new one, in the new one you then click on a bunch of stuff, then you go back to the other one and press 'Back' --- what happens? Stuff like that.
- Interviewer: So, there are lots of edge cases here but, ultimately, separate spaces have entirely separate history chains. And, yeah, it's moving things when it maybe gets a bit weird.
- Participant: But what if, say, I clicked on a bunch of stuff in a space, and, like, the most recent one I keep in that space, but the second to last one I move into a new space. Then that's broken the history of the space.

Interviewer: Yeah, that's another good question. So, I think possibly what you'd want to do there is re-open the same resource again in the space. I have a feeling you don't want to skip over it just because it's moved. But there's certainly an argument to be made that maybe you want to just skip over it in the history if you're going back past it.

Participant: Mhm, yeah. The forwards and back stuff is pretty cool though.

...

Interviewer: What are your general thoughts and feelings about this one?

Participant: Well, multiple desktops are always good, aren't they? But then again, you've got to maintain them. If you just sort of forget about one and the resources in it are no longer useful... you've just got to make sure you stay on top of managing them. But yeah, I think... yeah. It's standard really, isn't it? Multiple desktops is fine. I think the idea of moving back and forward is pretty cool, that would be useful. But how do you manage having so many windows open, as opposed to say tabs? ... If I've got tabs open, I can sort of organise things in terms of tabs. Like I can put my leftmost tab to certain stuff, and know from the order of them which things I've opened recently. Whereas I don't really see how that's going to happen here.

Interviewer: So, you're talking about within a space? Because on some level, there's a correspondence between tabs and spaces in terms of maintaining separate history chains.

Participant: Yeah. In a way, but it's a lot more effort than just opening a new tab.

Interviewer: Well that's another question. I'm painting the picture that this is relatively efficient and very easy to do---

Participant: But, just like clicking a link and having it open in a new tab, and you can immediately just go straight back. And then, yeah, I don't know...

Interviewer: Yeah. So I think potentially you could get the same efficiency out of just using keyboard shortcuts there, but I think you're right in identifying that you don't have this constant visual display of what the other open things are and that they're available, and that makes it a little harder to switch between spaces as opposed to tabs.

Participant: It seems like a big, new trail of thought. If I'm working on stuff, I might be working on a few different pieces of work --- like different modules. I would split those into different spaces. But it's rare that I feel the need, within my browser say, to split out a bunch of different stuff I'm reading to different spaces. So, yeah. I mean, this is still useful if you're working on different stuff, but...

Interviewer: Yeah. I think part of the thing here as well is that there are two sort of different uses to tabs. One of them is separating

chains of history, in which case you probably want to create a new space, because that's the only way you can do that here. And the other one is just fast switching between resources, which is maybe what you were getting at when you were saying "what do you do when you have lots of windows in one space?", right? What's the efficient way to switch between them? Because with tabs you can very easily see where they are and switch between them. And if we're talking about that second problem, that's something we might address in the storyboard after this one to some extent.

Participant: As far as I can tell, you're reusing spaces for both different work and different trails of thought on the same piece of work, which to me are different things. And I'm not sure I'd use them for the same thing. If I'm working on, like, logically distinct pieces of work --- like, I'm doing some crypto or I'm doing some functional programming --- I would have those in different spaces. But if I was just looking at different resources for one of those things, and it was a slight tangent, I can't see I would need a new space for that. But I can see the use of it if I'm working on different stuff --- the initial bit.

Interviewer: Yeah, that's interesting. I think I would maybe say there's an argument sub-tasks being their own, logically distinct things. But I think what you're talking about is more... if you like control-clicked a link to open it in a new tab, and then just do something that's just slightly tangential. So it's not a completely new sub-task, but it's just slightly tangential. And you think maybe spaces are a bit too heavyweight for that. Is that accurate?

Participant: Yeah. So, like, then I would have to keep flicking back and forth between spaces, because they're slightly related. Whereas, I can deal with the slight excess of tabs potentially... if I open too much stuff, or... at least I know where it all is. Whereas with spaces, I think of that as completely different like... almost like context switching. When I switch space, I want to be thinking about something completely different. That's how I think of it. Whereas that would be similar work that I want... and switching might, you know...

Interviewer: Okay, yeah. That makes a lot of sense.

Participant: But, yeah, I can see... you need something like this for what you're proposing. Some kind of separation of stuff.

Interviewer: So, with that in mind... what are your thoughts on solving the problem of history via a solution like this --- in which you have total separation between spaces --- and maybe a solution in which you have windows within a single space, but you provide the user some way to differentiate between different groups of windows that should behave in different ways.

Participant: But then I think that's just becoming into spaces again. And I think that would probably be more confusing. You're putting more onus on the user, having to constantly think about how to group their documents in this sort of temporary way, which I think is a bit... I wouldn't bother with that.

Interviewer: Okay. Well, I guess another possibility... what if you could have sort of sub-spaces, which are effectively like tabs. If you have spaces, and then you have sort of tabs --- each of which holds a whole desktop environment --- within each space. So you kind of get the same tab display you might get in Chrome, but within each space. Something like that, I don't know.

Participant: Potentially. But then it's like, where does it end? I feel like if you have too many options, you're just never going to utilise them in the right way. As soon as I start using too many desktops or whatever, I just don't. I end up just using one again. Like, at work when I used to have like three screens, I would try and maintain some kind of logic for what's on each screen, but as soon as you get too granular, you're just spending more time organising than doing productive work.

Window Search

Interviewer: What are your general thoughts and feelings on this one?

Participant: I'd probably just like... I don't know. If I'm in a space and I need to find a document, I generally know, like, what program I'm looking for and what the document looked like. So searching would be... I'm more likely to remember what it looked like than a portion of what I'm trying to find, I would say. If I was looking for a certain window on some Haskell problem, then I might not remember what I actually searched for to get it. I would probably just look at the windows and say "It was that one". Unless you could filter on say the application, maybe that would be useful. But then that would be analogous to just the three-finger swipe down on the Mac.

Interviewer: Okay. So you don't think you would remember a snippet of the title or the text or the author or something, that you could just have direct access through?

Participant: Uhh, not unless it was something I used a lot. I don't really know the author of anything. I just remember what the distinguishing features look like. So I don't think I would ever personally... unless I had like tonnes open. Like, I normally have quite a lot open and I just sort of scan and say "Okay, it's that one". Searching it would be, like, I don't know... if I could filter it well enough then it would be quick, but I would probably end up typing the wrong thing, and I would get confused and click on the wrong one... Also it depends how similar your resources are. Because if this is within a space, then all your resources might be pretty similar, so I don't know how specific your search criteria would have to be.

Interviewer: Okay.

Participant: Like, at the moment if I want to change window, I swipe up, look at which one it is, and just click on it. It takes me like a second. Even if I've got a load of stuff open. And especially with tabs, I know where the stuff is in the tabs --- I don't have to

- figure anything out.
- Interviewer: So, yeah. I think another part of this, though, is that often people don't have that many windows open because of tabs, right? Like, your text editor is going to have a bunch of tabs, and your Web browser is going to have a bunch of tabs. But actually, if you have all your stuff as separate windows --- and you're also in this big hypermedia system that spawns lots of windows --- you're potentially going to have a lot of windows that you need to filter through.
- Participant: Is that not just a flaw then of the fact that there's going to be too many windows? The fact you need this sort of searching thing? Could this not be mitigated by just having some form of tabs?
- Interviewer: Well, potentially, yeah. Tabs don't really work in a system like this, though, and tabs also aren't very well suited to certain kinds of tasks. They don't work because we allow resources of different file types, right? And it's a bit weird to group things by program if you're thinking in terms of links. So if you open a link and it then switches program and opens a new tab, there's this weird relationship between tabs in different programs.
- Participant: Why not just have everything, like... say you didn't want to have all these windows of different sizes overlapping. Just have one window, with each program running in a tab. So you don't have to have the same program in tabs.
- Interviewer: I mean, yeah, you could do that. What you're effectively of asking for there, though, is to replace windows as the fundamental building block in operating systems. That's definitely something you could do, but I think that probably comes with some fairly serious downsides. I don't think I would like it personally. But yeah, it's an option.
- Participant: That would be a much clearer way of me seeing the history. I wouldn't need to necessarily go back, I can just immediately... if I know "okay, I want to keep this stuff here but my trail of thought diverged here, so let's just click back to here". Otherwise you have to go and find the window... I don't know. I guess I use tabs a lot, so I can remember where stuff is without other cues.

Exploration Mode

[Interviewer explains related resources in Exploration Mode.]

Participant: So, are these related to the marker stuff, or... how are these related?

Interviewer: So, the idea is that there's some degree of intelligence here. It's not necessarily related to the currently enabled linkbases, but it potentially is. It's more like "here are some things we thought it would be useful for you to see that are related to the things you're currently reading". It may rely, for instance, on doing almost like a Google search for "what's the most interesting related content to this document?".

Participant: Okay, yeah.

...

Participant: Hmm... I guess this could be useful, but it's rare, like, that I look at suggested stuff. Because it's generally not intelligent enough to be useful or specific enough, is my problem. ... It would have to be really relevant, otherwise it would always just feel like ads or spam, so that's going to be quite hard. But yeah, it seems useful. This bit in particular seems really useful. [Points to transpointing pane.]

...

Participant: It would be kind of cool if you could just like jump straight into a new space. If you were like "Right, this changes everything --- let's take this into a new space". Because then that can make it more natural to create spaces and trails of thought, rather than you having to stop and think "Right, I'm going to make a new space now, which things do I need", whereas this is much more like branching, which I feel would be more useful.

...

Participant: I think it's a more natural way of browsing, almost... instead of trying to craft a search.

...

Participant: Could you not also have in here say a search bar that searches the whole Internet for resources? How would, say, starting a completely new trail of thought work?

Interviewer: Yeah, so that's not really encompassed in any of these storyboards, but I was thinking that it might be kind of like a Spotlight sort of search-type UI. But yeah, that might be a good fit to Exploration Mode too, if you can just search from there.

Participant 6

Link Behaviour

Participant: Quick question: what if there were, like, three links on the same line? Would there be an extra thick bar on the side, or would you just have one still?

Interviewer: So, that's a good question. It's something that hasn't really been answered in this prototype, but yeah I think you could certainly do something here... either with thickness, or making it look like there are multiple bars on top of one another, or colour or shape or something like that.

...

[Interviewer explains linkbase activation.]

Participant: So where do those links appear? Do they appear over the document he's already viewing?

Interviewer: That's exactly right, yeah. So you can see there are new link highlights and marks in the marker bar in the documents that were already open, and they adds to the links that were already there.

Participant: Ahh, I get you. So if you'd written about cookies and it mentioned cookies there... say there wasn't a cookies link before but he's written loads about cookies, he can link to his stuff about cookies.

Interviewer: Exactly, so the content of the document is identical, but the link structure is now different.

...

Participant: I have a question. What if he only wants a subset of the other dude's linkbase? Is there a way to do that? Say Victor's written about cookies and ice cream, and you only care about cookies. Is there a way to limit how much of his linkbase you incorporate?

Interviewer: So, that's a good question. I think in this prototype my answer would be that I don't think there is a way to do that. I guess it kind of cuts to what granularity should people make their linkbases at. Does Victor --- the guy who's linkbase it is --- does he maintain lots of small linkbases, or does he have one big linkbase with all his links?

Participant: I guess it depends how it's used by the users, yeah.

Interviewer: Yeah. It's a good question though, it might be valuable for people to be able to distinguish between different parts of a linkbase or something like that.

Participant: He could have a key to say "Oh these are about cookies", and then you could tick that you want to look at all the cookie links.

...

Participant: It seems useful. Although I'm not sure... why did you suggest the highlighting rather than normal links? Just for visibility, or...?

Interviewer: So... yeah, part of the reasoning was that with linkbases you potentially have many more links displayed than you did previously. And we don't want that to impact readability.

Participant: I would say that lots of highlighting, like, could impact readability even more. Could there be a way to, like, hide the links and then show them again while you're reading? Because, I don't know... if you see where the things are on the left, that's less invasive than having loads of like yellow highlighted bits, I find it quite hard to read highlighted text --- I'd rather just read a...

Interviewer: So, yeah. This highlighting here doesn't have to be in a bright colour, it might even be in grey or something like that. Or it could be like a relatively unobtrusive blue background highlight. But, yeah.

Participant: I feel like it definitely shouldn't be too obtrusive.

...

Interviewer: What are your thoughts on the way that link creation works?

Participant: It seems intuitive enough. Although, how would you communicate what button they need to hold to make links?

Interviewer: Yeah, I think that's a good question. I don't know.

Participant: Just some sort of preview when you first start using the system, or...?

Interviewer: Or maybe... because you probably also want some way of creating links where you detach the start bit from the end bit, in case the other window isn't visible at the moment, right? So, maybe the way of doing that --- when you do like a right click, 'Start Link', and then 'End Link' --- also has a key combination shown. And then it can be about whether you click and drag, or click once and then click again, or something like that. Something that seems relatively straightforward... but I don't know, yeah.

Participant: Mmm --- yeah, yeah.

...

Participant: This definitely seems useful. And I feel that making the link physically drawing a line feels intuitive. But you just have to teach it somehow.

...

Participant: And would this be running in some kind of program they start up, or how would it... is it something that's going to be running all the time in the background, or they start up like a...?

Interviewer: So, that's a good question. I had imagined that this just replaces your operating system, like this is just everything. Umm... but I think it's an interesting question --- do you want this all the time? I don't know.

...

Participant: You said for the video that you might be able to send it to a particular point, and even maybe to a particular end point as well. How... so would that then save the video somewhere, like a separate file? And where are all the links... are the links all Web-based still?

Interviewer: Yeah, so I think... I should highlight explicitly that I think all of these resources are potentially remote, just in the same way that you just access a URL and the resource just comes up today.

And so I think the way that links to a video like that would work is much like the way that YouTube time links work today, right? It just loads the resource at that point when you navigate to the link.
Participant: Okay, cool.

...

Participant: I hope it doesn't sound like I don't like it, because I do. It's just it's easy to say what you like, it's the idea itself. But then there are sort of the nuances.

Spaces

Participant: Quick question: would it work, like... when you go forwards, you can no longer go back and then forward to the previous thing? If that makes sense?

Interviewer: Uhh... can you give me an example?

Participant: So like in Word, if you make a change and then undo it, then you do something else, you can't go back to the previous change.

Interviewer: Okay, yeah. I think it would work like that, yeah.

Participant: Okay.

...

Participant: How are the spaces sorted along the top?

Interviewer: I think I would say it's chronological, but you can drag and reorder if you want.

Participant: Would there be a way... say you had like 50 spaces. Would there be a way to see them all at once, or would you have to go along five at a time?

Interviewer: That's a very good question. And it's a question I don't really have an answer to. But yeah, I think it's probably valuable, because you could potentially have a great many spaces.

Participant: Cool.

...

Participant: Again, this seems useful. The idea itself is sort of a no brainer, it seems useful being able to... yeah.

...

Participant: How far back would the history go?

Interviewer: Probably as far back as it can. Just the same as a Web browser today, you just store all the history.

Participant: Yeah, cool.

...

Participant: Say you're working on worksheet A... what did you call them, workspaces?

Interviewer: Just spaces.

Participant: Spaces. So you're working on space A, and you decide "Oh, that's interesting, but it's not entirely related... I'm going to make space B". Then you start working in space B. Then some time later, you decide you don't want any of this any more, and you decide to go "back, back, back, back, back". If you went into space A, you may have already done some more stuff in space A, so it wouldn't make sense to go back into the previous space. I think the spaces should be kept separate.

Interviewer: Okay.

Participant: So I think once you undo everything in one place, it should just delete it or say "You're out of stuff, do you want to delete it?".

...

Interviewer: What do you think about... I guess, because we don't have a Web browser, we can't really do the lightweight things of using tabs between different chains of history. What do you think about spaces compared to tabs in that way?

Participant: I hadn't even made the connection, but they are quite like tabs. But they're like... you can sort of spread out and see more things all at once. I think... it's nice to be able to have everything in one place, like sort of a hierarchy, but also I think that could very easily get messy, and you could be like "Oh I don't know what I've got in here anymore" once you have more than four windows or whatever. So you need some shortcut way of quickly seeing what you've got in a space. Like, I don't know, some sort of bar down the side... maybe you could even sort and group things, although I guess you've got the workspaces for grouping... yeah, again, it's a problem of granularity and how much you want to do it.

Interviewer: Okay, interesting. Hold that thought, because the next storyboard kind of concerns this topic.

Window Search

Participant: Would you prioritise one over the other, like the titles over the text?

Interviewer: I think so, yeah. Like, the typical kind of thing you might do in a good search algorithm.

...

Participant: This seems very useful. I have no real complaints either. Yeah... I can definitely see that this would be useful with the

spaces as well, because they might get cluttered after a while. I guess, at the end of the day the quality depends on how good the search algorithm is. Otherwise, I don't know if there's much to comment on. It's just search.

...

Interviewer: Do you think this solves a real problem, or is this just sort of a neat idea?

Participant: [Points to a large number of open windows.] That's a problem. If you have that much... I mean, I'm sure Jacob could find what he wanted in the end, but it's definitely going to save him time. And at the end of the day time is money or whatever, so it seems useful.

...

Participant: Would the search work... like, could you put in multiple terms? Or sort of use logic to filter your queries even further?

Interviewer: I think that would be useful, yeah.

Participant: Because like for Google you can do all the colon stuff, like using quotes to say "include this and this"... or even just like, if you could self-tag things while you're working, that might be useful. Although to be honest they probably have the word in somewhere anyway.

Exploration Mode

Participant: Firstly, would you make sure there aren't too many lines between the two documents? [Points to transpointing pane.] You don't want to get bogged down. Like, would you just do the two or three most important or most frequent connections? Because if the entire document is relevant, you don't want like a spiderweb of lines.

...

Participant: You said that once she's happy, she quits. Is there some way to save resources she particularly liked for later viewing? Because there's no guarantee that it's going to find the resources she found last time if she tries again.

...

Participant: Is there a way... say you feel that the six aren't really relevant. Is there a way to like refresh the search to find the next six most relevant, or something like that?

Interviewer: So, I think that's a very good question. Like, if we're employing some kind of intelligence, how can you tell it that it hasn't done a good job, or how can you direct it to do a better job?

Yeah, I don't know the answer.
Participant: Yeah, fair enough.

...

Participant: So if you clicked on say this one, would that open and fill the whole screen, or would it take the place of the middle one, with more links coming off of it?

Interviewer: You're not the first person to ask that... I think you could do it either way. The way I was imagining it working is that you click it and it exist Exploration Mode and opens as a new window. But certainly, yeah, you could have it so it becomes the new focus of Exploration Mode, so you can continue your exploration elsewhere.

Participant: Yeah. I think it would be useful to do both. So you can have a choice. If you like right click it and say 'Open out of Exploration Mode' or 'Carry on exploring through this one', or yeah.

...

Participant: Again, I think it's a good idea. There are just a few little things that I'm thinking about. I like it.

...

Interviewer: Does the fact that it's kind of a mode switch bother you at all? So if you want to check something, you kind of have to switch into a different space and then switch out of it if you want to go back?

Participant: Does it have to be a mode switch? Could it be, so, you press the button and it just shrinks and goes to the middle of the screen and other things pop out? You could still have all your other programs running in the background.

Interviewer: I mean, I guess... my thought is that if it's not really a mode switch, though, then what if you've done this and then you Alt+Tab to some different window. Then your focus is on a different window, but what you still have this weird thing still coming out and kind of weird stuff still going on? And you probably want to be able to exit it an go back to exactly where you want too. So to me it feels like it's kind of like entering a thing and exiting a thing.

Participant: Yeah, I think you're right. ... I don't think it's too bad, I think a mode switch makes sense after what you said. But yeah, definitely a way to save the links, and more freedom with what you want to do with the connections it makes.

...

Participant: Also, it could be cool to make that a permanent link. So if you like that connection to that document, you could, like, click on that and it turns into like a chain or something, I don't know. That

would be useful. So you could chain the ones you thought were useful, then do like a re-roll and get six new ones. And go "Oh that's good, I'll keep that one", then you can see all the chained ones at the end and start exploring properly. That might be cool.

Whole System

Participant: I definitely can see it being useful. It's just about how easy it is to learn, I guess? It seems very intuitive after you explained it to me, but not everyone is going to be able to sit down for half an hour and have someone explain it to them. You generally want tutorials to be quite quick and easy to find.

...

Participant: I think Exploration Mode is the most... new? Most different to anything I've seen or thought of before. So it's probably going to take the most thought and the most explanation.

A.3 High Fidelity Prototype Survey Responses

What are your general thoughts on the prototype? Does it make sense to you? Do you think it's well-executed?

Participant 1: Generally I would say it's very well executed, it makes sense and is easy to understand the user interactions throughout.

Participant 2: I really like the prototype, I can see a clear use-case for this form of system and how it could be integrated into my current workflow. It does explain the key systems well though there are some aspects that I have questions about, which I cover in the later questions.

Participant 2: I feel it does a much better job than any current system for linking documents especially from multiple source types. Currently we have hyperlinks on the Internet, but the great thing about this prototype design is the ability to link also between different sources like video, and documents, and having them all intertwined.

Participant 3: Yes the prototype is well executed and it all makes perfect sense.

Participant 4: Well thought out prototype, the flow of information makes sense. I do think it is well executed.

Participant 5: I love the video as a high-level prototype! Very professionally made and all functionality was made sufficiently clear by the explanations in the video. Nice job!

Participant 6: The system seems coherent and useful. Having a separate marker bar works well to bring emphasis to links without obstructing the actual content of the particular window.

Participant 6: I'm not sure I grasped the concept or use for pinning a link.

Participant 6: How would the system cope if multiple perspectives had a marker on the same line/link?

What do you like about the prototype?

Participant 1: I think the different link overlays (perspectives or whatever they are called) are a neat idea and work well to show different pathways depending on the requirement and knowledge of the user.

Participant 1: Neat to have the links appear at certain times in the video though, could this potentially work with audio and other time contextual sources?

Participant 2: It provided a clear demonstration of how the implementation would work.

Participant 2: It also provides a logical flow of how a perspective user would use the system whilst showing off its key features

Participant 2: I like the simplicity of the links and their versatility with different mediums (video, documents, web pages), and how the system works uniformly across all of them systems

Participant 2: It also seems a very logical system, that could be nicely implemented at an OS level.

Participant 3: Being able to change what is highlighted by changing the perspective is a nice feature. Having all the links on the page at once would be a bit overwhelming.

Participant 3: Keeping a log of actions on a system wide level is a cool feature which I can imagine being useful when doing something that requires many different applications.

Participant 3: The highlighting of links is also very subtle, which is nice. I like that you can link in videos as well.

Participant 4: The use of the marker bar in conjunction with changing perspectives a person with more knowledge could go through a different perspective than someone with little knowledge and find it just as useful. It would be extremely useful for learning.

Participant 5: I was a big fan of the subtle highlighting to prevent distraction while reading and the overall unobtrusive interface.

Participant 6: I like that perspectives can be individually toggled with multiple being active at a time using different marker colours. Having to switch back and forth between perspectives would become arduous and I can foresee the need to have multiple "perspectives" active at once for a particular workflow.

Participant 6: Toggling between perspectives appeared quick and simple allowing attention to remain on the workflow.

Participant 6: I like that the system builds on and reuses existing interface behaviours and so seemed simple to grasp.

What do you wish was different about the prototype?

- Participant 1: Slightly confusing given that you were showing wikipedia which traditionally has links to other pages. How does this work alongside the other linkbases? Are the built-in links simply an author defined linkbase for that page?
- Participant 1: What happens in multicolumned text wrt links in the marker bar? Could be a little confusing. Maybe a marker bar for every column but then that would be ugly, not sure how you'd address this.
- Participant 1: What happens if two linkbases have linked the same piece of text?
- Participant 1: I think the number of different colours could become slightly confusing trying to remember which linkbase each different colour link in the marker bar corresponded to. The colours didn't seem to match up with the colours of the associated link bases in the drop down where you can choose to add link overlay sources.
- Participant 1: It seems slightly confusing to have to sometimes click on one of the links in the marker bar on the side to fully highlight the corresponding link which only then you can see and then use. Similarly, what happens if you click items in the marker bar that don't correspond to hidden links?
- Participant 1: Would be useful to know what the sort of link is that I'm clicking, is it a video? Sound? As a user, I want to know what I'm about to open as I may be somewhere where I can't listen to sound for example. What happens if I click a matlab link and I don't have matlab? Is it worth even showing the link?
- Participant 1: Not quite sure how the dragging works to link documents. What happens if the document is obscured by an open window? How do I create a time contextual link in a video? Can I create a link to a specific section of video? Also, the actual demo showed you dragging to a small circle in the document, what is the rules around when these circles show up?
- Participant 2: I wished the prototype had information visible upon hovering over links, saying where they are going to, what perspective they are from etc.
- Participant 2: The generic grey highlighting on all links I would think could be problematic depending on the document etc. What happens if it has a grey background?
- Participant 2: Also, some hover-over style change, would be nice to show possible interaction. EG when your cursor hovers over an link, it changes to the Perspective colour etc? (If this feature was there, it wasn't made clear in the prototype).
- Participant 3: Only thing I can think of is the perspectives getting confusing, especially if there are many of them. Having descriptive names for them would probably help that however.
- Participant 5: 1) The number of different marker colours quickly became confusing --- an easily accessible key (perhaps as an overlay upon clicking a button on the toolbar or holding some combination of keys) would be useful.

Participant 5: 2) It might also be useful to be able to open internal links as a new resource/window (i.e. for a user to better preserve their place in an earlier part of the resource they're in) - apologies if this is already a feature!

Participant 6: I wish markers and links indicated the type of content that they linked to perhaps with a tooltip. The type of content you're navigating to shouldn't be a surprise you have to react to.

Participant 6: One of the main uses of the marker bar is to provide a constant view of all links on the current window. This use seems to get lost when applied to video content as marks only appear whilst there is a link. Perhaps something could be added to the seek/progress bar to show the same kind of information i.e. highlighting of the progress bar on sections where a link is present in the video.

What questions or suggestions do you have regarding pieces of the system not shown in the video?

Participant 1: Could be cool to maybe use the colours to indicate the type of document being linked to, e.g. red - video, blue - sound, green - document etc.

Participant 1: You didn't put anything about spaces into the video, was this intentional?

Participant 2: How is it dealt with the side colour markings, if there is more than 1 link on a line, and from different perspectives?

Participant 2: Was unclear where the new link made in the compression to informal overview example. As in where is it linked to in that document? The title etc?

Participant 2: What if the same word has multiple links from different sources? Does the top perspective become the default? Is there any other way to see the rest?

Participant 2: On the "JPEG codec example" page, there are 3 different visible colours of perspective; yellow, blue and black. It is later understood the perspective that is accurate that yellow represents, but is unclear where the black and blue perspectives originated from.

Participant 2: It was shown at one point of jumping back to a previous page that you had visited, however, it didn't make sense in that, couldn't I (the user) have just clicked back onto the page that was there (since it was still open). What would have happened to the other links i had then followed from that page etc?

Participant 2: Also it would be interesting seeing how the perspective sharing system works. What if I (the user) only want some of my links to be shared with X people? Can I make multiple personal perspectives, and share them individually etc?

Participant 4: In the likely event a person does not have a program required from opening a link there should be some coping mechanism or an alternative. Also maybe allow a feature that closes all open links related to a subject or related by how recent they were opened, it

- would not be so fun to close everything manually.
- Participant 5: 1) Where do the different perspectives come from, and how will they be ranked/chosen for a user to choose from in his/her session?
- Participant 5: 2) Could there also be a way some way to quickly display all links in a document, perhaps filterable by type? This would allow the user to quickly find the quantity of mathematical content in a file, or even provide a quick index of the mathematical knowledge required etc.
- Participant 6: Is there a permission system to perspectives? Perhaps I could publish a perspective as read only. New links created could simply be suggestions sent to the owner who has to vet them first. Will there be a more central interface to curate and manage "owned" perspectives.

Any other comments or thoughts?

- Participant 2: I like the simplicity of using colours to label what perspectives lines have links on, however I do fear the ambiguity that colours provide. EG if I disable the 1st perspective and enable another, would that then default to the yellow? Could become confusing, so maybe some consideration for that UI element could be done.
- Participant 2: Is there a system to remove accidental links? How would that work? Maybe right clicking and having a remove link option. What if there are multiple links though?
- Participant 2: How does linking a document to part of a web-page work? In the sense of, I create that link, I send X person my perspective, how do they then get access to that document? Does it then have to be uploaded somewhere by me? Does it just not provide them links to that user since they don't have that document, but then that would seem to contradict the linking ideology you put forward. Would all resources have to be available by networks? What security implications could that have, etc. My point here just being data availability and transfer-ability is questionable and needs considering.
- Participant 2: Overall, there are a few points (as raised above), that I would have questions with regards to implementation and usage of the system. However overall I would wholeheartedly believe this to be a very beneficial system, especially to me (as a student), and it is a very logical layout/links of data between different sources.
- Participant 4: I really like this idea for learning and seeing what other people thought was useful. It means you don't need to get distracted and have all resources available to you. Very nice!
- Participant 5: 1) The number of open windows could quickly become confusing, though I guess that is the responsibility of the user to manage!
- Participant 5: 2) A resource will only be ever be as good as the people who work with it and form connections for others to use. How will

such a system work at the start (i.e. when there are minimal, if any, links between resources)? And, perhaps more importantly, how are you going to police bad links (or even malicious ones)? Some kind of report system? Moderators? Apologies if this is outside the scope of what you're considering!

Participant 5: 3) Overall, awesome job!

A.4 Exploratory Study Informed Consent Form

Participant Consent Form

Exploring Multi-window Knowledge Work

Before you decide whether you wish to participate in this study, please take a moment to read the following information carefully.

This is an exploratory study to understand how people carry out information-heavy tasks using contemporary, multi-window computing environments. You have been selected for participation based on your background in the subject of computer science, but are by no means obliged to take part.

If you wish to participate in the study, it will consist of a single ~90 minute session. In the first ~45 minutes, using a computing environment of your selection, you will be tasked with understanding and annotating the structure of a provided hex dump to the extent that is possible within the time constraints. This is an intentionally difficult task with several possible outcomes. A directory of resources will be provided to aid this process, and are the primary resources that should be used in the completion of the task. The study is about how you work with these resources to understand the hex dump. In the latter ~45 minutes, we will discuss the details of this experience.

Participants are free to withdraw from the study at any time. By participating, you consent to any recordings or notes from the session being used in the larger project to which this study belongs. This data will be anonymised, stored securely, and will only be used for the purposes of this project.

I have read and understand the information above and agree to participate in this study.

Print Name: _____

Signed: _____ Date: _____

Researcher: Joe Savage <js2322@bath.ac.uk>

Supervisor: Leon Watts <l.watts@bath.ac.uk>

A.5 Storyboard Evaluation Informed Consent Form

Participant Consent Form

Advanced Hypermedia Storyboard Evaluation

Before you decide whether you wish to participate in this study, please take a moment to read the following information carefully.

This is a study to evaluate a handful of advanced hypermedia design ideas. You have been selected for participation based on your background in the subject of computer science, but are by no means obliged to take part.

If you wish to participate in the study, it will consist of a single ~45 minute session. Over this session, you will be guided through a total of four storyboards, each of which roughly sketches out a single idea. Following the presentation of each storyboard, a short discussion will take place to gauge your opinions and thoughts on the design. The study is primarily about these thoughts.

Participants are free to withdraw from the study at any time. By participating, you consent to any recordings or notes from the session being used in the larger project to which this study belongs. This data will be anonymised, stored securely, and will only be used for the purposes of this project.

I have read and understand the information above and agree to participate in this study.

Print Name: _____

Signed: _____ Date: _____

Researcher: Joe Savage <js2322@bath.ac.uk>

Supervisor: Leon Watts <l.watts@bath.ac.uk>