Chapter 11

Conclusions and Further work

Hypertext began as a vision of the complete integration of texts (Bush, 1945), and has since become the vision of the complete integration of all information in any media (including text, image, audio, and video). The depth and diversity of the Web has tremendous potential to provide such an integrated docuverse. This thesis has explored the issues and challenges surrounding the realisation of this potential through the process of Associative Writing — the authoring and publishing of *integrated* hypertexts which connect the writer's new ideas to the wider context of relevant existing work.

One of the first World-Wide Web browsing programs opened with a blank page ready for the links and jottings of the user (Berners-Lee, 1999). Since these early days, the Web has grown into the most successful and farthest-reaching hypertext system to date, linking at least 10^9 documents (Pennock et al., 2002). However, the Web owes some of this success to the popular NCSA Mosiac browser (the basis of many modern browsers), which shifted the early focus on contributing new ideas towards a 'read-only' Web, marginalising writing and linking to separate editing applications isolated from the browsing experience. Through systematically examining Web pages from the Internet Archive for evidence of Associative Writing, and by relating the Associative Writing process to a real-life dance performance analysis scenario, this work has shown that Associative Writing is an important and valid process (which should not be marginalised), and furthermore that some writers are overcoming the Web's one-dimensional content uploading model to write and publish (albeit limited) integrated hypertexts. In investigating the issues facing these writers, this thesis has identified a number of core challenges facing computer-based support for Associative Writing in the Web docuverse:

- The "lost in hyperspace" problem refers to "the tendency to lose one's sense of location and direction in a nonlinear document", and may be particularly applicable in integrated hypertexts in the Web where information from possibly diverse sources is connected by associative links.

Challenge	
Lost in hyperspace	•
Deep linking	•
Limited Web hypertext model	•
Link integrity	◦
Supporting the writing process	•

•	Addressed
◦	Partially addressed

Table 11.1: Response to challenges in current implementation of AWF.

- Copyright issues surrounding "deep linking" to existing content published on the Web could prevent writers integrating new contributions with it.
- Links, the basic building blocks for integrated hypertexts, are restricted by the Web's hypertext model to uni-directional, binary links, which must be embedded in documents 'owned' by the writer (a consequence of NCSA Mosiac's 'browse only' model): as a result the Web as a publication medium limits the expressive capabilities of the writer (reducing the range of what can 'be said').
- The Web's chaotic nature (documents are frequently edited, moved, or deleted in an ad-hoc manner) compromises link integrity, potentially permanently disconnecting integrated hypertexts from the existing context on which they build.
- Popular tools for writing new Web hypertexts do not adequately support the different activities involved in the Associative Writing process.

In response to these challenges, this thesis has introduced the Associative Writing Framework (AWF). AWF builds on existing open hypertext, Semantic Web, hypertext writing, and annotation approaches to provide a novel Web browser-based interface for supporting browsing and reading, annotation, link creation, and integrated writing. AWF has been designed and implemented to address each of the challenges identified above (Table 11.1):

- Although some supporters of the "lost in hyperspace" problem recommend that the use of associative links be minimised, or omitted altogether, this work takes the view that such restrictions are detrimental to the reader's experience, and therefore AWF enables writers to produce integrated hypertexts without restrictions on linking. In order to avoid the potential orientation problems arising from "muddled writing" (Bernstein, 1991), AWF takes Thimbleby's "engineering approach" (Thimbleby et al., 1997) by embodying writer orientation aids in the system design.

- In response to copyright issues surrounding "deep linking" to existing Web content, the AWF approach attempted to demonstrate some important advantages of (unrestricted) deep linking in the context of Associative Writing. Key observations included:
  - New integrated hypertexts can link directly to existing context rather than using hundreds or thousands of words to establish that context.
  - Integrated hypertexts demonstrate (using associative links) the reliability of the conceptual foundation being built on, and the innovation and significance of new ideas.
  - The bi-directionality of links in AWF's integrated hypertexts means that newer contributions become reachable from the older texts that they build on.
  - "Rigid hypertext makes a large hypertext seem smaller.. Complex and intricate structure makes a small hypertext seem larger, inviting deeper and more thoughtful exploration" (Bernstein, 1998a).
- AWF successfully addresses the limited Web hypertext model by adopting an open hypertext approach — hyperstructures created in AWF are managed and stored separately to document content, enabling the Web's notion of 'ownership' to be lifted (writers can annotate and link arbitrary documents), and facilitating better support for navigation (links are bi-directional and may be typed) and structure (links can share a source, giving the effect of multi-headed links). This approach therefore enriches the writer's expressive capabilities and supports the publication of integrated hypertexts in the Web.
- The challenge of link integrity in integrated hypertexts has been partially addressed by AWF. A *preventative* strategy has been applied to dealing with the "editing problem" as the writer writes and revises the new integrated hypertext. Extensions to the open RDF representation used by AWF to incorporate the Robust Locations (Phelps and Wilensky, 2000c) and Robust Hyperlinks (Phelps and Wilensky, 2000a) technologies have been demonstrated.
- In order to support the Associative Writing process, AWF introduces a novel interface and interaction designed to bridge the advantages of document-based (good for integrated context building) and map-based (good for context building and structure building, but not integrated) approaches. AWF's Annotate tool (integrated with the Internet Explorer browser) allows writers to mark significant material in existing content. Annotated content then becomes a potential link anchor which may later be connected to other annotations using AWF's Relate tool. Writers create links by 'drawing' connections directly onto the desktop; once captured, these links are displayed whenever the two endpoints are visible on the

desktop, following the movement of windows and scrollbars. By extending the functionality of the Microsoft Frontpage Web editor, AWF allows writers to build integrated hypertexts directly onto existing (marked and linked) context.

The ideas embodied by the framework have been evaluated in the specific real-world problem domain of dance analysis, in which a need for support in creating integrated hypertexts was identified. Initial indications were that the framework method is valid, and that continued work to promote and evaluate the more general applicability of AWF is worthwhile. The next sections consider some of the directions that this continued work could take.

11.1 Summary of Proposed Extensions to Completed Work

A number of extensions to the completed work reported in this thesis have already been proposed over the course of the previous chapters. This section therefore brings together and summarises these 'short term' directions.

11.1.1 Evidence of Associative Writing in the Web: Further Searches

Chapter 3 presented the results of two investigations which aimed to uncover examples of Associative Writing in the Web, and identified a number of issues that would need to be addressed in future searches. The number of 'false positive' results reported in the search of the Internet Archive can be decreased in the first instance by updating the search algorithm to correctly process preformatted content. Occluding other false positives from the result set may be more difficult since linked content regions can exhibit similar characteristics to functional regions and vice versa — it was therefore suggested that a Markov model could additionally be used to analyse the 'orderliness' of a region (where functional regions would be expected to exhibit a higher degree of orderliness). Also, the pseudo-random dataset used in the investigation may not provide a good representative sample of Web pages, so future searches should operate over a truly random dataset.

Alternative approaches were also considered. Web Connectivity Analysis studies linking across a collection of pages as a whole, rather than focusing on links in content regions of individual pages. The identification of "hubs" and "authorities" (Kleinberg, 1999b) could be a useful starting point for further analysis, since hub pages potentially contain many associative links. Measuring the "dominance" and "connectedness" (Jackson, 1997) of collections of pages may also help searches 'home in' on areas of the Web deserving more thorough investigation. Areas of interest would exhibit a high connectedness (high number of links) and low dominance (links distributed evenly across nodes).

11.1.2 Extensions to AWF

Chapter 10 compared the completed work on the AWF with other research, and in doing so proposed several extensions to the framework as part of a unified architecture. Although the current AWF prototype demonstrates the possibilities of bridging document-based and map-based writing approaches to support Associative Writing, many of AWF's components (including the structure service, web annotation, hypertext map, and link manager) were developed as minimal bespoke implementations for convenience and proof of concept. The anticipation of this direction therefore is that having demonstrated the strengths of the AWF approach, more fully-featured third party tools can be integrated into the framework. Such a unified architecture has the potential to provide advanced structure support and filtering (WebVise and Annotea), free-form annotation and information gathering (Multivalent Annotations and Hunter-Gatherer), emergent structure building (VKB and Tinderbox) synchronised with existing desktop visualisations, dangling link and failed location correction (Robust Hyperlinks and Robust Locations), and link and structure versioning (WebDAV).

11.1.3 Specific Extensions to AWF as a Dance Hypertext System

More specific proposals for extended AWF functionality (compared to the general extensions listed in the previous section) arose from the complex requirements of the dance analysis task, reported in Chapter 9. The positive evaluation results gathered suggest that AWF appears to provide, at the very least, an important testing ground for the intertextual analysis of dance in digital video form. In order for AWF to realise its full potential in this domain (and therefore facilitate a more in-depth evaluation study), it was proposed that the framework should be extended to incorporate features such as a recording/playback device to allow meaning makers to witness the construction of the hypertext (a feature which could be provided 'for free' by an integrated VKB component), filtering facilities to deal with heavily annotated documents (a feature which could be provided 'for free' by an integrated Annotea component), theme extraction (a feature which has been demonstrated by ScholOnto — see Section 11.2.1), and additional options for rendering the interpretation network built within AWF (for example, linear essays which could be printed out and handed in for assessment – see Section 11.2.3.2).

Furthermore, although the number of options given for defining the semantics of annotations (interpretations) and links (intertextual connections) in AWF was kept to a minimum for the purposes of the evaluation, it was proposed that AWF should be extended to embrace the wider range of types identified by dance researchers as being useful in dance analysis, and to make extensions or adaptations to these semantics intuitive.

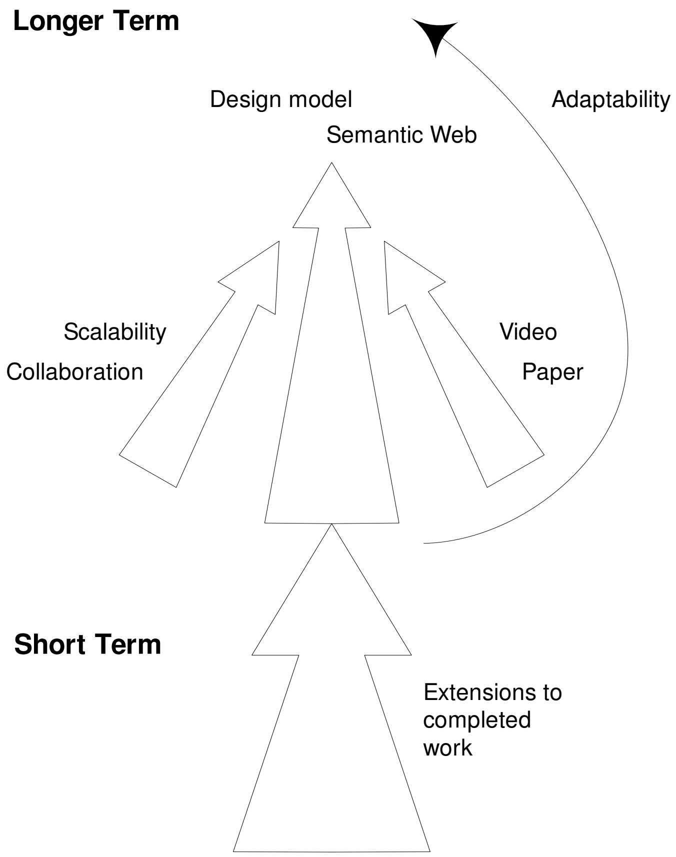

FIGURE 11.1: Road-map for future research.

11.2 Future Research Directions

This section proposes a number of potential future research directions for this work, which Figure 11.1 attempts to illustrate in the form of a 'road-map'[1]. In the short term, continued research will attempt to address the extensions to the completed work listed in the previous section. In the longer term, however, the core research direction adopts an increased focus on the role of Semantic Web technology in the Associative Writing process, with parallel explorations of different Associative Writing mediums, collaborative writing, and scalability increasing the range and impact of the core approach. A number of possibilities relating to adaptive integrated hypertexts are also suggested, offering some more open-ended avenues of exploration.

11.2.1 Associative Writing in the Semantic Web

Chapter 5 suggested that an initial response to the Semantic Web architecture and initiatives outlined in the chapter is that the focus is on *knowledge* — in initiatives like SHOE

[1]Note that this 'road-map' illustration is not provided merely as an abstract schematic, and should be interpreted as having a practical impact on the relative chronology, position, and contribution of each of the research directions proposed in this section.

and OntoBroker, Web documents serve merely to carry semantic information to facilitate machine understanding. In contrast, Associative Writing focuses on *links* as a means of achieving a deeper, *human* understanding of the document content by integrating new hypertexts with the existing work on which they build. However, the chapter concluded that more recent Semantic Web initiatives such as ScholOnto, OntoPortal, ESKIMO, COHSE, and WebVise+OHIF focus on both knowledge and links — knowledge about documents is not only used to find answers to specific questions, but also to help users understand and navigate documents more effectively — suggesting that these initiatives demonstrate a convergence between supporting *human understanding* in the information Web and *machine understanding* in the Semantic Web. This proposed direction would therefore further explore the theme of "machine-supported human understanding" by investigating the possibilities of Associative Writing in a knowledge-rich Semantic Web environment.

Chapter 5 introduced three major enabling technologies for the Semantic Web: knowledge representation (using machine-understandable metadata to describe resources), ontologies (describing interconnections between related resources), and agents (autonomous computer programs which collect and reason over metadata). By virtue of leveraging (and extending) the annotation templates specified by the W3C's Annotea project, AWF in fact already goes some way to implementing the first two layers of this simplified Semantic Web model. For each annotation created, AWF records metadata including the creator of the annotation, the URL of the resource annotated, and the semantic type assigned to the annotation by the creator. Similarly, metadata is also recorded for each link created. This metadata is described and exchanged between AWF components using the RDF syntax, providing the basis of a knowledge representation layer. The annotation and link templates themselves are part of the Annotea hierarchy, and are formally described using RDF Schema syntax, providing the basis of an ontological layer. This direction should therefore focus on the possibilities offered by allowing machine-based agents to operate across the metadata held by distributed AWF Servers. Recent Semantic Web initiatives (reviewed in more detail in Chapter 5) allow us to suggest some of these possibilities:

Inferencing facts SHOE, OntoBroker, and ESKIMO agents helped the user find out new facts about their community. For example, in a hypertext community context, ESKIMO could find answers to questions such as "what other *adaptive hypermedia papers* has *team X* produced?" and "who are the experts in *hypertext*?". In contrast, ScholOnto modelled user's interpretations of document content (a similar modelling approach is taken by AWF). Using sophisticated ontology inference rules, ScholOnto agents were able to help the user answer questions regarding intellectual lineage "Where did this idea come from, and has it already been done before?", the impact of ideas "What reaction was there to this idea, and has anyone built on it?", and inconsistencies "Is there contrary evidence to this claim?".

Uncovering perspectives By separating the conceptual representation of anchors (annotations) and links, AWF and ScholOnto support the co-existence of multiple interpretations by multiple users. ScholOnto agents were able to identify emerging perspectives (or "schools of thought") within the network of interpretation — common sets of concepts and claims on which a number of researchers had built their work — suggesting that similar derivations may be achievable by operating on AWF's metadata.

Structured navigation ScholOnto, OntoPortal, and ESKIMO demonstrated how an ontology-based Web portal can provide structured navigation of an information domain. COHSE and WebVise+OHIF demonstrated the possibilities of facilitating this navigation across the documents themselves by combining open hypermedia and Semantic Web technology. As an open hypermedia system, AWF may be in a position to leverage the advantages of this approach.

The extension of the existing AWF implementation to facilitate such a Semantic Web 'agent layer' has the potential to help the reader obtain a deeper understanding of the integrated hypertext, and in the case of the reader-as-writer, may facilitate the context building activity of Associative Writing (aiding retrieval). However, there may be some obstacles to the success of this approach. For example, Chapter 9 highlighted the concern that the formalisation of dance semantics in an ontology may leave little room for the more organic interpretation of intertextual connections surrounding a dance piece, calling for agents with 'fuzzy logic' capabilities. Additional Semantic Web layers (Berners-Lee et al., 2001) such as the proof layer (checking agent's reasoning) and trust layer (verifying that information has been provided by a specific trusted source — an issue addressed by TRELLIS) may also need to be considered.

11.2.2 Design Models for Integrated Hypertext

The success of an integrated hypertext may be directly related to "the author's ability to capture and organise the structure of a complex subject matter in such a way as to render it clear and accessible to a wide audience" (Garzotto et al., 1993). The process of designing hypermedia applications has been captured by several different models, the objective of each being to help writers (or more generally, businesses and organisations) produce well-organised applications or Web sites. Chapter 10 introduced the proposal that AWF be integrated with other approaches such as VIKI, VKB, and Tinderbox in order to facilitate emergent structure building within the framework: a hypermedia design model could provide additional methodology and/or guidelines for organising and structuring material during this activity. A design model for Associative Writing may also help writers overcome possible effects of the 'lost in hyperspace' problem by thinking in terms of higher level abstractions (Bernstein, 1991, P.J. Brown Position Statement).

This proposed direction would therefore investigate the possibilities of supporting some or all of the processes embodied by such models to help writers build well-structured, integrated hypertexts.

There exist a number of recognised hypermedia design models and methodologies, including the Hypermedia Design Model (HDM) (Garzotto et al., 1993), Relationship Management Methodology (RMM) (Isakowitz et al., 1995), Object Oriented Hypermedia Design Model (OOHDM) (Schwabe et al., 1996), Web Site Design Method (WSDM) (De Troyer and Leune, 1998), Web Modelling Language (WebML) (Ceri et al., 2000), and OntoWebber (Jin et al., 2001). Although each model exhibits its own advantages and disadvantages (Christodoulou et al., 1998), all of the models focus on imposing a macro-structure on a collection. The common modelling layers in this organisation process are the conceptual (or structural) layer, hypertext layer, and presentation layer. In broad terms, the conceptual layer captures the structure of the hypertext application domain, typically using an entity relationship model (HDM, RMM, WSDM, WebML), object-oriented model (OOHDM, WSDM), or ontology (OntoWebber). The hypertext layer provides the compositional and navigational structure of an application (the nodes and the different types of link between them), and the presentation layer captures the application user interface. Some models may further facilitate flexible customisation features in a personalisation layer (WebML, OntoWebber).

However, as (Carr et al., 2003) have noted, such organisation is undertaken at the level of a particular building block — an abstract data unit which may match a frame, paragraph, or region on a Web page. This is argued to be a fundamental problem because the components of content production for many Web sites (including integrated hypertexts) are the concepts and ideas embedded in the text itself. Some models (those based on HDM) go so far as to stipulate that links *should not* be anchored in the content of a data unit.

The Writing in the Context of Knowledge (WiCK) project[2] (Carr et al., 2003) attempts to address this problem by using Semantic Web techniques, proposing that the parallel requirements of macro-level site organisation and micro-level content linking could be satisfied by an interleaved model (Figure 11.2). Existing design models are represented by white layers (independent of the exact design model used); the WiCK micro-structure (greyed layers) provides mappings between the design model layers which consider the actual knowledge content of the documents. Ontologies are used to describe the relationships between concepts embedded in the document texts, and to expose these concepts to users (via a link service) in the appropriate context.

We have previously considered a research direction in which the Associative Writing activity takes place in the context of the Semantic Web, taking advantage of agent technologies in exploring, retrieving and understanding existing works and the new contri-

[2] `http://wick.ecs.soton.ac.uk`

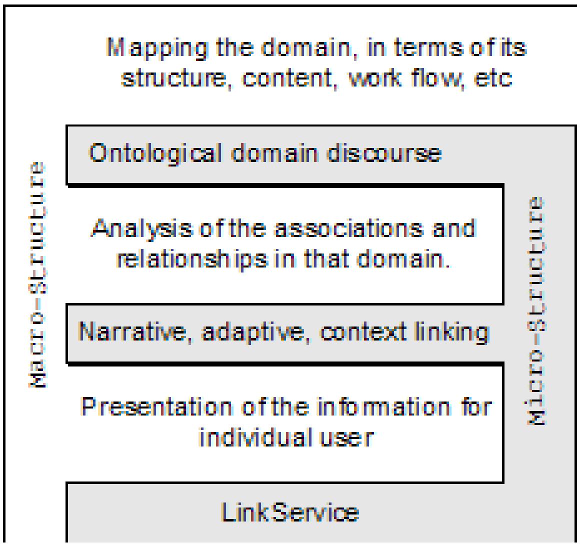

FIGURE 11.2: The WiCK interleaved design model.

butions that build on them. However, with the possible exception of CREAM's limited authoring component, such technologies do not yet help the writer actually produce (structure, write, and publish) the integrated hypertext itself (for example, for a writer to add their claims to the ScholOnto network, the writer must first produce a document describing the claims using a separate writing tool, and then create new concepts and claims for the work, rather than build directly on existing claims within the writing tool itself). Investigating the use of hypermedia design models in conjunction with these Semantic Web technologies to guide both macro-level organisation and micro-level content linking therefore represents a continuation of the core research direction (Figure 11.1).

11.2.3 Expanding the Medium of Associative Writing

Working in parallel to the core research direction (Figure 11.1), this section outlines two proposed approaches which attempt to expand the range of mediums available to the writer in the process of producing an integrated hypertext.

11.2.3.1 Hyper-video and Other Temporal Media

As the dance hypertext scenario described in Chapter 9 demonstrated, writers may require hypertext functionality in rich multimedia environments which incorporate many formats other than text and still images. It appears that this is not an isolated example; amongst these formats, audio and video media are becoming increasingly prevalent in hypermedia systems (Page et al., 2001). In the dance hypertext scenario, AWF allowed the meaning maker to attach an annotation to a specific segment of video/audio by interactively setting the start and end time offsets of the desired segment through a simple dialog (Figure 11.3). The annotation itself was then presented as a small icon

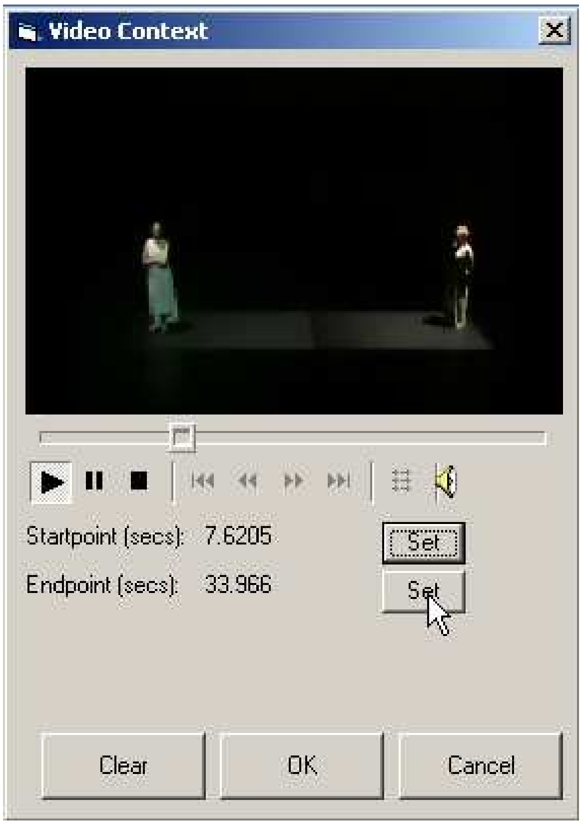

FIGURE 11.3: Interactively setting start and end time offsets of a video annotation.

next to the video object (Figure 11.4) — by clicking on the icon, users could watch the specific segment of video the annotation related to (albeit in a separate window). This simplistic approach has least two immediate drawbacks: (1) the dialog for setting video markers is currently limited — users are unable specify regions of interest within a particular frame or series of frames (for example, creating an annotation which describes the movement of one particular performer out of several visible in the frame at the same time); (2) an annotation (and links anchored to it) may be attached to a specific segment of the media, and yet its icon is visible for the entire duration of the media (and therefore may be displayed out of context). This proposed direction would therefore examine approaches which could be adopted to address these issues.

Despite the current drawbacks of video annotation in AWF, at least one approach seems to echo the annotation based interaction, taking the stance that just as annotations in textual media promote *active reading*, annotations in video promote "active watching" (Correia and Chambel, 1999). The AntV (Annotations in Video) system demonstrated this principle by allowing users to create and edit video annotations (stored separately in special files), and view annotations made by others (by sharing annotation files). Annotations could be applied to the entire video area or a specific region (for a specified duration), with the former being presented outside the video frame and the latter presented within the frame next to the annotated region.

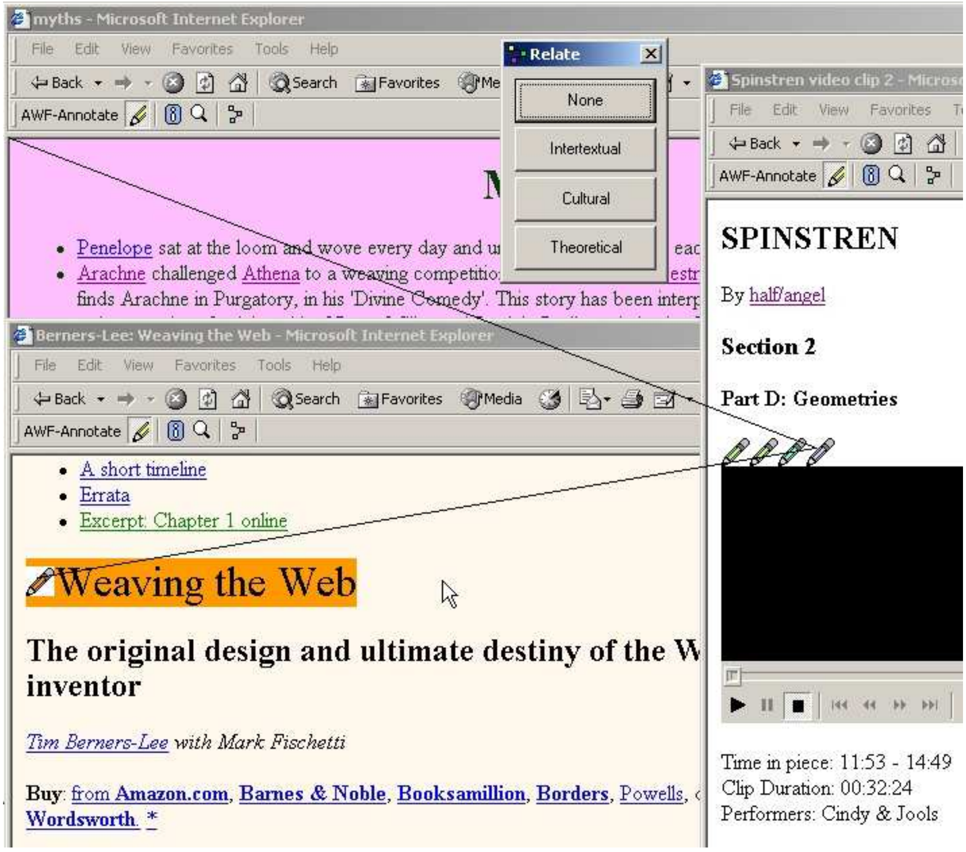

FIGURE 11.4: Display of annotation icons next to video object.

Other approaches provide more complex interactions with temporal media. For example, OvalTine (Smith et al., 2000) supported the collaborative authoring of "hypervideos" (Sawhney et al., 1996) — video streams containing *embedded*, user-clickable link anchors. OvalTine was designed to allow users to establish a link anchor on an *object* in a video stream (for example, a person's face); as the object subsequently moved around in the video field of view, it was automatically tracked by the system and the anchor region updated accordingly. Such an approach could be employed in AWF to, for example, attach annotations to a specific dance performer amongst many visible in the video frame without having to manually track their movement from frame to frame.

The Classroom 2000 trials (da Graça Pimentel et al., 2000) introduced the "linking by interacting" paradigm, in which live experiences of a recorded event, such as a lecture, were captured and used to form the basis of a hypertext ("linking by capturing"). The hypertext could then be augmented with further resources, such as class notes and results of group exercises, after the event ("linking by augmenting"). The linking by capturing process was facilitated using a video camera, microphones which captured both lecturer and class interactions, an electronic whiteboard and projector, and other displays (e.g. for browsing related web pages). Captured streams of timestamped events, including continuous video/audio and discrete sequences of slides, visited web pages, and handwritten notes, correspond to nodes of an automatically created hypertext. A similar approach could be used to capture responses and interpretations of video/sound media away from the user's computer (for example, at a theatre or concert) which would then form part of the writing context in AWF.

In terms of facilitating the temporal presentation of annotations and links in AWF, (Hardman et al., 1999) describe how the Amsterdam Model (Hardman et al., 1994) and the SMIL (W3C, 2001a) and HyTime (ISO, 1997) standards can be used to represent timing constraints in a multimedia presentation, including constructs describing "anchor time" (the display duration of a link anchor), "synchronisation arcs" (relationships which allow the start/end time of the destination component to be defined relative to a source component, allowing the presentation to remain synchronised even in the face of network delays), link behaviour (what happens to the playing presentation when the user follows a link to another presentation), and atemporal events (where part of a presentation continues to play whilst the user navigates through other parts). Similarly, HyStream (Page et al., 2001) demonstrates an open hypermedia approach to multimedia presentation, in which such constructs (expressed here using XML) are provided alongside the media stream as "continuous metadata" — the media player then displays the metadata with temporal relevance alongside the stream.

11.2.3.2 Physical Meets Digital

Implicit in the concept of hypertext is the assumption that people will read from computers rather than from paper (Price et al., 1998a). However, paper materials have many merits from a usability standpoint (Harper and Sellen, 1995; Mackay et al., 1995); they are familiar, portable, inexpensive, highly adaptable, and have very good readability characteristics (O'Hara and Sellen, 1997). As a result, there has been a long history of work attempting to create an electronic analog of paper in one form or another (Heiner et al., 1999). For example, we have already seen in Chapter 7 how the design of XLibris (Price et al., 1998b) and Multivalent Annotations (Phelps and Wilensky, 1997) systems was influenced by a desire to combine the advantages of paper-based reading with the power and flexibility of a computer. However, such devices may still be a long way from the convenience, flexibility, and portability of paper — "electronic reusable paper" is a long way off (Ditlea, 2001) — everyday annotation practices are mainly based on paper copies of documents (Sellen and Harper, 2002). This proposed direction would therefore examine approaches that rather that creating devices to replace paper, seek to add various digital capabilities to actual paper and pens, exploring how these techniques could facilitate a paper-based interface to AWF.

A number of approaches, including MEMO-PEN (Nabeshima et al., 1995), Digital Ink (Kasabach et al., 1998), Compupen[3], VPen[4], and Logitech's io[5], are able to capture

[3] `http://www.compupen.com/`
[4] `http://www.otmtech.com/vpen.asp`
[5] `http://www.logitech.com/`

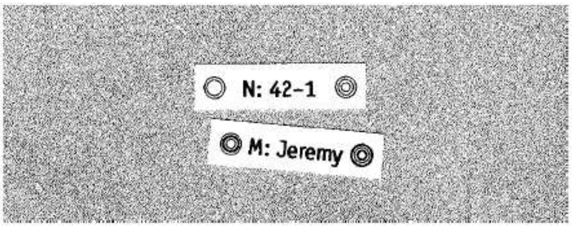

FIGURE 11.5: Example StickerLinks (Heiner et al., 1999).

the marks made on paper (e.g. handwritten notes, drawings) using a wireless electronic pen, and then transfer these marks to a computer or other electronic device. However, in the context of transferring annotations of printed paper documents to AWF, the system needs to not only record the marks made on paper, but also to 'understand' their significance in relation to the printed information on the printed page (*what* has been annotated?).

The PaperLink approach (Arai et al., 1997) is a step towards this goal, using a highlighter pen augmented with a camera to allow a user to link marks on paper to electronic documents. To create a link, the user highlights some text in the document (the anchor), chooses the 'hyperlink' command, and assigns electronic content (for example a document) to the anchor. The associated content can later be recalled on screen by moving the pen back over the highlighted text. StickerLinks (Heiner et al., 1999) go a step further, allowing users to create a link between two pieces of paper. Each page has its own sticker, which has the identifier of the page printed on it (Figure 11.5): to create a link, the user 'sticks' the StickerLink for the target page onto the source page. However, the source page has to be scanned back into the system to make the link electronic.

Anoto (Silberman, 2001) used a wireless electronic pen in conjunction with special 'digital paper' — paper with an almost invisible grid of dots printed across the entire surface. Tiny displacements of each dot enable the pen to uniquely identify its position on a map almost 1.8 million square miles in area. By mapping recorded positions to information stored in a central database, an electronic device can work out exactly what content on a page the user has marked. For example, a corporation may buy an 'area' of the map and use the dot patterns in a paper catalogue — to buy an item, the user simply ticks a box next to it. In an Associative Writing context, if a document is printed onto 'digital paper', a computer could work out exactly which content on the page the writer annotates using the electronic pen. By combining Anoto's 'digital paper' with XLibris' "ink anchors" (Price et al., 1998a), writers could potentially capture links by annotating related passages in different documents using a similar mark. XEROX's Intelligent Paper (Dymetman and Copperman, 1998) used a similar approach.

Several approaches go beyond the physical boundaries of paper to display peripheral annotations and links. By mounting a video camera and data projector above the user's physical desktop, Live Paper (Robertson and Robinson, 1999) and EnhancedDesk (Koike

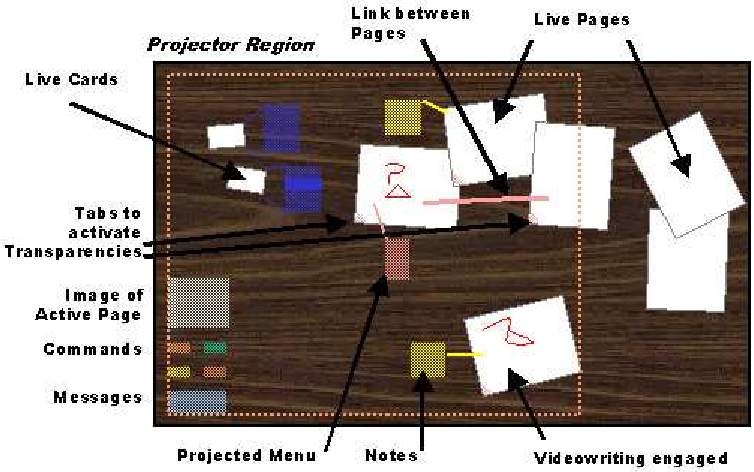

FIGURE 11.6: Conceptual Live Paper desktop (Robertson and Robinson, 1999).

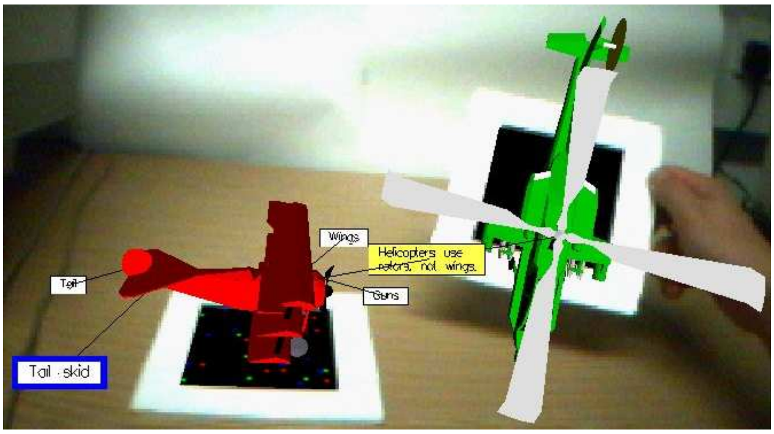

FIGURE 11.7: Displaying links between virtual 3D objects (Sinclair et al., 2002).

et al., 2001) tracked paper and hand movements, and could project applications, annotations, and potentially also links between related documents (Figure 11.6). For example, the Active Alice children's grammar tutor (Brown and Robinson, 1999) used the DigitalDesk (Newman and Wellner, 1992) to project lesson information (users could 'click' on a word with a special pen to view its usage) and quizzes (for example, "point to 2 verbs") alongside a paper copy of 'Alice in Wonderland'. (Sinclair et al., 2002) demonstrate how hyperstructures stored by an open hypermedia structure service can be displayed in a visual and tangible manner using 3D graphics. Special symbols printed on paper cards are recognised by a video camera, and an 'augmented' 3D reality is displayed on screen or through special goggles (Figure 11.7).

11.2.4 Challenges and Opportunities of Scale

Scalability is a key technical issue in hypermedia (Anderson, 1999); its importance in future work on AWF is therefore no surprise given this work's grand goal of helping writers to build an integrated Web. Although a number of open hypertext systems, including AWF, have been integrated with the Web (Anderson, 1997), simply providing access to the vast amount of information stored on the Web does not imply that such systems can scale to the same order of magnitude. This proposed direction would therefore examine some of the technical challenges which need to be addressed in working towards this goal. Anderson's experiences in enhancing the Chimera open hypertext system to support large hypertexts provide some useful starting points in which the resources to enable scalability are "readily available and easily applied" (Anderson, 1999).

XML Parsers AWF's components currently extract annotation and link metadata from the RDF interchange format using an XML validating parser. Anderson reported that such validating parsers construct a parse tree in memory in order to validate the document, and so do not scale well to large XML (or XML-based) documents. An event-based parser, which does not load parse tree into memory, should therefore be used instead (although validation must now be achieved through other means).

Relational Database To improve performance with large hypertexts, Chimera's link server was extended to use a relational database in place of the previous in-memory structure management routines; a storage component already adopted by the current AWF Server implementation.

Compound Hypermedia Operations A design goal for distributed hypermedia applications is to reduce the use of the network as much as possible, since networks operations are expensive (transfer time, error handling). By implementing compound operations that allowed multiple requests to be performed with one network operation, Chimera was able to increase its support for scalability. Network communications within AWF should therefore be examined and replaced with compound operations where possible.

User Interface The user interface of an open hypermedia system is critical to its success in supporting large scale hyperwebs. Changes to the Chimera interface were made to make the system responsive, informative, and designed such that the user maintained control of long duration operations. For example, before retrieving annotation metadata for a heavily annotated document, AWF should inform the user of the potential delay (perhaps suggesting filtering options — see below), and allow the operation to be cancelled at any time during transfer.

Filtering Searching and filtering facilities are critical if a user is to be able to understand and work in a large information space; Chimera was therefore extended to incorporate a mechanism to independently filter both anchors and links. When an operation produced a set of anchors or links, the result set was passed through a chain of filters which removed members not matching some criteria. Although

AWF currently does not provide a filtering mechanism, this feature has already been discussed in the context of the Annotea server. (Marshall et al., 1999) lists some filtering approaches for presenting multiple user's annotations of a single document:

- Overlay several sets of annotations from different readers, in order that the group can see multiple annotations at once (current AWF behaviour).
- Allow readers to access the annotations (and underlying text) apart from the document to create a compact representation of the passages that other readers have annotated ('meta-document' containing only annotated passages).
- Process different readers' annotations to produce a document which has passages emphasised where one or more readers annotated them, but preserves anonymity by not showing the original annotations. By looking at where people annotate, rather than how people annotate, areas of consensus can be identified — the annotations may create paths through a document that others may benefit from.

11.2.5 Collaborative Associative Writing

(Bouvin, 1999) identifies 4 levels of collaboration in Web augmentation tools — *personal* (relevant to a single user only), *shareable* (created structures can be shared, for example, by sending to another user or posting on a Web site), *asynchronous collaborative* (structures can be browsed or edited by a group of users taking turns), and *synchronous collaborative* (a group of users can collaborate through the structures in real time). As the dance hypertext scenario described in Chapter 9 demonstrated, AWF falls into the asynchronous collaboration category — although annotations and links created by a meaning maker would be immediately available to other members of the group, users received no notification that the interpretation network had changed. This proposed direction therefore would investigate how approaches to synchronous collaboration could be used to extend AWF's support for collaborative writers.

Many collaborative hypermedia tools, such as SEPIA (Haake and Wilson, 1992) and HyperDisco (Wiil and Leggett, 1996) have supported asynchronous and synchronous collaboration by propagating user events to every participating (and possibly distributed) tool. By monitoring such events, tools could make changes to their local presentations of the shared hypertext.

A recent example is the Arakne open hypermedia framework, which has been extended to support a range of collaborative working modes through the addition of a Session Manager (Bouvin, 2000). A session in Arakne was defined by a set of users, tools, and a coupling mode. Sessions were persistent so users could join and leave at any time —

rejoining a session would launch the hypermedia views relevant to that session, and depending on the coupling mode synchronise the views with the other session participants. Supported coupling modes were *uncoupled* (completely private), *loosely coupled* (shared creation and deletion of structure), and *tightly coupled* (completely synchronised interfaces). Such a facility in AWF could also be used by a single writer wishing to resume work where it was left earlier (for example, Web pages automatically reopened and repositioned on desktop). Users started by either creating a new session and inviting others to join, or by joining an existing session. To keep co-workers mutually aware of one another's presence, a continually updated "tickertape" display showed events generated by other users' actions (for example, "John left session *Report*"). In tightly coupled mode, these propagated events might result in more dramatic changes (for example, changing the current Web page). Every active session had an associated IRC (Internet Relay Chat) channel; discussions amongst co-workers could be persisted by using "chat notes" (IRC logs), which could then be linked to as any other hypermedia structure.

Arakne also supported the WebDAV protocol (Whitehead Jr and Goland, 1999), allowing co-workers to upload and lock documents residing on WebDAV compliant Web servers. Such versioning support may also be important in a collaborative setting — (Vitali, 1999) lists the advantages of versioning for hypermedia which include "keeping a good and reliable baseline makes authors more confident in doing experiments and trying out new development paths with their documents". Versioning may also provide an 'straightforward' solution to the problem of link integrity, as discussed in the context of AWF in Chapter 10.

(Haake, 1999) suggests that in addition to a group-aware structure editor (such as Arakne), an overview browser (providing co-workers with an overview of the current state of the workspace, and where people are working within it), search tool (to help find relevant parts of the shared workspace), and process browser (to help assess the past and make informed decisions about the future of the workspace) are important in supporting long-term synchronous and asynchronous collaboration.

11.2.6 Integrated Adaptive Hypertexts

A potential limitation of "static" integrated hypertexts built with AWF is that they provide the same page content and same set of links to all readers — even if the readership is relatively diverse (as in the case of the World Wide Web), the integrated hypertext will present the same static explanation and integrate the same resources for readers with widely differing goals and knowledge of the subject. By building a model of the goals, preferences, and knowledge of each individual user, and using this model throughout the interaction with the user, *adaptive* hypermedia systems can adapt the displayed information to the needs of that user (Brusilovsky, 2001). This proposed direction would therefore examine opportunities for combining AWF with aspects of adaptive hypermedia systems, in order to support writers in the creation of *integrated adaptive hypertexts* — integrated hypertexts in which the associative links are able to be adapted to the individual goals, preferences, background knowledge, and even perhaps the "context diary" and "context-of-interest" (Brown and Jones, 2002) of the readership.

The theme of adaptation to user requirements has already been touched upon at various points over the course of the preceding chapters. For example, QuIC (El-Beltagy et al., 2001) demonstrated how generic links could be applied to documents in context, according to the user's registered interests. The Connection Muse (Kendall and Réty, 2000) and StorySpace (Bernstein, 2002) writing systems, introduced in Chapter 6, demonstrated how writers could be supported in constructing adaptive hyperfiction. Both systems modelled the reader's "hyperspace experience" (Brusilovsky, 1996) during reading and adapted the navigation elements of the hypertext accordingly, for example, by making links visible or by changing the colour of links to reflect reader knowledge. Brusilovsky's taxonomy (Brusilovsky, 2001) defines two main types of adaptive hypermedia technology: content-level (adaptive presentation) and link-level (adaptive navigation support) — those systems falling under the latter category (including QuIC, Connection Muse, and StorySpace) seem most relevant to this direction. According to Brusilovsky's taxonomy QuIC represents an "adaptive link generation" approach — the system augments documents with generic links according to the user's interests and context. Conversely, Connection Muse and StorySpace demonstrate an "adaptive link hiding" approach; in both cases the writer defines how the structure of the hypertext will emerge, and this definition is used by the system to provide the adaptive navigation to the user. The adaptive link generation approach seems less well-suited to this proposed direction than adaptive link hiding, since Associative Writing focuses on links between specific ideas (as uncovered by the writer), rather than on generic links.

Other recent approaches to adaptive navigation support include AHA! (De Bra et al., 2002) and Auld Linky (Michaelides et al., 2001). In AHA!, an application consisted of a number of concepts (representing pages, high-level, or abstract entities). The user model (for each user) records a value for each concept (the knowledge level) — when a page is visited the knowledge level for that concept is increased (the knowledge of other concepts may also be increased, for example, reading a section contributes to knowledge of a chapter). The writer defines "requirements" — boolean expressions which operate on knowledge values — which determine when the user is "ready" to access a new concept, causing links to be hidden or revealed accordingly.

Auld Linky, a contextual link server, stored and served hyperstructures represented in the Fundamental Open Hypermedia Model (FOHM) (Millard et al., 2000). Context objects (interpreted by the server, and used to answer queries) and Behaviour objects (interpreted by a client, and used to facilitate the adaptation, for example, by describing how the user model should be updated) can be attached to any parts of the FOHM structure. The AWF Server API currently supports a `GetAnnotations` query, which returns a description of all the annotations for a specified document URL — the equivalent query to Auld Linky could also contain a Context object which, for example, specified that user was specifically interested in topic X. Only the annotations matching that context (i.e. those whose attached Context objects indicate that they are relevant to topic X) would be returned. (Bailey et al., 2002) demonstrated how FOHM could be used to represent various adaptive hypermedia structures (each with an associated Context object), including navigational links, tours (ordered objects), "levels of detail" (multiple representations of the same object, ordered from simplest to most complex), and concepts (multiple objects representing the same conceptual entity, for example, different media presentations of the same information): HA^{3}L (Miles-Board, 2003) used these structures to provide adaptive guided tours according to the reader's previous knowledge of the subject and preferred presentation medium.

Elements of the AHA! and Auld Linky approaches could be integrated into AWF to provide adaptive navigation support, although perhaps an important consideration in this respect is the interface mechanism for allowing the writer to specify the "behaviour" of the integrated hypertext. Such a mechanism could include assigning concepts or Context objects to annotations and defining the circumstances (relative to the user model) under which a link to an integrated resource should be displayed. Connection Muse and AHA! provide standalone form-based interfaces for specifying such behaviours, whereas in StorySpace adaptive links can be specified during the writing process — one possible line of investigation could therefore be extending the existing AWF Annotate and Relate components to support the specification of adaptive features on the desktop.

11.3 Concluding Thoughts: Looking Forward

This chapter has looked ahead to a future Associative Writing Framework, scaled to meet the demands of an ever expanding Web docuverse, supporting writers as they share and discuss their interpretations and contributions across mixed medias spanning the physical and digital universes. Autonomous Semantic Web agents operate across these networks of interpretation, augmenting writers' understanding of the global context of their work by uncovering new facts and perspectives, and providing navigation support through tangled webs of debate. Semantic Web technologies also play a role in helping writers to structure their integrated works using Web design models, and integrated adaptive hypermedia toolkits allow writers to specify how the work should be presented according to the varying goals, preferences, and knowledge of the reader.

Berners-Lee, reflecting on the difficulties of scaling human intuition and understanding on the Web, proposed that "our minds hold thousands of ephemeral tentative associations at the same time. To allow group intuition, the Web would have to capture these threads [...] as we work", concluding that "this all works only if each person makes links as he or she browses, so writing, link creation and browsing must be totally integrated. If someone discovers a relationship but doesn't make the link, he or she is wiser but the group is not" (Berners-Lee, 1999). Having presented in this thesis a framework for writing in the Web which achieves Berners-Lees goals for scaling intuition, these reflections hint at the exciting social implications of the wider application of this work: having spent several years filling up the Web with diverse and original content, will research efforts like AWF lead us into a new age of Web writing where new content is visibly and tangibly built on what has gone before rather than uploaded in isolation? In short, will we finally achieve the vision of hypertext's early pioneers in our integrated global information space?