

PADLs : supporting digital scholarship in digital libraries

Goh, Dion Hoe-Lian

2001

Goh, D. (2001). PADLs: supporting digital scholarship in digital libraries. *Aslib Proceedings*, 53(8), 325-337.

<https://hdl.handle.net/10356/91808>

<https://doi.org/10.1108/EUM0000000007065>.

Goh, D.H. (2001). PADLs: Supporting digital scholarship in digital libraries. *Aslib Proceedings*, 53(8), 325-337.

PADLs: Supporting Digital Scholarship in Digital Libraries

Author Information

Dion Goh

Division of Information Studies, School of Computer Engineering

Nanyang Technological University

Blk N4, #2a-32, Nanyang Avenue

Singapore 639798

Singapore

Tel: (65) 790-6290

Fax: (65) 792-6559

Email: ashlgoh@ntu.edu.sg

PADLs: Supporting Digital Scholarship in Digital Libraries

Abstract

This paper introduces digital scholarship, a process in which individuals perform all scholarly work electronically, working entirely with digital media. A proposal is made for patron-augmented digital libraries (PADLs), a class of digital libraries designed to support the digital scholarship of its patrons. PADLs not only provide facilities for search and retrieval of library artifacts, but also allow patrons to augment the library's collection with new artifacts such as annotations, original compositions and organizational structures. Finally, a prototype PADL (called Synchrony) providing access to digitized video segments and associated textual transcripts is described. Synchrony allows patrons to search its collection for artifacts, create annotations/original compositions, integrate these artifacts to form synchronized mixed text and video presentations and, after suitable review, publish these presentations into the digital library if desired.

Introduction

Digital library research is mostly focused on the development of large collections of multimedia resources and advanced tools for their indexing and retrieval. While these efforts are essential, it is important to recognize that the ultimate goal of a library, whether physical or digital, is to serve the scholarly needs of its users – whose objectives are not solely the retrieval of library artifacts. Users instead seek these artifacts (the items that constitute a library's holdings) in order to manipulate and combine them to produce new artifacts.

This observation is especially evident in scholarly (work-oriented) settings in which patrons peruse existing artifacts in order to produce new ones. Consider as examples, three commonly occurring scenarios:

- (1) a faculty member of a university would invariably seek library artifacts (such as books or journals) for the purposes of composing a journal article;
- (2) an information analyst working for a privately owned organization must acquire various artifacts in order to produce a report for a staff meeting;
- (3) a student assigned to produce a term paper must acquire and peruse library artifacts for its successful completion.

The Use of Library Artifacts

While these examples portray users of library artifacts in various situations, two common themes are apparent:

- (1) library artifacts are sought in order to complete a task – typically the production of a new information artifact such as a journal article, a staff report or a term paper

(2) these new information artifacts are disseminated – through a formal publication process for the journal article, through handouts and a presentation in the case of the staff report, or through the submission of the paper to an instructor in the case of the student’s term paper.

Studies of library artifact use support these observations. For example, Levy and Marshall [11] observed and interviewed a group of information analysts, their managers, information assistants, and technology providers in two organizations in order to gain insights into the use of libraries. While acquiring documents (artifacts) was a crucial component, this represented only an initial step in the analysts’ task. Once completed, these analysts would then annotate the documents as a means of interpreting them, produce new artifacts, and finally disseminate them. In addition, analysts would commonly share documents and other interpretative structures of documents with other analysts, as well as establish and maintain “reading rooms” which serve as collections of reference materials for the benefit of others involved in similar work.

Likewise, Stone [22] studied humanities scholars and identified five steps that scholars performed in their studies: (1) thinking and talking to others, (2) reading existing material on a topic, (3) studying original sources of information and making observations and notes, (4) drafting a document on what has been found, and (5) producing a final document based on the draft.

If library use indeed extends beyond search and retrieval, what types of activities do patrons perform? In a study of library use by O’Hara et al. [18], 25 PhD students in the arts and humanities at Cambridge University were asked to complete a diary of their document-related

research activities during a working day. Information recorded included the nature of the research activity, time taken, documents used, support activities performed (such as annotating), and place where the activity was conducted. At the end of the working day, the subjects were interviewed for approximately an hour. They were asked to elaborate on the information recorded in their respective diaries. Using the data collected, a model for document-related activities by library users was developed. The model characterized scholarly research as a complex process involving searching, information retrieval, reading, information extraction, annotation, review and writing new compositions. These processes were iterative in nature and occurred over varying periods of time.

These activities are similar to those found by Case [3] in interviews with 20 historians to determine their use of information. After searching from a variety of information sources, historians would make annotations and copies of the material, arrange and index the material according to their needs, and then produce an original work using the information gathered. Case also found that these activities were often performed concurrently within and across projects.

Four activities that occur over library artifacts may thus be identified. Firstly, they are **acquired** from a library's collection to solve some specific problem. This is typically performed through an iterative searching and/or browsing process [2, 12]. A second activity involves **organizing** the acquired artifacts to make better sense of the information at hand within the context of the prescribed task. Also known as *information triage* [15], patrons filter the artifacts to determine the relevancy of each artifact, as well as create various organizations for the artifacts to allow them to be used efficiently and effectively.

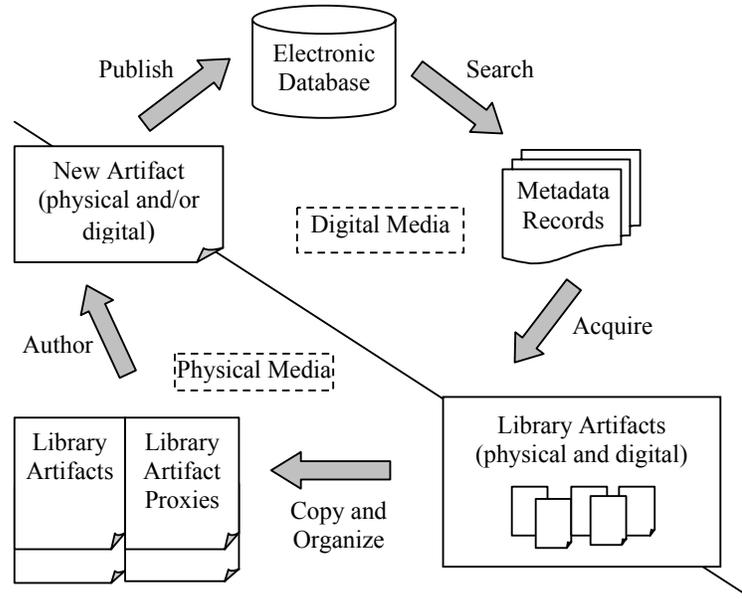
The third activity involves the **authoring** of new artifacts using the artifacts already acquired. Artifact types are varied and may range from annotations, to documents and organizational structures (such as indexes). Finally, artifacts are **published**, that is, the newly authored artifacts are disseminated. The audience of this artifact may be personal (for private use), public (for use by interested parties), or a selected group of users. Methods of publication are also varied and may occur through formal print channels (such as books and journals), presentations, informal handouts, and through the World-Wide Web.

Digital Scholarship

Traditional (as opposed to digital) libraries, with the majority of their holdings in physical form, typically promote a form a scholarship termed in this paper as *paper-based scholarship*. Here, physical media, predominately paper, play a major role in the scholarly use of library artifacts. For example, although patrons may use electronic databases to search for artifacts, the resulting metadata records point to both physical and digital artifacts, requiring patrons to switch between digital and physical domains in order to accomplish their tasks.

Figure 1 depicts paper-based scholarship as a cyclic set of transitions occurring in both the physical and digital domains. Artifacts (physical and digital) are located electronically through their metadata records. Since scholarship is (mostly) paper-based, copies of physical artifacts (or their proxies) are made for incorporation into the work process. Digital artifacts must also be converted to physical form before they are used [11]. These copies are then organized, and used to author and ultimately publish new artifacts which again may either be physical or digital. The work cycle is completed when the artifacts are incorporated into the library and metadata records are generated for them.

Figure 1. *Paper-based scholarship in traditional libraries*



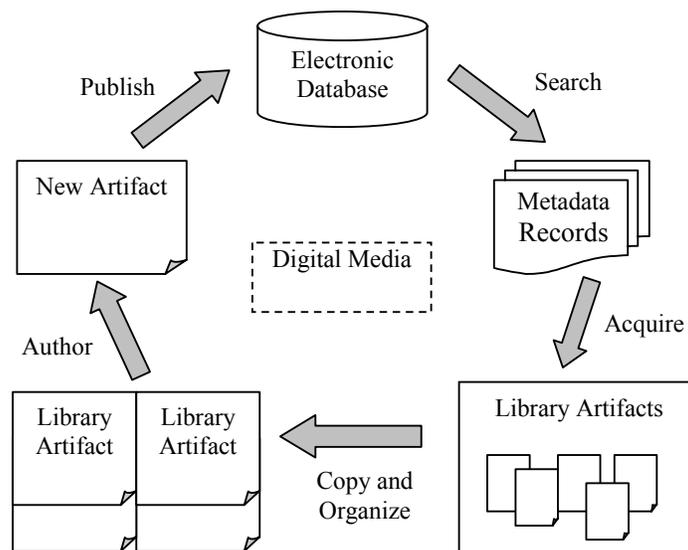
Digital libraries however provide new service opportunities to patrons as well as an expanded set of informational data types [17], and thus have the ability to promote *digital scholarship*. As shown in Figure 2, patrons can now perform their scholarly work electronically, working entirely with digital media. Using tools that interface with the digital library, patrons are able to search and acquire library artifacts, organize them to form coherent structures suitable to the task at hand, author new artifacts, and publish them electronically for future use.

Digital scholarship offers several advantages over paper-based scholarship. These include:

- (1) A single access point for library artifacts. Patrons are able to acquire all library artifacts in one location – at the computer. There is no longer a need for a two-step acquisition process in which patrons first search electronic records for artifacts of interest and then physically locate them.

- (2) New data types and new ways of access and manipulation. Digital media provide new opportunities for patrons to interact with library artifacts not previously possible with paper-based artifacts. Data types such as audio and video can now be used directly in the scholarly process. Patrons can search within artifacts, combine and edit portions of existing artifacts to form new ones, create links/associations between artifacts, and so on.
- (3) Shorter publication times. Paper-based artifacts typically take between 12 to 36 months from submission to publication excluding actual authoring time [5]. The digital medium has the potential to shorten such times by supporting online layout/formatting/editing, and electronic refereeing services, as well as removing the transitions between physical and digital media.

Figure 2. *Digital scholarship in digital libraries*



Patron-Augmented Digital Libraries

There is no doubt that traditional library models, in which searching is the main service provided to patrons and scholarly work is mostly paper-based, have utility. However, we postulate that in many instances, an expanded model of digital library services would benefit patrons. That is, digital libraries should provide services that encompass not only searching, browsing and retrieval, but an entire range of services that support patrons' digital scholarship from task inception to task completion.

The question becomes one of the types of services that should be supported. Returning to the discussion of artifact use, a plausible starting point would include services for acquiring and organizing library artifacts, together with services for authoring and publishing new artifacts. Hence, we propose *patron-augmented digital libraries* as a class of digital libraries that provide acquiring, organizing, authoring and publishing services to patrons. A patron-augmented digital library (PADL) is one whose holdings are enhanced by the digital scholarship of its users – both librarians and patrons contribute to the evolution of a library's holdings.

In the PADL model of use, librarians populate the digital library with artifacts that meet the goals of the library. At the same time, patrons may augment the PADL's holdings to meet specific needs through new artifacts such as documents, annotations or other organizational structures over the existing holdings of the library via the support services offered by the PADL. Often, the results of a patron's task (the newly authored artifacts) are deemed useful to the community at-large. When this happens, the patron may want to publish the artifacts for the benefit of others.

The term “artifact” used in our research refers to any information-bearing object that is accessible by a patron. Two major classes of artifacts are distinguished. *Information artifacts* are artifacts that contain information about a topic and are obtained either by librarians for the purpose of populating the library or by patrons who create and publish them into the library. Examples include electronic books, journals, and so on. *Patron-augmented artifacts* on the other hand, refer to artifacts produced by patrons and incorporated into the digital library after a review process. These may fall into three categories: (1) *structuring artifacts* which are used to organize other artifacts, (2) *annotations* which provide commentary and context to other artifacts, and (3) *original compositions* created by patrons. Patron-augmented artifacts become reusable information artifacts through the publication process.

It is important to note that while a PADL is designed as an environment for digital scholarship where patrons author and publish artifacts, a system of checks and balances must be in place to ensure the quality of the artifacts produced. For this reason, PADLs must include support for *publishing policies* that determine if an artifact considered for publication meets the goals and standards of the PADL. In other words, artifacts are subject to reviews, and these may be as stringent or flexible as necessary depending on the stakeholders of the PADL.

PADL Services

The facilities provided by a PADL are based on a model of digital scholarship termed ASAP [8]. This model suggests the need for tools that allow patrons to acquire artifacts from the PADL, organize these to put them into the context of the task, create new artifacts, and finally publish these new artifacts back into the PADL for future use. Hence, the minimal requirements for establishing a PADL would include the following services.

Storage and Retrieval

A fundamental requirement in all digital libraries is the support for services to store and retrieve artifacts, and in the case of PADLs these artifacts would encompass both information and patron-augmented artifacts. Two important features of a storage and retrieval service are: (1) the ability to accommodate artifacts of different multimedia types, and (2) the ability to deliver artifacts to patrons through browsing and searching modes.

Publishing

The publishing service functions as an intermediary between a patron who wishes to publish an artifact and the storage and retrieval service responsible for incorporating it into the PADL. A typical publishing service would acquire the artifact from the patron, obtain the necessary metadata for it (from the patron and/or analysis of the artifact), forward it for review, and upon acceptance, communicate with the storage and retrieval service for the purposes of storing the artifact in the PADL. Once again, it must be stressed that publishing policies must be implemented to ensure that published artifacts meet the standards and needs of the PADL.

Manipulation

The manipulation service is responsible for delivering the model of digital scholarship to the user, and provides the interface through which the patron interacts with the PADL. All user requests come from this subsystem and all results are returned to this subsystem. The tasks supported by this service include searching/browsing of artifacts, organizing/structuring of acquired artifacts, and authoring and publishing of new artifacts. While these tasks may be provided by separate tools, one advantage of a single tool functioning as the access point to

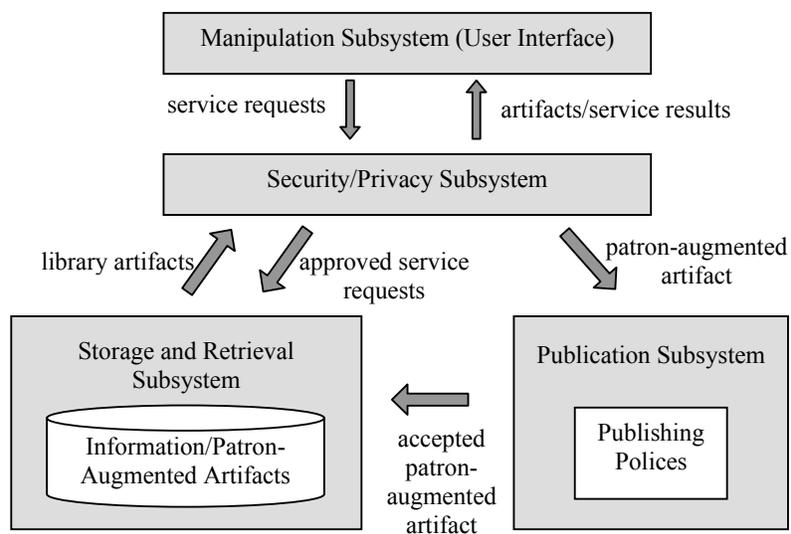
the entire PADL is the lower cognitive overhead required by patrons in learning and using the PADL facilities.

Security/Privacy

A PADL may be utilized by a large number of patrons, and as such, mechanisms must be present to ensure that artifacts, service requests, and personal information about patrons are secured from unauthorized access [1]. For example, a system for enforcing access rights is necessary to determine if a patron is able to manipulate (read/write/annotate/reference) published artifacts. Likewise, mechanisms are necessary to ensure that only a patron has access to his/her own personal artifacts and work areas in the PADL.

Figure 3 shows the conceptual architecture underpinning PADLs. Each service is supported by a separate subsystem that interacts with other subsystems in response to users' requests or actions.

Figure 3. *A conceptual architecture for PADLs*



Synchrony

Synchrony [7, 8] is a prototype PADL system that is designed for the purposes of digital scholarship. It allows patrons to search and retrieve artifacts from the library's collection, organize these artifacts to meet the needs of their tasks, author new artifacts, and publish these new artifacts into the digital library.

Synchrony's collection of artifacts consists of digitized videos of speeches given by former President George Bush (Senior) and their corresponding textual transcripts acquired in collaboration with archivists at the George Bush Presidential Library and Museum. The transcripts are full-text indexed at the paragraph-level and made available to patrons via standard query operations. In addition, each paragraph is associated with its streaming video segment, allowing patrons to view search results in text-only, video-only, or synchronized text and video formats.

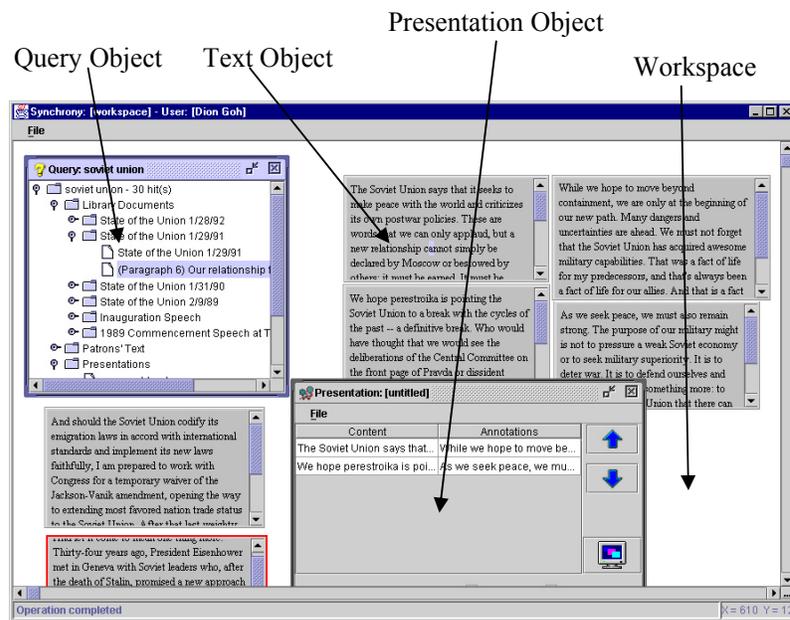
The collection also contains artifacts authored by patrons and these fall into three classes: *original compositions*, *annotations* and *structuring artifacts*. Original compositions are text-based documents that patrons author and publish into the digital library. Annotations are also text-based documents, but are designed to provide commentary and context to other artifacts. Presentations serve as structuring artifacts in Synchrony. These composite entities consist of sequences of artifacts, each of which may contain a video segment of a speech, its corresponding textual transcript and an annotation/original composition displayed in synchrony. Associated with each presentation is a table of contents that allows patrons to navigate to any sequence within the presentation. Artifacts contained within the presentations are referenced, not copied. This allows modifications made to individual artifacts to automatically propagate to presentations if desired.

Synchrony is so named because it allows patrons to author and publish synchronized text and video presentations.

The User Interface

Synchrony's user interface is patterned on a spatial metaphor and represents a large, 2 1/2 dimensional direct manipulation workspace in which patrons manipulate and organize objects of different types such as text and presentations. The interface is depicted in Figure 4 and consists of two major entities: the *workspace* and *library objects*.

Figure 4. Synchrony's user interface



The workspace forms the background of the interface and functions much like a physical desktop on which items are placed and a patron's tasks are performed. Library objects, that is

the information and patron-augmented artifacts in use by the patron, are positioned on this workspace. Objects may be arranged (by selecting and dragging an object on the workspace), resized (by selecting and dragging an object's borders) and visually altered (by modifying an object's properties such as color) by the patron to create information structures suitable to the current task. In addition, scrolling and panning are supported to allow patrons to view different portions of the workspace.

Library objects are the means with which a patron accomplishes his/her digital scholarship. They represent the information and patron-augmented artifacts as well as the results of a patron's tasks in the PADL. Library objects fall into four basic categories:

- (1) **Query objects** represent the results of a search, with each query object representing one result set. Queries are performed against information artifacts (speeches) and/or patron-augmented artifacts (original compositions, annotations and presentations) depending on the search options selected by the patron.
- (2) **Text objects** represent text-based information and may be of two content types: information artifacts (speeches) and patron-augmented artifacts (original compositions and annotations). Text objects allow editing if their underlying content types are editable. In Synchrony, published artifacts (those that are part of a PADL's collection) are not editable while unpublished patron-augmented artifacts are editable by those having the appropriate access rights. For editable text objects, text is typed directly on the objects themselves.
- (3) **Presentation objects** contain presentations authored by patrons and consist of sequences of artifacts each of which may contain a video segment of a speech, its corresponding textual transcript and/or an annotation/original composition displayed in synchrony. The contents of a presentation are displayed in tabular form, with each row corresponding to a

single sequence in the presentation while columns contain the types of artifacts in use within each sequence.

- (4) **Container objects** are workspaces within the main workspace and may contain query, document, presentation or even other container objects. While positioning may be used to divide a workspace, containers provide a more formal means of doing so, and are thus typically used to organize a workspace into various tasks and subtasks.

Synchrony shares common goals with digital library interfaces such as Artemis [23], DLITE [4] and NaviQue [6] in its support for an integrated, direct-manipulation environment for library-related tasks. In terms of design philosophy however, Synchrony is similar to VIKI [16] in that both systems derive their interfaces from the branch of hypertext/hypermedia systems known as *spatial hypertext* [14]. Spatial hypertext is characterized by the use of space in the creation and perception of structure. Whereas traditional hypertext systems employ explicit linking mechanisms to associate objects (e.g. unidirectional links between HTML documents) to create information structures, spatial hypertext systems describe associations among objects through space, that is, by geometrical relationships (e.g. proximity), visual characteristics (e.g. font size, color, shape), and recurrence (e.g. relative positioning of an object within a group of objects).

Studies have demonstrated the utility of such systems. For example, an analysis of Aquanet use (a collaborative hypertext tool) [13] found that for drawing relationships between objects, users preferred spatial positioning of objects to communicate structure rather than through predefined schemas (a collection of objects and relationship types). Further, in the Walden's Paths project [20], the spatial hypertext system VIKI has been used to some success in the authoring of paths - linear presentations of existing and new web pages.

A Scenario of Use

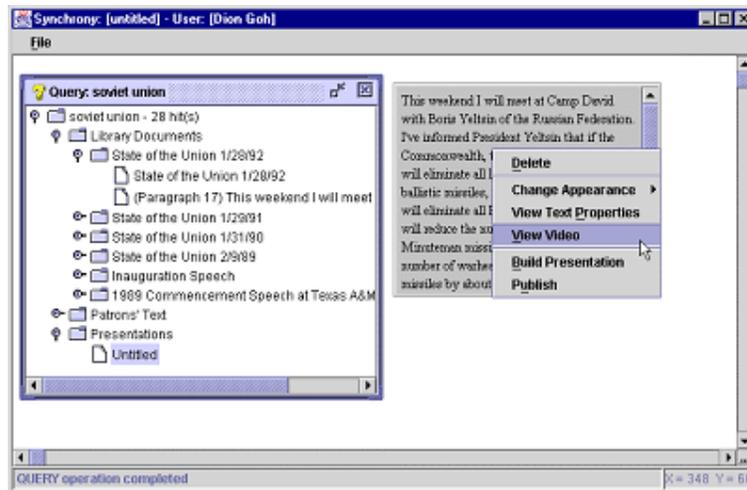
The following scenario illustrates how users may potentially use Synchrony and highlights the operation of the system. An educator is preparing a lesson about the Bush presidency and the Soviet Union for his history class and decides to prepare a multimedia presentation of speeches and press conferences given by George Bush on the subject from Synchrony's collection as a resource for his students.

After logging onto Synchrony, the educator is presented with an empty workspace. As this will be a new presentation, his first task is to locate relevant information by querying the PADL collection. He thus right-clicks at any point on the workspace to display a list of PADL services, and after selecting the query service, he enters the query (together with any options) in the dialog box that appears on the workspace.

When the query has been processed by Synchrony, a query object appears at the click location showing the results of the query. To view an artifact, the educator selects it from the query object, drags it onto the workspace and drops it at a desired location. Depending on the artifact type, a text object or a presentation object appears at the drop location. Figure 5 depicts the results of these actions.

After enough information has been retrieved, the educator's next step is to author the presentation. Synchrony simplifies the authoring process through a technique known as *incremental formalization* [21] which attempts to make a system understand informally represented information. This feature allows users to rapidly create presentations by first positioning document objects linearly within the workspace and then later specifying which objects to include into the presentation.

Figure 5. *Selecting and viewing artifacts*

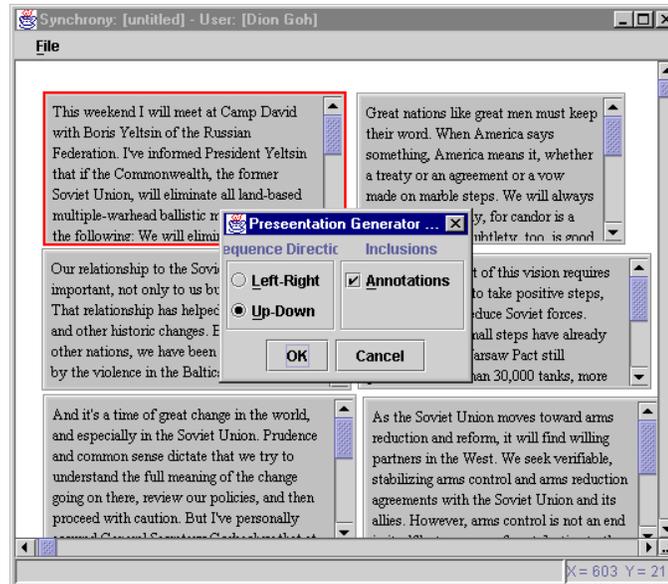


Returning to the scenario, the educator uses familiar drag-and-drop operations to assemble the text objects (which may include his annotations) to form two vertical adjacent list structures as depicted in Figure 6. He then invokes the presentation building service, causing Synchrony to automatically map these list structures to presentation sequences. In the current version, Synchrony assumes that the leftmost list contains video segments of speeches and their textual transcripts, while the adjacent list to its right is assumed to contain the corresponding annotation/original composition. In other words, sequences are mapped to the rows in the lists in a top-to-bottom manner while content is mapped to the columns. (Synchrony also supports a left-to-right mapping).

When Synchrony completes the mapping, a presentation object is displayed depicting the contents of the presentation in a tabular format (see Figure 4). In addition to providing a formalized representation of a presentation, the presentation object also allows patrons to

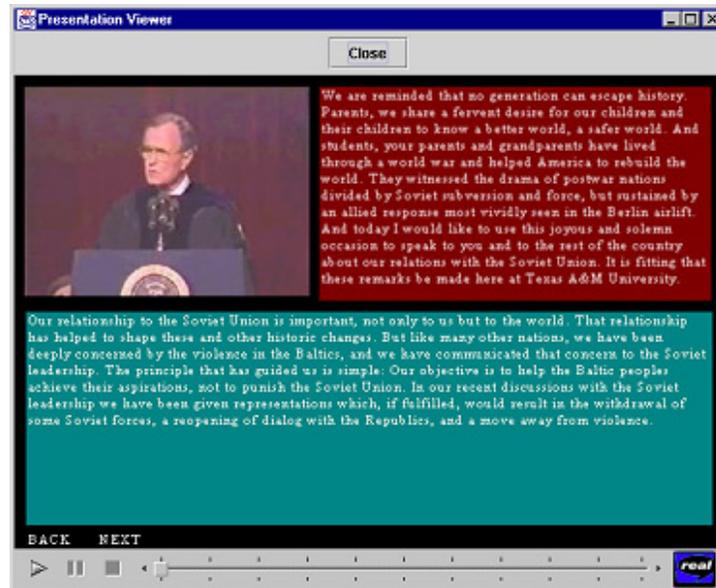
modify its contents. Patrons are able to add/remove sequences, add/move/remove content in any sequence, and shift the display order of sequences.

Figure 6. *Authoring a presentation*



When the educator is ready to view the presentation, he clicks a button on the presentation object. This causes Synchrony to assemble the sequences into a SMIL (Synchronized Multimedia Integration Language) [9] presentation and invoke a presentation viewer to display it. Figure 7 shows the presentation viewer. The viewer provides playback controls to allow patrons to play, pause, stop and seek. Each presentation sequence consists of three regions – a content region for displaying the text of a speech segment, a video region for presenting the associated video segment, and an annotation region for displaying associated annotations/original compositions.

Figure 7. *Viewing a presentation*



When the educator has finished authoring the presentation, he forwards it for review and possible publication by completing a form provided within Synchrony. Here, the educator provides the title of the presentation, a description, and an explanation of why the presentation should be published. Synchrony then uploads the completed form and presentation to the publication subsystem which stores them in a temporary holding area pending review. At this point, the educator's task is complete. He will later be notified through electronic mail about the outcome of his submission.

In the current version of Synchrony, all presentations submitted for publication are routed to a designated person (a librarian, a reviewer, an editor, etc.). Using a submissions viewing facility, reviewers may accept or reject submissions, or reroute submissions to other reviewers if necessary. When a submission is accepted or rejected, the author is informed via

electronic mail. Further, if the submission is accepted, it is indexed and incorporated into the PADL. To conclude the scenario, the author, upon receiving the acceptance message from the reviewer, informs his students about the presentation. The students may then begin their own Synchrony sessions, retrieve the presentation, and view and interact with it.

Note that for clarity, this scenario portrays the authoring process as a fixed sequence of tasks, that is, querying, organizing, viewing and publishing. In reality, Synchrony provides an environment in which these tasks may be performed in a fluid, iterative process. Patrons would move effortlessly among these activities depending upon the need at hand.

Implementation

Synchrony consists of a suite of client-server tools implemented mainly in Java together with two third-party applications. MG [24], a public domain full-text indexing and retrieval system, is used for the storage and retrieval of the textual content of speeches. Synchrony also utilizes RealNetworks' video server [19] for the delivery of streaming video and its implementation of the Java Media Framework [10] for the rendering of video segments and SMIL presentations.

Conclusion

Digital libraries must offer more than advanced collection maintenance and retrieval services since patrons often do not solely retrieve library artifacts for their own sake. In scholarly settings, patrons instead seek these artifacts to manipulate and integrate them to produce new artifacts. Traditionally, these activities have occurred mainly in physical media (predominantly paper), and as such, may be classified as paper-based scholarship. Digital

libraries however provide new service opportunities as well as an expanded set of informational data types, and when combined, have the ability to promote digital scholarship. Patrons are now able to perform their scholarly work electronically, working entirely with digital media.

We thus propose patron-augmented digital libraries as a class of digital libraries that support the digital scholarship of its patrons. A patron-augmented digital library (PADL) is one in which librarians and patrons both contribute to the evolution of the library's holdings. Librarians provide the seed material (information artifacts) to form an initial collection and maintain the collection while patrons augment the library with patron-augmented artifacts over the existing collection. To support this new role, a PADL departs from the traditional library model of service provision and supports authoring, structuring and publishing services in addition to search and retrieval.

Synchrony was developed to determine the feasibility of the PADL concept. The system provides access to a collection of digitized videos of speeches given by former President George Bush and their corresponding textual transcripts together with artifacts authored by patrons.

A pilot study was also conducted on Synchrony and results were encouraging [8]. In particular, the study found that attitudes toward PADLs in general were positive, provided that appropriate security and quality control (through publication policies) mechanisms were employed. However, because this study was performed in a laboratory setting with a small number of subjects, these results cannot be generalized. Consequently, in the next phase of Synchrony's development, we envision a larger-scale longitudinal study that will require

participants to use Synchrony to author and publish presentations for actual homework assignments over a semester. These results will be used to guide future work in the development of Synchrony and patron-augmented digital libraries.

References

1. Adam, N. and Yesha Y. Strategic directions in electronic commerce and digital libraries: towards a digital agora. *ACM Computing Surveys*, 28(4), 1996, 818-835.
2. Belkin, J., Oddy, R. and Brooks, H. ASK for information retrieval: Part I. Background and theory. *Journal of Documentation*, 38(2), 1982, 61-71.
3. Case, D. The collection and use of information by some American historians: a study of motives and methods. *Library Quarterly*, 61(1), 1991, 61-82.
4. Cousins, S., Paepcke, A., Winograd, T., Bier, E. and Pier, K. The digital library integrated task environment (DLITE). In: *Proceedings of the 2nd ACM International Conference on Digital Libraries, Philadelphia, 23-26 July 1997*. New York: ACM Press, 1997, 142-151.
5. Denning, P. and Rous, B. The ACM electronic publishing plan. *Communications of the ACM*, 38(4), 1995, 97-103.
6. Furnas, G. and Rauch, S. Considerations for information environments and the NaviQue workspace. In: *Proceedings of the Third ACM Conference on Digital Libraries, Pittsburgh, 24-26 June 1998*. New York: ACM Press, 1998, 79-88.
7. Goh, D. and Leggett, J. A spatial approach for the access, manipulation and publication of digital library artifacts. In: *Museums and the Web 1999, New Orleans, 11-14 March 1999*. Pittsburgh: Archives & Museum Informatics, 1999, 103-110.

8. Goh, D. and Leggett, J. Patron augmented digital libraries. In: *Proceedings of the Fifth ACM Conference on Digital Libraries, San Antonio, 2-7 June 2000*. New York: ACM Press, 2000, 153-163.
9. Hoschka, P. *Synchronized Multimedia Integration Language*. 1998. Available via <http://www.w3.org/TR/REC-smil/>.
10. Javasoft. *JMF Home Page*. 2001. Available via <http://www.javasoft.com/products/java-media/jmf/index.html>.
11. Levy, D. and Marshall, C. Going digital: a look at assumptions underlying digital libraries. *Communications of the ACM*, 38(4), 1995, 77-84.
12. Marchionini, G. *Information Seeking in Electronic Environments*. New York: Cambridge University Press, 1995.
13. Marshall, C. and Rogers, R. Two years before the mist: experiences with Aquanet. In: *Proceedings of the ACM Conference on Hypertext 1992, Milan, 30 November-4 December 1992*. New York: ACM Press, 1992, 53-62.
14. Marshall, C. and Shipman, F. Searching for the missing link: discovering implicit structure in spatial hypertext. In: *Proceedings of the Fifth ACM Conference on Hypertext, Seattle, 14-18 November 1993*. New York: ACM Press, 1993, 217-230.
15. Marshall, C., and Shipman, F. Spatial hypertext and the practice of information triage. In: *Proceedings of the Eighth ACM Conference on Hypertext, Southampton, 6-18 April, 1997*. New York: ACM Press, 1997, 124-133.
16. Marshall, C., Shipman, F. and Coombs, J. VIKI: spatial hypertext supporting emergent structure. In: *Proceedings of the 1994 ACM European Conference on Hypermedia Technology, Edinburgh, 19-23 September 1994*. New York: ACM Press, 1994, 13-23.

17. Nürnberg, P., Furuta, R., Leggett, J., Marshall, C. and Shipman, F. Digital libraries: issues and architectures. In: *Proceedings of the Second Annual Conference on the Theory and Practice of Digital Libraries, Austin, June 1995*. College Station: Hypermedia Research Laboratory, 1995, 147-153.
18. O'Hara, K., Smith, F., Newman, W. and Sellen, A. Student readers' use of library documents: implications for library technologies. In: *Conference Proceedings on Human Factors in Computing Systems CHI '98, Los Angeles, 18-23 April 1998*. New York: ACM Press, 1998, 233-240.
19. RealNetworks. *RealNetworks.com - The Home of Digital Media*. 2001. Available via <http://www.realnetworks.com/>.
20. Shipman, F., Furuta, S., Brenner, D., Chung, C. and Hsieh, H. Using paths in the classroom: Experiences and adaptations. In: *Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia, Pittsburgh, 20-24 June 1998*. New York: ACM Press, 1998, 267-276.
21. Shipman, F. and McCall, R. Supporting knowledge-based evolution with incremental formalization. In: *Conference Proceedings on Human factors in Computing Systems CHI 94, Boston, 24-28 April 1994*. New York: ACM Press, 1994, 285-291.
22. Stone, S. Humanities scholars: information needs and uses. *Journal of Documentation*, 38(4), 1982, 292-313.
23. Wallace, R., Soloway, E., Krajcik, J., Bos, N., Hoffman, J., Eccleston, H., Kiskis, D., Klann, E., Peters, G., Richardson, D. and Ronen O. ARTEMIS: learner-centered design of an information seeking environment for K-12 education. In: *Conference Proceedings on Human Factors in Computing Systems CHI '98, Los Angeles, 18-23 April 1998*. New York: ACM Press, 1998, 195-202.

24. Witten, I., Moffat, A. and Bell, T. *Managing Gigabytes*. Second Edition. San Francisco: Morgan Kaufmann Publishers, Inc., 1999.