Patron-Augmented Digital Libraries

Dion Goh

Department of Computer Science
Texas A&M University
College Station, TX 77843-3112, USA
E-mail: diong@cs.tamu.edu

John Leggett

Department of Computer Science
Texas A&M University
College Station, TX 77843-3112, USA
E-mail: leggett@cs.tamu.edu

ABSTRACT
Digital library research is mostly focused on the generation of large collections of multimedia resources and state-of-the-art tools for their indexing and retrieval. However, digital libraries should provide more than advanced collection maintenance and retrieval services since the ultimate goal of any (academic) library is to serve the scholarly needs of its users. This paper begins by presenting a case for digital scholarship in which patrons perform all scholarly work electronically. A proposal is then made for patron-augmented digital libraries (PADLs), a class of digital libraries that supports the digital scholarship of its patrons. Finally, a prototype PADL (called Synchrony) providing access to video segments and associated textual transcripts is described. Synchrony allows patrons to search the library 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. A study to evaluate the PADL concept and the usability of Synchrony is also discussed. The study revealed that participants were able to use Synchrony for the authoring and publishing of presentations and that attitudes toward PADLs were generally positive.

KEYWORDS: Patron-augmented digital libraries, digital scholarship, publishing, user interfaces

INTRODUCTION
Digital library research is mostly focused on the generation of large collections of multimedia resources and state-of-the-art tools for their indexing and retrieval. However, digital libraries should provide more than advanced collection maintenance and retrieval services since the ultimate goal of any (academic) library is to serve the scholarly needs of its users. Library artifacts, that is, the items that constitute a library's holdings, are often sought not as an end, but as a means to achieve a particular goal or set of goals.

Studies of library artifact use support this observation. For example, Levy and Marshall [13] 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, the analysts would annotate the documents as a means of interpreting them, produce new artifacts, and finally disseminate the new artifacts. In addition, analysts would commonly share documents and other interpretive 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.

Stone [24] 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. For a further example, see [18].

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. [19], 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. Using the data collected together with interviews of the students, 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.

Digital Scholarship
With the majority of their holdings in physical form, traditional (as opposed to digital) libraries typically promote *paper-based scholarship* in which physical media, predominately paper, play a major role in the scholarly use of library artifacts. For example, while patrons may use electronic databases to search for library 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 scholarly work as a cyclic set of transitions occurring in both the physical and digital domains.

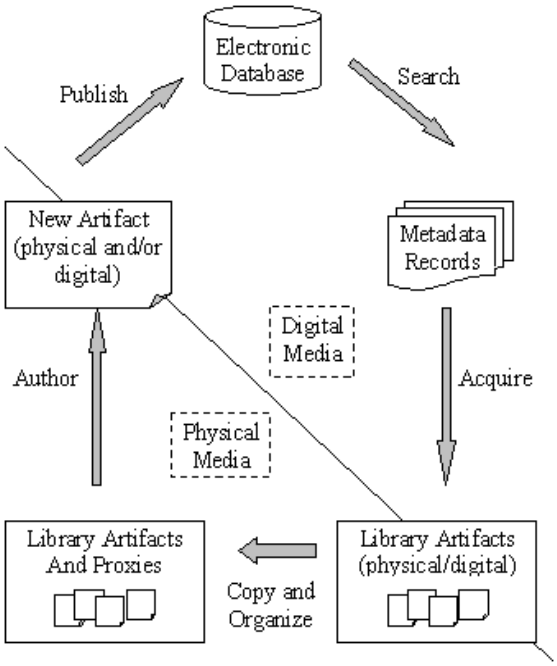

Figure 1. Paper-based scholarship in traditional libraries.

Artifacts (physical and digital) are located electronically via 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 being used [13]. 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.

Some types of multimedia information do not lend themselves to this model of scholarship. The canonical example may be video. Digital libraries however provide new service opportunities to patrons as well as an expanded set of informational data types [17], 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 (see Figure 2). 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.

Digital scholarship offers several advantages over paper-based scholarship.

(1) A single access point for library artifacts. Patrons are able to acquire all library artifacts in one location – at the computer (or PDA, etc.). 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 accessing and structuring information. Digital media provide new opportunities for patrons to interact with library artifacts not previously possible with paper-based artifacts. For example, data types such as video can now be used directly in the scholarly process. Further, 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 [6]. 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.

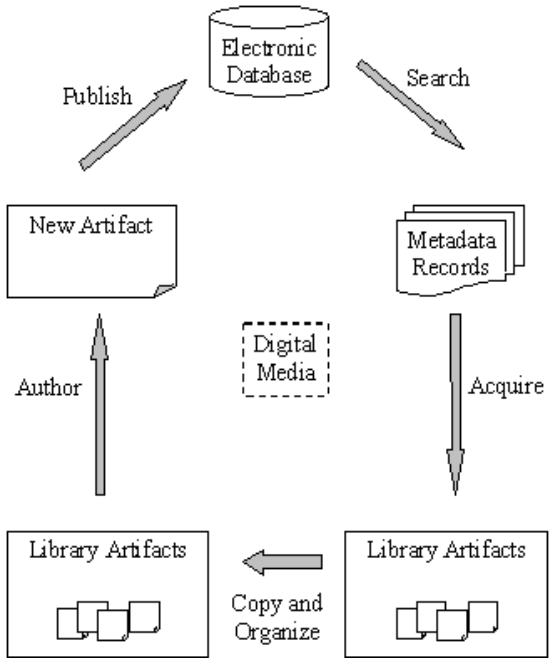

Figure 2. Digital scholarship in digital libraries.

The ASAP Model of Digital Scholarship

The discussion of digital scholarship suggests four phases that occur over library artifacts. A model of digital scholarship encompassing these four phases may be postulated, and for brevity will be known as the **ASAP** (**A**cquire, **S**tructure, **A**uthor, **P**ublish) **model** which is shown in Figure 3.

Artifacts are first acquired by patrons through searching and/or browsing of the digital library. Once acquired, they are structured (organized) to allow the patron to make better use of the information. Following this, patrons begin authoring new artifacts in the context of the ones already acquired. Finally, the newly created artifacts are disseminated – that is, published, using a variety of channels, both formal and informal. Once published, these artifacts are then available for use by other patrons.

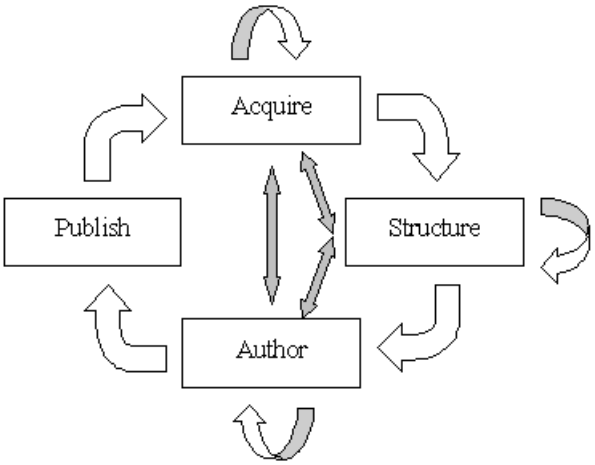

Figure 3. The ASAP model.

The progression between phases however may not be purely cyclic, in that patrons often move freely between phases, perhaps skipping a phase or returning to previous ones. This phenomenon arises because patrons often do not know the needs of a task at the outset, but must instead undergo an iterative clarification process to arrive at its successful completion [7].

Hence, added to this model are acyclic transitions between phases that need not occur in order, but instead depend on the needs of the patron at a particular point in time. In addition, studies of library use (described above) reveal the possibility that within each phase, an iterative process occurs as well, in effect creating mini-cycles within the main artifact use cycle (see for example [3, 15, 25]). Thus, the main **Acquire-Structure-Author-Publish** cycle describes the activities of patrons on library artifacts while within each phase, mini-cycles describe the cognitive processes that occur in order to complete that phase.

In this model the **Publish** phase is considered a terminal phase in that, once a new artifact has been published, it cannot be "unpublished" – returned to an earlier form. Consequently, the only way to modify a published artifact is to create a new version by beginning a new cycle or by removing it from the library.

Finally, an important aspect of the ASAP model not shown in Figure 3 is that phases may occur in parallel, allowing patrons to be in more than one phase at the same time [4, 5]. For example, in a digital library system that supports the concurrent execution of service requests, a patron may be acquiring artifacts through a background query process while structuring artifacts that are currently at hand.

Patron-Augmented Digital Libraries

While traditional library models (in which searching and browsing are the main services provided to patrons and scholarly work is mostly paper-based) have utility in the digital domain, this research postulates that, in many instances, an expanded model of digital library service would benefit patrons. In other words, *all* 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.

Assuming that patrons would indeed benefit from digital scholarship, the question becomes one of the types of services that should be supported. Returning to the ASAP model, a plausible starting point would include services for acquiring and organizing library artifacts, together with services for authoring and publishing new artifacts.

We propose *patron-augmented digital libraries* (PADLs) as a class of digital libraries that provide acquiring, structuring, authoring and publishing services to patrons. A patron-augmented digital library is one in which holdings are enhanced by the digital scholarship of users – both librarians and patrons contribute to the evolution of a library's holdings. In a PADL, librarians may initially populate the digital library with artifacts that meet the goals of the library. Patrons may then augment the holdings to meet specific needs through new artifacts such as documents, annotations or other organizational structures via the support services offered by the PADL.

The term "artifact" used in this paper 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 (see below). Examples include books and journals. *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 *information artifacts* created by patrons. Patron-augmented artifacts become reusable information artifacts through the publication process.

RELATED WORK

Due to space limitations, we briefly describe work from two related areas. Further discussion may be found in [9].

Publication Support in Digital Libraries

The University of California, Berkeley's digital library [27] focuses on diverse material related to the California Environment, and includes computer models, maps, photographs, videos and various datasets. In particular, the *multivalent document* paradigm [20] developed as part of this digital library project partitions a document into layers of distinct but related content. Functionality in each layer is provided by objects called *behaviors* that allow content to be manipulated. Among other features, layers and their behaviors allow users to attach additional content such as annotations to a document at any point in time.

The Global Digital Museum [26] is a federated museum and classroom that allows users to search and access multimedia information from various museums through a single user interface. Within a museum, two *textbook* types are defined. *Textbooks for teachers* are created by museum experts for teachers and explain the resources available for educational purposes. *Textbooks for students* are created by teachers as instructional resources for their students. In addition, users have access to *personal books* that allow them to maintain personal collections of museum materials as well as annotations.

Digital Library Interfaces

The Digital Library Integrated Task Environment (DLITE) [5] developed as part of the Stanford Digital Library project is one such example of a digital library interface. Central to DLITE is the concept of a *workcenter*, a place on the user interface that provides all the tools (or *components*) necessary for the completion of a specific task. An important aspect of DLITE is its use of direct manipulation. For example, a service may be invoked simply by selecting, dragging and dropping a request to its corresponding visual representation on the workcenter.

Similarly, the University of Michigan Digital Library's NaviQue system [8] provides an integrated workspace that allows users to issue queries in the same environment in which they are performing their tasks. In addition to querying, users are able to author text, as well as organize material on the NaviQue workspace using direct manipulation to clarify their conceptualizations and coordinate their search activities. NaviQue however differs from DLITE in that users are provided with an infinite 3-dimensional workspace provided by Pad++ [2] on which to perform their tasks together with panning and zooming capabilities.

Artemis [28] is a digital library interface designed to support K-12 student access and use of digital library resources. Artemis provides an integrated, direct manipulation environment called the *Persistent Workspace* for students to perform their information seeking tasks which may span multiple sessions. The workspace and its various components aim to create an inquiry-based learning environment for students by helping them to focus on the task at hand, conceptualize problems, generate search terms, and evaluate information retrieved.

Discussion

Although our research shares some similarities with the prior work reviewed here, several differences, notably in the area of information use, distinguish this work. While a few systems provide some level of support for authoring, their focus is on search and retrieval, with authoring as an aid to facilitate the search process. In contrast, authoring and publication with the goal of promoting digital scholarship are central activities in Synchrony. Support for this view is provided by the studies cited earlier which suggest that the search and retrieval of information artifacts are typically two of a set of tasks that patrons perform to achieve a larger goal - the authoring and publication of new artifacts.

Synchrony is also distinguished by its support for the authoring of composite structural entities in addition to simple/atomic entities such as text. It provides facilities for the authoring, publishing and display of presentations consisting of multiple synchronized media. These presentations are the digital "books" of the PADL.

SYNCHRONY

Synchrony is a prototype digital library system that implements the PADL concepts. Designed for the purposes of digital scholarship, Synchrony supports the ASAP model by allowing patrons to acquire artifacts from a digital library, structure these artifacts to meet the needs of their tasks, author new artifacts, and publish these new artifacts into the digital library. The notion of "scholarship" used in Synchrony is that put forth by the Stone [24] and O'Hara [19] studies described earlier.

Figure 4 shows the conceptual architecture underpinning Synchrony. The conceptual architecture provides an abstract view of the fundamental components of a PADL and their interactions. The manipulation subsystem is responsible for delivering the model of digital scholarship to the user and provides the user's view of the digital library. All user requests come from this subsystem and all results are returned to this subsystem. A security/privacy subsystem mediates all requests between the user and the PADL. This subsystem should provide for private, shared and published artifacts. The storage and retrieval subsystem is responsible for archival storage and retrieval of all artifacts, their various indexes and metadata. It is important for multiple media to be handled in a seamless fashion. The publication subsystem serves as an intermediary between patrons who wish to publish artifacts and the storage and

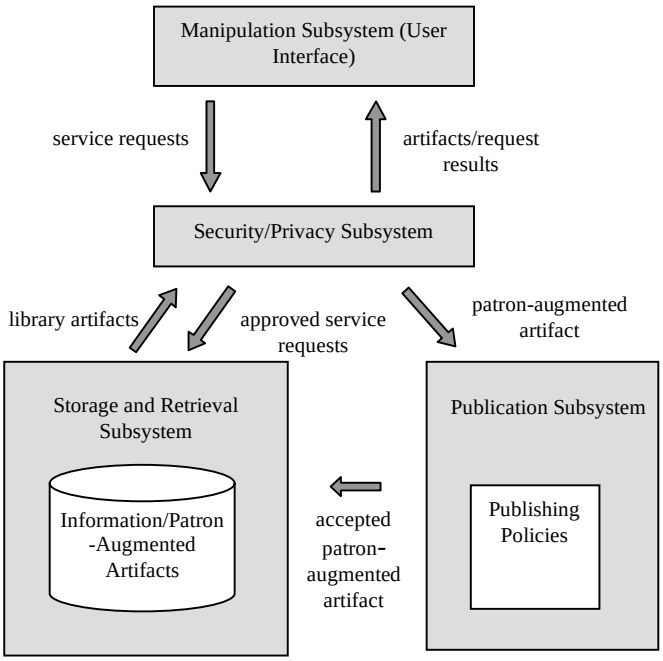

Fig. 4. Conceptual architecture of Synchrony.

retrieval subsystem. Publishing and reviewing policies should be separated from their implementation.

This conceptual architecture scales well from private PADL collections such as an individual's scholarly work, to journals such as ACM Transactions, to collections of professional society journals such as the ACM Digital Library, and finally to larger digital libraries consisting of multiple collections.

Synchrony's demonstration collection consists of digitized videos of speeches given by President George Bush 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 (books, papers, etc.). Annotations are also text-based documents publishable by patrons, but are designed to provide commentary and context to other artifacts. Presentations serve as structuring artifacts in the PADL. These composite entities [12] 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.

The Synchrony Client

The Synchrony client contains the user interface and is responsible for maintaining state information about a patron's current session, providing access to artifacts, rendering artifacts and executing service requests on behalf of a patron. Three user interface requirements guided the development of the Synchrony client:

(1) **Simplicity/Familiarity**. This requirement ensures that patrons with different levels of computer experience are able to utilize the resources of the digital library. This includes (1) leveraging familiar, existing technologies and tools so that patrons need not acquire and learn new tools and/or skills, (2) familiar interaction modalities such as direct manipulation (e.g. drag-and-drop), and (3) conceptually simple actions for querying, organizing, authoring and publishing.

(2) **Persistence**. Digital scholarship is an iterative process that occurs over an arbitrary amount of time. Consequently, the PADL must allow the results of patrons' tasks to be saved and restored as necessary.

(3) **Integrated work environment**. The Synchrony client supports seamless switching among various tasks that patrons perform. This integration not only reduces the number of tools required but also allows for a uniform method for accessing the resources of the PADL.

The 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 5 and consists of two major entities: the workspace and library objects.

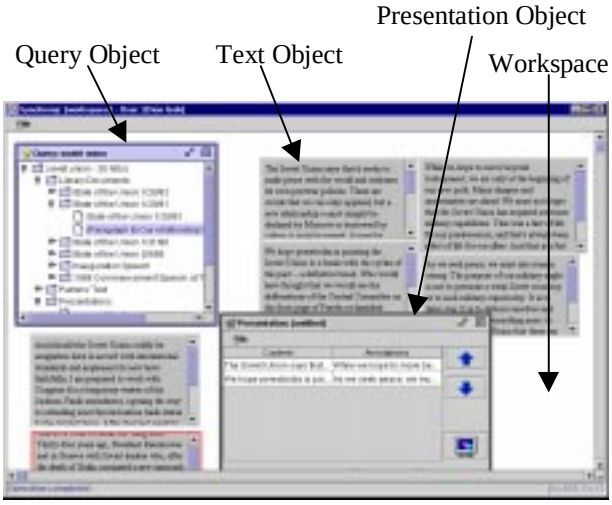

Figure 5. Synchrony's interface.

The background represents the *workspace* 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. The direct manipulation paradigm allows these objects to 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, because the size of the workspace is larger than the screen (essentially infinite in the X- and Y-axes), both 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 three basic categories: queries, texts and presentations.

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. Results of a search are

presented as a three-level hierarchy as shown in Figure 5. The first level of the hierarchy contains information about the query itself (the actual query issued and the number of hits). The second level divides the results into three categories of artifacts: information artifacts (speeches), text-based patron-augmented artifacts (annotations and original compositions) and presentations. Within the information artifact category, a third level further divides each speech into paragraphs matching the query. Queries may be compared through collateral display of the resulting query objects.

Text objects represent text-based information and may be of two content types: information artifacts (speeches) and patron-augmented artifacts (original compositions and annotations). Although content types are varied, these objects ultimately serve a similar purpose – the display of text-based information. Text objects are thus designed to look similar and provide similar functionality to reduce the number of interface objects users have to deal with. Patrons may, however, alter the appearance of text objects through size and color if differentiation is desired. Figure 5 shows six examples of text objects. There is no size limitation on text objects.

In addition to displaying information, 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.

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 synchronously. Figure 5 depicts a presentation object which represents presentations in tabular form. Each row corresponds to a single sequence in the presentation while columns contain the types of artifacts in use within each sequence. Two columns are available for each sequence: one for an information artifact (which also contains a reference to its corresponding video segment) and one for a text-based patron-augmented artifact (an original composition or an annotation).

The Synchrony Server

The Synchrony server is responsible for accepting and processing patrons' requests via the Synchrony client and consists of two tiers. Tier 1 functions as a PADL session manager through which all Synchrony clients must connect in order to gain access to the resources of the PADL. The session manager maintains a set of *PADL services* that are each responsible for processing a specific client request. When a service request is received from the Synchrony client, the session manager invokes the appropriate PADL service for execution.

In Synchrony's model of service provision, a PADL service may either be directly responsible for processing a client request, or may act as an intermediary between a third-party application that executes that request and the Synchrony client. In the latter role, the third-party application represents the second tier of the Synchrony server, and the PADL service communicates the request to it and relays the results of the execution back to the Synchrony client. In the current implementation, the PADL services supported encompass authentication, querying, artifact retrieval, presentation creation, workspace saving/restoration, and artifact publishing. Two Tier 2 applications are also employed. For storage and retrieval of text-based information and patron-augmented artifacts, MG [29], a public domain full-text indexing, retrieval and compression system is used. Streaming video segments of speeches, on the other hand, are delivered using the RealVideo server [22].

A Scenario of Use

The following scenario illustrates one of the potential uses of Synchrony and highlights the operation of the system. An educator is preparing a lesson about the Bush presidency and the Soviet Union for his political science class. As a resource for his students, he prepares a presentation consisting of selected video segments and textual transcripts from speeches and press conferences given by George Bush on the subject.

Upon launching the Synchrony client and successfully logging in, the educator is presented with an empty workspace. Note that in Figure 5, most of Synchrony's interface is occupied by the workspace with only a single menu item (the "File" menu) on the menu bar. Synchrony was designed to have all operations accessible through context sensitive menus via the right mouse button.

The educator's first task is to locate relevant information and, in the current version of Synchrony, this is accomplished by querying the PADL collection. He thus issues a query by right-clicking at any point on the workspace, selects the query service from the ensuing popup menu and enters the query (together with any options) in the dialog box that appears on the workspace. Synchrony supports four query modes: Boolean, ranked, author search and ID search. The use of the latter query mode will be discussed later in this section.

When the query has been processed, 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 read-only presentation object appears at the drop location. Figure 6 depicts the results of these actions.

When the educator determines that enough relevant information has been obtained, he begins authoring the presentation. In Synchrony, authoring is accomplished by positioning text objects to form two adjacent list structures [16] on the workspace. Returning to the scenario, the

Figure 6. Querying and viewing artifacts.

educator organizes the text objects to form two vertical adjacent list structures as depicted in Figure 7. To generate the presentation, he selects the starting sequence by clicking on the left- and topmost text object and, upon invoking the presentation building service, a dialog box appears (also shown in Figure 7). Here, the educator is provided with options for including annotations/original compositions as well as for confirming the direction in which the sequences are to be mapped.

Synchrony then maps these list structures to presentation sequences such that the leftmost list is assumed to contain 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 5). 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 7. 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) [10] presentation and invoke a presentation viewer to display it. Figure 8 depicts 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 8. The presentation viewer.

As patrons' ASAP tasks may take an indeterminate amount of time to complete, Synchrony supports task persistence through the saving and subsequent opening of workspaces. Assuming that the educator wishes to terminate his current Synchrony session and resume his work at a later time, he selects the save workspace service and enters a workspace name for his current session. Synchrony then saves all the library objects in the workspace including size, color and position information. To resume his work, the educator selects the open workspace service and enters the previously saved workspace name. Synchrony will then restore all the saved library objects in that workspace.

When the educator has finished authoring the presentation, he forwards it for review and possible publication by completing a form. Here, the educator provides the title of the presentation, a description, and an explanation of why the presentation should be published. This information, among other criteria, will be used to decide if the presentation will be accepted for publication. At this point, the educator's task is complete. He will later be notified through electronic mail about the outcome of his submission.

Note that while this scenario only discusses the publication of presentations, Synchrony also supports the publication of text-based patron-augmented artifacts (annotations/original compositions). In all cases however, the process of submitting an artifact for review is the same.

All presentations submitted for publication are routed to a designated responsible person (a reviewer, an editor, a program committee chair, etc.). Continuing with the

scenario, this reviewer begins a Synchrony session to check for new submissions. After logging in, the reviewer invokes the *submissions viewer* shown in Figure 9.

Figure 9. The submissions viewer.

The submissions viewer provides information about the artifacts submitted and their respective authors. Artifact information includes metadata as well as its content. Author information includes the name of the author, his/her e-mail address, the department/organization to which the author belongs and so on. The rationale for providing the latter is to assist the reviewer in establishing the identity and credentials of the author.

In addition, the submissions viewer allows the reviewer to accept or reject any submission. A text area is provided for reviewer comments on the submission and this will be sent to the author together with the publication status (acceptance or rejection).

Cooperative reviewing is supported through the rerouting of submitted artifacts. If a reviewer wishes, he/she may reroute the submission to other reviewers by selecting from a list of available reviewers. In this case, any comments provided will be sent to the selected reviewers when they view the submission. Using this facility, reviewers may thus solicit assistance/comments from other reviewers in order to make informed decisions regarding a submission.

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. The author is also provided with the artifact identifier which may be used during ID searches. 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. Of course, in this scenario, the reviewer may be the educator himself with publication accepted to the class PADL.

Note that for clarity, this scenario portrays digital scholarship as a fixed sequence of tasks – querying, organizing, authoring and publishing. In reality however, Synchrony provides an environment in which these tasks may be performed in a fluid, iterative process as prescribed by the ASAP model. Patrons move effortlessly among these activities depending upon the need at hand. Thus, this process of querying, organizing and authoring is repeated as many times as necessary, in any order, until the educator has all the material required. For example, if, while organizing text objects for the presentation, the educator realizes that more information is necessary, he can simply move to an unused portion of the workspace, issue further queries, create new annotations/original compositions, and incorporate these into the presentation by dragging-and-dropping them at the desired locations.

Implementation

Synchrony is implemented as a collection of Java 2 [11] classes with the exception of the MG and the RealVideo (Tier 2 server) applications. In addition to the standard Java 2 classes, the user interface of the Synchrony client is implemented using the Swing component set.

The Synchrony client relies on RealNetworks' implementation of the Java Media Framework [21] for the rendering of video segments and SMIL presentations and, therefore, currently runs only on Windows 95/98/NT platforms. The Synchrony server on the other hand is written using only the standard Java 2 classes.

PILOT STUDY

A pilot usability study was conducted with the aim of identifying future research opportunities in Synchrony as well as in the area of PADLs. The goals of this study were twofold:

(1) Elicit feedback from participants using Synchrony to identify any positive and negative features in its use as a tool for digital scholarship;
(2) Determine the participants' attitudes towards patron-augmented digital libraries by soliciting opinions regarding the use of the services of a PADL for digital scholarship.

These goals were accomplished using an experimental task and a post-experimental questionnaire.

Methodology

Participants. Five graduate students (and one professor), who were content experts in the area being studied, were recruited from the Speech Communications Department at Texas A&M University. All participants had at least five years experience using a personal computer and used computers nearly every day in their scholarly tasks.

Training. Prior to the commencement of the experimental task, participants were given a one-hour training session. Participants were taught to navigate the workspace, manipulate and organize library objects, issue queries, and author and publish presentations, annotations and original compositions. Following this, the participants were given a practice task that exercised the necessary features required to complete the actual experimental task.

Experimental Task. The experimental task was modeled after actual homework assignments. Participants were given 1.5 hours to complete the assignment which involved issuing queries, determining and retrieving needed speech paragraphs (possibly from different speeches) from the PADL, authoring original compositions/annotations, and integrating and organizing the relevant artifacts on the workspace to form presentations. Participants were also asked to publish their presentations and save their workspaces.

The speeches used were the four State of the Union addresses given by President George Bush between 1989 and 1992. The total number of paragraphs contained within the textual transcripts was 256. These four speeches total approximately six hours of video. Each task required a participant to critique a particular theme (determined by the professor) found across the four speeches. The areas covered were: (1) an examination of the policies involved in the president's war on drugs, (2) the evaluation of the president's role as the "Environmental President", (3) the evaluation of the president's role as the "Education President", (4) a commentary of the rhetoric on the economy during the recessionary years in the early 1990's, and (5) an analysis of the president's concept of a New World Order.

Post-Experimental Questionnaire. Upon submission of their presentations, participants were issued the post-experimental questionnaire and given 30 minutes for its completion. The post-experimental questionnaire consisted of 30 items rated on a 5-point Likert-type scale that ranges from Strong Disagreement (1) to Strong Agreement (5). In addition, participants also had the opportunity to express their opinions on 12 open-ended items. The questionnaire covered areas such as the participants' level of computer experience, the usability of the various features found in Synchrony (the design of the user interface, querying, organizing, authoring and publishing) and feedback for future improvements. In addition, one section was designed to elicit opinions about PADLs and their use and included questions such as whether participants liked the idea of publishing into a digital library, whether they would use publications by other patrons, and so on.

Results and Discussion

Results from the study indicated that aside from inexperience in using Synchrony due to time constraints and difficulties typically associated with prototype software, participants were generally able to use the system for the authoring and publication of presentations. Several participants, commented that Synchrony was easy to learn and use. The fact that Synchrony used direct manipulation as the mode of interaction probably contributed to its success. Further, access to many of Synchrony's features (such as querying and publication) were a mouse-click away, and this likely played a significant role in Synchrony's ease of use as well.

One participant, however, provided low ratings on Synchrony's method of authoring presentations and this may be attributed to the novelty of this authoring approach. Commercially available presentation authoring software commonly uses timelines (e.g. Director [14]) or flowcharts (e.g. IconAuthor [1]) and, as such, this participant might be unfamiliar with the "authoring through organizing" approach. On the other hand, the relatively higher scores given by the other participants may be attributed to the small number of presentation sequences (average of nine). It is possible that, for a larger number of sequences, Synchrony's approach might prove deficient. Clearly, future studies should focus on these issues. One study might be longitudinal in nature and require participants to use Synchrony to author presentations for actual homework assignments over a semester or more. Another study might compare authoring presentations using Synchrony against other software packages employing different authoring interfaces.

Several participants expressed reservations about viewing other patrons' publications due to issues of quality. Specifically, these participants would peruse a particular publication only if the author was an authority on the topic, and feared that a lack of control over which artifacts are published would only serve to diminish the value of the digital library. These legitimate concerns underscore the importance of establishing appropriate publication policies in a patron-augmented digital library. However, not only must such policies exist, they must be communicated to patrons to enable them to understand that the goals of the digital library and the quality of its holdings are not compromised through the publication of patrons' artifacts. In addition, all patron-augmented artifacts should contain metadata that describe the credentials of the author to help patrons make informed decisions about the quality of these artifacts. Information such as the name of the author, the organization to which he/she is affiliated, his/her job title and description, and a short biography of the author are relevant items that should be included in a patron-augmented artifact's metadata.

While publication policies ensure the enhancement and quality of a collection, mechanisms must also be put into place to encourage the participation of patrons. Two measures are envisioned. The first is the implementation of access controls that prevent illegal modifications and/or plagiarism of patron-augmented artifacts. These controls would promote a sense of ownership and encourage active involvement in the growth of the digital library. A second measure is to establish incentives that reward patrons for publications. These might include monetary benefits or system-wide recognition of exceptional contributions.

A major limitation of this study is its small number of subjects (five). This prevents generalizations from being made based on the findings. Nevertheless, the study has revealed important issues that may be used to guide future work in the development of Synchrony and patron-augmented digital libraries.

FUTURE WORK

An important aspect to be examined is the deployment of more facilities for navigating and visualizing the workspace. This is especially important because, given limited screen sizes, only a small number of library objects are visible to the patron. One enhancement to be considered is the integration of standard zooming capabilities into the workspace as is done in NaviQue. A more novel approach would be the development of nonlinear visualizations (e.g. multiple focus fisheye transformations) such as those used by VIKI [23].

Another area of work involves the development of a video annotation facility, including voice and video annotations. Currently, patrons may create annotations to speeches and original compositions but are unable to associate them with video segments directly. A video annotation feature would allow patrons to develop a more diverse variety of artifacts for the digital library. A related area of work is to allow the specification of segmentation granularity of speeches (currently preset at the paragraph level). Several approaches to relaxing this restriction are possible. The first is to return entire documents as search results and permit patrons to devise their own segmentation granularities, but this would require associating each word with its corresponding video segment to accommodate a variety of granularities. A second, more restrictive approach would be to provide preset granularity levels (e.g. sentence, paragraph or document) for patrons to select among. Another possibility is to retain the current scheme and use the video annotation facility for patrons to specify their desired segmentations, in effect creating new patron-augmented artifacts.

Increased support for presentation functionality is also a candidate for future work. Presentations in Synchrony are linear with no support for linking to other artifacts. A linking facility is envisioned that would allow patrons to create both textual and video start- and end-point links. Customization of presentation layouts is a related consideration. Presentations currently consist of three regions per presentation sequence – a content region, a video region and an annotation region. A useful feature to incorporate would allow patrons to determine the layout of each presentation sequence. Parameters that might be customizable would include position, size, color, graphical elements and so on. In short, such a feature should provide full access to SMIL layout commands.

Finally, it would be an interesting exercise to use Synchrony in other collections beyond political speeches. A possible candidate would be a digital music library consisting of lyrics and music. Synchrony would then be used to create presentations that combine various segments of different pieces of music or even to create new music. Several questions would have to be answered. For example, what constitutes a "paragraph"? Is the current text-based visualization for lyrics sufficient or is a score-based visualization necessary? Employing Synchrony in a variety of different collection types would test the broad applicability of its user interface and tools, and in the process inform research in the fields of digital libraries and user interfaces.

ACKNOWLEDGEMENTS

This research was supported in part by the Texas Advanced Research Program under Grant No. 010366-0041-1999.

REFERENCES AND CITATIONS

1. Asymetrix. 1999. *Asymetrix IconAuthor*. Available via http://www.asymetrix.com/products/iconauthor/.

2. Bederson, B., and Hollan, J. 1994. Pad++: A zooming graphical interface for exploring alternate interface physics. In *Proceedings of the ACM Symposium on User Interface Software and Technology* (UIST '94, Marina del Rey, CA, Nov.). ACM Press, New York, NY, 17-26.

3. Belkin, J., Oddy, R., and Brooks, H. 1982. ASK for information retrieval: Part I. Background and theory. *Journal of Documentation* 38, 2 (June), 61-71.

4. Case, D. 1991. The collection and use of information by some American historians: A study of motives and methods. *Library Quarterly* 61, 1 (Jan.), 61-82.

5. Cousins, S., Paepcke, A., Winograd, T., Bier, E., and Pier, K. 1997. The digital library integrated task environment (DLITE). In *Proceedings of the 2nd ACM International Conference on Digital Libraries* (DL '97, Philadelphia, PA, July). ACM Press, New York, NY, 142-151.

6. Denning, P., and Rous, B. 1995. The ACM electronic publishing plan. *Communications of the ACM 38*, 4 (Apr.), 97-103.

7. Engelbart, D. 1963. A conceptual framework for the augmentation of man's intellect. In *Vistas in Information Handling Volume 1*, P. Howerton, Ed. Spartan Books, Washington, DC, 1-29.

8. Furnas, G., and Rauch, S. 1998. Considerations for information environments and the NaviQue workspace. In *Proceedings of the Third ACM Conference on Digital Libraries* (DL '98, Pittsburgh, PA, June). ACM Press, New York, NY, 79-88.

9. Goh, D. 1999. *Patron-augmented digital libraries*. Dissertation, Texas A&M University, College Station, TX.

10. Hoschka, P. 1998. *Synchronized multimedia integration language*. Available http://www.w3.org/TR/REC-smil/.

11. Javasoft. 1999. *Java 2 platform, standard edition*. Available via http://www.javasoft.com/jdk/.

12. Leggett, J., and Schnase, J. 1994. Viewing Dexter with open eyes. *Communications of the ACM 37*, 2 (Feb.), 76-86.

13. Levy, D., and Marshall, C. 1995. Going digital: a look at assumptions underlying digital libraries. *Communications of the ACM 38*, 4 (Apr.), 77-84.

14. Macromedia. 1999. *Macromedia Director*. Available via http://www.macromedia.com/software/director/.

15. Marshall, C., and Shipman, F. 1997. Spatial hypertext and the practice of information triage. In *Proceedings of the Eighth ACM Conference on Hypertext* (Hypertext '97, Southampton, UK, Apr.). ACM Press, New York, NY, 124-133.

16. Marshall, C., Shipman, F., and Coombs, J. 1994. VIKI: Spatial hypertext supporting emergent structure. In *Proceedings of the 1994 ACM European Conference on Hypermedia Technology* (ECHT '94, Edinburgh, UK, Sept.). ACM Press, New York, NY, 13-23.

17. Nürnberg, P., Furuta, R., Leggett, J., Marshall, C., and Shipman, F. 1995. Digital libraries: Issues and architectures. In *Proceedings of the Second Annual Conference on the Theory and Practice of Digital Libraries* (Digital Libraries '95, Austin, TX, June). Hypermedia Research Laboratory, College Station, TX, 147-153.

18. O'Day, V., and Jeffries, R. 1993. Orienteering in an information landscape: How information seekers get from here to there. In *Conference Proceedings on Human Factors in Computing Systems* (CHI '93, Amsterdam, The Netherlands, Apr.). ACM Press, New York, NY, 438-445.

19. O'Hara, K., Smith, F., Newman, W., and Sellen, A. 1998. Student readers' use of library documents: Implications for library technologies. In *Conference Proceedings on Human Factors in Computing Systems* (CHI '98, Los Angeles, CA, Apr.). ACM Press, New York, NY, 233-240.

20. Phelps, T., and Wilensky, R. 1996. Toward active, extensible, networked documents: multivalent architecture and applications. In *Proceedings of the 1st ACM International Conference on Digital Libraries* (DL '96, Bethesda, MD, Mar.). ACM Press, New York, NY, 100-108.

21. RealNetworks. 1998. *RealSystem G2 SDK developer's guide*. Available via http://www.real.com/devzone/library/creating/rmsdk/index.html.

22. RealNetworks. 1999. *RealNetworks, the home of RealAudio, RealVideo and RealFlash*. Available via http://www.real.com/.

23. Shipman, F., Marshall, C., and LeMere, M. 1999. Beyond location: Hypertext workspaces and non-linear views. In *Proceedings of the Tenth ACM Conference on Hypertext and Hypermedia* (Hypertext '99, Darmstdt, Germany, Feb.). ACM Press, New York, NY, 121-130.

24. Stone, S. 1982. Humanities scholars: Information needs and uses. *Journal of Documentation 38*, 4 (Dec.), 292-313.

25. Streitz, N., Haake, J., Hannemann, J., Lemke, A., Schuler, W., Schutt, H., and Thuring, M. 1992. SEPIA: A cooperative hypermedia authoring environment. In *Proceedings of the ACM Conference on Hypertext* (ECHT '92, Milan, Italy, Nov.). ACM Press, New York, NY, 11-22.

26. Takahashi, J., Kushida, T., Hong, J., Sugita, S., Kurita, Y., Rieger, R., Martin, W., Gay, G., Reeve, J., and Loverance, R. 1998. Global digital museum: Multimedia information access and creation on the Internet. In *Proceedings of the Third ACM Conference on Digital Libraries* (DL '98, Pittsburgh, PA, June). ACM Press, New York, NY, 244-253.

27. UCB. 1999. *UC Berkeley digital library project*. Available via http://elib.cs.berkeley.edu/.

28. Wallace, R., Soloway, E., Krajcik, J., Bos, N., Hoffman, J., Eccleston, H., Kiskis, D., Klann, E., Peters, G., Richardson, D., and Ronen O. 1998. 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, CA, Apr.). ACM Press, New York, NY, 195-202.

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