Journal of Universal Computer Science, vol. 2, no. 9 (1996), 610-622
submitted: 19/9/96, accepted: 19/9/96, appeared: 28/9/96

Designing Digital Libraries for the Hyperliterate Age

Peter J. Nürnberg
(Center for the Study of Digital Libraries, Texas A&M University, USA
pnuern@csdl.tamu.edu)

Erich R. Schneider
(Center for the Study of Digital Libraries, Texas A&M University, USA
erich@csdl.tamu.edu)

John J. Leggett
(Center for the Study of Digital Libraries, Texas A&M University, USA
leggett@csdl.tamu.edu)

Abstract: Many Web and Internet technologies have traditionally been used to serve information across machines and among people. Recently, there has been a great deal of interest shown in using these information services to support *digital libraries*. Research in digital libraries is an interdisciplinary effort that must synthesize existing results from highly disparate fields. This paper examines two such contributing fields — information systems and orality-literacy studies — and applies them to a particular digital library domain, botanical taxonomic work. In trying to build digital libraries for botanical taxonomists, we show how two widely differing fields can each provide part of a solution neither can provide alone.
Categories: H.4.m, K.4.m

0 Description of the Paper

Many Web and Internet technologies have traditionally been used to serve information across machines and among people. Recently, there has been a great deal of interest shown in using these information services to support *digital libraries*. Research in digital libraries, as any interdisciplinary endeavor, is confounded by the fact that one must consider and synthesize from several fields, including those that are perhaps unfamiliar. On the one hand, we do not want to become experts in many fields — generally, things are complicated enough on their own without introducing the complexities, terminologies, and traditions of other fields. On the other hand, we must learn to understand some unfamiliar ideas if we ever hope to undertake interdisciplinary research[*] [Nürnberg et al. 1996b].

This paper seeks to tie together two very different fields — information systems and orality-literacy studies — that each have something to offer the digital library designer. As such, it will be the case that almost everybody will find some part of this paper unfamiliar. The authors have chosen an unconventional format for presenting this material. The paper contains two threads for the two fields it draws upon. Some sections belong in only one thread, while others belong in both. The paper can be read

---

[*] This paper is an expanded version of "Designing Digital Libraries for Post-literate Patrons," presented at the WebNet '96 Conference.

in different ways, but most people will find it easiest to read the thread with which they are more familiar first, in order to contextualize the material, and then delve into the other thread. Figure 1 below illustrates the organization of the paper. (Note: the information systems thread should be read from the left column and the orality-literacy thread from the right.)

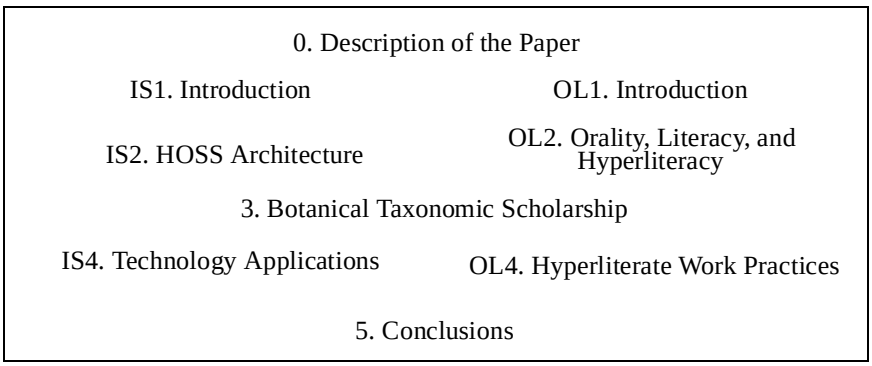

Figure 1: Organization of the Paper.

IS1 Introduction

For many reasons, archaic work practices of varying "inappropriateness" to modern scholarship linger on despite their known flaws. In information-intensive fields, considerable support for the development of new work practices can be provided by digital libraries and the technologies underlying them. In particular, advanced, distributed, computationally-oriented hypermedia systems, with their capability to support more fluid information structures, have often been proposed for use in fields where the mutable cognitive artifacts that scholars employ are known to be poorly reflected in the static artifacts produced by pre-electronic work practices for pre-electronic distribution methods.

In this thread, we discuss such a hypermedia system, named HOSS, and then relate it to a particular digital libraries problem — building digital libraries of botanical taxonomic information. We

OL1 Introduction

For many reasons, archaic work practices of varying "inappropriateness" to modern scholarship linger on despite their known flaws. In information-intensive fields, the derivation of possible new work practices can be suggested by differentiating those aspects of current practice that are archetypic to the problem addressed from those artifactual to the technologies currently employed. In particular, orality-literacy studies are here proposed for this purpose in fields where the mutable cognitive artifacts that scholars employ are known to be poorly reflected in the static artifacts produced by pre-electronic work practices for pre-electronic distribution methods.

In this thread, we propose an extension to orality-literacy studies, namely hyperliteracy, and then relate it to a particular digital libraries problem — building digital libraries of botanical

show how various aspects of the HOSS architecture can be applied to some of the archaic work practices in this field, yielding support for new, more natural practices.

IS2 HOSS Architecture

HOSS is a computationally-oriented hypermedia system [Nürnberg et al. 1996a]. It consists of a hyperbase layer, a structure processing layer, a metadata manager layer, and an application layer (see figure 2). Each of these will be briefly described below.

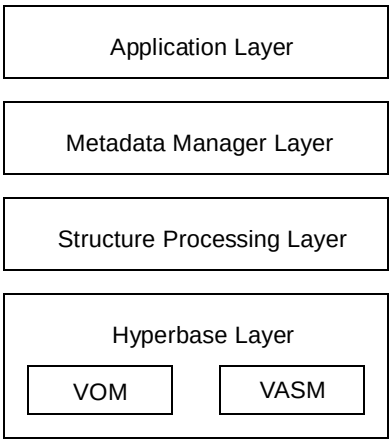

Figure 2: HOSS Architecture.

The main difference between HOSS and other hypermedia systems is that HOSS is an entire operating environment. It provides file system, memory management, and scheduling features. Other operating system functionality is provided by a SunOS 5.4 kernel. HOSS is best thought of as a hypermedia-aware operating system. An immediate result of this is that HOSS, as any operating system, admits an open set of application processes. Furthermore, just as all applications in a real-time operat-

taxonomic information. We show how various aspects of the theory of hyperliteracy can be applied to some of the archaic work practices in this field, yielding the characterization of possible new, more natural practices.

OL2 Orality, Literacy, and Hyperliteracy

Since the 1960s an interdisciplinary research area within the humanities known as orality-literacy studies has existed, concerned with differences in the modes of thought and expression exhibited by individuals in cultural situations which exhibit *primary orality* (where writing is not used as an adjunct to thought and memory) and those exhibiting *pervasive literacy* (where it has become indispensable for thought and memory).

OL2.1 Orality and Literacy

A seminal work in orality-literacy studies is *Preface to Plato* by classicist Eric Havelock [1963], whose starting point is Plato's attack on poetry in the *Republic* [Waterfield 1993]. Plato's proposal that poetry be banned from his ideal state, because it degraded the intellect, is found odd by many modern students of Plato, who would argue that encountering great poetry is an uplifting experience. Havelock sets out to examine what this apparent oddity in the philosopher's thought implies about the cultural situation of Plato's Greece.

Havelock contends the extensive ground of common knowledge and worldviews required by classical Greek culture was encoded in the great poems of the time, most notably Homer's epics. While today they are studied for their aesthetic value, to the ancient Greeks, these were a "tribal encyclopedia" of

ing system may take advantage of real-time awareness on the part of the operating system, all HOSS applications have immediate access to hypermedia functionality. The functionality of the hyperbase and (open) structure processing layer is available to all HOSS processes.

IS2.1 Hyperbase Layer

A HOSS hyperbase is a process with three threads: a Versioned Object Manager (VOM) and a Versioned Structure Manager (VASM). The VOM acts as a client of some Storage Manager that exists outside of the hyperbase. The VOM serves simple object and composite object abstractions. The VASM serves structural entities that can be specialized by structure processors (see below). These structural entities may connect data and other structural entities.

The Storage Manager may be an arbitrary database system. The current HOSS prototype uses Illustra [IIT 1995] as its Storage Manager. Previous versions have used other relational database systems as well as semantic database systems [Schnase et al. 1993a; Schnase et al. 1993b]. The VOM normalizes the interface to the arbitrary Storage Manager level and implements versioning.

A HOSS hyperbase is conceptually similar to other hyperbase systems [Leggett and Schnase 1994; Schnase 1993b; Shackelford et al. 1993; Schütt and Streitz 1990; Wiil 1993].

IS2.2 Structure Processing Layer

HOSS allows an open set of structure processors called *Sprocs*. Sprocs are heavy-weight processes that contain threads to load structure and execute behaviors. Sprocs are handled differently than other processes at the operat-

cultural ways and norms, much more accessible to the non-literate masses than written records. Poetry was also well-suited to the problems of information storage in a non-literate culture, namely retention in living memory and content-preserving transmission [Havelock 1963].

	Orality	Literacy
Ideas as... [Havelock 1963]	properties of concrete situations	abstract and eternal "things in themselves"
Socially relevant truths as... [Ong 1982]	mutable objects	fixed objects
Language use as... [Ong 1982]	requiring consideration of situation	manipulation of abstract placeholders

Table 1: Examples of Differences Between Orality and Literacy.

Poetic encoding had three advantages as a solution to the issue of stable transmission between living memories. First, it placed the material in a dramatic context, making it easier to remember than an abstract catalog. Second, it was able to employ the formal techniques of poetry to assist memorization. Since Greek poetry was based on rhythm of accent, the correctness of a phrase in context was reinforced by the necessity of maintaining this rhythm; the primacy of rhythm also allowed the body to assist the ear and mind through unconscious or semi-conscious sympathetic movement. Third, since the epics were recited and listened to as a form of recreation from childhood onward, remembrance of them was linked to pleasurable memories and the reflex to seek out good feeling.

In essence, recitation of the epics

ing system level for reasons of efficiency and programmer convenience [Nürnberg et al. 1996a]. All Sprocs are clients of the VASM. Sprocs are differentiated by the structural abstractions they serve.

One example of a HOSS Sproc is the Link Services Manager (LSM). The LSM manages "traditional" hypermedia structure — namely, inter-application linking structure. It provides functions to create, navigate, manipulate, and destroy structure between application data. In the case of the LSM, behaviors correspond to the semantics of particular navigational structure traversals. [Leggett and Schnase 1994; Schnase 1993b].

The Taxon Manager (TaxMan) provides a second example of a HOSS Sproc. TaxMan serves taxonomic structural and behavioral abstractions. Taxonomic structural abstractions are widely applicable. For example, botanical taxonomists can use these abstractions to model concepts such as family, genus, species, etc. to classify plant specimens. Linguists can use these taxonomic structural abstractions to develop linguistic taxonomies to represent the developmental histories of different languages.

Some examples of the taxonomic behavioral abstractions served by TaxMan (for botanical taxonomies) include structure querying (e.g. find all family taxa that contain four genera with only one species each) and "structure collapsing" (e.g. collapsing species, subspecies, section, etc. into the genera taxa and transferring the associations between specimen data and these collapsed taxonomic levels to their respective genera.)

A third example of a HOSS Sproc is the Annotation Manager (AnnoMan). AnnoMan allows clients to associate annotations with any object served by any HOSS process. It does this by building VASM structural elements between objects to be annotated and the annotation objects themselves. A simple example was able to induce in reciters and listeners an almost hypnotic state that assisted correct remembrance. It also encoded cultural knowledge situationally, allowing people in everyday life to act like the heroes of their favorite epics and thus act according to cultural norms. Both of these were anathema to Plato, who was promoting reflective thought on the nature of abstracts. Plato's literacy allowed him to encode knowledge externally as a thing "in itself" and allowed him to examine concepts and their abstract structures without forgetting them. Thus, Havelock concludes, arises Plato's excoriation of poetry as education method, as inhibitor of abstract speculation on the nature of the true, good, and beautiful. For our purposes, we note that Havelock showed the consideration of ideas as eternal "things in themselves" is an artifact of literacy, not an archetypic aspect of thought.

More artifactual properties of literacy are examined in another seminal work of the field, Walter Ong's *Orality and Literacy* [Ong 1982]. Below, we present two of these properties: the notion of written truth as permanent truth, and the notion that words are merely placeholders for abstract entities.

Today, it is common for material to be written down and remain unchanged for extended periods of time. If that material had some veracity when it was recorded, we tend to regard its "truth" as a permanent property that can be redemonstrated at any time. This is not the case with orally transmitted knowledge, which cannot be "recorded" except in living memory. As a result, material for which there is no call is forgotten, and changes to the material that give advantage will occur. This is well illustrated by the case of genealogical recitation in some African cultures. It has been shown that when a family's star falls, its genealogy is gradually forgotten, while if it rises, not only is the

ple of annotation objects are short text items corresponding to "post-it" notes. Of course, more complicated annotation objects (e.g. audio or video streams) may be created. In principle, this may sound very similar to inter-application linking as implemented by the LSM. However, there are several significant differences. The LSM records with which application an inter-application link "end-point" is associated, so that end-points of link navigation may be displayed in the context in which the association was made. Conversely, the AnnoMan requires that the requesting client be responsible for the display of the annotations on an object. Also, the LSM handles only object identifiers, requiring applications to load data objects in which association end-points reside if traversals to these end-points are made. The AnnoMan actually handles the storage and retrieval of annotation objects.

IS2.3 Metadata Manager Layer

Metadata managers are processes that primarily serve abstractions to other system processes. They build the abstractions they serve from abstractions served by other metadata managers, Sprocs, and hyperbases. Metadata managers can be viewed as abstract data types, exporting data and functional abstractions.

In some sense, Sprocs are conceptually similar to metadata managers. The key distinction between them is made at the operating system level. The fact that Sprocs are known in advance to handle structure and behaviors has certain implications for swapping, pre-fetching and scheduling [Nürnberg et al. 1996a].

genealogy retained, but other families start incorporating elements of the successful family's genealogy into their own. This mutability caused havoc when contact occurred between tribal justice systems (which incorporated the genealogies) and those of colonial powers. Colonial officials would write down the genealogies and attempt to present the written records at later proceedings, only to be told that they were not the "true" genealogies, which had changed with the tribal political situation. Revisionism is reality in primary oral cultures; the beliefs that the written retains its truth for all time and that, by extension, publication implies truth are artifacts of literacy.

Another aspect of primary oral cultures is their recording of knowledge situationally rather than abstractly. As a result, their members have difficulty thinking non-situationally and do not recognize abstractions the same way members of primary literary cultures do. Some examples come from fieldwork done with mostly non-literate Soviet peasants of the 1930's. When presented with syllogisms, such as "where it snows, bears are white; in the north, it snows; what color are the bears there?", they would respond with statements like, "well, all the bears I've seen are brown." Those with some literacy responded with statements like, "*by your words*, the bears would be white". When asked to select "miscategorized" items from lists like "axe, grain, sickle, hatchet", they would give responses like "well, you can use the sickle to harvest grain, but an axe would be better than the hatchet to cut trees, so the hatchet doesn't belong." Literate people would exclude "grain", because they see the abstract category "tool" rather than the uses of the tools. Ong points out that these syllogisms and word games force one to focus on the abstract concepts denoted by the words and exclude the

IS2.4 Application Layer

Application processes are user processes familiar from conventional operating systems. The nature of these processes is open. One example of an application that has been built is a WWW Common Gateway Interface (CGI) [Berners-Lee et al. 1992] program that acts as a client to the TaxMan and AnnoMan, allowing queries to be made over a taxonomic space, displaying the results, and allowing users to annotate the records displayed in answer to the query. Another example is a Motif/X [Nye 1988; Young 1990] client that allows graphic editing and manipulating of taxa.

IS2.5 Other Tools

A number of tools have been built for application, metadata manager, and Sproc construction [Nürnberg 1994]. The H toolkit provides a certain process model and inter-process communication primitives. These primitives are based on a message/port data model. An open set of IPC media implement a byte transport layer (similar to the OSI Network Layer 4 [ISO 1992]) on top of which H ports are implemented.

A tool called the Protocol Definition Compiler allows quick construction of servers by generating the necessary protocol libraries from high-level protocol specifications, similar to CORBA IDL compilers [OMG 1995]. Server (MDM or Sproc) programmers define the interface to their servers by building a PDC Language (PDCL) file. These files are then compiled into libraries implementing the calls that effect the reading and writing of service requests and their responses. Using the PDC can greatly reduce the effort needed to code a HOSS server.

The Generalized Process Template (GPT) provides a skeletal server that can

irrelevant details, while the kind of word game popular in primary oral cultures, the riddle, forces one to consider precisely these "irrelevant" details. The consideration that words are merely placeholders for abstract concepts is another artifact of literacy.

OL2.2 Hyperliteracy

Many believe that we are entering an era where electronic tools for storing and manipulating information will be considered indispensable for everyday thinking and remembering. Douglas Engelbart [1963] expressed this belief when he described a "certain progression of our intellectual capabilities", from *concept manipulation* (manipulation of concepts in the mind alone) to *symbol manipulation* (expression of concepts through language) to *manual external symbol manipulation* (manipulating linguistic symbols using writing) and finally to *automated external symbol manipulation* (manipulation of symbols using computers). Engelbart's second stage corresponds with the concept of "primary orality," and his third stage with "pervasive literacy." We extend the concept of orality and literacy by positing a new property of culture, *pervasive hyperliteracy* or simply *hyperliteracy*, corresponding to Engelbart's fourth stage.

Why posit hyperliteracy? Why should digital libraries researchers and developers concern themselves with orality-literacy issues at all? If we are indeed entering an era where automated external symbol manipulation tools have become prerequisites of serious thought, then the designers of such tools (and digital libraries are certainly one kind) should be interested in which aspects of thought are intrinsic to language-using human beings and which aspects are products of the use of non-electronic

be tailored through the addition of server specific protocol libraries (such as those generated by the PDC) and callbacks designed to dispatch the server specific messages. The GPT also handles various name service operations (such as registering service names or finding other servers) automatically.

writing, since some of the latter may decrease in strength or disappear altogether in the residents of this new era. As can be seen from the above, these artifactual properties are not trivial, and they are precisely the concern of orality-literacy studies.

Our goal is not to define precisely the exact properties of hyperliteracy: such exact forecasting is impossible. We do believe, however, that given the potential for differences in thought between literate and hyperliterate individuals, those who wish to design the most useful systems for hyperliterate users (such as future digital library patrons) should consider the areas in which these differences might occur.

3 Botanical Taxonomic Scholarship

A curious aspect of some scholarly work practices is that often these practices are known to depend on false assumptions or over-simplifications of a problem. For example, most economic models use admitted caricatures of the way real people make decisions regarding their economic future. In some cases, such as in economics, these false assumptions are taken as reasonable both because they produce good results and make the models tractable (whereas more realistic models may not be).

In other cases, however, these false assumptions are simply products of tradition, based in part on artifacts of old technology and literate mindsets. We take as one very specific example our experiences with botanical taxonomists. For several years, we have worked together with botanists to build a digital library of herbarium collection data. We have been able to observe or otherwise determine (through asking for explanations) several common current work practices that have changed as our botanist colleagues both gain access to new technology and re-evaluate those parts of their old technology that dictated how they did their jobs. As a particularly good example of a current work practice dictated by current technology, consider that there are journals that use taxonomies that everyone (including the journal editors!) acknowledges are outdated. The editors of the journal, however, are reluctant to correct the errors in this standard taxonomy, partly because some of the fixes are not universally agreed upon, but also because changing the taxonomy now would “invalidate” articles just published. The current common practice, then, is for researchers to carry out their work using a more realistic taxonomy, and then literally “uncorrect” their terms to match the journal standard.

For reference, a statement of the botanical taxonomic problem is given here. The object of taxonomic classification is, not surprisingly, the *taxonomy*. These taxonomies consist of *taxa*, which themselves consist of other taxa or *specimens*. Taxa are composed in a hierarchic fashion — that is, the taxonomy itself may be viewed as a tree, with specimens at the leaves of the tree. Taxa at different levels in the tree have

different names, such as family, genus, species, etc. We briefly describe three interesting problems we observed the taxonomists encounter in their current work practices.

Different groups of taxonomists produce different taxonomies, even if the specimen set examined is identical. Groups in which particular specialists work on a given taxon may show more detail in the expansion of that taxon, or different groups may use different measures of similarity when composing taxa, weighting various kinds of evidence differently. It seems contradictory to have multiple solutions to a classification problem.

Separate taxonomic groups produce separate taxonomies, which are then identified with the groups that produced them. This is despite the fact that the taxonomy may always be used in conjunction with other taxonomies, or that it is based on the prevailing attitudes in the community. It seems contradictory that a communally defined, communally used product is identified with a small set of taxonomists.

The products of the work are often taxonomies, not simply revisions to existing taxonomies. Whether updates or new full revisions, the products are viewed as closed, well-defined entities, representing an opinion of a group at some time. However, new evidence, new analysis methods, and new interpretations are constantly being introduced. It seems contradictory to produce a well-defined, static analysis of an ill-defined, dynamic phenomenon.

IS4 Technology Applications

Addressing the three examples of seeming contradictions in current work practices requires different supporting technologies than those present in the physical library. What is required here are new digital library elements and tools, not derived from physical antecedents. Of course, it is impossible to say what all of these technologies will be. This section outlines some possible technologies to begin to address these issues.

IS4.1 Single/Multiple Taxonomies

Two important capabilities that help address single/multiple taxonomy problems are structure management and versioning. Hypertext structure management abstracts the structure over objects from the objects themselves. Oftentimes, this takes the form of abstracting traversal or navigational structure from data to be navigated. However, the principle of structure abstraction can be

OL4 Hyperliterate Work Practices

Addressing the three examples of seeming contradictions in current work practices requires different artifacts than those present in the physical library with its literate artifacts. What is required here are new digital library elements and tools, not derived from physical antecedents [Nürnberg et al. 1995]. Of course, it is impossible to say what all of these artifacts will be. This section outlines some possible artifacts to begin to address these contradictions.

OL4.1 Single/Multiple Taxonomies

One artifact of literacy is the notion of single-valued, static truths [Ong 1982]. The work practice of developing and publishing taxonomies separately from one another is a particular instantiation of this artifact. The product of this work is a taxonomy, a “taxonomic fact” or truth, presented and interpreted as such. However, the notion of truth is changing

applied to *any* realm in which multiple structures may be applied to a given data set. This is precisely the case in taxonomic work. Different taxonomies (structures) are built over the same specimen (data) set. Because the TaxMan inherits the structure management abstractions of HOSS, including sets of structure elements and their associated behavior processes, it can use these to partition the taxonomic data into consistent taxa sets.

Because the TaxMan is implemented on top of HOSS, it inherits the versioning support for data, structural, and behavioral objects therein. This provides a natural way to model difference over time in a given taxonomy, as well as differences with respect to authority in the same time frame. Additionally, changes in the analysis of specimens (perhaps the addition of new images or new genetic information) can be added to the data set by versioning the appropriate specimen data object, thereby not invalidating taxonomies based on the older version of the object.

IS4.2 Ownership of Taxonomies

One important capability that helps address ownership of taxonomy problems is annotation support. An important aspect of maintaining and using community objects is annotating and sharing annotations over community objects. Such annotations can be used to judge the communal level of acceptance of a part of the community body of knowledge or other particularly noteworthy aspects. Moderately sophisticated access control, search facilities and filtering mechanisms over the annotation space should be provided. We have developed a HOSS Sproc named AnnoMan, which models sets of annotations as structure contexts and provides these features. Modeling annotations as

from the literate view of static and single-valued to the hyperliterate view of dynamic and multi-valued. Consider the Guides project approach to teaching history in which various persona contextualize history from a particular point of view [Solomon et al. 1989]. The "truth" of the matter is a space, in which various points of view are represented. This contrasts sharply with the notion of the authority of the book as conveyor of a single, coherent message as in the literate world [Chartier 1994]. Perhaps instead of viewing the primary goal of a taxonomist as the generation of a new taxonomy, which then must be related to previous and competing taxonomies by the consumer, the product may be viewed as a change to the existing body of knowledge. In fact, in essence, taxonomists *do* view the *purpose* of their work in this way, but the actual product of their work, the printed taxonomy, is only a means to this end. Reconciliation and contextualization is the responsibility of the consumer.

OL4.2 Ownership of Taxonomies

Literacy promotes the concept of idea ownership by the individual, even when the idea represents a communally held truth. In this case, taxonomies are identified with their producers or publishers. There is no way to recognize the contextualization of a taxonomy in itself. However, the notion of authorship is changing from owner of a document and by extension its ideas to recorder of ideas that are the product of several people, past and present. Consider an analogy from the business world — the growing role of the analyst [Reich 1991]. The analyst provides a filtering or ordering function for data that is oftentimes already available. Many new companies focus no longer in the production of information, but its compila-

structure in a hyperbase is straightforward — different structural elements (annotations) are laid over existing taxa and specimens (data), grouped into contexts, and managed by existing hyperbase software that can provide access control.

IS4.3 Definition of Taxonomies

Another important capability that helps address definition of taxonomy problems is computation over hypermedia structure. The nature of the information in taxonomic research may be open in the sense that the boundaries around it may be hard to define, especially outside of a particular context. However, dealing with documents that exhibit no sense of closure at all can be disorienting as well. What is needed is a way in which the open space can be viewed as only "partially" open — that is, enforcing some sort of boundaries appropriate in a context, but allowing these boundaries to be crossed or recomputed. One way in which to do this is to take advantage of computation over structure which dynamically generates closed sets of structure appropriate for a particular use.

tion. This reflects a situation in which the problem of information is what to do with the overabundance of it (the "information explosion"), and not how to find and retrieve data [Chartier 1994].

OL4.3 Definition of Taxonomies

One artifact of literacy is closure of ideas. The product of taxonomic work is a well-defined, discrete entity. Products no longer must be closed. They may exist as changing entities over time, with poorly defined borders. Consider Web sites with links to many other sites. These sites have no closure per se. Where one chooses to draw boundaries is contextually and individually defined. This is in opposition to the closure engendered by books and other written entities [Chartier 1994]. As above, one new possibility is a communally maintained set of taxa, with various notes, modifications, and addenda separately maintained over these taxa. The boundaries of the communal knowledge could only be determined by a given consumer at a given moment.

5 Conclusions

The information systems thread of this paper asserted the existence of new work practices in botanical taxonomic scholarship enabled by new technologies. The new work practices, however, were assumed to arise spotaneously due to problems found in current work practices.

The orality-literacy thread of this paper motivated why certain new work practices might arise in botanical taxonomic scholarship, but did not offer any particular ways to cope with them.

The digital library will have to support the new work practices of people. The changes in such practices must be identified. We extended orality-literacy to *hyperliteracy* in an attempt to characterize the changes. The new practices will have to be supported by new technologies. We showed systems and tools able to support the needs of one particular research community. The threads in this paper, therefore, must rely upon one another, one for motivation, the other for prototypic solutions. We see this as a microcosm of the digital libraries research field — a field in which results

from many different and dissimilar areas will need to be synthesized to produce the research necessary to redesign the tools with which people think.

6 References

[Berners-Lee et al. 1992] Berners-Lee, T., Cailliau R., Groff, J., Pollermann B.: "World-Wide Web: The information universe"; Electronic Networking: Research, Applications and Policy, 2, 1 (1992), 52-58.

[Chartier 1994] Chartier, R.: "The Order of Books: Readers, Authors, and Libraries in Europe Between the Fourteenth and Eighteenth Centuries"; Stanford University Press, Stanford, CA (1994).

[Engelbart and English 1968] Engelbart, D., and English, W.: "A research center for augmenting human intellect"; AFIPS Conference Proceedings, 1968 Fall Joint Computer Conference, San Francisco, CA (1964).

[Havelock 1963]. Havelock, E.: "Preface to Plato"; Belknap Press, Cambridge, MA (1963).

[Hicks 1993] Hicks, D: "A Version Control Architecture for Advanced Hypermedia Environments"; Ph.D. dissertation, Department of Computer Science, Texas A&M University (1993).

[IIT 1995] Illustra Information Technologies, Inc.; "Illustra User's Guide"; Illustra Information Technologies, Oakland, CA (1995).

[ISO 1992] International Standards Organization: "ISO 7498 (1992): OSI — Basic Reference Model — Second Edition"; Transaction Publishers, New Brunswick, NJ (1992).

[Leggett and Schnase 1994] Leggett, J., and Schnase, J.: "Viewing Dexter with open eyes"; Communications of the ACM, 37, 2 (1994), 76-86.

[Nürnberg 1994] Nürnberg, P. "Implications of an open, extensible, and distributed hypermedia information system architecture for inter-process communication subsystem design"; M.S. thesis, Department of Computer Science, Texas A&M University (1994).

[Nürnberg et al. 1995] Nürnberg, P., Furuta R., Leggett, J., Marshall, C., and Shipman, F.: "Digital Libraries: Issues and Architectures"; Proceedings of the Digital Libraries '95 Conference, Austin, TX (1995).

[Nürnberg et al. 1996a] Nürnberg, P., Leggett, J., Schneider, E., and Schnase, J.: "Hypermedia Operating Systems: A New Paradigm for Computing"; Proceedings of the Hypertext '96 Conference, Bethesda, MD (1996).

[Nürnberg et al. 1996b] Nürnberg, P., Schneider, E., and Leggett, J.: "Designing Digital Libraries for Post-Literate Patrons"; Proceedings of the WebNet '96 Conference, San Francisco, CA (1996).

[Nye 1988] Nye, A.: "Xlib Programming Manual for Version 11"; O'Reilly and Associates Inc., Sebastopol, CA (1988).

[OMG 1995] Object Management Group: "The Common Object Request Broker: Architecture and Specification" (1995).

[Ong 192] Ong, W.: "Orality and Literacy: the Technologizing of the Word"; Methuen, New York (1982).

[Reich 1991] Reich, R.: "The Work of Nations: Preparing Ourselves for 21st Century Capitalism"; A. A. Knopf, New York (1991).

[Schnase et al. 1993a] Schnase J., Leggett, J., Hicks, D., and Szabo, R.: "Semantic data modeling of hypermedia associations"; ACM Transactions on Information Systems, 11, 1 (1993) 27-50.

[Schnase et al. 1993b] Schnase, J., Leggett, J., Hicks, D., Nürnberg, P. and Sanchez, J.: "HB1: Design and implementation of a hyperbase management system"; Electronic Publishing: Origination, Dissemination and Design, 6, 1 (1993)35-63.

[Schütt and Streitz 1990] Schütt, H., and Streitz, N.: "HyperBase: a hypermedia engine based on a relational database management system"; Proceedings of ECHT '90, Versailles,

France (1990).
[Shackelford et al. 1993] Shackelford, D., Smith, J., and Smith, F.: "The architecture and implementation of a distributed hypermedia storage system"; Proceedings of the Fifth ACM Conference on Hypertext (Hypertext '93), Seattle, WA (1993).
[Solomon et al. 1989] Solomon, G., Oren, T., Kreitman, K.: "Using Guides to explore multimedia databases"; Proceedings of 22nd Hawaii International Conference on System Sciences, Kailua-Kona, HI (1989).
[Waterfield 1993] Waterfield, R. (trans.): "The Republic of Plato"; Oxford University Press, New York (1993).
[Wiil 1993] Wiil, U.: "Experiences with HyperBase: A multiuser hypertext database"; SIGMOD RECORD, 22, 4 (1993), 19-25.
[Young 1990] Young, D.: "The X Window System: Programming and Applications with Xt. OSF/Motif Edition"; Prentice-Hall, Englewood Cliffs, NJ (1990).

Acknowledgements

This research was supported in part by the Texas Advanced Research Program under Grant No. 999903-230.