

Chapter 5

Review of Related Work

This chapter reviews work in various areas related to the material covered in this dissertation, some of which is cited in other chapters. The work is divided into four sections: technologies for interoperability in networked information systems, historical overviews of digital library research, evaluations of the nature and the impact of the Web 1.0 to Web 2.0 transition, and digital libraries as sociotechnical systems.

Technologies for interoperability in networked information systems

The notion of “networked” information – information units inter-connected by various organizational paradigms – precedes the invention of the computer, Internet and the later World Wide Web. Raymond [474] provides a relatively complete history of the pre-digital historical origins of underlying concepts.

A fundamental underlying notion is *hypertext*, “connect[ing] text across more than one document boundary” [136], thereby breaking down the traditional atomistic notion of the document (e.g., book, monograph, etc.). The origins of this notion can be traced to the visionary, and often forgotten, early 20th century work of Paul Otlet [392]. More frequently cited as the origin of hypertext is the microfilm-based *memex* envisioned by Vannevar Bush in his famous post-World War II article [105] about harnessing wartime science for peacetime challenges.

The invention of the computer and the spread of desktop workstations interconnected by a common network – the Internet – provided the foundation for the realization of Otlet’s and Bush’s prescient ideas. The two most notable early pioneers and creators

of demonstrable examples of the power of inter-linked digital information are Ted Nelson, with his Xanadu system [384], and Douglas Engelbart [177, 178], who is perhaps best-known for his invention of the computer mouse.

Coincident to this, the idea of extending the notion of the library to computers and networks began to take form in the 1960's. Early manifestations include the visionary designs of Licklider [346], the development of library automation systems and the machine-readable catalog [78], and the invention of modern information retrieval [428]. These were followed by a number of pre-web or non-web-based digital or electronic library concepts, experiments, and applications. These include the rather detailed design of a digital library system (and a plan for a digital library research program) by Kahn and Cerf [256] and client-server-based systems such as Schatz's Telesophy [431], the System 33 Document Service from Xerox [411], the RightPages system [237] from AT&T Bell Laboratories, and the CORE project [183] from BellCore, Cornell University, OCLC, and others.

The introduction to this dissertation describes the reports and workshops that led to the contemporary (post 1992) digital library initiatives. As described the notion of *interoperability* – providing the user with a seamless experience as they use heterogeneous, distributed information services (discovery, access, browse, etc.) – has been a central aspect of digital library research. Paepcke, et al. [394] describe the issues and historical lineage of interoperability in the digital library community. As described there and in Lynch, et al.[354], interoperability exists across a spectrum. At the lowest level it provides minimal interfaces and tools with which humans can navigate and infer coherence across multiple systems. The more common intermediate form is *syntactic* interoperability whereby common protocols, metadata formats, and digital object exchange standards provide a modicum of coherence across

systems. Considerably more complex and still the subject of research is the notion of *semantic* interoperability, which according to Paepcke, et al. [354]

... deals with the ability of a user to access, consistently and coherently, similar (though autonomously defined and managed) classes of digital objects and services, distributed across heterogeneous repositories, with federating or mediating software compensating for site-by-site variations. It also extends beyond passive digital objects to actual services offered by specific digital library systems.

This section focuses on related work on syntactic interoperability in the digital library community, the primary locus of the work described in this dissertation, while touching on tangential results in the web community. This work is divided into five categories: federated interoperability infrastructures, modeling of compound digital objects, metadata standards and ontologies, repository architectures, and semantic modeling.

Interoperability infrastructures for federated digital libraries

As described by Leiner [330] a federated (or confederated, the term used by Leiner) infrastructure seamlessly links distributed library services and content. The Dienst architecture and protocol [151, 289] and its instantiation in the NCSTRL global digital library [153], described in Chapter 7, is an example of an infrastructure for federated digital libraries. The reader is directed to that chapter for more details.

A distinguishing aspect of Dienst is its foundation in the web architecture – HTTP, HTML, and URIs – thus allowing accessibility to its functionality through standard web browsers. Two systems contemporary to Dienst, also built on web technologies and focused on computer science content, were the Wide Area Technical Reports Service (WATERS) [356] and The Unified Computer Science Technical Report Index (UCSTRI) [471]. Both of these systems had lower functionality than Dienst. They were by and large interfaces to central indexing sites (no federated search, no

collection model) and provided simple document access through HTTP and/or FTP URLs (no document model).

A considerably more complete, and complex, interoperability digital library infrastructure, which arose from Phase I of the Digital Library Initiative (DLI-1), was the Stanford InfoBus [422]. The goal of the InfoBus was to “extend the current Internet protocols with a suite of higher-level information management protocols” [422]. The use of the root “bus” in the name InfoBus connotes the same meaning as with hardware “buses”, a pluggable infrastructure. InfoBus was implemented over a Java-based CORBA [396] foundation and was used as a vehicle for tying together a protocol layer for managing items and metadata, and services such as search, payment, and rights and obligations. While these experiments produced demonstrable results, the InfoBus was not used for any widespread production environments, perhaps due to the decreasing use of CORBA in favor of more web-based paradigms or due to the complexity of the interoperability paradigm.

Perhaps as a reaction to the complexity of their previous full-functionality interoperability effort, the follow-on Stanford Digital Library project under Phase II of the Digital Library Initiative (DLI-2) provided a more limited federation mechanism, the Simple Digital Library Interoperability Protocol (SDLIP) [393]. SDLIP limits functionality to “search middleware” – providing protocols and software for mediating search interactions among several information sources. It was proposed as a mid-point between the widely-used (in the library community) but quite heavyweight Z39.50 standard [352] and extremely light-weight single text box, web-crawler based search paradigm. SDLIP was the basis for experimentation among several DLI-2 projects. However, as we noted in [170, 171] federated searching is fraught with performance, reliability, and interface problems, making it considerably less practical for

widespread use than centralized indexing (such as those used by crawler-based search engines) or harvesting (using mechanisms like those defined by the Open Archives Initiative Protocol for Metadata Harvesting [303]). Thus, while interesting technically, SDLIP has not found widespread use.

Another interoperability result of DLI-1 was the distributed agent architecture from the University of Michigan [60]. This architecture relied on cooperating system modules (or agents) to mediate actions across heterogeneous process and tasks. The intention was complete flexibility since a new function could be added to a digital library system via the addition of a new agent that understood basic interaction protocols. Only a few prototype implementations were successfully constructed.

The relative simplicity, flexibility, and web integration of the Dienst architecture led to its extension and adaptation in a number of follow-on systems and architectures that extend its protocol and services. OpenDlib⁴⁴ [112], a product of considerable EU funding, is an architecture for linking digital library services. A even more direct descendent of Dienst is the DPubS⁴⁵ (Digital Publishing System) at the Cornell University Library, which is an open-source software system designed to enable the organization, presentation, and delivery of scholarly journals, monographs, conference proceedings, and other common and evolving means of academic discourse.”⁴⁶

A final piece of related work in the area of federated digital library infrastructure is the current effort to develop a “digital library reference model” [109]. This work grows out the substantial work on federated digital libraries in the U.S., Europe, and

⁴⁴ <http://www.opendlib.com/area1/project.html>

⁴⁵ <http://dpubs.org>

⁴⁶ <http://dpubs.org/about.html>

elsewhere (rooted in part in the Dienst work) and is an attempt to formalize the structure and characteristics of such systems and provide the theoretical basis for future interoperability efforts.

Modeling of Compound Digital Objects

The notion of a document model has been a focus of digital library research from the beginning. The work is motivated by the opportunities for an expanded notion of “content” or “document” once the constraints of physical media are eliminated. In general, the work has focused on the notion of a compound digital object, a container for aggregating data, metadata, rights information, administrative (e.g. logging) data, and other related data streams.

This dissertation describes the various aspects of my work that fall into this area. Chapter 7 describes the Dienst architecture, which includes a compound document model, visible through the protocol. Chapter 8 describes work in the Dublin Core Metadata Initiative, which included the so-called Warwick Framework, a packaging abstraction for multiple metadata formats. Chapter 9 describes Fedora, a repository system with a powerful document model that not only includes compound aggregations, but also allows dynamic disseminations (i.e., linkages of static data streams and distributed web services), and expresses semantic relationships among digital objects. Chapter 12 describes work in Open Archives Initiative – Object Reuse and Exchange, which defines a standard for identifying and describing compound objects in terms of web architecture fundamentals.

The roots of this work and related efforts lie in what is known as the Kahn-Wilensky Framework (KWF) [255], which was a result of the Computer Science Technical

Reports Project⁴⁷ (CSTR) . KWF defines at an abstract level the notion of digital objects, identified via uniform naming system known as Handles [449], that contains key metadata, multiple other packages of data and metadata, and possibly recursively containing other digital objects. The KWF does not cover implementation details. However, in an early paper [290, 449] we outlined these issues, and the KWF has had a considerable influence on our follow-on work described in the remainder of this dissertation.

Nelson, in his PhD work [382], created a related compound object architecture inspired by both the KWF and Dienst work. His notion of “buckets” exist in a “Smart Objects and Dumb Archives” model that pushes intelligence or functionality usually found in repositories or archives down into the object. The motivation is to enhance the long-term survivability and portability of the compound objects. The buckets architecture was leveraged in a number of other digital library experiments including one where bucket functionality was expanded to dynamically change inter-bucket relations according to user retrieval patterns [73].

Over the past decade a number of compound object formats have emerged in the digital library community. One of these is the Metadata Encoding and Transmission Standard (METS) [367] that arose from experimentation with the KWF in the Library of Congress National Digital Library Program (NDLP)⁴⁸. METS is XML-based and accommodates the encoding of various forms of metadata for a digital object including bibliographic, administrative, rights, and structural. METS is used as the default storage format for the popular DSpace [439] institutional repository software. One of

⁴⁷ <http://www.cnri.reston.va.us/cstr.html>

⁴⁸ <http://lcweb2.loc.gov/ammem/dli2/html/lcndlp.html>

the common criticisms of METS is that it requires the classification of object components into pre-specified metadata categories, limiting its flexibility. Also, for some applications its hierarchical rather than graph-based model is too restricted.

Other communities have created or adopted their own object formats: examples include IMS-LOM [4], from the Learning Objects community, and MPEG-21 DIDL [245], originally from the consumer electronics community and adapted to the DL environment by Los Alamos National Laboratory [47]. Although the syntax and application domain for these formats differ, they all have goal of combining descriptive, structural and administrative metadata to represent digital manifestations of “intellectual works”.

Despite their utility, these formats all share a common problem, which motivated our OAI-ORE [316] work described in Chapter 12. There is no clear mapping of these compound objects into the web architecture. The result is that agents and services, such as crawlers for search engines, are unable to interpret the contents of these compound object representations without special provisions, which are generally not implemented and are deemed undesirable (each special case interferes with the efficiency and scalability of the crawler). As a result objects represented in these formats are frequently invisible to widespread web search techniques.

Metadata standards and ontologies

The term “metadata” is used differently across a number of contexts, among them scientific data, software engineering, databases, and digital libraries. The use of the term here is restricted to the (digital) library context. A NISO report defines DL metadata [13] as “... structured information that describes, explains, locates, or otherwise makes it easier to retrieve, use, or manage an information resource”. In addition, it falls into three main categories: descriptive (sometimes also called

bibliographic), structural, and administrative (including rights metadata and preservation metadata).

Historically metadata has been positioned as a successor and supposedly simpler alternative to traditional library cataloging. Library cataloging standards include both the rule sets, the most widely used of which is the Anglo-American Cataloging Rules Version 2 (AACR2) [213], and machine encoding rules, including the widely used MACHine-Readable Cataloging (MARC) format⁴⁹.

As stated elsewhere in this dissertation, metadata has been both a popular and controversial subject in the digital library research community. Outside the immediate digital library research community (e.g., in the general Internet community) its utility as a general tool for resource discovery has been profoundly criticized [165, 434] and it has been generally rejected by virtually all popular crawler-based web search engines because of quality and integrity problems in favor of content-based and link-based indexing.

Nevertheless, the work of the Dublin Core Metadata Initiative⁵⁰ stands out as one of the most well known results and ongoing efforts of the Digital Library research efforts. Chapter 8 provides more details on Dublin Core, describes my personal work in that initiative including work on the so-called Warwick Framework [286], and includes criticisms of the Initiative and the use of Dublin Core metadata in general. The chapter also includes a description of an alternative more descriptive metadata ontology that is the result of my work on the Harmony/ABC research effort.

⁴⁹ <http://lcweb.loc.gov/marc/marc.html>

⁵⁰ <http://dublincore.org>

A complete review of related work in the area of digital library metadata would be inappropriately lengthy and out of scope for this dissertation. The interested reader is referred to the aforementioned NISO report [13] for more details. As an alternative, I will enumerate a few of the more popular and active metadata efforts and standards that coexist with Dublin Core:

- Text Encoding Initiative (TEI)⁵¹ is a standard for marking up electronic text with a focus on humanities scholarship. It includes the ability to incorporate bibliographic information into the header of the marked up text to assist in the discovery and use of the included text.
- CIDOC/CRM [142] is a formal ontology for describing the structure and relationships among objects, which developed in the cultural heritage community. It is a complex description format that encompasses temporal, part/whole, and epistemological relations.
- Metadata Object Description Schema⁵² (MODS) was developed by the library community as a simpler derivative of MARC that is more expressive than Dublin Core for the description of electronic resources⁵³. It is encoded in XML.
- <indec> [425] is a format developed with support by the European Commission to support the encoding of rights information for electronic commerce.

⁵¹ <http://www-tei.uic.edu/orgs/tei/>

⁵² <http://www.loc.gov/standards/mods/>

⁵³ The creation of the MODS initiative is ironic since the original motivation of Dublin Core was the description of electronic resources, and the qualification effort within DCMI was motivated by a desire for more expressivity than the core elements.

- ONIX (Online Information Exchange)⁵⁴ is a metadata format developed by publishers to distribute electronic information about their books to booksellers, other publishers, and other organizations involved in both industry transactions. ONYX is widely deployed and is the basis of cataloging in Amazon.
- MPEG Multimedia Metadata⁵⁵ is a format for describing the structure, content, and relationships in multimedia audiovisual objects such as pictures, music, audio, speech, or video.
- Functional Requirements for Bibliographic Records (FRBR) [3] is a model developed by the International Federation of Library Associations and Institutions (IFLA) in an effort to modernize the model underlying the catalog. It is based on an entity-relationship framework. One of its major contributions is the distinction between the abstract notions of works and expressions – intellectual concepts and their realization in different genres such as musical scores or screenplays – and the concrete notions of manifestations and items – the translations and material instances of these abstractions such as a book or DVD. FRBR has been used in a number of experiments with new cataloging tools [454]. However, with decreasing interest and commitment to the notion of the catalog in general (in contrast to the increasing influence of web-based search engines), there is some question about the viability of any cataloging model including FRBR.

⁵⁴ <http://www.bisg.org/documents/onix.html>

⁵⁵ <http://www.multimedia-metadata.info/>

Repository architectures

Chapter 9 describes our work on Fedora, an architecture that is both a digital object model and a repository architecture. The notion of a repository has been central to most digital library architecture and, in fact, as described earlier, digital library architecture can be described in general as “repository-centric”.

Despite its wide use, there is some uncertainty in the digital library community about what a repository actually is. According to the KWF [255] “a repository is a network-accessible storage system in which digital objects may be stored for subsequent access and retrieval.” This definition, taken at face value, implies containment. However, arguably, a repository, in the manner of Fedora, may either contain (i.e., store on its own discs) objects, or may reference surrogates or references to those objects or to parts thereof [405]. By this logic then, it makes more sense to think of the repository of a service interface (a set of APIs) for the deposit, access, and management of digital objects regardless of their location of storage.

Digital repositories have proliferated in the library community with the introduction of the notion of an *institutional repository*[143, 253, 351]. These are part of an effort by University and research organization libraries to capture intellectual output of resident faculty and researchers “upstream” and make it accessible in a manner independent of its publication in more formal (e.g. Journal) publication. The institutional repository movement can be seen as one part of a broader *open access* movement [247, 491] that promotes free and open availability to the results of scholarship so that those results can be mechanically harvested [314], reused, and become the inputs for new scholarly work.

Fedora coexists with a number of other institutional repository architectures including:

- DSpace [439] developed by MIT and Hewlett-Packard. Although DSpace lacks a number of advanced features in Fedora, such as a flexible object model and open API, its integrated packaging of a professional user interface with a repository system has made it the most popular architecture for institutional repository applications.
- ePrints⁵⁶ was developed at the University of Southampton specifically to promote open access to scholarly publication. Like DSpace it is a software package that includes the underlying repository, user visible interface, and workflow layer.
- AdoRE [461] developed at the research Library at Los Alamos National Laboratory is a digital repository focused mainly on archiving and preservation. It takes a “write-once/read-many” storage approach and makes generous use of OAI-PMH [465] for protocol-based access to complex digital objects.
- Greenstone [494] developed at the University of Waikato in New Zealand is a repository application and suite of software for building complete digital library systems. It has been specifically constructed for ease-of-use, ease-of-installation, and modest hardware requirements to make it useful for developing countries and empower them to cross the “digital divide”. Because of this focus, Greenstone is supported by UNESCO.

Semantic Models for Digital Libraries

Chapter 11 of this dissertation describes the use of Fedora and Semantic web technologies to build a digital library based on an “information network overlay”.

⁵⁶ <http://www.eprints.org/>

Chapter 12 also covers the use of Semantic Web technologies as a means of integrating digital library compound objects into the web architecture.

This work builds on and is related to the overall Semantic Web Activity⁵⁷ of the World Wide Web Consortium (W3C). The efforts of this activity are wide-ranging and documented completely on their web page and elsewhere [18, 22, 56, 185]. The aspects of this work most closely related to this dissertation are:

- The Resource Description Framework (RDF) [273] a data model for expressing triples, assertions of typed relationships between named subjects and either literals (e.g., strings, numeric values) or named objects.
- The RDF Vocabulary Description Language (RDFs) [91] a mechanism for using RDF to define vocabularies (entities and relationships) for use in other RDF descriptions.
- The OWL Web Ontology Language [43] an RDF-based language for publishing and sharing ontologies on the World Wide Web.

In addition to these core Semantic Web technologies, there are two additional outputs of the Semantic Web Activity of special relevance to the OAI-ORE work described in Chapter 12:

- Named Graphs [110, 111] are a mechanism for instantiating a set of RDF triples (a connected or unconnected sub-graph) as a first class Resource, with a URI.
- POWDER (Protocol for Web Description Resources) [23] is the focus of a W3C Working Group⁵⁸ charged with developing a standard for associating structured metadata with groups of web Resources.

⁵⁷ <http://www.w3.org/2001/sw/>

⁵⁸ <http://www.w3.org/2007/powder/>

Finally, there is considerable interest in leveraging the entire semantic web technology stack as a substrate for building digital libraries and knowledge domains [278]. One notable example in this area is the JeromeDL [277] a self-declared “Social Semantic Digital Library” that allows rich bibliographic description of library content and social activities (e.g., bookmarking, semantic annotations, knowledge sharing) over library content.

Historical overviews of digital library research

A number of other members of the digital library research community have written about the past and future of digital library research, albeit with a different focus than that used in this dissertation.

Lesk [333] and Arms [29] provide textbook-like overviews of the origins and contemporary (at the time of publication) state of the field. Both books were published before the emergence of Web 2.0 (although Arms issued a revised online edition in 2005 [27]), and therefore do not describe the effect of Web 2.0 on digital libraries.

Borgman, who has been cited numerous times throughout this publication, has reviewed and critiqued the process and products of digital library research in a number of publications and presentations. Her 1999 paper “What are digital libraries? Competing Visions” [82] contains an excellent overview of the different communities involved in the shaping of the field. She expands this somewhat in her 2000 book “From Gutenberg to the global information infrastructure: access to information in the networked world” [79], and in a more recent book [80] she places the digital library efforts in the context of the broader cyberinfrastructure initiatives. Her recent talk [77]

at the ACM Joint Conference on Digital Libraries covered subject matter closely related to the analysis of this dissertation⁵⁹.

In his 2004 paper [119] Y.T. Chien, one of the original program managers of digital library funding at the NSF recounted the successes of digital library research but argued for understanding of “disruptive technologies” that DL research needed to pay attention to if it was to remain relevant. These technologies included mobile communications, broadband, distributed storage and retrieval (now called “the cloud”), and E-payment. His analysis is similar to that presented here in that he warns of fundamental disruption of core assumptions of digital libraries in a changing information landscape and changing expectations and demands by users.

The July/August 2005 issue of D-Lib Magazine celebrating the 10th anniversary of that online publication [492], included a number of interesting analyses of digital libraries.

A paper by Stephen Griffin [219], the program manager at the NSF mainly responsible for funding both DLI-1 and DLI-2, enumerates the success of both initiatives but notes the problematic nature of short-term funding for an inherently long-term endeavor such as the curation of information. He argues for attention to content as part of future cyberinfrastructure efforts.

A paper by Clifford Lynch [353], whose work has been cited numerous times through this dissertation, argues that the era of digital libraries as a definable and useful research area may be over. He points out the importance of work in targeted areas, especially digital curation and preservation, with a particular focus on the products of eScience. Another theme of his argument, tangentially related to the material in this

⁵⁹ And, in fact, the talk was partially informed by personal communication with me.

dissertation, is the need for research at the “crossroads of technology and social science” to examine information issues quite outside the traditional library context. These include the issues of personal information management, the role of digital information in teaching and learning, and the integration of collaboration (social activity) and information. Regarding the final point, he states, “...at least some sectors of the digital library community have always found active work environments to be an uncomfortable fit with the rather passive tradition of libraries”. Notably, these “active work environments” are characteristic of Web 2.0.

A paper by Paepcke, et al. [398], with the intriguing name “Dewey Meets Turing”, covers territory that is quite similar in some aspects to the argument presented in this dissertation. In this paper they state:

The coalition between the computing and library communities had been anchored in a tacit understanding that even in the 'new' world there would be coherent collections that one would operate on to search, organize, and browse. The collections would include multiple media; they would be larger than current holdings; and access methods would change. But the scene would still include information consumers, producers, and collections. [398]

They continue to describe what they call “the cuckoo’s egg surprise”, the emergence of the Web, which disrupted these assumptions and

... not only blurred the distinction between consumers and producers of information, but it dispersed most items that in the aggregate should have been collections across the world and under diverse ownership. This change undermined the common ground that had brought the two disciplines [computer science and libraries] together. [398]

The authors assert in their conclusion that despite this “the core function of librarianship remains. The information must be organized, collated, and presented.”

While I support their assertion of a role for librarians (i.e., information experts) in the emerging information paradigm, their conclusion seems to validate that this role will

be quite outside the traditional library meme and therefore may be something quite different than librarianship.

Finally, a paper by Arms [26] in the same issue uses viewpoint analysis, a technique from software engineering, to critically analyze digital library research. In a manner similar in ways to the content of this dissertation, he asserts that too much of the development of and evaluation of digital libraries has been done from an organizational or institutional viewpoint – in particular from the library perspective, which evaluates success in terms of the organization’s prevailing values or memes. He argues for the work to assume a user perspective and states that digital libraries should be integrated into, rather than distinct from the “single unified Internet that we take for granted today.”

Impact of the Web 1.0 to Web 2.0 transition

I have asserted throughout this dissertation that the web has gone through a fundamental transformation known commonly as Web 2.0. Although there is an abundance of work investigating the nature, structure, and evolution of many Web 2.0 applications such as Wikipedia, Twitter, Flickr, and blog usage, there is surprisingly little scholarly work on the phenomenon as a whole and its broad impact.

A large proportion of the literature falls into the class of popular information technology or business literature. The first widespread use of the term was by Tim O’Reilly in [390]. Another O’Reilly published popular business-oriented text is by Nickull, et al., who enumerate the strategies that entrepreneurs should take in order to leverage Web 2.0 technologies for their businesses. A similar argument is presented by Li and Bernoff [343]. Another book by Tapscott and Williams [453] takes a similar approach, describing how the traditional rules of competition and information

secrecy are invalidated by the Web 2.0 notions of benefitting from the “wisdom of crowds.”

In the realm of more scholarly approaches, Beer and Burrows [45] examine the Web 2.0 phenomenon from a sociologist’s prospective. They point out three societal impacts of Web 2.0, the first and third of which are especially relevant to the issues of intermediation and control zone that are described in this dissertation: “the changing relations between the *production and consumption* of content; the mainstreaming of *private information posted to the public domain*; and ... the emergence of a *new rhetoric of 'democratization'*.” [45] (italics in original)”

This notion of democratization of information in the digital age and its affect on the political process is the focus of a growing area of scholarship. Two recent books that cover this area are by Winograd and Hais [493] and a set of essays by Bolter [72].

Another scholarly examination on the social effects of Web 2.0 and the *participatory culture* that it has enabled is a white paper “Confronting the Challenges of Participatory Culture: Media Education for the 21st Century” written for the MacArthur foundation by Jenkins and others [251]. Jenkins has also written two books on this subject [249, 250]. Although the white paper primarily focuses on education, it also describes the profound changes in the flow of information in modern culture because of the *interactivity* [249] of new media and information technology. According to Jenkins et al.: “Participatory culture is emerging as the culture absorbs and responds to the explosion of new media technologies that make it possible for average consumers to archive, annotate, appropriate, and recirculate media content in powerful new ways” [251]. The white paper argues that this participatory culture cuts across education, creative practices, community interactions, and politics and requires

a new way of training young people in their opportunities and responsibilities in this culture.

Katz's excellent volume of collected papers [258] is perhaps the best current source of expert thinking on the effect of the changing web on higher education institutions including libraries. The chapter by Katz himself [259] and another one by Wheeler [485] examine the new model of information sharing in the Web 2.0 meme and describe how university information technology and information (i.e. libraries) will need to make profound readjustments in the face of this phenomenon.

As described earlier, the library community has responded to Web 2.0 with the notion of "Library 2.0". This term was originally coined by Michael Casey, who maintains a blog dedicated to the idea [145] and was further popularized by Chad and Miller [113, 373]. Perhaps the best attempt to dissect the meaning of Library 2.0 is a paper by Maness [358]. He defines Library 2.0 as "the application of interactive, collaborative, and multi-media web-based technologies to web-based library services and collections." The major impact of this, in his opinion is that:

Library 2.0 is completely user-centered and user-driven. It is a mashup of traditional library services and innovative Web 2.0 services. It is a library for the 21st century, rich in content, interactivity, and social activity.

I do not criticize the motivation or underlying premises of Library 2.0. I share with many others the hope that libraries and their collections and librarians as information specialists will be integrated into the Web 2.0 framework. I do find it ironic, however, that much of the Library 2.0 concept seems to imply the disappearance of the library as a recognizable and distinct meme, and projects a reality in which it and its assets are just other participants in a world of "crowd sourcing". This may be good, but seems to have little to do with the traditional mission of the library as a "first choice" for information and scholarly activities.

Finally, in a paper in First Monday, Cormode and Krishnamurthy [138] describe the effects of Web 2.0 on network and server loads. They do not cover any social or institutional effects.

Digital libraries as sociotechnical systems

Chapter 4 of this dissertation focuses on some of the techniques and theories of Science and Technology Studies (STS) to examine the historical context of libraries and digital libraries as information infrastructures. I asserted earlier that the digital library research community by and large viewed DLs from a technical perspective, assuming that technical changes would take place without disruption to established institutional models and roles. Because of this, they failed to anticipate the “genie in the bottle” that would emerge as networked information technologies, especially those manifested in Web 2.0, were adopted and adapted by all levels of society. This “genie” has revealed a number of profound changes in the social and personal relationships to information, knowledge, and the institutions interweaved with them.

This argument falls into a class of critical analyses, rooted in STS, that view infrastructures and technologies as *sociotechnical* systems – “networks of technology, information, documents, people, and practices” [469]. Van House’s excellent review article [468] provides a thorough overview of this type of analysis covering theories including Actor-Network Theory, Social Construction of Technology, Symbolic Interactionism, Epistemic Cultures, and Activity Systems. The interested reader is referred to that review. This short section will focus on the application of this approach to libraries and information systems.

David Levy, formerly at Xerox PARC now at the University of Washington, has written several papers and articles that take a broad perspective of libraries in general and digital libraries in particular. In a 1995 paper [341] with Cathy Marshall, another

colleague at PARC, he looks at libraries through the perspectives of documents, technologies, and work practices. He argues that Digital Library research is at times too narrow-minded and states that “library developments ought to be grounded in a solid understanding of past and present practices. Without this, we risk losing still relevant structures and practices while maintaining an allegiance to mythical and irrelevant features of an unrealized past or an idealized present.”

In a later paper from 2000 [337] Levy argues for the importance of the library as an institution that has “come to symbolize and to exemplify the values we impute to the entire [information] circuit” and states that the development of digital libraries must account for this and not assume that the simple availability of information in digital form is de facto good or better.

Finally, in a 2003 chapter [338] he focuses on the nature of the document as a fundamental constituent of the library and states that any work in the area of digital libraries must account for the historical continuity of the document paradigm in its movement across technologies.

Levy’s colleague Cathy Marshall, now at Microsoft, has also written about the need to recognize both the notions of continuity and change in the transformation of libraries from physical to the digital [361, 362]. Her 2003 article is particularly relevant to the discussion in this dissertation about “control zones” because it describes the nature of boundaries that exist in libraries whether they are surrounded by walls are absent of walls on the Internet. She makes some astute comments about the information loss that occurs when boundaries are broken down, stating that “document disaggregation... has an effect on a user’s ability to interpret the content. A document is more than the sum of its parts.”

There are a number of instances of the use of activity theory as a vehicle for examining and especially evaluating digital libraries. All of the applications of this methodology cite the advantages of understanding the interactions among tools, users, objectives, and community rule constraints for full understanding of complex information systems.

In work described in both a chapter [442] and journal article [441] Spasser uses activity theory to evaluate the effectiveness of the Flora of North America Digital Library Project. He argues for the advantages of this technique because of the manner in which it reveals organizational issues and contradictions in the construction and use of digital library technology. Collins, et al. [132], Blacker, et al. [67, 68], and Boer, et al. [70] also employ activity theory as a tool for evaluation of digital libraries and information environments.

Another theory from STS, actor-network theory, has also been employed as an evaluative mechanism for understanding the effect of digital libraries. In a book chapter on the subject [467], Van House claims that actor-network theory is a useful evaluation tool because of the way it accounts for the complexity of digital libraries as a tool, a boundary object, a locus of multiple translations, and an “active participant in the creation and circulation of documents, images, and other kinds of inscriptions.”

The Social Construction of Technology (SCOT) is another approach for looking beyond technology and understanding its cultural and social context. As described by Van House [468]:

Different groups have different arrays of problems; each problem has an array of possible solutions. The SCOT descriptive model proceeds with the “sociological deconstruction” of the object of interest, showing the different meanings the artifact has for different groups, focusing on the problems and associated solutions that each group sees with respect to the artifact. SCOT contends that a technological artifact possesses

“interpretive flexibility,” revealed through the different meanings attributed to it by the different relevant social groups.

This notion of “interpretive flexibility” is a useful framework for understanding the historical overview presented in Chapter 2: that is, the differing manners in which the involved communities (funders, computing and information scientists, librarians) interpreted the application of network technologies to the library context, and, in fact, the vastly divergent manner in which the DL and web communities applied the same online technologies.

Other applications of SCOT in the context of digital libraries include Kilker and Gay [263], who used SCOT as a framework for evaluating the “Making of America” project, and for understanding users’ perceptions of the performance of digital library technologies.

O’Day and Nardi [388] introduce another holistic mechanism, information ecology, as a mechanism for evaluating the design of digital libraries. They state that “we believe that looking at the broad picture is more important than focusing only on the details of particular technology innovations. Even when a new technology is meant to serve a general purpose, exposure to the richness of users environments’ is a viable resource for design input and creativity.”

The work of Kling, whose name and research are virtually synonymous with *social informatics*, deserves mention in this description of related work. In [270, 271] Kling defines social informatics as the “interdisciplinary study of the design, uses, and consequences of information technologies that take into account their interaction with institutional and cultural contexts”. Kling applies this contextualizing framework to libraries and in particular to digital libraries in [269, 271]. In 1996, Bishop and Star [61] wrote about the utility of social informatics as a tool for the design and evaluation of digital libraries. Agre in [17] employs social informatics and states that “Every

technology is embedded in the social world in complicated ways, and this is particularly true for digital libraries, which are intertwined with the cognitive processes of a complex society. Unless our conceptualization of society stands on an equal footing with our conceptualization of the technology it uses, our analysis will inevitably be overwhelmed by myths”

Finally, Fuchs, et al. [199, 200, 412], in works cited earlier, describe the compound social effects of the web and in particular in its “version” transition. They ground their work as an outgrowth of social informatics, but take a notably Marxist perspective, describing the effect of web technologies and the democratization of information creation and distribution on the structure of capitalist society.

In conclusion, while there exists a body of work examining digital libraries as sociotechnical systems, none of that work examines the nature of those systems and their library-based information model in the context of the increasingly dominant sociotechnical information system, the web.