

Next-generation digital libraries

Daniel Greenstein
Director, Digital Library Federation
dgreenstein@clir.org

Abstract:

The article traces the development of digital libraries and asks what is next for a suite of activities (or is it an organizational form?) that is fundamentally important to the future of research, learning, and cultural engagement. It focuses largely on the experience of leading US research libraries as a spring-board to two discussions: about key challenges that digital libraries more generally may confront in the next few years, and about how national cultural, legal, and funding regimes may influence the digital library's history and its future course.

Introduction

The article traces the development of digital libraries and asks what is next for a suite of activities (or is it organizational forms?) that is fundamentally important to the future of research, learning, and cultural engagement. To bound discussion, it focuses on the experience of leading research libraries and, with the exception of its final section, on the experience of research libraries in the United States. The article makes no claims about the extent to which the experiences of US research libraries are shared by peer institutions abroad or by public, academic, and other library organizations. Although information professionals are notoriously bad at forecasting future developments, the article's single greatest purpose is to anticipate with reference to this relatively narrow developmental history, some of the key challenges that will occupy the digital library and its many users over the next several years. It also supplies a starting point for discussion about how different national cultural, legal, and funding regimes influence the digital library's developmental history and its future course.

The article itself is a work in progress. At present, it includes three sections that characterize the first generations of the digital library's development. The greatest concentration is on the second and third generations when the digital library evolves during the 1990s from experimentation to institutionalization. Two final sections are forthcoming. The first of these characterizes the next generation digital library that is only now beginning to take shape. It anticipates that libraries will take on fundamentally new and different organizational forms – forms that enable them to deliver the full range of information and information services that their users have come to expect. Only the next generation digital library will supply this comprehensive range of services by sharing out responsibility for them, not, as presently, by husbanding that responsibility as a closely guarded institutional prerogative. In effect, the article suggests that the era of the “universal research library” – that is, the library that meets its users' many information needs with collections and services that are maintained primarily in house – is all but over. The single greatest challenge for the next generation digital library will be organizational in nature as hitherto independent institutions forsake their insularity so they may continue to fulfill their historic cultural and scholarly objectives. In the second forthcoming section, the article extends itself geographically as well as chronologically. There it looks at how national differences influence the digital library's development to date and, crucially, the manner in which it will meet its future challenges.

Although the text throughout refers to “the digital library” as if it were a monolithic entity, it is taken as given that no such entity exists. Nor, it should be stated, does the author possess a whiggish or progressive view of history in which actors and institutions are seen to march in lockstep along the same linear path into the future. Digital libraries develop at vastly different rates, to be sure. They also develop differently. Early adopters can become so burdened by the weight, the cost, and the risk associated with their premature innovation, that they may be seen, with hindsight moving into slow-moving backwaters where they are overtaken by those caught up in a more rapidly flowing mainstream. Late starters, on the other hand, skip entire developmental phases evolving new organizational and functional forms in doing so. Although this generational approach has severe limitations as a means of describing institutional trajectories, it is valuable in capturing understandings and discussions as they evolve in a closely knit professional community whose constituent members are very well and often acutely aware of what others are up to.

Finally, it is worth acknowledging that the article's intention is to provoke discussion, still better, practical action. It is accordingly opinionated; in places, it may be pointedly so.

The first generation. Taming the bibliographic record

Of the digital library's generations, the first is the longest lasting and is taken up almost entirely with online catalogs. It is not by any means complete, although there is evidence of its substantial progress. Its products are well known and include MARC and numerous national and international cataloging cooperatives and services. The purpose of this article is not to dwell on these earliest trials. Suffice it to say that they leave the library community with a renewed (at least reinforced) legacy of cooperative endeavor and individual libraries with some experience (not always painless) of developing and managing large and complex systems environments. Interestingly, the online cataloging choices that individual libraries made, often shaped the rate and pace of their maturation. In more than a few instances, libraries that opted for online cataloging approaches that failed to stand the test of time – were left “solving” their cataloging problem a second and in some cases even a third time, and thus struggling hard to play catch up into the 1990s.

The second generation. Experimentation

A second generation of digital libraries is defined, for all intents and purposes, by the Internet; rather, by the library's effort to harness the Internet to its historic roles. It begins, naturally enough with the initial penetration of the public Internet in the early 1990s. Though innovative in many respects, the second-generation digital library is perhaps best characterized like early nineteenth-century European diplomacy, as old wine in new bottles. At least, initially, the library exploited the Internet as an additional means of delivering traditional services. First and foremost of course was access to the online catalog that was quickly retrofitted so that it could be accessed remotely via Telnet, gopher, mosaic, and later on by today's web-browsing technologies. The library did not stop, of course, at the catalog. Already by the early 1990s, it was beginning to amass a range of other digital information, often on hand-held magnetic media – abstract and indexing and other referencing services were already available as were electronic editions of some scholarly journals. The Internet did not create these information products. By supplying a ubiquitous means for their efficient distribution, however, it encouraged their proliferation. Providing access to an initially small but growing corpus of electronic reference materials substantially improved the digital library's technical competencies. The experience, however, did not yet even begin to transform the library's functions or its role.

A further hallmark of this second generation is initial experimentation with digital reformatting. Surrounded with lofty rhetoric that made reference to universal access to all human knowledge, to innovative scholarship and teaching, and to national and international digital libraries, some research libraries turned themselves to the task of providing online access to selected holdings. It is too soon to fairly judge the results of these early initiatives. Arguably, the risks involved for early adopters were so great that they were forced to make too bold claims for their work. Consequently few of the earliest digitally reformatted collections stood a chance of meeting the

expectations that grew up around them. Judged in light of initial expectations, then, the earliest online collections are, with notable exceptions:

- ❑ too small to support more than very casual kinds of browsing such as that associated with the hobbyist;
- ❑ too idiosyncratic to be integrated meaningfully into larger virtual collections; and
- ❑ too passive in their presentation to maintain a user's interest for very long.

Here, too, there was nothing potentially revolutionary for the library. Although there were repeated references to ways in which digital collections would fundamentally change the nature of scholarship, they were created in the image of the library, not that of the researcher, the teacher or the student. At least in most cases (and again with notable exceptions) the earliest digital collection were conceived as means of exhibiting selected rare and special materials and in that respect, to illuminate the historic and cultural significance of a particular library's holdings. This orientation toward exhibiting aspects of local collections perhaps goes furthest in explaining the shallowness and relative un-sophistication of many early online collections.

Judged against more modest (and, I would argue, more realistic) claims, early digital collections were remarkable successful. In fact, many such collections were developed more as technical experiments than as means of redefining scholarship or of developing or even exhibiting library collections. Mounted more often than not as one-off projects with one-time funding these "working laboratories" were self-consciously constructed to road-test new technologies, to establish core competencies, to acquire key infrastructure for the library, and to cultivate a cadre of new and differently skilled professionals.

The second-generation digital library is known as well for its quest after the killer application, or rather by the numerous and parallel hunts for said applications. Like the Holy Grail, these killer apps were elusive and appeared to different seekers in very different places – in data and metadata formats, network protocols, even in systems and system architectures. The logic of their appeal is simple enough to comprehend. Digital libraries, it was discovered early on, were really quite complicated and, let's face it, very hard to build and maintain. Complexity was only compounded by the fact that few libraries had more than a very small handful of appropriately skilled research and development staff. The "killer app" was the silver bullet solution that promised to catapult the library into a networked age without undergoing the fundamental restructuring, staff re-skilling, and soul-searching mission reorientation that information technologies seem to force on virtually every other organizational entity known to late twentieth-century civil society.

The second-generation digital library was also competitively disposed towards its peers. Competition was hardly new. Research libraries have always vied for endowments, for collections, and for position. Where they are embedded within academic institutions (as so many are) they are part of a broader and possibly even more aggressive competitive dynamic. The competitive disposition is only notable in the second-generation digital library because it ran contrary to the deep information and service sharing that network technologies permit. Perhaps competition at this stage can be explained again with reference to the risk that the early adopters embraced and the need they felt to justify that risk in terms of demonstrable innovation. Whatever the cause, second-generation digital libraries struggled to find distinctive furrows to

plough. Some attached themselves to systems or standards that they hoped would emerge in the ascendant. Others wedded themselves to killer collections that would redefine or comprehensively inform whole fields of inquiry. The culture of online experimentation, competitively oriented as it was, made standards development work extremely tricky. Few would question the promise that standards had to offer. In fact, in the early 1990s, widespread adoption of standards promised the same benefits that we think of today – interoperability, persistence, and predictability of our online collections and services. The difficulty, of course, was in agreeing “which standards”, a question which was conceptually equivalent to “whose standards”, and which consequently could not be addressed on the true merits of any particular case (Greenstein, 2001).

The competitive disposition of second-generation digital libraries goes some way in explaining the relatively slow progress that was made extending collaborative initiatives. There are, to be sure, some very interesting examples of shared investigations. Various metadata initiatives are exemplary in this regard and some have had a welcome and lasting impact. Still, the ubiquitous penetration of networked technologies opened up whole new horizons for deep resource sharing (with collection development for example) that were left almost entirely unexplored. Constrained organizationally from pushing technologies to their logical extent, second-generation digital libraries fell back on safer modes of cooperation in talking shops that explored new ideas, and in still further shared cataloging activities. The latter extended tried and tested techniques by sweeping up digital information objects along with the bibliographic ones, and exploring mechanisms for constructing virtual as opposed to union catalog databases.

Organizationally, too, the second-generation digital library had a number of distinctive characteristics, though these may be bounded nationally by different cultures and funding milieus. Where leading edge digital libraries took deepest roots within operational libraries (as they did in the US), the second-generation digital library was understandably and self-consciously run out of a skunk-works; that is, an experimental laboratory set apart from mainstream operational library services (including technical services in most instances) so that it could flourish in a more unconstrained environment more appropriate to its innovative aims. In countries like the UK where leadership in digital libraries was taken by national funding programs rather than in individual institutions, the same drift into the skunk works is also apparent. There, operational digital library services were located at universities and colleges that competed with one another to host them. These services sat alongside rather than in the traditional library operations that existed at host institutions thereby enjoying considerable protection against both the bureaucracy and local faculty demands that could combine to act as impediments to innovation.

The second-generation digital library, then, was an experimental form. It explored new opportunities and gathered new competencies. It did all this within very safe harbors of soft-money projects and other activities that were organizationally, financially, and even culturally set apart from traditional operational services. If anything, the second-generation digital library sought only modestly to add to rather than transform the library’s historic collections and services.

The third generation. Institutionalization and some local transformation

Where the Internet provided the second-generation digital library with additional means of supplying access to holdings and selected other data, it offered the third generation digital library a means of reinventing itself. From the mid- to late-1990s, third generation digital library websites grew in their depth and their breadth, emerging toward the end of the millennium as online spaces that embraced the full range of library collections and services. Further, they made locally managed collections and services indistinguishable from those offered up at a distance and by third parties. In a technological instant, the underpinnings of the research library, the foundation of its physical, cultural, organizational, financial, and professional form, were removed and replaced. Historically, the research library had evolved to bring together in one place information and the people and services necessary to organize, preserve, manage access to and support use of that information. Libraries had of course developed numerous means of sharing content and services across space – shared cataloging, interlibrary loan, and selected depositories are pre-eminent examples. Still, the importance to research and learning of place-based collections and services acted as a fundamental impediment to more deeply shared activities. The Internet's ubiquitous penetration questioned fundamentally the logic, cost effectiveness and utility of this historic *modus operandi*. At least it forced questions about the library's role and its organizational form in an era when information access no longer required physical proximity to information objects or the professionals that managed and knew so much about them.

Similar transformative tendencies became apparent in developing digital collections. Having acquired core competencies and technical understandings, the third generation digital library abandoned experimentation and the “build it and they will come” philosophy that characterized early digital collections. It focused instead on fully integrating digital material into the library's collections, and on developing (and supporting with core funding) the requisite panoply of policies, technical capacities, and professional skills. Work by Jewell, Pitschmann, and Smith demonstrate that this is true across a range of digital information whether sourced in local rare and special collections or not. Jewell shows the extent to which leading research libraries have routinized the highly complex processes involved in identifying, evaluating, negotiating access to and supporting use of electronic information that is commercially supplied by third parties (Jewell, 2001). Pitschmann demonstrates an evolving and highly sophisticated understanding of the pitfalls, opportunities, and real costs involved for libraries that choose to organize access to “free” external Internet resources through variously described subject gateways, portals and other linking services (Pitschmann, 2001). And Smith demonstrates that the digital libraries that are digitally reformatting materials from their collections in the late 1990s (and not all digital libraries are), are doing so strategically rather than experimentally, for example to profile selected rare and special collections, to support specific teaching and research needs, or, indeed, to manage and conserve selected general holdings (Smith, 2001).

Third generation digital libraries are also far less interested in killer apps. Instead they conceive of the digital library as a complex online service environment that is supported by systems, both local and global, each of which supplies specific functions and inter-relates with others in a way that can be represented in a modular architectural schematic (Powell and Lyon, 2001). The

model is not only sophisticated, it is both practical and economical. It permits greater freedom in the selection of service components and enables the library to manage and respond to technical change with greater facility. Relying upon a modular systems architecture, the digital library can select a new authentication service, for example, or integrate a new authentication technology without re-engineering its entire service environment. It can concern itself primarily (though non-trivially) with the authentication system and the API through which it communicates with others. This modular approach is also fundamentally liberating since it permits libraries to think creatively about how to build upon services supplied by others. The extent to which libraries are able to realize any part of the grander visions, however, depends almost entirely on their ability to transcend their historic organizational independence and insularity.

In the third generation, the digital library's approach to standards setting is shaped by these nascent aspirations after a more deeply networked future. In brief, it is becoming fashionable in this third generation for digital libraries to lead from the rear with respect to the standards and practices they adopt (whether for their objects or for their modular system components). That is, there is a discernible preference amongst digital libraries for claiming adherence to practices that are already vetted and endorsed by at least one, but preferably several, peer institutions, rather than for making bold claims for local innovations. This approach has been apparent this last 18 months in initiatives of the Digital Library Federation (DLF) – a consortium of 27 leading US digital libraries that invest together in shared research and development. In that time, members of the DLF have reached consensus on a model for negotiating access to commercial journals and databases (CLIR/DLF, 2001) and on the minimum requirements for digitally reformatted book and serial publications (DLF, 2001). DLF members have also facilitated agreement among selected libraries and publishers about the minimum requirements of a digital archival repository for electronic journals (DLF, 2000), and are working on a metadata encoding and transmission scheme which will act as a means for conveying information about the structural, administrative and technical characteristics of digital objects (METS, 2001).

The sea-change in attitude and approach is partly economic. As part of the maturation process, digital libraries transform interesting skunk-works projects into essential library infrastructure. At that stage, failures can no longer be written off as “learning experiences” gained at limited cost and subsidized largely (at least, hopefully) with external or soft money. It also reflects changing understandings of digital collections and digital library architectures. Where online collections are concerned, digital libraries have recognized their highly constrained ability to supply end-users with what they really want – enough online information to meet their specific and evolving information needs. In short, the premium that has been set on interoperability and ultimately on usefulness and usability is so high as to substantially outweigh any benefit that might accrue to the original inventor or even the early adopter. The drift towards modular systems architectures also has a part to play since systems modules (locally and globally arrayed) must interoperate at a very fundamental level. Libraries at this stage demonstrate a desire, then, to pool their collective uncertainties and to define, and then frame, a broad suite of practices as benchmarks in which they can all invest and upon which they can more safely and predictably build. Increasingly attempts to develop and codify digital library standards and best practices are tapping into this very fruitful seam. Examples are available in the DLF's work on a registry of digital masters (DLF, 2001a), in the Metadata Encoding and Transmission Standard (METS, 2001), and in emerging requirements for digital archival repositories (RLG, 2001).

So are efforts to surface common supply of some essential services that digital library's require but can ill-afford to supply individually. Good examples are available in community experimentation with name-resolution services, registers of digital masters, and, increasingly, with digital archival repositories. The logic in each case is the same as that underpinning the shared cataloging initiatives that characterized the very first generation digital libraries. In addition, it demonstrates a tendency in the third-generation digital library to focus as much effort as possible on collections and services that reflect local collection strengths, support local users' specific needs, and in general distinguish the library from sister institutions regionally, nationally, and even internationally. Competition amongst libraries isn't entirely eradicated. Rather, it is shifted to the look, feel, and functionality of the online environments through which libraries' virtual visitors encounter their collections and services. In an environment where libraries are increasingly distinguishing themselves there and, of course, in their special collections, there is no need whatever for them redundantly to build and support essential underpinning infrastructure that extends beyond shared catalogs to include data repositories, book stores, and other such services.

The third generation seems also to rediscover users. Users, at least, don't figure much in the antecedent experimental phase. Why should they? The library at that stage is experimenting with new technologies – a purely internal affair – or looking for additional means of giving users access to holdings catalogs, reference materials, and some journals – areas where users' needs are deemed to be well and comprehensively known. As the integration of new technologies begins to transform the library, and the possibilities that are open to it in the construction of innovative networked services, the need to engage users and to [re-] assess their interests and needs, is perceived as pressing. In addition, by the late 1990s, there is evidence to suggest that the proliferation of Internet-based information is fundamentally altering the expectations, behaviors, and preferences of library users generally. Accordingly, the third generation digital library needs to know about what users want from the networked library and about the role that users perceive for the library in a much larger constellation of networked information and service providers.

Interestingly, some of the library associations that take the lead in quantifying traditional aspects of library use have been relatively slow to respond to this new and pressing need. The reasons for this are complicated. For a start, the metrics are complex and difficult to agree. If you doubt this, I merely suggest that you sit down with a small number of colleagues to see whether you can agree what constitutes a use of a networked information object. Secondly, the library associations that are so well suited to this kind of work are also typically memberships and as such are driven by the consensus that in this case is very difficult to engineer. Further, the measures themselves threaten potentially to disrupt the organization by fundamentally altering the criteria on which it admits (and excludes) new members. Debate about e-metrics is quickly transformed into debate about what institutions should and should not be recognized as leading research libraries and are accordingly difficult to resolve.

One result is that some of the best analyses of user behavior and user need are taking place at the grass-roots level in what can only be described as a series of largely uncoordinated guerilla attacks that are mounted at the institutional level and by *ad hoc* and informal associations (Troll,

2002). Amongst the numerous and important revelations that emerge from these very fragmented efforts (of which some synthesis would be very timely indeed!) is one that demonstrates the extent to which users want to work in highly personalized and infinitely malleable online environments; that is, environments that present them with the information and services they actually need at any one time. The operational lessons for the library are twofold: seamless presentation of collections and services irrespective of where, by whom, or in what format they are managed; and deployment of user profiling technologies that enable users to construct a networked information environment in their own image. Both lessons, if taken seriously and reflected in new operational services, have potentially revolutionary implications for the library. The first integrates the library as one part of a globally arrayed network of information services in a way that questions fundamentally its historic organizational insularity. The second potentially obscures from the user's view, the library's importance as a portal to that global network.

The third-generation digital library also appears to take very seriously its users' needs and interests through its support for a suite of activities that have become known as e-scholarship. Although the phrase has a frustrating tendency to take on some new meaning every time it is used, definition is most easily pinned down with reference to initiatives that enable scholars to produce and disseminate "publications" with minimal intervention from third party commercial publisher. E-scholarship, in fact, emerges in the third generation digital library as something of a cottage industry; one that has all the trappings (and the term is used advisably in both its best-known senses) of handicraft production (there are notable exceptions in PubMed, JSTOR, HighWire, and one or two other ventures). Leaving to one side debate about what role research libraries (digital or otherwise) should and can effectively play in this proto-publishing activity, it is clear that it has very little to do with supporting innovative forms of scholarly communications. In fact, it has everything to do with relieving very real and increasing pressures that are being placed on libraries relatively inelastic acquisitions budgets (It is time, in my view, to have this debate, preferably with more reference to real data and less recourse to high rhetoric than has been known tried to this point.). Interestingly, by focusing so much attention on the economics of scholarly publishing – and let's face it, we are speaking here of a very traditional form of publishing that is merely moving into an online environment, rather than one that is being fundamentally transformed – second-generation digital libraries are overlooking other more promising opportunities for supporting innovative scholarship. A single example, will, I hope, suffice by way of explication. As guardians of our cultural and scholarly heritage, libraries are naturally poised to supply services that help researchers, teachers, and learners to navigate, find, inter-connect, interpret, and use information that is relevant in whatever form it exists. Additionally, they are poised to capture those interconnections and interpretations as they are made, and to manage them as new information in its own right – that is, information that adds to the combined corpus of cultural and educational knowledge. The question, really, is whether libraries have very much to gain from their current course through which they seek to diminish the publisher's influence over the market for scholarly journals by becoming publishers of such journals in their own right and in this respect by moving into a business where they have neither history nor competence. Alternatively, are libraries better off adopting a more offensive stance in respect to publishers' pricing pressures – one that that builds on and exploits traditional roles and competences to support whole new forms of scholarly communication while at the same time ensuring that cultural and scholarly collections are well and fully used?

Third generation digital libraries finally, take on a profoundly different organizational shape than their ancestors. Managed in a second generation from ghettoized (but, importantly, ring-fenced) skunk works, digital libraries evolve in the third generation in a way that sees them integrated into the fabric of the library. Whether managed through centralized units, coordinating support services, or as a network of alliances between distinctive library cost centers, the third-generation digital library emerges as a suite of activities that involves library professionals across the organization and is formally a part of library services as they are generally conceived. As metadata librarians, human factors researchers, and digital library production coordinators co-mingle more extensively with subject librarians and bibliographers, public service and reference librarians, systems librarians, and library administrators, and as digital library services become indistinguishable from others that the library offers, the digital library begins to fade from view, not because it is less important than it ever was, but because it has become part of the furniture.

Tentative summary

Without prejudicing sections of this document that have yet to be written and that deal with the fourth generation digital library and its evolution in different countries, it is safe, I think to take a retrospective look of digital library developments in US research libraries. In 1996, it was possible to identify a small handful of second- and an even smaller handful of third-generation digital libraries in the US. In characterizing the landscape on which they operated, Donald Waters, then the director of the Digital Library Federation, offered an organizational perspective. "Digital libraries", he wrote, "are organizations that provide the resources, including the specialized staff, to select, structure, offer intellectual access to, interpret, distribute, preserve the integrity of, and ensure the persistence over time of collections of digital works so that they are readily and economically available for use by a defined community or set of communities" (Greenstein, 2000).

Looking toward the next-generation digital library, it is perhaps appropriate to extend that definition, to emphasize a functional perspective that may allow us to transcend the organizational boundaries from which so much current digital library activity is struggling to escape. From that perspective, the digital library is part of a complex networked array of information services. As one star in a much wider constellation, it acts to extend the breadth and scale of scholarly and cultural evidence and to support innovative research and life-long learning. To do this, the digital library it mediates between diverse and distributed information resources on the one hand and a changing range of user communities on the other. In this capacity, it establishes "a digital library service environment" - that is, a networked, online information space in which users can discover, locate, acquire access to and, increasingly, use information. Although access paths will vary depending upon the resource in question, the digital library service environment makes no distinctions among information formats. Books, journals, paper-based archives, video, film, and sound recordings, and historical artifacts in addition to those aforementioned are equally as visible in the digital library service environment as are online

catalogs, finding aids, abstract and indexing services, e-journal and e-print services, digitized collections, geographic information systems, Internet resources, and other "electronic" holdings.

In constructing a digital library service environment, the library becomes responsible for configuring access to a world of information of which it owns or manages only a part. Accordingly, the digital library is known less for the extent and nature of the collections it owns than for the networked information space it defines through a range of online services. In the world of commercial publishing, aggregators compete on the basis of the value-added services that they layer on top of overlapping electronic collections. Similarly, digital libraries establish their distinctive identities, serve their user communities, emphasize their owned collections, and promote their unique institutional objectives by the way in which they disclose, provide access to, and support the use of their increasingly virtual collections.

The digital library service environment is not simply about access to and use of information. It also supports a the full range of administrative, business, curatorial, and educational functions as may be required to manage, administer, monitor engagement with, and promote and ensure fair use of its "collections" and services whether they are available in digital or non-digital forms, whether located locally or off site. The digital library service environment integrates (and interfaces with) information repositories that are characterized by open-access shelving, high-density bookstores and availability via inter-library loan, and include data services and digital archival repositories. It manages information about collections and items within collections often throughout their entire life cycle. It incorporates patron, lending, and other databases, and integrates appropriate procedures for user registration, authentication, authorization, and fee-transaction processing. The digital library service environment may also evolve into a networked learning space, providing access to and a curatorial home for distance and life-long learning materials. The digital library service environment is, in sum, an electronic information space that supports very different views and very different uses of a universe of networked information. It is designed for the library's patrons as well as for its professional staff and with an eye on the needs and capacities of those who supply it with information content and systems. It is built in the full knowledge that information technologies will continue to change rapidly as will our understanding of how to apply those technologies in support of research, learning, and cultural engagement.

References

CLIR/DLF. *CLIR/DLF Model License*. Online. Available: <http://www.library.yale.edu/~llicense/modlic.shtml>. April, 2001.

DLF. *Draft benchmark for digital reproductions of printed books and serial publications*. Online. Available: <http://www.diglib.org/standards/draftbmark.htm>. July, 2001.

DLF. *Minimum criteria for an archival repository of digital scholarly journals*. Online. Available: <http://www.diglib.org/preserve/criteria.htm>. May, 2000.

DLF. *Registry of digital reproductions of paper-based books and serials*. Online. Available: <http://www.diglib.org/collections/reg/reg.htm>. August 2001.

Greenstein, Daniel. *DLF strategy and business plan*. Online. Available: from <http://www.diglib.org/about/strategic.htm>. September, 2000.

Greenstein, Daniel. "On digital library standards. From yours and mine to ours." *CLIR Issues*. November/December 2001. Online. Available: <http://www.clir.org/pubs/issues/issues24.html#digital>. December 2001.

Jewell, Timothy D. *Selection and presentation of commercially available electronic resources: Issues and practices*. Online. Available: <http://www.clir.org/pubs/abstract/pub99abst.html>. July, 2001.

METS. *Metadata Encoding and Transmission Standard*. Online. Available: <http://www.loc.gov/standards/mets/>. August, 2001.

Pitschmann, Louis A. *Building sustainable collections of free third-party web resources* Online. Available: <http://www.clir.org/pubs/abstract/pub98abst.html>. August, 2001

Powell, Andy and Lyon, Liz. *The DNER technical architecture: scoping the information environment*. Online. Available: <http://www.ukoln.ac.uk/distributed-systems/dner/arch/dner-arch.html>. May, 2001.

RLG. *Attributes of a Trusted Digital Repository: Meeting the Needs of Research Resources*. Online. Available: <http://www.rlg.org/longterm/attributes01.pdf>. xxx

Smith, Abby. *Strategies for building digitized collections* Online. Available: <http://www.clir.org/pubs/abstract/pub101abst.html>. September, 2001.

Troll, Denise. *Survey of methods for assessing use and usability of online collections and services*. Online. Forthcoming: <http://www.clir.org/pubs/>. February, 2002.