



D-Lib Magazine pioneered Web-based Scholarly Communication

Michael L. Nelson
mln@cs.odu.edu
Old Dominion University
Norfolk, VA, USA

Herbert Van de Sompel
herbert.van.de.sompel@dans.knaw.nl
Data Archiving and Networked Services
The Hague, Netherlands

ABSTRACT

The web began with a vision of, as stated by Tim Berners-Lee in 1991, “that much academic information should be freely available to anyone”. For many years, the development of the web and the development of digital libraries and other scholarly communications infrastructure proceeded in tandem. A milestone occurred in July, 1995, when the first issue of D-Lib Magazine was published as an online, HTML-only, open access magazine, serving as the focal point for the then emerging digital library research community. In 2017 it ceased publication, in part due to the maturity of the community it served as well as the increasing availability of and competition from eprints, institutional repositories, conferences, social media, and online journals – the very ecosystem that D-Lib Magazine nurtured and enabled. As long-time members of the digital library community and frequent contributors to D-Lib Magazine, we reflect on the many innovations that D-Lib Magazine pioneered and were made possible by the web, including: open access, HTML-only publication and embracing the hypermedia opportunities afforded by HTML, persistent identifiers and stable URLs, rapid publication, and community engagement. Although it ceased publication after 22 years and 265 issues, it remains unchanged on the live web and still provides a benchmark for academic serials and web-based publishing.

CCS CONCEPTS

• Information systems → Digital libraries and archives.

KEYWORDS

D-Lib Magazine, World Wide Web, Digital Libraries, Scholarly Communication, Open Access

ACM Reference Format:

Michael L. Nelson and Herbert Van de Sompel. 2022. D-Lib Magazine pioneered Web-based Scholarly Communication. In *The ACM/IEEE Joint Conference on Digital Libraries in 2022 (JCDL '22)*, June 20–24, 2022, Cologne, Germany. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3529372.3530929>

1 INTRODUCTION

The [WWW] project started with the philosophy that much academic information should be freely available to anyone. It aims to allow information sharing

within internationally dispersed teams, and the dissemination of information by support groups. [12]

Internet-based digital libraries in support of scholarly communication (or “electronic libraries” as they were frequently known as prior to 1994) predated the popularity of the web; some of the well-known examples include: “Knowbots” [53], the CORE electronic journal project [69], Netlib [26], xxx.lanl.gov [34], Computer Science Technical Reports Project (CS-TR) [51], Wide-Area Technical Report Server (WATERS) [73], and the Langley Technical Report Server (LTRS) [87]. The NSF-funded Digital Library Initiative (DLI, 1994–1998) which co-occurred with the rapidly increasing interest in the web and was accelerated by the late 1993 release of the NCSA Mosaic browser marks the point where the web began to deliver on Berners-Lee’s original vision for academic information.

These early digital libraries were moved to the web and embraced an access approach based on HTML landing pages that provided descriptive information pertaining to a scholarly resource and links to the actual content. The content itself was typically available in PDF or PostScript format, not HTML. This was a necessary choice because the digital libraries largely consisted of digitized paper-born documents or papers that originated with communities at the digital forefront that used TeX/PostScript for rich manuscript markup, including features such as formulas and graphs that were not supported in early HTML. Traditional publishers that moved their holdings to the web followed the same approach, either because they were inspired by early digital libraries or because it was easier to migrate their print-oriented workflows to PDF rather than to HTML publishing. Although the drivers that led to this approach no longer stand in the way of publishing scholarly contributions in HTML, the access paradigm based on HTML landing pages and PDF papers remains dominant to date. Scholarship is accessible by using HTTP as a transport mechanism, but with few notable exceptions (e.g., BioMed Central,¹ Public Library of Science,² PeerJ³) and despite promising explorations (e.g., [19, 109]), scholarship is still not web native. This is a regrettable state of affairs, especially since pioneering efforts in the early days of the web effectively demonstrated the feasibility of publishing scholarship in the web’s language HTML. In this paper, we zoom in on D-Lib Magazine that became a prominent communication venue in the Digital Library community that we were part of and also pay attention to similar efforts (Ariadne, First Monday, The Journal of Digital Information) that communicated content in the same scholarly realm.

In July, 1995, the Corporation for National Research Initiatives (CNRI) published the first issue of D-Lib Magazine (www.dlib.org). D-Lib Magazine was the most visible and impactful component of



This work is licensed under a Creative Commons Attribution International 4.0 License.

JCDL '22, June 20–24, 2022, Cologne, Germany
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9345-4/22/06.
<https://doi.org/10.1145/3529372.3530929>

¹<https://www.biomedcentral.com/>

²<https://plos.org/>

³<https://peerj.com/>

the D-Lib Forum⁴ administered by CNRI and funded by the Defense Advanced Research Projects Agency (DARPA). In July, 2017, D-Lib Magazine published its 265th and final issue, bringing to a close a successful 22 year run that saw it evolve into an entity around which the entire digital library (DL) community coalesced. D-Lib Magazine was itself an innovation: it was published in HTML only and thereby encouraged exploration in scholarly publishing with hypertext and hypermedia, it was open access with no article processing charge so it reached a broad community, its “magazine” focus and initially monthly publication schedule facilitated community building in a pre-blog and pre-social media world, and it found the elusive middle ground between researchers and practitioners.

During its 22 year run, D-Lib Magazine offered several opportunities for self-reflection for both the magazine and the community at large. In 2000, Bill Arms surveyed the first five years [9]. In 2005, a ten year anniversary special issue⁵ was published with contributions from many of the central figures of D-Lib Magazine and the DL community at large [52, 132]. The 20 year anniversary had a more muted tone, with only the issue’s editorial marking the event [65]; perhaps because editor Larry Lannom knew the time for the final editorial was not far off [66]. On the 25 year anniversary of the first issue, we reflected on the impact of D-Lib Magazine, both for the information that it conveyed as well as a proof-of-concept for many DL and web concepts and technologies that we enjoy today [90].

In this paper we explore two threads for assessing the legacy and impact of D-Lib Magazine: first, the tools, techniques, and modalities of Web-based scholarly communication pioneered by D-Lib Magazine that have recently been embraced and extended by pioneers from the larger scientific community, and second, the topics covered in the first issue of D-Lib Magazine as concepts that remain central to the Digital Library (and thus the Joint Conference on Digital Libraries (JCDL)) community. D-Lib Magazine is important not only for *what* we, the Digital Library community, shared, but also in *how* it enabled us to communicate.

2 D-LIB MAGAZINE AS A PUBLISHING EXPERIMENT

From the editorial of the first issue [31]:

The magazine is itself an experiment in electronic publishing, which fulfills its communication function for the Digital Library Forum by testing the limits of writing in and for a wholly networked environment. We have no – and propose no – print analogue, and we will be most intrigued by substantive articles that take advantage of the power of hypermedia while retaining the strengths of traditional, print publishing.

D-Lib Magazine was unique in many respects. First, although it clearly billed itself as a “magazine”, it quickly became a venue where original research was published. Second, although it initially offered additional services and categories, the real innovation came about because it embraced HTML, and only HTML, as the publication medium. HTML allowed the articles themselves to take advantage of a rapidly evolving medium, including links and multimedia in a

way PDF-primary publications could not. Finally, with the vantage of 25 years, the decisions made in how D-Lib Magazine would be structured and maintained compare favorably to other Web-based publishing peers which began shortly after D-Lib Magazine.

2.1 More than a magazine, even if not quite a journal

Although early issues had unsuccessful experiments with HyperNews [16] for comments as well as a separate “technology playpen” / “technology spotlight” section [132], these features were eventually subsumed within the HTML publishing experiment itself, and D-Lib Magazine’s primary unit of currency became its articles. From 1995 through 2017, D-Lib Magazine published 265 issues and 1062 articles (D-Lib Magazine actually defined and evolved many different categories of contributions [131], but we refer to entries available from the title index as simply “articles”). The issues were published monthly through June, 2006 (with the July/August issues published simultaneously as “7/8”), and it switched to bimonthly publication from July/August 2006 through July/August 2017. D-Lib Magazine aimed for “articles that are 1,500 to 3,000 words in length and seldom accept articles in excess of 5,000 words” [132].

Figure 1(a) shows the number of articles published each calendar year, and Figure 1(b) shows the average number of words per article for each calendar year. From Figure 1(b) we can see that although switching to bimonthly publication in 2006 reduced the number of articles per year, it did not halve it. Even though in 2017 D-Lib published only four issues (instead of six), the total number was only slightly down from 2016, perhaps indicating clearing the queue of remaining articles for the year.

Figure 1(b) shows a trend of shorter articles in the first three years, and then finally hitting its stride in 1998, perhaps corresponding with the acceptance of the format by both authors and editors. From 1998 on, the values fluctuate (we are unsure of why 2009 has a low value) but it is not until the last six years (2012–2017) that the word count approximates the early peak from 1998.

Even though it was always “a magazine, not a peer-reviewed journal”, and did not have an editorial board like a conventional journal (though it did have an advisory board,⁶) D-Lib Magazine had a significant impact in the conventional literature and served as a de facto journal. A ten year anniversary analysis (from 2005) showed that D-Lib Magazine had acquired 147 citations from the ACM/IEEE Joint Conference on Digital Libraries and its predecessor conferences [132]. A more detailed authorship and citation analysis showed over 1300 citations in the first 15 years [92].


2.2 Other contemporary peers

There were other contemporary experiments in on-line publishing from generally the same community as well. We will periodically compare D-Lib Magazine to three contemporary peers: Ariadne, First Monday, and the Journal of Digital Information (JoDI), and all of whom pioneered idioms and designs embraced by much later open access, HTML-centric serials such as BioMed Central, PLOS One, and PeerJ.

⁴<http://www.dlib.org/forum/note.html>

⁵<https://doi.org/10.1045/july2005-contents>

⁶<http://web.archive.org/web/20000226003334/http://www.dlib.org/forum/advisory-board.html>



(a) Total articles published per year (1995 and 2017 were incomplete years).

(b) Average words per article per year.

Figure 1: D-Lib Magazine 1995–2017

Ariadne is an online magazine that began publishing in 1996 and is still publishing (78 issues since 1996). It was similarly not peer-reviewed, aimed at practitioners, and was initially funded by the Joint Information Systems Committee (JISC, since renamed Jisc), a UK activity that can be considered roughly analogous to the USA DLI program. Ariadne also had an HTML focus from the very beginning.

First Monday began in 1996 as a monthly peer-reviewed journal, and is still being published. It began as an HTML/web publishing experiment by commercial publisher Munksgaard, but was sold to its editors in 1999, who maintained it as open access [95]. It is now hosted at the University of Illinois at Chicago.

The Journal of Digital Information (JoDI) began as a peer-reviewed journal in 1997, and ceased publication in 2013 after irregular publication of 46 issues. It was originally hosted at the University of Southampton, and then moved to Texas A&M University and Texas Digital Library.

3 D-LIB MAGAZINE'S INNOVATIONS IN WEB-BASED SCHOLARLY PUBLISHING

As the initial editorial makes clear, D-Lib Magazine was an ongoing experiment in “electronic publishing” itself, and as a result was an early adopter and proof-of-concept for a lot of conventions and techniques that are now best practices in the community.

In the sections that follow, we will review some of the innovations that D-Lib Magazine explored and relate them to their contemporary peers.

3.1 HTML

Perhaps most importantly, D-Lib Magazine was always published in HTML – and only in HTML: there was never a parallel PDF version. Submissions were encouraged in MS Word,⁷ but the editors handled the conversion to HTML themselves. Adopting an HTML-only publishing strategy seems obvious in retrospect, but considering the limitations of HTML ca. 1995 (cf. HTML5 [1] today) this was a bold strategy. Despite the dominance of the PDF in the scholarly publishing ecosystem, the HTML format allowed authors to experiment with multimedia and interactivity extensions not possible with PDFs. Quoting from the October, 1995 editorial, one gets a glimpse of the willingness to explore the boundaries of what an HTML-only publication could be [32]:

You will see that the stories have varied in their treatment of images, for example, in the background color, and even in the organization of the text itself. But I do not believe that these individual treatments posed a problem for our readers, partly because the stories are unified by subject, partly because the medium is itself experimental and preconceptions are fairly few, and partly because in each case, the structure of the story reinforces and extends its informational content. Thus, the highly visual story that the Informedia team wrote on indexing video⁸ subtly embodies the notion of frames in its file structure. It offers readers multiple paths through the material and cues through buttons not unlike the signage found in museums and airports,

⁷<http://web.archive.org/web/20000613151426/http://www.dlib.org:80/dlib/author-guidelines.html>

⁸<http://www.dlib.org/dlib/july96/07wactlar.html>

and through menus that other writers for the magazine have also employed. In the same issue, the Netlib authors used a classic, straightforward narrative approach with an internal menu to explain the complex structure of a library of mathematical software.⁹

As authors, we appreciated the editors' willingness to explore what new features were possible in an HTML scholarly publication. For example, in our 1999 article about the Universal Preprint Service [117], we included screen cams to show the now defunct ups.cs.odu.edu digital library in action. Those screen cams were stored in .exe format and would thus likely require emulation to run now, but those animations (stored at dlib.org) would not have been possible in a PDF. Another of our articles from 2002 used animations, but this time in a more web-friendly and standard MPEG format [85]. In a 2005 article [17], we did not use animations, but did have 377 images linked from the article, a feat that would have been unwieldy at best in PDF. Our last article in D-Lib Magazine used JavaScript to make annotated hyperlinks in the article actionable, thereby serving as a demonstration of how "Robust Links"¹⁰ could work in practice [119].

Ariadne was always HTML only. For its first three issues, First Monday utilized parallel HTML and PDF formats, but the PDF format was just the HTML file saved to PDF. While both Ariadne and First Monday utilized hyperlinks, we do not recall instances of allow more common hypermedia, including videos and executable, custom JavaScript libraries, and hundreds of hyperlinked images. JoDI also used a hybrid approach, allowing for HTML, PDF, or both, depending on the authors' preferences.

Although they could not have known it at the time, being directly on the web with only HTML proved to be a boon for early web crawlers, such as search engines and web archives. Although both search engines and web archives successfully process most non-HTML formats now, for a long time they had a strong preference for HTML documents only [76, 78, 110]. This ran counter to the prevailing digital library architectural thinking of the time, where crawlers would interact with dozens or hundreds of complex repositories instead of crawling documents directly from millions of simple HTTP servers (see 4.2).

3.2 Open Access

Another groundbreaking innovation for D-Lib Magazine was that it was open access before that term was even coined, with the authors retaining their copyright, and D-Lib requiring neither subscriptions for readers nor article processing charges from the authors. This ensured it reached a wide audience, both authors and readers, but it also resulted in chronic funding problems after the expiration of the initial grants that supported the D-Lib Forum ended. In an editorial for the ten year anniversary issue [52], Robert Kahn said:

Producing a high quality magazine on the net each month turned out to be somewhat less difficult than I would have expected, due almost entirely to the quality of the editorial staff and the willingness of the readership to contribute interesting articles. Funding the continued production of the magazine has

been, perhaps, its biggest challenge. While the initial funding from DARPA covered most of the early costs, DARPA was unable to continue the support indefinitely. Subsequent funding from NSF helped greatly, but covered perhaps half the ongoing costs, with CNRI picking up the other half.

Although subscriptions and author fees were considered [64], they were never implemented. In 2007, the "D-Lib Alliance" membership organization was created [130] that assisted with funding, but the final issue in July 2017 acknowledged that decreased financial support was part of the reason for ceasing publication [66]:

Financial support for the magazine has waned over recent years, the number of unsolicited high quality articles thrown over our transom has declined, and the very phrase 'Digital Libraries' has gone from sounding innovative to sounding a bit redundant. In short, it seemed like time to make a graceful exit.

All three peers, First Monday, Ariadne, and JoDI, were also open access, although the earliest iterations of JoDI required readers to establish an account on the server. Unfortunately, this limited web archiving of the earliest versions of JoDI, since web archive crawlers are unauthenticated and thus could not crawl the site. This restriction was not removed until 2000, and it is at that time that the web archives were able to begin crawling.

3.3 No landing pages

An innovation that resulted from open access HTML-only publishing was D-Lib being the first venue to have its handles (and later DOIs, to be discussed further in section 3.6) resolve to articles themselves, not a landing page describing the article. By eschewing landing pages and PDF, which remains the choice of most (pay-walled) journals, D-Lib Magazine was able to subtly reinforce that its content was part of the Web, and not something separate, to be downloaded via the Web. The ability to link and provide embedded multimedia enables the scholarly object to enjoy the same advances (and risks, such as link rot [47, 56, 75]) as the rest of the web.

Since Ariadne was always HTML-only, it too never utilized landing pages. But both JoDI and First Monday used landing pages, the latter continuing to use them long after support for parallel HTML and PDF formats was abandoned.

3.4 Persistence of content and layout

Another significant decision was to fix the template and formatting of past issues, and not reformat earlier issues with updated templates. Updates were only made in the cases of errata and corrigenda.¹¹ D-Lib Magazine updated their design as tools and experience allowed, but the first issue looks the same today as it did more than 25 years ago, thereby serving as a monument to the best practices of the time. Indeed, the live web version of the first issue and the web archived version of the first issue are indistinguishable (Figures 2, 3, 4).

¹¹Although we thought we remembered this policy being explicitly stated somewhere, we could find no record of it. In emails with former editors Larry Lannom and Cathy Rey, neither could recall such a document. The closest we could find was "Once the issue has been released, only vital corrections or changes will be made to the file. These changes will be noted and dated at the end of the file." in the Author Guidelines: <http://web.archive.org/web/20000613151426/http://www.dlib.org/dlib/author-guidelines.html>.

⁹<http://www.dlib.org/dlib/september95/netlib/09browne.html>

¹⁰<https://robustlinks.mementoweb.org>

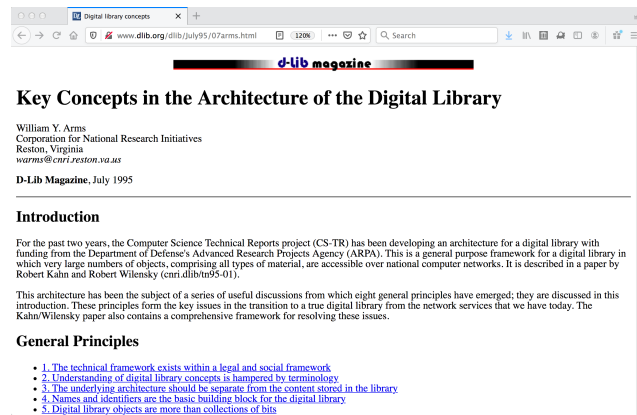


Figure 2: D-Lib Magazine, live web: <http://www.dlib.org/dlib/July95/07arms.html>.

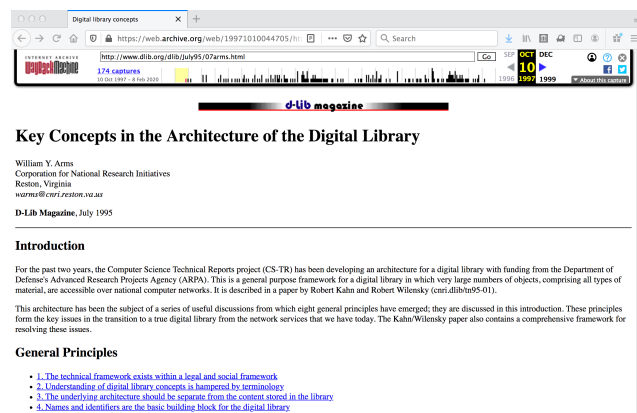


Figure 3: D-Lib Magazine, archived in 1997: <https://web.archive.org/web/19971010044705/http://www.dlib.org/dlib/July95/07arms.html>.

```
$ date
Wed Dec 8 13:25:20 EST 2021
$ curl -s http://www.dlib.org/dlib/July95/07arms.html | md5
3cc0fb32a7fe8f1f4de9a40aa5069cfe
$ curl -s https://web.archive.org/web/19971010044705id_/
http://www.dlib.org/dlib/July95/07arms.html | md5
3cc0fb32a7fe8f1f4de9a40aa5069cfe
```

Figure 4: Using curl to download both the live web version and first archived version (from 1997 and in “raw” format, via id_) and show they produce the same md5 hash.

This dedication to maintaining the original layout is unmatched by D-Lib’s peers: Ariadne, First Monday, and JoDI all changed their content and layout over time. In First Monday’s case, it was subtle since the original page was simply encompassed in an HTML frame when it changed publishers, and though the inner frame appears similar to the original content (i.e., ignoring the parent frame), they are not byte equivalent and the content has transformed over time.

3.5 URL persistence

Not only did D-Lib keep their HTML and style intact, but thanks to an ongoing commitment from CNRI all of D-Lib Magazine’s issues are still available on the live web, with no changes in their URIs since the fourth issue (October, 1995)¹². Although we have long known “Cool URIs Don’t Change” [13], the reality is that most do (e.g., [55]), and persisting over 5,000 URIs¹³ for more than 25 years is an accomplishment in itself.

Ariadne has changed publishers a few times, and although the domain has always been ariadne.ac.uk, the URLs within the domain have changed over time. First Monday changed URIs (from firstmonday.dk to simultaneouslyfirstmonday.org and a path within journals.uic.edu) as it changed publishers. While JoDI was active, it transitioned from the University of Southampton (jodi.ecs.soton.ac.uk and journals.ecs.soton.ac.uk (the former no longer resolves) to Texas A&M University and Texas Digital Library (journals.tdl.org).

3.6 Persistent identification

Although D-Lib Magazine was published as a conventional serial, it also embraced persistent identifiers (initially handles and later DOIs) for individual articles (owing from the computer science technical report heritage of CNRI’s technology), which facilitates the disaggregation of serials into articles that are directly and persistently identifiable, which reinforces them as being “on the Web” as first-class citizens.

Handles [63] were part of the technical infrastructure for digital libraries built by CNRI. Although frequently considered a URN implementation [102], they are not registered as URN namespaces¹⁴ and their status as URIs remains unresolved [113]. D-Lib Magazine used handles beginning with the first issue, and, as the DOI effort matured, D-Lib Magazine was naturally an early adopter, starting in January, 1999 [129]. DOIs are a proper subset of handles, so all DOIs are handles (e.g., the DOI [doi: 10.1045/july95-arms](https://doi.org/10.1045/july95-arms) identifies the same resource as the handle [hdl:cnri.dlib/july95-arms](http://hdl.handle.net/1045/july95-arms)), but not all handles are DOIs.

Beginning in 2013, D-Lib Magazine adopted the Crossref guidelines for displaying DOIs in references [20], displaying them as full HTTP URLs, instead of the neither URN nor URL approach of, for example, [doi: 10.1045/june2001-reich](https://doi.org/10.1045/june2001-reich).

Of its peers, only First Monday uses DOIs (we believe it adopted them in 2013). The DOIs for First Monday resolve to the landing pages, whereas the DOIs for D-Lib Magazine resolve to the article itself. This subtle distinction has implications: having a DOI resolve to the article itself, especially for an article in HTML, reinforces that the article is a first-class object on the web, and not just a PDF or other non-native web format that simply uses the web as for transport. As serials embrace HTML-centric approaches (e.g., BioMed Central, PLOS One, and PeerJ) having DOIs resolve to articles instead of landing pages is becoming more common.

¹²The first three issues were published at <http://www.cnri.reston.va.us/home/dlib.html> (cf. <https://www3.wcl.american.edu/cni/9507/6207.html>), and it was not until the October, 1995 issue that www.dlib.org was adopted (“Please note that D-Lib has a new address: <http://www.dlib.org>” – <http://www.dlib.org/dlib/october95/10contents.html>).

¹³<https://www.google.com/search?q=site:dlib.org>

¹⁴<https://www.iana.org/assignments/urn-namespaces/urn-namespaces.xhtml>

3.7 Metadata

Given the community in which it originated, it should come as no surprise that D-Lib Magazine paid attention to the provision of metadata about its articles. It actually did so in ways that demonstrate an understanding of the ways of the web, specifically that information consumers on the web include machines. Starting in January 1999, the HTML META tag is used to convey the DOI of articles as a string, not HTTP URI. Additionally, starting in March 1999, the HTML LINK element is used with a “metadata” relation type to point at descriptive metadata for an article:

```
<LINK REL="metadata" HREF="03maly.meta.xml">
<meta NAME="DOI" CONTENT="10.1045/march99-maly">
```

The metadata uses the Dublin Core vocabulary and is expressed in XML¹⁵ that adheres to a DTD¹⁶ defined by D-Lib Magazine. Along with the Dublin Core best practice work by the Consortium for the Computer Interchange of Museum Information [123] around the same time, this DTD was one of the early efforts to formally express syntactic constraints for Dublin Core serializations. The Dublin Core metadata shows the DOI, again, expressed as a string but also includes the HTTP URI of the article. The choice to link to a Dublin Core metadata document is interesting in light of the efforts around the same time to support embedding the metadata in HTML, which led to the publication of an RFC [57] in December 1999. Most likely the D-Lib Magazine editors were aware of this work but, maybe, their library backgrounds led them to choose an approach that treated metadata as a stand-alone, first class citizen on the web. First Monday made another choice and started embedding Dublin Core metadata in HTML with the first issue of 1999.¹⁷ Going by web archived articles, by the early 2000s, neither JoDI¹⁸ nor Ariadne¹⁹ were providing descriptive metadata for machine consumption.

3.8 Mirror sites

Another innovation D-Lib Magazine embraced was the use of site mirrors allowing users in Europe and Asia to interact with geographically closer mirrors for faster response. That approach to address bandwidth limitations was common at the time and is now solved via content delivery networks (CDNs). Three of the D-Lib mirrors are still functioning, down from a peak of five.²⁰ In addition to the utility the mirrors provide, they were also presumably intended as demonstrators for more advanced Handle resolution techniques, such as being able to resolve to one of multiple URLs [94]. To the best of our knowledge, none of D-Lib’s peers ever employed mirror sites.

4 D-LIB MAGAZINE: THE FIRST ISSUE

Conveying the scope and impact of the entire corpus of D-Lib Magazine is beyond what we can do here. Instead, we can examine

¹⁵For example: <http://www.dlib.org/dlib/march99/maly/03maly.meta.xml>

¹⁶<http://www.dlib.org/dlib/dlib-meta01.dtd>

¹⁷For example: https://web.archive.org/web/20000817043514/http://www.firstmonday.dk/issues/issue4_1/slowinski/index.html

¹⁸For example, from 2004: <http://web.archive.org/web/20040803111530/http://jodi.ecs.soton.ac.uk/Articles/v05/i01/Brown/>

¹⁹Forexample, from 2001: <http://web.archive.org/web/20010622191945/http://www.ariadne.ac.uk/issue27/johnston/>

²⁰<http://web.archive.org/web/20150224045836/mirror.dlib.org/about.html>

the first issue and what it foresaw in both digital libraries and the web in general.

The first issue had 14 “Clips and Pointers” – announcements, deadline reminders, calls for participation, requests for proposals and papers, and brief updates. Although email lists served these functions (and still continue to do so), this announcement and awareness function of a magazine has largely been replaced by blogs and social media. One no longer expects to learn of calls for papers or requests for proposals in a magazine, and event summaries are now easily discoverable via search engines with far more precision than those of the mid-1990s (e.g., Lycos [74]). For example, our conference report for the 2003 Joint Conference on Digital Libraries (JCDL) was published in D-Lib Magazine [84], but JCDL 2020 is best reviewed in blogs [46] or Twitter.²¹

The first issue had three articles, then carried under the heading of “stories and briefings”, reflecting the early position of a “magazine” and not an online journal. In fact, they were summaries of existing conventional reports and publications, and presumably not intended as publications in their own right:

- (1) “Metadata: the foundations of resource description” – a summary of the OCLC/NCSA Metadata Workshop [127] that produced the Dublin Core Metadata Element Set, which continues today as the Dublin Core Metadata Initiative (dublin-core.org).
- (2) “An agent-based architecture for digital libraries” – a description of the distributed agent architecture explored in the University of Michigan Digital Library (UMDL) [11]; the University of Michigan was one of six participants in the first NSF/DARPA/NASA Digital Library Initiative (DLI).
- (3) “Key concepts in the architecture of the digital library” – an introduction to and contextualization of what would become known as the “Kahn-Wilensky Framework” (KWF) [49], part of which included handles [63], upon which Digital Object Identifiers [93] are implemented.

As tentative steps in this new publishing experiment, all three articles are single authored (though they summarize multi-author publications), are relatively short, and have limited figures and references. Although D-Lib Magazine would soon evolve into a venue where original research was published (e.g., a 1999 editorial estimates that half of the contributions described original research [8]) and essentially functioned as an online journal, it was edited and never refereed. This produced a well-known problem: if you wanted your material to reach a wide audience, it needed to be in D-Lib Magazine, but if you wanted academic “credit”, it needed to be in a conventional journal or refereed conference proceedings. In the time before Google, Google Scholar, CiteSeer, Microsoft Academic et al., this was a binary choice. Now it is possible for authors to gain the imprimatur of a quality journal or conference proceedings, and at the same time leverage the permissive attitude regarding pre-prints and e-prints of many publishers (e.g., ACM) to ensure that articles are discoverable and freely available.

The final portion of the first issue that we will explore is a forward looking piece in which the editors solicited feedback from five prominent figures in the then emerging DL community about setting the research agenda for the field. With the vantage point

²¹https://twitter.com/search?q=%23JCDL2020&src=typed_query&f=live

of time, we assess how well the community has achieved their collective vision from 1995.

4.1 Dublin Core

The first article in the first issue of D-Lib Magazine, “Metadata: The Foundations of Resource Description” [126], is a summary of the OCLC/NCSA Metadata Workshop Report, which resulted from the workshop in Dublin, Ohio, only four months prior (March, 1995) [127]. The Dublin Core Metadata Element Set (DCMES, or “Dublin Core”) was still forming at this point, with only 13 metadata elements, not the final 15, defined, and “DCMES” becoming the Dublin Core Metadata Initiative (DCMI) Terms. While the DCMI has gone on to issue over 70 specifications,²² today’s DCMI Terms can trace their origin to the 1995 Metadata Workshop and the original DCMES. The impact of Dublin Core is far beyond what we can cover here, but a search for “dublin core” in Google currently yields over 11M hits, and a similar search in Google Scholar yields over 98K hits.

Dublin Core would form its own community, complete with its own governance and document series. But D-Lib Magazine would continue to be a venue for conveying the status of Dublin Core [24, 58, 115, 128], and other related Web metadata efforts, such as PICS [80] and its progeny, RDF [79], and IEEE LOM [27].

While Dublin Core is abundantly used for the description of assets in a variety of content management systems,²³ continues to this day to play a role in web-based discovery, co-existing with similar formats such as the Open Graph Protocol [42] and Schema.org [39], yet facing some significant competition from the latter when it comes to Search Engine Optimization [45].

4.2 DLI and DLI2

The second article, “An Agent-Based Architecture for Digital Libraries” [15], is a high-level summary of the University of Michigan Digital Library (UMDL) project, one of the original six NSF/DARPA/NASA Digital Library Initiative (DLI) projects. The DLI ran from 1994–1998, so the 1995 article only summarizes the earliest results.

The architectural details of the UMDL are academically interesting, but the real value in 2020 is reading the article as a time capsule of 1990s perception of the Web, DLs, and DL architecture. A quote from near the beginning of the article describes a scenario that we have since seen come to pass:

The WWW, while it probably contains more information than any single traditional library, is arguably not as useful as a traditional library because it lacks these services (particularly organization and sophisticated search support). No one is dismantling their libraries because of the WWW yet.

The envisioned architecture focuses heavily on agents, which navigate a distributed, heterogeneous tapestry of distributed repositories on behalf of the user. The model of distributed search was dominant in early DL architecture thinking, and was reflected in the design of search protocols like Z39.50 and WAIS, as well as DLs such as WATERS [73], NCSTRL [22], NTRS [88], and many other

examples.²⁴ The DL commitment to distributed searching on the Web culminated in the STARTS protocol [35], and dissatisfaction with the state of distributed searching DLs (cf. [85, 99]) was at the heart of the Universal Preprint Service prototype that demonstrated metadata harvesting and centralized searching, a design decision that would inform OAI-PMH [62, 117, 118] and later DLs based on it (e.g., [3, 71, 86]).

Typical of the time, the UMDL design is fully committed to distributed search and crawling, with personalized agents handling the foraging and negotiation with the various repositories (similar to CNRI’s Knowbots [53]). After more than 25 years and with a post-Google perspective, we can now see that most meta-search / distributed search architectures have been retired, in part by the hegemony of Google-style crawling and searching²⁵. In the end, (logically) centralized architectures won. HTTP servers were (and are) broadly distributed, but the complexity of crawling and indexing turned out to be centralized. The search engines now dictate to the web servers how to expose and structure their site, instead of the anticipated model where sites instructed the best way to access their holdings.

The DLI ran from 1994 through 1998, and its \$30M supported six projects,²⁶ each of which is summarized in the one year anniversary issue (July/August 1996) of D-Lib Magazine.²⁷ DLI2 ran from 1999–2005, and the \$55M from the NSF, DARPA, National Library of Medicine, Library of Congress, NASA, and the National Endowment for the Humanities (with additional participation from the National Archives and the Smithsonian Institution) supported 36 projects. Despite the importance of the DLI and DLI2 funding efforts in the early days of the web, very little about the funding programs remains on the live web outside of what is hosted at dlib.org. Sites like www.dli2.nsf.gov and www.cise.nsf.gov/iis/dli_home.html are no longer on the live web, with only a single page left at nsf.gov to mark 12 years of research and \$85M in funding²⁸. Although the former pages are accessible in web archives²⁹ and the individual, specific technical contributions resulting from the DLI work are widely described in the broader literature, D-Lib Magazine was a prime venue for program-level reflection [9, 36–38, 44, 67, 70, 91].

4.3 KWF and DOIs

Digital Object Identifiers (DOIs) are the most visible contribution from the “Kahn-Wilensky Framework”, discussed in an article from the first issue of D-Lib Magazine entitled “Key Concepts in the Architecture of the Digital Library” [7]. It was a summary by Bill Arms of “A framework for distributed digital object services”, which would later be known as the “Kahn-Wilensky Framework” (KWF) [49]. Although it is just an abstract framework and not tied to a specific implementation, the KWF has had a significant impact on

²⁴Bill Arms indirectly notes that such architectural decisions trace back to 1991 [10], see also Bill Mischo’s reflections on federation [81].

²⁵Unified Computer Science Technical Report Index (UCSTRI) [43], a computer science DL that independently crawled and indexed anonymous FTP sites is the first known example of the architecture that successful DLs like CiteSeer [33] and Google Scholar [124] would employ.

²⁶https://web.archive.org/web/19981202064413/http://www.cise.nsf.gov/iis/dli_home.html

²⁷<http://hdl.handle.net/cnri.dlib/july96>

²⁸https://www.nsf.gov/news/news_summ.jsp?cntn_id=103048

²⁹For example, https://web.archive.org/web/*/http://www.cise.nsf.gov/iis/dli_home.html and https://web.archive.org/web/*/http://www.dli2.nsf.gov/projects.html

²²<https://www.dublincore.org/specifications/dublin-core/>

²³<https://lov.linkeddata.es/dataset/lov/>

the architectural design of digital libraries, especially concerning the identification and structure of “digital objects” and their relationship with the repositories in which they reside. In 2006, we edited a special issue of the International Journal on Digital Libraries (IJDL) on “Complex Digital Objects” [89], which also featured a reprint and of the KWF along with commentary from Robert Kahn [50].

The KWF provided the architecture for the initial CS-TR project [51], which in combination with WATERS [73] formed the basis for the Dienst protocol and NCSTRL [22] as a distributed digital library for computer science technical reports. Lessons learned from Dienst were incorporated in OAI-PMH [62, 117, 118]. The KWF also had an impact in the Dublin Core community, resulting in the Warwick Framework [60], which was later extended with “distributed active relationships” [21], which itself later evolved into Fedora [97]. The management of Fedora and DSpace [114] were merged into Duraspace in 2009 [82], and in 2019 LYRASIS absorbed Duraspace. The separate Fedora and DSpace open source products continue to be offered.

KWF specified the role of repositories in mediating access to their digital objects via the Repository Access Protocol (RAP). Over the years, numerous papers have been published, many of them in D-Lib Magazine, that pertain to RAP, including [6, 96, 101]. The design of RAP was repository-centric [119] and explicitly decoupled the protocol for expressing interactions with digital objects from the transport protocol used to transfer interaction requests between client and server. Such a choice was not uncommon in the days preceding the dominance of the web and its now omnipresent HTTP protocol and can, for example, also be observed in the design of OAI-PMH [119]. Many transport protocols (TCP, SMTP, FTP, Gopher, HTTP, IIOP, etc.) overlapped in time and there was a predisposition to viewing them as impermanent, interchangeable; something on which one built richer, domain-specific protocols. Also, Roy Fielding did not publish his dissertation about Representational State Transfer (REST), which made the resource-centric [119] semantics and potential of HTTP explicit, until 2000 [28, 30]. By that time, the state of thinking and practice in digital libraries had already diverged from that of the web. In this way, RAP was the initial manifestation of an architectural fault line between digital libraries and the web that continues to this day.³⁰ As a prominent example, the FAIR Digital Object effort [23] that fits under the broad umbrella of the European Open Science Cloud³¹ program and is supported by activities of the Research Data Alliance,³² considers two approaches³³ to devise rich interactions with digital objects that are stored in cooperating repositories. One aligns with the RAP line of thought [25]. The other embraces a webby approach and advocates leveraging a range of HTTP-based standards that have become available over the years, including the Open Archives Initiative Object Reuse and Exchange (OAI-ORE) [61] for the representation of digital objects as aggregations of web resources, the Memento Framework [120, 121] for temporal versioning of web resources, Linked Data Platform [112] and the Fedora API [5] for

CRUD operations on digital objects and their constituent resources, Web Annotation [103, 105], RO-Crate for packaging digital objects [108]. It is interesting to note that several of these specifications were co-authored by people with roots in the Digital Library community.

4.4 To the editor: What’s needed in future research?

The last part of the first issue we would like to review is a page entitled “To the editor: What’s needed in future research?”³⁴, in which the D-Lib editors polled five prominent digital library researchers and administrators and asked them to briefly identify and discuss areas that warranted further research. Although the purpose was to generate discussion about a near-term research agenda (and perhaps establish D-Lib Magazine’s credentials through adding additional voices to the first issue), this page now serves as a time capsule and allows us to reassess progress in the field since 1995. In the page the editors stated “Now, we would like to know what you think; send [us] your thoughts, reactions, and comments”, which we now do 25+ years later.

“An interoperable national and global information web” – Barry M. Leiner, Deputy Director, ARPA/CSTO. The world-wide web technology, infrastructure, and protocols that Barry Leiner credited for an explosion in the availability of accessible and viewable information have persisted and have given rise to a global networked environment that we can hardly imagine living without anymore. Over time, numerous open standards have been specified to support interoperability beyond the basic level provided by the core ingredients of the web, HTTP and HTML. In the Web 2.0 era, some of these acted as catalysts for the frictionless creation of value-added services across web platforms, both for and not for profit. But since interoperability is not a significant concern for companies that want to protect their turf or establish monopolies, a trend has emerged to support rich access by means of bespoke APIs rather than open protocols, significantly increasing the investment for the development and management of services that require cross-platform interactions. This trend has become so prominent that platforms routinely claim to be interoperable because they expose a self-defined API, and, to add insult to injury, touting its RESTful-ness while many times it is not [29].

In the digital library community, the dream to achieve interoperability based on open standards remains alive and actively pursued. Despite the aforementioned ongoing debate regarding which path to take – repository-centric or resource-centric – most community driven specification efforts of the past decade have chosen to embrace the ways of the web, many times aiming for approaches that have applicability beyond the digital library community. In the realm of technologies to support digital libraries of multimedia information on which Barry Leiner zoomed in, prominent examples include the Fedora API [5] that leverages the W3C Linked Data Platform recommendation [112], the W3C Web Annotation recommendations [103, 105], and the specifications that resort under the International Image Interoperability Framework³⁵. Because of

³⁰An excellent review of the complex relationship between DLs and the Web is Carl Lagoze’s 2010 dissertation, “Lost Identity: The Assimilation Of Digital Libraries Into The Web” [59].

³¹<https://www.eosc-portal.eu/>

³²<https://www.rd-alliance.org/>

³³<https://github.com/GEDE-RDA-Europe/GEDE/tree/master/FAIRDigitalObjects/FDOF>

³⁴<http://www.dlib.org/dlib/July95/07messages.html>

³⁵<https://iiif.io/technical-details/>

its growing global adoption by GLAM institutions, especially the latter stands as a testimony that rich interoperability for distributed resource collections is effectively achievable. But other promising specifications that aim for the same holy grail are struggling for adoption, and, many times, lack of resources is mentioned as a reason. While that undoubtedly plays a role, it did not stand in the way of rapid adoption of protocols that have emerged from large corporations, such as the Google-dominated [4] schema.org³⁶. This consideration re-emphasizes that a core ingredient of a successful interoperability specification, and hence of achieving an interoperable global information web, is a large megaphone, either in the guise of commercial power or active community engagement [104].

“Integration between electronic and non-electronic forms of communications and publications” – Ann L. Okerson, Director, Office of Scientific & Academic Publishing, Association of Research Libraries (ARL). The dichotomy between print and digital resources that Ann Okerson describes was a major concern in the mid nineties. Despite large-scale digitization efforts (see below) and the exponential growth of born-digital materials, analog collections will remain. But these worlds are no longer perceived as being radically distinct because the analog world has largely been absorbed by the digital one. This did not necessarily happen by making both discoverable through library OPACs. Rather, it has become commonplace to cater to the crawl-driven discovery paradigm of major search engines by exposing resource descriptions for materials of both types of collections to the web using Search Engine Optimization techniques such as the Sitemap protocol and, more recently, schema.org. For many analog GLAM collections doing so requires making traditional catalog systems web savvy, a task that is still ongoing. So while “electronic and non-electronic forms of communications and publications” will remain parallel worlds for the foreseeable future, the percentages have shifted, as Lannom’s note from the final D-Lib editorial makes clear: “the very phrase ‘Digital Libraries’ has gone from sounding innovative to sounding a bit redundant.” [66]

“Foreground’ information stores, or personal digital libraries” – William L. Scherlis, Senior Research Computer Scientist, Department of Computer Science, Carnegie Mellon University. The personal digital libraries envisioned in this piece have not become mainstream. Instead creators have embraced a myriad of web productivity portals to share both their intellectual artifacts and daily activities. As a result, assets created by individuals and information pertaining to their comings and goings are distributed across the web, to such an extent that both research (e.g., [40]) and development³⁷ efforts have considered approaches to aggregate it into a personal environment that provides a concise representation of the self on the distributed web. In the realm of scholarly communication, the experimental myresearch.institute³⁸ effort tracks, collects, and archives assets created by researchers in a variety of web portals, including GitHub, Slideshare, Wikipedia. The plethora of APIs used by these portals and the lack of support for protocols such as W3C ActivityPub [125] and W3C ActivityStreams2 [111] make such an aggregation task far from trivial. The result of gathering the distributed information could be considered a proxy personal digital library. But

³⁶<https://schema.org>

³⁷for example: <https://github.com/LockerProject/Locker>

³⁸<https://myresearch.institute>

maybe the days of the actual personal digital library are still to come. Motivated and frustrated by the monopolies certain web portals have established over the years, and the concerns regarding privacy and data abuse that result, the Decentralized Web movement is aiming for alternatives, with a focus on giving individuals back control over their personal assets. As part of this movement, the Solid³⁹ effort led by Tim Berners-Lee introduces the notion of a pod [14], a personal storage space that complies with a stack of open standards and allows its owner to grant or deny applications and users access rights to stored resources. Clearly, these pods are conceived as a technology that can help information producers manage foreground information, to put it in the words of William L. Scherlis.

“Diversifying and access” and “the distribution chain” – Paul Evan Peters, Executive Director, Coalition for Networked Information. As far as “diversifying and access”, we believe Peters is criticizing researchers’ emphasis on the “attributes of the resources and services”, as manifested in Leiner’s assessment, as well as the prevalent distributed searching paradigm described in section 4.3, and instead we should be focusing on the “attributes of the users and uses”, which we interpret to be in harmony with Scherlis’s vision of personal DLs. Most siloed repositories have been flattened and exposed for crawling by search engines. Google, for example, showed little appetite for even simple protocols such as OAI-PMH [41, 77], and officially retired support for it in 2008 [83].

Initially we struggled with understanding Peters’s description of “the distribution chain.” Eventually we decided that part of the ambiguity is that he is describing something that is so common now but for which the language did not exist in 1995, resulting in a terminology gap that we had to bridge to understand what he meant. When he says we need to focus on “closing the gap between creators and users of resources and services”, we understand that to be an admonishment that the point of DLs should not be merely the automation of the existing publication process. Unfortunately, 2020 still resembles 1995, with the distribution chain largely paralleling the value chain, just now with PDFs instead of paper. Instead of reenvisioning / reengineering the scholarly communication process (e.g., [122]), we have a confusing array of open access options (“gold”, “green”, “hybrid”, etc.⁴⁰), and by retaining the publisher at the center of distribution, they still fail to address the broader needs of scholarly communication (e.g., [18]).

We have now long had the ability to “link creators and users” as Peters’s calls for, but have lacked the collective will to make the transition [19, 116]. Web 2.0, blogs, and social media provided some hope initially, but as noted in section 4.3, platforms have since moved away from Atom and RSS in favor of bespoke APIs that provide more functionality at the expense of interoperability. Add to this Elsevier buying various platforms that “link creators and users” (e.g., SSRN [98], Mendeley [72], bepress [107], and Peters’s vision, while technically feasible, seems no closer than it did in 1995.

“Retrospective capture of content” – James Michalko, President, The Research Libraries Group, Inc. Michalko, in a statement that aligns

³⁹<https://inrupt.com/solid>

⁴⁰https://en.wikipedia.org/wiki/Open_access

with Okerson's, observes "[t]here's a major opportunity and demand for the retrospective capture of content", but in 1995 "[t]here are few service bureaus that can do the scanning and capture of maps, manuscripts, and other primary research materials." However, Michalko was writing at a time when these projects either did not exist or were just beginning: JSTOR [106], Google Books [68], Open Content Alliance [48], Internet Archive [54], HathiTrust [133], National Digital Newspaper Program [2], museum mass digitization projects [100], etc. Mass digitization of primary research materials remains incomplete, and it is not always clear how the digitized models will fit within a search engine-centric model of crawling and searching. But the momentum is there, and what is digitized exists at a scale that we could only dream of in 1995.

5 CONCLUSIONS

The web began with a vision of freeing academic information, and in the earliest days the paths of both digital libraries and the web were highly intertwined. D-Lib Magazine was published from 1995 through 2017. During this time, it helped shape the digital library community, via the information published in articles, the ancillary awareness and informational updates now largely provided by social media, and as an ongoing experiment in web(-only) publishing. Although it ceased publication in 2017, the entire site is still on the live web as an unchanged time capsule, and as a serial it still accrues many citations. The 1,062 articles and 5,000+ web pages available at www.dlib.org offer many opportunities for reflection about the DL community, but we took the first issue as our point of reference.

Of the three articles in the first issue, all were summaries of work described elsewhere. However, only the article about the DLI-funded UMDL project summarized information in conventional, peer-reviewed publications. The other two articles, about Dublin Core and the Kahn-Wilensky Framework, summarized "unpublished" (i.e., grey literature) reports, providing a more formal and citable surrogate for standards, practitioner, engineering work that was crucial in the early days of DLs, which the conventional, peer-reviewed publication venues would largely ignore.

D-Lib Magazine would cease publishing due in part to an unsustainable funding model, the maturity of the field, and the rise of blogs and social media. However, its role in shaping the then emerging DL community is hard to overstate, both in terms of the research and technology it described in its article, but also embracing and experimenting with the web, HTML, and hypermedia in general in a way that that was unmatched by its peers.

Given the perspective of more than 25 years, one would be hard pressed to retroactively construct a more prescient first issue: Dublin Core is an ongoing initiative and suite of standards that DLs and the general web still employ today, the \$30M from the DLI bootstrapped the DL community and eventually gave us Google, and the Kahn-Wilensky Framework have influenced the design of repositories (e.g., Fedora), interoperability (e.g., OAI-PMH), and provided the proof-of-concept to help launch the DOI ecosystem that provides a fundamental level of interoperability across scholarly publishing.

Furthermore, the open-access, HTML-centric scholarly serials such as BioMed Central, PLOS One, and PeerJ owe an intellectual

debt to the pioneering work of D-Lib Magazine, its editors, and its authors.

ACKNOWLEDGMENTS

We would like to thank everyone who contributed to D-Lib Magazine's success, first of which is Robert Kahn, who applied the resources of CNRI (and DARPA) to make D-Lib Magazine possible. We appreciate the work of all the editors through the years: Bill Arms, Amy Friedlander, Barry Leiner, Bonnie Wilson, Peter Hirtle, Allison Powell, Larry Lannom, and Cathy Rey.

We are grateful that D-Lib Magazine existed during the maturation of the field of digital libraries as well as during the pivotal points in our respective careers.

REFERENCES

- [1] P. Aas, S. Dixit, T. Eden, S. Faulkner, B. Lawson, S. Moon, and X. Wu. HTML 5.3, 2018. W3C Working Draft, 18 October 2018. URL: <https://www.w3.org/TR/html53/>.
- [2] H. Agu era. The US National Digital Newspaper Program: Thinking ahead, designing now part I: Program development. *IFLA Publications*, 122:79, 2006.
- [3] H. Anan, X. Liu, K. Maly, M. Nelson, M. Zubair, J. C. French, E. Fox, and P. Shivakumar. Preservation and transition of NCSTRL using an OAI-based architecture. In *JCDL '02: Proceedings of the 2nd ACM/IEEE-CS Joint Conference on Digital Libraries*, pages 181–182, 2002. URL: <https://doi.org/10.1145/544220.544256>.
- [4] M. Andrews. Time to end Google's domination of schema.org, 2020. URL: <https://storyneedle.com/time-to-end-googles-domination-of-schema-org/>.
- [5] B. Armintor, E. C. D. Lamb, S. Warner, and A. Woods. Fedora API Specification 1.0, 2018. Candidate Recommendation 21 November 2018. URL: <https://fcrepo.github.io/fcrepo-specification/>.
- [6] W. Arms, C. Bianchi, and E. Overly. An Architecture for Information in Digital Libraries. *D-Lib Magazine*, 3(2), 1997. URL: <http://hdl.handle.net/cnri.dlib/february97-arms>.
- [7] W. Y. Arms. Key concepts in the architecture of the digital library. *D-Lib Magazine*, 1(1), 1995. URL: <http://www.dlib.org/dlib/July95/07arms.html>.
- [8] W. Y. Arms. From the Publisher: Scholarly Communication, Digital Libraries, and D-Lib Magazine. *D-Lib Magazine*, 5(4), 1999. URL: <https://doi.org/10.1045/april99-editorial>.
- [9] W. Y. Arms. Editorial: The European Connection. *D-Lib Magazine*, 6(7/8), 2000. URL: <https://doi.org/10.1045/july2000-editorial>.
- [10] W. Y. Arms. A viewpoint analysis of the digital library. *D-Lib Magazine*, 11(7/8), 2005. URL: <https://doi.org/10.1045/july2005-arms>.
- [11] D. E. Atkins. The University of Michigan Digital Library Project: The Testbed. *D-Lib Magazine*, 2(7/8), 1996. URL: <https://hdl.handle.net/cnri.dlib/july96-atkins>.
- [12] T. Berners-Lee. WorldWideWeb - executive summary. <https://archive.md/Lfoppj>, 1991.
- [13] T. Berners-Lee. Cool URIs don't change, 1998. URL: <http://www.w3.org/Provider/Style/URI.html>.
- [14] T. Berners-Lee. Socially aware cloud storage, 2009. URL: <https://www.w3.org/DesignIssues/CloudStorage.html>.
- [15] W. P. Birmingham. An agent-based architecture for digital libraries. *D-Lib Magazine*, 1(1), 1995. URL: <https://hdl.handle.net/cnri.dlib/july95-birmingham>.
- [16] D. A. Black. Introducing HyperNews. *Linux Journal*, (27), 1996. URL: <https://www.linuxjournal.com/article/1281>.
- [17] J. Bollen, M. L. Nelson, G. Manepalli, G. Nandigam, and S. Manepalli. Trend Analysis of the Digital Library Community. *D-Lib Magazine*, 11(1), 2005. URL: <https://doi.org/10.1045/january2005-bollen>.
- [18] B. Bremb. Scholarship has bigger fish to fry than access, 2019. URL: <http://bjoern.brembs.net/2019/10/scholarship-has-bigger-fish-to-fry-than-access/>.
- [19] S. Capadisli. *Linked Research on the Decentralised Web*. PhD thesis, University of Bonn, 2020. URL: <https://csarven.ca/linked-research-decentralised-web>.
- [20] Crossref Revises DOI Display Guidelines. Crossref. <https://www.crossref.org/news/2011-08-02-crossref-revises-doi-display-guidelines/>, 2011.
- [21] R. Daniel, Jr and C. Lagoze. Extending the Warwick Framework. *D-Lib Magazine*, 3, 1997. URL: <http://hdl.handle.net/cnri.dlib/november97-daniel>.
- [22] J. R. Davis and C. Lagoze. NCSTRL: Design and deployment of a globally distributed digital library. *Journal of the American Society for Information Science*, 51(3):273–280, 2000. URL: [https://doi.org/10.1002/\(SICI\)1097-4571\(2000\)51:3%3C273::AID-ASIE%3E3.0.CO;2-6](https://doi.org/10.1002/(SICI)1097-4571(2000)51:3%3C273::AID-ASIE%3E3.0.CO;2-6).
- [23] K. De Smedt, D. Koureas, and P. Wittenburg. FAIR digital objects for science: From data pieces to actionable knowledge units. *Publications*, 8(2), 2020. URL: <https://doi.org/10.3390/publications8020021>.

- [24] L. Dempsey and S. L. Weibel. The Warwick Metadata Workshop: A Framework for the Deployment of Resource Description. *D-Lib Magazine*, 2(7/8), 1996. URL: <https://hdl.handle.net/cnri.dlib/january96-weibel>.
- [25] DONA Foundation. Digital Object Interface Protocol Specification Version 2.0, 2018. URL: https://www.dona.net/sites/default/files/2018-11/DOIPv2Spec_1.pdf.
- [26] J. J. Dongarra and E. Grosse. Distribution of mathematical software via electronic mail. *Communications of the ACM*, 30(5):403–407, 1987. URL: <https://doi.org/10.1145/22899.22904>.
- [27] E. Duval, W. Hodgins, S. Sutton, and S. L. Weibel. Metadata Principles and Practicalities. *D-Lib Magazine*, 8(4), 2002. URL: <https://doi.org/10.1045/april2002-weibel>.
- [28] R. T. Fielding. *Architectural Styles and the Design of Network-based Software Architectures*. PhD thesis, University of California, Irvine Department of Computer Science, 2000. URL: <http://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm>.
- [29] R. T. Fielding. REST APIs must be hypertext-driven, 2008. URL: <https://roy.gbiv.com/untangled/2008/rest-apis-must-be-hypertext-driven>.
- [30] R. T. Fielding and R. N. Taylor. Principled design of the modern web architecture. *ACM Transactions on Internet Technology (TOIT)*, 2(2):115–150, 2002. URL: <https://doi.org/10.1145/337180.337228>.
- [31] A. Friedlander. From the Editor: A Word (or Two) of Welcome. *D-Lib Magazine*, 1(1), 1995. URL: <https://hdl.handle.net/cnri.dlib/july95-friedlander>.
- [32] A. Friedlander. From the Editor: Where Are We in Space? *D-Lib Magazine*, 1(4), 1995. URL: <http://hdl.handle.net/cnri.dlib/october95-friedlander>.
- [33] C. L. Giles, K. D. Bollacker, and S. Lawrence. CiteSeer: An automatic citation indexing system. In *Proceedings of the third ACM conference on Digital libraries*, pages 89–98, 1998. URL: <https://doi.org/10.1145/276675.276685>.
- [34] P. Ginsparg. First steps towards electronic research communication. *Computers in Physics*, 8(4):390–396, 1994. URL: <https://doi.org/10.1063/1.4823313>.
- [35] L. Gravano, C. Chang, H. Garcia-Molina, and A. Paepcke. STARTS: Stanford proposal for Internet meta-searching. In *Proceedings of the 1997 ACM SIGMOD international conference on Management of data*, pages 207–218, 1997. URL: <https://doi.org/10.1145/253260.253299>.
- [36] S. M. Griffin. NSF/DARPA/NASA Digital Libraries Initiative: A Program Manager's Perspective. *D-Lib Magazine*, 4(7/8), 1998. URL: <https://hdl.handle.net/cnri.dlib/july98-griffin>.
- [37] S. M. Griffin. Digital Libraries Initiative - Phase 2: Fiscal Year 1999 Awards. *D-Lib Magazine*, 5(7/8), 1999. URL: <https://doi.org/10.1045/july99-griffin>.
- [38] S. M. Griffin. Funding for Digital Libraries Research: Past and Present. *D-Lib Magazine*, 11(7/8), 2005. URL: <https://doi.org/10.1045/july2005-griffin>.
- [39] R. V. Guha, D. Brickley, and S. Macheth. Schema.org: evolution of structured data on the web. *Communications of the ACM*, 59(2):44–51, 2016. URL: <https://doi.org/10.1145/2857274.2857276>.
- [40] A. Guy. *The Presentation of Self on a Decentralised Web*. PhD thesis, University of Edinburgh, 2017. URL: <http://hdl.handle.net/1842/29537>.
- [41] K. Hagedorn and J. Santelli. Google Still Not Indexing Hidden Web URLs. *D-Lib Magazine*, 14(7/8), 2008. URL: <https://doi.org/10.1045/july2008-hagedorn>.
- [42] A. Haugen. The Open Graph Protocol design decisions. In *International Semantic Web Conference*, pages 338–338. Springer, 2010. URL: https://doi.org/10.1007/978-3-642-17749-1_25.
- [43] M. V. Heyning. The unified computer science technical report index: Lessons in indexing diverse resources. In *Proceedings of the Second International World Wide Web Conference*, pages 535–543, 1994. URL: <http://archive.ncsa.uiuc.edu/SDG/IT94/Proceedings/Day/vanheyning/paper.html>.
- [44] P. B. Hirtle. Editorial: A New Generation of Digital Library Research. *D-Lib Magazine*, 5(7/8), 1999. URL: <https://doi.org/10.1045/july99-editorial>.
- [45] howard-go. Dublin Core vs Schema.org: A Head-To-Head Metadata Comparison, 2015. URL: <https://seopressor.com/blog/dublin-core-vs-schemaorg-metadata-comparison/>.
- [46] Y. Jayawardana and G. Jayawardena. Joint Conference on Digital Libraries (JCDL) 2020 Trip Report, 2020. URL: <https://ws-dl.blogspot.com/2020/08/2020-08-16-joint-conference-on-digital.html>.
- [47] S. M. Jones, H. Van de Sompel, H. Shankar, M. Klein, R. Tobin, and C. Grover. Scholarly context adrift: three out of four URI references lead to changed content. *PLoS one*, 11(12):e0167475, 2016. URL: <https://doi.org/10.1371/journal.pone.0171057>.
- [48] B. Kahle. Announcing the Open Content Alliance, 2005. URL: <http://www.ysearchblog.com/archives/000192.html>.
- [49] R. Kahn and R. Wilensky. A framework for distributed digital object services. Technical Report cnri.dlib/tn95-1, CNRI, 1995. URL: <http://www.cnri.reston.va.us/home/cstr/arch/k-w.html>.
- [50] R. Kahn and R. Wilensky. A framework for distributed digital object services. *International Journal on Digital Libraries*, 6(2):115–123, 2006. URL: <https://doi.org/10.1007/s00799-005-0128-x>.
- [51] R. E. Kahn. An Introduction to the CS-TR Project, 1995. URL: <http://www.cnri.reston.va.us/home/describe.html>.
- [52] R. E. Kahn. Ten Years of D-Lib Magazine and Counting. *D-Lib Magazine*, 11(7/8), 2000. URL: <https://doi.org/10.1045/july2005-editorial>.
- [53] R. E. Kahn and V. G. Cerf. An open architecture for a digital library system and a plan for its development. Technical report, Corporation for National Research Initiatives, 1988. URL: <https://www.cnri.reston.va.us/kahn-cerf-88.pdf>.
- [54] J. Kaplan. Over 1 million digital books now available free to the blind and print-disabled, 2010. URL: <http://blog.archive.org/2010/05/06/over-1-million-digital-books-now-available-free-to-the-blind-and-print-disabled/>.
- [55] M. Klein and M. L. Nelson. Moved but not gone: an evaluation of real-time methods for discovering replacement web pages. *International Journal on Digital Libraries*, 14(1-2):17–38, 2014.
- [56] M. Klein, H. Van de Sompel, R. Sanderson, H. Shankar, L. Balakireva, K. Zhou, and R. Tobin. Scholarly context not found: one in five articles suffers from reference rot. *PLoS one*, 9(12):e115253, 2014. URL: <https://doi.org/10.1371/journal.pone.0115253>.
- [57] J. Kunze. Encoding Dublin Core metadata in HTML, Internet RFC-2731, December 1999.
- [58] C. Lagoze. The Warwick Framework: A Container Architecture for Diverse Sets of Metadata. *D-Lib Magazine*, 2(7/8), 1996. URL: <https://hdl.handle.net/cnri.dlib/july96-lagoze>.
- [59] C. Lagoze. *Lost Identity: The Assimilation Of Digital Libraries Into The Web*. PhD thesis, Cornell University, 2010. URL: <https://hdl.handle.net/1813/14813>.
- [60] C. Lagoze, C. A. Lynch, and J. Ron Daniel. The Warwick Framework: A container architecture for aggregating sets of metadata. Technical Report TR96-1593, Cornell University, 1996.
- [61] C. Lagoze, H. Van de Sompel, P. Johnston, M. Nelson, R. Sanderson, and S. Warner. ORE Specification and User Guide - Table of Contents, 2008. URL: <http://www.openarchives.org/ore/1.0/toc>.
- [62] C. Lagoze, H. Van de Sompel, M. L. Nelson, and S. Warner. The Open Archives Initiative Protocol for Metadata Harvesting, 2002. URL: <http://www.openarchives.org/OAI/openarchivesprotocol.html>.
- [63] L. Lannom. Handle system overview. *ICSTI Forum*, (30), 1999. URL: <http://www.icsti.org/forum/30/>.
- [64] L. Lannom. D-Lib Funding. *D-Lib Magazine*, 12(7/8), 2006. URL: <https://doi.org/10.1045/july2006-editorial>.
- [65] L. Lannom. Editorial: Twenty Years and Counting. *D-Lib Magazine*, 21(7/8), 2015. URL: <https://doi.org/10.1045/july2015-editorial>.
- [66] L. Lannom. Editorial: The End of an Era. *D-Lib Magazine*, 23(7/8), 2017. URL: <https://doi.org/10.1045/july2017-editorial>.
- [67] R. L. Larsen. Whence Leadership? *D-Lib Magazine*, 11(7/8), 2005. URL: <https://doi.org/10.1045/july2005-larsen>.
- [68] H. Lee. 15 years of Google Books, 2019. URL: <https://www.blog.google/products/search/15-years-google-books/>.
- [69] M. Lesk. The CORE electronic chemistry library. In *Proceedings of the 14th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 93–112, 1991. URL: <https://doi.org/10.1145/122860.122870>.
- [70] M. Lesk. Perspectives on DLI-2 - Growing the Field. *D-Lib Magazine*, 5(7/8), 1999. URL: <https://doi.org/10.1045/july99-lesk>.
- [71] X. Liu, K. Maly, M. Zubair, and M. L. Nelson. Arc-An OAI Service Provider for Digital Library Federation. *D-Lib Magazine*, 7(4), 2001. URL: <https://doi.org/10.1045/april2001-liu>.
- [72] I. Lunden. Confirmed: Elsevier Has Bought Mendley For 69M–100M To Expand Its Open, Social Education Data Efforts, 2013. URL: <https://techcrunch.com/2013/04/08/confirmed-elsevier-has-bought-mendley-for-69m-100m-to-expand-open-social-education-data-efforts/>.
- [73] K. Maly, J. French, E. Fox, and A. Selman. Wide area technical report service: technical reports online. *Communications of the ACM*, 38(4):45, 1995. URL: <https://doi.org/10.1145/205323.205330>.
- [74] M. I. Mauldin. Lycos: Design choices in an internet search service. *IEEE Expert*, 12(1):8–11, 1997. URL: <https://doi.org/10.1109/64.577466>.
- [75] F. McCown, S. Chan, M. L. Nelson, and J. Bollen. The availability and persistence of web references in D-Lib Magazine. In *5th International Web Archiving Workshop (IWA'05)*, September 2005. URL: <http://www.iwaw.net/05/papers/iwaw05-mccown1.pdf>.
- [76] F. McCown, N. Diawara, and M. L. Nelson. Factors affecting website reconstruction from the web infrastructure. In *JCDL '07: Proceedings of the 2007 conference on Digital libraries*, pages 39–48, New York, NY, USA, 2007. ACM. URL: <https://doi.org/10.1145/1255175.1255182>.
- [77] F. McCown, X. Liu, M. L. Nelson, and M. Zubair. Search engine coverage of the OAI-PMH corpus. *IEEE Internet Computing*, 10(2), 2006. URL: <https://doi.org/10.1109/MIC.2006.41>.
- [78] F. McCown and M. L. Nelson. Recovering a website's server components from the web infrastructure. In *JCDL '08: Proceedings of the 8th ACM/IEEE-CS Joint Conference on Digital Libraries*, pages 124–133, 2008. doi:<http://doi.acm.org/10.1145/1378889.1378911>.
- [79] E. Miller. An Introduction to the Resource Description Framework. *D-Lib Magazine*, 4(5), 1998. URL: <https://doi.org/10.1045/may98-miller>.
- [80] J. S. Miller. W3C and Digital Libraries. *D-Lib Magazine*, 2(11), 1996. URL: <https://hdl.handle.net/cnri.dlib/november96-miller>.

- [81] W. H. Mischo. Digital libraries: Challenges and influential work. *D-Lib Magazine*, 11(7/8), 2005. URL: <https://doi.org/10.1045/july2005-mischo>.
- [82] C. M. Morris. Doing So Much More: The Fourth Annual International Conference on Open Repositories (OR09). *D-Lib Magazine*, 15(7/8), 2009. URL: <https://doi.org/10.1045/july2009-morris>.
- [83] J. Mueller. Retiring support for OAI-PMH in Sitemaps, 2008. URL: <https://webmasters.googleblog.com/2008/04/retiring-support-for-oai-pmh-in.html>.
- [84] M. L. Nelson. Report on the Third ACM/IEEE Joint Conference on Digital Libraries (JCDL). *D-Lib Magazine*, 9(7/8), 2003. URL: <https://doi.org/10.1045/july2003-nelson>.
- [85] M. L. Nelson and B. D. Allen. Object persistence and availability in digital libraries. *D-Lib Magazine*, 8(1), 2002. URL: <https://doi.org/10.1045/january2002-nelson>.
- [86] M. L. Nelson, J. R. Calhoun, and C. E. Mackey. The OAI-PMH NASA technical report server. In *JCDL '04: Proceedings of the 4th ACM/IEEE-CS Joint Conference on Digital Libraries*, page 400, 2004. URL: <https://doi.org/10.1145/996350.996466>.
- [87] M. L. Nelson and G. L. Gottlich. Electronic document distribution: Design of the anonymous FTP Langley Technical Report Server. Technical Report NASA TM-4567, NASA Langley Research Center, 1994. URL: <http://hdl.handle.net/2060/19940023070>.
- [88] M. L. Nelson, G. L. Gottlich, D. J. Bianco, S. S. Paulson, R. L. Binkley, Y. D. Kellogg, C. J. Beaumont, R. B. Schmunk, M. J. Kurtz, A. Accomazzi, and O. Syed. The NASA technical report server. *Internet Research: Electronic Network Applications and Policy*, 5(2):25–36, 1995. URL: <https://doi.org/10.1108/10662249510094768>.
- [89] M. L. Nelson and H. Van de Sompel. IJDL special issue on complex digital objects: Guest editors' introduction. *International Journal on Digital Libraries*, 6(2):113–114, 2006. URL: <https://doi.org/10.1007/s00799-005-0127-y>.
- [90] M. L. Nelson and H. Van de Sompel. A 25 year retrospective on D-Lib Magazine. Technical Report arXiv:2008.11680, 2020.
- [91] A. Paepcke, H. Garcia-Molina, and R. Wesley. Dewey Meets Turing: Librarians, Computer Scientists, and the Digital Libraries Initiative. *D-Lib Magazine*, 11(7/8), 2005. URL: <https://doi.org/10.1045/july2005-paepcke>.
- [92] T. K. Park. D-Lib Magazine: Its First 13 Years. *D-Lib Magazine*, 16(1/2), 2010. URL: <https://doi.org/10.1045/january2010-park>.
- [93] N. Paskin. DOI: Current Status and Outlook. *D-Lib Magazine*, 5(5), 1999. URL: <https://doi.org/10.1045/may99-paskin>.
- [94] N. Paskin. The DOI Handbook, 2019. URL: <https://doi.org/10.1000.182>.
- [95] M. Paulli. First Monday – the first of its kind: a 15-year-old open access journal about the internet, catering to a diverse audience. <http://web.archive.org/web/20111028201636/http://www.oastories.org/2011/09/denmarkus-journal-first-monday/>, 2011.
- [96] S. Payette, C. Blanchi, C. Lagoze, and E. A. Overly. Interoperability for Digital Objects and Repositories: The Cornell/CNRI Experiments. *D-Lib Magazine*, 5(5), 1999. URL: <https://doi.org/10.1045/may99-payette>.
- [97] S. Payette and C. Lagoze. Flexible and extensible digital object and repository architecture (FEDORA). In *ECDL '98: Proceedings of the Second European Conference on Research and Advanced Technology for Digital Libraries*, pages 41–59, 1998. URL: https://doi.org/10.1007/3-540-49653-X_4.
- [98] G. H. Pike. Elsevier buys SSRN.com: What it means for scholarly publication. Technical Report SSRN 2963709, 2016. URL: <https://doi.org/10.2139/ssrn.2963709>.
- [99] A. Powell and J. French. Growth and server availability of the NCSTRL digital library. In *Proceedings of the fifth ACM conference on Digital libraries*, pages 264–265, 2000. URL: <https://doi.org/10.1145/336597.336696>.
- [100] Primary Research Group and others. *Survey of Library & Museum Digitization Projects*. Primary Research Group Inc., 2014.
- [101] S. Reilly and R. Tupelo-Schneck. Digital object repository server: A component of the digital object architecture. *D-Lib Magazine*, 16(1/2), 2010. URL: <https://doi.org/10.1045/january2010-reilly>.
- [102] P. Saint-Andre and J. Klensin. Uniform Resource Names (URNs), Internet RFC 8141, 2017. URL: <http://tools.ietf.org/html/rfc8141>.
- [103] R. Sanderson. Web Annotation Data Model, 2017. W3C Recommendation 23 February 2017. URL: <https://www.w3.org/TR/annotation-protocol/>.
- [104] R. Sanderson. Keynote: Standards and communities: Connected people, consistent data, usable applications. In *2019 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*, pages 28–28, 2019. URL: <https://www.slideshare.net/azaroth42/standards-and-communities-connected-people-consistent-data-usable-applications>.
- [105] R. Sanderson, P. Ciccarese, and B. Young. Web Annotation Data Model, 2017. W3C Recommendation 23 February 2017. URL: <https://www.w3.org/TR/annotation-model/>.
- [106] R. C. Schonfeld. *JSTOR: a history*. Princeton University Press, 2012.
- [107] R. C. Schonfeld. Elsevier acquires bepress, 2017. URL: <https://scholarlykitchen.sspnet.org/2017/08/02/elsevier-acquires-bepress/>.
- [108] P. Sefton, E. Carragáji, S. Soiland-Reyes, O. Corcho, D. Garijo, R. Palma, F. Copens, C. Goble, J. M. Fernández, K. Chard, J. M. Gomez-Perez, M. R. Crusoe, I. Eguinoa, N. Juty, K. Holmes, J. A. Clark, S. Capella-Gutierrez, A. J. G. Gray, S. Owen, A. R. Williams, G. Tartari, F. Bacall, and T. Thelen. RO-Crate Metadata Specification 1.0, 2019. URL: <https://w3id.org/ro/crate/1.0>.
- [109] D. Shotton, K. Portwin, G. Klyne, and A. Miles. Adventures in semantic publishing: Exemplar semantic enhancements of a research article. *PLOS Computational Biology*, 5:1–17, 04 2009. URL: <https://doi.org/10.1371/journal.pcbi.1000361>, doi: <https://doi.org/10.1371/journal.pcbi.1000361>.
- [110] J. A. Smith, F. McCown, and M. L. Nelson. Observed web robot behavior on decaying web subsites. *D-Lib Magazine*, 12(2), 2006. doi: [doi:10.1045/february2006-smith](https://doi.org/10.1045/february2006-smith).
- [111] J. M. Snell and E. Prodomou. Activity Streams 2.0, 2017. W3C Recommendation 23 May 2017. URL: <https://www.w3.org/TR/activitystreams-core/>.
- [112] S. Speicher, J. Arwe, and A. Malhotra. Linked Data Platform 1.0, 2015. W3C Recommendation 26 February 2015. URL: <https://www.w3.org/TR/ldp/>.
- [113] S. Sun, S. Reilly, L. Lannom, and J. Petrone. Handle System Protocol (ver 2.1) Specification, Internet RFC 3652, 2017. URL: <http://tools.ietf.org/html/rfc3652>.
- [114] R. Tansley, M. Bass, D. Stuve, M. Branschofsky, D. Chudnov, G. McClellan, and M. Smith. The DSpace institutional digital repository system: current functionality. In *JCDL '03: Proceedings of the 3rd ACM/IEEE-CS Joint Conference on Digital Libraries*, pages 87–97, 2003. URL: <https://doi.org/10.1109/JCDL.2003.1204846>.
- [115] H. Thiele. The Dublin Core and Warwick Framework: A Review of the Literature, March 1995 - September 1997. *D-Lib Magazine*, 4(1), 1998. URL: <https://hdl.handle.net/cnri.dlib/january98-thiele>.
- [116] H. Van de Sompel. Paul Evan Peters Lecture: Scholarly Communication: Deconstruct & Decentralize?, 2017. URL: <https://www.slideshare.net/hvdsomp/paul-evan-peters-lecture>.
- [117] H. Van de Sompel, T. Krichel, M. L. Nelson, P. Hochstetbach, V. M. Lyapunov, K. Maly, M. Zubair, M. Kholief, X. Liu, and H. O'Connell. The UPS prototype: An experimental end-user service across e-print archives. *D-Lib Magazine*, 6(2), 2000. URL: <https://doi.org/10.1045/february2000-vandesompel-ups>.
- [118] H. Van de Sompel and C. Lagoze. The Santa Fe Convention of the Open Archives Initiative. *D-Lib Magazine*, 6(2), 2000. URL: <https://doi.org/10.1045/february2000-vandesompel-oai>.
- [119] H. Van de Sompel and M. L. Nelson. Reminiscing about 15 years of interoperability efforts. *D-Lib Magazine*, 21(11/12), 2015. URL: <https://doi.org/10.1045/november2015-vandesompel>.
- [120] H. Van de Sompel, M. L. Nelson, and R. Sanderson. HTTP framework for time-based access to resource states – Memento, Internet RFC 7089, 2013. URL: <http://tools.ietf.org/html/rfc7089>.
- [121] H. Van de Sompel, M. L. Nelson, R. Sanderson, L. L. Balakireva, S. Ainsworth, and H. Shankar. Memento: Time Travel for the Web. Technical Report arXiv:0911.1112, 2009. URL: <https://arxiv.org/abs/0911.1112>.
- [122] H. Van de Sompel, S. Payette, J. Erickson, C. Lagoze, and S. Warner. Rethinking scholarly communication: Building the system that scholars deserve. *D-Lib Magazine*, 10(9), 2004. URL: <https://doi.org/10.1045/september2004-vandesompel>.
- [123] J. H. van der Starre, T. van der Werf, G. Schaaf, and P. Winn. Guide to best practice: Dublin Core (DC 1.0 = RFC 2413), 1999.
- [124] A. Verstak, A. Acharya, H. Suzuki, S. Henderson, M. Iakhiaev, C. C. Y. Lin, and N. Shetty. On the shoulders of giants: The growing impact of older articles. Technical Report arXiv:1411.0275, Google, 2014. URL: <https://arxiv.org/abs/1411.0275>.
- [125] C. L. Webber, J. Tallon, E. Shepherd, and A. G. E. Prodomou. ActivityPub, 2018. W3C Editor's Draft 23 January 2018. URL: <https://w3c.github.io/activitypub/>.
- [126] S. Weibel. Metadata: The foundations of resource description. *D-Lib Magazine*, 1(1), 1995. URL: <https://doi.org/10.1045/july95-weibel>.
- [127] S. Weibel, J. Godby, E. Miller, and R. Daniel. OCLC/NCSA Metadata Workshop Report, 1995. URL: http://www.oclc.org:5046/conferences/metadata/dublin_core_report.html.
- [128] S. Weibel, R. Iannella, and W. Cathro. The 4th Dublin Core Metadata Workshop Report. *D-Lib Magazine*, 3(6), 1997. URL: <https://hdl.handle.net/cnri.dlib/june97-weibel>.
- [129] B. Wilson. From the Editor: An Introduction and an Invitation. *D-Lib Magazine*, 5(1), 1999. URL: <https://doi.org/10.1045/january99-editorial>.
- [130] B. Wilson. Introducing the D-Lib Alliance. *D-Lib Magazine*, 13(3/4), 2007. URL: <https://doi.org/10.1045/march2007-editorial>.
- [131] B. Wilson. What's in a Name? Categories of D-Lib Content Defined. *D-Lib Magazine*, 14(3/4), 2008. URL: <https://doi.org/10.1045/march2008-editorial>.
- [132] B. Wilson and A. L. Powell. A Tenth Anniversary for D-Lib Magazine. *D-Lib Magazine*, 11(7/8), 2005. URL: <https://doi.org/10.1045/july2005-wilson>.
- [133] J. York. This library never forgets: Preservation, cooperation, and the making of HathiTrust digital library. In *Archiving Conference*, pages 5–10, 2009.