Integrating Digital Libraries and Electronic Publishing in the DART Project

Gordon Dahlquist
+1 212-854-7421
dahlquist@columbia.edu

Brian Hoffman
+1 212-854-0295
bh311@columbia.edu

David Millman
+1 212-854-4284
dsm@columbia.edu

Research & Development Group
Academic Information Systems
Columbia University

ABSTRACT
The Digital Anthropology Resources for Teaching (DART) project integrates the content acquisition and cataloging initiatives of a federated digital repository with the development of scholarly publications and the creation of digital tools to facilitate classroom teaching. The project's technical architecture and unique publishing model create a teaching context where students move easily between primary and secondary source material and between authored environments and independent research, and raise specific issues with regard to metadata, object referral, rights, and exporting content. The model also addresses the loss of provenance and catalog information for digital objects embedded in "born-digital" publications. The DART project presents a practical methodology to combine repository and publication that is both exportable and discipline-neutral.

Categories and Subject Descriptors
H.3.7 [**Information Storage and Retrieval**]: Digital Libraries – collection, dissemination, standards, system issues, user issues.

General Terms
Design, Standardization, Management.

Keywords
Digital Libraries, OAI, Metadata, Electronic Publishing.

1. INTRODUCTION
The Digital Anthropology Resources for Teaching (DART) project integrates the content acquisition and standardized cataloging initiatives of a federated digital repository with the development of scholarly publications and the creation of digital tools to facilitate classroom teaching. A product of EPIC (Electronic Publication Initiative at Columbia[1]), DART is implemented as a union between the traditional perspectives of the library and the scholarly publisher.

This paper describes the DART project's technical architecture and presents its integrated implementation of repository and electronic publishing initiatives. Following this is a discussion of issues raised by our approach (metadata issues, detailed object referral, rights management, services to external systems), and an outline of plans for future work, specifically with regard to expanded authoring environments and exporting formats.

Funded jointly by the National Science Foundation in the U.S. and the Joint Information Systems Committee in the U.K., DART is a partnership between Columbia University and the London School of Economics and Political Science in which postdoctoral Fellows, working closely with senior faculty and specialists in digital libraries and learning technology, are developing and using new models and resources for teaching anthropology. This is an area in which, as reinforced in the recent Digital Library Federation report "Digital Library Content and Course Management Systems: Issues of Interoperation" [15], there is an urgent need for real world interoperability demonstrations. While the focus of the existing repository is in the field of anthropology, the DART model presents a practical methodology to combine repository and publication that is both exportable and discipline-neutral.

The scope of the digital repository is established by area librarians and scholars, working with editorial staff to curate content selection, describe hierarchies, rights, provenance, and other metadata, and utilizes harvesting protocols such as OAI-PMH [9,18] to acquire targeted records and resources. The project then employs postdoctoral teaching Fellows, working within the EPIC publishing environment with editorial and technical staff, to apply teaching-related metadata, annotation, text, etc. to repository material to create self-contained digital teaching tools such as online syllabi, complex learning objects, and curriculum models.

JCDL'05, June 7–11, 2005, Denver, Colorado, USA

[1] The Electronic Publishing Initiative at Columbia is a collaboration among Columbia University Press, the Columbia University Libraries, and Columbia's Academic Computing R&D, which creates and maintains a number of scholarly digital publications, including Columbia International Affairs Online (CIAO), Columbia Earthscape, Gutenberg-e, and Columbia American History Online (CAHO).

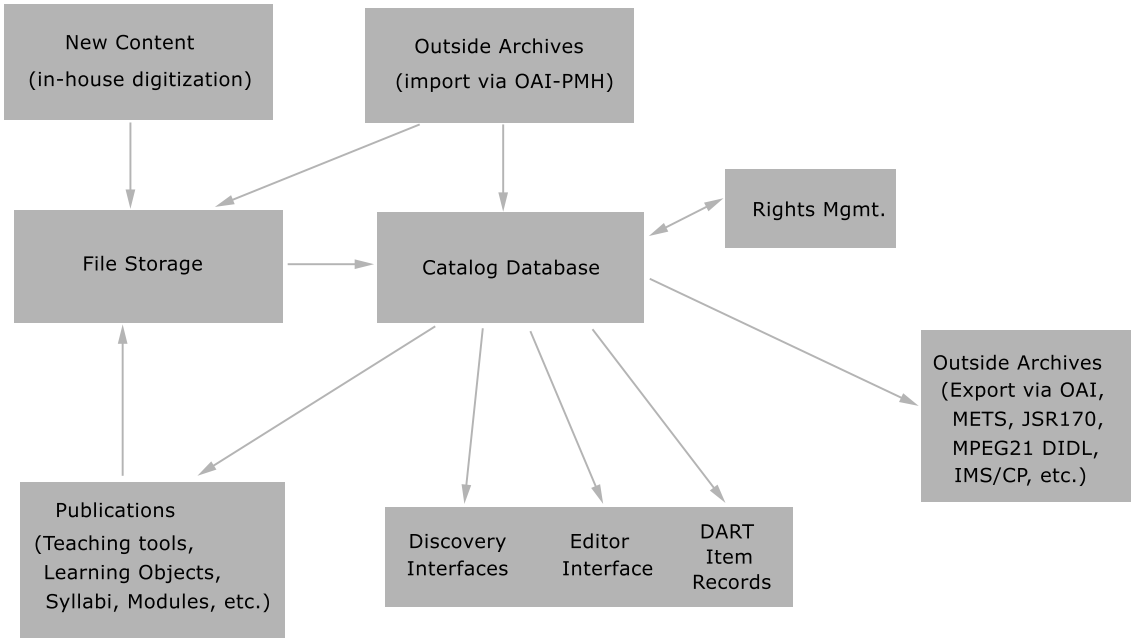

Figure 1. DART Publishing Model

Because these teaching tools emerge from and retain links back to the larger DART repository, students are introduced to a specific context within a given learning object, while remaining free to examine those same resources within the unconstrained context of the entire collection. This unique combination puts students into a relationship where they benefit from the added value of editorial and pedagogical structure without sacrificing the unfiltered access to a traditional library collection crucial to their own independent research. Because these publications and learning objects emerge from and lead back into the larger collection, DART offers an environment where undergraduates are given the ability to make the transition to graduate-level research methods in a way not available in most (digital or non-digital) secondary-source learning materials.

By synthesizing technologies and practices developed by both electronic publishers and digital librarians, the DART project is able to both create digital publications from a rich repository, and ensure scholarly access to canonical metadata records and resources across that repository. The model also addresses a problem for many 'secondary source' electronic publications, and indeed across the Internet: the loss of provenance and catalog information for digital objects embedded in digital publications.

2. SERVICES AND PRODUCTS

In this section we describe the services and workflows that are responsible for building and maintaining the digital library, the fabrication of "born digital," library-integrated teaching tools, and the distribution of these to students at Columbia, the London School of Economics, and institutions that may license our content to support their own classroom teaching.

2.1 Catalog

DART maintains catalog records for each item in its digital library. An item's **DART record** contains, at a minimum, Dublin Core metadata [2], links to content bitstream(s), and metadata that identifies the item's source(s). A record may link to more than one bitstream. For instance, unless it was obtained at a low resolution, an image will have a thumbnail copy, a viewing copy, and an examination copy: small, medium, and large. Though DART aims to acquire bitstreams for each item in the library, there are cases where technical or legal factors make this impossible. In these cases, a DART record will link to bitstream(s) stored on other systems.

End-users (such as students), DART staff, and DART teachers all use DART records to see metadata and access bitstreams. For end-users, the catalog functions as a read-only resource that is accessed through bibliographical links woven into DART teaching tools. For DART teachers and editorial staff, the catalog functions as a read/write interface to an "asset library," and allows them to create and edit metadata, upload bitstreams, and organize their workflow.

The DART library contains items that were digitized in-house and items that were copied from remote collections. While teachers and editorial staff can discover content in any remote collection or website, collections that expose their metadata through protocols like OAI are our favored sources. The project's digital librarian uses the catalog to present collections that do this for our teachers and editorial staff. Regular record-harvesting cycles then keep the catalog supplied with **candidate records**. When teachers and editors search or browse the catalog for items to stitch into their teaching tools, these candidate records appear interlaced with DART records. This makes a fruitful discovery space within the

catalog itself, and is analogous to viewing items in stock and items in a warehouse at the same time.

After a teacher or editor flags a candidate record to indicate that she desires to acquire that item for teaching, a new DART record is automatically created and assigned a unique identifier, its DART ID. The newly-created DART record pulls metadata from the candidate record, according to pre-defined mappings between the candidate record's format and the DART format. The candidate record is then duplicated and associated with the DART record, since candidate records are overwritten with each harvest, potentially changing. Finally, bitstreams are copied from their remote location to the library whenever possible. Though a script facilitates this last process for some of our remote sources, it requires close oversight or manipulation, because the OAI protocol is not normally used for bitstream harvesting (most OAI DC Identifier elements point to HTML pages with links to bitstreams, discussed in 3.2 below).

The candidate records make the catalog a good place to discover new things that are likely to be relevant to a teacher's interests, because of the librarian's general discrimination and awareness of teachers' syllabi. From a publisher's standpoint, including candidate records in the catalog is a matter of efficiency, since they allow publishing concerns to influence the teachers' and editors' discovery of new library assets. Given that many different remote items may equally satisfy a single pedagogical role (e.g., a photograph of Gandhi), one that already has some metadata and is derived from an accessible, reliable collection is the better choice. Since the administrative overhead associated with acquiring assets is virtually always a limitation for online educational publishing, an efficiency like this is ultimately a matter of quality, too.

End-users are presented read-only copies of DART records that have complete metadata and links to (local or remote) bitstreams. DART records derived from candidate records include two links that map this derivation for end-users (and external systems, see 3.4). The first link is to the resource from which DART's item derives: in OAI records, the DC:Identifier element. The second link is to the duplicated candidate mentioned above, which we preserve in the event that the original resource or its metadata are altered after the DART derivation is made.

2.2 Teaching Tools

DART teaching tools include online syllabi, topic-organized modules, interactive web sites, simulations, and other structured assemblages of smaller digital objects. DART teachers and technologists at the London School of Economics, for instance, have created highly experiential teaching tools that use video, primary sources, and interactivity to encourage students to think critically about how ethnographies are made. Some teaching tools are designed to be used outside the classroom, as assignments; others can be used in the classroom as group exercises, experiments, or as presentations. Each teaching tool's form is determined by a teacher with particular pedagogical aims, in consultation with an editor and an instructional website designer. While some formats, e.g., online syllabus, invite a Content Management System (CMS) authoring solution, the complexities of games, interactive maps, and animations have convinced us that no single authoring tool can offer the kind of range of learning object forms our project hopes to encourage.

A teaching tool may relate to many items in the library, as an item in the library may be deployed in many teaching tools. A teaching tool may simply cite or refer to an item's DART record, or may itself be a derivation of the intellectual content contained in an item, e.g., a collage. In collage-like constructions, an item's intellectual or artistic content gets physically built into a born-digital object. Our project has set out to find ways to design these synthetic born-digital objects in ways that expose their relation to digitized items in the library. We believe that many other born-digital learning objects (and other digital objects) could also, with a small amount of attentiveness during their production, allow users to "dissolve" collage-like syntheses of intellectual/artistic content and discover sets of discrete items. However, in almost all current electronic publishing practices, when an item "goes into" a born-digital object during its production, a user cannot "go out" of the final product and examine its primary sources.

We have developed, and continue to refine, a navigational idiom to express the relationships between teaching tools and their constituent items' DART records. At or near the point in a teaching tool where it connects to an item in the library, a hyperlinked icon indicates to the end-user the sign of a DART record. When an end-user follows a link to a DART record, she may choose to read an item's metadata, view an examination copy, move laterally to other DART records, explore the collection from which DART derived the resource, or return to the teaching tool.

These relationships between a teaching tool and items in the library—between a secondary source and its constituent content—are critical to the project's effort to produce "narrative" relationships among primary sources, narratives that do not become self-contained, closed objects and that do foster exploration of—and deconstruction of—academic/pedagogical syntheses of primary content.

Themselves digital objects, DART's teaching tools are tagged with the same metadata fields as the digitized items discussed above, and can be found in the catalog, alongside digitized items.

2.3 Digital Library Interfaces

We recognize that exposing items through the contexts of teaching tools, while adding value to a library, does not eliminate the need for services that allow users to browse and search the library's catalog. We provide users with these services with a variety of pages that we call, collectively digital library interfaces. Currently, these interfaces are simple browse and search services, which are adequate for exploration of the relatively small DART library. Over time, these library interfaces will evolve into more complex services, such as faceted interfaces. Metadata makes the content divisible by controlled-vocabulary keywords, genre and format, geographical coverage, data, complexity, etc. Library items can also be arranged according to the classes/semesters they have been used in.

2.4 Exports to Outside Discovery Services and Learning Management Systems

Digital objects in the DART library can be exposed to other library catalogs and learning environments through the export of OAI and/or METS [8] records, as well as DART's local record

format. DART can also pass along the records it harvests from other libraries (candidate records). In so doing, we seek to fully articulate the relations (e.g., of derivation) between items in our library with the items in other libraries. This has made the OAI provenance tag especially interesting to us.

The OAI provenance tag allows versions of an object's metadata to travel across systems, and articulates metadata's origin and history. While this element underpins versioning control for metadata, it does not address the versioning problems that a re-purposer like DART introduces by copying and altering objects themselves. Nor does it address the parallel problem of re-purposed metadata, as when the description of a derived object is extracted from the original's metadata.

We seek to address the first of these problems by offering OAI requesters the dart_xdc metadata schema, which provides, in addition to Dublin Core metadata elements, our derivedFrom tag. This tag contains the sourceObject tag, which contains a datestamp and an identifier (mapped from the DC identifier in the candidate record's metadata). It also contains a description of the derivation process, which may involve resizing or cropping an image, extracting a clip from a video, or editing text. In designing the derivedFrom tag, we have looked to the role the OAI provenance tag performs with respect to metadata, and created a parallel structure for the digital objects themselves.

The second problem introduced by the use of OAI in asset acquisition and repurposing arises when DART uses a candidate record to copy a bitstream from a remote location, and then populates the new derivative version's record with some of the metadata from the candidate record. This metadata is then altered by an editor or area specialist, an occurrence that is highly similar to the alterations, made by service providers, that occasion 'true' values in the altered attributes of provenance tags. To use the OAI provenance element to record such a relation would, in our case, confuse derivation with the life history of the archival original (e.g., format migration). Including such a relation seems to be a crucial part of tracking provenance, yet there is no unambiguous way to do this through the OAI provenance tags. While the potential exists to express metadata provenance alongside object provenance in the OAI metadata element, this does not eliminate the problem of harvesters making incorrect inferences, from the absence of provenance tags in the OAI 'about' container, in attributing authorship to a record's metadata.

2.5 Access

The DART library is protected by the Shibboleth system [14]. Administrators can offer access to specific sets of items based on users' affiliations with groups or classes at their learning institution. Students at the London School of Economics, Columbia's partner in the DART project, currently access items in the digital library using their university ids and Shibboleth authorization.

2.6 Rights Management

Part of the DART editorial process includes rights management—the pursuit, and documentation of, permission to distribute, aggregate, or alter intellectual content.

Rights management is handled by an application developed at Columbia that is based on concepts traceable to both the ODRL

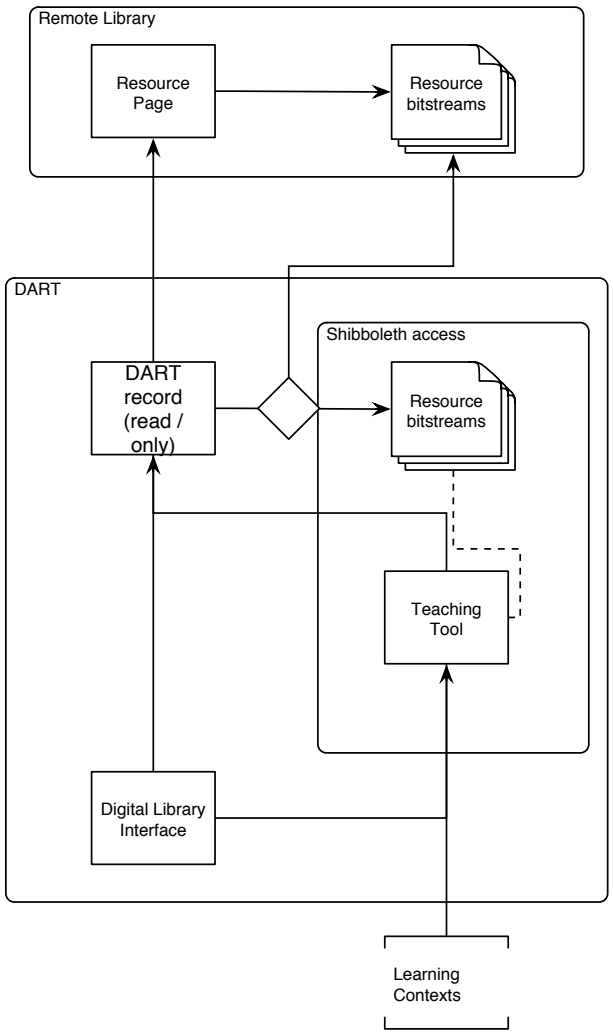

Figure 2. Simplified DART Information Architecture

[10] and the FRBR recommendation [4,17]. In addition to modeling relationships between containers of intellectual or artistic content, this application has become central to many electronic publication efforts at Columbia. It keeps track of item histories, centralizes metadata for sources (both works and agents), and describes usage agreements and contracts, including conditions and restrictions of use.

Editorial staff use the rights management application to manage their workflow as they vet the contents of the library. This work proceeds in parallel with the production of teaching tools and digital library interfaces, which actualize the relationships the rights researchers track and sanction. Depending on a number of administrative factors, rights research may occur before, during, or after the production of teaching tools. And, as the distribution requirements for completed tools evolve, rights management continues, negotiating alterations in usage permissions as necessary.

2.7 Implementation and Technology

The DART catalog data currently reside in XML files that are stored in proximity to the bitstreams that make up the library's content. The static DART records viewed by students and other end-users are generated with XSL stylesheets that execute scheduled transformations on the XML files. The editable interfaces for teachers and editors are run by PERL CGIs. As the project scales, a database and Java web interface will replace this system, which has to date had the advantage of making workflow and data-model changes relatively easy. The rights management application is already running as a production level Java/database application.

PERL scripts manage regular requests for record harvests from selected repositories that expose their catalog data through OAI. The XML records acquired through this process are used to populate the catalog with candidate records. When a teacher sees a candidate record in the catalog and wishes to acquire that item, a script creates a new DART ID and pre-populates an XML file with metadata from the candidate record, using an XSLT stylesheet to map from OAI to the in-house schema we use for DART items' XML files. So far, bitstreams have been copied manually by technical staff and stored with the XML file (see discussion below).

Teaching tools are built by both hand-coding XHTML and using XSL/XML publishing. A teacher, editor, and learning object designer collaborate to determine the navigation structure, pedagogical aims, and use cases for each learning object.

3. DISCUSSION

Through our work on DART we have identified a set of tensions between the needs of authors who assemble, manipulate, and reuse digital materials and the curatorial and preservation priorities of libraries.

3.1 Metadata Custodians in Distributed Environments

As a matter of course, metadata records harvested for DART retain all original metadata as well as provenance links to source records and source bit-streams. In some instances, the requirements of DART's publishing process necessitate changes to the original source metadata, resulting in a new record specific to DART. In one example, a user of DART requests the acquisition of an actual digital object from a harvested candidate record. A new record is then created, with a DART ID number assigned to the digital object (as it becomes a new resource within the DART repository). The original metadata is mapped against the in-house metadata requirements and any empty fields are populated by DART area librarians and editors, including applicable rights information. All of this, while the object itself remains identical to the original, results in a different metadata record expressing DART's editorial and pedagogical standards. Further, by retaining provenance information and links back to the original, preserved object, the DART-specific record exists as a local parallel to, but not a replacement of, the original metadata record. In another instance, the acquired digital object may be altered for publication (e.g., an image resized into thumbnail, view, and examination copies). In this case the DART metadata record, while again enhanced with local requirements, provenance and rights information, and links to the original source material, serves as a wholly unique record for the newly published versions.

3.2 Digital Object Referral Semantics

A reference to a digital object in a catalog record may refer to the literal bits of an image file, or it may refer to a web page with that image embedded in it, or it may refer to a thumbnail version of the image, etc. As a result, it is usually necessary to manually inspect the object, perhaps navigate further to obtain the most desirable version, and then make a copy that may be further manipulated for a particular use. There is no standard semantic convention for describing digital objects with enough precision to automatically retrieve, or refer to, a version suitable for a particular use in a publication or teaching resource.

Efforts to solve this problem have taken several forms. The Fedora repository system [5] has the notion of object "behaviors," which allow references to be qualified with additional parameters that can indicate a particular manipulation to be carried out by the repository before delivering the object. Fedora is designed for very wide flexibility in the actions that can be specified. But there is no standard or consensus on what actions ought to be made available and how they should be specified.

For preservation and stability reasons, the current DSpace [1] system emphasizes the use of persistent identifiers to retrieve objects. As object formats change over time, the identifier remains the same and it resolves to a web interface offering further selection of the currently appropriate bitstream. Inventing a semantic scheme to resolve directly to bitstreams would amount to the same situation as with Fedora above and may, in fact, pollute the preservation philosophy implied by the preferred use of persistent identifier.

A repository that serves objects embedded in the METS format could presume that the receiving system will process the metadata carried by the METS wrapper and that this information will be sufficient to enable automated manipulation of the objects. There is not yet a wide enough body of practice to understand the interoperability implications clearly.

Researchers involved in these and related efforts are aware of these issues and progress is being made. For example, Fedora can import and export METS. And DSpace is actively working on an export capability for METS.

3.3 Towards a More Distributed Rights Management

As publishers, we are concerned with intellectual property rights surrounding use, derivation, and distribution for the items in our library and our teaching tools. We are therefore interested in the development of methodologies and vocabularies for documenting and exchanging information about intellectual property rights with the distributed sources of our learning content. Namely, we can imagine the expression of requests to use, permission to use, conditions, requirements, signatures, and modifications to original

agreements over a distributed network of libraries and publishers of digital materials.

Though this paper has stressed the scholarly value of our bibliographical links and catalog records, we feel that our information architecture is also a positive step towards articulating intellectual property relationships between distributed resources. The bibliographical links woven into our publications preserve some essential provenance information in hyperlink form, allowing users to navigate "through" and backwards in the production history of a portion of content. We will continue to investigate the possibilities for promoting scholarly references by preserving the relationship between a remote source and an asset's appearance in publication. We will also continue to observe, and where possible apply, the work currently being done in the world of shared rights documentation.

Standards and examples of structured rights expression already exist, such as the ODRL vocabulary [10] and the Creative Commons [3]. The RoMeo Project [12] has begun to establish a set of requirements for a rights management standard appropriate for archives and digital libraries.

3.4 Services for External Systems

We recognize that the DART Library should ideally operate in the broader digital landscape that includes a variety of other systems, tools, and presentation strategies, including course management systems, aggregation services, authoring tools, content management systems, and other digital libraries. As discussed in McLean and Lynch [16] and in Flecker and McLean [15], this is a highly complex environment in which communication among different organizations, communities of use, and software systems is not yet well understood.

The DART metadata and content will therefore be exposed in several ways. Currently we serve metadata in OAI and object content with metadata in METS. We expect the initial customers for this data will be other library catalogs and digital library systems.

We are also exploring content integration with both the local Columbia course management system and with the Sakai Collaboration and Learning Environment [13]. In both of these systems, the architectures for interoperation with content repositories are not entirely finalized. In the local course management system, known as Courseworks, the motivation for a separate content repository was as an extension to its authoring functions. In other words, the primary purpose of the repository was a more persistent area for content export. Several architectural assumptions of that system are being reconsidered so that it can better accommodate operations in the opposite direction, i.e. discovery and importing of content from library repositories.

By exposing the DART catalog through multiple standard interfaces we enable Library items to be easily discovered by other systems. By exposing DART content through such interfaces we hope to maximize others' ability to reuse our material in a variety of settings.

We have not yet refined our export plans into standard subsets of the data. Current exports serve the entire library, including data for objects created entirely locally, derivative objects imported and manipulated, and all metadata, whether generated locally or imported.

4. FUTURE WORK

4.1 Learning Objects/Authoring Services

At this point the publications generated from the DART repository reflect the specific needs of individual scholars, and have been limited to online syllabi and basic curriculum modules. As the project develops, our intention is to tailor publication strategies equally to scholars developing materials "inside" DART (that is, through local authoring toolkits enabling affiliated scholars to create materials more readily with an expanded range of design), as well as "outside" the project (that is, making the repository available to other scholars with authoring initiatives of their own). While this will serve to more rapidly expand the repository and make its resources available to a broader community of users by opening DART to outside authors, the intention is also to avoid the limitations of any particular authoring technology.

The greatest strength of EPIC lies in its experience as a publishing organization. While the initial stages of DART's development have concentrated on the establishment of a digital repository, the next phases will see a shift toward the more systematic creation of complex learning objects, as well as commissioned teaching case studies by selected area scholars. The intention is to measure whether these various publishing initiatives can translate into a sustainability model able to support the continued growth of the repository, as well as provide a more robust methodology to support the DART model for other disciplines.

4.2 Export

We plan additional object and metadata exports in the IMS Content Packaging [6] format and possibly MPEG21 DIDL [11] to follow. We are also exploring a content exposure service based on the JSR170 Content Repository for Java API [7].

As we track the Sakai development, we look forward to a stable interface between the Sakai environment and external content repositories within the next few months.

4.3 Extending Rights Application Services

Our rights management process does not currently support the exchange of rights-related metadata between organizations. Though most content and descriptive metadata originate with our content sources, virtually all rights-related metadata is the result of arduous research by DART project staff. The contractual transactions that authorize our use of intellectual property are only recorded in our rights application's database after these transactions have been executed through phone calls, e-mails, and faxes. Clearly, tighter integration of digital library methodologies with rights expression languages would further both our interests as publishers and those of collections whose content we seek to repurpose.

Future releases of our rights management application will include other organizations in our intellectual property workflow. As standards develop for the expression and exchange of rights metadata between organizations, our rights application will need to adapt; we believe it is likely that the services we offer in the near future to administrators at remote locations will serve as a temporary and transitional step toward true distributed rights documentation.

5. CONCLUSION

The DART project integrates digital repositories and electronic publishing in a way that invokes a host of attendant issues related to each in particularly dense combinations. Starting with the acquisition of resources and metadata records, locally digitized or harvested through OAI-PMH, the DART workflow also tracks provenance, usage rights, and includes discrete links to both bitstreams and source metadata. As learning objects are published from the DART repository these issues can become interrelated with those of access control, as well as specific metadata, file-types, and publication needs associated with different authoring environments, content management systems, course management systems, and other digital libraries as resources within the DART repository are in their turn made available for export.

By locating its digital repository within a framework of electronic publishing, the DART project emphasizes both the curated quality of its collection and the varied publications that emerge from it, developed without keying production to a specific, potentially limiting authoring technology. DART's open range of original content developed within the organizational protocols of a federated archive may serve as an opportunity for basic sustainability, as well as a model of cyberinfrastructure that is easily distributed and discipline-neutral.

6. ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 0229076. DART is a collaboration among the Columbia University Department of Anthropology, the Electronic Publishing Initiative at Columbia (EPIC), and Academic Information Systems (AcIS). It is a partnership with the London School of Economics and Political Science, funded jointly by the NSF and by the Joint Information Systems Committee (JISC), in the U.K., under the Digital Libraries in the Classroom Program. The authors acknowledge the contributions of the DART project team, Kate Wittenberg, Nicholas Dirks, Ann Miller, Rashmi Sadana, and Gustav Peebles at Columbia; Charles Stafford, Steve Ryan, Chris Fuller, Caroline Ingram, Stephen Bond, Luke Freeman, Jerome Lewis, and Henrike Donner at the LSE. We would also like to thank Robert Cartolano at Columbia, Charles Blair at the University of Chicago, and Peter Morgan at the University of Cambridge for their valuable help.

7. REFERENCES

[1] DSpace Project. http://www.dspace.org/

[2] Dublin Core Metadata Initiative. http://www.dublincore.org/

[3] Creative Commons. http://creativecommons.org/

[4] IFLA Study Group on the Functional Requirements of Bibliographic Records. Functional Requirements of Bibliographic Records: final report. K.G.Saur, München, 1994. http://www.ifla.org/VII/s13/frbr/frbr.pdf

[5] Fedora Project: An Open-Source Digital Repository Management System. http://www.fedora.info/

[6] IMS Content Packaging. http://www.imsglobal.org/content/packaging/

[7] JSR-000170 Content Repository for Java Technology API. http://www.jcp.org/aboutJava/communityprocess/review/jsr170/

[8] Metadata Encoding and Transmission Standard (METS). Library of Congress. http://www.loc.gov/standards/mets/

[9] Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH). http://www.openarchives.org/

[10] The Open Digital Rights Language Initiative (ODRL). http://odrl.net/

[11] Organisation Internationale De Normalisation ISO/IEC JTC1/SC29/WG11 Coding Of Moving Pictures And Audio. MPEG-21 Digital Item Declaration Language (DIDL). http://www.chiariglione.org/mpeg/standards/mpeg-21/mpeg-21.htm

[12] Rights MEtadata for Open archiving (Project RoMEO). http://www.lboro.ac.uk/departments/ls/disresearch/romeo/

[13] Sakai Project. http://www.sakaiproject.org/

[14] Shibboleth Project. http://shibboleth.internet2.edu/

[15] Flecker, D. and McLean, N. eds. Digital Library Content and Course Management Systems: Issues of Interoperation: Report of a Study Group. July, 2004. http://www.diglib.org/pubs/cmsdl0407/

[16] McLean, N. and Lynch, C. Interoperability between Library Information Services and Learning Environments -- Bridging the Gaps: A Joint White Paper on behalf of the IMS Global Learning Consortium and the Coalition for Networked Information. May, 2004. http://www.imsglobal.org/digitalrepositories/CNIandIMS_2004.pdf

[17] Tillett, B. *What is FRBR?: A Conceptual Model for the Bibliographic University.* Library of Congress Catalog Distribution Service, Washington, DC, 2003. http://www.loc.gov/cds/FRBR.html

[18] Van de Sompel, H. and Lagoze, C. The Santa Fe Convention of the Open Archives Initiative. *D-Lib Magazine.* 6(2), 2000.