Collaborative Librarianship

Volume 8 | Issue 1 Article 7

2016

Shared Print Analysis Tool at the Colorado Alliance of Research Libraries

George Machovec
Colorado Alliance of Research Libraries, george@coalliance.org

Machovec: Shared Print Analysis Tool at the Colorado Alliance of Research Libraries

Shared Print Analysis Tool at the Colorado Alliance of Research Libraries

George Machovec (george@coalliance.org)
Executive Director, Colorado Alliance of Research Libraries

Abstract

The Colorado Alliance of Research Libraries has launched the Alliance Shared Print Trust and is in the process of developing a shared print analysis tool. The system allows libraries to compare themselves with other libraries that have added their MARC records so that they can easily and quickly determine what records are unique or held in common with other libraries. The comparison system is built on open source tools and has been embedded in the Gold Rush framework. The author provides a brief overview of other shared print analysis tools.

Keywords: Shared print, Consortia, Gold Rush, Library content comparison, Colorado Alliance of Research Libraries

Alliance Shared Print Trust

The Alliance Shared Print Trust is a collaborative effort of the Colorado Alliance of Research Libraries (the Alliance) through which a group of academic libraries in Colorado and Wyoming have committed to coordinate their long-term retention of print resources. As with many other regional initiatives, the goals of the project are multifaceted but focus on assisting libraries in making better decisions about what to discard or put in storage.

With the transition from print to digital collections, academic libraries want to make responsible decisions about their legacy print holdings to ensure that access to important materials is not lost. Even though there is a huge growth in the amount of material available online, much has not been digitized. Many reasons to retain strong regional print collections have been cited.[1] These factors, among others, were certainly influential in the establishment of the Alliance Shared Print Trust, with different libraries citing one or more key factors for their reasons to participate. The primary motivations tend to vary by library depending on local needs and interests.

- Resource sharing (access) – Many ebooks are locked down by contract and are only accessible to those that license them. Printed materials have a long history of being sent via interlibrary loan and through various resource sharing tools. Maintaining this fundamental function of libraries has strategic value that should not be lost.
- Preference – Many patrons still prefer the print format. This is particularly true in some disciplines and in different use cases where reading or referencing the print format is preferred.
- At risk materials – Many regional materials and other specialized items may not be widely held and the preservation of the print format helps preserve the scholarly record.
- Public Relations – One of the key benefits of a collaborative program is to allow libraries to balance competing space needs for collections, services, and studying. Many users are disturbed by the shift from print to digital and the anger over the culling of collections can be mitigated through collaborative shared print programs.
- Building programs – New library buildings or major remodeling efforts are huge motivating factors to reduce the footprint of print collections. Most efforts in these areas recognize that the growth of print

> collections has greatly declined and new space needs to center around patron and campus demands.

The Colorado Alliance of Research Libraries began in the 1970s as a consortium with a focus on collection development. In the pre-Internet era, this took the form of libraries pooling money to purchase expensive sets (e.g. Gmelin, Sadtler Standard Spectra, microform sets) and housing them at a local library but allowing sharing of these materials that otherwise might not circulate. In the 1980s the consortium focused on building its own integrated library system (called the CARL System), which was later sold in the 1990s. In the 2000s and later, the Alliance has focused on the development of a regional union catalog (Prospector), collaborative e-resource licensing, and other technology projects. The Alliance currently has 15 member libraries (14 academic libraries and Denver Public Library) although many of its initiatives expand well beyond this group to other libraries in the region. For example, the Prospector union catalog now extends to 44 libraries and is still growing.[2]

The Alliance had talked for many years about a collaborative strategy for storing print materials, and in 2014 appointed committees to draft a variety of policies. Rather than reinvent the wheel, the group closely examined other existing programs and used much of their material as a starting point for the Alliance Shared Print Trust. The Alliance collected sample agreements from other consortia doing similar programs to help inform the committees' work.[3] Participation in the Print Archive Network Forum (PAN), which is managed by the Center for Research Libraries, provided valuable networking.[4] Since 2009 PAN has held semi-annual pre-conferences prior to each American Library Association conference at which regional initiatives are highlighted and local experts can share experiences.

Through collaboration with colleagues at other consortia the Alliance learned that attempting specific shared print initiatives should be done after policies have been set. Some organizations tried to develop the policies while doing specific initiatives, which meant that policy development was greatly slowed while the details of particular projects caused natural delays. During 2014 and 2015, the Alliance Shared Collection Development Committee (SCDC) took the lead in policy development and eventually developed four documents, which were approved by the Alliance Board of Directors in 2015. Each library was then asked to sign the broad Memorandum of Understanding (MOU), which indicated interest in the program. Signing the MOU did not obligate that library to any immediate action, but as specific initiatives developed, libraries committed to keep materials for 25 years and to editing MARC records to disclose intent and codify each decision in the 583 field.

Some key characteristics of the Alliance Shared Print Trust are that the program is voluntary (no library is forced to participate) and distributed (materials may be stored in a library, local storage facility or shared storage facility), and that libraries may selectively participate in specific initiatives as defined by the group.

Following the OCLC Shared Print Metadata Guidelines has been an important tenet in the Alliance Shared Print Trust Disclosure Policy.[5] These guidelines were developed between 2010 and 2012 through a broad community-based effort and define how retention decisions may be codified and communicated through three key areas. The local implementation of these principles may vary in specific projects and the Alliance guidelines leave some flexibility as to specific implementation issues. At a very minimum, Alliance libraries need to note in the 583 field in MARC records items that have been committed for retention. This note will appear in the local catalog and the regional Prospector union catalog. While not required, The Alliance strongly encourages that these commitments also be shared with OCLC or other registries. The OCLC Shared Print Metadata Guidelines recommend:

- "Define separate OCLC Institution Symbols to identify print archived titles in facilities and full-service libraries."
- "Enter holdings-level print archives data in MARC Holdings records (OCLC Local Holdings Records, LHRs)."

Machovec: Shared Print Analysis Tool at the Colorado Alliance of Research Libraries

- "Use the 583 Preservation Action Note to describe specific characteristics of the print archives action(s) for each set of holdings."

Up until this point, the member libraries have approved four specific policy documents to launch the program.[6]

1. Alliance Shared Print Trust MOU – The broad framework document that has been signed by participating sites.
2. Alliance Shared Print Trust – Circulating Monographs
3. Alliance Shared Print Trust – Serials Policy
4. Alliance Shared Print Trust – Disclosure Policy

Since the Colorado Alliance was able to stand on the shoulders of many other people and organizations that had developed similar policies, the development of the Colorado program was quite efficient at the theoretical phase. Work by Sam Demas, College Librarian Emeritus at Carleton College, has helped guide some programs and was very helpful in determining the elements to include in the broader framework MOU.[7] Several Colorado Alliance library members are also members of the Greater Western Library Alliance (GWLA) and the Western Regional Storage Trust (WEST), which is managed by the California Digital Library (CDL). Previous policies developed by these groups helped guide the development of the broad Serials Policy.[8] Referencing the excellent work done by ConnectNY (a regional academic consortium in New York), GWLA, and WEST helped keep the Alliance in concert with other initiatives.[9]

The Alliance Shared Print Trust has focused initially on circulating monographs since these have the potential for significant space savings in some of the libraries if they can be analyzed at a large scale. Doing title-by-title determination for monograph retention is not practical for larger libraries, so the SCDC recommended investigating or developing tools for initial and ongoing analysis. Specific programs and initiatives will begin later in 2016 based on the needs and interests of participating libraries. Although the majority of discussion has focused around academic libraries, Denver Public Library, which is recognized as a major research library in the region, is an Alliance member with important collections that will factor into the Alliance Shared Print Trust. The Alliance libraries that are members of the WEST program are already making serial commitments through that initiative. These commitments will factor into future serial programs to be developed through the Alliance Shared Print Trust and are one reason why some of the early interest in the Alliance program has been on monographs.

The Alliance reviewed the marketplace and identified several excellent products for comparative collection analysis. However, the cost of the various commercial solutions was out of reach for the Alliance libraries for both the initial and ongoing phase. Since the Alliance had a long history of software development, the members decided to develop a local solution for comparing and analyzing library collections rather than to use a commercial counterpart. This development would be done at the consortium office and it also opened the door for extending the use of the tool to other consortia at a reasonable cost.

Review of Existing Tools

Beyond broad policy considerations, one of the major challenges in shared print programs is how to quickly and easily identify materials that may be candidates for storage or weeding. Doing projects at scale typically cannot be accomplished manually because of the large number of items involved. This means that libraries must use software that has been designed to provide deep analysis into a collection in comparison to others participating in a regional or national project. The Colorado Alliance did a review of some of the extant software and found some excellent solutions available.[10]

There is a range of commercial and open source projects and services available, each with the concomitant features and costs (whether in terms of actual expenditures for commercial solutions or sweat equity for open source solutions).

AGUA – In July 2014 the Western Regional Storage Trust (WEST) announced the initial release of their AGUA service.[11] As part of this program AGUA has employed gap analysis software from JRNL to analyze serial holdings gaps (see below). This service focuses on serials analysis for the WEST initiative and has a variety of characteristics including:

- Offering collection comparison reports for serials,
- Viewing Archive Holder title list proposals and commitments,
- Applying and iterating regional selection criteria,
- Prioritizing serials with specific characteristics, and
- Comparing proposals with other regional archives.[12]

Intota Assessment – This collection analysis service was developed by ProQuest to do qualitative and quantitative analysis of both monographs and serials. In addition to shared print programs, Intota Assessment can also be used in prospective collection analysis for projects such as looking at an existing print collection and comparing it to commercial ebook collections. In a 2015 pilot program with the Statewide California Electronic Library Consortium (SCELC) this software was used to do some analysis of selected members' collections. Due to this SCELC project, ProQuest enhanced the software with such characteristics as:

- Creating record match points beyond ISBN,
- Developing peer analysis reports including circulation data for those in the pilot, and
- Comparing pilot project holdings to all other SCELC libraries using OCLC data (and OCLC number match points).[13]

Greenglass for Groups – Sustainable Collection Services (SCS), which is now wholly owned by OCLC, has perhaps one of the most mature sets of tools and consulting services for shared print programs on the market. Long recognized as the gold standard for shared print analysis, the software now has full access to OCLC metadata so that libraries may analyze their content from the WorldCat knowledgebase. Items not in OCLC may be uploaded (to be added to WorldCat) and circulation data from local integrated library systems can be added. Examples of the rich suite of reports available include:

- Same edition and any edition overlap analysis by state,
- Circulation analysis with recency of use,
- Retention of use analysis after commitments have been noted in the 583 field, including a report of commitments by library, and
- Automatic addition of retention commitments to the WorldCat database.[14]

PAPR – The Center for Research Libraries (CRL) has played a leading role in shared print programs and has developed the Print Archive and Preservation Registry (PAPR) with an initial focus on providing a central clearinghouse for serial print retention commitments from various initiatives. The registry includes a directory of many of the major initiatives, the ability to download lists of archived serial titles from many of the regional programs, and statistics about various programs. The PAPR program has worked closely with the California Digital Library (CDL) on their AGUA project. Some of the features of the program include:

- Serial holdings, gaps and conditions report by program,
- Targeted collection comparison service for serials, and
- Selected aggregated statistics.[15]

JRNL – The Journal Retention and Needs Listing (JRNL) is an open source serial analysis tool, which has been optimized for gap analysis of serial runs for participating libraries. [For more information on JRNL, see pp. 22-28 in this issue of *Collaborative Librarianship*.[16]] The software was

Machovec: Shared Print Analysis Tool at the Colorado Alliance of Research Libraries

developed at the University of Florida and supports a number of major initiatives such as the Scholars Trust, which is a collaborative print journal archiving project from the Association of Southeastern Research Libraries (ASERL) and the Washington Research Library Consortium (WRLC).[17] JRNL is also used with the Florida Academic Repository (FLARE) program and WEST.[18] The JRNL tool supports such features as:

- Tracking archived titles,
- Identifying gaps or missing volumes in series, and
- Recording holdings, circumstances of storage and physical condition of serial runs.

ReCAP and Iron Mountain – A recent entrant into the shared print world is a new collaborative of ReCAP and Iron Mountain. Iron Mountain is best known for off-site tape storage and is developing a high density, robotically controlled and climate friendly off-site storage for library materials. ReCAP is developing open source middleware to support union catalog functionality including requesting features. One of the major targets for this service is shared print programs. Although the middleware in development by ReCAP has limited collection analysis features at this time, it could become an interesting solution down the road.[19]

Gold Rush Library Content Comparison System

The Colorado Alliance of Research Libraries has long had a history of software innovation and development. In 2003, it developed a suite of software called Gold Rush, which offered four major elements – link resolution, A-Z journal interface, an electronic resource management system (ERMS), and a journal comparison tool. The Alliance developed this platform at a time when these types of tools were only beginning to be introduced in the commercial marketplace and were very expensive. Some Alliance libraries used the Gold Rush suite but the service was also offered to libraries outside of the consortium and there are still many libraries in North America that use some module or another of the service.

One of the original components of the Gold Rush service was a journal comparison system that is now called "Gold Rush Decision Support" and which includes metadata from publishers, aggregators, and indexing/abstracting services. The comparison tool was built in MySQL and comparisons were done via ISSN or eISSN. Much of the original metadata provided by various vendors was "thin" and came in delimited files with few good match points except the ISSN/eISSN. When the KBART metadata standard was formally approved by NISO in early 2010 (NISO RP-9-2010) many vendors upgraded their metadata distribution to the new standard for link resolvers and ERMS systems. However, the lack of additional substantive match points still meant that Gold Rush continued to use ISSN-based comparisons. The journal comparison module in Gold Rush is widely used and has a number of valuable features:

- The system maintains both full-text as well as citation-only journal entries so that users can analyze databases for both the full-text and indexing components,
- Since indexing/abstracting databases are included in the system it is possible to compare indexing only databases (e.g. Scopus versus Web of Science) but also how indexing/abstracting services compare with aggregator databases
- The system operates in real-time and users can compare one-to-one or many-to-many in the same simple interface, and
- It is a multi-tenant cloud-based system with over 1,700 regularly updated title lists.

The Gold Rush platform made an ideal framework for the Colorado Alliance to begin to develop a library catalog comparison tool for the Alliance Shared Print Trust. However, in this case the system would not be loading thin KBART metadata but full MARC records from library catalogs. This opened the door for many new features and options.

The system is tentatively named the Gold Rush Library Content Comparison System.

Technical Infrastructure

After an initial attempt to develop the new MARC-based Library Content Comparison System in MySQL the developers realized that SQL was not scalable to complex real time comparisons of tens of millions of MARC records. The title lists loaded on the Gold Rush Decision Support journal comparison system rarely had lists over 100,000 records so comparisons, even when comparing many lists at once, happened very quickly. But when comparisons were attempted for MARC records it became clear that new technology was needed.

The obvious solution was to use Apache Solr, which is used by many of the major Web-based systems and for big data. The vendor describes Solr as "highly reliable, scalable and fault tolerant, providing distributed indexing, replication and load-balanced querying, automated failover and recovery, centralized configuration and more. Solr powers the search and navigation features of many of the world's largest internet sites."[20]

One of the advantages of using Solr is that many bibliographic systems, including many of the modern library discovery layers, use Solr/Lucene as the underlying layer for searching large bibliographic systems, many of which have over a billion records. Ingesting the MARC records requires a tool to disaggregate a binary MARC record, which is easily exported from most integrated library systems. The Google Books project developed such a tool, which it made open source and was adopted for the project and made available through GitHub where people and organizations can deposit open source software.[21] Charting software from Highcharts.com was selected as an excellent tool for visualizing output from comparisons and can be used at no cost for non-profit organizations. It is used by many of the top Web companies (e.g. Facebook, Twitter, Yahoo, Yandex, VISA, Verizon). The Highcharts library is written in JavaScript and was an easy way of adding interactive charts to this new service.[22]

The Colorado Alliance of Research Libraries, which is also the host of the Prospector union catalog and other services, hosts the entire system in its data center.

The first phase of development of the Library Content Comparison System used MARC records from the Prospector union catalog. At first glance this appeared to be an ideal solution. The server was housed at the consortium data center and the Alliance was already exporting MARC records to share metadata with library discovery services - Summon and EBSCO Discovery Service (EDS). It provided easy access to the content from 44 libraries in the region since metadata were being added in real-time. Prospector is based on the INN-Reach union catalog product from Innovative Interfaces, Inc. and has over 14 million unique MARC records. A successful beta version of the project was begun with Prospector metadata but the developers soon realized that direct deposit of metadata from local libraries was needed to fulfill the longer-term project goals.

There were many reasons to cease using the union catalog data for this project. First, not all libraries contributed all metadata to Prospector, meaning that comparisons in the new tool would therefore be incomplete. Second, Prospector uses a "master" MARC record for each bibliographic entity that was derived from the best record available (encoding level and precedence table). Although this is fine for a union catalog, libraries may want to get back their own metadata, rather than a generic MARC record, after doing an analysis. Finally, circulation and some other types of metadata are not available in the union catalog but only in local integrated library systems. If usage analysis were to be incorporated in future releases of the service then metadata would need to come from local library systems. For all of these reasons, the Alliance decided to switch to data extracted from local catalogs.

Deposit

Machovec: Shared Print Analysis Tool at the Colorado Alliance of Research Libraries

Obtaining MARC records from each local site required a space where the content could be deposited. This was accomplished through establishing a secure FTP (SFTP) login for each site that wanted to contribute records. To make the system more secure, each contributing site provides an IP address of the workstation or server from where the records would be transferred. The firewall is opened to the one address from each contributing library.

Two directories are available for each site when making a deposit. One directory is for "full" exports of library catalogs, meaning that when the records are processed they replace the existing metadata for that site. A library would use this method if it were the first time it is loading its catalog. A second directory is for "updates" and when metadata are added there it is concatenated to the extant record set for that library.

Libraries may do "updates" as often as once per day as processing is done each night. Replacing full record sets is done as needed, but libraries are encouraged to only do this on a monthly or quarterly basis. At present, MARC record deletes are handled by just replacing the entire library catalog.

Match Key

Since the project might include MARC records from all library types (academic, public, and special) as well as for materials in all formats (e.g. monographs, serials, media, government documents) the system uses a single generic match key for matching bibliographic records within the system. ISBNs and ISSNs do not exist for all records, including many older records that were created before ISBNs and ISSNs existed. OCLC numbers also could not be used since many of the records were from other cataloging sources and did not exist for all records.

After examining several match keys used in other projects and services, the Alliance developed a single match key that incorporated many elements in the MARC record. The key was developed to work with both older MARC records as well as modern RDA-compliant records. The building and indexing of the match key is part of the metadata loading process and is the slowest part of ingesting records. But once records have been ingested, the comparisons within the Solr framework happen in real time. There is no perfect match key, so it is periodically adjusted to improve matching as the system is tested and used. Although there has been no effort to create multiple match keys that are optimized for different purposes, this type of development theoretically would be possible.

Portions of selected MARC fields that build the match key include:

Author
    100 $a
Title
    245 $a $b
General Media Description
    245 $h
Publication Year
    260 or 264 $c
Pagination
    300 $a
Edition Statement
    250 $a
Publisher Name
    260 or 264 $b
Type of
    '_' Leader
Title Part
    245 $p
Title Number
    245 $n

Core Functionality and Features

The system is designed to allow the selection of one or more libraries to compare; after the initial comparison is done the user may use facets or searches to tailor the result set to what is needed. Facets may be created from any field(s) in the MARC record. The initial facets include format (e.g. monograph, serial, map, microform, musical score), publication date (with a date ranging feature), language, subject heading, geographic region (from geographic subdivisions), LC call number, and some other technical facets.

The entire MARC record is also indexed in a keyword index, but special indexes have also been created for title, author, publisher, and subject headings. Additional indexes can be created, but the generic keyword index should cover

many specialized fields. A search box is provided and users may add a search either before or after comparison-libraries are selected. (See Figure 1.)

If multiple libraries are included in the comparison the system analyzes not only what is held in common between the two sets but will also analyze that records are held at two sites, three sites, etc. To allow real time processing of comparisons, some thresholds have been built into the software to optimize performance. (See Figure 2.)

After selection of a data set to view, the system displays brief records and allows online sorting. Data may be exported in a variety of formats including MARC21 (binary MARC), MARC XML, and a delimited format (for loading into Excel or other tools). It is also possible to view full MARC records online and if a record is held by more than one library the user can view MARC records from any holding library by clicking on the institution name in the brief record.

When a library selects to export a data set, one that includes just their data (e.g. what is uniquely held by my library in the comparison set), its local MARC records are returned. (See Figure 3.)

Searches may be saved and re-executed at any time. This is particularly important for complex queries involving many facets and libraries. Result sets may be viewed on the screen and users may view their MARC records or those from other sites when there is overlap. Sorting by relevance, date, title, and author may be selected for on-screen viewing but it is expected that once a brief analysis of a few records is done, that data will be exported for the desired application. Exported metadata are stored on a web-accessible server, and users receive an email with a link to where the file resides. At the present, exports have been limited to 200,000 records. This can be expanded if needed, but limits were put on the number of records exported so that huge data sets are not all exported at once from many users only to unexpectedly fill-up server space.

The primary use case for the analysis tool is for shared print programs where a library can analyze its holdings with other partner libraries and then generate result sets for batch updating back into local systems. It is important that a library's local MARC records are returned so that libraries can batch update these records. Once a record set has been selected, the library can export local MARC records and do global updates in MarcEdit or the library's loader program. In a typical shared print program, retention decisions are codified in a 583 field and this analysis tool will assist in creating sets of records that are uniquely owned by a local library in comparison to selected libraries.

Future Functionality and Unexpected Uses

As with any system, software is never finished. The system is currently in an active state of development and libraries have suggested a list of additional features. Additional facets can be added as needed for project-specific needs. Any field from the MARC record can be made into a facet. Some of the requested updates include:

- Additional call number options. Currently there is a facet for LC call numbers, but some public libraries have asked for a Dewey call number facet. A SuDocs call number facet will also be added. Call number ranging is also in the process of being added to assist libraries in working on focused areas of their collection that cannot easily be done with other searches and facets.
- The addition of branch level comparisons, which is of particular interest to public libraries, but academic libraries with multiple branches might also want to compare overlap between branches.
- Circulation data that have been contributed from local integrated library systems will assist in determining what should be weeded or stored, based on use.
- Adding some conspectus-like features to do more specific call number range analysis and comparison.

After deploying the comparison tool, some discussion has emerged about other uses beyond

Machovec: Shared Print Analysis Tool at the Colorado Alliance of Research Libraries

shared print programs for weeding and storage decisions. Some of these include:

- Using the tool when an institution adds a new program to see how the local collection compares to an institution with a similar program.
- Loading a specific set of titles that are under consideration for weeding or storage to determine what is unique in that particular set.
- Performing quick exports of data sets for participation in other cooperative programs.
- Analyzing a collection for accreditation or membership in another organization.
- Remodeling and building projects are often drivers for needing to reduce the footprint of physical collections. This tool will assist in that while allowing a library to retain access to materials with other libraries in the region.

Conclusion

The development of the Gold Rush Library Content Comparison System has been a rewarding consortial development project that holds great promise for not only the Alliance Shared Print Trust but also other groups. The software has been designed to be scalable with the Solr/Lucene architecture and can theoretically include as many libraries as are interested. Collaboration with other consortia is currently underway and several pilots have been launched. The system is flexible so that other groups can be added in their own Solr cores (separate instances), if desired, so that attention can be focused on that particular initiative.

Figure 1. University of Colorado at Boulder vs University of Denver. Limited to call number "H" for imprints since 2010.

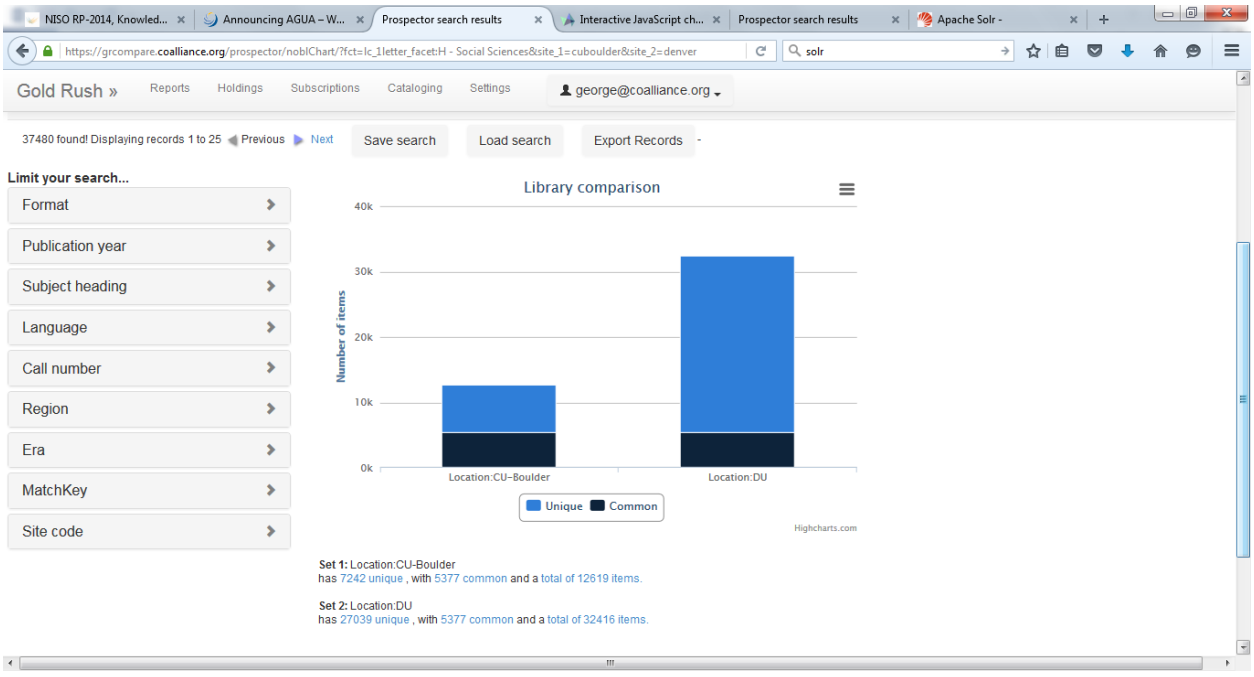

Figure 2. University of Colorado at Boulder vs 8 other libraries in the region. Limited to call number "H" for imprints since 2010.

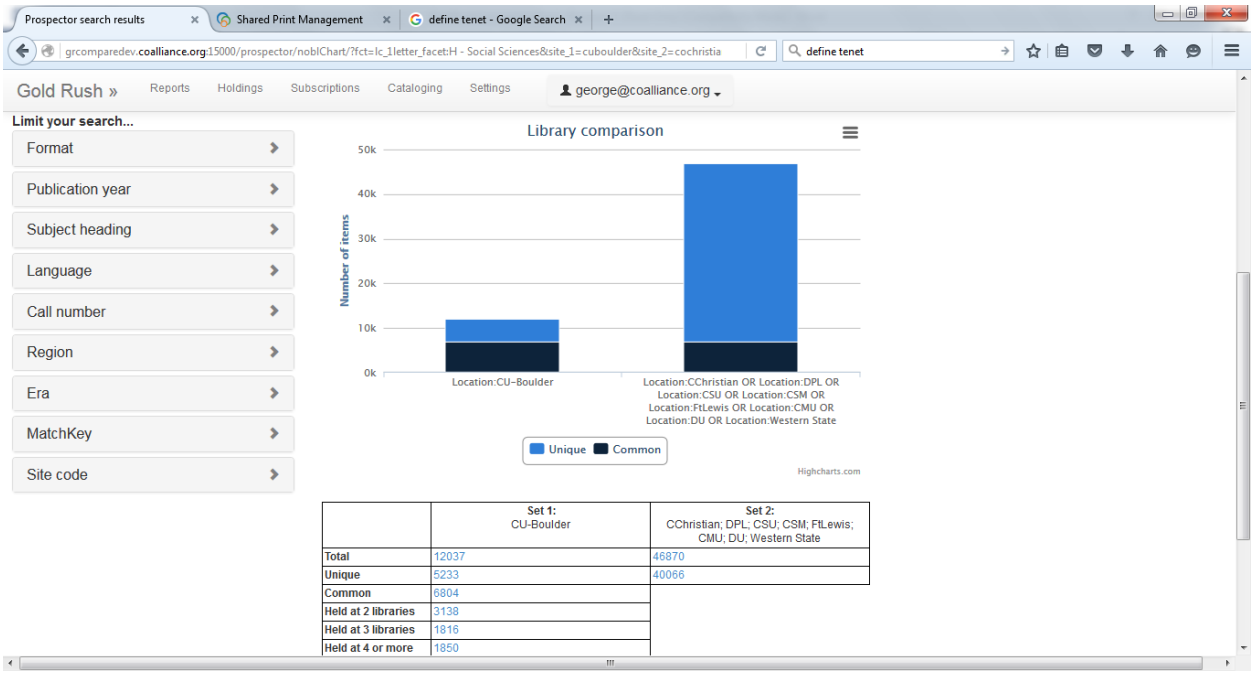

Figure 3. Viewing a record set before export.

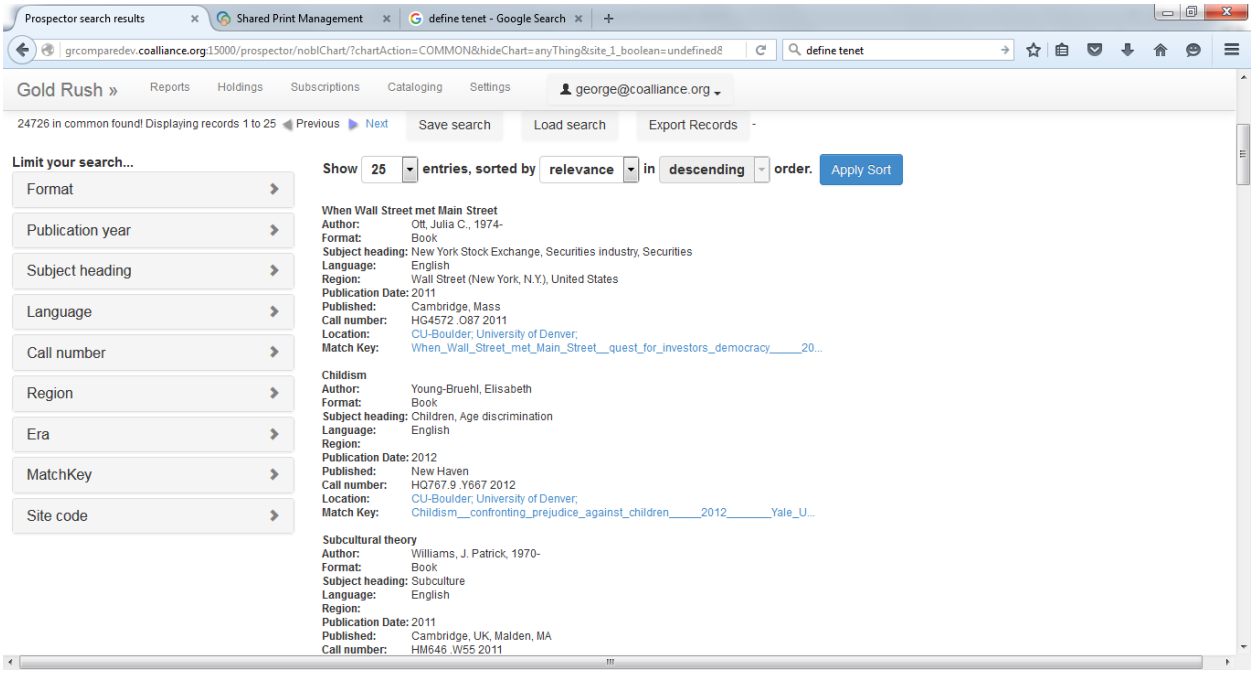

Endnotes

[1] Susanne K. Clement, "From collaborative purchasing towards collaborative discarding: The

Machovec: Shared Print Analysis Tool at the Colorado Alliance of Research Libraries

evolution of the shared print repository," *Collection Management* 37, no. 3 (2012): 153; Emily Stambaugh, "Heading West: Circling the wagons to ensure preservation and access," *Against the Grain* 22, no. 5 (2010): 18–22.

[2] "Prospector," Colorado Alliance of Research Libraries, accessed March 19, 2016, http://prospectorhome.coalliance.org/.

[3] "Shared Print Archiving Policies," Colorado Alliance of Research Libraries, accessed March 19, 2016, https://www.coalliance.org/shared-print-archiving-policies.

[4] "Print Archive Network Forum (PAN)," Center for Research Libraries, accessed March 19, 2016, https://www.crl.edu/programs/print-archive-network-forum-pan.

[5] "Shared Print Management, Shared Print Metadata Guidelines," OCLC, accessed March 19, 2016, https://www.oclc.org/en-CA/services/projects/shared-print-management.html.

[6] "Shared Print Archiving Policies," Colorado Alliance of Research Libraries, accessed March 19, 2016, https://www.coalliance.org/shared-print-archiving-policies.

[7] Sam Demas, "Curating Collective Collections – Policies for Shared Print Programs: Questions to Address in Writing a Memorandum of Understanding" *Against the Grain* 26, no. 1 (2014): 72.

[8] "WEST: Western Regional Storage Trust," California Digital Library, accessed March 19, 2016, http://www.cdlib.org/services/west/.

[9] "Shared Print Archiving Policies," Colorado Alliance of Research Libraries, accessed March 19, 2016, https://www.coalliance.org/shared-print-archiving-policies.

[10] George Machovec, "Shared Print Archiving—Analysis Tools," *Journal of Library Administration* 54, no. 1 (2014): 66-76.

[11] "Announcing AGUA – WEST's Decision-Support Information Center," California Digital Library, accessed March 19, 2016, http://www.cdlib.org/cdlinfo/2014/07/09/announcing-agua-wests-decision-support-information-center-2/.

[12] Emily Stambaugh, "What to Archive Next? Decision-Support Tools for Shared Print Journal Archiving," accessed March 19, 2016, https://www.crl.edu/sites/default/files/event_materials/7_Stambaugh.pdf.

[13] Jane Burke and Rick Burke, "Shared Print Collection Assessment (Monographs): The Use Case from SCELC," accessed March 19, 2016, https://www.crl.edu/sites/default/files/event_materials/5_%20Burke%20SCELC%20PAN%20Pres.pdf.

[14] Ruth Fisher, "GreenGlass for Groups: Visualizing the Shared Collection and Modeling Retention Services," accessed March 19, 2016, https://www.crl.edu/sites/default/files/event_materials/6_FisherGreenGlass%20-%20PAN%20-%20January%208%2C%202016_0.pptx.

[15] "Print Archives Registration Registry (PAPR)," Center for Research Libraries, accessed March 19, 2016, http://papr.crl.edu/.

[16] Judith C. Russell and Benjamin F. Walker, "JRNL: Journal Retention and Needs Listing — A Software Tool for Managing Print Journal Archives," *Collaborative Librarianship* 8, no. 1 (2016): 22-28.

[17] "Welcome to Scholars Trust," Scholars Trust, accessed March 19, 2016, http://www.scholarstrust.org/; "Cooperative Journal Retention," Association of Southeastern

Research Libraries, accessed March 19, 2016, http://www.aserl.org/programs/j-retain/.

[18] "Florida Academic Repository," accessed March 19, 2016, http://cms.uflib.ufl.edu/flare/Index.aspx.

[19] Jacob Nadal and Jeremy Suratt, "New Off-Site Storage Technologies," accessed March 19, 2016, https://www.crl.edu/sites/default/files/event_materials/10_NadalALAMW-PAN-NewOffSiteTechnologies-Final.pdf.

[20] "Apache Solr 5.5.0," Solr, accessed March 19, 2016, http://lucene.apache.org/solr/.

[21] "GitHub," accessed March 19, 2016, https://github.com/.

[22] "Highcharts," accessed March 19, 2016, http://www.highcharts.com/.