From Abstract to Virtual Entities: Implementation of Work-Based Searching in a Multimedia Digital Library

Mark Notess[1], Jenn Riley[1] and Harriette Hemmasi[2]

[1]Digital Library Program, Main Library, Indiana University
Bloomington, Indiana, USA 47405
{mnotess, jenlrile}@indiana.edu
[2]Executive Associate Dean, Indiana University Libraries, Indiana University
Bloomington, Indiana, USA 47405
hhemmasi@indiana.edu

Abstract. Libraries of digitized multimedia content provide access to virtual entities. In the case of music, where there are frequently many different performances, editions, and arrangements of a given work, the Variations2 metadata model, links all instances of a work to an abstract work record, thus yielding superior search capabilities to digital library users. This paper summarizes the motivation for addressing the music metadata problem and describes the Variations2 search user interface, which is based on our work-centric, FRBR-like metadata model.

1 Introduction

The Variations2 Indiana University Digital Library is a large test-bed development and research project funded in part by Phase 2 of the Digital Libraries Initiative, with support from the National Science Foundation and the National Endowment for the Humanities [1]. This paper reports on the state of the Variations2 test-bed software, describing in particular the search user interface. We begin by reviewing the motivations for attempting an improved environment for music search. Some of these motivations are common to other digital library efforts; others are specific to issues associated with music. We then describe our implementation of a search user interface and the current state of our system.

2 Background

Motivations for the Variations2 approach to searching come from at least two directions. First, Variations2 shares in larger library and digital library issues associated with virtualization. Second, music information offers unique challenges, challenges which have not always been met well by existing solutions.

2.1 Virtualization, Abstraction and New Metadata Models

This paper springs from the junction of two simultaneous developments: library virtualization and catalog entity abstraction.

Digital libraries provide a level of disembodiment of library materials. Digital materials have a reduced physicality in at least three respects. First, patrons cannot pick a digital item off a shelf and hold it in their hands. Digital library contents are less tangible. Second, the collocation of items in a collection need no longer be spatial in a physical sense. Hence the term *virtual*, while not synonymous with *digital*, is often used to describe digital libraries. Third, reduced reliance on physicality also becomes evident as users seek content (i.e., works) rather than containers (e.g., the "red book," the "CD with a picture of a dog"), influenced at least in part by the MP3 phenomenon where users tend to think in terms of "tracks" or individual works.

Over the last several decades, librarians have been reconsidering cataloging models. To a large extent, reconsideration has been driven by the development of cooperative cataloging models and the consequent need for common practices brought about by such systems as OCLC's WorldCat [2] and RLG's Union Catalog [3]. Such efforts also afford opportunity beyond mere consistency towards fundamental improvements to the overall model. One such improvement effort is the Functional Requirements for Bibliographic Records (FRBR) effort from the International Federation of Library Associations and Institutions (IFLA) [4].

FRBR seeks to improve upon the existing paradigm of MAchine Readable Cataloging (MARC, [5]) bibliographic and authority records, the paradigm used by cooperative cataloging efforts such as OCLC's. The MARC-based paradigm stores information about the physical item in a bibliographic ("bib") record. It also has authority records for such information as work titles, people's names, and subject headings—these records help ensure consistent and unique naming. However, MARC-based implementations often provide no *linking* even at this most fundamental level, between the record types. For example, a cataloger will find the name authority record for a book's author but may not have any way to reference that authority record explicitly within the bib record or enact global changes across the system. Instead, the authoritative name for the author is copied separately into each bib record. The *raison d'être* for authority records is to help catalogers produce consistent bib records efficiently.

In contrast to MARC, FRBR uses an entity-relationship approach to provide strong linking between records. For example, in FRBR, an *item* (e.g., a copy of a book) is an exemplar of a *manifestation* (e.g., all books with the same ISBN), which embodies an *expression* (edition) of a *work* (the abstract entity representing the original intellectual or creative content). This strong linking can be used to provide both collocation and a coherent disambiguation path for users.

The FRBR specification has been used as the basis for some system development. FRBR-based projects include FRBR support within the VTLS Virtua system [6], the AustLit Australian Literature Gateway [7], RLG's RedLightGreen [8], and OCLC WorldCat's Fiction Finder [9].

When the Variations2 Indiana University Digital Music Library project began more than three years ago, we determined to develop a metadata model that would support a greatly improved search interface for music [10]. The weaknesses of MARC-based music cataloging are well documented (see, e.g., [11]; we review them briefly in the next section). While not based directly on FRBR, the Variations2 metadata model nonetheless bears a strong resemblance (Table 1).

Table 1. Variations2 and FRBR Compared

Variations2		**FRBR *Rough* Equivalent**
Entity	**Description**	
Work	abstract concept of a musical composition or set of compositions	Work
Instantiation	recorded performance of a work (audio) or edition of a work (score)	Expression[1]
Container	physical item or set of items within which one or more instantiations are present (e.g., a CD or CD set, a score)	Manifestation[2] (physical embodiment of an expression—a release or edition)
Media Object	digital sound file(s) or score image(s)	Item[3] (an actual copy of a manifestation)
Contributor	individuals or groups related to a work, instantiation, or container (e.g., composers, performers, conductors, producers, ensembles)	Two Entities: – Person – Corporate Body

Notes:

1. In FRBR an expression can be manifested multiple times; in Variations2, instantiations are unique to a container, even if two containers reflect the same performance.
2. "A manifestation may embody one or more than one expression" [7, p. 13]. The Variations2 Container, however, is less abstract, having some amount of item-level descriptors.
3. The FRBR item refers to a copy in a collection; the Variations2 media object is a digitization of a container. Thus in FRBR, there are potentially many items for a manifestation; in Variations2, there is only digitization of a container, even if multiple media objects are needed to capture all the container's content.

Like FRBR, our system is *work*-centric, being influenced by the work of both Velucci [12] and Smiraglia [13]. We have implemented a digital music library, Variations2, based on that metadata model, have deployed the system in our music library, and have seen increasing usage over the past year and a half.

2.2 Finding Music in a MARC-based OPAC

Online Public Access Catalogs (OPACs) are the primary means by which library users access library collections. OPACs offer searching of bibliographic records (almost always) in the MARC bibliographic format, and under certain circumstances provide to the user a list of authorized and unauthorized (i.e., cross-referenced) names, titles, or subjects from MARC authority records. Despite many advances that have been made to OPACs since library catalogs first went online, searching for musical materials in OPACs can still be problematic, due to both OPAC design and to the structure and contents of the MARC bibliographic record itself.

Library catalog records are created by a convergence of a number of different standards. The MARC Bibliographic format prescribes the fields, subfields, and indicators used to mark what type of information is being recorded. The basic descriptive information that is contained in the MARC record is copied from the item being cataloged and is formatted according to the Anglo-American Cataloguing Rules (AACR2 [14]). "Access points"—other descriptive information formatted in a standard way, not directly copied from the item being cataloged—are similarly selected and formatted according to AACR2 rules. Subject headings are chosen from controlled lists, most often the Library of Congress Subject Headings (LCSH).

The MARC format and its associated data content standards provide precision to bibliographic data. Encoding of information in bibliographic records, for example, allows the distinction between works *by* a person and works *about* a person, while still providing for a connection to be made between them by using the same form of the name in both places. The catalog of MARC records provides both a *descriptive* function—reproducing exactly what is on a physical item allowing users to access titles or authors they've seen—and a *collocating* function—grouping bibliographic items representing the same authors, subjects, and, to some extent, works.

One challenge to music searching is the MARC record's focus on a "static physical artifact" [11, p.2]. The data in a MARC record describe a bibliographic item as a whole, not necessarily any specific part of it. This is problematic because items held in a music library, especially sound recordings, often contain multiple works. Thus there is often no way for a user to know, for example, which of the performers listed in a record is connected with a given piece on the recording being described.

The Nature of an OPAC. The OPACs in use in libraries of all sizes today are typically one part of large Integrated Library Systems (ILSs) used for automation of many library services, including acquisitions, cataloging, circulation, and patron billing. OPACs from different vendors also have vastly different native functionalities, and are customizable by the library implementing them. Search and browse success in an OPAC relies heavily on the design and implementation decisions for an individual installation in addition to the nature and structure of the bibliographic data in the MARC records it contains.

Indiana University's IUCAT, based on the Sirsi Unicorn ILS, is a fairly typical example of a modern OPAC. Keyword searching in a large number of fields from the MARC bibliographic record is available, as is browsing and searching on fields (actually groups of fields from the MARC record) labeled as author, title, subject, series, periodical title, and medical subject. Basic Boolean operators and term truncation are available. The OPAC performs reasonably well for simple bibliographic searches on author, title, and in many cases topical subject terms, but less well for more specialized queries essential to music searches, such as for chamber pieces with specific instrumentation. Cataloging rules place names of instruments in multiple fields within the MARC record, sometimes in the uniform title, and others as part of a subject heading. But these fields do not use terms for instrumentation in a consistent manner, so a keyword search of the entire record on instrument names will not find all relevant records in the catalog, and will at the same time add many irrelevant records. One partial solution developed was the creation of a dedicated field in the MARC record for instrumentation, but this field is rarely used, due in part to the amount of time it takes to add this information to bibliographic records, but also largely due to the fact that almost no OPACs, including Indiana University's, allow this field to be indexed for searching or display it to users.

Collocation by Work. Collocation by *work* is one of the functions of cataloging wherein OPAC designers and consequently OPAC users often do not succeed. MARC and AACR2 provide basic work collocation through a mechanism called the *uniform title*. This is a specially-formatted title that is intended to be added to all bibliographic records representing the same work. All records describing the same musical piece, whether in score or recording, have the same uniform title. There are also additions to music uniform titles that indicate, among other things, whether a record is for an arrangement of a musical work, a part of a musical work, or a musical setting of a textual work. The uniform title is intended to allow connections to be made between multiple versions (in the case of music – performances captured on sound recordings or publications of scores) of the same work and its variations.

But this work collocation function is often not readily available to the average library catalog user. First, uniform titles are not present in all bibliographic records. Cataloging rules governing appropriate use and the presence of records created before the uniform title achieved its present form are among the reasons a uniform title may be missing from a given record. Second, many OPACs don't make full use of uniform title for display

purposes. Many catalogs provide basic grouping capability on the first part of the uniform title (the actual name of the work), but then fail to meaningfully use the other parts of the uniform title that indicate format, arrangement or selection, and the like. Similarly, most library systems do not use the semantic links between whole works and their parts that uniform titles provide [11, p.4]. Since the title for a larger work is present in the uniform title for one of its parts (e.g., opera and aria), a library system could use this implied link to return records for a whole work when a user enters a query for one part of it. Current OPACs on the whole do not recognize this link, and thus fail to retrieve the larger work when a part is searched. It also offers no solution for the previously described problem of the record describing a bibliographic item while the user is in search of a musical work.

3 Implementation

In this section, we describe the current (version 2.1.1) Variations2 search user interface, including the options available on each of the four tabs (basic, advanced, keyword, and browse). We also describe how the disambiguation process varies depending both on what fields the user fills in and the actual content of the digital library.

The Variations2 software is cross-platform (Windows and Macintosh), implemented as a Java application. While the search interface could have been implemented in a web browser, the other features of Variations2 (audio player, score viewer, etc.) would not have worked as well within a browser window, so we decided to implement the entire application as separate Java windows. The technical architecture of Variations2 is beyond the scope of this paper, but a description may be found in [15].

3.1 Search Tabs

The Variations2 search window (Figure 1) is the default initial window displayed by the application after users log in. The search window is divided into two sections: the search tabs, where users specify their search criteria, and the results pane, where the results of the search are displayed. The results pane has a row of controls above it for forward/backward navigation, canceling a search in progress, or changing the display of the results by sorting or filtering.

Basic Tab. The search window defaults to the Basic tab, which provides five fields for search criteria specification.

- Creator/Composer (like *author,* but music is different)
- Performer/Conductor (critically important for music)
- Work Title (often different from the name of the container)

– Key (two drop-down lists: key letter and mode, e.g., A, minor)
– Media format (drop-down list with various types of recording and score formats)

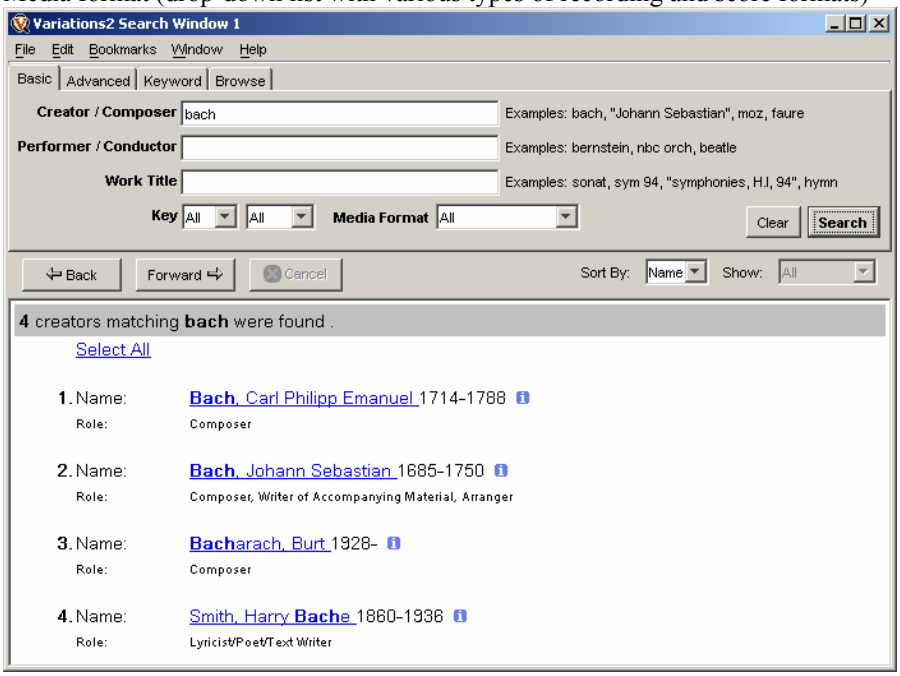

Fig. 1. Search Window, Basic Tab

In all of the text entry fields in the search interface, the following properties apply.
– Case insensitivity
– Partial words are matched by default, e.g., “beeth” will find Beethoven
– Quotation marks permit searching for the exact word or phrase
– Other punctuation and diacritics are ignored

Advanced Tab. The advanced tab offers the same fields as the basic tab, with the following additions.
– Recording/Score Title (i.e., container title)
– Other Contributor (e.g., arranger, producer)
– Publisher
– Subject Heading

Keyword Tab. The keyword tab offers two fields.
– Keywords(s) – accepts parentheses and the Boolean operators *and*, *or*, and *not*

– Media format (drop-down list with various types of recording and score formats)

Browse Tab. The browse tab (Figure 2) offers browsing of the entire collection. Users select one of the following "browse by" options.

– Creators (composers, poets, lyricists, etc.)
– Works
– Performers
– Recording albums/score volumes

Users can initiate a search either by pressing the Enter key on their keyboards while they are in one of the text fields, or by clicking on the *Search* button.

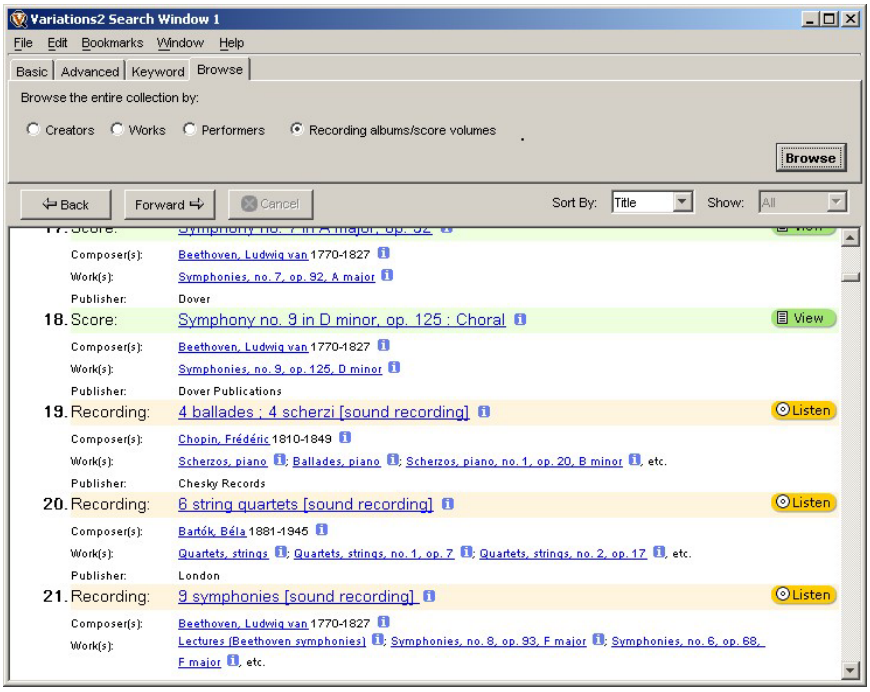

Fig. 2. Search Window, Browse Tab

3.2 Results Display and Interaction

The results display area uses a Java Swing component to render HTML. Descriptive text is black, hyperlinks are blue, and there are also buttons of various colors. Figure 3 shows

a part of the Figure 1 results display. In the gray box at the top of each result set is a description of the results that follow. The main entry (first line) for each result is in a larger font, and the matching part of the string (if any) is bolded. The iconic button indicates detailed information is available.

Fig. 3. Search Results Detail

Results are not paged: all results are returned. If there are not any results for a given search, the results pane indicates which criteria matched something in the database so users can broaden their search appropriately. Sample "zero results" output is given in Figure 4.

NOTHING MATCHED ALL OF YOUR SEARCH CRITERIA. Try changing or removing search elements.
- 1 Creator matches "beethoven".
- NO Performers match "mazel".
- 64 Work Titles match "symphony".

Fig. 4. "Zero Results" Feedback

The results pane control buttons work as follows.

- The *forward* and *back* buttons work like the buttons in a web browser: *back* displays the previously displayed search results, changing the tabs and search criteria at the top of the window as appropriate. *Forward* moves in the opposite direction through the results stack.
- *Cancel* stops a search in progress.
- *Sort By* allows users to change the ordering of the displayed search results. The choices depend on the record type currently displayed. For example, the list of creators in Figure 1 may be sorted alphabetically by name or role. The list of containers in Figure 2 may be sorted by title or (first listed) composer.
- *Show* allows filtering of containers by media type (score or recording).

Whenever a View or Listen button is present in the search results, clicking that button, or clicking the title on that same line, will launch the Variations2 score viewer or audio player, as appropriate.

An alternative navigation mechanism is available from a right-click popup menu (Figure 5). In this example, right-clicking on the score name offers two choices: opening the score in the score viewer (the default behavior had the link been clicked) or viewing detailed information about the score.

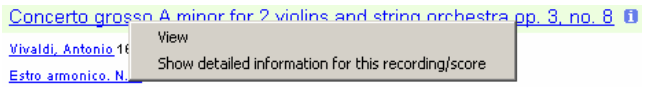

Fig. 5. Right-Click Menu Navigation

Right-clicking on "Vivaldi, Antonio" in Figure 5 also gives two options: getting detailed information about Vivaldi, or launching a new search for works by Vivaldi.

Viewing Record Details. Users may request record details either by using the popup menu or by clicking on the button. Record details are displayed in a separate window, also using HTML and having both internal and external links. Internal links allow convenient navigation within a container record. External links provide record details for referenced records, bring up an audio player or score viewer, or provide links to external resources.

Disambiguation. The search logic in Variations2 provides step-by-step disambiguation during searches. Disambiguation steps are inserted in the search process when all of the search criteria can be satisfied by a variety of results, but

– a name used as search criteria matches more than one individual or collective name in the database, *or*
– a work title used as search criteria matches more than one work title in the database.

There is a set sequence to the disambiguation. In the "worst" case, a user specifies an ambiguous *creator*, *performer*, *work title*, and *other contributor*. First the user is presented with search results listing all the matching creators where the other criteria also have matches. After selecting the desired creator, the user is presented with the list of all matching performers who perform works by the selected creator, the other criteria still matching, etc. In this worst-case scenario, the user is not presented with media links until the fifth set of search results. Typically, however, only one or two disambiguation steps are required. If, at any disambiguation step, users want to see all the results without having to disambiguate, they can click the "Select All" link (Figure 3).

3.3 Current State

Cataloged content in Variations2 is somewhat limited at present. In March 2004, the digital library contained records for 1500+ works, 1300+ contributors, in support of 282

containers (262 recordings and 20 scores). The collection grows in response to pilot project needs, development team testing needs, and an overall goal of broadening the collection.

Variations2 is installed on approximately 120 computers in the music library. Any person with an IU login can come to the music library and use Variations2. While the primary mechanism for online access to music at IU is still IUCAT and Variations (our previous-generation digital music library [16]), Variations2 is available for general use.

4 Conclusions and Future Work

This paper documents the current Variations2 digital music library search user interface as a user-centered, FRBR-like alternative to traditional MARC-based OPACs as mechanisms for finding music. We have carried out multiple lab-based and field-based evaluations; results will be published separately. The short summary of our evaluation results is that we found no fundamental flaws with the user interface or the design of the metadata model. Such problems as were uncovered seem addressable by relatively non-invasive user interface improvements.

Variations2 is a continuing research project. Among the search-related features planned for future releases are the inclusion of themes and incipits in the search interface, initially as a means for users to distinguish works but eventually as a mechanism for limited content-based searching of music. We also plan on adding search fields for instrumentation, genre (e.g., jazz, pop, rock), musical form (e.g., song, symphony, opera), and style (e.g., baroque, romantic). To the current audio recording and scanned score formats we plan to add encoded scores. We are also considering implementing a web-browser-based search interface.

Variations2 is designed as a distributed solution, for use by multiple institutions. The current implementation is more monolithic, based on the collection of a single institution. As we evolve Variations2 to fulfill its distributed promise, we will have to consider how a distributed "union" catalog can be used within the search interface (while ensuring only authorized access to the digital content!). Only by addressing barriers to distributed deployment can we develop the cooperative cataloging community necessary to support re-cataloging and thereby a future existence for our metadata model and software.

Acknowledgements

We wish to acknowledge Mary Wallace Davidson for her contribution to the Variations2 metadata model, as well as Don Byrd, Jim Halliday, and Ryan Scherle for designing and implementing the search functionality. The authors also thank Don Byrd, Mary Wallace

Davidson, Jon Dunn, and Jim Halliday for their helpful reviews of this paper. This material is based upon work supported by the National Science Foundation under Grant No. 9909068. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

1. Variations2: the Indiana University digital music library project. http://variations2.indiana.edu/.
2. OCLC WorldCat web site. http://www.oclc.org/worldcat.
3. Research Libraries Group (RLG) web site. http://www.rlg.org.
4. IFLA Study Group on the Functional Requirements for Bibliographic Records. Functional requirements for bibliographic records. K.G. Saur, Munich, 1998.
5. Library of Congress. MARC Standards web site. http://www.loc.gov/marc.
6. VTLS announces first FRBR user. Biblio Tech Review (24 January 2003). http://www.biblio-tech.com/btr11/S_PD.cfm?ArticleID=496&DO=A.
7. Ayres, M., Kilner, K., Fitch, K., & Scarvell, A. Report on the successful AustLit: Australian Literature Gateway implementation of the FRBR and INDECS event models, and implications for other FRBR implementations. International Cataloguing and Bibliographic Control, 32 1 (January/March 2003), 8-13.
8. Revolutionizing the catalog: RLG's RedLightGreen project. http://www.rlg.org/redlightgreen/
9. FictionFinder: a FRBR-based prototype for fiction in WorldCat. http://www.oclc.org/research/projects/frbr/fictionfinder.htm
10. Minibayeva, N., and Dunn, J.W, A digital library data model for music. In Proceedings of JCDL '02 (Portland, Oregon, 13-17 July 2002). ACM Press, New York, NY, 2002, 154-155.
11. Hemmasi, H. Why not MARC? In Proceedings of ISMIR 2002: The Third International Conference on Music Information Retrieval (Paris, France, 13-17 October 2002). IRCAM—Centre Pompidou, Paris, France, 2002, 242-248.
12. Velucci, S. L. Bibliographic relationships in music catalogs. Scarecrow Press, Lanham, MD, 1997.
13. Smiraglia, R. P. The nature of "The work." Scarecrow Press, Lanham, MD, 2001.
14. Anglo-American cataloguing rules, 2nd ed., 2002 revision. American Library Association, Chicago, 2002.
15. Variations2 Development Team. Variations2 system architecture: report for the project midterm review. Technical report. http://variations2.indiana.edu/html/v2-architecture-findings-nov2002.
16. Variations website. http://www.dlib.indiana.edu/variations.