

GREENSTONE AS A MUSIC DIGITAL LIBRARY TOOLKIT POSTER/DEMONSTRATION

David Bainbridge
Dept. of Computer Science
University of Waikato
Hamilton
New Zealand

Sally Jo Cunningham
Dept. of Computer Science
University of Waikato
Hamilton
New Zealand

J. Stephen Downie
Graduate School of Library
and Information Science
University of Illinois
Urbana-Champaign, IL, USA.

1. INTRODUCTION

Greenstone is an open source digital library system that has developed and matured since its inception in 1995. Today it is used in over 60 countries, with a strong emphasis on humanitarian aid. The software is also used as a framework for research in other fields such as human computer interaction, text-mining, and ethnography. This article provides a summary of Greenstone's uses to date with music documents. First we discuss incorporating musical formats into the Greenstone system; then we describe provision for searching and browsing in a music collection.

2. DIGESTING SOURCE DOCUMENTS

Central to the Greenstone design is a generalised and extensible mechanism for importing documents, called plugins. Here we describe four such plugins designed for music formats that span a range of representations: symbolic notation, raw audio and sheet music. Plugins are written in Perl. Once a base plugin has been established for a new document representation, such as symbolic music, much of the work for supporting additional formats of a similar form can be carried out through external applications, with the Perl code wrapping things together.

The first music plugin written was RogPlug, an internal format we have historically developed. With this established, developing MIDIPlug was straightforward. MIDIPlug inherits from RogPlug, and uses a modified version of the open source music notation software RoseGarden (www.rosegardenmusic.com) to convert MIDI data to this format. The data is then streamed through to the underlying Rog plugin.

Greenstone currently does not support direct content retrieval of raw audio, and so incorporation of MP3 files in Greenstone is through the text metadata contained in a file's ID3 tags. The role of MP3Plug, therefore, is to

extract this information, which it does through the use of a Perl package MP3::Info, available from the Comprehensive Perl Archive Network (CPAN). In an extension to this plugin, artist and title metadata can be used to retrieve related images from the web (albeit heuristically). These can then be used in the digital library to evoke additional interest in the work, such as the collaging technique described in Section 4.

Sheet music is our final form of music plugin. Here we have linked the plugin with CANTOR, a research-based Optical Music Recognition (OMR) tool [1], to extract a preview image of the first staff system, plus any accompanying text above it. In constructing collections based on this plugin we follow Greenstone's existing page image model that supports full-size and preview images, and navigation by page number. In an experimental version of the plugin, we fed the (possibly erroneous) symbolic music notation extracted by the OMR tool into RogPlug to build a version where direct content queries were supported [2].

3. FORMS OF QUERYING

Any text or metadata extracted from the supported music formats can be used to build indexes for querying: search by title field of MP3 files, for example. To extend the searching paradigm in Greenstone, MELDEX [2], has been added as an extension to the runtime system. Moreover, music and text queries can be combined. Figure 1 shows one application of its use. Based on a collection of MIDI files, a user is formulating a query based on "Gold" appearing in the song's title and singing an excerpt of the song they remember. After a short delay upon pressing the search button, a ranked list of tune titles is displayed (not shown) with an audio icon next to each for playback.

4. BROWSING

Collaging is a visualisation technique whereby a series of images are displayed over time at random locations on the screen. Images already in the visualisation fade over time; new images are added on top of what is already there. If at any stage a displayed image sparks some interest, the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page.
© 2004 Universitat Pompeu Fabra.

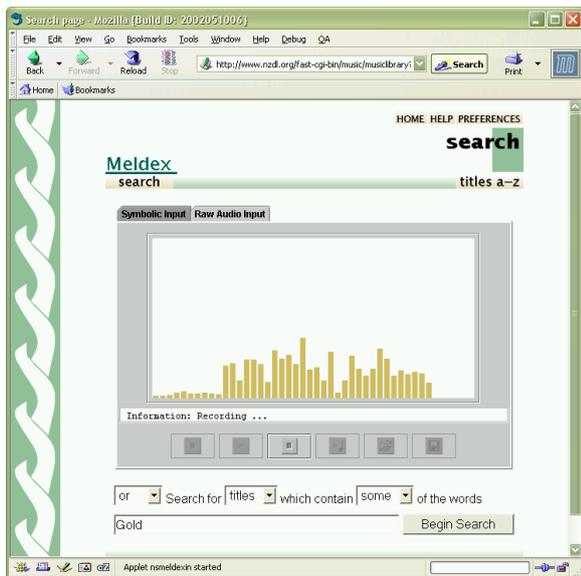


Figure 1. Direct content querying of symbolic music notation.

user can click on it and have a new window open up displaying the context of that image. We have adapted this idea as a browsing facility in Greenstone using images located through the MP3Plug extension. The new window that opens up when an image is clicked on is a metaphor for the selected record, and shows two images associated with the track (fashioned on the idea of opening up the CD cover) while the audio is played through an embedded MP3 applet. Collaging in a digital music library is the subject of a longer submission to ISMIR 2004.

PHIND, an interactive hierarchical phrase indexing technique already implemented in Greenstone, is a novel form of browsing that makes for a compelling facility in a digital music library. When the digital library is built, it is a simple matter to direct the text/metadata extracted by the music-based plugins into the PHIND module. Accessing the “phrases” link in a collection based about the Beatles and entering “love” as a starting term yields a list including “love me do,” “love you to,” “real love,” “words of love” and so on. From this list the user can explore usage of a phrase, and variants, in all documents it appears in.

5. HETEROGENEOUS DATA

Another strength of the Greenstone framework is working with a heterogeneous mix of source documents. As a concluding example, consider extending The Beatles collection alluded to above with MP3 files, MIDI files, Guitar tablature in the form of plain text files, and web pages containing lyrics and discography information (see Figure 2).

The navigation bar at the top of this figure gives the user access to:

Searching. Direct audio querying of MIDI files plus text and metadata searching over all documents;

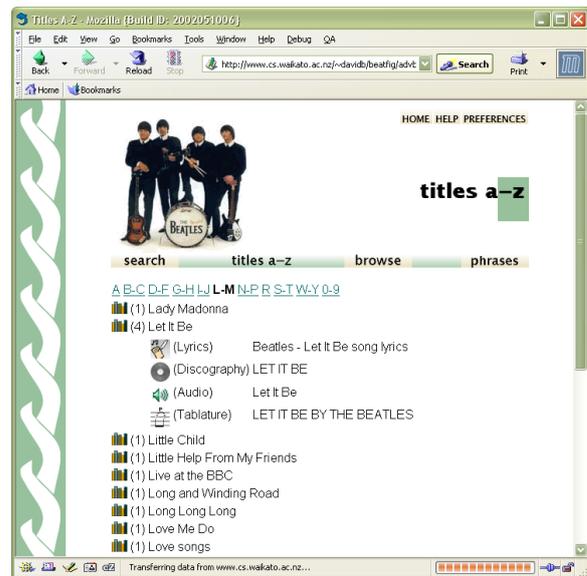


Figure 2. Browsing heterogeneous documents by title.

Titles A-Z. Documents in the collection clustered by title and sorted alphabetically;

Browsing. Documents clustered into groups reflecting document type: audio, lyrics, tablature, or discography;

Phrases. The PHIND phrases mentioned above.

The figure shows the L–M section of titles, where different media types with the same title have been automatically grouped together. Opening up the cluster for the song *Let It Be* we see that there are four items in the collection: lyrics, discography, audio and tablature. Clicking on one of these links yields the respective document.

As part of the poster/demonstration, a variety of collections that illustrate the various aspects described in this paper will be given.

6. REFERENCES

- [1] D. Bainbridge and T. Bell. A music notation construction engine for optical music recognition. *Software Practice and Experience*, 33(2):173–200, 2003.
- [2] D. Bainbridge, C. Nevill-Manning, I. Witten, L. Smith, and R. McNab. Towards a digital library of popular music. In *The 4th ACM conference on Digital Libraries*, pages 161–169, 1999.