Center for Research in Electronic Art Technology
University of California, Santa Barbara
http://create.ucsb.edu

The Open Music Network Infrastructure (OMNI)

Abstract

This proposal describes the Open Music Network Infrastructure (OMNI), an Internet-based music service that aims to provide music content providers with a new forum in which to attract music consumers, enabling the so-called "second music industry." The OMNI system consists of a content provider interfaces, a large-scale artificial-intelligence-assisted "smart" music/sound database, and listener services that allow users to select musical selections based on their personal taste. The most unique feature of OMNI relative to other web-based music services is this use of a smart indexing and search component in the database, which facilitates little-known musicians finding an audience that would like their songs. This document is aimed at a semi-technical reader.

Introduction

Computers, digital audio technology, and wide-area networks are used in all aspects of music performance, production, and distribution today. These new technologies have not changed very much about the way we select and access music, however, nor have they had any impact on the business model that funds music production and distribution. In the words of computer music visionary F. Richard Moore, "the new face of computer-based music is as shallow as it is broad. [...] The field of music as effectively 'absorbed' [technology in music] with little effect on what music is produced or our understanding of it" (*Computer Music Journal*, 20.1 p. 40).

It is the intention of the *Open Music Network Infrastructure* (OMNI; the name is a place-holder) to provide a new web-based framework to make it easier for unsigned musicians (i.e., those without a traditional commercial distribution network) to find an audience. To do this, we also have to provide an economic model that opens the market up to smaller music producers. The OMNI software includes a web-accessible music/sound database, a set of tools for music producers, and a "smart" consumer user interface that provides new services for listeners.

There are several related efforts already underway; these range from pay-per-view web sites for independent musicians, to new forms of music licensing. The OMNI initiative aims to learn from these projects, and to provide a scalable solution that can serve as a clearinghouse for new music.

Motivations

Our interest in the development of flexible and scalable web-accessible sound/music archives is driven by several factors and several target user groups. The first of these is the needs of the "second music industry" that could be enabled through the use of web-based technology and new legal and financial models for music distribution.

Another primary motivation is for the preservation of world cultural heritage and language diversity. A popular cliche asserts that digital networks are transforming the world into a "global village." Cultural globalization is also, however, leading to serious loss of languages and dialects, oral traditions, and local music styles.

The Components of OMNI

The OMNI system consists of three basic components: (a) front-end applications for music content providers; (b) a large-scale "smart" multimedia database; and (c) several kinds of end-user (music consumer or scholar) applications that access the database.

Special database administrator tools are also to be provided.

The simplest producer-side front-end is a web site where content providers can up-load music in any of several encodings. Extensions to this might add features such as multi-format storage, and intelligent encoders that can "listen" to a piece of music and choose the optimal parameters for the encoders (see the discussion of the FASTLab technology below).

The central issues in the design of the OMNI database are scalability, mirroring, and automatic indexing. To make a system such as OMNI truly useful, we need to enable listeners to locate music in the database based on their personal preferences, rather than on genre classification or producer-provided annotation. We propose using the FASTLab technology developed at CREATE in partnership with Predixis, Inc. for this. We will present FASTLab below, and believe that it is of central importance to the success of OMNI.

The consumer-side client programs can be as simple as the basic web sites used in today's music access services, or they can encompass sophisticated user profiles and automatic music selection. As it is one of OMNI's primary goals to expose listeners to new music providers, smart indexing and flexible search techniques are of central importance. Some users (e.g., musicologists) might be allowed to add their own annotational information, while standard users need flexible modes of accessing the OMNI database's contents.

The Figure on the right shows a simplified schematic view of the components of OMNI.

The OMNI Content Provider

For the first stage development of OMNI, we assume a moderately computer-literate content provider such as the average semi-professional musician of today. These users are knowledgeable at their own tools and need a simple web-based user interface for up-loading material to the database and for providing the basic annotational data such as title, performers, copyrights, etc.

The current standard for low-bandwidth music distribution is MP3, but this will soon give way to higher-quality encoding schemes such as AAC (advanced audio coding). We hope to support both standards within OMNI, and to encourage providers and listeners to opt for the higher-quality schemes.

An advanced OMNI content provider interface would support multiple encodings of the same material, as well as supporting "smart" encoders that use the FASTLab technology for better compression. We could also allow providers to interact with the system's analysis and statistics component.

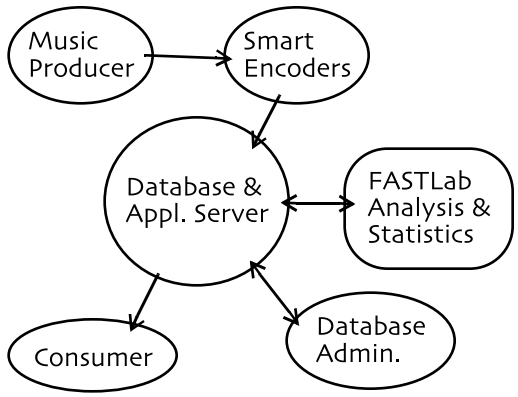

Figure: The OMNI Framework

The OMNI Database

The core of OMNI is a sound / music database that can be accessed via the Internet. To make the system easily accessible, we have to develop web-based front-end applications that access the database and allow users to add their own comments, translations, links, and other annotations of the content. Users can search the database using sound input ("query by humming"), by similarity to other materials ("sounds like"), or by annotational information ("other songs by this artist").

Any number of servers can exist, and can be made open-access or set up to limit access to the members of a specific user community. We plan to build the OMNI server based on off-the-shelf object database, web server, and application development technology, and to integrate components from our previous

projects such as Paleo, HPDM, FASTLab, and Siren (see the CREATE Web site for information on these projects).

There are many options for packaging the content, allowing users to develop profiles and play lists, and providing for "administrative filtering" of the database.

What sets OMNI apart from traditional music databases, and enables its use for the "second music industry," is the use of analysis-based annotation provided by the FASTLab technology. In effect, FASTLab "listens" to musical pieces being added to the OMNI system and creates a multidimensional database with up to 100 features of each selection. This will allow users to find music that they like, even if that music is new in the database or is classified in a genre with which the users is unfamiliar.

The OMNI User

OMNI's goal is to enable the "second music industry," so it is important that users be motivated to listen to unfamiliar artists. To this end, we propose a system that maintains complex user preferences and supports queries such as "play more that sounds like this," or "play a slow country waltz with multivoice harmony and no banjos." This is the primary goal of the FASTLab system and the MusicMagic service that uses it (see below).

The OMNI Economic Model

There are several efforts to provide a new licensing model for music distribution. For example, the German LinuxTag OpenMusic group has developed a series of copyright and distribution licenses for unsigned musicians (see http://openmusic.linuxtag.org). The details of the economic and legal models for OMNI are outside of the scope of this document.

FASTLab Music Analysis

To make OMNI useful, we need a better indexing and search mechanism than that provided by the current crop of web-based music services. We have been working on this together with Predixis, Inc. since 1999 in the FASTLab project and the MusicMagic web application.

The FASTLab system is an analysis tool for use in music and sound databases, preference-matching applications, and other cases where a multidimensional characterization of a musical selection is needed. It uses a number of different analysis techniques to extract features from musical selections. The present version is oriented towards the analysis of popular music songs stored as MP3 files, and the generation of an XML database for statistical analysis and user preference matching.

The goal of the FASTLab software is to derive a number of high-level musical properties from recorded selections so that these properties can be used to match, compare, and contrast pieces of music, either in a database query, a user-specific music selection system, or related applications.

FASTLab is a flexible and open system that uses a range of sound analysis techniques in the time and frequency domains. It also uses several second-level heuristic programming techniques such as rule-based matching, fuzzy logic, and categorization and regression trees (CART). The framework is extensible both in terms of the low-level analysis and feature extraction techniques it uses, and in terms of the second-level pattern recognition, matching, and classification tools that are available. (The name of FASTLab comes from the first names of its three developers: Frode Holm, Alex Kouznetsov, and Stephen Pope.)

The software components of FASTLab are:

- Input decoder
- Rhythm/tempo analyzer
- Spectrum analyzer
- Linear prediction analyzer
- Partial and formant tracking
- Stereo correlation and phase derivation
- Filter module
- Bass pitch tracker
- Instrument signature ID
- Style Matcher
- CART tree analysis/matching
- XML output
- Debugging GUI

Each of these subsystems is implemented by one or more C++ classes.

A separate document describes the details of the FASTLab software; there are several screen shots of the internal debugging screens and prototype end-user screens in the Appendix of this proposal.

MusicMagic pairs the FASTLab analysis engine with a sophisticated set of multidimensional statistics routines that can match a set of user selections with features in the musical feature database. In the simplest scenario the user is presented with a number of excerpts (ten or so), and evaluates them (i.e., like / don't like). Given this input, the system searches for statistical correlations, for example that the user does not like accordions, or likes two-part vocal harmony, for example. This preference data is collected independent of any musical style or genre, i.e., the user is never asked if he / she likes hip-hop. The more preferences we have collected from a user, the better we can characterize his / her preferences. It is also possible for a user to select "surprise me" as an option.

As mentioned above, we have also discussed the use of FASTLab analysis for optimizing MP3 (or AAC) encoders in which the system analyzes a selection and tunes the encoder properties to achieve the maximum compression with the minimum of audible artifacts.

OMNI and LoCAA

A separate project proposal from CREATE describes a system called the *Low-Cost Audio Archivist*, or *LoCAA*, which has a strong relationship with OMNI. Although it addresses different user groups and applications, LoCAA shares several of the OMNI components, including the database and end-user interfaces.

What separates LoCAA from OMNI is the content provider "client" interface. The LoCAA client will be a low-cost (less than $1000 as-delivered) portable recording / distribution input client or "field recording appliance. We envision a portable, low-power computer (palm-top, tablet, or notebook-sized), running easy-to-use applications for sound recording, editing, compression, and wireless up-loading to the server. The LoCAA client will allow non-expert users to record, edit, compress, and distribute non-commercial speech or music content from anywhere in the world.

The content that is up-loaded from these recorders will be stored in a family of Internet-accessible sound / music database servers (like the OMNI servers) that support on-going extension and annotation. This will allow musicologists, linguists, and end-users to add comments, translations, and arbitrary links related to the sound / music content. See the LoCAA proposal for more details.

Project Plan

In the interest of getting the OMNI service up and running soon, we can simplify the requirements in the first phase, and add new features and interfaces later. For the first phase, we would separate platform integration, application development, and user interface prototyping into parallel efforts.

This version will assume a computer-savvy content provider with some sort of personal computer and Internet connection. We estimate that we can have a simple provider interface and smart database ready for deployment in less than six months after project start (or more depending on how we handle the licensing or recreation of the FASTLab components).

In the second phase, we can then add features such as the smart encoders, more flexible end-user interfaces, more powerful FASTLab-based indexing, and database mirroring. During this time, we might also dedicate some resources to the LoCAA components.

Why UCSB?

The CREATE Lab at UCSB has been the home of a number of R&D projects that feed directly into the OMNI effort, including work on ease-of-use user interfaces, analysis and feature extraction for sound / music databases, audio encoding, sound synthesis pro-

gramming languages and applications, wide-area distributed systems, and Internet multimedia streaming (see the CREATE R&D white papers at http://create.ucsb.edu/wp for details of our work on these issues).

Members of our team have also participated in several commercial projects that involve web-based multimedia streaming, sound feature extraction, and complex web-based services.

We are also phase 1 partners on the CalREN-II high-speed network, with Gigabit Ethernet fiber coming into our facility.

Lastly, CREATE is the birthplace of the Digital Media Innovation Initiative (DiMI), a state-wide R&D funding program that provides dollar-for-dollar matching grants for partnerships between California-based corporations and UC-based researchers.

Contact Information

Stephen Travis Pope, CREATE, UCSB
Tel: (+1-805) 967-2621
Email: stp@create.ucsb.edu

Appendix: FASTLab/MusicMagic Screen Shots

The following screen shots are taken from the FASTLab and MusicMagic systems and are included here by permission of Predixis, Inc., Monrovia, California.

Debugging GUI showing the waveform display in the upper pane and the operation of the rhythm/tempo analyzer in the middle two panes. The lower graphical pane shows that the song is in strict 4/4 time.

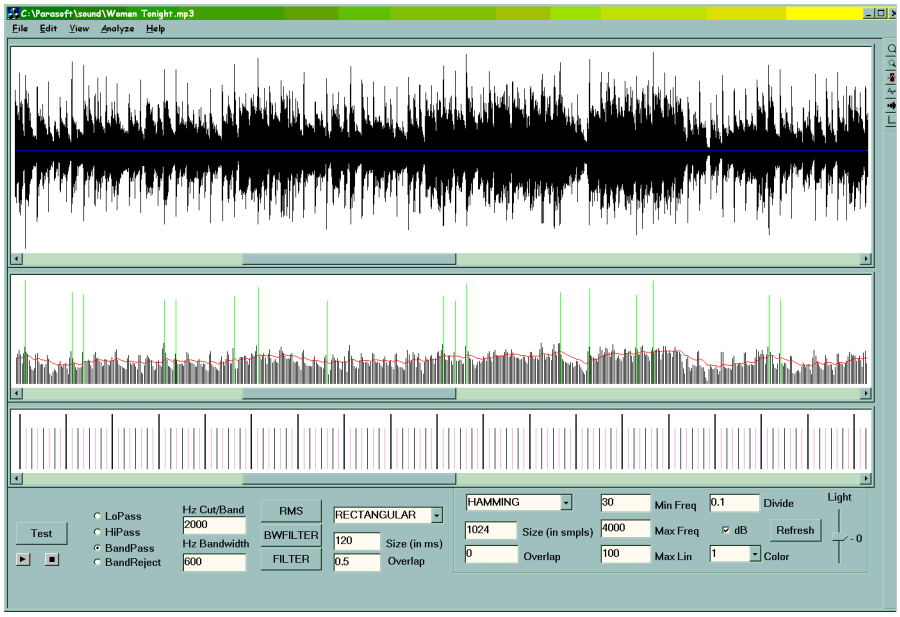

Debugging GUI showing the spectrum tracker pane. In this case, the system can tell that the sound consists of several instruments playing in Western tonal harmony.

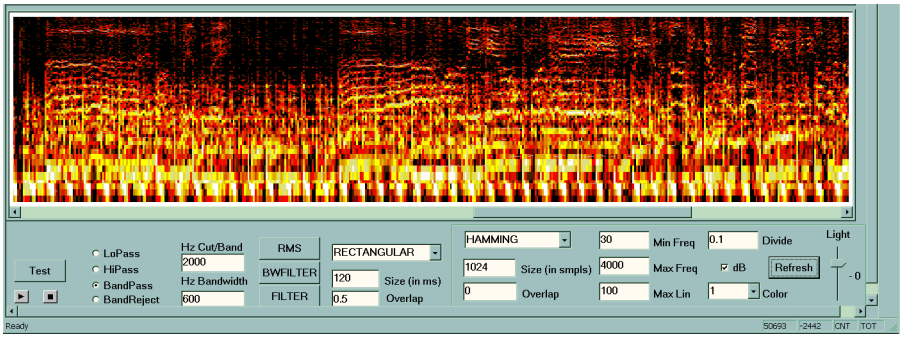

Analysis data comparison view showing some of the FASTLab-derived properties in the left-most list pane and the values for two songs in the middle two panes.

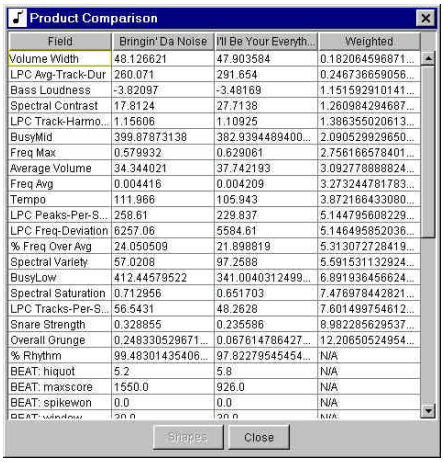

The categorization and regression tree (CART) editing view is used to train the style classifier expert system.

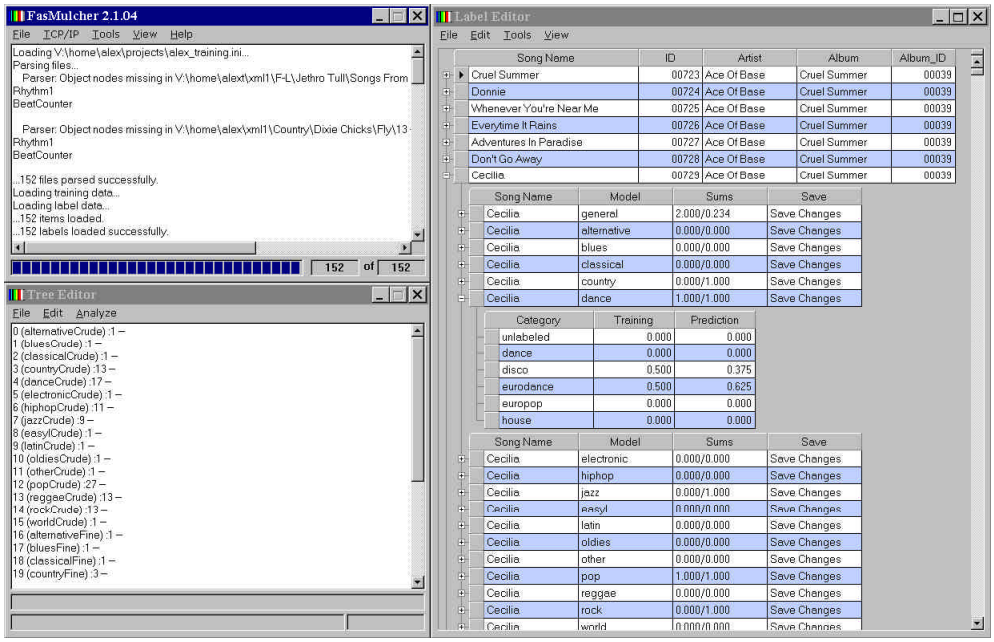

"MusicMagic" listener's playlist view showing suggested songs that match a give selection (i.e., a "sounds like" query).

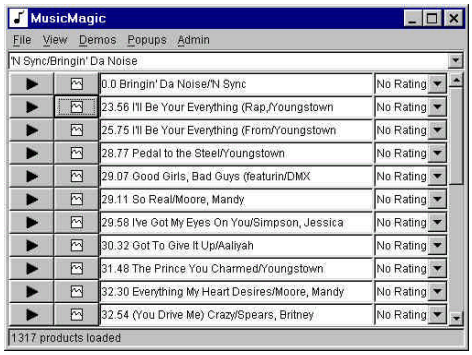

"MusicMagic" playlist view showing user input for preferences (used to refine the system's statistical model of the user).