Music Analysis, Retrieval and Synthesis of Audio Signals MARSYAS

George Tzanetakis
University of Victoria
PO Box 3055, STN CSC
Victoria, BC, CANADA
gtzan@cs.uvic.ca

ABSTRACT

Marsyas is an open source software framework for music analysis, retrieval and synthesis with specific emphasis on Music Information Retreival applications. It has been in development for 10 years and has been used for a variety of projects in both academia and industry in several countries. Based on a novel dataflow architecture named implicit patching it provides a variety of existing processing modules for digital signal processing, machine learning and audio input/output that can be combined at run-time to form complex dataflow networks expressing audio processing algorithms (black-box functionality). In addition it allows the easy addition of new processing modules that need to be compiled for performance purposes. Finally Marsyas is designed with inter-operability in mind and provides various mechanisms for communicating with other software including bindings to the run-time functionality in scripting languages (Python, Ruby), run-time data interchange with MATLAB, support for the Music Instrument Digital Interface (MIDI) protocol and Open Sound Control (OSC) for communicating with controller devices, and infrastructure for easy interfacing to the GUI components of the Qt toolkit.

Categories and Subject Descriptors

H.3.1 [**Information Storage and Retrieval**]: Content Analysis and Indexing Methods

General Terms

Algorithms, Theory, Experimentation

Keywords

sound analysis, music information retrieval

1. INTRODUCTION

Music has always been transformed by advances in technology. Examples of technologies that transformed the way music was produced, distributed and consumed include musical instruments, music notation, recording and more recently digital music storage and distribution. Recently portable digital music players have become a familiar sight and online music sales have been steadily increasing. It is likely that in the near future anyone will be able to access digitally all of recorded music in human history. In order to efficiently interact with the rapidly growing collections of digitally available music it is necessary to develop tools that have some understanding of the actual musical content. Music Information Retrieval (MIR) is an emerging research area that deals with all aspects of organizing and extracting information from music signals.

In the past few years, interest in Music Information Retrieval (MIR) has been steadily increasing. MIR algorithms, especially when analyzing music signals in audio format, typically utilize state-of-the-art signal processing and machine learning algorithms. The large amounts of data that is processed together with the huge computational requirements of audio processing can stress current hardware to its limits. Therefore efficient processing is critical for building functional MIR systems that scale to large collections of music and eventually to all of recorded music. Moreover, MIR is an inherently interdisciplinary field with practitioners with varying degrees of computer and programming expertise (examples of fields involved include musicology, information science, and cognitive psychology). Therefore it is desirable for MIR systems to support multiple hierarchical levels of usage and extensibility. These issues make the design and development of MIR systems and frameworks especially challenging.

MARSYAS (Music Analysis, Retrieval and SYnthesis for Audio Signals), is an open source audio processing framework with specific emphasis on building MIR systems. It has been under development since 1998 and has been used for a variety of projects both in academia and industry. The guiding principle behind the design of MARSYAS has always been to provide a flexible, expressive and extensive framework without sacrificing computational efficiency. Addressing these conflicting requirements is the major challenge facing the software engineer of MIR systems.

2. RELATED WORK

Music Information Retrieval is a new area of content-based multimedia information retrieval. Although there was sporadic earlier work, a good reference starting point is the first international conference on MIR (ISMIR) which was held in 2000. These conferences (ISMIR) have been a forum for bringing together music researchers, audio engineers, computer scientists, musicologists, librarians, and music industry (`http://www.ismir.net`). MIR with audio signals typically requires signal processing and machine learning algorithms in order to achieve tasks such as classification, similarity-retrieval and segmentation.

MM'09, October 19–24, 2009, Beijing, China.
ACM 978-1-60558-608-3/09/10.

The current version of Marsyas evolved from MARSYAS 0.1, a framework that focused mostly on audio analysis. One of the motivating factors for the rewrite of the code and architecture was the desire to add audio synthesis capabilities and was influenced by the design of the Synthesis Toolkit (`http://ccrma.stanford.edu/software/stk/`). Other influences include the powerful but more complex architecture of CLAM (`http://clam-project.org/`) , the patching model and strong timing of Chuck (`http://chuck.cs.princeton.edu/`). The default naming scheme for controls is inspired by the Open Sound Control (OSC) protocol (`http://opensoundcontrol.org`).

The idea of dataflow programming has been fundamental in the design of MARSYAS. Dataflow programming has a long history. The original (and still valid) motivation for research into dataflow was to take advantage of parallelism. Motivated by criticisms of the classical von Neumann hardware architecture dataflow architectures for hardware were proposed as an alternative in the 1970s and 1980s. During the 1990s there was a new direction of growth in the field of dataflow visual programming languages that were domain specific. In such visual languages programming is done by connecting processing objects with wires to create patches. Successful examples include Labview (`http://www.ni.com/labview/`), SimuLink (`http://www.mathworks/com/products/simulink/`) and in the field of Computer Music Max/MSP (`http://www.cycling74.com/products/max5`) and Pure Data (`http://puredata.info/Puckette,2002`).

3. EVIDENCE OF IMPACT

The impact of Marsyas has been steadily increasing especially since 2007 where a more comprehensive website was launched and the core team increased to approximately **4-5 developers**. Since November **2007** when the new website was launched it has received 18236 (approximately **30 visits/day**) from **114 countries**. The software framework has approximately **300-400 downloads** per month since **2007** an increase from less than 50 downloads/month in 2003.

Marsyas has been used for a variety of projects in academia and industry several of which are described in `http://marsyas.sness.net/about/projects`. Industrial applications include prototyping the patented audio fingerprinting technology of Moodlogic Inc. which has been used by more than **100,000** users to link mp3 audio files to a large database (**1.5 million songs**) of metadata and a gender identification system for voice messages designed for Teligence Inc. that is processing approximately **25,000 voice recordings per day**. Although Marsyas is available under the GNU Public Licence (GPL) it can also be provided with commercial licenses that facilitate development of closed-source applications.

Academic projects include emotion recognition in music (Greece, Canada), novel interfaces for browsing music (Japan), a dancing music robot (Portugal), and a multi-model search engine for music in YouTube videos (Singapore). A large number of publications have been written about Marsyas (6), using Marsyas (31) and citing Marsyas (12). More details can be found at `http://marsyas.sness.net/about/publications`. Algorithms based on Marsyas are regularly submitted to the Music Information Retrieval Evaluation Exchange (`http://www.music-ir.org/mirex/2009/index.php/Main_Page`) where they exhibit state-of-the-art performance while being orders of magnitude faster than other submissions.

4. EXPLORING MARSYAS

Marsyas is a large open source software framework that can be used in a variety of different ways therefore there is no single easy path to learning the software and its capabilities. Detailed downloading, compiling and installation instructions are available at `http://marsyas.sness.net/docs/manual/marsyas-user/Source-installation.html`. A taste of different applications, projects and usage scenarios can be found by looking at the videos and webdemos (`http://marsyas.sness.net/about/videos`, `http://marsyas/sness.net/about/webdemos`). The documentation consists of a user manual, developer manual, library reference (automatically generated from the source code) as well as a cookbook that contains short code examples for simple tasks. The tour chapter of the user manual provides information about how to run several examples of the tools included with the Marsyas distribution for demonstration purposes.

5. CONTRIBUTORS

As with most open source software the contributors to Marsyas have fluctuated over the years. Currently there is a core team of approximately **4-5** developers as well as a more extended community of approximately **20-30** regular users who occassionally contribute bug fixes and small amounts of code. The current core development team and their current (June 2009) affiliations consists of:

- George Tzanetakis - University of Victoria, Victoria, Canada - *gtzan@cs.uvic.ca*
- Luis Gustavo Martins - Catolica University, Porto, Portugal - *lgustavomartins@gmail.com*
- Luis Teixeira - Fraunhofer Centre for Assistive Information and Communication Solutions (AICOS), Porto Portugal - *luis.f.teixeira@gmail.com*
- Graham Percival - National University of Singapore Singapore - *gpermus@gmail.com*
- Mathieu Lagrange - TELECOM ParisTech, Paris, France - *mathieu.lagrange@enst.fr*
- Steven Ness - University of Victoria, Victoria, Canada - *sness@sness.net*
- Emiru Tsunoo - University of Tokyo, Tokyo, Japan - *tsunoo@hil.t.u-tokyo.ac.jp*
- Stefaan Lippens - Ghent University, Ghent, Belgium - *stefaan.lippens@gmail.com*

More details about previous developers and more extended lits of users and projects can be found at the marsyas webpage `http://marsyas.sness.net`.