ISMIR 2008 Tutorial

MUSIC INFORMATION RETRIEVAL IN CHUCK REAL-TIME PROTOTYPING FOR MIR SYSTEMS AND PERFORMANCE

Ge Wang
Stanford University
Center for Computer Research in Music and Acoustics
ge@ccrma.stanford.edu

Rebecca Fiebrink
Princeton University
Computer Science
fiebrink@princeton.edu

Perry R. Cook
Princeton University
Computer Science, also Music
prc@cs.princeton.edu

ABSTRACT

We present a hands-on ISMIR tutorial focused on the free, open-source ChucK programming language for music analysis, synthesis, learning, and prototyping. Our goal is to familiarize the ISMIR audience with ChucK's new capabilities for MIR prototyping and real-time performance systems, and to stimulate discussion on future directions for language development and toolkit repository contents.

1. OUTLINE OF CONTENT

Part 1: Introduction to ChucK
This is a hands-on introduction to ChucK, in which participants with laptops can code along as we introduce the language's key syntax, the unit generator patching system, timing and concurrency model, and on-the-fly programming and prototyping practice using the miniAudicle development environment. Participants will understand at a high level how ChucK has been used for composition and performance (including with the Princeton and Stanford Laptop Orchestras) and music technology education at Princeton and Stanford.

Part 2: Audio analysis in ChucK
We demonstrate how to perform audio analysis in the language, including applying spectral transformations, extraction of standard MIR features, and designing low-level signal processing and feature extraction from scratch in the language. We will have fun applying ChucK's on-the-fly and rapid prototyping to analysis, demonstrating how feature extraction algorithms and parameters can be changed during the processing of live audio.

Part 3: Classification and learning in ChucK
We lead the group through experimentation with applying ChucK's classification and learning algorithms to off-line and real-time audio. For example, we use built-in capabilities to rapidly construct an instrument classifier and an artist classifier. Again, we stress the rapid prototyping capabilities of the language, and miniAudicle on-the-fly development and ChucK's standard classifier building blocks, as well as demonstrate how users can modify built-in classifiers and implement their own. We discuss exciting issues in applying classification to real-time performance, including on-the-fly learning.

Part 4: Discussion and Q&A
At this point, we discuss larger issues of applying MIR in live performance, and if possible take questions from participants regarding how they might implement ideas they have using ChucK, potential future directions for MIR-related ChucK development, as well as using MIR for live performance.

2. AUDIENCE

The intended audience includes people who meet any of the following criteria:

- People who work on real-time systems, including those for score following, computer accompaniment, real-time beat tracking, etc.
- People who would like to learn how to port their non-real-time code into a real-time system, to be used by performers and composers
- People who are interested in an alternative MIR prototyping framework that supports rapid feedback and a combination of low- and high-level programmability
- People who teach courses on computer music, music technology, signal processing, or other areas related to MIR who would like to learn about ChucK as a teaching tool
- People who have heard about the Princeton and Stanford laptop orchestras, would like to know more about the technologies we use, and are curious how MIR tools have been and could be applied in this type of setting.

We do not expect participants to have had any previous ChucK experience, though programming experience in some language will be helpful.

3. PRESENTER BIOS

Ge Wang is an assistant professor at Stanford University's Center for Computer Research in Music and Acoustics (CCRMA). His research interests including interactive software systems for computer music, programming languages, sound synthesis and analysis, music information retrieval, new performance ensembles (e.g., laptop orchestras) and paradigms (e.g., live coding), human-computer interaction, audio over networks, and education at the intersection of computer science and music. Ge is the chief architect of ChucK, a founding developer and co-director of the Princeton Laptop Orchestra (PLOrk), and founder and director of the Stanford Laptop Orchestra (SLOrk). He performs with PLOrk/SLOrk, as well as with Perry Cook in a live coding duo, and with Rebecca Fiebrink in a duo exploring new performance, software tools and paradigms for expressive music-making, and great food.

Rebecca Fiebrink is a PhD student in Computer Science at Princeton University. She has worked with Ge Wang and Perry Cook to implement audio analysis and machine learning in ChucK. She holds an MA in Music Technology from McGill University, and she has been involved in the ISMIR community for several years, with research interests spanning music classification and recommendation, applied machine learning, audio fingerprinting, and computer music programming and interface design. She is currently a co-director of PLOrk and a flute and laptop performer.

Perry R. Cook attended the University of Missouri at Kansas City Conservatory of Music from 1973 to 1977, studying voice and electronic music, eventually (but only after working as a sound engineer and designer from 1976-1981) getting degrees in both Music (1985) and Electrical Engineering (1986), whereafter he went to Stanford CCRMA, getting a Masters and PhD in EE in 1990, continuing as Technical Director of CCRMA and also consulting/working for NeXT, Media Vision, Xenon/Chromatic, Interval Research, and other companies which he still worked with some after he joined Princeton University as a Professor (not at first, he started there as an Assistant Professor, but then they let him Associate with the Professors, then let him Profess himself) of Computer Science with a joint appointment in Music, doing stuff such as publishing lots of papers and books, getting a Guggenheim Fellowship in 2003, and founding the Princeton Laptop Orchestra with Dan Trueman, which just received a MacArthur Digital Learning Initiative grant. Perry also sings.