TEACHING STEM CONCEPTS THROUGH MUSIC TECHNOLOGY AND DSP

Youngmoo E. Kim, Alyssa M. Batula, Raymond Migneco, Patrick Richardson, Brian Dolhansky,
David Grunberg, Brandon Morton, Matthew Prockup, Erik M. Schmidt, and Jeffrey Scott

Drexel University
Electrical & Computer Engineering
{ykim, batulaa, rmigneco, patrickr, bdol, dgrunberg,
bmorton, mprockup, eschmidt, jjscott}@drexel.edu

ABSTRACT

Music is an integral part of high school students' daily lives, and most use digital music devices and services. The one-week Summer Music Technology (SMT) program at Drexel University introduces underclassmen high school students to music technology to reveal the influence and importance of engineering, science, and mathematics. By engaging participants' affinity for music, we hope to motivate and catalyze curiosity in science and technology. The curriculum emphasizes signal processing concepts, tools, and methods through hands-on activities and individual projects and leverages computer-based learning and open-source software in most activities. Since the program began in 2006, SMT has enrolled nearly 100 high school students and further developed the communication and teaching skills of nearly 20 graduate and undergraduate engineering students serving as core instructors. The program also serves to attract students from backgrounds under-represented in engineering, math, and science who may not have considered these fields.

Index Terms— Education, signal processing, music

1. INTRODUCTION

The Summer Music Technology (SMT) program at Drexel is a novel educational experience, designed to introduce 9th- and 10th-grade students to modern music technology to reveal the influence and importance of engineering, science, and mathematics. Music is an integral part of these students' daily lives, and the vast majority use digital music devices and services and possess large personal libraries. By engaging our participants' affinity for music technologies, we hope to further motivate and catalyze curiosity in science and technology. The week-long curriculum emphasizes signal processing concepts, tools, and methods through hands-on activities and individual projects and leverages computer-based learning and open-source software in most learning modules.

The program, initiated in 2006 as part of an NSF CAREER award, has enrolled nearly 100 students and recently completed its fourth session. The curriculum has been revised each year with new material and includes significant contributions from graduate and undergraduate engineering students, who form the core team of instructors for the program, further developing their communication and teaching skills. The program also serves to attract students from backgrounds under-represented in engineering, math, and science who may not have previously considered further study in these fields.

Current topics covered in the program are as follows: learning how music is recorded and produced, transduction of signals to sound using speakers, the wave-based nature of sounds, creating new sounds and sound effects, musical instrument acoustics, using novel interfaces for music manipulation, digital signal representations, and music information retrieval. These modules are designed to introduce students to core signal processing concepts without relying on background knowledge. The program is primarily inquiry-based and strives to demonstrate and teach new concepts through hands-on activities and projects rather than formal lectures.

2. BACKGROUND

High school students today have access to enormous databases of digital music content and, having a seemingly limitless appetite for music, many are frequent users of products and services incorporating highly-sophisticated signal processing (iPods, YouTube, and digital audio editing software). These techniques, tools, and aesthetics permeate almost all forms of modern music today, yet the current generation remains largely unaware of the science, engineering, and computing within that heritage. Thus, it is unsurprising (though somewhat ironic) that while we are witnessing an explosion of interest in digital music content creation and manipulation, there are two educational trends that run counter to this. First, it is difficult to retain interest in the fields of science, technology, engineering, and mathematics (STEM) after high school, particularly among women and minorities [1]. Second, in-

structional time devoted to music and arts education in public schools is decreasing [2].

One proposed method of motivating interest in STEM is to connect everyday technology to mathematical and scientific concepts, placing them in a context familiar and appealing to students. The INFINITY Project [3] pioneered a collaboration between universities, high schools, and industry leaders to establish a project-based engineering curriculum at the high school level in order to motivate and attract students. Digital music curricula have also been demonstrated to be a positive and engaging educational gateway in public and at-risk high schools in other countries. For example, digital music workshops in British schools bring together the youth culture of musical "chatting" (British rapping) with free music-creation software usable on their school computers [4].

It is widely believed that musical affinity and engineering and mathematical prowess are closely linked. Howard Gardner's well-known pedagogical work on the theory of Multiple Intelligences suggests that music may be one of the core methods through which children relate to the world and thus learn from their experiences [5]. Similarly, Jeanne Bamberger's research has shown that study of coherent musical structure can provide contexts for better learning of basic and intermediate math skills, and that this advantage is enhanced when taught through interactive multimedia of learning software [6]. Long before computers became a viable tool for modern music creation, the urban Hip-Hop movement of the 1980s set an important musicological precedent. Inner-city youth without formal music education found new ways to compose and create music by re-purposing music technology, first in the analog domain (record turntables) and later using digital technology (audio sampling) [7].

The SMT program attempts to address the disturbing trends in both STEM and music education using a cross-disciplinary approach. Our program has attracted students with no prior interest in a technology career, but who enjoy music technology. We believe that introducing such students to STEM concepts in the context of music has broadened their perspectives to include new career paths involving STEM.

3. EQUIPMENT

In order to complete the activities, students are provided with all of the materials and equipment that they need.

3.1. Hardware

Each student is provided a laptop to use during the program. They also each receive a USB drive so that they can store their files on it and have them on any computer. We choose to do this, rather than having students share computers, so that each student has an opportunity to do the entire activity. This makes it more hands-on for the students and prevents one group "leader" from doing all of the work.

3.2. Pure Data software

Pure Data[1] (Pd) is an open-source real-time programming environment for audio, video, and graphical processing initially developed by Miller Puckette (who also created the similar, commercially available Max/MSP application). We chose to base much of our curriculum on this software primarily because of its real-time audio processing capabilities, which enables direct manipulation of algorithms and parameters and provides instantaneous feedback to students. Pd's graphical programming environment is also much easier for students to comprehend without prior programming experience. The graphical programs (called "patches") allow students to visualize signal flow within an algorithm. Since Pd is open-source and freely available, interested and motivated students can download it to their computers at home, allowing students to continue working on their projects overnight. Some continue to use the application even after the SMT program concludes.

Because Pd is used extensively in this program, we developed basic tutorials to teach the students the fundamental Pd components, such as arithmetic operators, connections and signal flow, plotting, and sound input/output as shown in Figure 1. These tutorials were incorporated into some of the early curricular modules and the fundamental concepts were reinforced in other modules throughout the week.
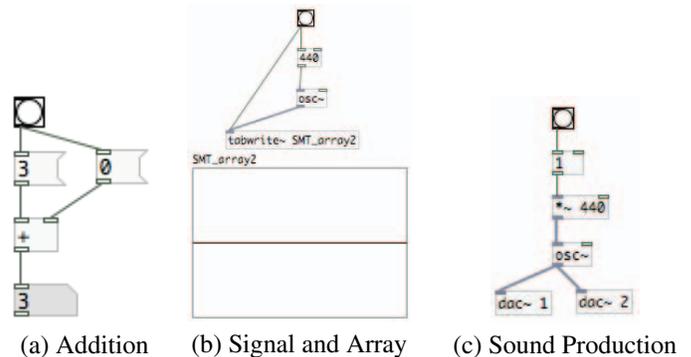

(a) Addition   (b) Signal and Array   (c) Sound Production

Fig. 1. Example Pd patches used in the tutorials.

4. DESCRIPTION OF MODULES

In general, the SMT modules are designed to promote hands-on, inquiry-based learning. Whenever possible, the activities lead students to discover concepts for themselves rather than being fed the "right" answers. Students also give presentations on their projects and findings, as shown in Figure 2.

In this section, we describe the modules presented in the most recent (2010) program iteration, which included 24 student participants. Each lesson plan is presented in terms of
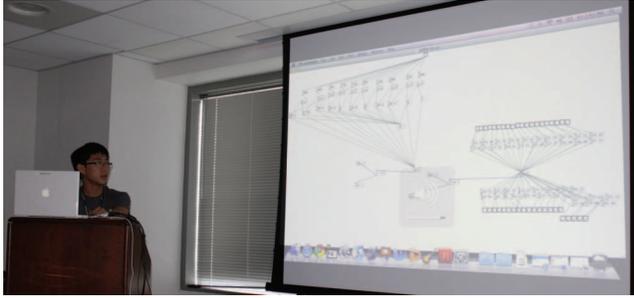
---

[1]http://puredata.info

Fig. 2. A student presents his work on a Pd patch.

the educational objectives, required materials, procedure, and student feedback. The full lesson plans are available online[2].

Basic evaluations were collected after each module, asking each student to provide their responses to the following questions (using 5-point scale):

- How much did you learn from this activity?

- How much did you enjoy this activity?

The overall mean responses to these questions (spanning all activities) were 3.92 and 4.16, respectively, providing context for the individual module ratings.

4.1. Sound transduction and speaker building

In this activity the students work in groups to build a working speaker from household materials and drive it with an amplified audio signal to produce sound.

Educational objectives
After this activity, students should understand how the speaker acts as a transducer, converting electrical energy into mechanical energy. They should also know the basic principles of magnetic fields and electromagnetic induction.

Required materials
Most of the materials are common household items including cardboard boxes, paper/styrofoam plates, magnets, tape, and sandpaper. The amplifiers were custom built by the SMT staff to be able to drive the low impedance coils produced by the students.

Procedure
Groups first verify the physical principles behind speaker operation by moving a magnet about a pre-made coil that is connected to an oscilloscope. A pair of students then assemble a coil by wrapping the magnetic wire around a cylindrical object. The group mounts the coil to a paper or styrofoam plate and mounts the speaker in the housing (cardboard box).

Students then proceed to connect the coil to the amplifier and hold a permanent magnet near the voice coil. Afterwards, they evaluate the sound produced and the instructors

give feedback about how to improve their coil to make a better sound. This could also be made into a competition in order to see who can design the best performing speaker.

Participant feedback
The participants' responses for learning (4.0) and enjoyment (4.2) for this activity are close to the overall program averages. Qualitatively, the students were impressed that sound could be produced from such basic materials.

4.2. Waves and sounds

This activity allows students to explore the nature of sound as well as how basic periodic signals can be created and viewed using a function generator, oscilloscope, and the Pd application. Students are also exposed to the frequency limits of human hearing and the concept of harmonics (overtones).

Educational objectives
After the activity, students should understand that sound is a pressure wave moving through a medium (air) and how to estimate the parameters of a periodic signal (amplitude and frequency/wavelength). Furthermore, participants should recognize the importance of sinusoids (sine wave) and how to create a complex, single-pitch tone using a sum of sine waves.

Required materials
This module requires an oscilloscope, a function generator, a speaker, and a laptop, preferably one for each student. A Pd patch designed for the activity allows the students to draw their own waveforms and the ability to listen to and measure the resulting sounds.

Procedure
Students are given a brief introduction, asking them to develop a working definition of sound and its properties using real-world examples (musical instruments, loud and quiet noises, and digital music players). Students use a function generator to output a tone to a speaker and then view and measure the signal on an oscilloscope. They use this configuration to attempt to determine the limits of human hearing by adjusting the frequency value of the function generator. Multiple frequencies are also played simultaneously to demonstrate the creation of complex tones.

The students then reverse the configuration, using the speaker as a microphone and viewing the resulting signal on the oscilloscope, revealing that a microphone consists of essentially the same components as a speaker, with the transduction occurring in the opposite direction.

Finally the students use a Pd patch on their laptops to draw one period of a wave. These signals are made periodic and the resulting spectrum is displayed as the sound is output over headphones. The interface invites them to explore how different wave shapes affect the strengths of the harmonics present and the overall timbre of the sound.

---

[2]http://music.ece.drexel.edu/smt/curriculum

The mean response to the assessment question on learning (4.3) was higher than the program average (second-highest overall among modules). The ratings for enjoyment (4.1) were close to the overall SMT average.

4.3. Echo and sound design

In this activity, students create a software simulation of an echo, which they use to explore real-world acoustics.

Educational objectives
At the conclusion of this lesson, students should understand that echoes are a combination of direct sound and delayed copies and that the delay time is a function of distance and the speed of sound. They should also gain some intuition regarding the effect of the delay time and feedback model parameters on the resulting sound.

Required materials
This module employs a custom Pd patch that allows students to experiment with echo model parameters. It also serves as a starting point for their echo simulation.

Procedure
The lesson begins by using GarageBand to experiment with shifting copies of a sound and listening to the echo-like result. They then move outside of the classroom where a real-world echo can be heard to provide students with a sense of the echo's timing. Returning to the classroom, students use Pd to explore the echo simulation.

Students start with a single tap echo and timing control to discover the linear nature of delay time and "sense of distance." Students then increase the complexity of the model by creating more realistic reverberation using multiple taps. The final phase reverts to a single tap with feedback, which sums attenuated delay copies of the sound with the original input. Multiple reflections of a sound source are a physical property of reverberant environments and implementation of this using feedback is key to producing richer simulations with a much simpler signal flow. The parameters derived through simulation are then compared to those of the original physical space.

Participant feedback
This activity rated higher than average in student learning feedback (4.4), but was tied for the lowest score in terms of enjoyment (3.8). Several students indicated that the introduction and manipulation of feedback was the key to their understanding and exploration of realistic echo simulation.

4.4. Musical instrument acoustics

This activity, divided into two lessons, examines sound production in musical instruments. Students use simple tubes to explore the concepts of resonant frequencies and filtering. They also examine common musical instruments to determine how they produce sound.

Educational objectives
Students should understand the inverse relationship between a tube's length and its natural resonant frequencies and how the resonances of acoustic systems naturally filter a complex excitation source. Students should be able to apply these concepts to acoustic musical instruments and have a basic understanding of how the distribution of spectral and temporal energies relates to the timbre of the instruments.

Required materials
Students use PVC pipe, a speaker, an amplifier, laptops, a microphone, acoustic instruments, and two custom Pd patches. The first patch provides a signal generator and spectrum analyzer and the second allows students to record sound input and visualize the waveform and frequency spectrum.

Procedure
Students play a tone through a speaker connected to one end of a tube and record the sound at the opposite end. They find the tube's resonances by adjusting the tone's frequency and noting the peaks in the frequency spectrum of the tube output. They then measure the length of the tube and use a formula to compute the theoretical resonant frequencies in order to verify experimental results. Students then take a tube with a different length and drive it with white noise, identifying the peaks in the spectrum corresponding to the resonances.

In the second part of the activity, small groups of students work with a musical instrument to determine how the sound originates and propagates. They explore how the shape, material, and structure of the instrument affects its sound. They also record the instrument and view the waveform and frequency spectrum using a Pd patch, noting how musical pitch and frequency spacing are related. Finally, each group presents their findings for their instrument. In the discussion, the instructors seek to highlight the relationship between the spectral and temporal envelopes of an instrument and the resulting timbre.

Participant feedback
In terms of learning, these were some of the lowest rated activities (3.6 and 3.7 for the two parts, respectively). They also rated as less enjoyable than other activities (3.9 and 3.8). Students enjoyed learning about the instruments, but not presenting their findings. They were also surprised that the tubes could selectively "filter" resonances from noise.

4.5. Analog and digital representations

This session uses a Pictionary-style game show format to teach students how computers store and reproduce continuous audio signals. Students are asked to convey continuous signals to one another using only discrete information.

Educational objectives
After the activity, students should understand why computers cannot directly store continuous information and how such

signals are discretized via sampling. They also obtain insight into the effect of sampling rates on digitized sound quality.

Required materials
The activity requires two video projectors and two large white boards to project and draw upon. Each team receives a packet with the waveforms used in the activity. A custom application developed using Adobe Flash is used to project the grids, "ground truth" waveforms, and game timer.

Procedure
A brief introduction reviews the continuous nature of acoustic waves and why computers cannot record these signals directly. Students are divided into two equal sized teams and paired with a team-member (a "transmitter" and a "receiver"). The game show is divided into rounds in which a transmitter-receiver pair competes against a pair from the other team.

Every round, an empty grid is projected onto each whiteboard. The transmitter from each pair describes the round's waveform to their receiver abiding by round-specific guidelines (e.g., using only words, only numbers, etc.). The receiving partner draws the waveform on top of the projected grid. Whichever team draws the most accurate waveform, as judged by the instructors, wins the round.

Initially, students describe the waveforms using strictly qualitative terms. In the middle rounds, students can only name coordinate pairs that the waveform touches. In the final rounds, students are restricted to a method that mimics sampling: moving from left to right at fixed intervals (representing different sampling rates), students state the y-value of the wave. At the conclusion of the game, students discuss which of the methods is most accurate and representative of how a computer stores sound. The advantages and tradeoffs of higher vs. lower sampling rates is also discussed.

Participant feedback
This activity was significantly revised from the previous year's activity, which involved individual parings as opposed to a group game. Students reporting learning slightly less than the program average (3.8, down significantly from the previous year's score of 4.8). This may be due to the team nature of the game, which provided students fewer chances to participate. Students did, however, rate this as the most enjoyable activity (4.6, up from 4.4 in the previous year).

4.6. Music interfaces

This module explores the use of nontraditional interfaces and computer-based instruments in music. Students discover new ways of interacting with music using the interfaces to manipulate sound in Pd.

Educational objectives
After the activity, students should gain an appreciation of the wide variation in input controllers and some of their basic characteristics, such as degrees of freedom and whether they are continuous or discrete. They should also be able to view raw sensor input (in Pure Data) and use simple mathematical operations to adjust the values for use in sound synthesis.

Required materials
This activity uses five unique interfaces connected to a common Pd sound generation patch for consistency. The interfaces used are: a Nintendo Wii controller, an Xbox Guitar Hero controller, a Korg KP3 Kaoss Pad, a Mercurial STC-1000 pressure sensitive touch pad, and an iPod Touch multiple point controller. It also uses the students' laptops to obtain accelerometer sensor readings.

Procedure
Students are first introduced to common interfaces (such as a music keyboard) and then shown how other electronic devices can be used as musical interfaces. Basic interface concepts are discussed such as degrees of freedom and discrete (buttons and switches) versus continuous (knobs and sliders) controls.

Students are divided into small groups, which are given an interface to discover how it interacts with a sound synthesis Pd patch. In particular, they investigate the types of controls used by the interface (discrete or continuous), how many degrees of freedom each control possesses, and how those controls specifically affect the sound synthesizer output.

In the second part of the activity, students use their laptop's sensors (such as the accelerometer and touchpad) to control a provided Pd synthesizer patch. Students must determine appropriate methods for connecting each of the five sensors to the synthesizer and scaling the raw sensor input into reasonable ranges to achieve the desired results.

Participant feedback
The students' scores for learning (4.0) and enjoyment (4.2) for the first half of the activity are very close to the overall program averages. For the second part, the learning feedback score (3.6) was the lowest overall, while the mean enjoyment score was slightly above average (4.3)

4.7. Music information retrieval

The emerging field of Music Information Retrieval (Music-IR) includes such topics as music recommendation and playlist generation. This module demonstrates the methods used in several popular music recommendation services (Pandora[3], iTunes Genius[4], and Last.fm[5]).

Educational objectives
Students should appreciate the ambiguity of the song "similarity" problem and the motivation for content-driven (signal processing-based) systems. They should also understand the methods behind several music recommendation services.

---

[3] http://www.pandora.com
[4] http://www.itunes.com
[5] http://last.fm

Required materials

A data set of twenty popular songs are provided on a USB flash drive, and computers and headphones are also required. Handouts are provided to each of four groups with the instructions to generate their playlists.

Procedure

Each group is given the same set of twenty songs. Given an initial seed song, each group creates a playlist using the method of a popular service provider (except for the manual group). The final playlist for each group consists of ten songs and no songs may be repeated.

Students in the *Manual Group* listen to each song and agree on the best ten-song playlist. Students in the *Features Group* (Pandora) each rate six songs according to a set of ten music features selected from Pandora's *Music Genome Project*. To determine the group playlist they compute the average feature ratings for each song and, starting with the seed song, quantitatively determine the most similar song to the current song (e.g., using Euclidean distance).

Those in the *Tags Group* (Last.fm) each assign descriptive tags to six songs using any words or phrases they feel best describe the music. They compile the tags for all twenty songs and construct a playlist by picking the songs with the most tags in common. Each student in the *Collaborative Filtering Group* (iTunes Genius) is assigned a group of four-song sets to manually arrange into the best possible playlists. The master playlist is generated by choosing songs that appear most frequently after the current song in the mini-playlists.

After the activity the groups present their playlists and discuss issues with the methods used. The manual group presents their playlist first and explain their choices (generally based on key, modality, and tempo), which sets the stage for evaluating the other playlists. We motivate content-based methods, discussing the scale of large music services.

Participant feedback

The Music-IR activity was changed significantly from previous years, which involved the class collaborating to compile a manual playlist and comparing that to one generated by content driven methods. The mean learning score was less than the previous year's (3.9, down from 4.2), but was average for the current cohort. Students rated this activity as enjoyable as the previous year (both 4.5) and one of the most enjoyable activities overall.

5. DISCUSSIONS AND FUTURE WORK

These modules represent our advances developing a program to interest high school underclassmen in STEM fields. We have refined and expanded upon these modules for four years and will continue to improve upon the lessons in future programs. Every year, we have updated the activities in order to keep them relevant to students' current interests. Now that the modules have been tested for several years, we intend to make these activities available to high school teachers so they can be integrated into the classroom setting.

This is the first year that we have used Pd almost exclusively in all of the modules. Previously we used a variety of web-based and Adobe Flash activities [8]. Switching to a single environment allowed us to further integrate the different sections and enabled students to attain mastery of a single programming environment rather than receiving cursory exposure to multiple tools and languages.

In addition to improving the modules themselves, we are also striving to embed more assessment into the activities. Since most activities are computer-based, we also have the opportunity to use our software to collect low-level assessment data as they work, such as how long students spend on each part of the activity or saving snapshots of their work at preset intervals. Short computer-based pre- and post-quizzes for the students would also provide insight into how much they learned throughout the week. Having students present their findings or asking them questions throughout the activity are also options for obtaining more feedback.

6. REFERENCES

[1] Higher Education Research Institute, "Degrees of success: Bachelor's degree completion rates among initial stem majors," 2010.

[2] Center on Education Policy, "From the capital to the classroom: Year 4 of the No Child Left Behind Act," 2006.

[3] Scott C. Douglas, "The infinity project: Digital signal processing and digital music in high school engineering education," in *IEEE Workshop on the Applications of Signal Processing to Audio and Acoustics*, New Paltz, NY, 2001.

[4] John Finney and Pamela Burnard, Eds., *Music education with digital technology*, Continuum, 2007.

[5] Howard Gardner, *Multiple Intelligences: New Horizons*, Basic Books, New York, 2006.

[6] Jeanne Bamberger and Andrea diSessa, "Music as embodied mathematics: a study of mutually informing affinity," *International Journal of Computers for Mathematical Learning*, vol. 8, no. 2, pp. 123–160, 2004.

[7] Virgil Moorefield, *The Producer as Composer*, MIT Press, 2005.

[8] Y. E. Kim, T. M. Doll, and R. V. Migneco, "Collaborative online activities for acoustics education and psychoacoustic data collection," *IEEE transactions on learning technologies*, 2009, preprint.