

BlockyTalky: Tangible Distributed Computer Music for Kids

R. Benjamin Shapiro, Annie Kelly

University of Colorado Boulder

Matthew Ahrens

Tufts University

Benjamin Johnson, Heather Politi

Boulder Valley School District

Rebecca Fiebrink

Goldsmiths, University of London

Abstract

Computer music research realizes a vision of performance by means of computational expression, linking body and space to sound and imagery through eclectic forms of sensing and interaction. This vision could dramatically impact computer science education, simultaneously modernizing the field and drawing in diverse new participants. In this article, we describe our work creating an interactive computer music toolkit for kids called BlockyTalky. This toolkit enables users to create networks of sensing devices and synthesizers, and to program the musical and interactive behaviors of these devices. We also describe our work with two middle school teachers to co-design and deploy a curriculum for 11- to 13-year-old students. We draw on work with these students to evidence how computer music can support learning about computer science concepts and change students' perceptions of computing. We conclude by outlining some remaining questions around how computer music and computer science may best be linked to provide transformative educational experiences.

Introduction

The creation of new digital musical instruments exemplifies how computational practices can be simultaneously creative, expressive, and technologically rich. Much research and creative work in computer music investigates new ways for computational systems to enable musical expression and interaction. Reciprocally, the task of creating interactive systems for performing computer music functions as a sort of “extreme human-computer interaction,” pushing the limits of existing computing systems and demonstrating possibilities for technological advancement (Tubb 2016). Despite this boundary-pushing relationship between Computer Science (CS) and music, the innovations of the computer music community have not yet become a widespread part of computer science education.

We believe that computing education could benefit from integrating topics central to computer music performance into CS curricula, and from adopting the computer music community’s understanding of computing as an expressive, creative domain.

First, computer music is a uniquely strong application domain through which we could modernize computer science curricula. Computing education research scholars as well as industry have bemoaned the absence of contemporary computing topics like distributed systems and concurrency from computer science education (Patterson 2006). Recent standards documents like the ACM-IEEE Computer Science Curricula 2013 (ACM/IEEE-CS Joint Task Force on Computing Curricula 2013) mandate these topics’ inclusion within CS education pathways, but compelling examples of how to do so are currently lacking. In contrast, many computer music performance systems rely on timing-sensitive distributed computing, concurrency, and networking.

Next, computer music also offers a chance to combat pervasive and insidious misperceptions of computer science and computer scientists. American students and adults have negative stereotypes of science and scientists, often believing that they are socially distant, dangerous, workaholic, peculiar, irreligious, and missing fun in their lives (Mason, Kahle, and Gardner 1991; Losh 2010). They tend to hold similar

stereotypes of computer scientists (Martin 2004; Carter 2006; Grover, Pea, and Cooper 2014). These stereotypes can be challenged by educational approaches that bring scientific methods together with topics that are relevant to people's lives and that showcase the wide range of possible pro-social impacts of STEM (Science, Technology, Engineering, and Mathematics), including computing (Grover, Pea, and Cooper 2014). Given that music is often both highly social and creative, CS education experiences that draw on computer music could drastically change students' perceptions of computing.

Integrating music into computer science education may likewise be helpful in broadening participation in CS by women, ethnic or cultural minorities, and low socioeconomic people, all of whom are under-represented in CS. Integrating computer music into computing curricula could improve participation both by combatting the misperceptions described above and by creating more pathways into computer science—as well as computer music itself—for students who are already musicians.

In this paper, we describe the creation of a distributed and physical computer music systems-building and performance toolkit for youth aged 10 and up. We have designed this toolkit to enable youth to create new digital musical instruments and other interactive music systems, with the aims of engaging them with a variety of computing concepts as well as challenging them to think more positively about the creative potential of computers and their own ability to do computing. We then describe our recent work with two middle school teachers—one a math teacher, and one a music teacher—to co-design and deploy a curriculum for 11- to 13-year-old students. Our work with these students provides evidence of ways in which computer music can support learning about computer science concepts and change students' perceptions of computing. We conclude by outlining some remaining challenges and questions around how computer music and computer science might best be linked to provide transformative educational experiences.

Design Considerations

We have constructed a toolkit for youth to create computer music systems that they can use in collaborative music performance. Our tool, called BlockyTalky enables users to construct a variety of physical, sensor-rich interfaces, and it supports the combination of different student-made technologies in distributed systems (e.g., sensors communicating wirelessly with software sequencers and synthesizers).

Systems that integrate real-time sensing, programming logic, sound-making, and networked communication are quite common in computer music, yet existing programming tools designed for youth do not target this type of design. For example, Scratch is a popular software system for kids to use to create games, stories, and animations. Scratch enables users to write programs by dragging pieces of code onto sprites in order to define their behaviors. Scratch allows users to play sounds and even to build modest physical interfaces for triggering sound (via MaKey MaKey). However, the sound programming features are more suited to adding sound to games than they are to composing and performing music. EarSketch (Freeman et al. 2014) is another educational programming environment, and one that combines music with CS education. It enables users to programmatically sequence and manipulate synthesized sounds and samples. However, it is not designed for collaborative real-time performances or for using physical inputs to shape sound synthesis.

We wished to replicate certain successful aspects of tools like Scratch and EarSketch; for instance, their browser-based interfaces drastically reduce the complication of system use in schools. However, we also recognized that we would need to borrow heavily from design and engineering patterns that are common in the design of new digital musical instruments such as those featured at the International Conference on New Interfaces for Musical Expression (NIME) or the Guthman Musical Instrument Design Competition. Such instruments can often be understood as assemblages of the following types of components: sound synthesis methods that offer real-time control over their parameters, sensing hardware that obtains real-time information about

performers' physical actions, software that controls the parameters of sound synthesis in response to sensor data and/or algorithmic processes, and networking components that support distributing sensing and sound-making (and possibly also communication and synchronization between multiple performers). As engineers, we wished to create the first toolkit that enabled young learners (aged 10 and over) to create these types of components and connect them to make new systems for live music performance.

Furthermore, as educators, we wanted to empower users of this toolkit to learn about relevant computational concepts. For instance, BlockyTalky requires students to design their own high-level networking protocols—messages passed between different devices for synchronization and control. Students often must discover and adapt to constraints such as latency in sensing, communication, and synthesis. This is essential for students to learn as they design instruments that allow them to improvise during performances using physical inputs or modification of code. Furthermore, building instruments for live performance can lead students to discover how different configurations of sensors support different types of physical interaction.

Implementation

Over the past three years we have iteratively refined the BlockyTalky toolkit in order to meet the above engineering and educational goals. BlockyTalky is open source and runs on low-cost single-board computers like the Raspberry Pi. Typically, we equip these boards with “shields” that allow the use of modular sensors and actuators, including child-friendly LEGO Mindstorms and the slightly more complex Seeed Studio Grove components.

Each BlockyTalky device runs a server that provides users a complete web interface for configuring and programming its hardware. Drop-down menus enable users to declare what kinds of sensors are connected to each input port, and the system provides real-time sensor readings to help users plan, monitor, and troubleshoot their designs.

A variety of programming blocks enable users to define musical motifs (including synthesized notes and samples), to send messages between devices, and to create event-handlers for inputs received from physical sensors or for messages received over the network. The programming model assumes that users will typically compose and enact performances using several physical devices at once, with some devices handling sound synthesis and others handling physical inputs from users; user-created asynchronous messaging protocols coordinate activity across these devices. Figure 1 shows a typical configuration along with code for that configuration.




Figure 1: Typical project architecture and code

Users can configure BlocklyTalky synthesizers to synchronize their clocks to one another by choosing one synthesizer to serve as a reference clock, and programming the other synthesizers to subscribe to that clock (using a `Sync to` block). Then they can use `Wait for` blocks to specify the timing of note synthesis or sample playback. For instance, multiple synthesizers can synchronize their actions by `Waiting for` the same event, such as the next downbeat.

The block-based programming interface is implemented using Google’s Blockly toolkit (Google Developers 2016). Users’ block programs are transpiled into a textual Domain-Specific Language (DSL). The DSL provides convenient abstractions around common complexities for physical computing and network programming. For example, the block program



becomes

```
when_sensor "PORT_1" == 1 do
  send_message("tickle", "elmo")
end
```

The `when_sensor` macro manages hardware state information that is necessary to detect and dispatch events, while the `send_message` function encapsulates peer discovery, serialization, and transmission of messages over UDP or TCP.

This functionality is implemented in JavaScript and the Elixir functional programming language, which runs on the Erlang virtual machine. Erlang's actor-based architecture enables us to quickly add new capabilities to the system, such as support for new hardware, networking protocols, and user interfaces. The Phoenix web framework (which serves both static web content and streaming realtime communication between BlocklyTalky devices and the browser over WebSockets) enables us to provide live information to users about device and sensor states and network communications.

Results with Youth and their Teachers

We have used BlocklyTalky with youth in a variety of different educational settings, including two month-long computer music units in middle school classrooms, two multi-week computer music summer camps (Shapiro et al. 2016), many short 1–2 hour workshops, and in 5- to 10-minute interactions during outreach events. Here, we present results from the two most recent iterations of our approach, which were enacted in two middle school classes in the Rocky Mountain region of the United States. All data presented in this paper were collected under the supervision of our university's Institutional Review Board.

Co-Designing Classroom Implementations

We partnered with two middle school teachers to design four- to six-week computer music units that they would deploy—with our support—in their classrooms. In the United States, middle school students are typically 11–13 years old. Mr. Johnson is a music teacher who has 18 years of teaching experience. Ms. Politi is primarily a math teacher, though she also teaches courses on computing and design, using technologies like Scratch, Arduino, and 3D Printing. She has taught for nine years. We were introduced to Mr. Johnson by another researcher, who had previously partnered with him to integrate electronics design into his music composition class. We met Ms. Politi at an outreach event focused on connecting teachers to CS Education resources.

Our work with students in Ms. Politi's and Mr. Johnson's classrooms began in March 2016. However, our co-design (Roschelle & Penuel 2006) work with Politi and Johnson began months earlier, when we began meeting approximately weekly to play with BlocklyTalky, brainstorm possibilities for classroom learning, draw connections between each teacher's pre-existing goals for student learning and approaches to teaching, to create rough lesson plans for using BlocklyTalky within their classrooms, and even to modify aspects of the BlocklyTalky programming language to better match Mr. Johnson's music teaching methods. Once the classroom implementations began, we continued to meet regularly to discuss how things were proceeding, and to make—or adjust—plans for subsequent days of teaching.

Ms. Politi and Mr. Johnson had distinct but overlapping goals for using BlocklyTalky in their classrooms. Mr. Johnson planned to use BlocklyTalky in his course on music theory and composition, a course in which students already had experience composing and performing music of their own creation using acoustic instruments. He hoped that computer music would offer students a new context to apply and further develop their composition skills, as well as push them to strengthen their teamwork skills. Students in his class had previously used Finale NotePad software to support their music

composition, and had even built their own acoustic musical instruments, but had not previously worked together to compose and enact joint performances.

In contrast, Ms. Politi had never taught, or even played, music before our work together, so some of our co-design sessions involved introducing her to basic musical concepts and vocabulary (e.g. the diatonic scale, how pitch can be described in terms of note or frequency, relationships between melody, harmony, and rhythm). She was eager to participate because she believed that creating interactive musical devices could offer a challenging and interesting design domain for her students, and that it would offer them an opportunity to further develop the programming skills they had developed with Scratch. As such, while Mr. Johnson was primarily excited about how doing computer music with BlockyTalky could enhance students' learning of music, Ms. Politi was primarily motivated by expanding students' "computer power."

Sharing Knowledge: Algorithmic Composition for BlockyTalky

We drew upon Mr. Johnson's knowledge of teaching music theory to address a challenge that had arisen in our prior work teaching non-musician youth to create interactive computer music systems. Many students who participated in our previous workshops and camps have been excited to re-create pop songs that they are familiar with (Shapiro et al. 2016). We have found this to be a double-edged sword: On one hand, it is exciting to many students to be able to recreate music that they already enjoy, to make something that sounds good even without formal knowledge about music composition. On the other hand, we have frequently found this to be a time-consuming dead-end for students, and we have come to call it the "Pop Music Recreation Trap." Students who succumb to this trap tend to become intensely focused on one-to-one replication of pop music, and do not explore re-arrangement or improvisation around melodies and rhythms in their chosen songs. Their work focuses more on fidelity of recreation rather than creative expression. Because re-creations are necessarily linear and

non-interactive, students doing them tend to write programs that consist of one or two very long, non-branching procedures. They tend not to explore programmatic methods for flow control, event handling, and synchronization of musical motifs that would be useful in improvisation, thereby limiting their CS learning.

Because Ms. Politi's students were mostly non-musicians, we were concerned about encountering the same computationally weak Pop Music Recreation Trap as we had encountered with other non-musician participants. Fortunately, in his prior work teaching music theory, Mr. Johnson had developed an algorithmic approach to composition which he calls "Johnson's Recipes." These simple rules govern the construction and integration of melody, harmony, and rhythm, and they can be presented in a way that is accessible to middle school students. The rules are articulated in terms of scale degrees (in the classroom, these were often referred to as "finger numbers") so that they can be applied to composition in any scale (e.g., in C Major, a 2 corresponds to *D*). The rules for melody composition are as follows: 1) Start and end on 1; 2) Use steps frequently; 3) Use skips sparingly. Here are two examples of melodies composed using these rules:



After they compose melodies, students can complement them with harmonies.

Johnson's Recipes initially specified the rules for harmony composition as follows:

1. Use the I, IV, and V chords from the same key the melody is in. The chord is named from the root note of the chord. For instance, I is 1-3-5, IV is 4-6-1, V is 5-7-2.
2. For each note in the melody, choose a chord that contains that note.

- Do not use a IV chord immediately following a V chord. It may be necessary to backtrack, changing the chord on a previous note in order to meet this constraint.

Upon seeing this, computer scientists on our research team were excited to realize that musical rules could be represented as state machines; in particular, Nondeterministic Finite Automata (NFA) offer a simple way to do so. Therefore, we rewrote the rules listed above as:




Figure 2: An NFA diagram to help composers choose harmonizing chords which follow classical Western rules. Each state corresponds to a particular chord.

This representational shift offered a way to depict musical conventions with a representation that is computationally powerful. The NFA diagram in Figure 2 also affords easy representation of the different “rulesets” associated with different musical genres (e.g., adding an additional edge from V to IV in the above state machine transforms it to allow for Blues harmony). Indeed, the power and flexibility of NFAs and other state machines have often been exploited by algorithmic composition systems in computer music community (e.g., Choi and Wang 2010). In an educational context, an NFA diagram offers the further benefit of affording easy checking of student work for conformance with a specific genre. This property of the diagram, of being useful as a tool for being able to generate harmonies as well as to verify their acceptability within a given genre of music, is emblematic of how such diagrams are generally used in CS theory courses and, as such, enables a deeply musical activity (songwriting) to also be a computationally rich one.

Our co-design work with Mr. Johnson and Ms. Politi made extensive use of Johnson’s Recipes. We used them to write songs together, to construct and program

networks of BlockyTalky devices to play those songs, and then designed lesson plans to teach students to do the same.

Unit Organization

Due to their differing goals and the different prior knowledge of their students, Ms. Politi and Mr. Linden opted to work together with us to create two different sets of plans for what to do in their classrooms, including both what the *products* of student work would be and what the *process* of creating those products would be.


Week	Mr. Linden's Class	Ms. Politi's Class
1	Intro to BlockyTalky	Intro to BlockyTalky
2	Small Projects, Show & Tell	"Throw Away Project", Teach Melody/Harmony
3	Plan & Play	Building final projects
4	Build	Final Presentations/Display
5	Finalizing the Projects	
6	Performances	

We also planned to sequence these classroom implementations such that Mr. Linden's class would go first, and Ms. Politi's afterward. This allowed us to manage equipment scarcity (we did not have enough sensors or synthesizers for both classes to work simultaneously) as well as debug any issues that arose in combining BlockyTalky with Linden's Rules in a school-day classroom (all of our previous work with youth had taken place in summer workshops or other informal educational contexts).

Student Work


All consenting students were video- and audio-recorded as they worked together in groups. We also retained time-coded snapshots of all code that students saved on their

BlockyTalky devices. In both Ms. Politi’s 24-student class and Mr. Johnson’s 21-student class, the students showcased their projects at the end of the workshops. Students created a variety of projects, including:




Coldplay Mashup

A mashup of two Coldplay songs, *Clocks* and *Viva La Vida*. The group of girls used buttons to trigger particular parts of the song as they performed.



Catspresso


A 4x5 grid of buttons that could trigger different combinations of melodies and instruments.





Fireball

A carnival style game that utilized a light sensor, a motor, and a button. Players score points by tossing balls into a cup.





Rubik's Cube Competition Accompaniment

3 boys try to solve their Rubik's cubes while music plays. The tempo increases the longer they take to solve their cubes.

Table 1: Project descriptions, network diagrams, and project images. In the network diagrams, orange nodes are devices with sensors, blue nodes are synthesizers, purple edges are student-constructed messages, and green edges are clock subscriptions

These projects involved students designing, building, and programming a variety of hardware and software structures. Table 1 describes a sample of student projects from both classes (9 projects total in Politi's and 5 in Johnson's) and shows their final system architectures and physical construction. The architectures differed strikingly depending upon the kinds of functionality and eventual user interaction the students hoped to support. For example, the *Coldplay Mashup* system supported a group performance in which all of the performers manipulated sensors connected to a single device, *Demeter*, which in turn dispatched messages to the various synthesizers in the system. In contrast, the *Rubik's Cube Competition Accompaniment System* had a strikingly parallel network structure, one in which each Rubik's competitor interacted with a distinct two-node system consisting of a sensor (e.g. *Mystique*) and a synthesizer (e.g. *Transparent*). Prior work suggests that this intra-class variation in solution structure is associated with students' attention to peers' creative ideas (Deitrick, O'Connell, and Shapiro 2014), and greater success at problem solving (Barron, 2003).

Their projects differed in their visual aesthetics (see Table 1). For instance, the *Coldplay Mashup* project was a spartan assembly of computing hardware, unadorned by decoration. This was consistent with the group's plans for them to use it within a performance of their own, in contrast to other groups who created projects for others to use, such as *Catspresso*, and *Fireball*.

These examples illustrate the kinds of projects that students created, as well as the complexity of the distributed systems that students designed and implemented to create them. Our qualitative data, consisting of audio and video recordings of students' working conversations with each other, their teachers, and our research team, offer us insight into the processes through which students created these systems. Specifically, they enable us to observe the processes of composition, design, programming, and problem-solving that produced these projects. We now present and interpret a pair of such conversations. These conversations were typical of students' conversations in both

classes, and as such are representative of the kind of computational problem solving that the BlockyTalky experience engendered.

Example 1: Reasoning About Message Passing While Debugging Silence

The following conversation took place between two students, Carter and Anthony, early in the work of the *Catspresso* group. They are trying to figure out why one of their motifs does not play after they press a button.

Anthony: Does it even play?

Carter: When port 2 is pressed, send message “ostinato”... now that should work, but nothing is playing. Is this speaker on? Yeah. That’s the thing we’ve been struggling with. We’ve never been able to play the ostinato from *Ned* and the melody from *Beeps* at the same time.

Here, Carter is explaining the problem to his groupmate: when they push the button connected to *Walle*, they want to have music play on two different synthesizers—*Ned* and *Beeps*—but it is only playing on *Beeps*. After the boys struggle with this for a while, a research team member, AK, comes over to help. The boys explain the situation to her, and she proceeds to help them debug their code. They start to discuss whether or not the synthesizers are synchronized with one another:

Carter: ...So on *Beeps* we have ‘sync *Beeps* to *Ned*’. And then—we don’t have any syncing stuff on this [*Ned*] because we only need one I think.

AK: Yeah you’re right. Can I see what happens when you press both the buttons? Does anything play at all?

Carter: When we press one the melody does, and when we press two it receives the message but it doesn’t start playing.

Carter’s explanation of the problem includes subtle details that evince a nuanced understanding of event dispatching and inter-device communication in BlockyTalky. Their program on *Walle* defined distinct event handlers for each button, with each one sending a message to a different synthesizer. This arrangement has numerous places for

potential failure, including the electronic connection between the buttons and *Walle*, the configuration of *Walle* (defining the types of inputs that were connected to its various sensor ports), programming of the event handlers on *Walle* (including choosing what events to respond to and what to do in response, which in this case was to send messages to two other devices when a button was pressed), defining event handlers on both *Ned* and *Beeps* that respond to the specific messages that *Walle* is programmed to send, and programming musical motifs on *Ned* and *Beeps*. For the boys' system to function as desired, all of these elements must be done correctly.

When Carter notes that "it receives the message but it doesn't start playing," he is drawing upon several different pieces of knowledge. First, that pressing the button on *Walle* should send a message to *Ned*, that *Ned* should receive that message, and that when *Ned* receives the message, it should make sound. By pointing out that *Ned* does receive the message but does not make sound, he is pointing out that he knows that the message has been sent and has been received, and is locating the problem in the handling of that message-receipt event. He was able to determine this point of breakdown because he watched *Ned's* realtime event log, which displays a stream of information about all network traffic the device receives. Using this log information enabled him to troubleshoot the system relatively precisely, ruling out the *Walle* or the network as sources of error. We believe that this is a remarkably sophisticated act of debugging for a student who only days before had never done any network programming whatsoever, and it illustrates how building computer music systems with BlockyTalky not only offers pathways into computer music, but into learning about interactive and networked systems more broadly. Ultimately, AK helped the students to locate the bug in their system (they had not defined the key of the motif they were trying to play), and the boys created the complex *Catspresso* system depicted above.

Example 2: Synchronization Challenges

In this next conversation, a group of girls is trying to create a mashup of two Coldplay songs. They are experiencing problems getting the melody and the ostinato to synchronize properly; one seems to start before the other.

Ivanna: Why is it not working?

Nevaeh: Do you still have the “wait”?

Ivanna: Nope.

Nevaeh: I think maybe you pressed the button too soon.

Nevaeh asks Ivanna if she has the “wait” in her code, this is in reference to a block of code called `Wait for` that is important to use when synchronizing two or more devices. Synchronizing synthesis in BlockyTalky involves combining three different programming language features: 1) The `Sync to` block, which configures a device to subscribe to another device’s clock; 2) The `Wait for` block, which delays subsequent synthesis until a specified beat arrives (e.g. a downbeat or a 2 in a four-count); 3) The `Set Tempo` block, which specifies the tempo of a synthesizer in beats per minute. Using the `Sync to` block alone is not enough to assure synchronized playback across devices; `Wait Fors` must also be used. Because all BlockyTalky devices default to 120 beats per minute, using a `Set Tempo` block is only necessary for distributed synchronization if something other than the default tempo is desired, in which case all participating synthesizer nodes must use it.

In this instance, the group does not have `Wait For` logic in their code. When Nevaeh learns this, she responds by telling Ivanna that she may have pressed the button too soon. This is a reference to a common strategy that students developed for synchronization without programming: physically pressing multiple buttons (programmed to control different synthesizers) simultaneously. This can lead to approximate synchronization in simple musical systems, particularly those where all synthesizers use default tempos and where no `Wait Fors` (such as to rest to the start of the next measure) are used. However, this approach only works approximately

(groups of novice musicians are rarely good at simultaneous button pressing), and not at all when using motifs that include any scheduling (i.e. `Wait For`).

After this exchange, the girls test their song again. This time the two synthesizers happen (by chance) to start at the same time. The synchronized start enables them to notice that one device seems to be playing at a faster speed than the other.

Ivanna: Why are they going different speeds?

Nevaeh: Check your tempos.

Ivanna: Tempo 120, and...tempo 120.

Jasmine: There was one note that was too fast.

The girls believe that one synthesizer was playing faster than another, and check that both synthesizers are configured to use the same tempo. In actuality, the girls' tempo problem is that they had made an error while programming the motifs on the two devices, with one motif a single beat longer than the other. This is a peculiar instance of the Pop Music Recreation Trap. On one hand, the girls' mashup involved arranging pieces of two different Coldplay songs, and so it was musically rich work, involving modifying the keys and tempos of those different pieces to make them compatible. However, because their re-arrangements also involved fidelity to extended portions of the two original Coldplay songs, the girls' programs also included very long motifs corresponding to those song fragments. These were considerably longer than the motifs that the students tended to compose for themselves. This caused them to end up with code on each device that was just one very long list of notes to play, making it incredibly hard to debug off-by-one errors in note durations. This led to the frustrating situation in which the girls repeatedly attempted to program a distributed BlockyTalky system using the `Sync to`, `Wait For`, and `Set Tempo` blocks, but were unable to achieve the results that they desired. They gave up on programming the system to be as synchronized as they desired. Ultimately, their performance of the Coldplay mashup

involved choreographed button pressing to work around their motif length mismatches, and non-use of programmatic synchronization features.

Evaluation of Student Impacts

In addition to capturing data (audio/video recordings and code) about the processes of students' work, we surveyed students about their enjoyment of their experience and their beliefs about computer science and music. Surveys captured information about whether the students believed that CS and music are enjoyable fields, ones that they can be successful in, and ones that people of various ethnicities and genders can excel in. Surveys also asked students whether they intend to persist in music and/or computing. While most of these survey items were Likert-scale, we also asked free response questions about what students thought the best parts of their experiences were, what they thought needed improvement, and what they believed that they learned.

All of the students in Mr. Johnson's class agreed to participate in our surveys, and submitted signed parent/guardian consent to do so. However, only a minority of Ms. Politi's students returned signed parent/guardian consent forms. Out of concern for how this limited participation rate might bias our sample from Ms. Politi's class, we opted to only analyze the survey data from Mr. Johnson's class.

Survey Results

Students were generally positive about their experiences. On a Likert-scale survey, 18 of 21 students reported that they somewhat or strongly liked their experiences, with *strongly liked* being the most frequent response. Three members of the research team open coded (Strauss & Corbin 1990) the students' free responses to the prompts "The best thing about this workshop was...", "Describe, list or draw three things that you learned at this workshop", and "I would improve this workshop by..." They iteratively converged upon the summary of students' answers that appears in Table 2.

Best Things		Things Learned		Improvements	
Programming	8	Programming	19	More time	5

Creating	8	Tools	7	Instruction	5
Group work	3	Building	6	Programming	5
Interest	3	Computer Music	6	Participation	3
Music	2	Group work	4	Music	2
Problem solving	2	Troubleshooting	3	Perfect	2

Table 2: Summary of students' free response answers, with counts of students (out of 21 students in total) who mentioned each topic in their responses.

Overwhelmingly the students mentioned programming as something that they learned and enjoyed. The students also highlighted enjoying the process of creating and learning about using the tools and building. It is interesting to note that even though this was a music class, for all questions music was mentioned less than other categories. This could be because the novelty was not there, given that music was the primary study in this classroom setting.

Students' pre- and post-participation responses to our attitudinal questions are shown in Figure 3. Though there is considerable support for researchers treating Likert-scale data as interval data (Norman 2010; Sullivan & Artino 2013), there is also some debate about the appropriateness of doing so, with some arguing that they should be treated as ordinal data only (Allen & Seaman 2007; Jamieson 2004). We have chosen to respect both perspectives by graphically presenting the distribution of responses to each survey item (depicting its ordinal character) as well as presenting descriptive and inferential statistics (an interval-oriented approach).




Figure 3: Students’ survey response distributions. Black dots are item means.

Students generally had small positive shifts in their view of computational and musical participation, though in many cases the survey items were limited by strong ceiling effects. In addition to computing means for all items, we took the interval perspective one step further, conducting paired Student’s *t*-tests on all survey items. We

found significant ($p < .05$) positive changes on the items about CS Grades and Good Jobs, the students seeing themselves persisting in music, and the computational excellence of people of all genders and ethnicities queried (except for girls, for which ceiling effects made significant change all but impossible). We take this as a very positive sign that participation in programs like ours can ameliorate negative stereotypes about who can succeed in computer science.

Summary of Outcomes

The outcomes of this work suggest that creating interactive computer music systems can support creative engagement and meaningful learning. Students' projects reveal an openness to combining computing technology with music in innovative ways. They built systems that can be performed as single-player musical instruments as well as systems whose approaches to real-time interaction with music defy easy categorization. Designing, debugging, and refining these systems required students to become proficient with computational concepts beyond those encountered in learning programming, including reasoning about networked communication and synchronization. Students also designed and refined their projects with explicit attention to how the technology would integrate into human contexts of use (e.g., using music to encourage faster Rubik's cube solving, or using people to synchronize to make up for shortcomings of the technology). Furthermore, survey data suggest that students generally felt positive about the experiences of using BlockyTalky, and that computer music may have the potential to improve perceptions of computing and combat stereotypes.

Challenges and Open Questions

The many workshops we have run with youth have revealed several common challenges to supporting meaningful youth engagement with the design of computer

music systems. As we look toward future possibilities for engaging youth in computer music creation, we also see a number of open questions ripe for exploration within the computer music community.

Embracing Pop While Fostering Musical and Computational Richness

Leveraging students' prior knowledge is a central principle of learning environment design (Bransford, Brown, and Cocking 1999). In computer music education, non-musicians' prior knowledge of pop music may be a pathway into deeper musical participation, both by kindling interest in participation and by offering conceptual and cultural resources for students to use within their learning. Further, students who are accomplished musicians may wish to use their knowledge of music theory to rearrange or enhance pop works. Prior work, such as in the EarSketch project, demonstrates the potential value of pop music for computer music education. We should embrace this potential. Yet our work shows how using pop music can also be fraught.

Ultimately, we desire to create computer music learning experiences that are both musically and computationally rich, but the "Pop Music Recreation Trap" described above can trap students in activities that are neither. Focusing on programmatically reproducing existing songs can lead to uninteresting systems, which are not highly interactive and do not encourage computational ingenuity. Further, transcribing existing pop music into programmatic form can be error-prone for both novice and experienced musicians. We were frustrated to see how such errors led the Coldplay group described above to abandon sophisticated distributed programming techniques, which might have in turn facilitated a more musically satisfying performance that was less dependent on human synchronization.

Embedding music theory education into these activities is one promising strategy to combat the Pop Music Trap. In our recent workshops, "Johnson's Recipes" not only provided practical means through which non-musicians could compose songs, but also explicitly encouraged students to understand themselves as musical creators. One

might further imagine using Mr. Johnson's rules and the derived NFAs not just as tools for teaching music theory, but as representations that could be embedded and manipulated within the BlockyTalky programming environment to support higher-level control over algorithmic music generation.

Or, new tools could help students work more effectively and accurately with existing pop music. Programming tools to help students reason about how their composed musical motifs fit together over time could have helped the Coldplay group quickly identify their error and move on to other activities.

More broadly, we believe that additional design research is needed in order to understand how new learning experiences can offer students the benefits of drawing upon their pop music knowledge while avoiding the pitfalls that we have described.

The Challenge of Adapting Existing Tools for Real-Time Music-Making

Building new educational technologies using existing hardware and open source software can help to minimize engineering effort, provide easier pathways for teachers to gain proficiency in the tools, and lower costs for schools. However, support for real-time, expressive interaction with audio has not been a goal of commonly existing platforms. Balancing engineering effort, cost, and suitability for real-time music-making is an ongoing concern for BlockyTalky.

Our current iteration of BlockyTalky uses Sonic Pi for sound sequencing and synthesis. Sonic Pi works well for sound synthesis, even on the relatively low-end hardware of the Raspberry Pi. It also has an existing community of educators and students. However, Sonic Pi is designed for live coding, not for performances that involve real-time control over sequencing and synthesis. It maintains a long buffer in order to avoid skipping, which means that there is often a lengthy delay, up to a second, between when a user manipulates a sensor and when the sound they hear changes. This can be confusing to users, and the Sonic Pi project has no concrete plans

to address this problem. We are currently addressing synthesis latency by replacing Sonic Pi with a new synthesis engine better geared toward real-time performance.

The cost of sensors and the need for associated “shields” or “capas” for connecting sensors to processors are also important considerations in the development of BlockyTalky. The current cost of a single BlockyTalky device with LEGO or Grove-compatible shield is \$100US, making the cost of a set out of reach for most school programs. We are currently collaborating the BeagleBone Foundation to port BlockyTalky to the BeagleBone Green Wireless, a \$40 board that is also compatible microcontroller hardware targeting the computer music community (e.g., McPherson and Zappi 2015). This cost reduction has great potential to drive advances and wider adoption of educational technologies like BlockyTalky.

Should We Move Beyond Programming, and How?

BlockyTalky is robust, cheap, and low-latency enough to support a range of rewarding and engaging music-making activities. The drag-and-drop programming environment is also a usable tool for kids with no programming background, provided that they are open to experimenting with a wide variety of sounds and effects or that they have the musical knowledge to be able to translate their ideas efficiently into symbolic representations (e.g. lists of note names and durations). However, when the educational aims are to teach about musical instrument building, creative expression, design, and collaborative music performance—rather than about procedural programming or musical note reading—might other modes of software design be better suited to these aims? For instance, previous work has shown that building new musical instrument mappings using supervised learning—providing examples of human motions along with the musical outcomes to match those motions—can facilitate a more efficient, satisfying, and embodied approach to design (compared to programming) for professional composers (Fiebrink 2011). Might such techniques also allow kids to translate their musical instrument ideas into real systems? Or, might techniques for

symbolic transcription of sung melodies or automatic harmony generation (Simon, Morris, and Basu 2008) speed up the process of “writing” programs that mimic pop songs? Might helping kids easily realize the creative limitations of mimicry at an earlier stage of their work with technology encourage them to explore new ideas?

What Should Youth Computer Music Education Look Like?

Setting aside the question of how to embed computer music topics into computer science education, what should *computer music education* look like for youth? How could (or should) music education itself change to incorporate computer music ideas and practices? The potential benefits of expanding musical curricula to encompass computer music topics range from increasing the relevance of music education to youth who are most excited about musical genres that rely heavily on digital production practices, to facilitating music-making by youth with disabilities through bespoke digital instruments, to making a politically expedient argument for supporting music education because of its STEM content. But the risks include suggesting that music education is valuable only insofar as it aids in teaching “serious” or “economically important” STEM subjects, or exacerbating disparities between well-off schools with ample resources to invest in digital music equipment and those without them. We are excited about the potential benefits of early computer music education despite these risks, and one of our research aims is to engage the computer music research community more broadly in these questions.

Acknowledgements

We thank the National Science Foundation (CNS-1418463 & CNS-1562040), the NCWIT Academic Alliance Seed Fund, and LEGO Education for funding this work. We also thank our numerous collaborators, including Elise Deitrick, Joe Sanford, Paul Lehrman, Elena Cokova, Catherine Gao, Hilary Dwyer, Zach Lamb, Susana Gomez-Burgos, Lila Finch, and Monica Bolles.

References

- ACM/IEEE-CS Joint Task Force on Computing Curricula. 2013. "Computer Science Curricula 2013." ACM Press and IEEE Computer Society Press.
- Barron, B., 2003. When Smart Groups Fail. *The Journal of the Learning Sciences*, 12(3), pp.307-359.
- Bransford, J. D., A. L. Brown, and R. R. Cocking. 1999. *How People Learn: Brain, Mind, Experience, and School*. National Academy Press.
- Carter, L. 2006. "Why Students with an Apparent Aptitude for Computer Science Don't Choose to Major in Computer Science." *Proceedings of the 37th ACM SIGCSE Technical Symposium on Computer Science Education (SIGCSE '06)*: 27-31.
- Corbin, J. M, and A. Strauss. 1990. "Grounded Theory Research: Procedures, Canons, and Evaluative Criteria." *Qualitative Sociology* 13.1: 3-21.
- Choi, H. and G. Wang. 2010. "LUSH: An Organic Eco+Music System." *Proceedings of the International Conference on New Interfaces for Musical Expression*, pp. 112–116.
- Deitrick, E., O'Connell, B., & Shapiro, R. B. 2014. "The Discourse of Creative Problem Solving in Childhood Engineering Education." *Proceedings of the 2014 Annual Conference of the Learning Sciences*, pp. 591–598.
- Deitrick, E., R. B. Shapiro, M. P. Ahrens, R. Fiebrink, P. D. Lehrman, and S. Farooq. 2015. "Using Distributed Cognition Theory to Analyze Collaborative Computer Science Learning". *Proceedings of the 11th Annual International Conference on International Computing Education Research (ICER '15)*: 51–60.
- Fiebrink, R. 2011. "Real-time Human Interaction with Supervised Learning Algorithms for Music Composition and Performance." PhD dissertation. Princeton University.
- Freeman J., et al. 2014. "Engaging Underrepresented Groups in High School Introductory Computing through Computational Remixing with EarSketch." *Proceedings of the 45th ACM Technical Symposium on Computer Science Education (SIGCSE '14)*: 85-90.

- Ginsburg, H. P. 1997. *Entering the Child's Mind*. Cambridge University Press.
- Google Developers. 2016. Google Blockly. Accessed December 17, 2016.
<https://developers.google.com/blockly/>
- Grover S., Pea R., and S. Cooper. 2014. "Remedying Misperceptions of Computer Science Among Middle School Students." *Proceedings of the 45th ACM Technical Symposium on Computer Science Education (SIGCSE '14)*: 343-348.
- Jamieson, S., 2004. "Likert Scales: How to (Ab)use them." *Medical Education*, 38(12), pp.1217-1218.
- Kelleher C., R. Pausch, and S. Kiesler. 2007. "Storytelling Alice Motivates Middle School Girls to Learn Computer Programming." *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*:1455-1464.
- Losh, S. C. 2010. "Stereotypes about Scientists over Time Among US Adults: 1983 and 2001." *Public Understanding of Science*, 19(3): 372-382.
- McPherson, A., and V. Zappi. 2015. "An Environment for Submillisecond-Latency Audio and Sensor Processing on BeagleBone Black." *Audio Engineering Society Convention 138*.
- Martin, C. D. 2004. "Draw a Computer Scientist." *Working Group Reports from ACM ITiCSE on Innovation and Technology in Computer Science Education (ITiCSE-WGR '04)*: 11-12.
- Mason, C. L., J. B. Kahle, and A. L. Gardner. 1991. "Draw-a-Scientist Test: Future Implications." *School Science and Mathematics* 91, no. 5: 193-198.
- Norman, G. 2010. "Likert Scales, Levels of Measurement and the 'Laws' of Statistics" *Advances in Health Science Education*. 15(625). doi:10.1007/s10459-010-9222-y
- Patterson, D. A. 2006. "Computer Science Education in the 21st Century." *Communications of the ACM* 49, 3 (March 2006): 27-30.

- Roschelle, J. and Penuel, W.R., 2006, June. Co-design of Innovations with Teachers: Definition and Dynamics. In *Proceedings of the 7th International Conference on Learning Sciences* (pp. 606-612). International Society of the Learning Sciences.
- Shapiro, R. B., A. Kelly, M. Ahrens, and R. Fiebrink. 2016. BlockyTalky: A Physical and Distributed Computer Music Toolkit for Kids. In *Proceedings of the International Conference on New Interfaces for Musical Expression*.
- Simon I., D. Morris, and S. Basu. 2008. "MySong: Automatic Accompaniment Generation for Vocal Melodies." *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*, pp. 725-734.
- Strauss, A. and Corbin, J., 1990. "Open Coding." *Basics of Qualitative Research: Grounded Theory procedures and Techniques*, 2(1990), pp.101-121.
- Sullivan, G. M., & Artino, A. R. 2013. Analyzing and Interpreting Data from Likert-Type Scales. *Journal of Graduate Medical Education*, 5(4), 541-542.
<http://doi.org/10.4300/JGME-5-4-18>
- Tubb, R. H. 2016. "Creativity, Exploration, and Control in Musical Parameter Spaces." PhD dissertation. Queen Mary University of London.