

Pianissimo

The Digital Piano Assistant

TEAM 12

Maung Aung (mza), Kevin Dou (khd),
Peter Wei (pwei), & Diana Zhu (dianaz)

18-549: Embedded Systems Engineering
Carnegie Mellon University
Spring 2015
Professor Anthony Rowe

Abstract

Pianissimo is a versatile, inexpensive solution to learning how to play a piano. It is a system that uses audio processing and a laser key highlighting system to teach users how to play notes, chords, and songs. It works with all pianos and keyboards, and is unobtrusive, eliminating the need to modify your instrument or to go buy a keyboard with all the bells and whistles just to have access to a digital tutoring system. Unlike electronic keyboards, which typically come with only a few preprogrammed songs, Pianissimo is capable of tutoring a user through any song because it uses MIDI files as its music model. It is better than watching tutorial videos because Pianissimo offers immediate feedback to the user and because it aids the musician through their own instrument, promoting familiarity and confidence by direct interaction. Pianissimo is capable of detecting both single notes and multiple notes, and will project a laser marker onto the next key(s) to be played. It sits on top of a piano or keyboard via a flexible tripod and includes several calibration options so that users will be able to adjust the system to match their instrument. Once calibrated, users can start to use Pianissimo to learn how to play their piano or keyboard with ease.

Table of Contents

Project Description.....	2
Design Requirements	3
Functional Architecture.....	3
Block Diagram	3
Description	4
Design Tradeoff Studies.....	4
System Description.....	5
Project Management.....	9
Schedule.....	9
Team Member Responsibilities.....	9
Budget	9
Risk Management.....	10
Conclusions	11
Lessons Learned.....	11
To Do Differently Next Time.....	13
Future Work.....	13
Related Work.....	14
Similar Products	14
Our Competitive Edge	15
References	15

Project Description

Learning to play the piano has a very expensive barrier of entry, with the cost of private lessons totaling an upwards of a hundred dollars per hour. As a result, self-taught musicians have been on the rise, supported by the prevalence of video tutorials on YouTube and other online resources. We want to create a product that aids the musician directly—through their instrument—in ways videos cannot. This device will be specifically designed for the piano and will attempt to teach any chosen piece of sheet music to the fledgling musician.

A MIDI file is presented to the device which will turn the file into information it can use. The device is then placed on top of the piano, with no additional hardware needed and out of the way of the musician. The player first calibrates the device, adjusting for offset and key width. Afterwards, the player selects from a list of songs. The device runs through the song with the player by listening to the input of the notes currently being played and highlighting the next sequence of notes to be played via laser projection directly onto the piano keys.

Design Requirements

- **Functionality**
 - must be able to accurately identify the current keys being pressed on the piano
 - must be able to accurately highlight one or more keys on the piano in a well-lit room
 - must be able to highlight keys in a sequence
- **Usability**
 - must not interfere with the actual playing of the song
 - must not interfere with the sound from the piano, nor modify the sound of the piano
 - must not obstruct the view of the keys from the player
- **Timing**
 - must be able to detect the current notes being played within 350 milliseconds
 - the device should project the next sequence of keys within 250 milliseconds

Functional Architecture

Block Diagram

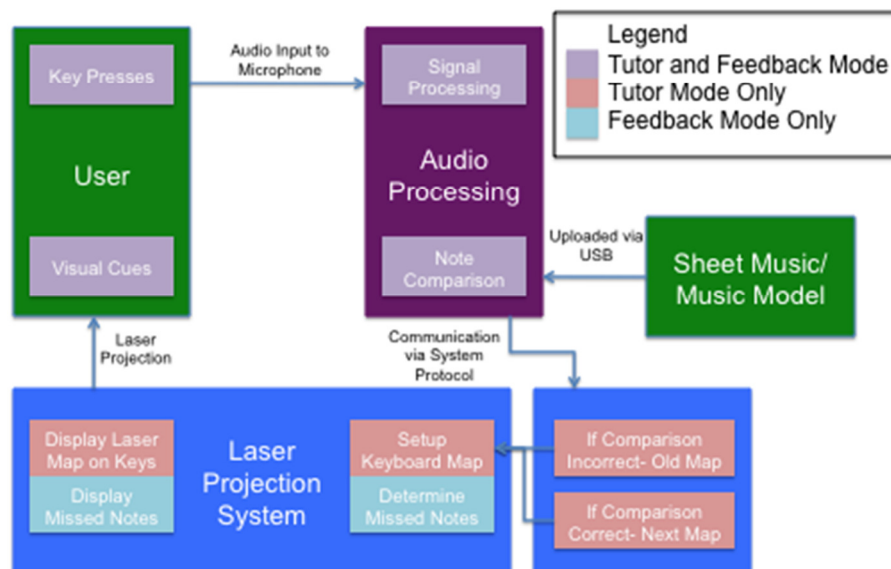


Figure 1 – The Block Diagram

Description

As illustrated in Figure 1, above, our project is composed of two main systems that communicate with the user as well as with some predefined MIDI file that the user wishes to learn or perform. The two main systems are the audio processor (acting as the main sensor module) and the laser projection system (acting as the main actuator module). The systems will act together to tutor a user in learning a piece of music through two stages: tutor mode and feedback mode. The two modes are separate but require the same responsibilities from the audio processor; however, they require different responsibilities from the laser projection system.

The audio processor is responsible for interpreting the key presses made by the user. Once the signal is received, the signal is processed and then compared with the user's predefined sheet music in order to determine the correctness of the user's key presses. In tutor mode as well as in feedback mode, the audio processor's responsibilities remain the same.

The laser projection system is responsible for displaying to the user mode specific output. The system will receive whether or not the user has pressed the correct keys. If so, the system should generate the key map for the next set of keys and let the laser projection system display that on the piano keys. If not, then the laser projection system will display the same map but highlight the incorrect keys. The overall system will continue to output the same key map until the user has correctly presses all the keys.

Design Tradeoff Studies

Note Recognition

The ability for our device to recognize the notes that users are playing is central to its role as a piano tutor. We have identified two methods in which we may be able to detect notes: by listening to audio through a microphone and identifying the different frequencies of notes or by using a thin bar of switches placed near the top of the keys that will trigger whenever a particular key is pressed. The microphone approach would allow for much cleaner packaging and, subsequently, be much more versatile with different pianos. If our device can recognize notes through a microphone, then our final packaging can be a singular, compact box that sits on top of the piano and coaches users through songs. The bar of switches approach, on the other hand, is clunkier. It will be difficult to design to fit pianos of many different sizes and will be bothersome to carry around. However, the advantages that the bar of switches has over the microphone approach are speed and accuracy. Instead of taking the processing time to guess which key was pressed from the microphone input, the bar of switches can know the key exactly and immediately. The bar of switches may also be able to house LEDs that could provide feedback and coaching to the user. Considering these tradeoffs, we have chosen to go with the microphone approach. Not only is the microphone approach cooler and more polished, but we believe we may be able

to mitigate its disadvantages with well-executed code performing the fast Fourier transform. The bar of switches will be our fallback.

Key Lighting

In order to provide feedback and coaching to our users, we will light up piano keys. To do this, we have considered four options: use a normal projector, use a spinning high-velocity laser such as the ones seen in alarm clocks that are able to project the current time onto the ceiling or wall, use a laser project, or simply use LEDs for each individual key. The LEDs are the simplest and most straightforward approach but they lack a certain coolness and will be cumbersome to place onto the keyboard (unless we are forced to go with the bar of switches approach for note recognition). The use of a normal projector was our original thought, but we grew concerned about its size, its cost, and its visibility and resolution in a well-lit room. Focusing a projector would also be very sensitive to placement in terms of its distance from the keys. This is why we began considering the third option as our main choice. With a spinning high-velocity laser, we have the coolness factor that the LEDs lack, we are less encumbered than with a normal projector and we can worry less about brightness and focus in well-lit environments. The spinning high-velocity laser is, however, much more technically challenging than our other two options, so we will keep the possibility of using a laser projector in mind as our backup plan. The laser projector gives us the ability to project numbers onto keys without having to worry about focus and brightness like with a normal projector and it has fewer moving parts than a spinning laser system, but it isn't quite as cool.

System Description

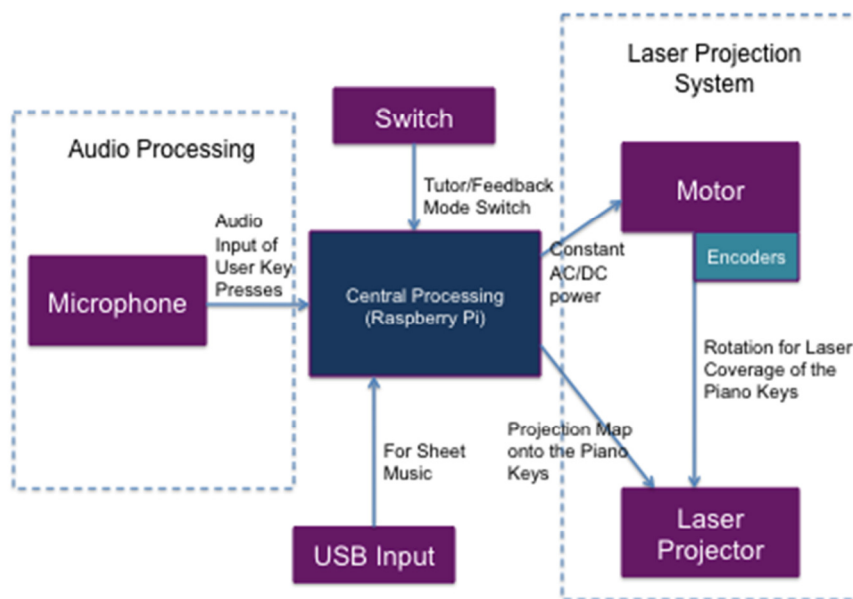


Figure 2 – The Physical System

Description

As seen in Figure 2, above, the physical architecture can be split into the audio processing components, the laser projection system components, and a few additional components that are manipulated by the user before using the system.

The switches are used by the players to provide input for the device. The 6 buttons are power, restart, back, left, right, and select. The button presses are recognized by the system and used to advance to different states.

The audio processor, depicted below, will utilize a microphone to detect the audio signal coming from the piano. The microphone will send the audio data directly to the central processor for processing.

The laser projection system, as shown in Figure 4, will utilize a motor with its shaft spinning at a constant rotational speed; connected to the motor shaft will be an array of lasers to project onto the piano keys. The lasers will each receive a projection map of angles. Each set of angles will denote a range between which the laser will shine on a specific key. If the laser is on during that range of angles, then the specific key will be illuminated by the laser. The motor encoders will provide a reference for which angle the lasers are pointing at during any given time.

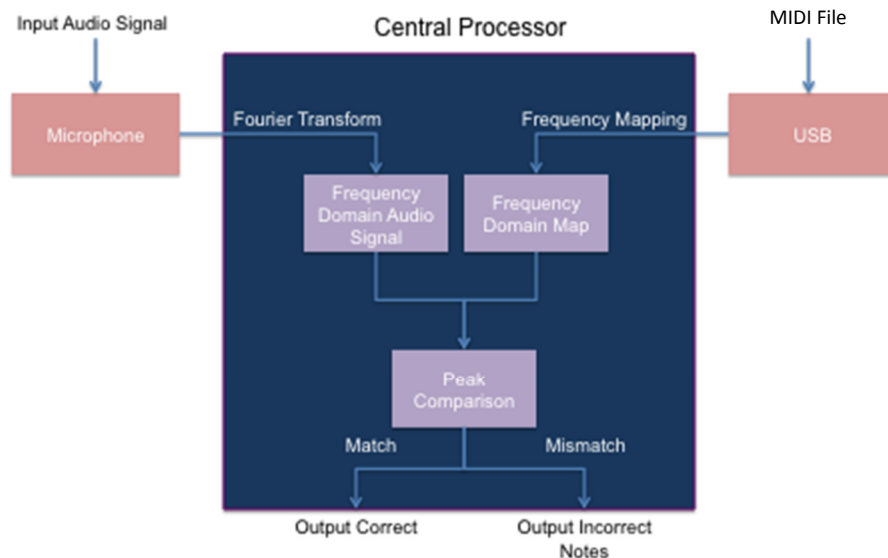


Figure 3 – The Audio Processing System

The input audio signal will be transformed into a frequency domain audio signal, which will then be compared to the frequency map of the sheet music. The sheet music notes will be converted into their frequency domain counterpart for comparison. The peaks will be compared, and if the frequency maps match, then the system will output that the

keys pressed by the user are correct. If the peaks do not match, the output will show the incorrect notes.

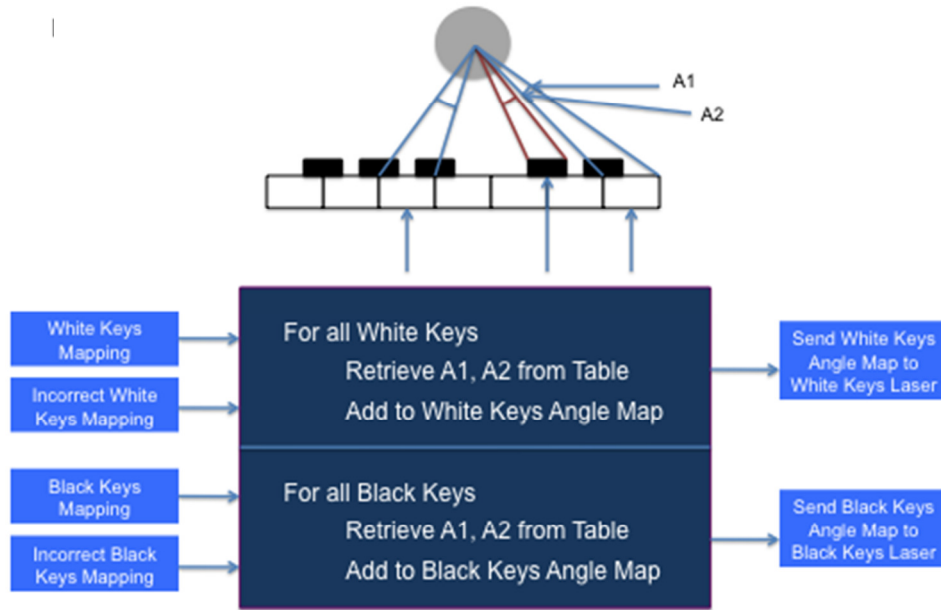


Figure 4 – The Laser Projection System

The incorrect notes will be input into the intersystem protocol. If there are no incorrect notes, then the sheet music will advance a time step and a new map of the notes will be generated. This will then be split into white key maps and black key maps for the different lasers. If there are incorrect notes, then the sheet music will not advance in time and the incorrect notes map will be added onto the original note map. This will generate 4 maps: a correct white key map, a correct black key map, and incorrect white key map, and an incorrect black key map.

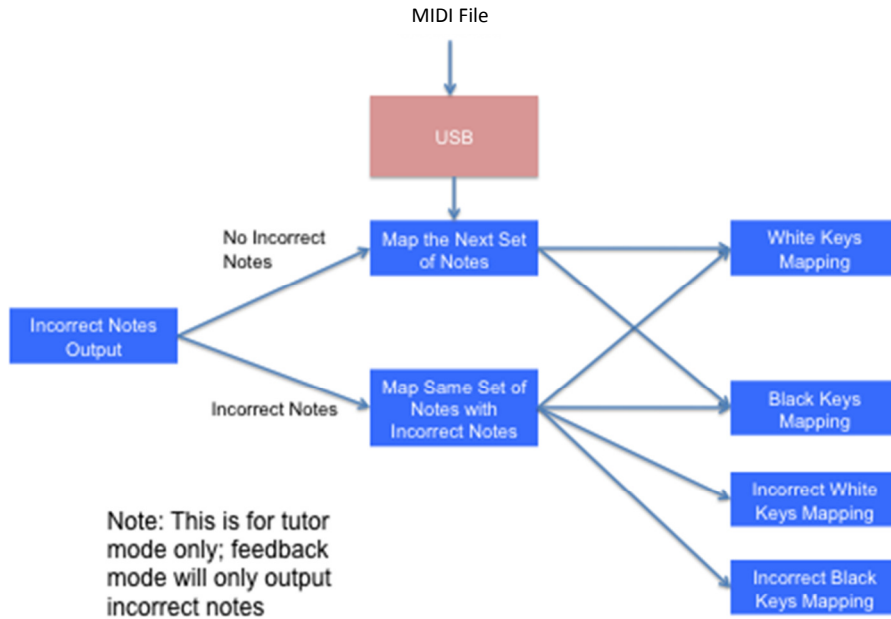


Figure 5 – The Intersystem Protocol

All key maps will pass through one more layer of abstraction. Each key has a set angle range which if the laser is on during that range, will illuminate the specific key. The angles will be retrieved from a table and an angle map will be generated and sent to each laser. The lasers will then display the correct and incorrect key presses on the actual piano keys.

Project Management

Schedule

Start Date	End Date	Goals
2/16	2/18	Proposal
2/19	2/20	Set up website for website check 1
2/21	3/16	Create a system for sound input with a mic Project onto keys – Spinning laser beam with a 3D printed fan mount for a mirror, control of laser beam spacing
3/17	3/23	Note recognition with audio Be able to modify width of projection with a lens
3/24	3/30	Sync audio with projection
3/31	4/6	Create/implement teaching mechanism
4/7	4/13	Read in sheet music
4/14	4/20	Improvements based on user testing
4/21	4/22	Final Demo - Fine tune based on feedback from demo 6

Team Member Responsibilities

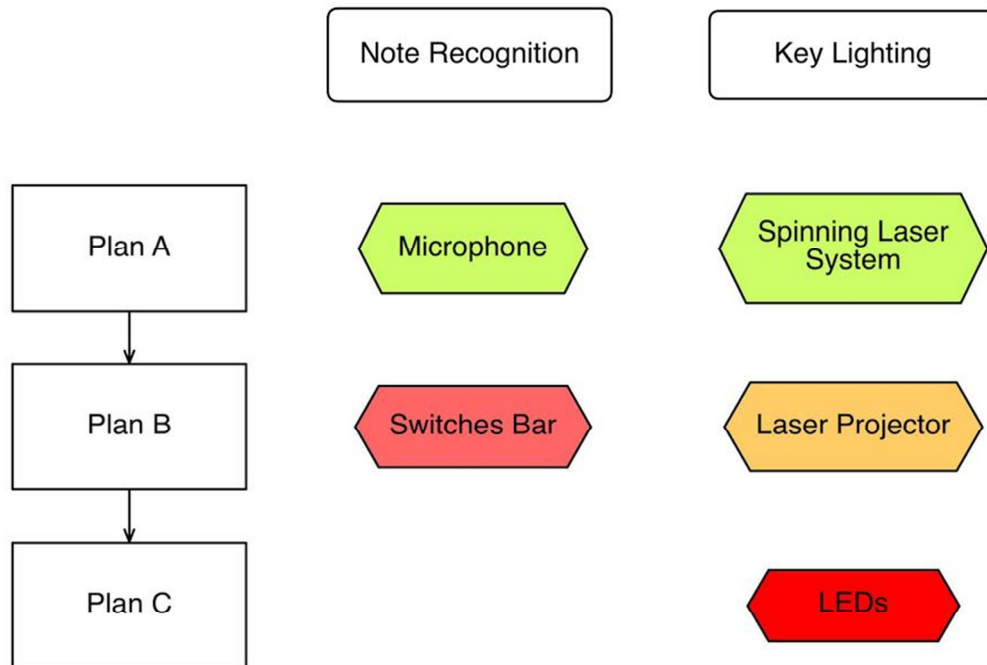
Major Functions	Responsibility	Team Member
Key Lighting	Motor Component	Maung Aung
	Laser Components	Maung Aung Diana Zhu
Note Recognition	Microphone Setup	Peter Wei
	Signal Processing	Peter Wei
Central Processing	Embedded Processor Setup	Maung Aung
		Diana Zhu
		Peter Wei
Intersystem Protocol	Protocol Management	Peter Wei
		Kevin Dou
Website	Webmaster	Diana Zhu

Budget ^[1]

Part Description	Cost Estimate
Lasers and Motor Components	\$74
PCB & Components	\$340
Microphone	\$72
Raspberry Pi	\$46
Casing & Stand	\$103
Power & Batteries	\$58

Risk Management

In order to manage and mitigate risk, we have backup plans for both of the two main components of our device. The potential problems we foresee with the microphone approach are the speed of the fast Fourier transform and its accuracy at determining notes if multiple keys are pressed at the same time. If the audio processing with a microphone doesn't work, then we will fall back to using a bar of switches to determine which keys are being pressed. The use of a spinning laser system to provide feedback and coaching is technically challenging to build from scratch, so we plan on building a quick prototype early on to test the concept, and fall back on using a laser projector in case our prototype fails or comes up short. We will design these two main components—note recognition and key lighting—to be modular and to have a common interface so that either component will be exchangeable in the case of needing to enact one of our backup plans.



Conclusions

Lessons Learned

Miscellaneous


There are several miscellaneous lessons we learned throughout this project. Some of them include: 1.) measure twice, cut one in reference to laser cutting the enclosure, 2.) even if subsystems are working properly, it takes some time to put it all together, 3.) true real-time operating systems are necessary when the margin of error for timing is very sensitive, 4.) hardware issues are just as important as software issues, and they're sometimes harder to correct for, and 5.) it's very difficult to get a powerful laser from a reputable source.

Jitter Problem

One big problem we encountered was an inconsistent fan rotation speed. Based on the output of the fan's tachometer and from conversions to a numerical RPM, it was clear that the fan's speed was inconsistent and often spiked much higher or much lower than the average speed. As a result, the laser projection onto the keys would jitter, sometimes causing misses. Here are the calculations for our estimates of the margin of error from fan jitter, and proof that we have accounted for jitter through software:

Let h be the shortest distance from the laser point on the mirror to the piano keys. Let μ be the average number of microseconds it takes for the fan motor to make exactly one full revolution, and let σ be the magnitude of the deviation from μ expressed in microseconds. Then the deviation in radians σ_θ is given by $2\pi(\sigma/\mu)$.

At the least affected angle—straight down—the jitter ϵ_{\min} as a result of σ_θ is given by the formula $\epsilon_{\min} = 2h(\tan\sigma_\theta) = 2h(\tan(2\pi(\sigma/\mu)))$. At the most affected angle, assuming that the total angle for all the keys upon which we want to project the laser less two keys (for easier calculations) is $\pi/2$, the jitter ϵ_{\max} as a result of σ_θ is $\epsilon_{\max} = h(\tan(\pi/4 + \sigma_\theta) - \tan(\pi/4)) = h(\tan(\pi/4 + 2\pi(\sigma/\mu)) - \tan(\pi/4))$.



In our tests, we used $h \approx 25$ cm. We observed that, on average, $\mu \approx 47,500$ microseconds and $\sigma \approx 500$ microseconds for a fresh 9V battery.

So, we have that $\sigma_\theta = 2\pi/95$. Then, we have the following:

$$\epsilon_{\min} = 3.312 \text{ cm}$$

$$\epsilon_{\max} = 3.547 \text{ cm}$$

In the final version of our code, you will observe that our jitter is not as extreme as the possible values above. This is because we correct for jitter in our code, by using the actual time delta of the previous fan motor revolution to scale our drawing offsets and variables before drawing.

Octave Error, Repetition Error, and Multi-Note Error

The audio processing system in our project ran into three major obstacles. The first of these obstacles, octave error, is a tricky hurdle to deal with due to the nature of the Fast Fourier Transform (FFT). The idea is that when a frequency such as 440 Hz (A4) is played, the FFT will also see peaks in amplitude at the halves and the doubles, 880 Hz (A5) and 220 Hz (A3). Because of the imperfections of audio receiving and external noise, the Fourier Transform may actually detect a note that is one octave higher or lower than what is actually played.

To combat this problem, we implemented an additional check for whether or not a note was played correctly. If the reference note is not detected as being played, we search the octave above and below; if the amplitude is close to the reference note amplitude, we make the assumption that octave error occurred and we consider that note as being played. While this may cause problems if the player is playing one octave higher or lower, we decided that with the laser guidance, correct usage should not cause a player to miss the correct notes by an octave.

The second problem is repetition error. The difficulty with thresholding the amplitude to detect a note has two problems: firstly, louder or softer notes will behave differently to an absolute threshold, and secondly, repeated notes may be falsely detected. The reason for this is that if the first instance of a note has enough amplitude, the next sample may also have amplitude above the threshold and thus show a double note when in fact only one was played.

In order to solve this problem, we took a time derivative to detect a played note. If the note amplitude is significantly higher than the previous sample's amplitude, then we assume a note has been played again. This way, a note can have a peaking amplitude multiple times in a row and be detected as repeated notes, whereas a note that has a high amplitude initially and decays over time will be detected as a single note, regardless of how high the initial amplitude is. We give a lower bound threshold for a played note, so that silence will not be detected as notes being played.

The third problem is multi-note error. Without an absolute threshold, there is no way to know exactly how many notes were played. To combat this, we took the number of notes (call this number n) that *should* be played (from the sheet music) and take the highest n peaks. If these notes align correctly with the sheet music notes, then we have successfully played the right notes. While there is potential for error, this system actually performed quite accurately in quiet environments and even in environments with ambient noise.

To Do Differently Next Time

Real-Time Operating System

Raspbian is not a pure Real-Time Operating System; thus, the software was not quite able to perform at the performance required. The most time critical section of our project is the laser drawing; the software should break out of its routine and start drawing at exactly the same place else the laser will have jitter. In our system, the software is interrupted at irregular times, which we suspect is a major cause of our inability to fix the jitter through software. The audio processing can be run in the background, as it is not so time-critical as the laser drawing.

Photodiode vs. Photoresistor

In our current design of Pianissimo, we use a photoresistor to determine the location of the laser as rotates and to indicate when to start drawing the key map. We have an ADC attached to the photoresistor and, together, they act as an interrupt signal to our microcontroller. Photoresistors, however, suffer from latency between their exposure to light and their decrease in resistance, “usually around 10 milliseconds,” [2] which we believe may contribute to our jitter problems. To improve upon this, we would use a photodiode instead since it has a much better response time. [3]

Future Work

Black Keys

An important feature that is missing from our project is the ability to draw on the black keys (any song that requires flats or sharps would require our system to be able to accommodate black keys). In order to draw on the black keys, we would require a stronger laser that can clearly show up on the black keys. One engineering challenge that we would inevitably run into is the elevation of the black keys; on the sides of the piano, the laser may not be able to reach the white keys, as the raised black keys would block them above a certain angle. Thus, it would be important to have two lasers, one for the white keys and one for the black keys in order to maintain the ability to draw on all the keys.

Better Calibration Method


We used a linear model to calibrate the angles for each key with respect to the laser. However, our model was only able to accommodate a little more than an octave around the center of the laser. In the future, we would develop a more accurate model to map to the keys accurately.

Faster Audio Processing

The audio processing can return the notes played every 0.3s to 0.35s, which is fast enough for some songs and amateur players, but would have difficulty keeping up with faster songs. The limiting factor is the Fast Fourier Transform; not only does the accuracy of the note detection decrease with faster times, the FFT calculation does not go significantly faster. The sampling is meant to be every 0.25s, meaning that the FFT calculation is in fact slowing down the feedback by up to 40%. With a faster processor, one may be able to achieve less computation time, and may be able to reduce the note capture interval to 0.2s, which is significantly faster and would be fast enough for many more songs.

Related Work - Similar Products

Projected Instrument Augmentation system (PIANO) [4]



The PIANO system, developed at the University of Ulm in Germany, consists of a large screen attached to an electric piano. The screen displays upcoming notes for users to play a la Rock Band or Guitar Hero, with colored bars for different fingers that are styled for different "ornamentations" such as trills. If a user plays an incorrect note during a song, the PIANO system lights up the offending key(s) in red. While this concept is very similar to the ultimate goals of our project, it is different in that the PIANO system does not work with different pianos. Instead, the screen is specifically paired with a particular electric piano, so users cannot purchase the PIANO system to teach them how to play on a piano that they may already have in their possession. Our product, on the other hand, is designed to fill this niche.

Yamaha EZ-200 Lighted Keyboard ^[5]

The Yamaha EZ-200 Lighted Keyboard is a portable, touch-sensitive keyboard with 61 lighted keys to teach users how to play the piano. It comes with 100 built-in songs and is representative of many of the lighted keyboards available in the market today. These keyboards, although capable of teaching users how to play the piano in a fashion similar to our own product, are limited in three ways that our product is not. First and foremost, lighted keyboards are almost always electric but many piano enthusiasts prefer the sound and feel of a real, acoustic piano ^[6]. Secondly, the number of songs that such an electric keyboard is capable of teaching is always limited to a fixed number. And, third, the actual teaching component of lighted keyboards is not capable of working with any other piano or keyboard; users that have already spent a fortune buying an expensive piano will need to make another investment buying something that they already have.

Our Competitive Edge

The main competitive edge that our device has over our competition is that it will be compatible with most, if not all, existing pianos. There is no need to modify a piano that users have already purchased and the impact and presence of our product will be minimal. This allows our users to pick and choose to their preference. Instead of having to invest more money in a lighted keyboard that they don't need, they can get started learning right away. In addition, our product will also be portable and versatile. It will be capable of reading in any sheet music and will be able to teach users how to play a practically unlimited number of songs.

References

- [1] “[18-549] Parts Ordering.” Google Drive. Web. 11 May 2015.
<https://docs.google.com/spreadsheets/d/18p_4EFcZiUi-l_oZQEwQci7KDGW0IT5SFyvj9Y4RDI/edit?usp=sharing>.
- [2] “Photo resistor.” *Resistor Guide*. Resistor Guide. Web. 11 May 2015.
<<http://www.resistorguide.com/photoresistor/>>
- [3] “Photodiode.” *Wikipedia*. Wikimedia Foundation Inc, 30 Apr. 2015. Web. 11 May 2015.
<<http://en.wikipedia.org/wiki/Photodiode>>
- [4] Marks, Paul. "Real Rock Band: Play Piano like a Pro with Light Keys." *New Scientist Tech*. Reed Business Information Ltd, 2 Oct. 2013. Web. 16 Feb. 2015.
<http://www.newscientist.com/article/dn24296#.VNz_V_nF9zs>.
- [5] "EZ-200 Lighted Keyboard." *Yamaha United States*. Yamaha. Web. 16 Feb. 2015.
<http://usa.yamaha.com/products/musical-instruments/entertainment/lighted_key_fret_instruments/ez_series/ez-200/>.

[6] "Digital vs. Acoustic Pianos." *Piano Teachers Federation*. Piano Teachers Federation. Web. 16 Feb. 2015. <http://pianoteachersfederation.org/Digital_vs_Acoustic>.