Harmony Innovations Inc.

providing technical solutions that harmoniously integrate music and technology

[POST MORTEM]
Maestro™
December 16, 2010

December 16, 2010

Mr. Mike Sjoerdsma
School of Engineering Science
Simon Fraser University
8888 University Drive, Burnaby, BC
V5A 16S

Re: Post Mortem of a Sheet Music Scanner Device (The Maestro™)

Dear Mr. Sjoerdsma,

Please find enclosed Harmony Innovations Inc.'s post mortem for a portable sheet music scanner. This device will be a prototype for a range of new handheld music recognition technologies. The goal of this project is to build a handheld device that converts sheet music into sound.

The purpose of this post mortem is to unambiguously describe the technical aspects of the Maestro™ in its proof of concept phase only, with some aspects of the production models being discussed.

This document outlines the current state of the proof-of-concept device, future plans for the device, deviation from originally planned design specifications, test results, and deviations from the budget and schedule. An individual review of the development process by each team member is also included.

Please feel free to contact me personally with any questions you may have at 604 649-3346, or by email at sle@sfu.ca.

Sincerely,

Sean Edmond

Sean Edmond
Chief Executive Officer
Harmony Innovations Inc.

Enclosed: Post Mortem for "The Maestro™" music education aid

2010

Post Mortem

The Maestro^{TM}

Portable Sheet Music Scanner and Player

MISSION STATEMENT

At Harmony Innovations, the future is our passion.

By combining musical education with the technological advances of today, Harmony Innovations strives not only to enhance the quality and accessibility of musical education for all students; but also to provide support and technology as a partner to many up and coming artists.

Technology has enormous potential to enhance our lives and this is the guiding principle behind Harmony's comprehensive approach to musical education.

Date Submitted: December 16, 2010

Submitted to:
Dr. Andrew Rawicz (ENSC 440)
Michael Sjoerdsma (ENSC 305)

Project Team:
Sean Edmond (301026670)
Nikola Cucuk (301033241)
Veronica Cojocaru (301055896)
Cris Panaitiu (301032665)

Table of Contents

List of Figures

List of Tables

Glossary

CPU	Central Processing Unit
IC	Integrated Circuit
I/O	Input/Output
IQM	Image Acquisition Module
fps	frames per second
LCD	Liquid Crystal Display
LED	Light-Emitting Diode
MIDI	Musical Instrument Digital Interface. A digital music storage standard, in use across all of the North American music industry
OCR	Optical Character Recognition
PCB	Printed Circuit Board
SMF	Standard MIDI File
SRAM	Static Random Access Memory
SPI	Serial Peripheral Interface
USART	Universal Synchronous Asynchronous Receiver-Transmitte

1. Introduction

Our team first met in August of 2010. Throughout the month, we met several times to discuss potential projects before deciding on “Maestro™”, a sheet music scanner and player. Since September, we have been working to develop a prototype, which we demonstrated formally on December 14th. This document summarizes our learning experiences, and outlines our project’s evolution over the last few months.

“Music education is the epitaph of human achievement and embodies the desire to understand that which we do not understand through that which we do” – Socrates

The effort to express through music is by no means exhausted today, especially with the advent of new technologies through which music can be more easily created. Musical education, however, has not benefited sufficiently from the progress made in digital music technologies: most learning is still done through a process of demonstration by a teacher and then practice by the student.

Maestro™ will consist of a simple user interface that will allow the user to configure the device. A camera will allow the user to scan and buffer the sheet music that they want to playback. Once they have finished scanning, they can listen to the music, which will be outputted to a headphone jack.

2. Current State of the Project

Over the past 4 months, the Maestro™ team has done a phenomenal job at bringing this project together. This section describes the progress made towards accomplishing the goal and designing the final device.

2.1 Overview

The system design of the Maestro™ was accomplished as described in the design specification: the EVK1100 development board provided the processor, display, inputs and connections to interface with the camera and sound modules. The sound module communicated to the processor using the SPI interface and outputs sound via a headphone jack as seen in Figure 1 below.

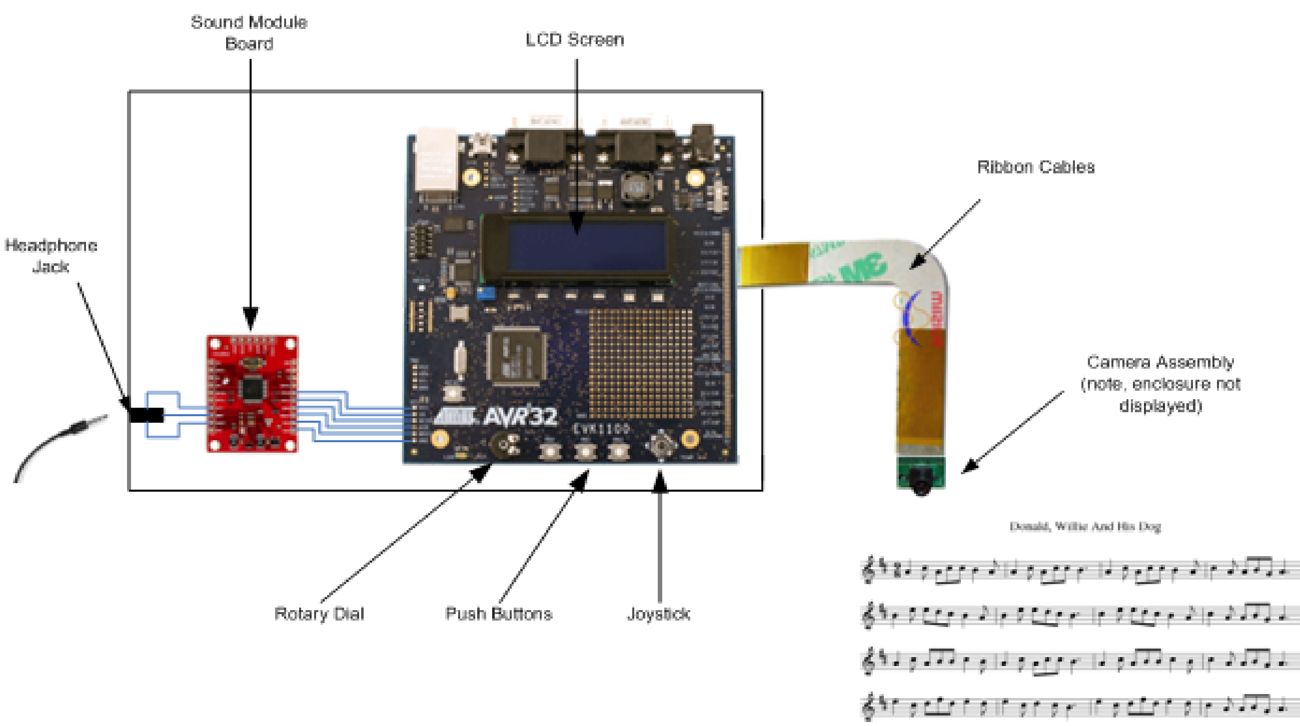

Figure 1: Maestro™ Mock-up

System integration went as planned and the Maestro™ is capable of operating as a standalone device. The LCD screen displays the menu, the options and the current state of the device. The user can interact with the menu by using the joystick for navigation and pushbuttons for selection.

If scan mode is selected, the user can use the camera assembly to scan in sheet music (line by line) by sliding the camera enclosure along the line of music to be scanned. A pushbutton is used to start and stop the scan and the processing happens after a line is scanned in. The corresponding notes are written into a MIDI file in the SDRAM. Users can then scan in more lines of music or listen to the MIDI already stored in the device. Scanning in additional sheet music will append the new music to the already existing MIDI file. The user can also erase the MIDI stored onto the device and start scanning fresh music in.

2.2 Hardware and System Integration

The completion of the camera module portion of the project had significant hick-ups. The final camera implementation used an RS232 interface, which operates much slower than the parallel interface of the camera initially purchased. The speed of the camera interface severely limited our polling frequency and also required us to use lowest resolution image format 160x120 (to limit the amount of data transferred). The polling frequency limited the speed of our scan (about 10 seconds for an entire bar of music). The advantage of this camera was that it came packaged with a computer user interface to allow for capturing images with a GUI on a personal computer. This allowed us to obtain test images for development of the algorithm before the embedded interface was set-up.

The sound module implementation was relatively seamless. However, when integrating the sound module, our camera stopped working. This integration bug was the result of not considering the overall clocking scheme of the project.

The user interface and top level integration required getting the LCD, push buttons and SDRAM of the board integrated. This task also entailed setting up a debug interface so that captured images from the microcontroller could be displayed on a computer when stepping through the microcontroller code. Integrating the camera, sound module and algorithm into the top-level code had its own unique challenges. The final version of the project still suffered from some shortfalls due to errors in our tools. We were not able to download the release version of our software, meaning that the code executing on our prototype is the debug version (which runs significantly slower than the release version).

Due to various reasons, it took longer than expected to get the hardware working and integrated, thus slowing down the development of the software. If we were to undertake a similar project in the future, we would make sure to decide on and order parts as soon as possible to give us the best chance of getting the hardware done in time to perfect the algorithms and the software.

2.3 Optical Character Recognition (OCR) Software

The optical recognition software to convert the scanned images into music was prototyped in Matlab. Development on this portion was done initially with screen shots of sheet music on a computer. Using these screen shots, the algorithm performed with 100% accuracy. However, the actual images from the camera were of much poorer quality and suffered from spherical aberration from the camera lens. As a result, many portions of the algorithm had to be enhanced or re-written. Eventually, the algorithm was ported to C so it could run on the microcontroller. In this process of porting many bugs were introduced. We were also face with the challenge to maintain 2 versions of code. When stepping through our C implementation of the algorithm on the embedded platform, we found the algorithm had about 60-80% accuracy in detecting:

- quarter notes
- half notes
- eighth notes
- eighth beams
- whole notes
- eighth rest
- quarter rests
- bar lines

- pitch

We were unable to start debugging the algorithm with fully integrated code until December 11th. If given more time, this algorithm could have been perfected. Sadly, music sounds almost unrecognizable unless it`s exactly the same as the actual song. Another issue with running the algorithm on the microcontroller is that it ran significantly slower, and space was limited. Stack overflows resulted in our code getting "hung up", requiring us to write our code more efficiently. This current algorithm relies on the images being captured without rotational or skew error.

2.4 Deviations from the Original Design

The main deviation from the specifications that affected the performance of the device was that we were unable to fully complete and debug the OCR software to the requirements of the prototype model as outlined in the functional spec. In that document we had outlined the various requirements for the proof-of-concept model, of which R89, R94, R95 and R97 (which have to do with the symbols that will be supported and recognizes) were not achieved as described in the previous section. As well, good accuracy in the recognition of the various notes has not yet been achieved, and further work needs to be done to develop the algorithm for it to work with music straight out of a music book, and not specifically printed out for the purpose of this project.

The user interface also differs slightly from what was described in the specifications. The main menu that was implemented was as follows:

```
Main Menu:
Play
Scan
Set Tempo
Set Time Signature
Set Key Signature
Erase MIDI file
```

Currently, the "Set Time Signature" option does nothing as the time signature error checking has not yet been implemented in the algorithm. "Play", "Scan", "Set Tempo" and "Set Key Signature" work as described in the design spec. "Erase MIDI File" has been added because we wanted to allow the user to continue scanning in music and appending to the currently stored MIDI file. This would be the default operation, thus users needed to have a way to clear the MIDI file from the device if they wanted to start fresh.

Initially, we wanted to implement volume control using the rotary dial, but due to time constraints, we had to prioritize and it was decided that this functionality was not necessary for the proof-of-concept model.

The user interface also greatly suffers from the fact that we were unable to use our original choice of camera: the camera that we ended up using to acquire the images for processing is much slower and the user is forced to slide it slowly in order to get enough images for the OCR to work properly. Also, the development board we used was rather big and we weren't able to implement a hand-held or battery powered device. This can improved in further development of the system.

Most of these deviations were due to time constraints and not to any flaws in the design of the device itself, thus we are confident that with more time we can get the device fully functional.

3. Future Plans

We feel our proof of concept has shown potential for this product to be a viable business venture. However, turning this into a marketable product would require improvements in several aspects of the project.

Firstly, we would need to design a better mechanism for capturing the sheet music. We would need to use a parallel interface, requiring the use of dedicated hardware for handling interrupts instead of firmware as in the current revision. Ideally, this parallel interface would have its own dedicated memory controller to interface with shared RAM. Considering that the bulk of our processing time is consumed in the mosaicing algorithm, we would benefit from to use of a linear scanner. Using a linear scanner would also eliminate the spherical aberration issue from using a camera with a lens.

Most of future work would need to focus on refining the music OCR algorithm. In addition to improving its current capabilities, this software would also need to consider the "peripheral features" in music (such as slurs, accidentals, bowing marks, etc). In addition, the current algorithm does not make use of the timing information of the music. This timing information could provide valuable in a more statistical determination of the "best match" of features between bar measures.

Histogram matching to isolate and identify the features is not be the best way to proceed since there are many ambiguous cases that can impede correct detection. This is especially relevant when additional features of sheet music need to be considered. As such image processing on a 2-dimensional image will need to be designed and integrated into the system. In order to speed up processing time, it is worth considering if custom hardware should be designed to perform some of the necessary processing.

The vs1053b is adequate in meeting our needs. Future revisions could make use of this chip's ability to synthesize different instruments. Future revisions would also have to implement a volume control with this chip in case individuals are using headphones rather than speakers with a separate volume control.

Maestro™ would also have to be fabricated as a hand-held device. Since a considerable amount of processing capability is required for the music OCR, this device would need to use a high speed processor. All components would need to be customized and placed on a PCB. An efficient lighting scheme for the scan area would also have to be implemented to save on power requirements.

The interface can be improved to allow connection to a computer and exporting of the generated MIDI file. Also, it may be useful to allow the user to store more than one MIDI file on the device and to choose which file to listen to in playback mode. More control over the playback mode such as pausing in the middle of the track should be implemented.

The Maestro™ team had yet to meet to decide if we would like to pursue this project more seriously. The future plans of this project are to be determined. One possibility is that the OCR software can be redesigned and incorporated into an already existing system that has all the necessary hardware components for the implementation, like an iPhone or a similar smartphone.

4. Final Schedule

A visual representation of our project schedule is shown in by the Gantt Chart (Figure 2). This chart illustrates the expect amount of time to complete each task estimated at the beginning of the semester and the actual timeline is outlined in Figure 3. Important project deadlines and deliverables are outline in (Figure 4).

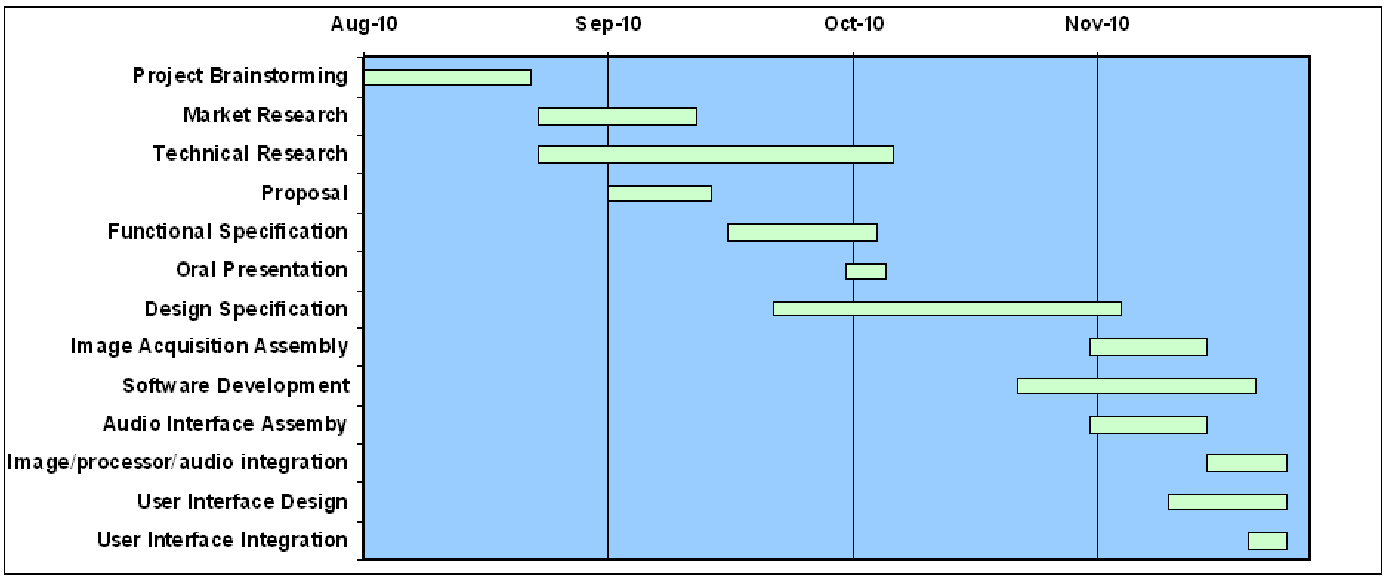

Figure 2: Original Project Schedule (from the proposal)

As you can see, the hardware assembly and integration took a lot more time than expected. This was due to the latency in parts arriving as well as unexpected interactions between hardware blocks, such as the initial camera not being able to properly communicate with the processor on the development board. All the time spent into making that camera work didn't amount to anything and we had to start again with a new camera.

Even though we started the software development earlier than we had originally planned, various setbacks kept us from achieving an optimum build. The camera module needed to work before we could develop the software to work with the actual images.

The sound module also set us back, since what we considered to be a rather simple task was actually more complicated than expected, nonetheless it should have been completed much earlier than it actually was (a few days before the demo).

Also, in the initial diagram didn't set aside time for debugging and improvements, possibly because it was assumed that any remaining time after integration would be spent testing and debugging. However, we were behind schedule in integration and we only had 3 days to test and debug the system and most importantly the OCR software now running on the processor. Most of that time was spent fixing issues resulting from running complicated software on an embedded system, such as running out of memory for variables (we only had 30KB of that) and stack overwriting making us have to optimize the code for it to run properly. We had very little time for debugging and improving the algorithm itself, which is why the music recognition had issues during the demo.

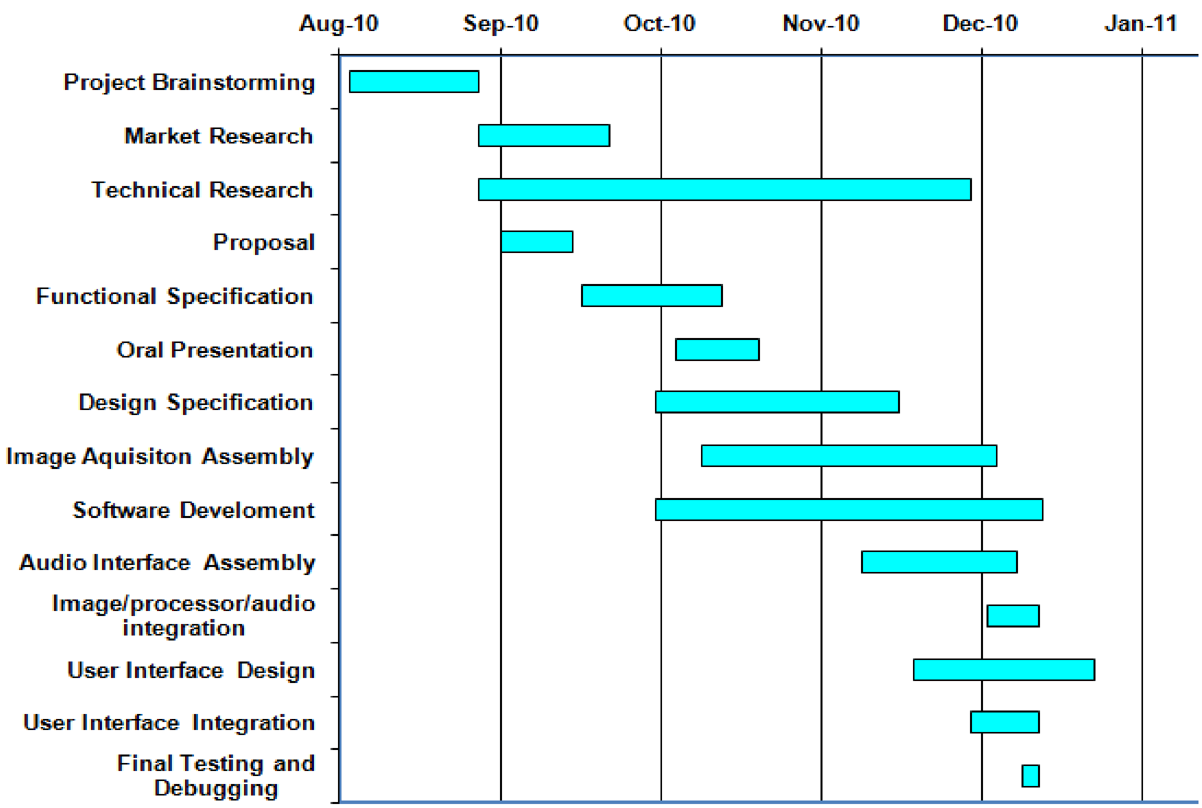

Figure 3: Actual Project Schedule

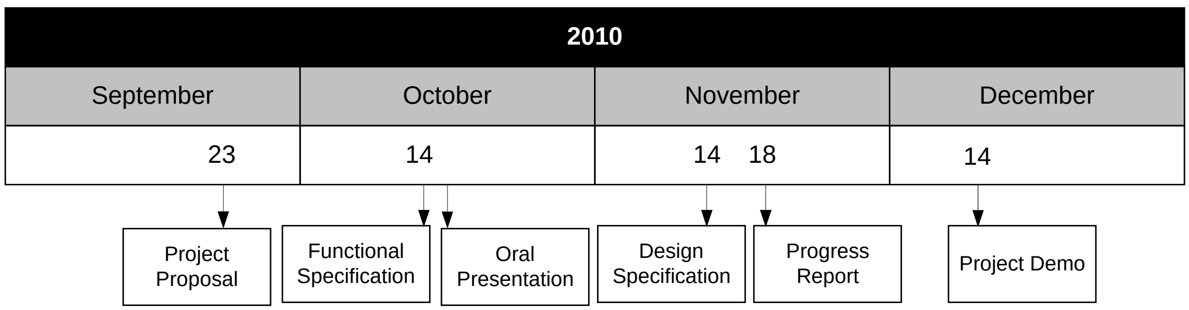

Figure 4: Project Deliverables

5. Final Budget

Our final project cost is outlined in the table below (Table 1). Significant portion of the expenses represent the three camera modules acquired. Since our initial camera module turned out not to be compliant with our EVK1100 microprocessor evaluation board we had to switch to our alternative camera module. Further expenses for the camera module have been introduced when we fried a camera module. Never the less, the total project cost was still within the original estimated budget of $800.

Table 1: Total project budget

Equipment List	Cost
Microcontroller & EVBD	$129
Sound module	$44.93
Camera_A Camera_B1 Camera_B2	$60 $105 $105
Microcontroller programmer	Free
Electronics	$160
Miscellaneous	$150
Total Cost	**$753.93**

6. Group Dynamics

This section outlines how the members of Harmony Innovations Inc. worked as a team. Furthermore, it outlines how the tasks were allocated and the technical difficulties encountered with each associated task. The interpersonal difficulties are discussed as well as the recommendations for future projects with respect to group dynamics. Below is a top level general overview of the group sub-block interactions and how the members interacted (Figure 5).

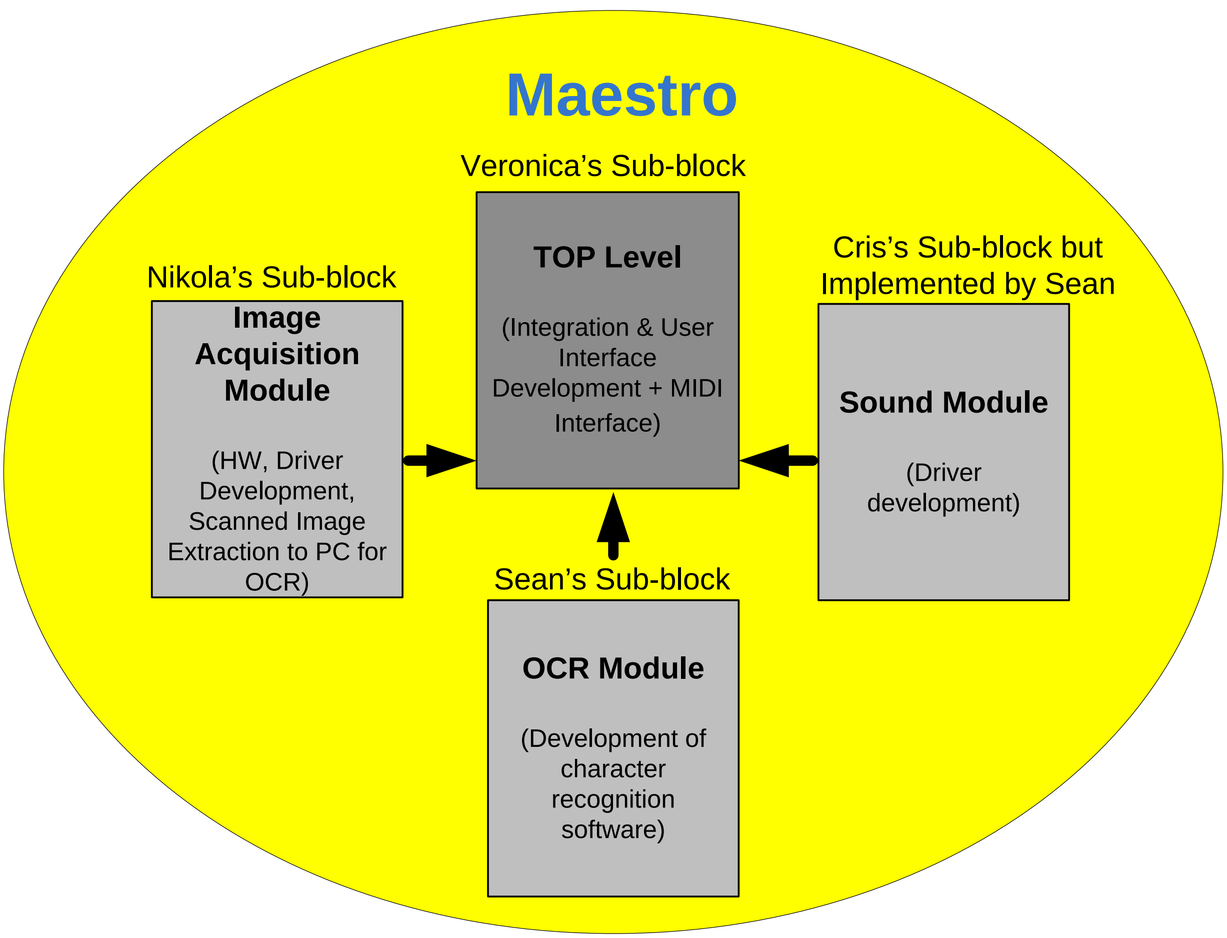

Figure 5: Maestro™'s Team Interaction Diagram

Figure 5 shows the blocks of work that were allocated to each individual in the team The choices of tasks were based upon the skills of each group member and their interest in taking up a particular task. This diagram held true for a majority of the project. Close to the end of the course there was a lot of resources shifting towards the blocks that needed help. For example, each sub-block was completed its owner would help out Veronica to implement the code at the top-level.

Our project management system worked well until we encountered difficulties with resource sharing and a member that was unable or unwilling to complete his section of the project. We had shifted many resources towards helping him with that task, including allowing him full use of one of our two development boards and JTAG cables used to program the boards. Unfortunately, positive results were not forthcoming, and the sound module had to be taken over by the rest of the group (Sean was the one responsible for finally making it work with Veronica helping to debug the code).

Unfortunately, this did not allow for any spare time at the end for debugging and code review, which was sorely missed.

7. Individual Reflection

Nikola Cucuk

Was in charge of the image acquisition module, top level integration support, and developing a portion of a suitable test environment for the OCR. Having so many options to choose for the image acquisition module, two cameras have been short listed and acquired. Extensive time was spent at the beginning of the semester just to get familiar with the interfaces and adequate image acquisition strategies, such as focused illumination on the scanning surface.

Our first camera module C3188A was ideal for out project. Firstly, image acquisition was fast, 60 fps, which is crucial for error reduction of the OCR software. Because more image samples will help the algorithm distinguish different characters, user would be able to scan the sheet music at faster speeds and the image stitching would be less error prone. Second, the user interface was straight forward and required little coding. However, we have made a group decision to ditch this camera module because interfacing it with our Atmel CPU required us to choose between, supporting high speed interrupts from the camera module or using the SPI interface which is used to interface with the sound module and the LCD display. Our initial goal was to handle camera interrupts using a SW interrupt handler but the speed of the CPU was not fast enough to capture adequate interrupts. For example, we were able to capture every other pixel and as such the resulting image was shifted and not adequate for our OCR SW.

Our second camera module uCAM has similar characteristics to our initial module. However, it is capable of outputting only 2.5 fps at the resolution suitable for our OCR SW. This was due to the slow RS232 interface between the camera module and the micro controller, but this was a price we had to pay such that we can get the project up and running to test our OCR software. Since the camera module was the big bottle neck of this project we rushed this sub module development early on into the semester. By doing so we got adequate images acquired using a custom illumination system. However, this milestone slipped almost three weeks and as such delayed the OCR debugging stage. This slip was due to the fact that one of the uCAM modules got fried and we had to order another one. We have worked independently on these sub-sections, for the most part, and the integration of the sub section was relatively straight forward.

The most important lesson I have taken away from this experience is, perhaps, the benefit of working in a team. My group members come from other specializations than myself, and that was helpful when I encountered difficulties. Another crucial lesson is the advantage of parallel programming and its ability to speed up code development.

Sean Edmond

Reflecting back on the project, I'm very satisfied with how the project came together. I assumed the role of group leader earlier on in the project. After all, as this was my idea I had an invested interest in seeing this come to fruition. I took the role of managing the different specifications, and ensuring that the different components could be integrated (relatively) seamlessly. I selected the second camera and the sound module for the project, and architected the system from a higher level. It was an extremely valuable experience to be put in the role of managing and motivating the group, especially in the earlier stages. Nik and Veronica did an amazing job with their tasks and that the project wouldn't have happened without the tremendous effort that they put in and without their continuous communication.

I was responsible for the music OCR software. Knowing almost nothing about image processing, this was an extremely valuable experience. The initial OCR software was prototyped in Matlab because it was a quick and efficient way to develop the algorithm. Not having actual test images captured from the camera, I tried my best to develop these algorithms using screen shots. My algorithm was separated into three major components:

- determining the overlap between adjacent images
- isolating the "features" in the music for recognition
- performing a "black histogram" comparison with these features to determine what type of note
- determining the associated midi pitch based on the feature
- useing the midi driver functions developed by Veronica to create a playable MIDI file

The resulting code was a set of about 20 different functions and thousands of lines of code. Perfecting the method of black histogram comparison was challenging, especially as I was trying to detect complicated features such as beams and rests. It was also extremely challenging to make this algorithm scalable so that different sized music can still be recognized with the same algorithm.

My knowledge of music, and my experience in software development assisted a lot in this task. However, my experience in algorithm development is limited and I should have employed more statistics in my algorithm, and fundamentals of image processing. I also feel that there are too many limitation with my histogram comparison method, and that other options should be have been explored.

Through our testing, we found that the algorithm performed at 60-80%. We managed to get an output from our device that sounded a lot like happy birthday! More time for testing would have allowed us to find the bugs and deficiencies in the algorithm.

A huge frustration in the project was managing the sound module portion of the project. Cris agreed to take on this initially and showed little motivated in selecting a part. I eventually chose the vs1053b because it seemed easy to interface with the EVK1100. Cris ordered this module 2 weeks after our internally set group deadline. He also demonstrated poor understanding of this device when writing his section of the design spec. I re-wrote this portion of the design spec, intending to give Cris an easier reference than the 80 page datasheet. Although Cris put in effort to try and get this working, he managed to output very little on his own (none of his code was used in the final project). I feel that I gave him a task outside of his skill set, but I also feel that he didn't focus enough in trying to understand the tool, the vs1053b chip and the SPI driver packaged with the AVR tools. I was always available for help, but he rarely met outside of project meeting time. I feel his task was manageable for the 3 months time frame and think with a more focused engineering approach he could have done this on his own. Towards the end of the project the group gave up trying to include Cris as we knew it would be easier to complete the work without him. Getting the sound module to work was the result of my continuous investigations throughout the project, and about 4 hours of work with Veronica last Saturday.

Working with Veronica and Nik was great. Since the OCR software was the bridge between images with the MIDI generation, I had to work with both of them quite a bit throughout the project. They both exceeded my expectations, and it was a pleasure to work with them.

Veronica Cojocaru

My major role in the project was to design the user interface and integrate all the hardware and software components into a working system. Initially, I was helping with the algorithm development, but as Sean took ownership of the OCR software, I shifted to more of the top-level design. Early in the project work, I had to communicate to all the group members to figure out how the components would work and fit together. When the development board arrived, Sean and I had to figure out how to use the tools to interface with the board and how the development environment worked. It turns out that the tools were very buggy and error-prone, and this was a constant source of frustration. We couldn't switch to newer versions of the tools, since they didn't work properly with both the EVK1100 development board and the sound module.

We managed to improvise and I started development on the user interface: programming the use of the LCD screen for the menu and the buttons for the user inputs. I put in wrapper functions for the components that were not yet developed to make integration easier when the time came. I tested the user interface and made sure that the components worked properly. When Nik finished getting the camera working, we scheduled time to meet and work on integration. Over the course of a few days, the camera was working properly with the development board and we also set up a debug interface in order to extract the images captured by the camera to allow Sean to test his algorithm with actual data. Next, the sound module was integrated into the system, which took some debugging as we were having issues with the different clock frequencies used by the camera and the sound module.

Time is NOT on your side during this project! Initially we had planned that we would get all the hardware working and integrated a couple of weeks before the demo, and the remaining time would be spent testing, debugging and improving the system. However, setbacks kept us from achieving this milestone and it wasn't until 3 days before the demo that the whole system was integrated and functioning. This left very little time for algorithm improvement and debugging – in the end the music recognition is only about 80% accurate. You need to be aware of the bottlenecks in your design process and try to accommodate as much as possible. In terms of this, there was not much that we could have done differently besides start working on the project sooner and ordering parts earlier.

Overall, the experience of working on this project was very positive. Besides the technical skills I acquired and improved, I also learned a lot about problem solving, the design process and group dynamics. Speaking of group dynamics, your entire experience in this course can be defined (for the good or the bad) by the individuals in your project group – so choose them wisely. Late nights of desperately trying to get things to work can be much less painful if the effort is shared with other dedicated group members. However, it is very frustrating when you are putting in every bit of time you can spare and all the effort you can reasonably manage into getting the project to completion, and another person in your group is not even putting in the minimum amount of work required to get their part functioning.

Nonetheless, the experience taught us a lot about project management and how to adapt to this sort of situation and shift workloads and priorities to accomplish your goals on time. I certainly appreciate how professional Sean and Nik were during the course of the project and it certainly couldn't have got as far as it did without their effort and dedication.

Cris Panaitiu

This project is supposed to be the accumulation of several years of engineering experience, all applied towards solving one problem. It is a challenge to apply all of one's skill sets to practical design problems within constraints like time and budget, which limit the engineer's avenues to solving a problem. It is a further challenge to work within a group where each member has different skills and most importantly different ideas about how to solve the problem. An engineer will know how to solve technical problems, but also work with the team mates in the group.

Throughout the project it was important to have a team leader that assigned deadlines and tasks based on how each team member seemed to work best. Although there was no formal selection process or clear definition of team leader, Sean ended up fulfilling this position's responsibilities mostly throughout the project's duration. A team leader should be able to understand every other member's duty well enough that he can do it himself; this being so that the leader can set realistic deadlines and announce decisions that have the support of the group. Sean set deadlines for the group to follow for documents and progress and meetings and this contributed to an organized working environment.

I have not had the opportunity to apply very many of my skills to this particular project, having been assigned a task that requires embedded programming. It would have been more beneficial to the project if it had been determined earlier that the work I had been assigned was progressing very slowly but the tasks that were assigned at the beginning were maintained and not changed throughout the semester.

The meat of the project was algorithm design and may have been better suited to my abilities; I have had experience working in matlab in many courses and individual projects and in retrospect can say that learning a whole new skill set properly during the course of two and a half to three months was not as easy as it sounded at the beginning.

I can differentiate my experience in this course from other courses with major projects (such as ensc 427 and other similar courses like 425 and 428) whose topic is at the discretion of the student. While in other projects, the team dynamics were not so well documented, they can be compared to those in this project. The division of tasks was similar, there was often a leader, and there were deadlines (and thus crunch time at the end). The only difference that I felt personally was how the group adapted to stress and failed results, as well as slower than expected progress. In other groups, it has been my experience that when a team mate struggles, it is productive to work together with that team mate thoroughly (to the limit of your possibilities) on his/her task until he/she is able to complete the task. This is my personal approach to large projects in a group. The extent of working together with that struggling team mate, the more they will learn, and you as well. I have experienced positive results with this not only because the work gets done faster and gets done well, but positive relationships can be formed this way and future project work is facilitated. I did not feel that our group environment was such; rather each individual worked on their task and communicated only results to each other. While I cannot fairly have an expectation that any group member do my work or spend their time to teach me or work with me because this is an unwritten rule of good team dynamics and the extent to which it is applied is up to the discretion of the individual.

The other members of my team were very talented and motivated for this project and it was a good experience having the opportunity to work with them. I have drawn many lessons about working in a team and communicating effectively and whereas I would have liked to pick up some of the technical side better, the reason this course is fulfilled.

9. Conclusion

In short, at this point in time the Maestro™ is not a marketable product as the algorithms used for optical character recognition are underdeveloped and in their infancy due to their complexity. Full scale tests and further time is needed to develop more accurate optical character recognition software. However, the current prototype is a great platform to further develop and debug these algorithms with great ease. The current tests results that we obtained prove the concepts feasibility. This possibly means that with further development our product will be able to recognize sheet music and play it accurately. A lot of development time was spent on acquiring clear and adequate images and playing the MIDI files, but the OCR software needs more enhancements such that we can finalize our prototype and start developing a marketable product. We believe that this product is capable of scanning sheet music and we would like to peruse our goal of correctly playing the scanned sheet music in the future.