

Availability Prediction Based on Multi-Context Data

PROJECT PLAN

Group: sdmay19-33

Client: Goce Trajcevski

Team members:

Justice Wright - Report Facilitator

Shane Impola – Scribe

Noah Chicchelly – Meeting and Communications Facilitator

Nick Schmidt – Software Systems Engineer

Tristan Anderson – Network Systems Engineer

Brendon McGehee – Hardware Systems Engineer

Team email: sdmay19-33@iastate.edu

Team Website: <https://sdmay19-33.sd.ece.iastate.edu>

Version 1.0

Revised: September 28, 2018

Table of Contents

1 Introductory Material	4
1.1 Acknowledgement	4
1.2 Problem Statement	4
1.3 Operating Environment	4
1.4 Intended Users and Intended Uses	4
1.5 Assumptions and Limitations	4
1.6 Expected End Product and Other Deliverables	5
2 Proposed Approach and Statement of Work	5
2.1 Objective of the Task	5
2.2 Functional Requirements	5
2.3 Constraints Considerations	6
2.3.1 Non-Functional Requirements	6
2.3.2 Standards	6
2.4 Previous Work and Literature	7
2.5 Proposed Design	7
2.6 Technology Considerations	7
2.7 Safety Considerations	9
2.8 Task Approach	10
2.9 Possible Risks and Risk Management	11
2.10 Project Proposed Milestones and Evaluation Criteria	12
2.11 Project Tracking Procedures	12
2.13 Test Plan	13
3 Project Timeline, Estimated Resources, and Challenges	
3.1 Project Timeline	14
3.2 Feasibility Assessment	14
3.3 Personal Effort Requirements	15
3.4 Other Resource Requirements	15
3.5 Financial Requirements	16
4 Closure Materials	
4.1 Conclusion	16

4.2 References	16
4.3 Appendices	16

List of Figures

Figure 1: Technology Considerations

Figure 2: Task Approach

List of Tables

Table 1: Project Timeline

Table 2: Personnel Effort Requirements

List of Definitions

MySQL - My Structured Query Language

AWS - Amazon Web Services

VM - Virtual Machine

1 Introductory Material

1.1 Acknowledgement

The Availability Prediction team would like to thank Dr. Trajcevski for all of his technical and project advice given throughout the duration of this project.

1.2 Problem Statement

Currently in a restaurant atmosphere there are multiple data sources that could be better used to predict a better recommendation for table wait time.

Our goal is to create a device that can collect and organize data from a restaurant environment such as, how long people take to eat, how many people are sitting at a certain table, and what tables are being used. The data can then be analyzed to answer many questions about the restaurant.

1.3 Operating Environment

The expected operating environment of this project will be inside of restaurants. The system should not be exposed to any harsh weather conditions because it will remain inside in a room temperature environment. The only condition that the system might be exposed to is some dust and debris over a long period of time.

1.4 Intended Users and Intended Uses

Our intended final project users are going to have a wide range of technical knowledge. The individuals that are going to be using this are both customers and employees of the restaurant that it will be implemented in. Although customers and employees will be using this in an app based form, the purpose of the app for each will be very different.

The employees will be using this app to input restaurant data such as when the food has been given to the table, when the food has been removed from the table, and when the check has been given to the table. This information will then be used to better predict table wait time, which brings me to customer use. The customer will use the app to see available tables as well as receive restaurant wait times.

1.5 Assumptions and Limitations

Assumptions:

The final product will only be implemented indoors - the hardware of the project does not need to stand up to any harsh weather conditions.

Hardware component will not have any power limitations - the hardware component will have access to electrical outlets.

Limitations:

Sensors need to be unnoticeable inside a restaurant - The sensors used to collect data need to be implemented into a restaurant setting without impeding the normal operation of the restaurant.

The application needs to be usable by all types of technical backgrounds - The app that we will be implementing in our project can not be overly difficult to use because people of all technical backgrounds will be utilizing the application.

1.6 Expected End Product and Other Deliverables

The final product of this project will include sensor with a microcontroller to collect data about tables, a networking component to organize and analyze the collected data, and an app component for either an employee or a customer. The components need to be able to communicate with each other but they do not need to communicate individually. The sensor component needs to be able to communicate with the networking component and the networking component needs to be able to communicate with the software component.

2 Proposed Approach and Statement of Work

2.1 Objective of the Task

The objective of the project is to design a product that will enable consumers and producers of restaurants to easily access real-time data about seating availability in a specific restaurant. The product will be a mobile app and possibly web app available to customers and employees. The team goal is to implement the functionality required to have a usable app with this information. This will include hardware installed in seating areas and a fully developed server and database to communicate with our application.

2.2 Functional Requirements

The functional requirements will be split into customer use, and employee use. There will be two forms of the application for each of these users. Below is a list of the customer's function requirements. The expected use case is to find a location you're interested in visiting, check for generic wait times of that day, or click on more specifics for a top down view of available seating, or an estimated time of arrival for a given table.

- Choose a Location
- View an estimated wait time for that time of day
- View available seating
- View an estimated wait time for a specific table
- Real time updates on when tables become available (push notifications)

Employee use cases are going to be a lot of the same, but with a few extra perks. Employees will be able to see specific state's of dining that a customer is in, as well as what was ordered:

- Choose a Location

- View an estimated wait time for that time of day
- View available seating
- View an estimated wait time for a specific table
- Real time updates on when tables become available (push notifications)
- View where a specific table is at in their meal (food ordered, food delivered, etc)
- View what was ordered at a specific table
- Override / shutdown specific functionality of the app (Table marked as out-of-order)

2.3 Constraints Considerations

2.3.1 Non-Functional Requirements

Availability - The app must be up 24/7 for estimated times at specific hours. Even if closed, information for future events is vital.

Data Integrity - Making sure our data is always accurate over the span of the app's lifetime. Allowing businesses and consumers the ability to trust the app and use it flawlessly.

Fault Tolerance - The app will continue to work if one system goes down. For example, if our live view of seating fails to work, you can still get an estimate of the wait time for that time of day.

Scalability - Restaurants are different shapes and sizes. The app must be able to accommodate a large or small number of seats, seating arrangements, and real-time data at smaller, or huge scales.

Usability - The app should be easily accessible by clients as well as employees. An intuitive way to view the information you want, immediately.

2.3.2 Standards

The standards of the project will be built with a few assumptions. For one, we will use MySQL that has standards built-in. However, with the data release of people, we will be working under the assumption that the data is of our own use. In the scope of the project, that is not one that we are going to be focusing on. Under reasonable circumstances we will be releasing data that is known by at least an employee of a restaurant, nothing that isn't "public" to the restaurant.

The customer use-cases of the app will also be limited to very generic information and ETAs, rather than giving specific data about what a table is or is not eating, or how much their bill is. This is trivial for standards and practices. Employers will not release that information, but have it at their use to better help customers with information they may request.

2.4 Previous Work and Literature

The app will resemble a mix up of a few different entities that are already in the market today. The first one is google:

Google takes information from restaurants to estimate the popular times of a restaurant. They do other things like general prices, reviews, etc. But in relation to our app, google only touches on one aspect, the estimated “popular times” of a given restaurant. Our app will have this same feature but implemented in a more specific way. Google uses the generic approach of taking users locations from their phones and aggregating the data. This can make a statistic fluctuate if people are there to eat, eat for a long time, just buy something nearby, etc. Our app seeks to perfect this feature and make it more accurate.

The second thing on the market that our app might resemble is event seating services:

You can now buy tickets to specific seats for a large event, or movie, and see top down views of the seating arrangements with what seats are what price, and if they are taken or not. This is the same process we want, but for restaurants. Specifically for wait times on specific tables, or seeing how busy it is in real time, from information given from us, the restaurant. Rather than aggregated data from location pings of mobile phones.

We are not following previous work for this. This means we have to implement all aspects of functionality. Starting with hardware implementation to give us the data we require to implement our functional requirements within the app.

2.5 Proposed Design

Our design will be relatively simple.

1. Find a way to gather the required data (Are people at a table, where are they within their meal, etc.)
2. Communicate that data to a server and database for organization and availability
 - a. We will have dynamic and static queries for both consumer and employees
 - b. We will have the ability to speak to the hardware, as well as the front-end app.
3. Implement the front-end application for employees and customers to use in real-time.

Our scope is relatively well defined. There are a few design alternatives to use/non-use cases. And there could be widened scope by adding in security and data-release-standard requirements. These are all possibilities depending on how quickly things move throughout the Fall semester of 2018.

2.6 Technology Considerations

There are three essential subsystems to our overall design, each with their own technological considerations: the data collection system, the data analytics system, and the client-side application, all of which are pictured below.

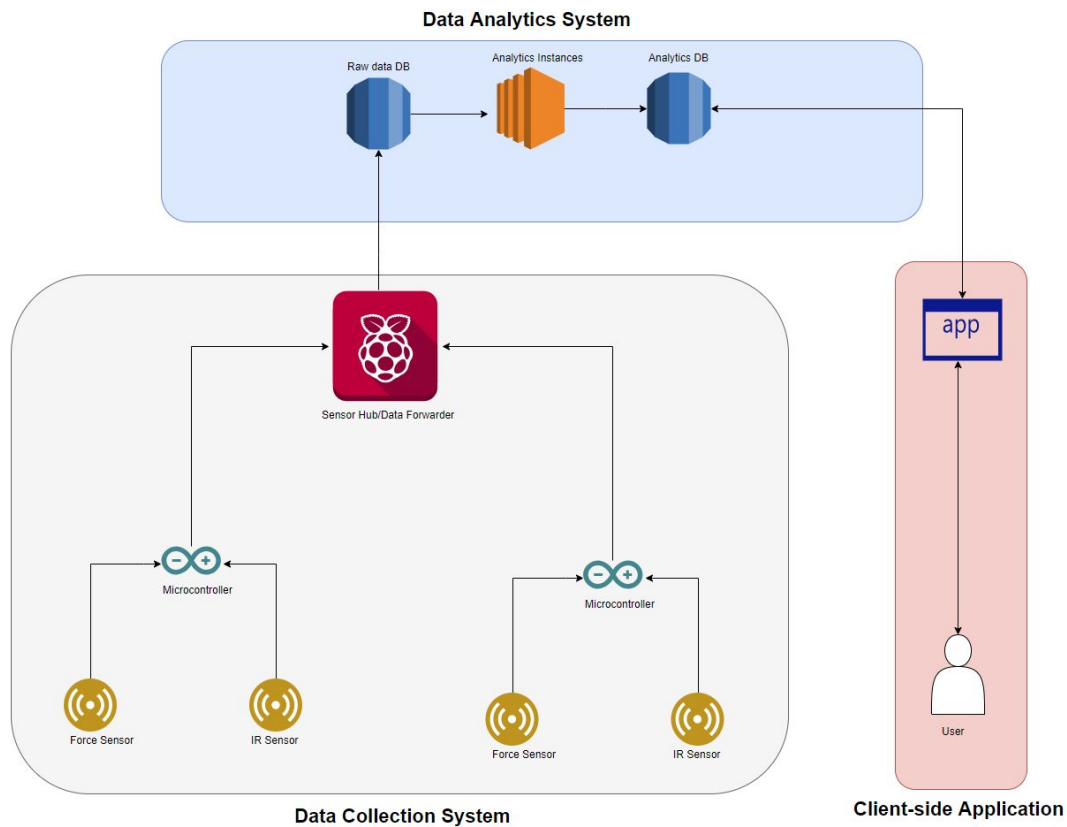


Figure 1: Technology Considerations

For the data collection system we needed multiple sensors to help determine occupancy of a table. Force sensors (load cells) were the original consideration, to be placed beneath the seats, thus triggering upon the sitting down of a guest. While this is a worthwhile datapoint it was determined further sensors would be necessary as both a precaution against false positives and a way to handle booth seating. Because of the nature of booth style seating, it is unsustainable to put enough force sensors in the seating area to accurately detect a seated occupant as weight is distributed more widely around a booth.

To overcome this shortcoming the decision was made to introduce IR sensors to the seating environment. IR sensors operate outside of the physical seat, which means they do not share the weakness of the load cells. In addition, seating is much more linear in a booth scenario, meaning IR blocking is much easier to achieve in such a scenario, so the addition of IR actually functions best in our previously worst visibility situation.

Both sensors are configurable to be controlled by a microcontroller which will both provide power and input necessary to the sensors as well as collect and transfer data from the sensors. In this instance we chose to use the Elegoo UNO R3 because we have team members who

have experience using them, and they're the cheapest microcontroller commercially available with our desired chip, merging familiarity with lower production cost.

From here the sensor data needs to make its way to our data analytics clusters. We decided against outfitting each microcontroller with networking capability due to concerns over scalability concerns with both cost overloading our servers with too many connections in production. Instead we opted to aggregate the data from each microcontroller into a raspberry pi which bridges the data to our analytics system. This decision was made part to due accessibility since many team members already own and were willing to contribute pi's to the project as well as the price of a single pi with connectivity coming in at less price/unit than equipping each Uno that a pi could manage with wifi modules.

For the data analytics system we knew we needed multiple databases to store our sensor data both pre and post analytics as well as instances for performing the real-time analytics. This could be achieved in a multitude of ways, and the real decision here was where to host this system. Ultimately it was decided to use Amazon Web Services (AWS) for this system. This was decided because AWS, in addition to being the industry standard at this point, doesn't require us to operate or own our own hardware, and has free offerings we felt we would stay inside for the scope of this project making costs non-existent for this system. Additionally were our product to go into production AWS scales well if we ever escape the data usage of their free environment and offers many security tools and compliance aids that could be vital to expansion.

Now that we have our environment to host our analytics in, we required a means to access the end data and relay that information to our customers. Given that many restaurants already employ tablets in some form for their seating systems it was decided fairly early that some form of mobile application. The decision point here is what platforms will we support (iOS, Android) and what language/tools will we use during development. The decision to use Flutter was made for multiple reasons. Flutter is an open source project created by Google that is used to create mobile apps for both iOS and Android. The open source aspect means again that there will be no cost of creating the application, of which itself could eventually be sold to increase profitability of our system. The fact that both iOS and Android development are supported means that we can produce a product that should fit any companies existing infrastructure. Additionally Flutter uses Dart as its primary language, which while no one on our team has much direct experience with, is an object-oriented programming language similar to Java which our entire team has experience with.

2.7 Safety Considerations

The ideal testing scenario for our equipment is a mock restaurant booth where patrons are unaware they are even interacting with our product. Because of this there become several safety considerations.

Sensors will need to be able to withstand and function under the weight of a fully grown adult. We rely on load cells to detect occupancy of seats, and a load cell breaking would both render our product useless, pose as a threat to our customer in the form of debris potentially puncturing their skin, as well as pose a fire risk if the wires were to come loose during the sensors destruction. To ensure the safety of our customers care was taken to inspect the datasheet for max load capacity and compare that against our expected customers. Additionally the load sensors would be put through numerous tests with clients of varying sizes before ever reaching production, and if a defect was discovered the product would not launch until an alternative sensor was found. The connection integrity of the wiring would also be checked routinely during testing, and done by Brendon, who has the most experience with proper circuit configuration and maintenance.

Restaurant scenarios also introduce the possibility of food spills (both liquid and solid) being introduced into our system. To reduce the risk of shock and failure the microcontrollers have been designed to be mounted outside of the typical dining area, adjacent to the booth where spills are very unlikely to occur. All sensors have been designed to be placed within seats, tables, or other existing dining infrastructure, as to add a layer of abstraction between themselves and the dining surface, thus reducing the risk of direct contact.

2.8 Task Approach

Our approach to solving this problem will have three main components. The first component is the hardware, which will take the form of the various sensors that are embedded into the client environment. We need our sensors to be able to collect various bits on information about the occupants of a given table, so we need to find types of sensors that will be able to collect to most robust information possible. The data points that they generate should be able to be used in multiple predictions, and they should be relatively straightforward to implement. We decided to use both IR and force sensors, and pair each table with an Arduino Nano with internet connectivity to gather the data generated by the sensors at the table. This approach allows us to gather sensory information locally at each table before sending and relevant information directly to the server. There will be no need for a central computer or server for the restaurant to maintain in which all sensors will be connected. Ideally, this approach should be cheap and simple to install, while requiring minimal maintenance.

The second component of our solution are the web services. This is the part that will be doing all of the data aggregation, storage, analysis and prediction, and broadcasting of the processed information. We decided to use AWS for our web services because it meets our needs while being cloud-based, meaning clients won't need to keep and maintain server hardware in-house. This also has the effect of reducing the initial cost of implementing our solution significantly. There will be two components to our usage of AWS, the server and the database. The server will be used to aggregate any information sent to it by the sensors and use it to update the database. It will also periodically perform analysis on the contents of the database, making predictions that will be stored until the next one is made. The server will also handle data requests from employee and customer phones. The database will have relational tables that will

store all the information from each sensor type and for each table, and an input table where unprocessed entries will be buffered. This approach allows us to make queries about information specific to one table, which might be more useful to guests, or for information about many different tables, which would be more useful to employees. The web services can be visualized in Figure 2.




Figure 2: Task Approach

The third component of our solution is the user interface, the apps that employees and guests will use to get predictions. We need the user interface to be easily accessible to both types of users, so using mobile applications is what we've decided to do. The apps will need to be able to have updates pushed to them by the server and pull updates on demand. Using mobile apps also eliminates the need for restaurants to buy proprietary hardware for customer use, akin to the buzzers commonly seen today. Theft or damage of the user interface will no longer be a concern because customers will be using their own devices.

2.9 Possible Risks and Risk Management

Risk: There persists the possibility that our current sensor array won't provide enough data points to allow for accurate prediction analytics to be performed. As our weakest area of exposure is data analytics there's no one on our team that has a good enough grasp to rule this out.

Risk Level: Medium

Mitigation: Our design to use a series of microcontrollers to control and feed data from individual sensors leaves room for expansion. We are not currently using nearly every input port available to the controllers, and could add sensors to each system very easily. If the need for inputs ever exceeds our microcontroller we can simply scale up to a larger microcontroller.

Risk Level Post-Mitigation: Low

Risk: Our team has no experience with Dart or AWS which are cornerstones to the software side of our project. There lies a risk that this inexperience may be insurmountable and hinder progress and push back deadlines.

Risk level: Low-Medium

Mitigation: We have fallbacks that we are more familiar with for both Flutter and AWS if the learning curve starts to hinder our progress in a meaningful fashion. The ece department offers VM and database services that we have utilized for previous classes, and could use to replace AWS if need be, four team members have experience with these systems. Three team members also have experience coding directly for android using java and Android Studio if Flutter proves to be a choke point. Both of these replacement opportunities are highly viable and still are free offerings, meaning using these as fail-state fallbacks would only incur costs of time.

Risk Level Post-Mitigation: Low

2.10 Project Proposed Milestones and Evaluation Criteria

The first milestone would be to gather data from our sensor and process it through a microcontroller. This will be verified by configuring the controller to read and save data it reads from the sensor, purposefully triggering the sensor, and examining the data to ensure the sensor was triggered and the controller was aware the trigger.

The next milestone is to send this sensor data to a database for storage. To verify this, the sensor system consisting of the individual sensors, and microcontrollers will be hooked up to a raspberry pi, configured to transmit the sensor data to a database in AWS. We can then query that database in real time, and observe data being populated, verifying this milestone.

A third milestone will be configuring instances to consume the data in the database and perform predictive analytics with the data. To verify this we can load a dummy database with known data, and have our analytics instance perform analysis of this data. We can by hand determine the outputs we expect to see and confirm this milestone is achieved if the results match.

Another milestone will be interfacing with our web instances to observe the output of the analytics in a client-facing application. This can be verified by starting with a blank database, triggering the sensors, and verifying that the app updates corresponding to the presence of new data.

2.11 Project Tracking Procedures

Trello and git and discord and weekly reports

2.12 Expected Results and Validation

The expected end result is a hardware system capable of relaying sensor data to backend analytics system that can deliver real-time predictability models to a front-end application used by the customer that meets all the requirements and goals outlined in section 2.1.

Validation is built into the design process, as each new analytic algorithm will need to be individually verified for correctness, each piece of hardware or sensor incorporated will be checked for reliability of data/functionality, and each feature of the client-side application will need to have accompanying use-case tests.

However to ensure that the final product is more than just the sum of its parts, validation for the final deliverable product will be performed and reported according to the test plan outlined below.

2.13 Test Plan

Hardware Tests

FR.1: This is a test program for the an Arduino Nano that prints out sensor values in real time to test if the sensors function properly.

Test Case: Test if the sensors are outputting data as expected.

Test Steps:

1. Start the test program.
2. Manipulate the sensors by hand.
3. Watch the printed output and determine if they are as expected.

Expected Results: The printed output should change corresponding to the manual changes in the state of the sensor.

FR.2: This is test software that will take a set of raw input from the sensors and process it in order to determine if the arduino is processing sensor output properly.

Test Case: Test a set of data on both the arduino and a computer

Test Steps:

1. Run the arduino program, enabling a debug option that will send all inputs to the computer.
2. Simulate the data processing on the computer.
3. Compare the result of the arduino program to the simulation.

Expected Results: The output from the arduino should match that of the simulation.

Server Tests

FR.3: This is a connectivity test that will ping the server and await a response to determine if the server is online and responding to requests.

Test Case: Test that the server responds to requests as expected.

Test Steps:

1. Send a request to the server test address.
2. Wait for a response.
3. Compare the response to a preloaded value.

Expected Results: The response from the server should match the preloaded value.

FR.4: This is a database test that will test the functionality of the database.

Test Case: Test that the database handles inputs and requests as expected.

Test Steps:

1. Send a request to the database test address.
2. Wait for the server to perform a set of automated queries on the database and respond with the result.

Expected Results: The result should indicate that the tests were successful.

App Tests

FR.5: This is a set of test functions on the apps that implement the same functionality as described in tests FR.3 and FR.4 to determine if the app is communicating properly with the server. See those tests for details.

3 Project Timeline, Estimated Resources, and Challenges

3.1 Project Timeline

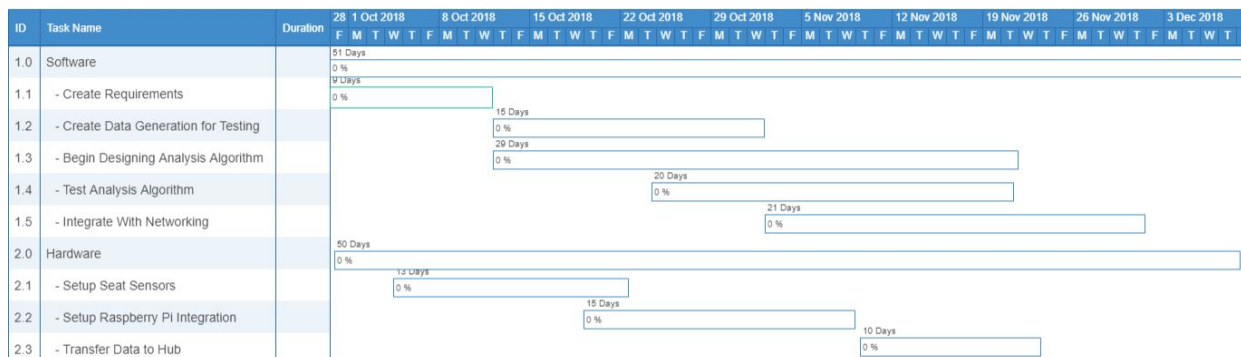


Table 1: Project Timeline

3.2 Feasibility Assessment

By the end of this project we expect to have an app capable of showing a live count of the current wait time for seating. We plan to include a version for both customers and owners. The customer version will show available tables and the live count. The owner version will contain detailed information and analytics about wait times at various times of the day and wait times for specific tables/foods.

On the hardware side of things we will have a system built to show a proof of concept. There will be at least one working table with a transmitter and seating sensors to show that the idea works. We will insure this design is modular and relatively simple to install so that the system is usable in other environments and able to be installed.

Some problems for this project would include difficulty with managing multiple tables in a scalable matter. Data being transferred and stored could become overwhelming. Additionally, formulating an algorithm to dig through such large amounts of data to calculate the wait time could become difficult.

3.3 Personnel Effort Requirements

Task	Description	Time
Test Data	Create a program to generate bogus data for use in later testing.	10 hours
Create Requirements	Create a list of the requirements needed for the software.	5 hours
Design Algorithm	Design the data analytics algorithm.	20 hours
Test and Improve Algorithm	Run algorithm through data to improve algorithm. Test accuracy of the algorithm. Possible machine learning.	20 hours
Integrate with Networking	Integrate with the networking tools and begin grabbing/submitting data to the network.	20 hours
Setup Seat Sensors	Set up seat sensors with chairs.	20 hours
Setup Raspberry PI Integration	Integrate with Raspberry PI and transmit get data usable format.	10 hours
Transfer Data	Get data transferring and entered into database/java	15 hours

Table 2: Personnel Effort Requirements

3.4 Other Resource Requirements

- Tables for testing
- Chairs for testing
- Raspberrypi/Arduinos for data collection and analysis
- Transmitters to send data from tables.
- Pressure sensors to detect occupancy.

3.5 Financial Requirements

Per Table:

- 10 nrf24l01+ transceivers with antenna. (\$1.19 each, \$11.98 total)
- 4 Infrared Proximity Sensor (13.95 each, 55.80 total)
- 4 Elegoo Nano (\$5 each, \$20 total)
- 1 Elegoo Uno (\$10.86)
- 4 Pressure Sensor (\$14.99 each, \$59.96 total)

Per restaurant:

- 1 Raspberry Pi 3 Model B (\$39)
- 1 nrf24l01+ transceiver with antenna. (\$1.19)

4 Closure Materials

4.1 Conclusion

Our goal is to design a product that will enable restaurant operators and customers to easily access real-time data about seating availability. The data will be gathered through various occupancy sensors placed around restaurants. This data will be processed to make assumptions on when seats will be available, and how customers use those seats. The data will be easily accessible on the users mobile device. If we are successful, our project will greatly improve the customer experience by reducing wait times, and it will improve logistics for operators by simplifying customer data.

4.2 References

(This will be built up over the course of the year)

nRF24L01+

<https://www.nordicsemi.com/eng/Products/2.4GHz-RF/nRF24L01P>

Arduino Nano

<https://components101.com/microcontrollers/arduino-nano>

Pi

<https://www.raspberrypi.org/documentation/hardware/raspberrypi/README.md>

4.3 Appendices

(This will be built up over the course of the year)