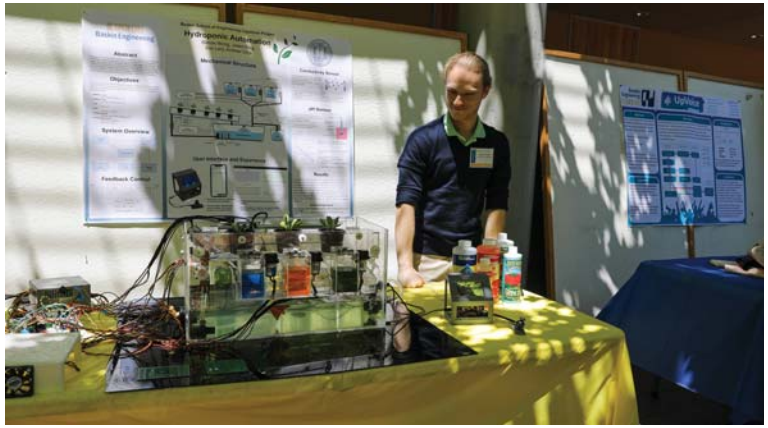

Baskin Engineering UC SANTA CRUZ

PARTNERS' DAY

CORPORATE SPONSORED
SENIOR PROJECTS
AND SENIOR DESIGN SHOWCASE

2019 - 2020 Program

Introduction

This publication highlights the ninth year of the Corporate Sponsored Senior Project Program (CSSPP) at the Baskin School of Engineering.

CSSPP provides students with a unique opportunity to experience working on real-world engineering projects as part of their undergraduate education. Throughout the academic year, students interact with teammates and hold frequent meetings with their sponsors, getting feedback on the solutions they have developed and guidance on the work in progress. By working with mentors at corporate partner companies, students learn important skills, take on interesting challenges, and begin to understand what it means to be a professional engineer.

We appreciate our corporate sponsors for their willingness to support this year-long program, mentor our students and provide them with challenging projects to work on. We also appreciate our students, who have worked hard and have enriched our lives through their energy, intellect and determination. This has been an unprecedented year as both students and sponsors have had to adjust their operations to work within the shelter-in-place regulations associated with the COVID-19 pandemic. Students worked remotely, and even those living locally had no access to campus in the spring quarter. Others worked from locations around the globe, and as campus facilities and equipment were not accessible, students had to move quickly to adapt their projects. We appreciate the flexibility and creativity of our partners and our students as they have worked through this challenging time.

The publication also includes this year's Senior Design Program Projects from student teams in Computer Science & Engineering and Electrical & Computer Engineering working on faculty/student initiated projects. As with our CSSPP projects, students working on our senior design program projects had to be agile and creative in pursuing their projects to completion in the face of limited access to critical resources and to one another. We commend them for their resilience and ingenuity.

Alexander L. Wolf
Dean
Baskin School of Engineering

Baskin Engineering UC SANTA CRUZ

ACKNOWLEDGMENTS

We would like to acknowledge and thank the faculty, teaching assistants, and staff who have been so instrumental in the Corporate Sponsored Senior Project Program:

SENIOR DESIGN FACULTY

CORPORATE SPONSORED SENIOR PROJECT PROGRAM 2019-2020

Patrick Mantey
Director, Corporate Sponsored Senior Project Program

Richard Jullig
Lecturer, Electrical & Computer Engineering

Tela Favaloro
Lecturer, Electrical & Computer Engineering

David Harrison
Lecturer, Computer Science & Engineering

TEACHING ASSISTANTS & GRADUATE RESEARCH ASSISTANTS

CORPORATE SPONSORED SENIOR PROJECT PROGRAM 2019-2020

Akila de Silva • Arindam Sarma • Faeze Brahman

Golam Muktadir • David Kooi

Flash Memory Summit

We would like to extend a special thank you to the Flash Memory Summit for their support this year.

https://flashmemorysummit.com/

CORPORATE SPONSORED SENIOR PROJECTS PROGRAM STAFF

Alexander Wolf
Dean, Baskin School of Engineering

Frank Howley
Senior Director of Corporate Development, Baskin School of Engineering

Lynne Sheehan
Network Administrator, Facilities/Machine Room, Baskin School of Engineering

Bob Vitale
Director of Laboratories & Facilities, Baskin School of Engineering

Richard Sabala
Director of Laboratories and Facilities, Baskin School of Engineering

Russell Evans
Senior Development Engineer, Baskin School of Engineering

Jeffrey Jue
Speical Agreements Officer, Industry Alliances & Technology Commercialization

Kelsey Evans Dewiler
Development Assistant, Baskin School of Engineering, University Relations

Stacy Wilson
Administrative Assistant, Baskin School of Engineering

Karyn Skemp
Communications Director, Baskin School of Engineering

James McGirk
Communications & Marketing, Baskin School of Engineering

Sponsors

SPECIAL THANKS TO OUR SPONSORS for your generous support of our Corporate Sponsored Senior Projects Program. Your time, experience and financial support were beneficial to our students and the success of their Senior Design Projects.

Low Latency Gaming

Lilian Gallon, Xudong Guo, Cody Hartsook, Howard Ng, Collyn Noda, Alan Vasilkovsky

Baskin Engineering UC Santa Cruz | Cisco

Abstract

The Internet is not built for gamers. Internet service providers and massive corporations are constantly competing to get their information across faster than the rest. This means that connection issues that you may experience when gaming online are entirely out of your control. Our goal is to find a way to remedy this problem and create a better online experience.

An Internet Freeway?

Just like a freeway is built on top of existing roads for the purpose of expediting *vehicular* traffic, we want to create an **Overlay** which builds on the existing infrastructure of the internet for the purpose of expediting your *internet* traffic!

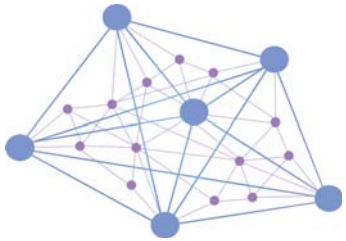

Jargon Check

Here are some of the metrics we use to measure internet connection performance:

- **Latency** is the time that it takes for a packet, a chunk of data, to go from its source to its destination.
- **Jitter** is a measurement of the change in latency over time.
- **Packet loss** is a count of how many packets fail to reach their destination.

Overview

Our new **Overlay** network consists of four modules: **Client**, **Server**, **Mesh** and **Controller**. The **Client** and **Server** modules are applications which reside on the client and server machines, respectively. They watch for packets coming from a game or application. Instead of allowing these packets to follow their intended route, they are snatched up and sent to our network of **Mesh** nodes. The **Mesh** nodes act like relays, forwarding the packets along from node to node along a predetermined route. The **Controller** module is the brain behind the entire operation. It identifies optimal routes for packets to take, keeps an updated status of our **Mesh** nodes, and stores data so that it can be accessed by any module in the **Overlay**.

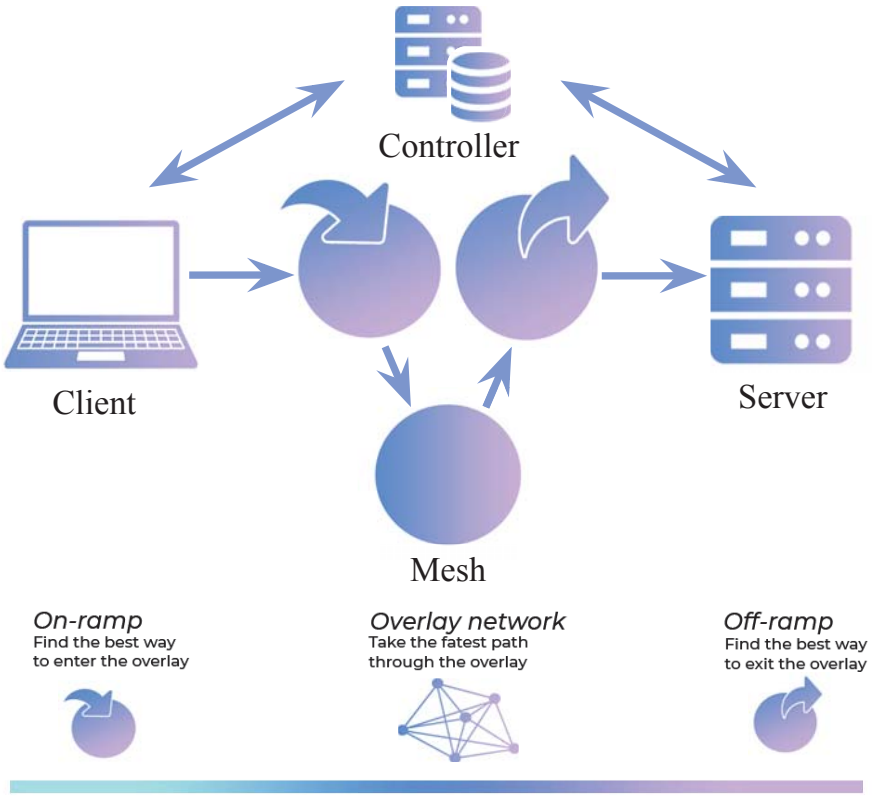

Acknowledgments

We would like to thank Alon Bernstein and Sangeeta Ramakrishnan of Cisco Systems, as well as our former teammate Jasleen Kaur, instructor Richard Jullig, and teaching assistant Arindam Sarma for all of their unwavering support throughout the entire process.

Results

In the end, we created a working prototype of our **Overlay**, complete with all four modules. We also created two separate methods to gain access to the **Overlay** network: A Software Development Kit which developers can add directly into their games, as well as a native desktop application for the gamers.

A Scenario

You live in Los Angeles and you are playing the hot new online game: *First Person Shooter III*, which has a server located in Dallas. According to Cisco's projections, here are some of the improvements that you may experience **with** and **without** using our **Overlay** network**.**

Average Latency

33.8 ms	38.2 ms

Average Jitter

354 µs	414 µs

Average Packet Loss

0.03%	0.06%

The Effects

Lower latency, jitter and packet loss mean that gamers have to worry less about connection issues during crucial gameplay moments. All they have to do is make sure they are on top of their game.

Cisco recognizes that this new technology is useful for more than just gaming. Countless other applications and industries could benefit from these improvements in internet speed, connectivity and reliability.

Wireguard Connection Manager

Cameron Saifizadeh, Jonathan Huey, Sinclair Liang,
Sanjeet Jain, Kailas Krivanka

Abstract

The increased usage and integration of the internet for all tasks large and small has led to an increased concern over security. A Virtual Private Network (VPN) provides a secure connection for users to help address these concerns. Wireguard is a new VPN protocol that promises a smaller codebase and stronger performance. We have used this VPN protocol to develop our own connection manager microservice that will route users to servers and provide each user an isolated namespace. This ensures the security of a VPN connection while making configurations simple for customers.

Overview

Cisco Umbrella is the division of Cisco focusing on cloud security. Their domain encompases many products such as DNS-layer security, and a cloud delivered firewall. One of the products they provide is a cloud-based VPN service.

A VPN is a common network security construct that protects a user's traffic as it travels across the public internet to its destination. It acts as a trusted intermediary between a computer and the internet and securely encrypts all traffic sent between the user and the VPN endpoint. This prevents attacks such as traffic snooping and interception because all the user's traffic appears to be sent only to the VPN. Its actual destination or content cannot be determined.

We investigated Wireguard on behalf of Cisco Umbrella which hopes to develop a new VPN service based on this new protocol. We created a VPN connection manager that uses this new protocol as a prototype of this new architecture.

Implementations

The main implementations that we focused on were Wireguard Go, Wireguard Rust, Wireguard Linux, and Boringtun. After extensively testing each implementation, we eventually decided to go with Wireguard Linux because it was the easiest and the most convenient to use. Also, each of the other implementations came with some issues.

Wireguard Rust implementation was incomplete and failed to build multiple times while running Wireguard Go implementation had issues spinning up and running a connection.

Boringtun implementation was implemented successfully on the client. However, we chose Wireguard Linux due to our familiarity with it after using it for the VPS component

Approach

The **connection manager** is a intermediary that passes the client's information to any of the Wireguard servers that it routes to. It then spins up a VPN interface that the client can use to establish a connection.

The Connection Manager receives information necessary to verify a user through a POST request made by a client. This information is then processed for any server permissions. If the user has no permissions, the Connection Manager will route to a generic server list using an algorithm to choose. Once a server is chosen, the manager generates a per session Wireguard key and sends a POST request to the server.

The **server** creates a new interface for the user. The connection provides an isolated environment for the user on the server using Docker containers. These containers are all run on separate ports so that multiple users may connect to a single server. Each server has a teardown protocol that will remove any interfaces that are no longer in use by a client.

The **client** service can configure an interface using the response returned by the Connection Manager and make a connection to the server. This process streamlines a connection using only one command.

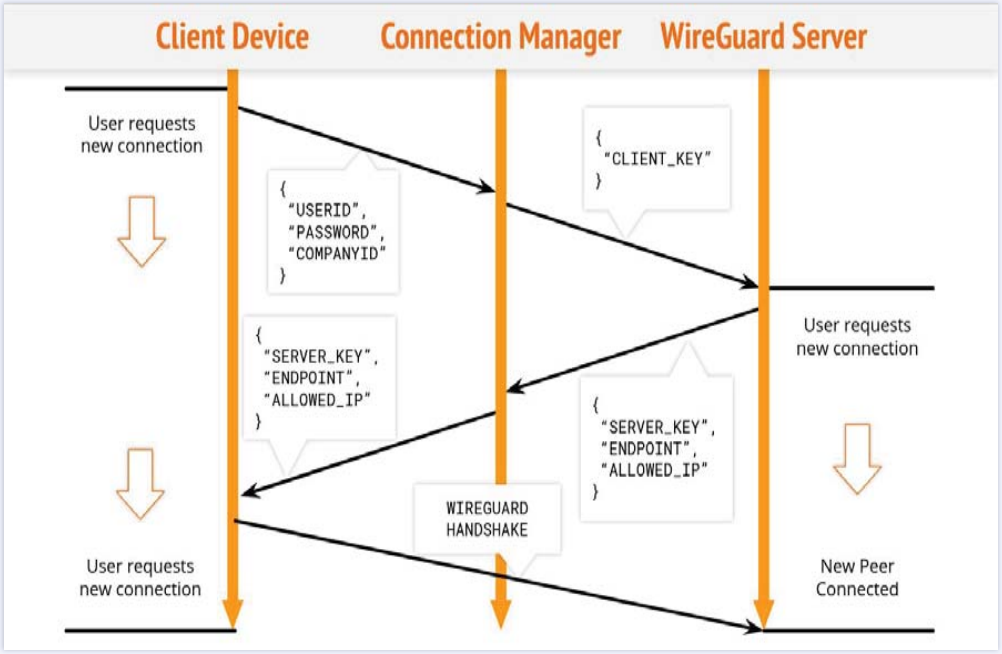

Results

We were able to successfully create a VPN connection manager for multiple simultaneous users that can route to multiple servers. This was tested with up to 3 servers. Users are authenticated with a username and password, have persistence in our database, and can use a single command to initiate a new Wireguard connection.

Moving Forward

We realized during the project that the architecture we created cannot run easily on commercial cloud platforms. This is because the VPS module requires kernel level access to networking APIs that are not always available in virtualized environments like Google Compute Engine.

Given more time to work on the project we would like to condense and combine some of our components and redesign the VPS service so that it could be deployed more easily.

Special Thanks

Cisco Umbrella Team: Kyle Mestery, Greg Duraj, Adrian Oliver, Bryan Hong, Leo R. Augusto
CSE 115 Faculty: Richard Jullig, Arindam Sarma

UC SANTA CRUZ
Baskin Engineering

Capstone Project

Data Studio Performance Tuning

Jinxiao (Jason) Song, William Kudsk, Michael Zhang,
Chuangbo Tong, Haofan Wang

ORACLE Labs

Abstract

The goal of this project is to take Data Studio, a web-based notebook platform currently in development, and create a performance testing suite that can identify bottlenecks in queries to its database. Using Java Microbenchmarking Harness, we created a testing suite that extracts information related to how much time specific queries take to complete. This information is used to identify performance bottlenecks, and to guide development of future optimizations.

Approach

- Understand Oracle Data Studio's system and data structure
  - Set up Oracle Data Studio in the working environment with **Ubuntu**[1]
  - Study **Spring Boot**[2], JMH, and Hibernate
  - Use Hibernate as the object relational mapping tool between Data Studio objects and our simplified version of Oracle database.
- Integrate benchmarking tool: **JMH**[3]
  - Add JMH dependency in **Gradle**[4] and create custom Gradle task
  - Import JMH tools into benchmark testing suite
  - Use these tools to benchmark the database operations
- Develop benchmark testing suite
  - Write tests for database Create, Read, Update, and Read operations
  - Use Benchmark annotation from JMH to observe result.
  - Create tests for custom database functions made by Oracle developers
- Use the result of the benchmark to find bottlenecks

Overview

Web-based notebooks are word processing web-applications which can compile code to display numerical and graphical output using various programming languages. Data Studio is a web-based notebook platform that can help with data visualization for data scientists, who deal with many data daily. Because of this, Data Studio has to handle a lot of data, as users can create and share notebooks between each other. As Data Studio is still in development, Oracle developers do not have insight into the performance of database components, as they lack the right tool to do so. Our goal involved building a testing suite that could give us valuable performance related data, and to then identify potential bottlenecks to provide possible optimizations.

Architecture

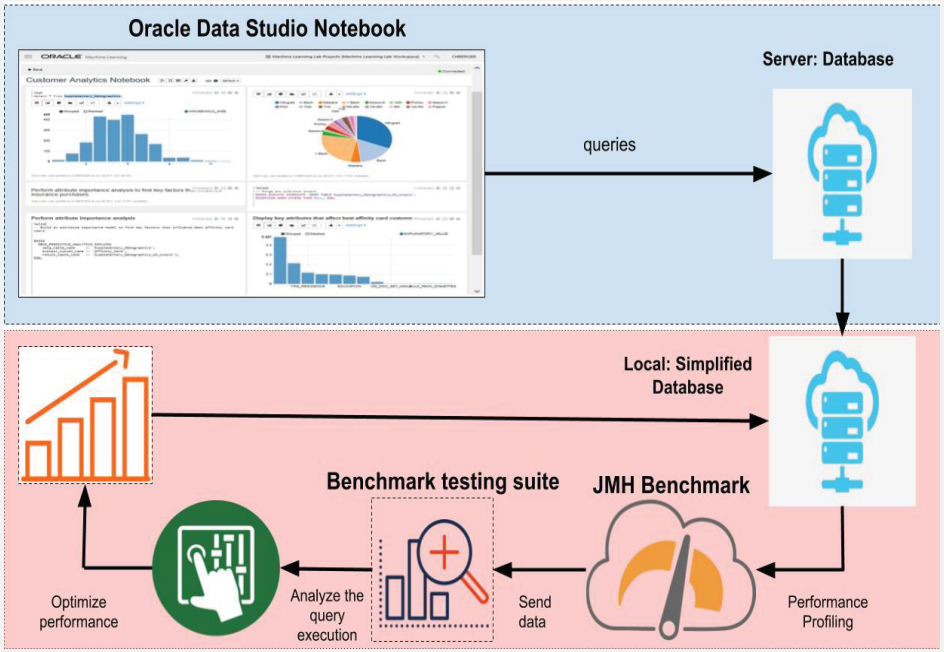

Technology used:

1. **Ubuntu:** an open source linux system
2. **Spring Boot:** an open source Java-based framework that provide Java Persistence API
3. **JMH**: a lightweight java benchmarking tool
4. **Gradle**: an open-source build automation tool

Results

We benchmarked database queries for 6 different Data Studio entities:

- Notebook, Job, Comment, User, PermissionMapping and Permission
- For each entity, we benchmarked create, read, update, and delete operations.
- Additionally benchmarked custom operation specific to certain entities.
- **Example: Notebook entity**

  Benchmarked functions involving the object's relationship with Users, Permissions, and filterNotebook with user tag.
- Used SQL scripts and different batch sizes of entities saved to the local database to best emulate performance in of data studio in practice
- Investigated ways to achieve optimizations especially for the costly findAll queries

Conclusion

With this project we created a benchmarking platform used to optimize database transactions in Data Studio. This platform is currently being integrated with Data Studio to give Oracle developers a tool for detecting bottlenecks and improving performance.

Acknowledgments

Thanks to:

Daniel Langerenken	Mentor
Dr. Richard Jullig	Instructor
Faeze Brahman	TA
Chandranil Chakrabortti	TA

Capstone Project

A Service to Parquet

Phu Le, Xiaobin Wu, Yibo Guo

Abstract

This project is intended to provide a cloud service that provides a variety of tools that can aid Oracle clients in storing and processing their data. Switching to Parquet is made easier with the file conversion library and data processing tools like filtering, searching, and sorting can transform that same data and store it on the cloud. Users can access this service on the Oracle cloud and will easily be able to streamline their data analytics process.

Approach

Service: The core of the service provides a way for a client to communicate with the server. This was accomplished with Helidon, an Oracle open-sourced library.

Library: After communicating with the server, the Helidon service calls on the library portion that provides all the functionality, including data-processing operations and cloud access. This required using the Apache Parquet library and the Oracle Cloud Infastructure SDK.

Overview

As data sets increase, storage and data analytics require optimizations. When Oracle customers store data, analyzing the data involves extracting, transforming, and loading the data, called the ETL process. While this process is costly and inefficient, Oracle's solution revolves around a file format called Apache Parquet. Designed as a columnar storage, Parquet files can offer efficient and complex data processing. However, current solutions handling Parquet are not-so-lightweight. Oracle hopes to create a cloud service for Parquet that is both lightweight and easy-to-use in order to streamline the ETL process.

Dataset	Size on Amazon S3	Query Run time	Data Scanned	Cost
Data stored as CSV files	1 TB	236 seconds	1.15 TB	$5.75
Data stored in Apache Parquet format*	130 GB	6.78 seconds	2.51 GB	$0.01
Savings / Speedup	87% less with Parquet	34x faster	99% less data scanned	99.7% savings

Architecture

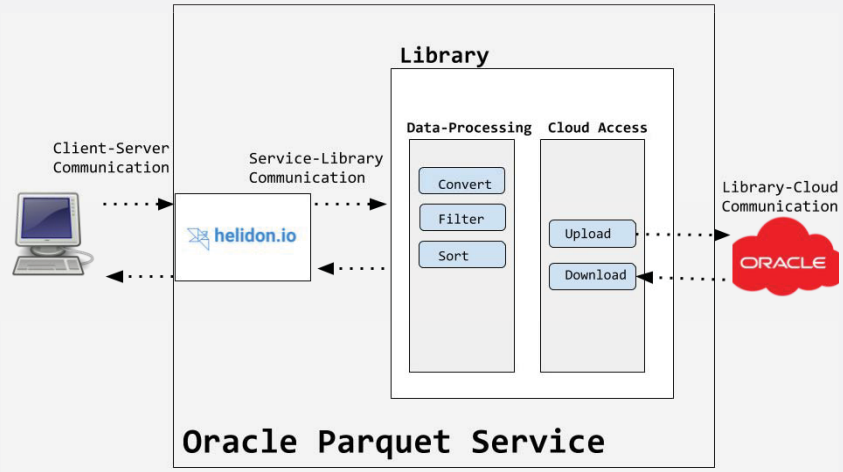

Acknowledgments

Daniel Langerenken – Vishal Vaddadhi – David Hernandez

Results

Clients who use our service have access to a variety of tools, including file conversion, row filtering, column filtering, and cloud access. Compared to current solutions, our service makes data processing simpler. With our time constraints, we used Apache Drill as a shortcut to implement some functionality. Additionally, we did not reach our intial goal of deploying our service on the cloud.

Benchmark

Benchmarking our native solution with the Drill implementation, we found that our native solution performed better than Drill's. We also see that that improved performance is more noticeable for smaller files.

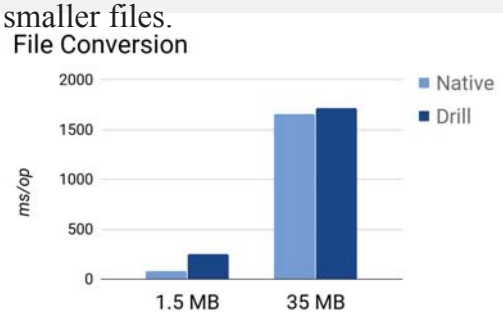

Accelerating Key In-memory Database Functionality with FPGA Technology

Jordan Leggett, Andrew Hartwell, Akihiro Mizusawa,
Robert Mushkot, Vadim Pelyushenko

Abstract

In-memory databases supporting online transactional and analytical workloads rely on *differential updates* to manage writes while maintaining the majority of data in an optimized read-only structure. This necessitates the merge of a write optimized delta partition with the read-only main partition. The purpose of this project is to develop an FPGA using OpenCL that can reduce this bottleneck on the database.

Overview

SAP uses their **H**igh-Performance **An**alytic **A**pplication (*HANA*) In-Memory Database to quickly provide data to over 400,000 customers using their numerous cloud services. *HANA* requires the use of a *delta merge* operation to merge recently written data from a delta storage partition into a main storage partition. The delta merge operation requires extra memory usage and longer data access times, which is expensive for both SAP and any clients using HANA. With OpenCL, we can program an FPGA to build custom hardware tailored around data and pipeline parallelism and optimal performance of the delta merge operation. SAP's previous FPGA research without OpenCL achieved a bandwidth of 5-10 GB/s, however that prototype is outdated and the technology used is no longer supported.

Technologies

- **FPGA:** A reprogrammable hardware device that allows flexible, reconfigurable computing and rapid prototyping of hardware designs. FPGAs are capable of data/pipeline/task parallelism.
- **OpenCL:** A programming language designed for parallel computing on hardware accelerators. OpenCL provides a framework for software developers who want to work with hardware.

Approach

The diagram to the right represents the structure of the HANA database, with the delta merge operation split into 4 OpenCL kernels each running simultaneously :

Delta Store: Keeps track of database writes and uses dictionary compression (D_D) to store data in the delta column (C_D). This structure is optimized for transactional processing.

Main Store: A read-only partition that uses a sorted dictionary (M_D) for data compression in the main data column (C_M). This structure is optimized for analytical processing.

Merge: The FPGA recodes the main store column Value IDs to map to new values based on a provided dictionary while utilizing pipeline parallelism.

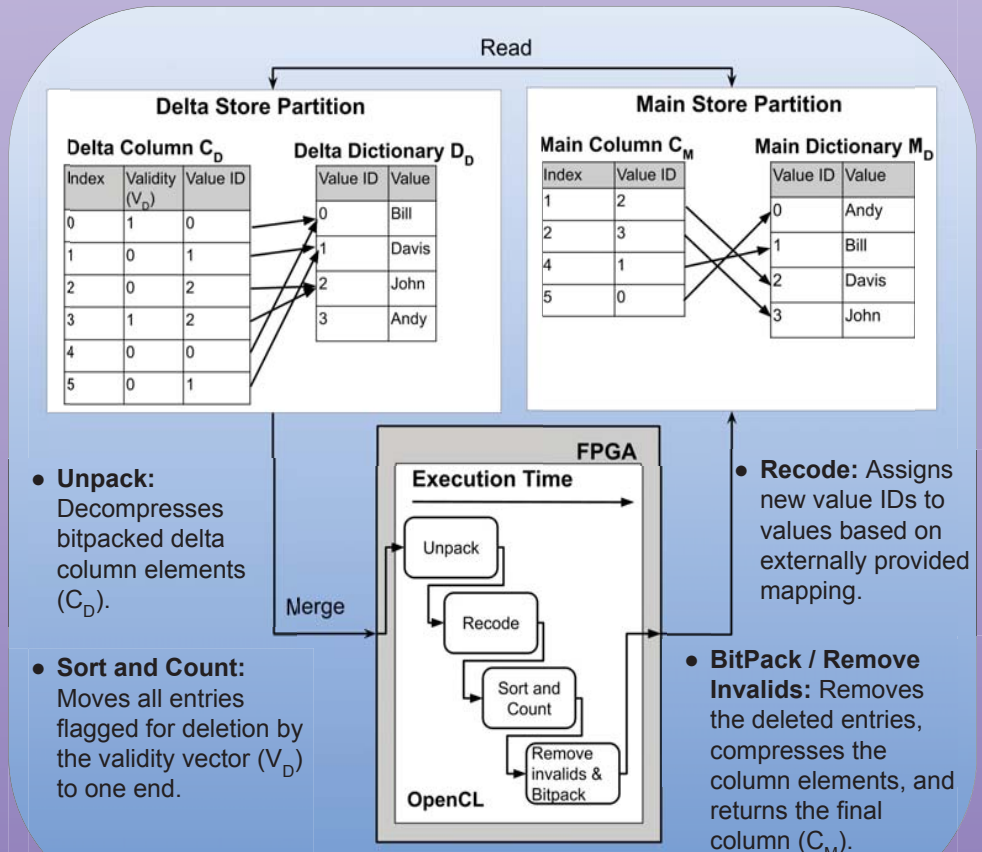

- **Unpack:** Decompresses bitpacked delta column elements (C_D).
- **Sort and Count:** Moves all entries flagged for deletion by the validity vector (V_D) to one end.
- **Recode:** Assigns new value IDs to values based on externally provided mapping.
- **BitPack / Remove Invalids:** Removes the deleted entries, compresses the column elements, and returns the final column (C_M).

Results

- Delta merge operation in C++.
- Initial OpenCL FPGA design with bandwidth of 48 MB/s.
- Used loop optimization techniques and utilized channels for data pipelining and fast kernel communication.
- Several FPGA designs with OpenCL and achieved bandwidth of up to approximately 2.9 GB/s.

Conclusion

The delta merge is a core function of SAP's HANA database, therefore efficiency is critical. We have identified several sections with optimization potential which could increase our throughput in hopes of surpassing the 5-10 GB/s achieved by SAP's previous research. Our project will serve as a prototype for a future version implemented into SAP's HANA.

Acknowledgments

We would like to thank the following for all of their help and guidance throughout this project: John McGlone (SAP), Jody Glider (SAP), Suleyman Demirsoy (Intel), Christian Faerber (Intel), Ryan Tekerlek (Intel), Dan Nguyen (Intel), Akila de Silva (UCSC), and Professor Richard Jullig (UCSC).

UC SANTA CRUZ
Baskin Engineering

Capstone Project

Classmates Elasticsearch

Kristopher Rollert, Kai Schniedergers,
Michelle Slaughter, Chuanshi Zhu

classmates

Abstract

The core of Classmates' web application functionality is reliant on their search engine that queries through Classmates' tens of millions of student records. Classmates has been using Apache Solr for its search engine since 2012 and the engineers feel that it is time to upgrade to something more modern. Solr has been causing performance issues as well as limiting the functionality of Classmates' search engine. With this project, we prototyped Elasticsearch's more modern functionality to reduce query latency, improve result accuray, and set up new features like search completion.

Approach

Classmates had no prior tools for Elasticsearch so for our conversion, we were given csv files with data and built the rest from scratch. We:

- Converted and merged ~100 GB of csv files into compact json files using Apache Spark
- Used Docker to stand up multiple "nodes" of Elasticsearch on an Amazon EC2 instance
- Created Python tools to send data to Elasticsearch containers so it can be indexed
- Built our one index and three index approaches and analyzed the speed, efficiency, and accuracy to figure out which would be optimal
- Created a query autocomplete feature using a neural network and log files of previous searches

Overview

Classmates has 207 GB of data on old classmates, 371 MB of data on schools and 287 MB of data on old yearbooks, which is searched over 100k times a day.

After years of using Apache Solr, the engineers at Classmates decided it was time to upgrade their system to Elasticsearch (ES). ES provides distributed queries, better functionality around filtering, and many more features that fit better with Classmates' system design. Our team was in charge of switching systems. This required moving over the Solr data, comparing Elasticsearch versus its wrapper software: Appsearch, and configuring the final setup.

Architecture

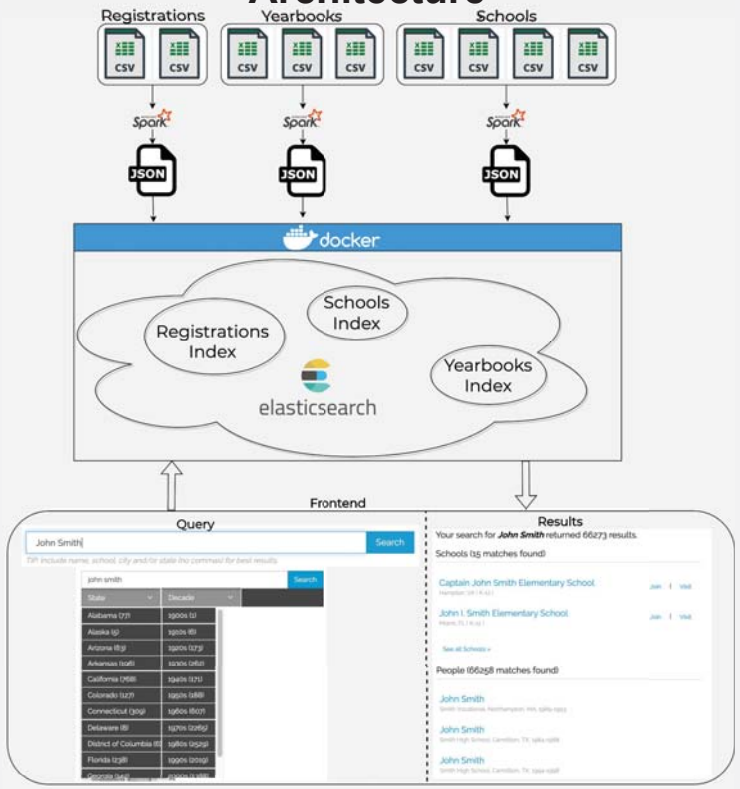

Acknowledgments

Thank you to our sponsors and mentors for all your help!
Sponsors (Classmates): Yalin Yesiltas & Payal Patel
Mentors (UCSC): Richard Jullig & Akila de Silva

Evaluation

There are two options for Elasticsearch. The one index option is to have schools, registrations,and yearbooks in one big index for searching. The three indices approach is to have an index for each section. In order to choose whether we use one index or three indices, we had eight different sections to evaluate their performance. They are name search, facets, auto-suggest, index size, spell correction, average query response time, development effort and relevance. We wrote our own python scripts to calculate scores for many of the sections.

Results

Result Highlights	One Index	Three Indices
Index Size	171,908 docs (2.7 GBs)	1,678,792 docs (830.5 MBs)
Avg Query Time	9.5 ms	5.0 ms
Relevance	13.78	12.99

In most of the data we compared, the three indices approach performed significantly better. Though the relevance score was lower, we found that the actual results were more relevant.

Conclusion

With the tools we've created, engineers will be able to seamlessly convert their search engine from an outdated version of Apache Solr to Elasticsearch. Our tools will allow engineers at Classmates to stand up Elasticsearch containers with Docker, as well as port and upload data from their existing Solr data. Our tools will help modernize one of the most vital functions of the Classmates.com website: searching.

Capstone Project

UC SANTA CRUZ Baskin Engineering

Learning Storage Networks

Hewlett Packard Enterprise

Paripinya Srijunyanont	John Hornbuckle	Joseph Rodrigues
Daniel Panasyu	Rohith Bollineni	Rahul Mahendru

Abstract

Learning Storage Networks is enterprise software intended to be used by HPE and their customers to collect data and analyze networks automatically, without need for manual input or intervention until actual erroneous behavior is detected. This will reduce both time and money spent to deal with data loss and hardware damage.

Approach

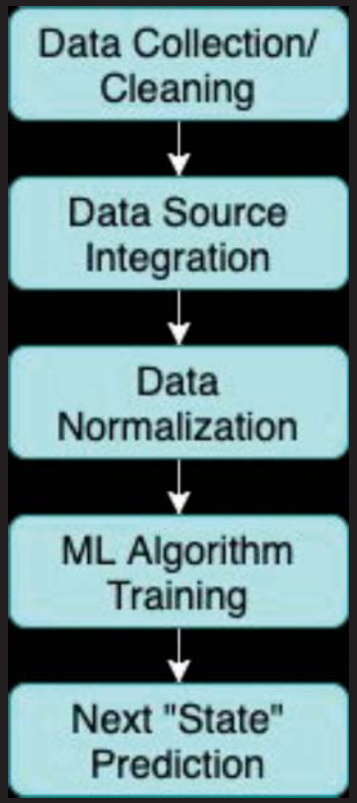

Overview

Hewlett Packard Enterprise(HPE) offers storage services to other corporations to fill today's massive demand of data storage spaces. Specifically, they offer storage area networks (SANs) to allow for large scale storage, tracking, and analysis of data. These networks are made up of interconnected Fibre Channel switches, that deteriorate over time and are sometimes manufactured defective. The magnitude of these networks prevents easy manual error detection, and thus, there is much time and money to be saved for companies like HPE with the use of machine learning to predict switch anomalies before they happen.

Storage Area Network

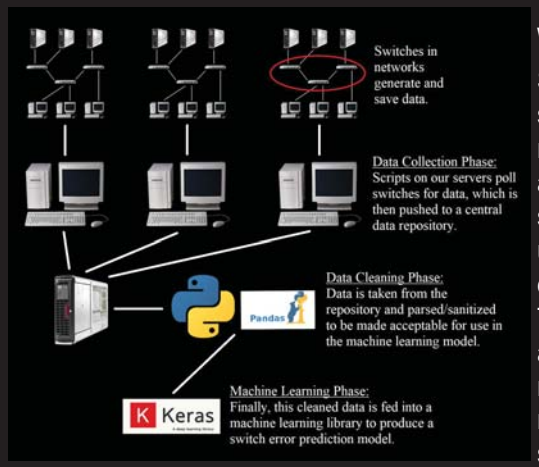

What are Switches?

Switches connect storage, servers, and other switches in a network called a fabric. Our architecture is comprised of a script running on a server that uses a single switch to collect data from fabrics of switches. This collected data is then analyzed via a machine learning model which predicts whether or not a switch will enter an error state.

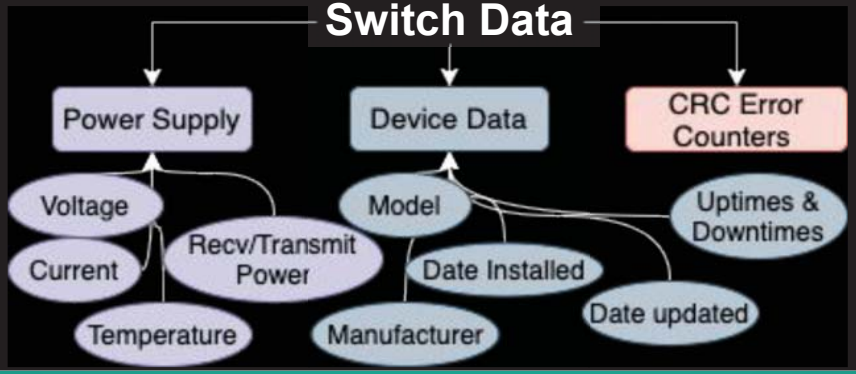

Results

- Automated data collection of Brocade switch data.
- Trained an ML model that is able to predict 15 future errors with 55% accuracy.
- Can notify network operators of possible risks before they occur.
- Prevented errors could potentially allow for a significant amount of costs to be avoided.
- Final planned deliverable: Software that runs on HPE servers, occasionally querying switches for data related to power supply, uptime, and CRC errors with the ability to email warnings to network operators when it is probable that the switch is going to malfunction.

Acknowledgments

We would like to extend a special thanks to Ayman Abouelwafa and Salil Gokhale for facilitating the process of communication between HPE and our team, as well as providing technical assistance whenever necessary. We deeply appreciate the care and effort which Professor Richard Jullig as well as our TAs Chandranil Chakrabortii and Golam Muktadir provided us to support the entire process of this project.

Anthem

Rob Currie
Paula Alves
Robert Baertsch

Model Team
Harshitha Arul Murugan | Aman Prasad

Understanding Healthcare Data

Data Team
Cassidy Norfleet | Brendan Reilly-Langer

Baskin Engineering

Richard Jullig
Scott Davis
Golam Muktadir

Abstract

Compared to other wealthy countries, the U.S spends a disproportionate amount on **healthcare** with the **gap widening every year**.

We **utilized machine learning tools** to create models that predict a patient's future health conditions given their previous medical history.

Anthem, one of the largest healthcare companies in the U.S, is exploring the use of these models in order to provide **better care** for the patient and at a **reduced cost** for the care provider.

Key Terms

Synthea: Creates statistically accurate synthetic medical records curated by doctors for the purpose of research.

One-Hot Encoding (OHE): In a month, if a medical code was in a patient's history, the entry is 1 or 0 otherwise.

Docker: a platform to package up applications and code for the purpose of reliably running on different machines

Python libraries

Numpy: Contains numerical tools for computations.

Pandas: Helps organize data and perform calculations.

Pytorch: Creates, trains, and tests machine learning models.

Overview

The **goal** of our project is to use machine learning to **predict** the factors that lead to the **diagnosis** of a patient with **congestive heart failure. Anthem** can use these **predictions** to preemptively form a **treatment plan** and inhibit the progression of a patient's disease.

Machine Learning Terminology

Recurrent Neural Network (RNN)

Uses predictions from previous inputs to influence the current input's prediction.

Convolutional Neural Network (CNN)

An image classification technique used to find unique patterns in different images.

Word Embeddings (Word2Vec)

Vectors of numerical values assigned to words or individual components in a given text.

Architecture

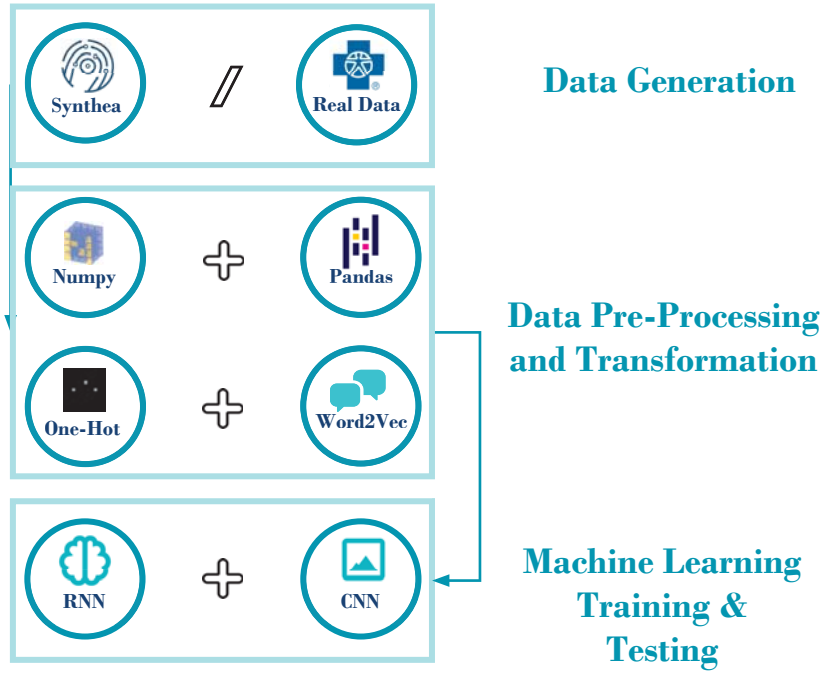

Approach

Because access to actual patient data is highly restricted, we used **Synthea** to generate artificial medical records. These records are used to create sparse **One-Hot Encoded** matrices **(OHE)** where each row represents one month of a patient's medical history and each column represents a unique medical code.

Using **OHE** matrices we trained a **CNN** and an **RNN**. We also created **word embeddings** for each of the medical codes present in a patient's history to train a **CNN**.

After developing successful models, we packaged our models in a **docker** container and sent it to **Anthem** to train on real anonymized patient data.

Results & Conclusion

For the **OHE** matrices built from the synthetic data, we were able to create a **CNN** that predicted outcomes at an accuracy of **96.36%**, and an **RNN** that predicted at an accuracy of **82.61%**.

In addition to OHE matrices, we also used the word embedding approach, which resulted in an accuracy of **79.34%** for the **CNN** and **71.53%** for the **RNN**.

Anthem then trained our models on the **word embeddings** processed from real data, which achieved an accuracy of **86.33%** on the **CNN** model and **83.72%** on **RNN** model.

UC SANTA CRUZ
Baskin Engineering

Painless Healthcare

Emerson Christie, Yousef Dost, Neel Apte, Victor Ye

Accolade

Abstract

The Consumer Health App allows users to connect a variety of different data sources to their Accolade health profile. This data is collated, stored in the cloud, and then visualized on our web application, giving the user instant access to all of their healthcare data such as their biometric and activity data in real time.

Approach

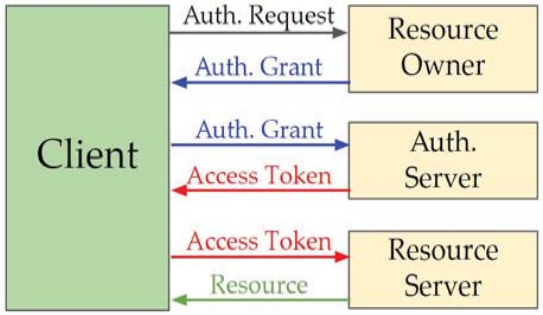

Consumer Health Client: From the front end, the user connects with third party data sources through an **OAuth** protocol flow (as seen above), in which the user is prompted to enter their login information.

Cookies: Once authenticated, an access token is returned for us to perform API requests. A refresh token is also returned and stored in a cookie for future polling.

Lib Modules: All pulled data point values are then parsed and formatted into FHIR compliant data structures via predefined mappings of data values and FHIR attributes.

AWS Cloud Services: The parsed data is then stored in DynamoDB and pushed to the **ACP Service**, which is Accolade's API for updating and accessing a patient's profile.

Overview

Accolade is a company that provides personalized health-and-benefits solution to improve the experience, outcomes and cost of healthcare for employers, health plans and their members. Many people have an abundance of healthcare data being collected from IoT (Internet of Things) devices like heart rate and blood sensors, as well as from doctor encounters. This data contains key insights into an individual's overall health that could be used for recommending doctors and predicting health issues. While this data is valuable many people either don't know it exists or don't put it to use. The consumer health app solves this problem by collecting healthcare data from a variety of data sources, making it available to the patient and Accolade; which designs machine learning algorithms aimed at finding solutions to some of healthcare's biggest questions.

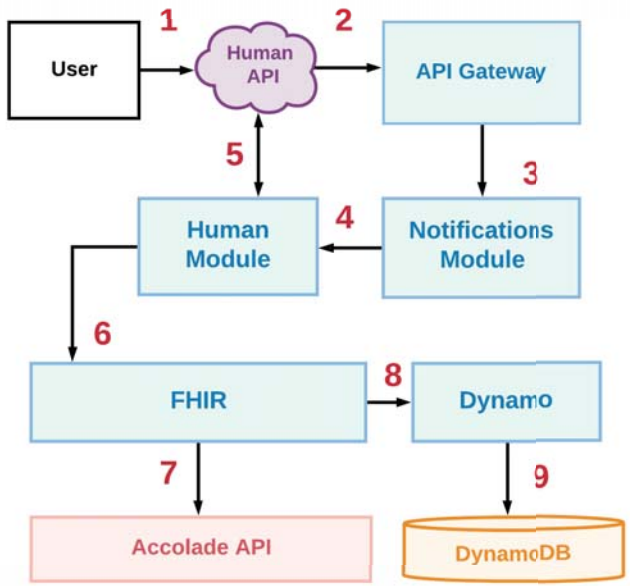

Acknowledgments

We are deeply grateful to our TA Arindam Sarma and Professor Jullig for their encouragement and support, along with our Accolade sponsor Gregory Ayre for his guidance and direction.

Methodology

Accolade's existing database utilizes the Fast Healthcare Interoperability Resources (**FHIR**) specification, which is an accepted industry standard data format for electronic health records. We created adapter classes for each data source which parsed and mapped the data to FHIR structures as seen below.


```
"blood_glucose": {
  "category": [...],
  "code": {...},
  "effectiveDateTime": "2014-01-23T22:48:19.858Z",
  "meta": {...},
  "performer": [...],
  "resourceType": "Observation",
  "status": "final",
  "subject": {...},
  "valueQuantity": {
    "code": "mg/dL",
    "system": "http://unitsofmeasure.org",
    "unit": "milligram per deciliter",
    "value": 90
  }
}
```


Results

We integrated Google Fit and Human API endpoints with Consumer Health App so that they automatically poll, parse, and push the most up-to-date data. Furthermore, we were able to create a conclusive health data platform that provides the user an effortless and painless way to view their health data.

Conclusion

In future implementations, we want to include more API integrations—pharmacies such as RiteAid, Walgreens, and CVS for medication data, and Teledoc for scheduling doctor appointments. We also want to implement a module that will generalize and streamline the flow for third party OAuth integration and data transduction.

UC SANTA CRUZ
Baskin Engineering

Capstone Project

VoIP Case Management

By Andrew Leamy, Allan Wong,
and Tarun Sivakumar

DATACARE

Abstract

VoIP Case Management is a web application designed to optimize workflow for managing insurance cases. It does this by neatly packaging information on said cases into organized profiles, which are stored in the host server using a **Relational Database**.

The application's primary features revolve around making VoIP (Voice over Internet Protocol) calls and routing a phone call to the client's internet browser. This is accomplished using the API libraries provided by the third party service **Twilio**.

Additionally, this application features **Case Prediction**, in which the server can predict which case an incoming call is associated with. The predicted case profile will then be automatically pulled up for the case manager, preventing the need for him or her to look up necessary information themselves.

Approach

The host server for this application runs in what is known as the **LAMP stack** (Shorthand for Linux, Apache, MySQL, and PHP). The process by which a call is handled - illustrated in the center diagram - can be summarized below:

1. A call is made to a phone number registered with Twilio.
2. Twilio requests the host server for instructions on how to respond.
3. The host server runs Case Prediction to find which case the call is associated with.
4. This information is displayed to the user.
5. The user accepts or declines the call
6. The server responds to Twilio, sending instructions to either connect the call or record a voicemail.
7. Twilio executes the given instructions.

Overview

The management of workers' insurance cases is a complex mire of bureaucracy and communication between many different professions. **Case Nurses** are usually put in charge of such cases, often managing multiple cases at once. How can we help improve the workflow of these nurses?

DataCare is the leading provider for software in the workers' insurance industry. Their mission is to provide innovative software solutions that improve the efficiency of patient care and recovery. DataCare has succeeded at lowering service costs for Fortune 500 companies, third party administrators, and insurance groups alike along this process.

Our software, **VoIP Case Management**, aims to be a proof of concept for the next level of workflow optimization for DataCare's products; that is, the integration of VoIP calls between the client browser and cell phones, along with other features such as case prediction and voice transcription.

Architecture

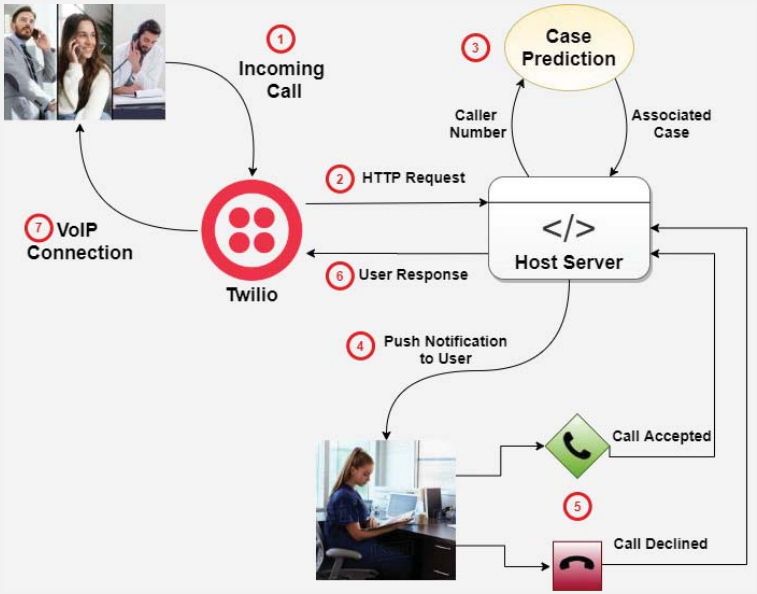

Acknowledgments

The team would like to thank the following for providing feedback during the development of the project:

- Paulo Franca, CTO of DataCare
- Chris Van, Engineer at Datacare
- Golam Muktadir, class TA
- Professor Jullig

Key Features

Case prediction is performed first by matching the caller number with a contact stored in the database. The host server then scans through the database, selecting the most relevant case based on the caller's own association with it, and how recently the case has been interacted with.

A **push notification** is sent to the client browser, which pulls up and displays quick reference data of the associated case to the user. The user - the Case Nurse managing this case - chooses to accept or decline.

Should the user decline, the server instructs Twilio to record a **voicemail**, and create a **transcription** of the voicemail. Both of these are downloaded to the host server, where the user can access and play them on command.

Results

Despite a large number of obstacles faced during development - primary among them the California stay at home order - the team has succeeded in integrating most of the features outlined by our sponsor. Namely: support for VoIP calls, support for voicemail and transcriptions, a functional relational database to store case data, and case prediction.

Conclusion

VoIP Case Management serves as a functional prototype and proof of concept. Features planned for potential new releases include call priority sorting, conference calls, voice-to-text notes transcription, as well as improvements to the case prediction algorithm.

UC SANTA CRUZ
Baskin Engineering
Computer Science Capstone

Conversational Chatbot-based Diabetes Assistant

Hoa Nguyen, Nicholas Henderson, Jade Bergamini, Zuzanna Dubanowska, Kevin Zhang, Eduardo Villa

glooko

Introduction

DiaBuddy is a proof of concept for an expansion to the Glooko® app. We want to make it easier to manage and learn about diabetes.

Oh, what's Glooko®? And why are you doing this?

With Glooko, you can sync, share and review data from a majority of available diabetes devices. Data is displayed in easy to understand charts and graphs. Dia-Buddy lets you access some of these features in an even easier way, through a simple voice interface. It also provides educational resources to help utilise this information.

Huh… how'd you do that?

It's all in a chatbot. You simply ask it to log your blood sugar levels, or ask it for highs or lows in your readings, and it will do that! It's like talking to your Siri or Alexa.

You can also ask questions about diabetes, and it will respond with answers curated from online sources*. It even recommends articles based on your frequently asked questions.

*Glooko® and DiaBuddy are not diagnostic or treatment substitutes. Please consult your provider or doctor for proper medical advice. All sources used in answers were found through online credible sources.

How does it work?

A user can type or use voice commands and the bot will respond accordingly. It supports Q&A for most popular questions about diabetes (50+ distinct questions), and voice operated data management (for logging and reading a user's data).

"What's a normal blood glucose level for women?"

"What was my highest blood glucose level from last week?"

"Recommend me an article on diets."

Additionally, the app features a newsfeed-like educational tab that provides the user with articles. The content of the tab is dynamically updated based on which topics the user asks the chatbot about most frequently.

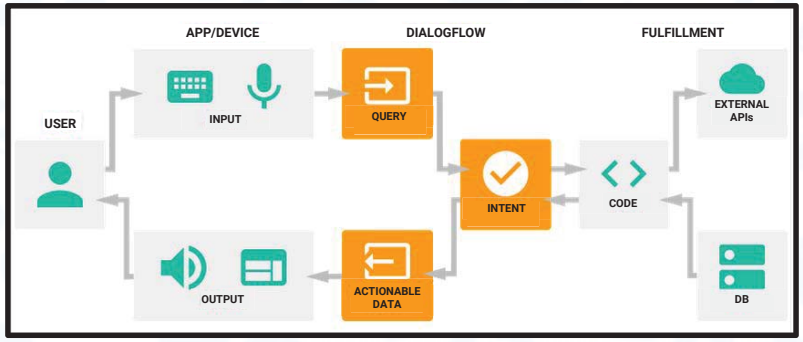

Results

The voice interface allows you to log your blood glucose levels, which previously took 4 clicks, with one sentence. Users can also avoid false information found online by utilizing the EDU tab to have access to reliable articles.

This chatbot prototype provides an excellent foundation for the development of more complex features, due to its easy to use nature and its friendly UI.

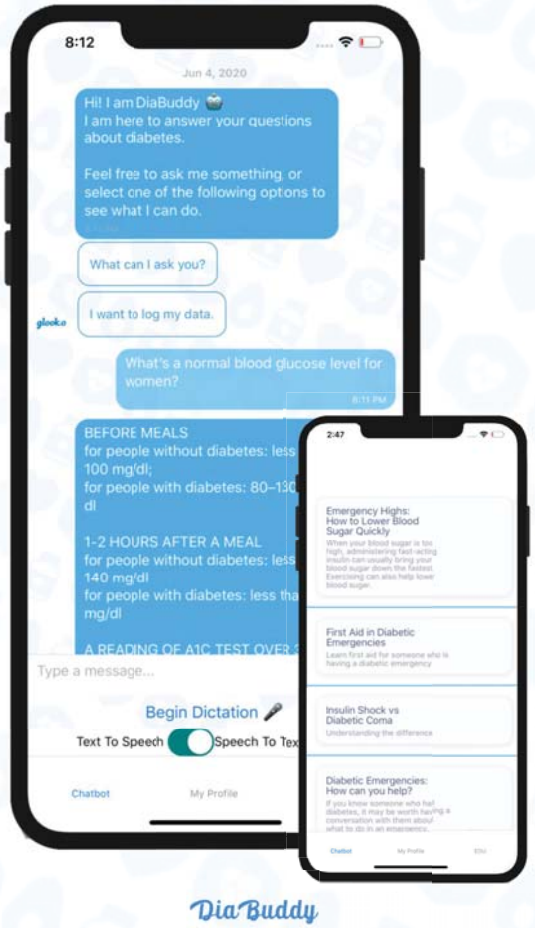

DiaBuddy

Acknowledgments

None of this would have been possible without the help of our generous sponsors at Glooko®: Kathleen McGarraugh, Veera Mukkanagoudar, Jeff Chang, and Pushkar Singh.

Additionally, we want to thank Akila, Arindam, Professor Patrick Mantey, and Instructor Richard Jullig for guiding us through this.

Capstone Project

Alexa-Enabled Universal Remote

Nixon Duong, Nikhil Punathil, Peter Eskraus

Abstract

The Alexa-Enabled Universal Remote leverages existing accessibility options within Amazon's Alexa voice assistant and pairs it with the hardware flexibility of a Raspberry Pi to control infrared (IR) and radio frequency (RF) enabled devices. The universal remote learns the signals of existing remotes and sends these signals to control devices hands-free by voice.

Approach

We use the Alexa Skills Kit SDK to build a custom skill. For this, we create a set of **intents** or actions that a user may perform using the skill, along with a set of **utterances** or words and phrases that may invoke said intents. In our case, we created "set up" intents for the training process and a "use" intent to transmit the recorded signals.

For hardware, we use a **Raspberry Pi 3 B+**, with an **IR/RF transceiver**, a module that allows us to transmit and receive IR signals, connected to the general-purpose I/O (**GPIO**) pins. The Linux program **LIRC** is used to send the appropriate IR signal when an intent is recognized.

Overview

A significant amount of people with disabilities use both regular and special-purpose devices that are controlled by handheld infrared(IR) and radio frequency(RF) remotes. Although these devices have a huge personal impact, they are restricted to those who could use handheld remotes. Voice is an empowering interface that could work around this problem. However, market solutions are low in volume and relatively expensive. The Alexa-enabled Universal Remote project aims to increase the number of devices that can be controlled by voice by recognizing the signals of existing remote controls. The cost-effectiveness of retrofitting existing devices, combined with the inexpensive hardware and the ease of not having to replace entire existing setups would make this product one of the easiest ways to enhance accessibility.

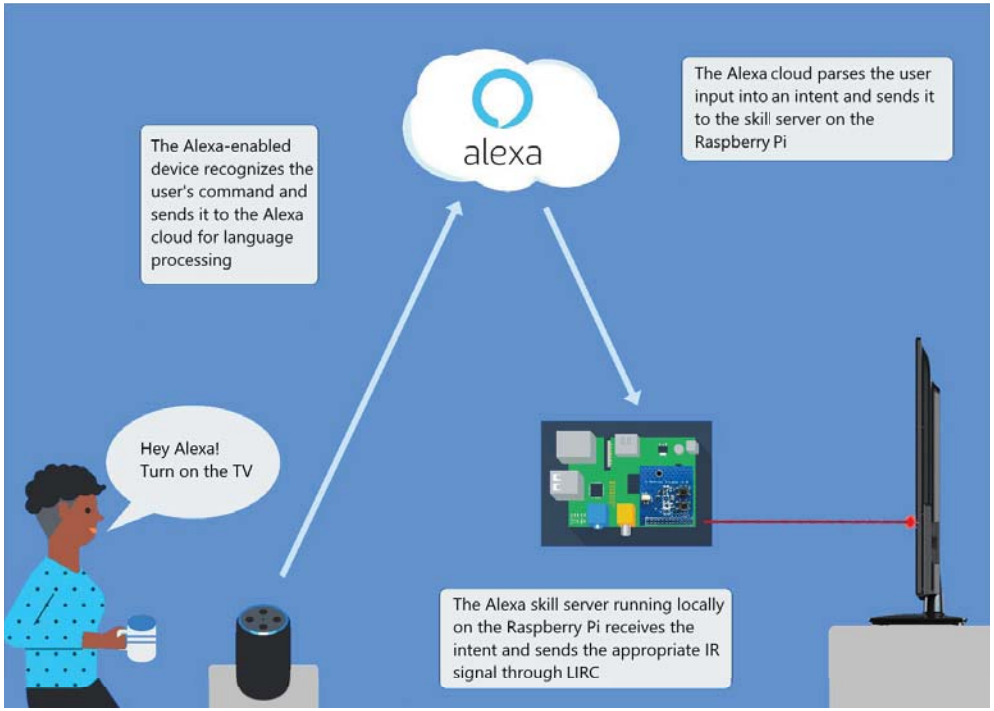

Training Remotes

The use of the **LIRC** program to parse and transmit IR signals came with the added advantage of easy access to their **community-sourced database** of remote configuration files that support an extensive number of devices.

The product also offers an **alternative interface** that allows the software to remember the signals of remotes through a **training routine**. While this does require the physical use of remote controls, which many among our target audience may not be able to do, it is a one-time process and is intended to be a **fallback** for rare situations.

Conclusion

The current version of our product allows users to record and use any remote control that uses infrared for communication. Future plans for the project include implementing a mesh network like system for the hardware, adding RF support and adding the ability to remember button sequences.

Acknowledgments

We would like to thank **Ken Karakotsios** and the other members of the **Alexa for Everyone team**, our instructor **Richard Jullig** and TA **Arindam Sarma**. We would also like to acknowledge our former teammates **Vincent Thai** and **Tongze Wang** for their initial contributions.

UC SANTA CRUZ
Baskin Engineering

Capstone Project

Alexa Modular Adapter

Anon Cheewakarn, Christopher Gunter, Kenneth Mai

amazon alexa

Abstract

The core feature of the Alexa Modular Adapter is to establish a connection between common switch control appliances and Alexa, a virtual assistant by Amazon, to improve the quality of life for those with physical disabilities. The installation process of the adapter has been voice-enabled, making the only physical interaction being the process of connecting wires to the adapter.

Approach

Software: The Alexa Voice Service(AVS) skills development kit is how Amazon devices process voice input. AVS is locally run on the Raspberry Pi (RPi). When given a voice command, the RPi sends that up the Alexa Cloud to process what has been said. Then, the cloud sends back what action the user has requested to smart home handlers. These handlers support features such as power, toggle, range, and mode for a device. Based on this directive, specific general purpose input/output(GPIO) pins are signaled.

For the modular adapter to engage in offline activities such as connecting to the internet via voice commands, the modular adapter uses an offline Speech-To-Text transcriber called DeepSpeech. DeepSpeech was developed by Mozilla to take in audio and transcribe it into text.

Hardware: As shown in the diagram, the toy is connected to an audio jack port, and the table is connected to an RJ45 port. To support these ports, a breakout module is connected to the RPi's GPIO pins. However, the pins on a RPi operate at 3.3V. To protect the RPi from any damage, relays are used between the pins and connected devices. Using relays, the RPi is able to control devices safely.

Furthermore, the adapter can detect whether a device has been plugged into the audio jack or ethernet(RJ45) port. This is done through GPIO pins that are pulled down and detect for any change(i.e. when a device has been connected).

Overview

The Alexa Modular Adapter strives to improve the quality of life for those with physical disabilities. This is done by voice-enabling appliances that are controlled with physical switches. Examples of supported devices include common household devices. The team decided to use a motorized desk and an adapted toy to prototype with. One aspect we emphasized was to make the setup process as hand-free as possible. To accomplish that, the adapter has a plug-and-play feature that allows users to setup and control their connected devices mostly with voice.

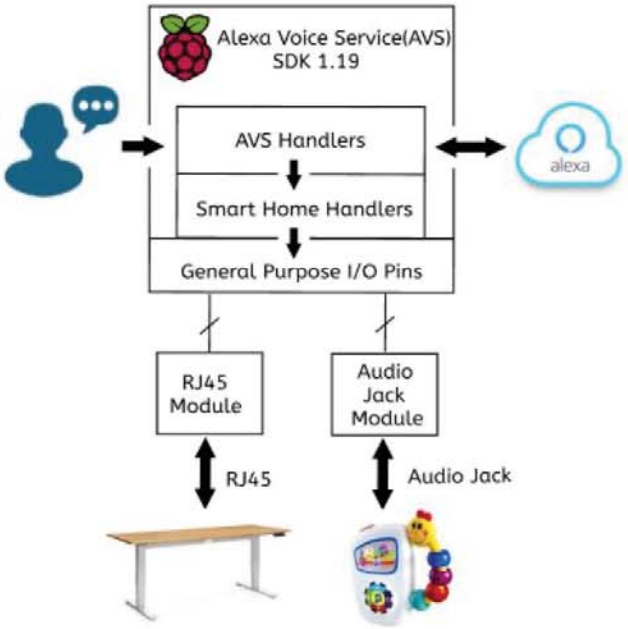

Acknowledgements

David Frerichs and the *Alexa for Everyone* Team

Instructor Richard Jullig and the Teaching Assistant Arindam Sarma

Former team members, Bryan Jimenez and Cagan Bakirci

Installation with Voice

In order to establish communication with the Alexa Voice Service, the Adapter must be connected to the internet. Our implementation allows the user to use voice commands only to have the adapter link to a Wi-Fi network. Using a speech-to-text library called DeepSpeech, the user prompts the product saying, "Find Wi-Fi Networks," the user then selects one of the listed Wi-Fi networks and then spells out the password to that Wi-Fi network.

Results

The Alexa Modular Adapter prototype can currently support appliances that either use an RJ45 or a 3.5mm Jack to control its components.

Furthermore, the adapter can simply be setup by connecting it to compatible appliances and use the voice agent to finish the process.

After installation, the adapter can control appliances using voice commands such as "Turn __ on/off", "Change __ to mode X", and "Set __ to N unit."

Conclusion

By enabling a voice-user interface on the motorized standing desk, the user can adjust its height to the one that suits the task they wish to perform without too much hassle.

The motorized desk and the musical toy demonstrate the feasibility of a voice-enabled adapter for devices with RJ45 or audio jacks. Another class in pretty close reach would be devices that use IR/RF controllers, which have been worked ony by the other Alexa Accessibility team.

CSE 123 A/B Senior Design Project: H2Eyes

Baskin Engineering

Corey Dickson, Linda Luu, James Melvin, Riley Silva, Venkat Vesta

UC SANTA CRUZ

Project Purpose

Need:
Visually impaired swimmer require assistance when swimming to assure their safety.

Goal:
Minimize the require supervision of visually impaired people to swim safely.

Aesthetic Kick-board Model

Target Market:

Name: Karen (Physical therapist)
Context: Plans to use the device to help visually impaired patients with orthopedic training

Name: Harry (Special Olympics athlete)
Context: Visually impaired athletic swimmer. Uses device to stay straight in his lane.

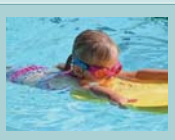

Name: Suzie
Context: Casual swimmer who is visually impaired. Uses device for improvement of aquatic exercise experience.

Lane Tracking

Standard lap swimming pool layout. The device will track the black lines shown running along the bottom of each lane.

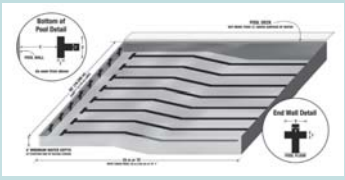

$$[L1_W, L2_W, g_W] = R_W^C * [L1_C, L2_C, g_C]$$

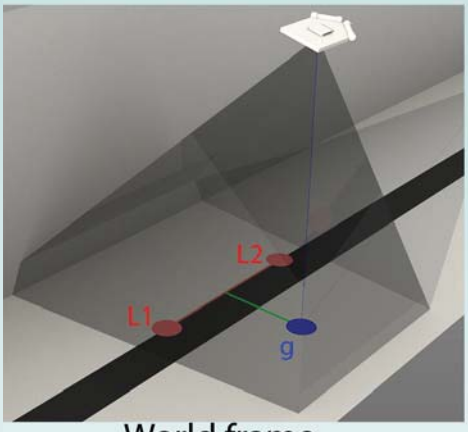

World frame

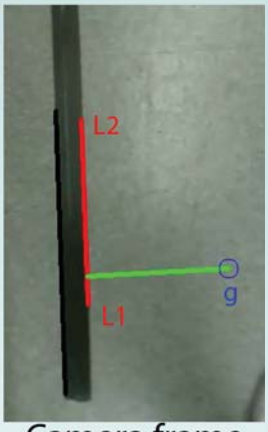

Camera frame

$$\text{Correction vector} = g_W - (\overline{g_W L1_W} \cdot \overline{L1_W L2_W})$$

Hardware

Our apparatus gathers data from the MPU9250 IMU and camera and then guide the user (swimmer) down the lane. If the user veers off course, this system will notify them what direction to turn via the vibration motors.

Preliminary prototype used to test sensors and algorithms used for line detection

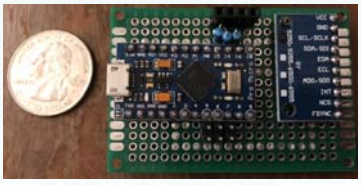

Arduino Pro Micro with IMU

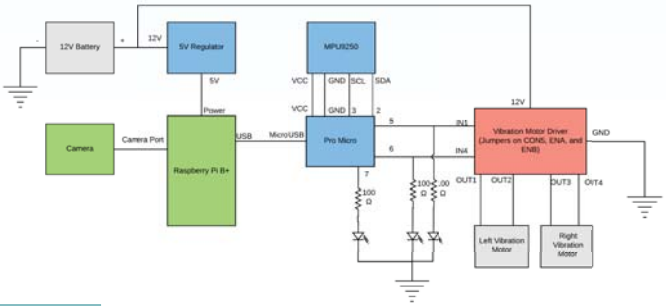

1. Load a frame from live video taken by camera and load rotation matrix from IMU
2. From image, find lines on the bottom of the pool. Obtain L1, L2 from the image and g from the IMU
3. Calculate correction vector
4. Send feedback to user via haptic motors

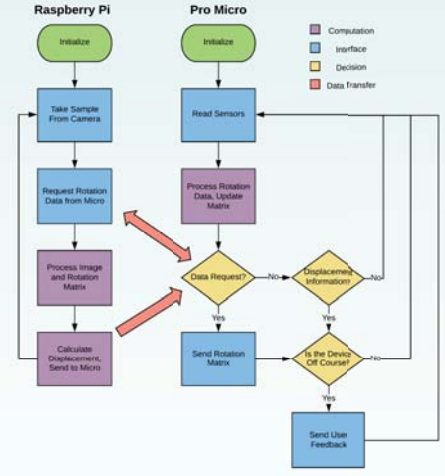

Next Steps

- Add end of lane detection functionality
- Test with end users
- Consult with end users to improve form factor and functionality
- Move device to wearable technology

Acknowledgements

We would like to thank The Inn at Pasatiempo for generously allowing testing of this device to be conducted at their pool facility. Furthermore, we would like to thank Kresge College for funding this project.

Braille sticker here

Baskin Engineering UC SANTA CRUZ

Capstone Project

Off Road Spotting Drone

NingHao He, Ryan Steinwert,
Manveer Randhawa, Zhaoheng Chen

DAIMLER

Abstract

The ability of drones to get a different perspective of an environment becomes useful in off-roading, where the success of the vehicle depends on choosing the perfect route through the terrain. Through building an off-road spotting drone, Mercedes hopes to achieve an engaging and safe off road experience for their SUV market.

Overview

Mercedes is committed to improving the experience customers have with their vehicles. In order to bring more convenience and utility to the user they sponsored our team to build a drone integrated with their SUVs, to provide real-time information on the terrain that a driver cannot directly see and improve off-road safety. We utilized a drone with multiple sensors to achieve:

- A 3D mapping of the environment.
- An algorithm to find the best path through the terrain ahead.
- An intuitive app display showing obstacles and the best path to the user.

Challenges

Finding and learning about the combination of technologies such as drone firmware and best path algorithms required to put together this project was our team's greatest challenge. Ultimately we gathered a strong blend of techniques to give us a stable and reliable system.

Approach

We decided to build our drone instead of buying one because of the added control over the drone's features and abilities.

Drone

- PX4 drone firmware
- LIDAR sensor for scanning

Server

- Raspberry Pi Server, running A* best path algorithm
- QGroundControl open source ground control

Display

- React Native app
- WebGL rendering used for creating the depth map

Architecture

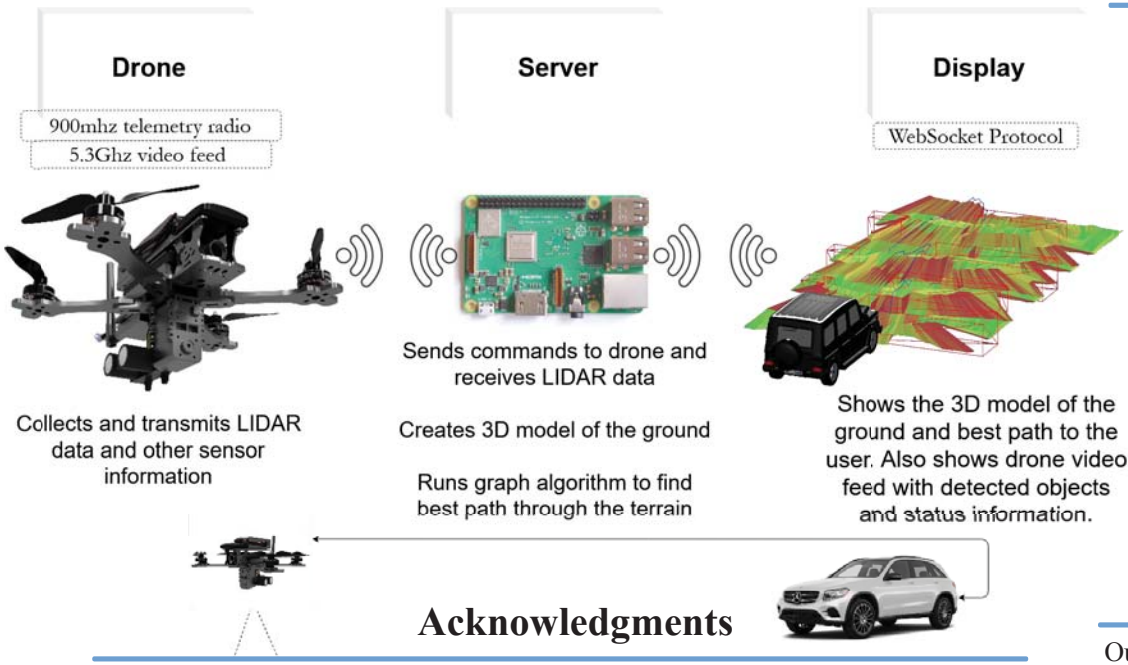

Results

- Built a drone capable of taking off and landing from a trunk, scanning the ground in front of the vehicle and returning the LIDAR data back to the server.
- Built a server for communicating between the display and drone.
- Designed an app that is able to show a 3d model of the ground and live video from the drone with obstacle overlay, and help the user navigate the terrain.

Conclusion

Our drone system creates a 3d map of the environment in front of the vehicle and finds the best path through the terrain.

Acknowledgments

We would like to thank our sponsor Mercedes, Jeff Bertalotto, Kevin Gee, Utkarsh Gautam, our TAs Faeze Brahman and Chandranil Chakrabortti, and Professor Jullig for their support and mentorship.

Indoor Robotics

Performance of Automotive Sensors in Indoor Settings

Babak Farahmand, Daniel Luft-Martinez, Mayowa Borisade, Paige Riola, Shealtiel Mulder, Walter Condori

Special thanks to:

Shuhei Takahashi, **Josh Frankfurth**, **Jolton Dsouza** & **Rob Tucker** of Continental, Our professor **Richard Jullig**, Our TAs **Chandraniel** & **Arindam**, And **Veronica Hovanessian**!

Abstract

We experimented with a fusion of automotive lidar, radar, and camera sensors in order to evaluate their potential for indoor use. These sensors are integrated with a Husky UGV robot; this setup allows us to collect data for tasks such as object detection & classification. We then evaluate the accuracy of our results by comparing our finalized object classifications to the actual environment.

Overview

Continental is an automotive parts manufacturing company with a variety of **sensors** suited for tasks such as collision detection and adaptive cruise control. With an expanding robotics market, Continental is interested in marketing these same devices for use in **indoor environments**. Our goal is to analyze the performance of Continental sensors, in indoor settings, and determine the **best combination of their sensors** for indoor robotics use.

Challenges

We began with the pretense of having physical access to the hardware. However, we faced many setbacks in acquiring and accessing the hardware, including legal delays for on-campus hardware access, and and campus closure due to COVID-19. Each new setback changed our workflow, shifted our goals, and reset our progress. We only managed partial hardware access towards the project's end. In response, we shifted focus to our data processing methods.

Approach

Data from our sensors is transferred to our middleware, ROS (Robot Operating System), located both on the Husky and the Jetson TX2, an external AI-computing device. ROS consolidates our sensor data for processing multiple data streams simultaneously. Each of the sensors publishes it's data to ROS through executables called nodes, which follow a publisher-subscriber pattern. We then perform radar-lidar object detection and live video classification, with the results published back to ROS. Ultimately, we're combining the data to map object labels to a 3D model of the environment within ROS. We can then determine the accuracy of the model by comparing it to the actual environment.

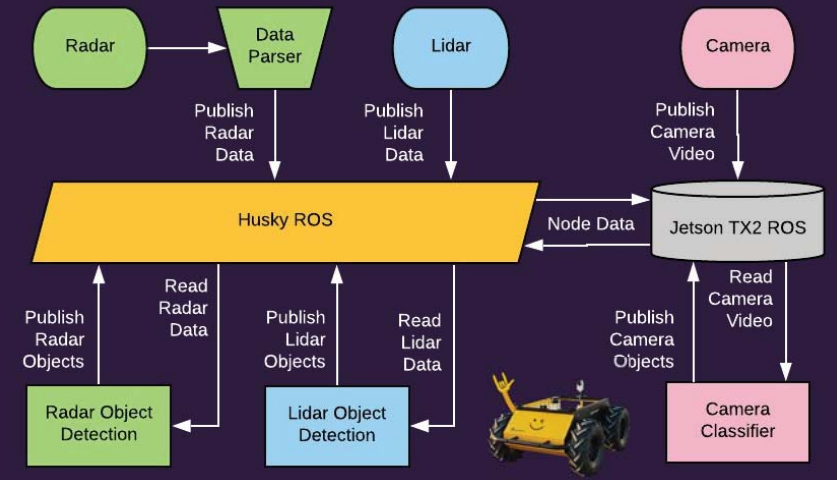

Edge Detection

Object Classification

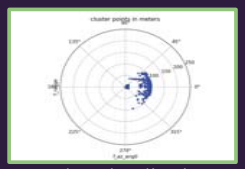

Radar Visualization

Conclusion

We were able to accomplish a variety of processing methods. For the camera, we implemented live video edge detection and object classification. In addition, we achieved video data streaming within ROS, which we can apply these processes to. For radar, we were able to parse the raw data directly from our radar sensor, produce a 2D visualization of the data, and detect objects. For lidar, we implemented 2D mapping within a simulated environment.

Given more time, we would integrate all of our processes into a single ROS-centered system, and produce a full 3D visualization of our processed data within ROS. Upgrading to the most recent ROS release would simplify the integration of our Python 3 code with our overall system.

Testing System for Respiratory Devices: Human Cough and Mucus Production Simulator (Googun)

Lisa Mcphillips (Bioengineering), Natalie Luttrell (Bioengineering), Owen Bosley (Bioelectronics), Patrick Mogianesi (Bioelectronics), and Todd Nabonne (Bioelectronics)

Project Background

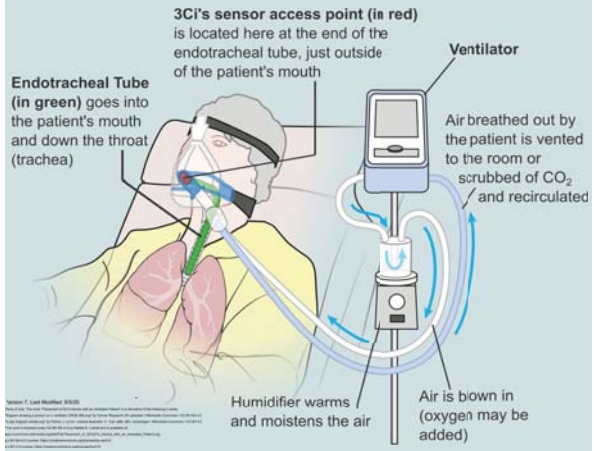

The Googun is a corporate-sponsored project aimed at building a system for stress-testing medical devices in development, by providing a cheaper, safer alternative to human test subjects. Our Googun offers pseudo-realistic mucus production and cough behavior, similar to that of an average person, making it ideal for testing respiratory devices.

Our client, Certus Critical Care Inc. (3Ci), is developing prototypes that have sensors located near the end of an endotracheal tube. They need to test how their sensors are affected by the secreted mucus from the patient at this location, as they do not yet have data on the short and long term effects of mucus on their sensors.

The Team

Left to right: Patrick M, Natalie L, Lisa M, Todd N, and Owen B.
Photo taken 1/6/2020

Our Solution

The solution is split into two parts: Mucus Injection and Cough Production. Both are controlled by a microcontroller (MCU) which receives system commands from the user interface.

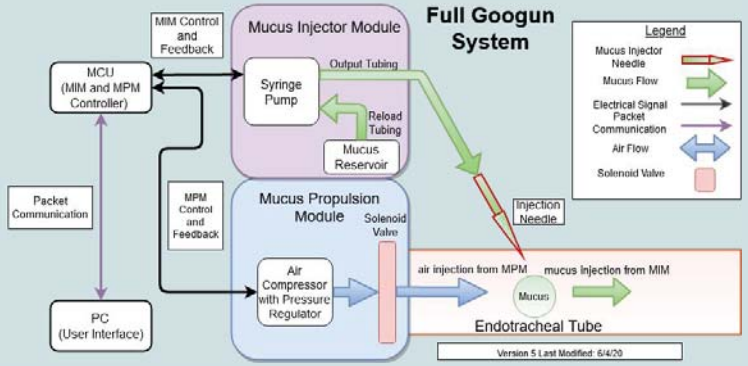

Mucus Injection Module (MIM)

The MIM functions as a substitute for human mucus production and injects controlled volumes of artificial mucus into the endotracheal tube.

Mucus Output Characteristics:
-Minimum resolution of 5 µl.
-Maximum output rate of 2 ml/min.

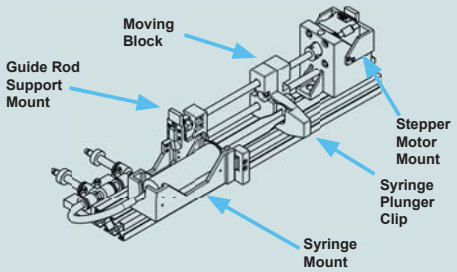

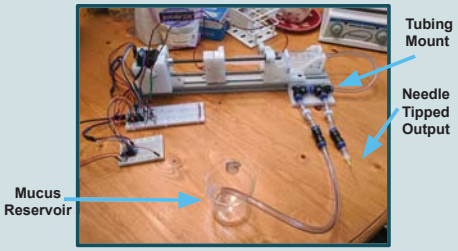

Mucus Propulsion Module (MPM)

The MPM simulates characteristics of a human cough, and functions to push the mucus in the endotracheal tube towards the client's sensor.

Air Output Characteristics:
- Human-like air volume output up to 750mL, at airflow rates between 160-300 mL/s.
- Human-level pressure range of 40-170cmH_2O

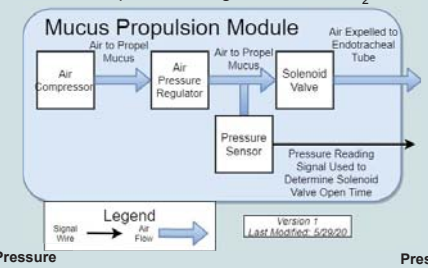

User Interface

The client required an interface to control the MIM and MPM modules. Our solution involved a graphical interface that monitors the status of the modules and allows for the user to send instructions on operation of the modules.

Googun Control Graphical Interface

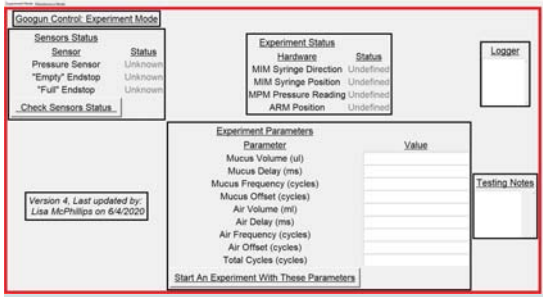

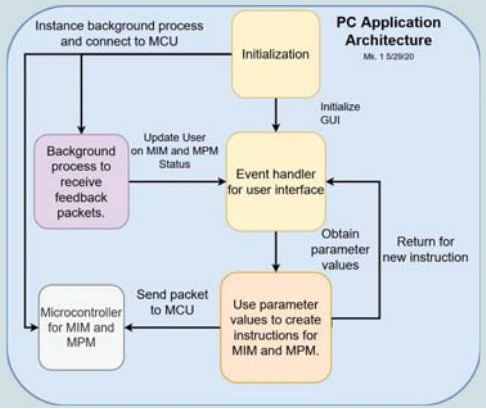

Conclusions

We successfully fabricated and characterized a Googun system. Unfortunately, due to the COVID-19 pandemic, we were not able to fully vet the system with verification tests. Due to social distancing laws, we were not able to test the systems together. The pandemic also prevented us from obtaining a working ventilator or an endotracheal tube.

Acknowledgements

We would like to thank David Poisner from Certus Critical Care Inc, professors Petersen and Favaloro, Joe Cox, and our TA Sushmita Joardar for making this all possible.

Intuitive Auto Irrigation

"Bringing residential irrigation into the 21st century with active sensing and real-time forecast data"

Sam Aiken (Data Logging and Forecast Lead), Brian Naranjo (User Interface & Water Delivery Lead),
Grant Skidmore (Project, Software, & Wireless Communications Lead), and Henry Tuckfield (Sensors & Power Systems Lead)

Problems with Automated Irrigation

According to a study performed by the EPA in 2017, "Outdoor water use accounts for 30 percent of household use, yet can be much higher in drier parts of the country and in more water-intensive landscapes."[1] Today, most irrigation systems still rely on outdated timer-based operation, resulting in inefficient water usage and overwatered plants. The days of irrigation systems filling gutters with unused freshwater needs to come to an end; it is about time residential irrigation systems get updated with all-new technology.

Reducing Water-Waste using Sensors

The Intuitive Auto-Irrigation (IAI) solution utilizes sensors placed near plants of interest. Through wireless communication, these sensor nodes relay readings of soil moisture and light level to a base station (the central hub) which accumulates all sensor and weather forecast data to determine when to trigger water delivery. Using the collection of data, the IAI system avoids unnecessary watering, effectively leading to a more sustainable method of general landscaping and residential gardening practices.

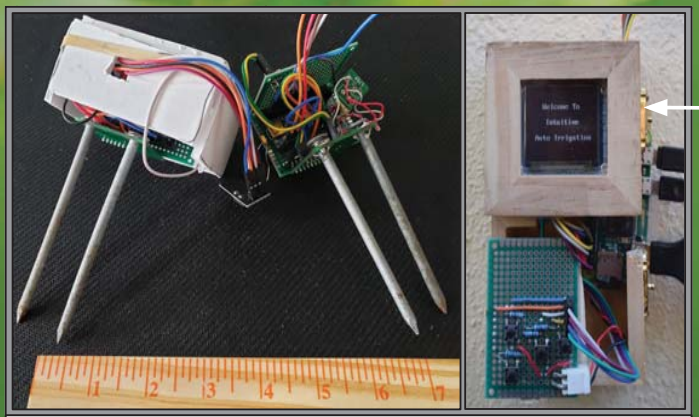

Figure 4: *(Left)* Two sensor nodes *(Right)* Central hub connected to user interface.

Central hub is controlled with buttons on the user interface

References

[1] https://www.epa.gov/watersense/how-we-use-water
[2] https://wasatchgardens.org/resources/item/258-drip-irrigation

System Design

- **Central Hub:** Wall-powered master node which receives real-time data from sensor nodes and weather forecast information to determine when to water. Key features include:
  - **User interface:** quick and easy configuration of the system
  - **Water delivery:** precise control of irrigation using latching solenoid valves
  - **Weather forecasting:** prevention of irrigation based on precipitation chance
  - **Data logging:** data stored locally and on a database for easy access to users
  - **Wireless communication**: seamless transfer of data between sensor nodes and central hub over a 400ft minimum range makes setup less restricting

- **Sensor Nodes**: Multiple battery-powered subsystems featuring several sensors, each node is found at a watering subjects' location. These nodes collect data and wirelessly transmit this data back to the Central Hub.
  - **Equipped Sensors:** continuous monitoring of plant conditions using custom soil moisture sensors and light-level sensors.

Figure 3: Sensor node with its protective covering placed in an ideal use-case location.

The sensor node's 4.5" probes dig into the soil to record soil moisture levels

Water is delivered when conditions are met and forecast predicts no rain

Day and night cycles indicated by red/yellow light level indicators

Potential Impacts

Our system proves the merits of sensor-based irrigation by effectively watering plants under ideal conditions and reducing unnecessary water waste. This precise irrigation method would save users time and resources while accurately monitoring soil moisture content for a range of gardening applications. Through this system, we hope to create a gardening culture built on technology and sustainability.

Soil Moisture Sensor Regression Modeling

In order to prevent unnecessary watering, the IAI system utilizes a custom soil moisture sensor at each node to determine the optimal time to water. For accurate soil moisture sensing, we use a power series regression model to accurately map soil moisture content to resistance across the sensor probes. Figure 1 shows this trend hold across two unique sensors and three sub-samples from a given location.

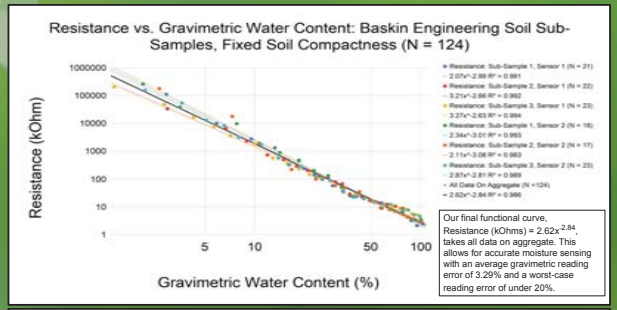

Figure 1: Resistance between moisture sensor probes and relative soil moisture, measured by gravimetric water content.

Validation Testing & Results

Figure 2 shows data gathered during a 5-day field test that involved two sensor nodes and one master node. In the figure, the blue and green lines show the gravimetric water content for Node 4 and Node 5 respectively; the yellow and red lines show the ambient light level measured in lux for Node 4 and Node 5 respectively; and the vertical black lines indicate when water delivery was triggered. The test data showed the system watering when specific environmental conditions were met, proving its ability to provide irrigation only when plants need water.

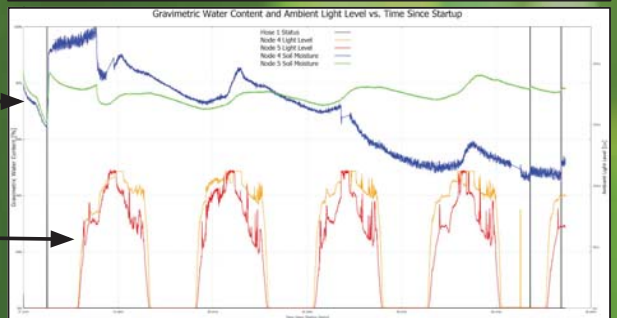

Figure 2: Testing results of a 5-day test to show correct operation of the full system in a realistic environment.

ARbot: Automated Recycling Robot

Recycling technology with the user in mind

Viet Nguyen (RE), Martin Contreras (BME), Clayton Tan (CE), Neha Dhayanand (CE), & Logan Fansler (RE)

The Problem: A Higher Standard For Recycling

Santa Cruz Resource Recovery Facility's "clean" paper waste stream.

- Recycling contamination levels in the US range between 10-25%
- Recycling contamination levels at UCSC is about 70%[2]
- U.S. loses $1 billion per year from contaminated recyclables
- Some counties have resorted to shutting down recycling operations
- Others have resorted to burning recyclables altogether
- Burning recyclables has led to respritory illnesses for nearby neighborhoods

About ARbot

The purpose of ARbot is to help Material Recycling Facilities (MRFs) meet the current 0.5%[1] contamination threshold imposed on global recycling operations. The ARbot team aims to deliver a sorting technology, designed to meet the industry's cost-sensitive demographic.

How ARbot is different?

- Provides a low-cost product
- Works alongside the facility and its workers
- ARbot is easily deployable
- ARbot is custom designed for each client

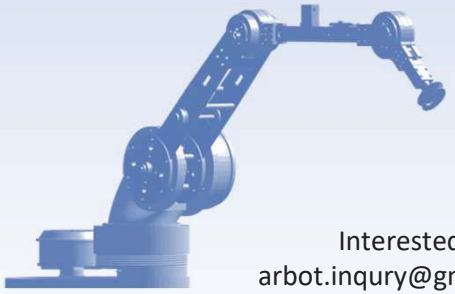

Interested?
arbot.inqury@gmail.com

Our Procedure

Detect & Track Contaminants

Use state of the art computer vision detection and tracking algorithms to identify contaminants.

Compute Trajectories

Plan and predict expected target location

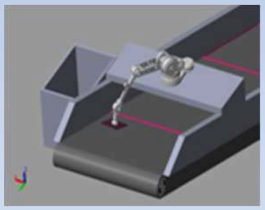

Execute Sorting Task

Complete sorting trajectory to remove target from waste stream

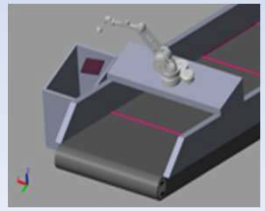

References

[1] T. Eng. Could the chinese national sword inspire global recycling innovation? [Online] https ://recycling.tomra.com/blog/chinese-national-sword-inspire-global-recycling-innovation

[2] L,Johnson. Campus Mixed Recycling Landfilled Since November [Online] https://www.cityonahillpress.com/2019/03/01/campus-mixed-recycling-landfilled-since november/#:~:text=UCSC%20sends%20six%20tons%20of,is%20about%2070%20percent%20contaminated.

We are very pleased to include posters for the Senior Design Showcase that were done without industry sponsors. Some of these projects were instigated and/or sponsored by research at the Baskin School while others were created by students with the assistance of faculty mentors and TAs.

Capstone Project

Edu Plastic Pollution

Cristian Castillo & Anna Kovneva

Abstract

Edu Plastic Pollution is an interactive education web application made in collaboration with the students at the IDEASS lab, a UCSC program that brings together students of different disciplines to tackle projects dealing with social, ecological and economic issues. Edu Plastic Pollution deals with the growing problem of plastic pollution by educating the general public on the life cycle of plastic as it deteriorates from a macro to a micro to a nano stage, and its detrimental effects on the environment, wildlife, and human health in each stage. Edu plastic pollution takes an interactive and engaging approach to educating the public on issues dealing with plastic pollution.

Our Approach

To create an interactive and engaging experience we focused heavily on the user experience. Our website is broken down into the three life stages of a plastic: the macro, micro and the nano.

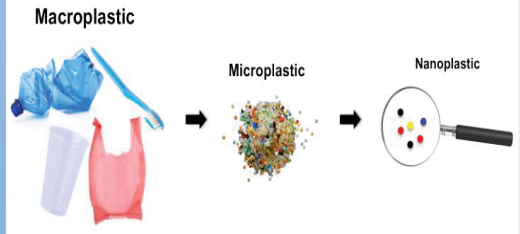

- We showcase how a macro can be recycled, trashed or dumped. If trashed it ends up in a landfill. If recycled it ends up on a boat to China or a landfill, but if dumped then it can eventually end up in our oceans, and then in our food.
- A micro can be found in things such as clothes, facial scrubs, or toothpastes and end up in our oceans, and eventually consumed as nano particles in seafood.

Overview

Over the last few decades, global consumption and production of plastics have skyrocketed. An estimated seven billion tons of plastic over the last seventy years accumulated worldwide, with merely half just within the last two decades. Plastic Pollution is a rapidly growing problem since the general public contributes to it daily. Edu Plastic Pollution aims at bringing awareness to plastic pollutants and their detrimental effects. Current plastic pollution educational websites can be disengaging and do not focus on all 3 stages of a plastic pollutant. Our interactive application makes it easy and fun to learn about the life cycle of plastic pollutants, and how you can do your part to help in fighting plastic pollution.

Macro, Micro, Nano Path Breakdown

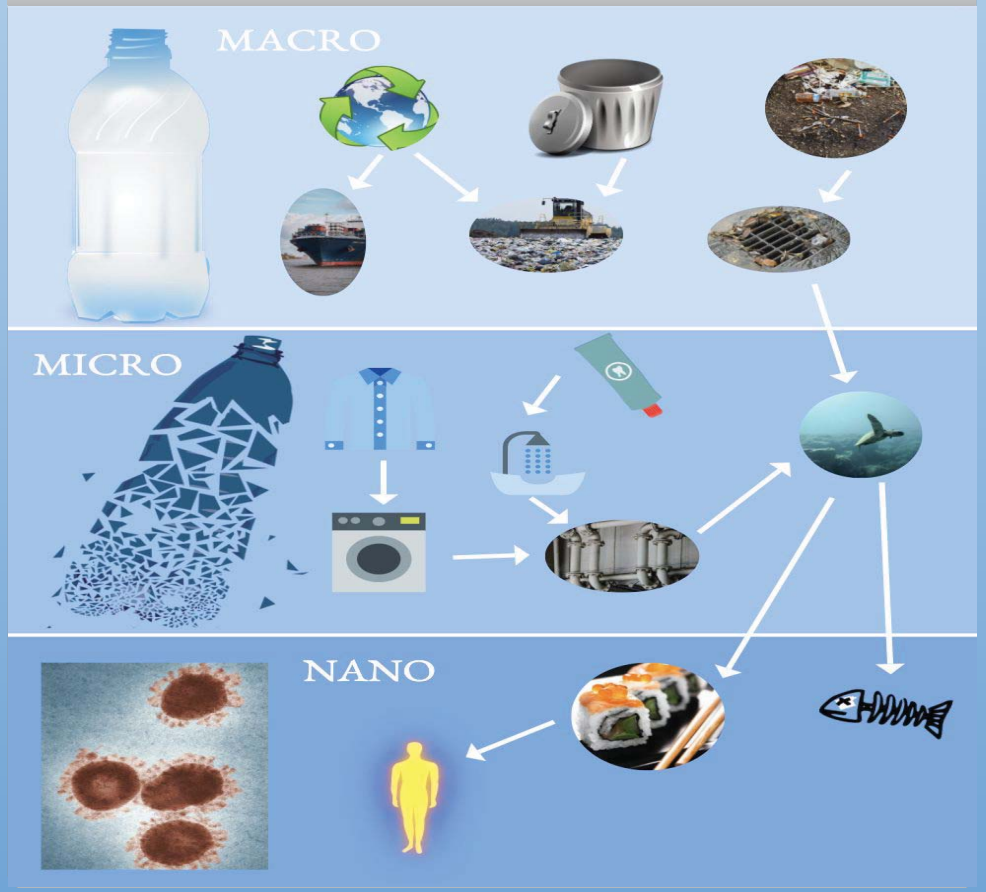

Technology Stack

- React
- React-Spring
- HTML
- CSS
- GitHub Pages
- Node.js

Results

The resulting product is an engaging and interactive website that makes learning about plastic pollution easy and fun for users of all ages. Our website offers:

- Three interactive informative story paths
  - macro, micro, nano paths
  - 3D interactive story
- Animated statistical data graphs
- An activism route
  - including links to activism organizations

Conclusion

Edu Plastic Pollution's goal of creating a web application to help bridge the education gap on plastic pollution was accomplished. The user's experience has been revolutionized making it easier and efficient to learn by filling in the educational void surrounding plastic pollution and its detrimental effects. Through interactive learning, we instill pertinent information that helps decrease plastic pollutants contributions over time. Of course, there is still more work that can be done to create an even richer experience.

Acknowledgments

Thank you to Alexis Hooper, Benji Weaver, Kavika Peiris from the IDEASS lab for the collaboration. Thank you to Collin Hurst, our Winter quarter group-mate, for his efforts on the project as well. We also thank our professor Richard Jullig and TA, Akila De Silva for supporting and giving us guidance throughout the development stages.

Capstone Project

Marine Plastics Monitor

Sean Gibson, Karishma Shah,
Orlando Chavez, Raeeka Yusuf

UC SANTA CRUZ Baskin Engineering

CLEAN OCEANS INTERNATIONAL

Abstract

Marine Plastics Monitor is a platform sponsored by **Clean Oceans International (COI)** where people can use a protocol to survey debris found on their local beaches. People can use a data sheet and submit it on the website, or use the mobile app, which sends data directly to the website. The website shows analysis of the surveys. This year, we made updates to the website and mobile app to make it more user-friendly, match them to the data sheets provided by COI, and allow users to edit and delete their surveys.

Website

- Submit beach surveys
- Create, edit, and delete own survey to correct mistakes in survey
- Map of beaches with data
- Visual representations of data

Mobile App

- Conduct surveys on any mobile device offline and send to website online
- Tutorial on how to conduct a survey
- Users can click on types of debris in survey to see example pictures of debris
- Dynamic styling so interface is same on all platforms

Surveys

SRS: along top of beach is 100m "spine". Select four "ribs" 5m wide, and collect data on debris in rib from spine to ocean.

AS: entire team records remaining debris on beach

MDS: in each rib, pick 1 spot to put 1 ft^2 of sand in sieve. Record debris that collects.

Overview

The goal of our sponsor, COI, is to allow people around the world to use their protocol to collect data on debris found on beaches and submit the data to the website for analysis. The purpose of this is to study trends and see if different recycling methods have a positive impact on the amount of debris that collects on beaches. COI provides a data sheet to collect the data and submit it to the website. The website and mobile app have both been updated to reflect recent updates to the data sheet. The protocol calls for three sub-surveys: **Surface Rib Scan (SRS)**, **Accumulation Survey (AS)**, and **Micro Debris Survey (MDS)**.

Website Data Analysis

Total debris per survey on beach

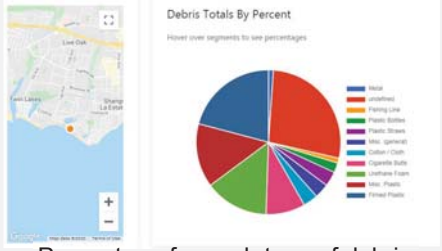

Percentage for each type of debris

Mobile App Survey Form

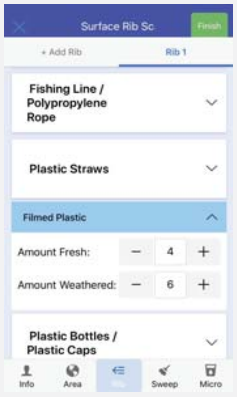

Mobile App: SRS

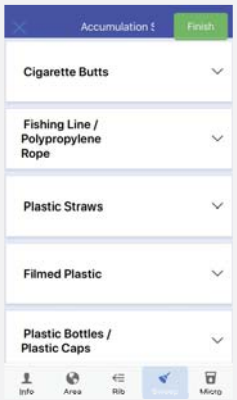

Mobile App: AS

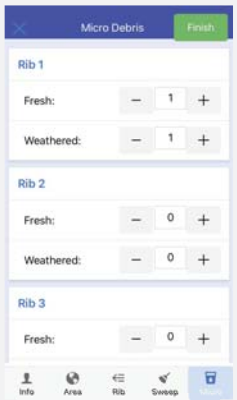

Mobile App: MDS

Acknowledgments

We would like to thank our sponsors from COI, David Schwartz and Jim Holm. We would also like to thank Dr. Richard Jullig and teaching assistants Scott Davis and Golam Muktadir.

Software

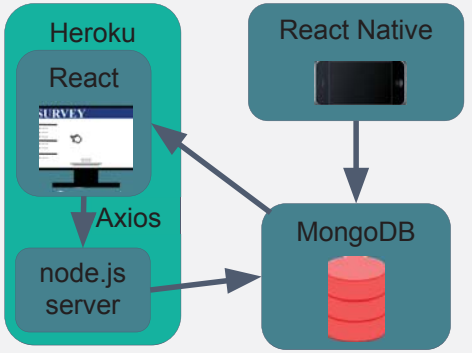

- React: code for website
- React Native: code for mobile app
- Heroku: Hosts website
- MongoDB: database
- Axios and node.js: server to send and receive data from database

Major Improvements

- Created a slideshow tutorial on how to conduct surveys
- Updated website and mobile app to reflect changes to survey made by our sponsors
- Updated server API to analyze micro debris in surveys
- Fixed bugs so users can view, edit, and delete surveys
- Added photos of types of plastic in app
- Changed mobile app design to have dynamic styling

Conclusion

Future plans:

- Publishing mobile app on Google Play and App Store
- Integrating Marine Plastics website with main COI website

Capstone Project

Long Marine Stranding Map

Ashley Pauley, Azizkhuja Asomiddinov, Harrison Fox
Maya Apotheker, Natalie Wilson, Pavel Yakovlev

Abstract

The stranding map is a web application to visualize marine mammal strandings of Santa Cruz county. This tool's features include multi-attribute filtering, marker clustering, heatmap layer, time series animation, and data upload. This allows both the Long Marine Lab and the public to look at stranding data in a more intuitive and effective way.

Approach

- **Firebase** stores the stranding data. This can be updated easily when new strandings are reported.
- **Mapbox** displays all the marine mammal strandings on the map with their GPS coordinates and creates the heatmap view.
- **Algolia** allows us to achieve fast multi-attribute querying for the filter.
- **Netlify** hosts the web app and some of the backend code.

Overview

To drive many of its research projects, the UCSC Long Marine Lab (LML) collects data from examining strandings: an event in which a marine mammal has washed ashore, living or deceased. With this tool, lab technicians can view the data they've collected over the years on an interactive map. As a scientific tool, the stranding map opens up new avenues of data analysis for the lab and provides geospatial visualization for LML's research in marine life, coastal conservation, and climate change. Making the map public can improve collaborative planning efforts and promote public involvement.

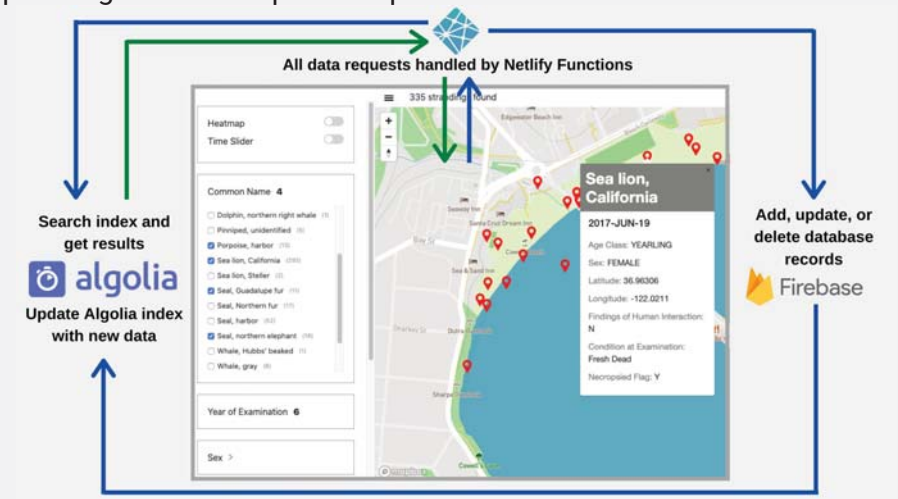

Check it out at lmlstrandingmap.netlify.app

Acknowledgments

Thank you to our sponsors at the Long Marine Lab: Dr. Robin Dunkin and Amber Diluzio. We would also like to thank Professor Jullig and our TA's Akila De Silva and Scott Davis that have provided support and mentorship throughout this project.

Features

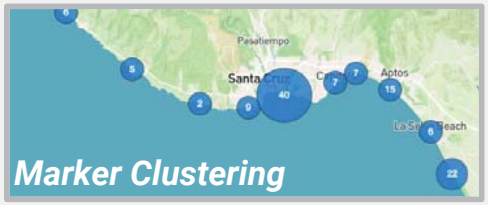

- Groups data points together and shows count of standings when zoomed out

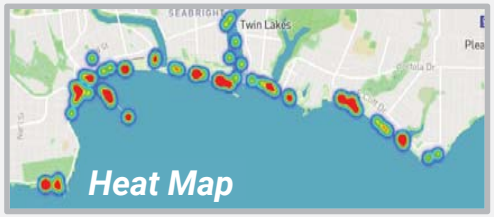

- Highlights density and magnitude of strandings in a particular area

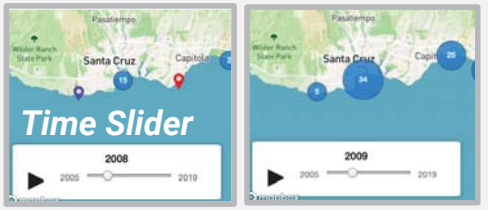

- Illustrates how stranding occurrences change over years

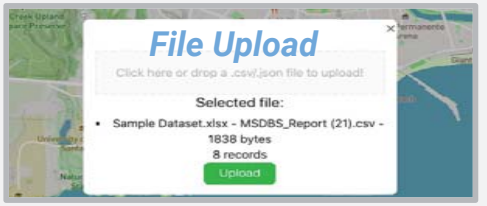

- After logging in, lab technicians can upload new strandings to the database

UC SANTA CRUZ
Baskin Engineering

Capstone Project

Land Trust Management

Yuhao Deng, Matthew Wirtz, Zihan Yang

LAND TRUST *of* SANTA CRUZ COUNTY

Abstract

Land Trust Mobile is a mobile app that allows users to register for permits, report conditions, access company content, and view maps of the preserved regions.

Land Trust Web is a complimentary website for the land trust office to view and process permit requests and reports.

Approach

The **mobile app** is for those visiting the nature regions. The **web app** is for the data those visitors generate.

For the mobile application:

- **Home** directs to the existing Land Trust services (become a member, donate, social media and other content).
- **Permit** emails the user a San Vicente permit; adds an entry to the database.
- **Report** sends the entered data to the database (see bottom right corner).
- **Maps** uses GoogleMaps to show the location of listed regions.

For the web application:

- **Reports** and **Permits** each pull from the database to display their data.

Overview

Land Trust of Santa Cruz County is committed to protecting, maintaining, and sharing natural areas. This land trust's non-profit work consists of constructing coastal bike routes, creating forest trails, and keeping natural places safe from industrialization. In order to extend the capabilities of Land Trust of Santa Cruz and enhance the experience for visitors of these regions, we developed a mobile and web application.

Architecture

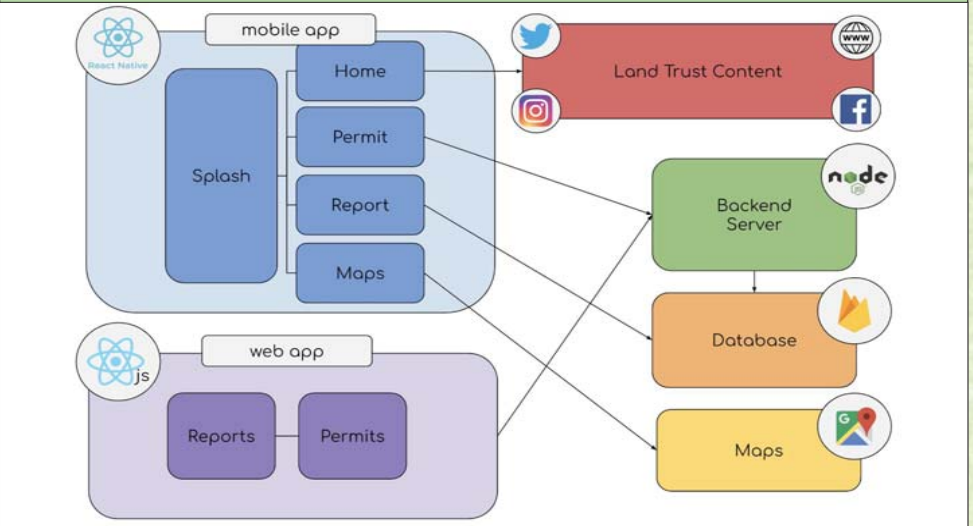

Acknowledgments

Land Trust of Santa Cruz County:
Brendan Quirk
UCSC:
Dr. Richard Jullig, Akila de Silva

Conclusion

We've extended the capabilities of Land Trust of Santa Cruz County by developing a cross-platform mobile app, and a web app. The mobile app will be in stores soon! Future improvements could include permit support for other regions and push notifications.

Santa Cruz Regions

Hiking trails:
- Byrne-Milliron Forest
- Glenwood Open Space Preserve
- Star Creek


Places for birding:
- Antonelli Pond
- Watsonville Slough Farm

Results

The **mobile app** style mirrors the land trust website.

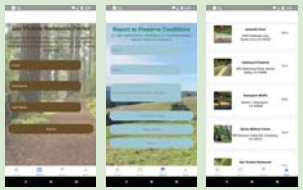

Forms to register for a permit and report conditions; list of maps.

The **web app** visualizes the mobile app data for the company to review.

UC SANTA CRUZ
Baskin Engineering

Rent-A-Driveway (R.A.D)

Korbie Sevilla, Henry Dang, Scott Melero,
Adithya Somasundaram, Ivy La

Abstract

Rent-A-Driveway allows the community to post and rent driveways for parking; making it easier to connect drivers in crowded areas to homeowners with open spaces available! This is accomplished through the R.A.D mobile app which uses a combination of Google API services to search for and create parking spaces.

Why a Parking App?

In the United States, the average driver spends approximately **17 hours per year** searching for parking at a cost of $345 per driver.

At UCSC, **2898 students purchase parking permits**, roughly 600 students are on the **waitlist** for those permits, and about **20,000 parking citations** occur each academic year. The demand to increase **parking accessibility** can be addressed by monetizing and opening access to driveways.

How R.A.D Works

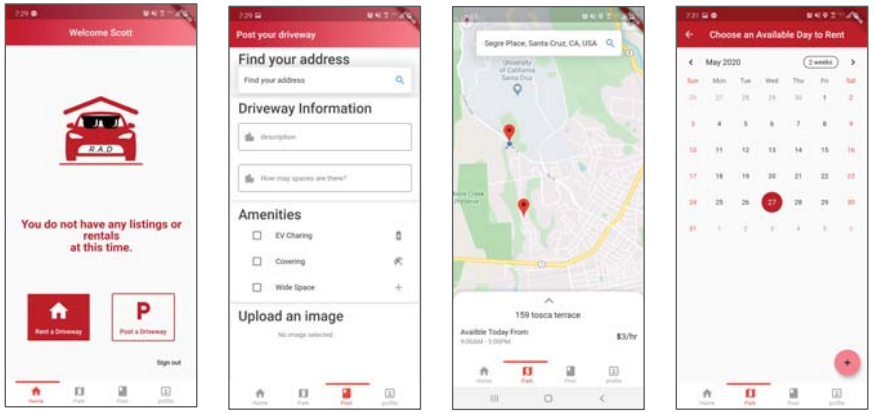

Landing screen | **Post your driveway** | **View open driveways** | **Schedule a time to rent**

System Architecture

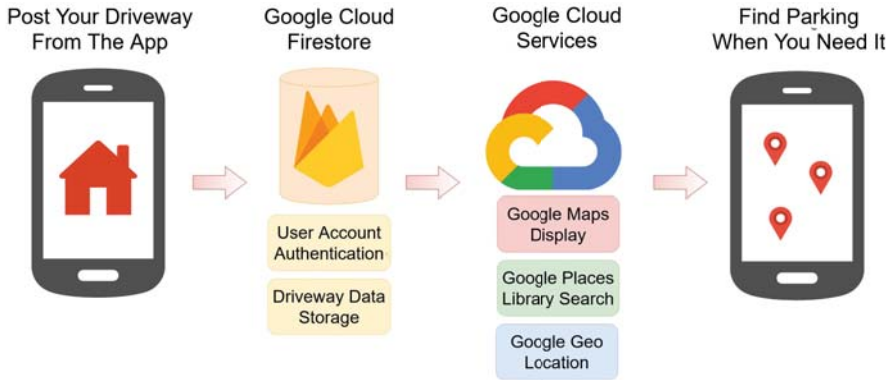

Our application was designed with Adobe XD and Supernova Studio.
It uses a Google Firebase backend, a Flutter front-end, and Google cloud services to handle the functionality.

Features

- **Post your driveway** for users to rent, bringing more money to your wallet.
- **Interactable map** to search and filter through available parking spaces and vehicle specific utilities such as electronic charging.
- **"Park now"** Find available spaces near you, for you, in real time.
- **Smart pricing algorithm** that suggests a spots worth based on time-of-day demand.
- **Driveway Schedule** to reserve parking for a later date.

Future Work

Next year, R.A.D will implement security features such as transactions and proof of residency.

Many thanks to:

Richard Jullig,
Police Chief Oweis Nader,
Faeze Brahman,
Scott Davis,
Golam Muktadir

UC SANTA CRUZ
Baskin Engineering

Capstone Project

ResearchConnect

Faraz Rashid, Cecilia Li, Liren Wu

ResearchConnect

Abstract

ResearchConnect is a platform for professors to advertise research opportunities for both undergrad and graduate students also providing incoming grad students with possible advisors.. Professors post research opportunities and filter applicants with questionnaires. Students can browse through currently available research opportunities to apply and find helpful information on how to approach professors.

Approach

We first wanted to familiarize ourselves with the new technologies and existing code.

Frontend:

- React.js
- Bulma

Backend:

- Node.js
- Express.js
- MongoDB/Mongoose

Once this was done, we wanted to add features to help students find professors easier by providing more data on professors. We also wanted to add a place for students to share their research experiences.

We also wanted a way for administors to manage some aspects of the database through the website alone.

Overview

Right now, incoming grad students does not have a platform to use to find out which professors are interested in being advisors and how to contact these professors. And both grad and undergrad students do not know where to find research opportunities. We provide such platform that solves both issues at once. This is a continuation of a past senior project where previously the focus was to target undergrad students. For this current iteration, we wanted to focus on incoming graduate students and help them find suitable advisors. ResearchConnect also aims to be a platform with updated information on faculty members and a forum base for students to share their experiences.

Architecture

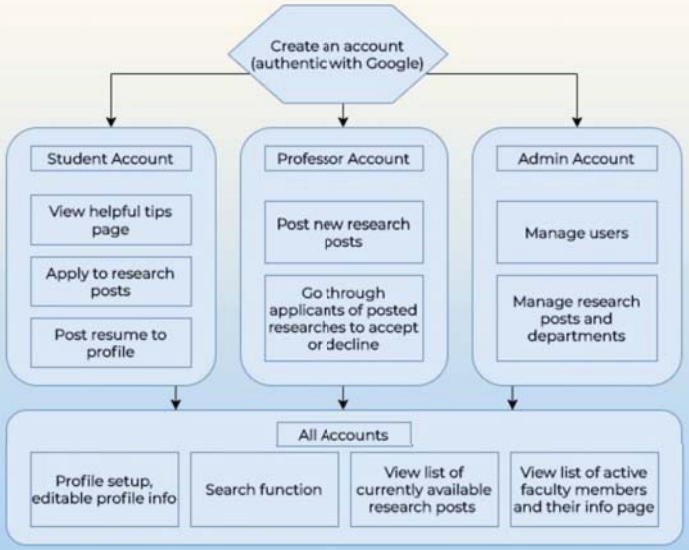

Acknowledgments

Special thanks to Patrick Mantey and Muktadir for helping come up with ideas to advance this project.

Challenges

- Difficulties in gathering useful and up-to-date information on professors.
- Needed professors and students to use and test product.
- Inconvenient to meet with teammates and faculty members
- Need to register under UCSC domain

Results

Admin Accounts:

- Manage users (ex. assign roles, remove users)
- Manage research posts and departments (ex. edit/remove)

Students Account:

- List of professors with bios and links to their research publications.
- Helpful tips page

General:

- Clearer and more efficient UI

Conclusion

We were able to provide a platform to help graduate students find advisors based on active research opportunities or research interests. As of now, Research Connect has it's basic functions available but could be improved on more.

UCSC Bus Tracker

Brian Thyfault, Radomyr Bezghin, Rizzian Tuazon, Nathan Lakritz
CSE 115B/C - Software Design Project
Spring 2020

Abstract

Currently, UCSC students do not have a way to track campus buses. There have been many attempts in the past, but none were designed to last and only worked for a couple years. Students need bus data to efficiently navigate around campus. As it stands, the current bus system without accessible live tracking is unreliable.

Our iOS app provides much-needed bus tracking alongside other helpful features. This is being done using an old hardware setup on campus used to locate the buses with radio base stations. With this data, we provide a clean and modern app to live-read the GPS location of on-campus buses and get an accurate ETA to bus stops.

Schedule & ETAs | Route Highlighting | Info Page

Features

- Real-time bus location tracking with distinct bus icons
- User location and points of interest give context to the map
- Bus orientation feature allows users to easily identify the direction of travel
- An ETA Table for every bus stop makes it easy for users to catch buses on time
- Colored bus routes help users who are new to the campus know exactly which bus to take to get to their desired destination
- Info panel lets users quickly learn how the bus system works

Approach

The hardware team that we were collaborating with was responsible for producing a whole new tracking system that includes: new on-bus tracking devices and upgrades of base stations. Together with our iOS application, these changes will ensure that students get reliable bus information whenever it's needed.

Mapbox: This is the core API our application is built on. It provides essential functionality for mapping points of interest and dynamically displaying the buses on a map.

MySQL & PHP: We are using data provided by the hardware team, that is stored on a mySQL database.

Radio Base Stations: Are acting as a middleman by collecting the data from the buses and sending it to the database.

Architecture

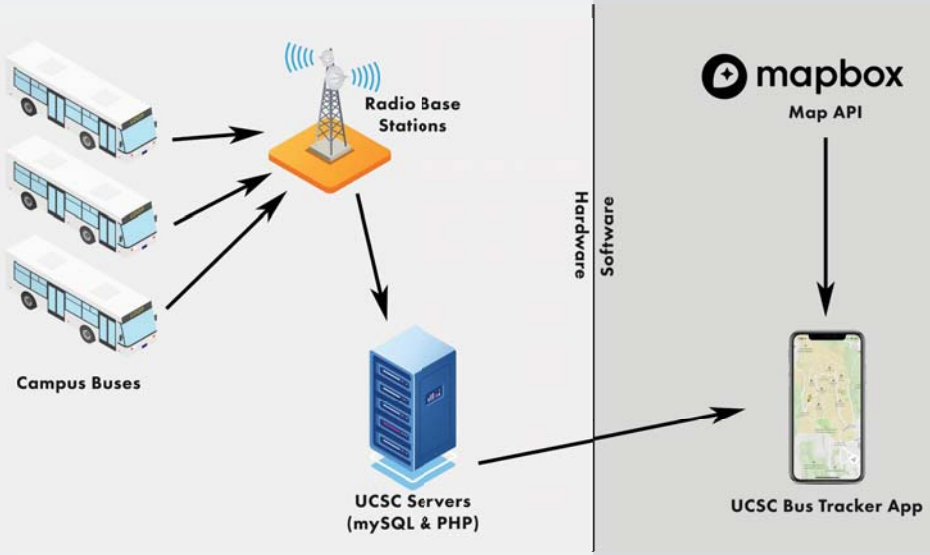

Conclusion

Our team was able to create a mobile application that has a great-looking and intuitive UI as well as functionality that will help students navigate this school.

Hardware team succeeded in creating new and improved tracking devices, but its integration has been delayed due to COVID-19 concerns. Installing newer hardware onto the buses will conclude a full revamp of the UCSC Bus Tracking System which hopefully can happen in the Fall.

UCSC Bus Tracking System 3.0 (BTS 3.0)

Team: Katelyn Young, Gavin Haight, Farinaz Rezvani, and Alexander Zuo

Jack Baskin School of Engineering-UC Santa Cruz

Introduction

Students & faculty at UCSC require accessible, reliable, and accurate information to create an efficient schedule for moving around, and off (to/from) campus.
We aim to provide an integrated bus tracking system that will be accessible to all campus bus users. Bus Tracking System 3.0 (BTS3) will address the cost issues encountered with the previous systems, create a scalable maintenance plan for TAPS, and add an Estimated Time of Arrival (ETA) feature at major bus stops.

Our Product Compared to our Competitors:

1) **Compared to off-the-shelf tracking systems such as GMV Syncromatics, our system is:**
   a) Cheaper
   b) Without monthly fees for cellular service
   c) Equipped with a route sign
2) **Compared to Bus Tracking System 2 (BTS2), our system is:**
   a) Cheaper
   b) Compatible with all types of campus buses
   c) Easier to implement
   d) Longer lasting
   e) Properly documented (Technical and user manual for TAPS)
   f) Equipped with an ETA feature

Why is a bus tracking system necessary for UCSC?

We surveyed **288 UCSC students** in order to understand the needs of the campus' residents. We were interested to know the dependency of students on the buses to commute around the campus and which features are more desirable to the students: a digital ETA board at each bus stop or a mobile-phone application. We also cared about the effect of a reliable bus tracking system on the mental and physical health of the students and staff.

96.5% of UCSC bus users are stressed out by not knowing when the bus is coming

92.1% of UCSC affiliates use campus busses on a daily basis

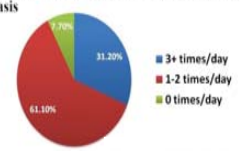

71.1% of UCSC affiliates would like to know when the night core is arriving

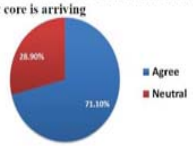

High Level System Architecture and Data Flow

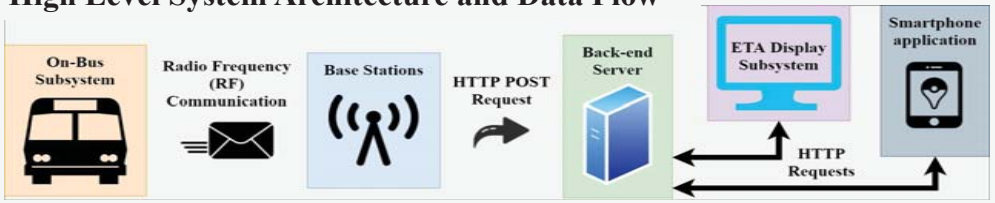

On-Bus System

Responsible for acquiring GPS data, displaying bus routes, and sending bus data to the base station subsystem.

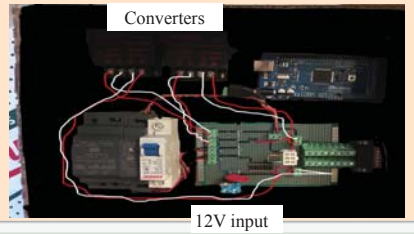

Base Stations

Responsible for receiving bus data and sending it to the server There are five base stations on campus to ensure full data coverage of the campus.

Major Components:
- **Radio** - Receives data from buses.
- **Raspberry Pi** - Processes data from the radio and sends it to the back end server.

Back-end Server

Responsible for data storage and interfaces the data-collecting subsystems with the user interface subsystems. Major components include:
- **MySQL database server** for storing bus data
- **FreeBSD web server** for interfacing with the database server

Smartphone Application*

*This is implemented by a CSE Capstone Team.
The app is responsible for providing all bus data in a user-friendly fashion featuring:
- → real-time location map
- → ETAs for each bus stop
- → direction of buses

ETA Display System

Responsible for displaying bus ETAs for the main routes such as UPPER, LOOP, OUT AT BASE.

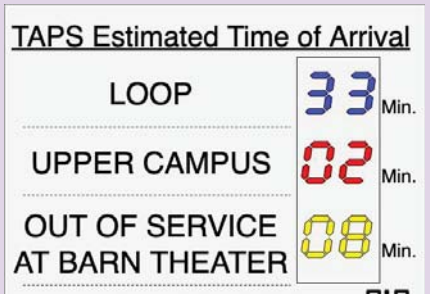

ETA Algorithm Design

There are two different approaches to the ETA algorithm by two different teams: one based the MapBox Direction API (CSE team) and one based on a machine learning (our ECE team). The machine learning method calculates the ETA by averaging historical travel and stop time data.

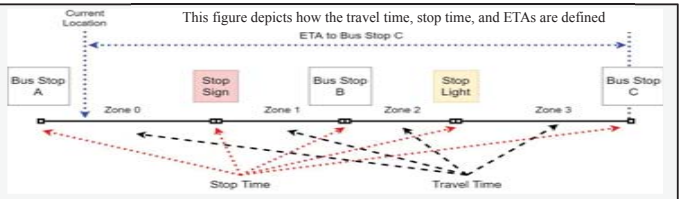

Consists of two main systems:
Travel and Stop Time Calculation and Storage:
- Travel time - the time to travel between defined locations on the route (bus stops, stoplights, and stop signs) called "stopping points"
- Stop time - the time a bus is stationary or near stationary at each stopping point

ETA Calculation: ETAs are the summation of the averaged travel times and averaged stop times based on the current zone and the destination bus stop.

Results

On-Bus System Improvements to BTS2:

- The system now operates on a **12-V** power system, allowing us to include bike shuttles and cutaway buses used for the night core route.
- The route sign cost decreased from **$800** to **$92.21+labor cost**, while still providing the main routes: LOOP, UPPER, BASE, CORE.
- The tracking part of the system now **operates independently** from the route sign. This feature is important for buses that do not need a route sign like the bike shuttles.

Base Station Improvements to BTS2:

- The system no longer uses any custom parts, allowing for easier installation and maintenance for the system.
- New micro SD cards have been chosen for the Raspberry Pi to increase the lifespan from 4 years to an estimated lifespan of more than 4 years.

Addition of the ETA Display System:

- The ETA display retrieves ETA data from the web server via WiFi.
- The ETA display can display 3-5 different routes, which can be changed by the client via the web server.
- ETA colors change from blue to yellow to red as the ETA counts down to indicate approaching buses. The three colors are distinguishable for people with any of the four types of color vision depicted in the left figure below. Arrows indicate the selected colors in each color wheel.
- If the system disconnects from WiFi, the timestamp of the disconnect is logged in the web server and a red 'X' is set on the ETA display.

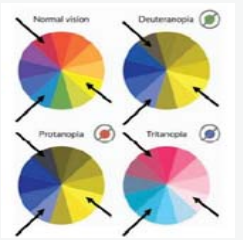

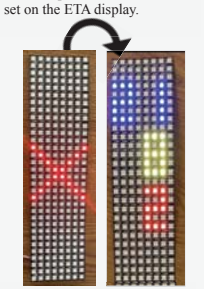

Addition of the ETA Algorithm:

The framework for the algorithm is implemented in web server and is currently collecting travel and stop time data from the BTS2 system. Proposed ETA accuracy specification was not met, but the accuracy will improve over time, as more data is collected.

Conclusion

Each system within the BTS3 design was successfully constructed according to their respective requirements. However, due to limited access to UCSC campus, we were unable to do the full system field testing. A Standard Operating Procedure (SOP) along with the Technical and User Manual will be provided to TAPS for future testing and implementation.

Acknowledgments

We would like to thank you the following people for their help:
- Kerry Veenstra, developer and implementer of BTS2
- UCSC Transportation and Parking Services (TAPS)
- Tela Favaloro, Stephen Petersen, Azzam Qureshi

SAWbots: A Platform for Affordable Minimally Invasive Surgeries

Remotely controlled miniature robot that traverses liquid surfaces with no motors or moving parts

Developed by Dr. Yanik's Research Lab in conjunction with Nic Van Oss, Allie Hunsinger, Phil Canete, Summer Alherz, and Alex Bakaleynik

Baskin Engineering UC SANTA CRUZ

There is currently a lack of minimally invasive procedure (MIP) technology in the medical industry. This type of technology significantly decreases the size and number of incisions needed to operate, which has been shown to improve patient outcomes when compared to traditional surgery. SAWbots provide the optimal solution by providing MIP technology at a decreased cost, thereby making safer surgery more accessible to the medical field. This project improves upon a novel form of propulsion in fluid in compliance with our stakeholder's requirements. The team used Yannyk Bourquin and Jonathan M. Cooper's *Swimming Using Surface Acoustic Waves* as a starting point for our research and benchmark for our results.

System Overview Block Diagram

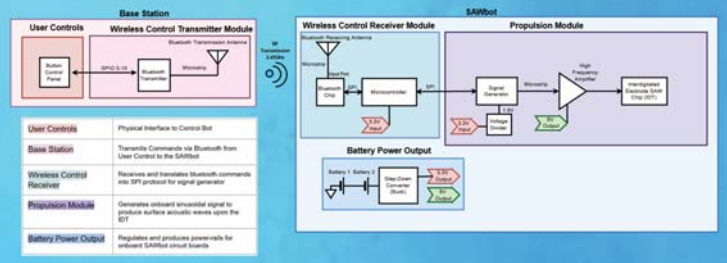

Assembly

Expanded 3D model of surface acoustic wave SAWbot

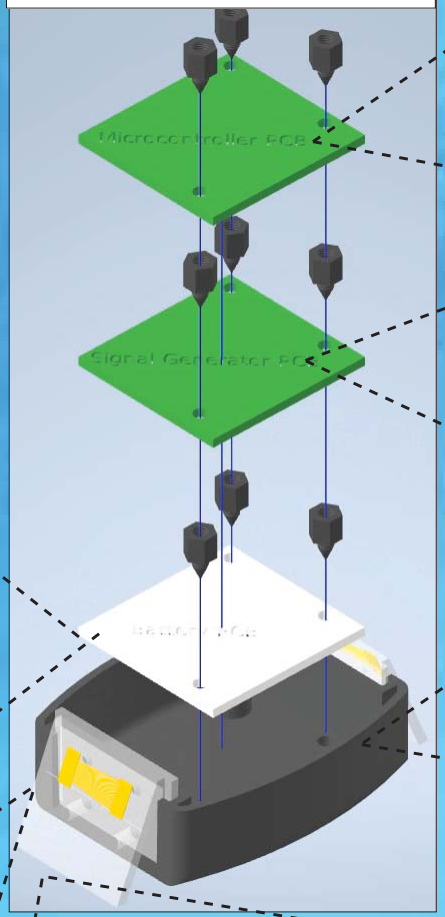

Onboard Microcontroller

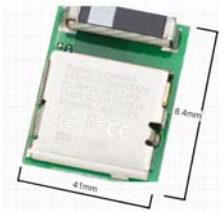

- Recieves inputs from offboard control panel
- Controls signal generator via SPI to move bot
- Bluetooth 5.1 Low Energy based communication - 10x more power efficient than WiFi

On Board Signal Generation

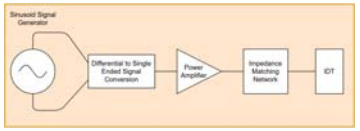

Propulsion Module Block Diagram
The AD9913 Digital Synthesizer generates a variable high-frequency sinusoid: this waveform is amplified to provide electrical power to excite the IDT

RF Power

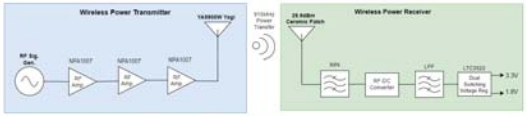

Block Diagram of Simulated Wireless RF Power System
RF Signal Amplified, Transferred to Receiver, and then Converted to DC Power

Battery Power

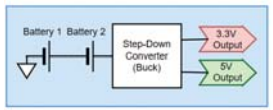

Block Diagram of Battery System
Battery Output Sent through Dual Converter to Provide DC Power Rails

Chassis

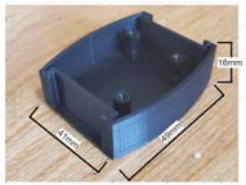

- 3D printed in biodegradable PLA
- Waterproofed using spray-on coating
- Designed for rapid, modular prototyping

IDT SAW Generation

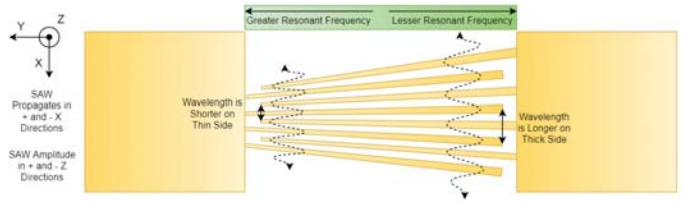

Interdigital Transducer: IDTs excite SAWs in piezoelectric substrate in response to AC voltage. IDTs with Slanted Fingers provide control of the y-axis location of the SAW beam through varying the frequency of the AC input. Thus, by varying our electrical input frequency, we can generate thrust at different locations and steer the robot

SAW Streaming Force

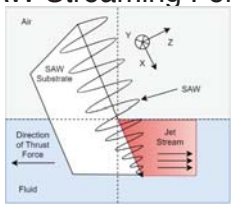

SAW Streaming: When a SAW enters a fluid interface, it attenuates and leaks energy into the fluid, generating a streaming force

Final Design Prototype:

The final prototype houses three circuit boards. These boards enable communication with the base station, provide on-board signal generation, and deliver battery power to the bot. The IDTs are mounted to provide maximum thrust. The bot is expected to generate 1mN of thrust which is 10x more efficient than ultrasonic thrusters.

Physical Bot
Isometric view of the physical bot
Length: 48.5 mm Width: 40.6 mm
Height: 16.5 mm

UV-Clean | Autonomous UV-C Sanitation Robot

Christian Deguzman, Mattiana Lang, Mary McNeil, Ethan Santos - Bioengineering Department

Current Sanitation Methods are Ineffective Against the Spread of COVID-19

- Waste hours of manual labor
- Use harsh chemicals
- Are not effective due to user error

4.48M COVID-19 cases worldwide illustrate a need for better sanitation methods [1]

An autonomous germicidal UV-C device will drastically improve our sanitation methods

Germicidal UV-C Destroys DNA of Targeted Pathogens, Rendering the Replication Process Inert

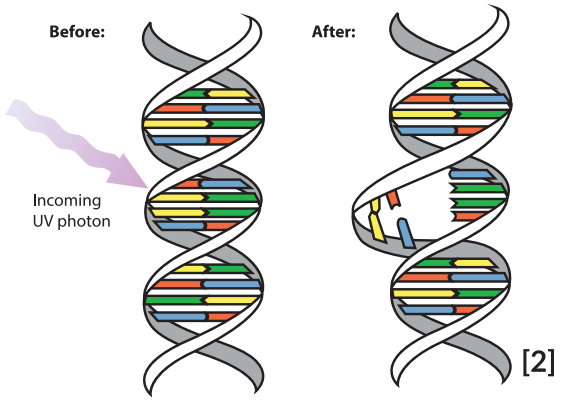

[2]

This characteristic of UV-C light provides a tool for minimizing the spread of disease

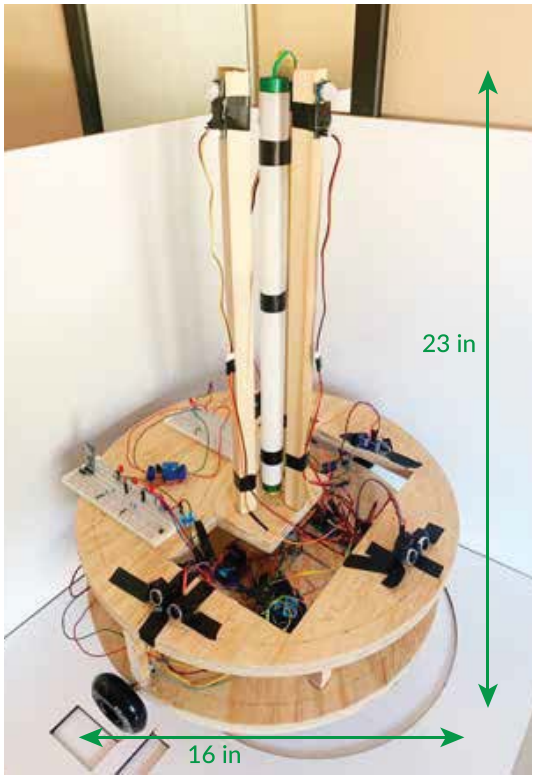

OUR SOLUTION: UV-CLEAN

- **Sanitizes 99% of common household germs: Salmonella, E. Coli, Staphylococcus Aureus, Influenza [3]**
- **Predicted to sanitize 90% of COVID-19 pathogens, based on the projected Coronavirus UV-C dosage requirements [4]**
- **Completes sanitation process without user presence or intervention**
- **Passive infrared sensors detect motion for user safety, preventing UV-C exposure**
- **Includes Android app for remote shut-off**
- **Navigates with obstacle avoidance**
- **Stationary mode for smaller rooms**
- **Lightweight (7 lbs) for easy handling**
- **Rechargeable battery powered**
- **Comparatively low cost**

Pathogen Testing Results

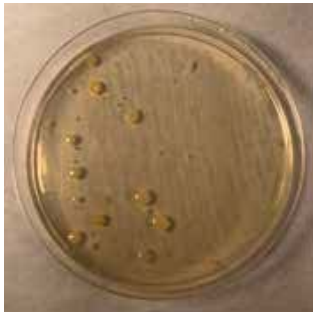

Pre-sanitation: 60 colonies

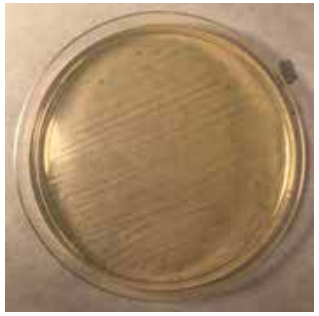

Post-sanitation: 10 colonies

[1] COVID-19 Dashboard by the Center for Systems Science and Engineering (CSSE) at Johns Hopkins University (JHU). (n.d.). Retrieved May 30, 2020, from https://gisanddata.maps.arcgis.com/apps/opsdashboard/index.html

[2] What is UVC – Ask an Expert. [Online]. Available: http://www.cisu-vc.com/ask-an-expert/what-is-uvc. [Accessed: 15-May-2020].

[3] "Application Note #12 Ultraviolet Light Disinfection Data Sheet," Clordisys. [Online]. Available: https://www.clordisys-.com/pdfs/misc/UV Data Sheet.pdf.

[4] Kowalski, Wladyslaw & Walsh, Thomas & Petraitis, Vidmantas. (2020). 2020 COVID-19 Coronavirus Ultraviolet Susceptibility. 10.13140/RG.2.2.22803.22566.

Capstone Project

RREES: SEADS Vault

Nathaniel Tjandra, Alex Bistagne, Joseph Csoti, Guangyang Chen

Abstract

When the power grid fails, the failure can be life threatening. The **Resilient Renewable Electric Energy Systems (RREES)** Lab aims to deliver power quality data fast enough to keep the failure from becoming an emergency, detailed enough to identify the failure, and comprehensive enough to locate the failure. While the **Smart Energy Analytic Disaggregation System (SEADS)** handles the collection of an overwhelming quantity of data (~5.5GB per day per device uncompressed), our team built the **SEADS Vault** to:

- identify important **events** in the SEADS data,
- **notify** homeowners of events,
- inform grid managers of event **locations**

Overview

How Outages are traditionally handled
- A user calls their provider to manually report.
- Perform a scheduled/monthly maintenance.
- Voltage is measured at the substation

How we want to change it
- Automated Event Detection and notification
- Isolate problem areas and fix faulty equipment
- Voltage is monitored at the consumer's house

All these factors provide better system health
- Monitors on the Consumer level leads to higher accuracy and can be used to link common events to faulty equipment.

Approach

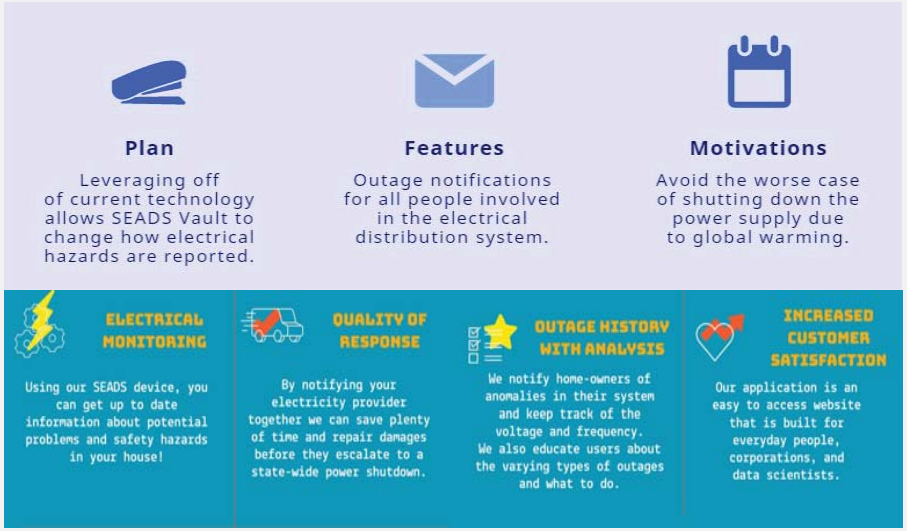

Architecture Diagram

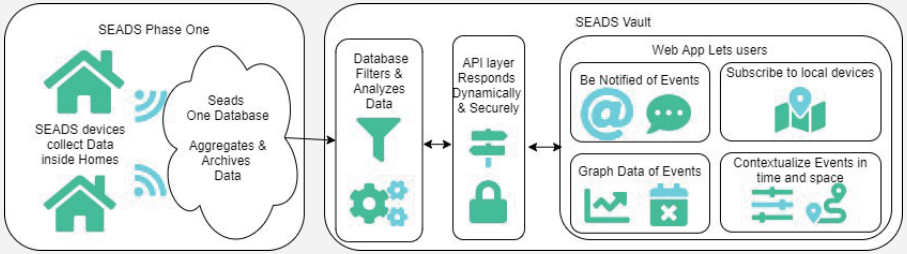

Technology used includes React, Flask, and PostgresSQL for development. We incorporated as third party services, such as Auth0, Mapbox, and Twilio; in order to handle account management, geolocation, and notifications. Finally, we used Skyhook for partitioning on each SEADS Device.

Acknowledgments

- *Brian Zhao: Phase 1 Design*
- *Patrick Mantey: Mentor and Sponsor*
- *Jeff LeFevre: Creator of Skyhook and Design*
- *Michael Choi: Building the SEADS Hardware*
- *Faeze Brahman: Managing Weekly Scrum*
- *Akila De Silva: Managing Weekly Scrum*
- *Richard Jullig: Bi-Weekly Check instructor*

Issues

Downed power lines can cause safety issues like wildfires.

Here, a tree falls down in an rural area and is not detected for a long time.

This causes a power outage and a road obstruction (at intersection of Highway 35. and Hutchinson Rd).

SEADS Vault Features

1. Have an application deployed to the web.
2. Display real time live voltage to end user.
3. Notify consumers of potential outages in interested areas via email and text.
4. Flag defective equipment for repairs ASAP to improve safety and improve reporting.

Conclusion

We hope to establish a strong foundation for future innovators to build new features using SEADS and extend the user base outside of Santa Cruz.

For instance, researchers and scientists could use the data to develop machine learning models.

Wildfire Detection Drone

Team: Antone Bajor, Janelle Chen, Skyler Ow-Bearing, Evan Plummer
Majors: CSE ECE RE BME

Baskin Engineering UC SANTA CRUZ

Problem

- In California 150k acres of land burn annually; the effects of global warming will only exacerbate this.
- The cost of damages in 2018 alone is a staggering $400 billion.
- Wildfires can be better suppressed if they have just started; they exponentially develop into larger fires that can become hard to contain.

2018 Camp Fire in Butte County

Current Technologies

Fire Spotting Plane

Butte County Lookout. Smoke and fire detection cameras.

- Helicopter and airplane patrols are able to survey remote areas, however each flight costs over $4,000
- Fire lookout towers with automated camera systems incur low upkeep cost, yet they're immobile
- Satellites scan massive areas, however they are not useful for early detection

Our Solution

- High resolution IR imaging from a low altitude allows for the early detection of small fires (5'x5') while RGB is useful for verification by an operator
- Autonomous flight leaves no room for human error, the exact region of interest is patrolled, and the collection of data is consistent for every flight

Our development drone with fire detection payload

IR image of a 5'x5' charcoal fire hovering at 350' Above Ground Level (AGL)

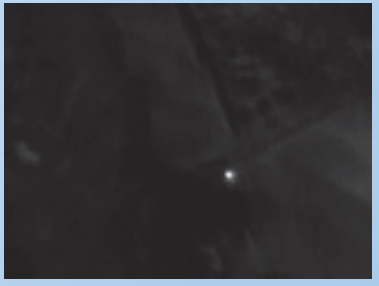

Key Benefits of Our Design

Cost:

- The cost of a new system is under $5,000 and each flight is the price of charging a battery
- A modular system means that maintenance is simply the replacement of broken parts
- Early detection of smaller fires leads to lower suppresion costs and less destruction

Automation:

- Hands free operation frees up first responders for other duties
- Because an operator is unnecessary, beyond line of sight operation is possible
- A competent flight algorithm reduces probability of complications such as crashing

Accessibility:

- Low learning curve allows firefighters to properly deploy the drone with minimal training
- Can independently deploy the drone
- Easy transportation of the drone allows surveillance of any area near a road

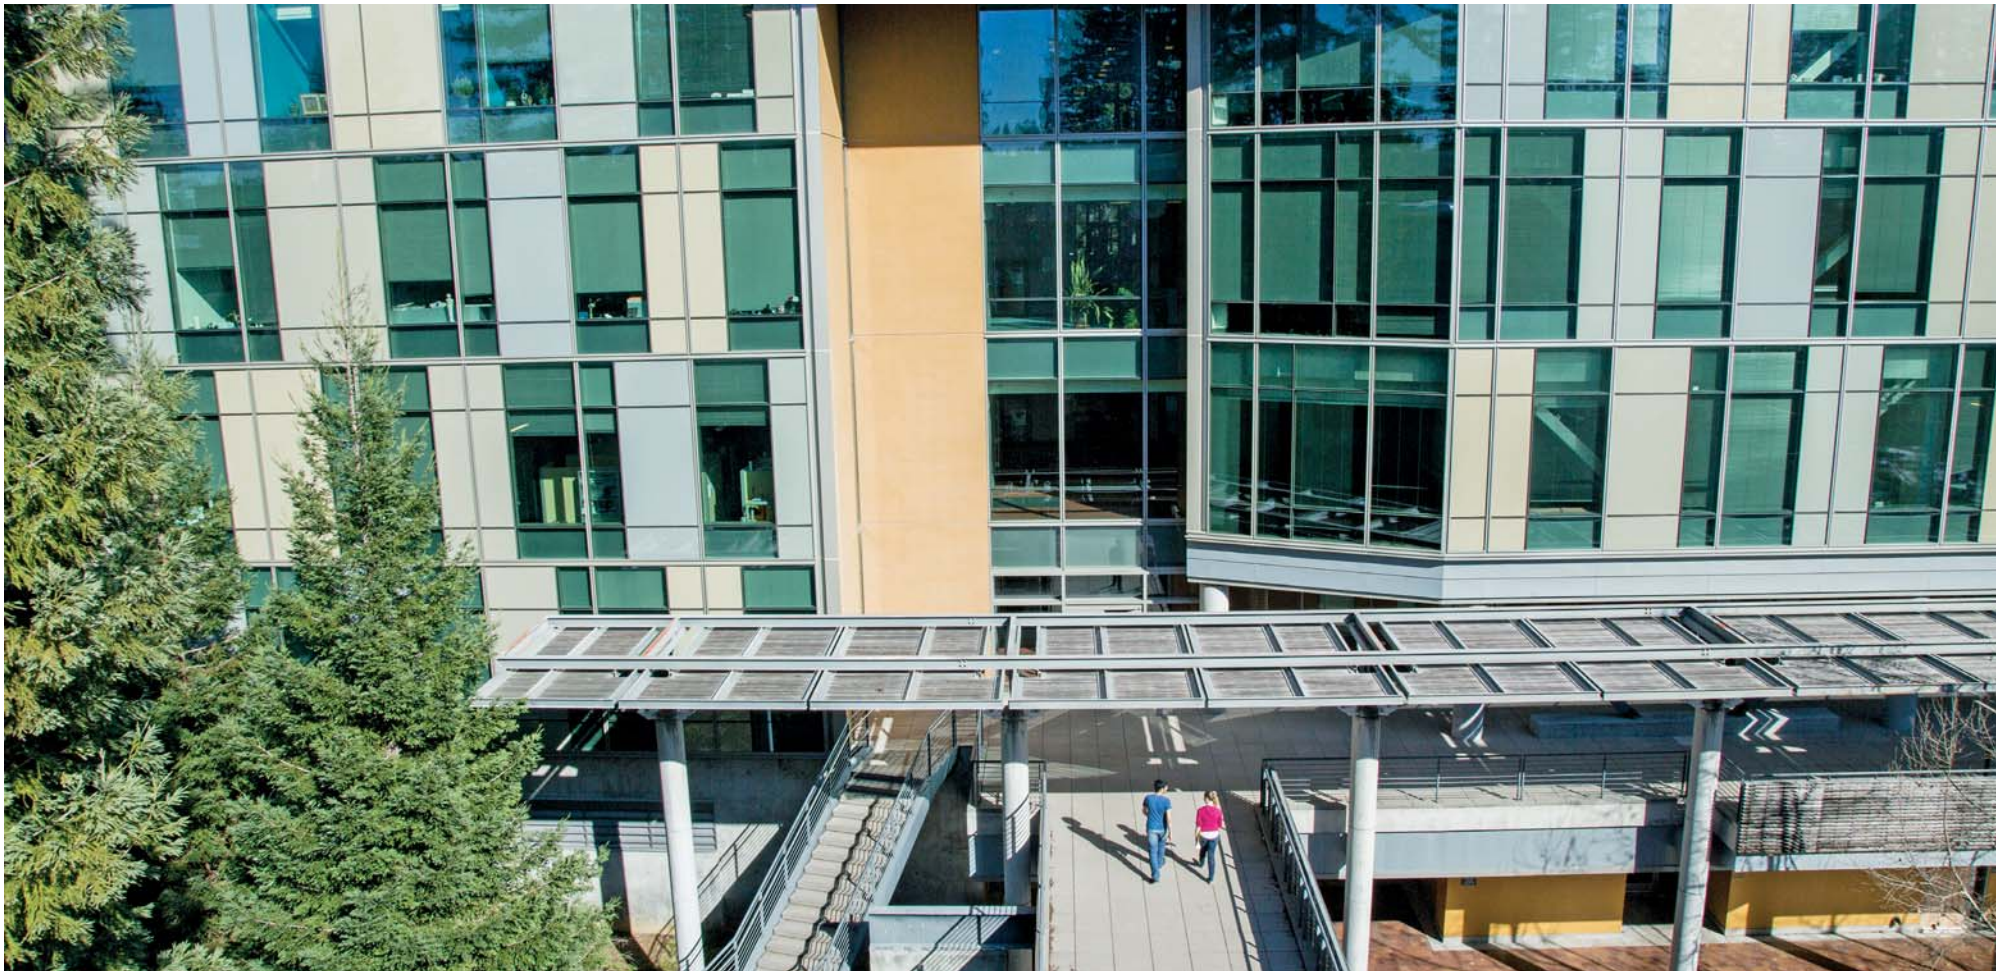

Contact us

Patrick Mantey
Director, Senior Design Capstone

Jack Baskin Endowed Professor, Computer Engineering, Emeritus

mantey@ucsc.edu

Frank Howley
Senior Director of Corporate Development

fhowley@ucsc.edu

Visit our website
csspp.soe.ucsc.edu

Baskin Engineering
UC SANTA CRUZ

facebook.com/BaskinSchoolofEngineering
twitter.com/UCSC_BSOE
instagram.com/ucscengineering
LinkedIn: University of California, Santa Cruz - Jack Baskin School of Engineering