Captivates: A Smart Eyewear Platform for Ambulatory Physiological Measurement Capture

by

Patrick Chwalek

B.S., University of Illinois (2015)
M.S., Georgia Institute of Technology (2018)

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
in partial fulfillment of the requirements for the degree of

Masters of Science in Media Arts and Sciences

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2020

Author .....
Program in Media Arts and Sciences,
School of Architecture and Planning,
September 1, 2020

Certified by .....
Joseph A. Paradiso
Associate Academic Head
Alexander W Dreyfoos (1954) Professor of Media Arts and Sciences
MIT Media Lab

Accepted by .....
Tod Machover
Academic Head
Professor of Media Arts and Sciences
MIT Media Lab

Captivates: A Smart Eyewear Platform for Ambulatory Physiological Measurement Capture

by

Patrick Chwalek

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
on September 1, 2020, in partial fulfillment of the
requirements for the degree of
Masters of Science in Media Arts and Sciences

Abstract

Smart environments amplify the lived experience within a space by learning and adapting to the preferences of the users within them. Unfortunately, these preferences are dependent on internal cognitive states of users, a characteristic that is not available to be sensed robustly across contexts. In addition, many sensing systems that are designed to sense users within specific contexts but often fail to adequately address privacy concerns. Captivates are a wearable eyeglass platform that attempts to fill this gap in sensing, allowing for the capture of physiological signals that have shown correlation to specific cognitive states. The system was designed with a few diverse sensing technologies that as a whole, can create a robust cognitive sensing architecture that won't impede a user's experience. Captivates are designed to be used across contexts and throughout the day, collecting samples of your physiological responses that can better inform individualized models of a user's internal states, leading to a more responsive environment that takes the human element into consideration.

Captivates are able to continuously sample and stream data through a wireless network that is mesh-capable, allowing for applications that require scale with limited infrastructure. The sensing modalities include face temperature, blink rate, head pose, and 3D location but can be extended or reduced for other applications. The system was built using recommendations from eyeglass manufacturers on how to create a robust smart eyewear system that shares the aesthetic of more traditional eyeglasses. The system was initially designed to aid in data collection for cognitive modelling efforts but is also a platform for others to use for similar or orthogonal researcher efforts. Apart from the sensing channels, several touch points and LEDs exist on the device to allow for applications that require user input or actuation.

Thesis Supervisor: Joseph A. Paradiso
Title: Alexander W Dreyfoos (1954) Professor of Media Arts and Sciences

This thesis has been approved by the following committee members:

Reader .....

Cynthia Breazeal
Associate Director
Associate Professor of Media Arts and Sciences
MIT Media Lab

Reader .....

Sanjay Sarma
Vice President for Open Learning
Professor of Mechanical Engineering
MIT

Acknowledgments

First and foremost, I wish to thank my advisor, Professor Joseph A. Paradiso. Without the liberties and breadth of sensor knowledge offered under his direction, this project could not have been completed to this degree. I am grateful for the opportunity of being part of his group and of the greater MIT Media Lab. I would also like to thank my readers, Professor Cynthia Breazeal and Professor Sanjay Sarma for all their advice as I progressed through this project.

A special acknowledgment goes to David Ramsay who I started this project with and have continued to work with ever since. His perspective on cognition modelling and his drive to push design constraints has shaped this project to what it is today. David primarily worked on the mechanical design of the glasses while offering insight into the electrical design and analysis. I could not have wished for a better collaborator to work with throughout this project.

I additionally like to thank the members of Responsive Environments for aiding me throughout this project. Specifically, thank you to Brian Mayton and Mark Feldmeier for tolerating all of my electronics-related questions and helping in debugging some of the hardware for this project. Additionally, most of the early stage work was done for and during the MIT Media Lab's Research at Scale Program so I would like to thank the program's organizers, Andrew "bunnie" Huang and Jie Qi. Both of them spent a considerable amount of time planning the entire trip and finding us relevant factories to tour and collaborate with.

I would like to thank my girlfriend, Beata Bednarczyk, who has been with me far longer than the duration of this project and has offered me support throughout. I cannot imagine being able to complete this project without her, especially through the pandemic crisis of 2020.

Lastly, but definitely not least, I thank my friends and family, who helped get me past all the stepping stones to become who I am today.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Environments are described as smart when there exists at least one device that amplifies the experience of being within them. This can take the form of a voice assistant that can respond to your queries, a thermostat that learns your average temperature preference on a minute-by-minute resolution, or even a lighting system that actuates to predefined settings. These systems learn things about you, e.g., your voice patterns, average work schedule, bedtime trends, and/or even some of your specific preferences. However, what they don't learn is how your preferences might change depending on mood. The ideal preferred temperature in your house, lighting levels and colors across spaces, and maybe even actuated scents will all be recommended by these smart systems, perhaps knowing if you had a stressful day at work or spent the last hour enjoying a walk with friends. We need smart environments to not only have temporal resolution but also cognitive resolution. Now how do we go about building this resolution through sensing? The only robust way is for long term sensing across contexts which would allow for a preference model to be built that can expand the resolution of these smart environmental systems, making them more sensitive to the human element.

When browsing academic articles and consumer products that have the potential for this type of sensing, not many systems exist that can be procured, are transparent to how they work for simple system integration, and/or robust enough to be used throughout the day. What is needed is a device that can measure you and your lived experiences but does not impede your natural way of living. Take for instance SenseCam [7], a simple device made by Microsoft to aid in memory rehabilitation. The device takes snapshots through a neck worn device of dynamic events in front of a user that can later be reviewed as a summary of one's day. Although large by today's standards, the device was capable of being used throughout an entire day,

and the summary that it could provide is the type of data that would be useful for a system that is attempting to estimate your internal state. However, while it is true that a system like this, coupled with frequent emotion self-assessment questions, can inform and shape a preference model, it would be far less socially invasive to flip the sensing around— sense the user's changing physiology instead of the stimulus and allow them to self-assess how their day is going.

Many academic studies exist that use single sensing modalities and controlled stimuli to measure physiological responses; you can find measurable responses in eye blink behavior, temperature across the body, body movement, and electric potentials across the skin. The way of measuring these signals isn't novel but creating a platform that can do it robustly and comfortably is. We need to understand how these signals change out in the dynamic world, and the only way to do that is to have a sensing system with you at all times. And to do that, a device needs to be made to blend in with the user, so as to not impede on the lived experiences and become burdensome. As far as we know, no such device exists that is readily available for researchers to use. As mentioned in Section 2.2, many studies in academic journals are either using custom systems that are not able to be taken out of a lab setting or commercial systems that are just not tailored for long-term, across-context use. Of the systems that advertise they can be used throughout long days, the sensing modalities are just not robust enough to the dynamic motion that comes about from everyday life.

Smart eyeglasses are a popular place to put sensing technology and have become ubiquitous in the tech world, from wearable displays to heart monitors. They also have become popular in the research community, because they are an ideal place on the human body for continuous monitoring of environmental and physiological signals. Unfortunately many smart eyeglasses are designed with an engineering-first design, often leading to bulky implementations that most people would not find comfortable to wear for long durations. This limits a researcher's ability to conduct naturalistic, long duration studies. Further, available smart eyeglasses do not allow researchers to connect their own devices (e.g., sensors) to the platform for data exfiltration and signal processing. Instead, researchers must create their own custom sets of glasses— a large time investment for an artifact that may only last the duration of the experiment.

It's also important to consider a user's privacy in the design of such wearable systems that are intended to be worn for long durations. These systems not only need to be physically comfortable but also transparent in what data is being recorded, how it is being stored, and who it is being shared with. For commercial systems, its fairly common for companies to use the stream of data from their current users to

improve their algorithms but this data can also be used for reasons that the users may be unaware of (e.g., targeted advertising). As described in Section 2.2.3, there are a few systems that currently exist that raise a few privacy concerns that erode the trust between the users and the devices that are intended to improve their overall day-to-day experiences. Without this trust, it's impossible to create systems that have a symbiotic relationship with their users.

In this work, we seek to create a system designed for ambulatory physiological signal monitoring for use in creating better-informed responsive environments. We also seek to create a system that other researchers are able to tailor for their own applications. In chapter 3, we describe our design considerations in making a scalable, extensible, comfortable, and manufacturable system in a traditional eyeglass form factor that inconspicuously blends in‘to the user's environment. In chapter 4, we discuss a preliminary sensor validation trial and potential promise of using this system to create an individualized model of a user's cognitive state.

Chapter 2

Literature Overview

2.1 Non-contact Physiological Measurements related to Cognitive State

Using physiological sensor data to estimate a user's cognitive state is not a novel concept, but in research it is often done in very controlled research settings and rarely done in-the-wild. We seek to explore the space of sensor systems that do not require direct contact to the skin, a physically less-invasive technique that results in a more natural user experience and a more robust system for measurements on-the-go. Contact-based sensor techniques are a popular method for measuring physiological signals but the use of these techniques is not ideal for measurements in dynamic environments. A few contact-based techniques for measuring internal electrical activity are electrocardiogram (ECG) for measuring the heart, electroencephalogram (EEG) for measuring the brain, electrooculography (EOG) for measuring eye movement, and electrodermal activity (EDA) sensing for measuring changes in skin conductance. All of these methods have been related in some way to cognitive state estimation [8, 9, 10, 11, 12] but in many of the papers, these measurements are taken in controlled settings and oftentimes, researchers struggle with noise artifacts from internal and external sources [13, 14, 15]. More importantly, because of the mechanical skin-electrode coupling, a large source of noise is that resulting from user motion, which makes this type of sensing technique non-ideal for measurements throughout a person's dynamic day [16].

2.1.1 Face Temperature

As summarized in table 2.1, several studies have found that our face temperatures vary with internal state changes (e.g., fear, joy, anxiety, etc.). Specifically, in [17], it was observed that nose temperatures decrease with the increase of cognitive load due to blood flow restriction when autonomic nerve activity increases. In that study, cognitive load was controlled by a series of reading tasks of different difficulties and Stroop Tasks (i.e., matching color of a word with the word itself) at varying speeds. However, the apparatus in that study used to collect these measurements was an expensive thermal camera that required the user to be a fixed distance away from the screen. This is ideal if the stimulus will always be located in a single location but for mobile use, a fixed setup requires multiple thermal cameras, is prone to occlusion, and is cost prohibitive.

Emotions								
	Stress	Fear	Startle	Sexual arousal	Anxiety	Joy	Pain	Guilt
Regions								
Nose	↓	↓		↑		↓		↓
Cheeks			↓					
Periorbital			↑	↑	↑			
Supraorbital			↑		↑			
Forehead	↓↑	↓		↑	↑		↓	
Maxillary	↓	↓	↓				↓	↓
Neck–carotid			↑					
Nose	↓							
Tail		↓					↓	
Fingers/palm		↓					↓	
Lips/mouth				↑				

Table 2.1: Overview of the Direction of Temperature Variation in the Considered Regions of Interest Across Emotions [6]

An alternative is to fix a face temperature sensing device onto the user, similar to [18]. In [18], the researchers fixed a passive infrared radiation sensor (i.e., thermopile) to a pair of glasses and measured the nose and forehead temperatures— the forehead temperature was used as a ground truth since it was thought to remain stable to internal temperature fluctuations. A ground truth is important because face temperature varies from internal and external events (e.g., convection due to wind). However, in [17], the forehead temperature was found to increase with increased cognitive load, so a more suitable reference point would need to be found that is robust to fluctuations during cognitive loading tasks. In [19], a patent is discussed on using a thermopile to measure the temple region on the face, since it was found to be a great estimate of a person's internal body temperature due to the large superficial temporal artery. Such technology is used in hospitals since it was found to be close to the accuracy of, and far less invasive than, rectal thermometry: the most accurate method for measuring

internal temperature in a clinical setting [20].

2.1.2 Blink Rate and Eye Gaze

As described in [21], there are four types of eye blinks: reflex blinks, voluntary blinks, non-blink closures, and endogenous blinks. Reflex blinks are involuntary responses to potentially injurious stimuli while voluntary blinks are self-initiated. Non-blink closures are slow closures of the eye for events such as the onset of sleep. Endogenous blinks are different than the others due to the "absence of an identifiable eliciting stimulus." These types of blinks are theorized to relate more to "a broad spectrum of information processing variables and both general and momentary task demands and can reflect the effect of these during the performance of a cognitive task."

To support this theory, in [13], test subjects are placed in front of a computer and the visual and auditory stimuli are modulated with a quiet, visually-blank stimulus interleaved for a baseline. The findings of that study was that with a stimulus that increased cognitive load, a decrease in blink rate was observed. Furthermore in [22], users were given arithmetic tasks at varying difficulties while their blink rate and galvanic skin response (GSR) were measured; GSR has already been shown to relate to cognitive load [23]. It was found that eye blink rate over a task showed to be indicative of task difficulty, with blink rate decreasing as the difficulty of the task increased.

Given how changes in blink rate can be related to changes in cognitive load, it would be valuable to incorporate eye blink sensing into a wearable that can be used with other sensing modalities for robust load estimation. It's not common to find blink rate sensors on existing wearables, but there are a few reported in research. In [24], the authors use a near infrared diode to illuminate the eye and a phototransistor sensitive in that spectrum to detect changes in reflectance when the wearer blinks. Interesting enough, you can also detect the severity of the blink (i.e., hard vs soft) to open a second feature dimension for exploration.

2.1.3 Head Motion Dynamics

One way to measure blink rate and face temperature robustly while the user is freely able to move is to design a wearable system that can be worn on the face. This conveniently also allows for adding an inertial sensor to measure head pose and for user activity prediction. In [25], a 20-participant study showed that head movement is "indicative of cognitive load and discriminative between different task types, as

well as exhibiting some sensitivity to the instant of task change." For that experiment, the tasks varied from watching a stimulus to physically writing on a piece of paper, requiring very different physical dynamics. The researchers found that higher cognitive loading tasks resulted in less movement which reflected "more unconscious movements related to thinking about the problem."

Apart from directly linking head dynamics to cognitive load, in [26], it is shown that gaze direction can be estimated by just knowing head pose. The researchers found that "head orientation was a sufficient indicator of the subjects' focus target in 89% of the time". Given some a priori knowledge of the environment, it would be possible for a system to predict where a person's gaze is fixated by just knowing their head pose and 3D location together with head movement dynamics; traditionally, this is usually done using higher computationally demanding systems, such as camera-based eye tracking.

2.2 Collecting Physiological Data for Cognitive Analysis in Non-Laboratory Settings

2.2.1 Daily Life

As mentioned before, contact-based measuring techniques for cognitive state estimation are a popular approach, especially for the consumer market. Several products exist that claim can help improve your attention, learning, and aid in relaxation. Muse [27] is an EEG device that is designed to help train in meditation, but claims it "makes it easy to access and use brainwave data, inside and outside the laboratory and in real world environments." Their device is a bit more practical for long-term use by using dry electrodes that do not require additional conductive gel for better coupling [28] but that still doesn't solve the issue of motion artifacts. Regardless, MUSE no longer supports access to the raw data, and third party applications are required to record the raw signals [29]. There are more open source methods of collecting raw EEG signals, but the designs aren't suitable for on-the-go applications [30].

There are several commercial smart eyewear devices that are packed with a suite of sensors for researchers to use. The Google Glass would have to be one of the earliest smart eyeglasses that have been heavily used by researchers across the academic spectrum, from analyzing head motion and blink frequency for activity recognition [31] to augmented-reality based indoor navigation [32] to even surgical applications

[33]. They were popular because they were one of the first smart glasses that were released that could survive a user's regular day while also providing an SDK for researchers to develop with. However, wide scale consumer adoption of Google Glass never really happened because the design was technology-first and not socially acceptable to wear— they looked too much like smart eyeglasses. So, what's needed is a pair of smart eyewear that enables physiological data collection over an entire day over multiple subjects while focusing more on a glasses-first design to create something that's socially acceptable to wear. A few recent examples exist in the consumer space that fall into a more glasses-first design direction. Vue [34] is a pair of glasses for activity tracking that offer wireless bone conduction audio for discrete listening. They are designed to look like your average pair of glasses with all the electronics densely embedded within the plastics. Focals by North [35] takes a similar approach by offering activity tracking but also offer a projected display onto one of the lenses for notifications and short information intake. Unlike Google Glass, the display is not visible from the outside since the projection is done on the internal face of the glasses, allowing the glasses to retain a traditional design aesthetic. Both of those systems are sensor-packed and follow a more consumer friendly design approach but, unfortunately, offer no ability for researcher's to grab any of the raw sensor signals or to interface their own sensors.

There are various open-source hardware acquisition platforms in literature that exist. In [36], a device is presented that "contains an electrocardiography circuit, two electrodermal activity channels that implement a time-multiplexing technique to remove inter-channel coupling interference, and a skin temperature module." This design uses a non-contact radiation measurement for temperature and then electrode-based techniques for ECG and EDA. There results are promising with the caveat being that their system is a single circuit-board that isn't designed to be worn and is tethered to the computer so its more of a tool to be used in the lab, similar to OpenBCI. This is often the case for most of these platforms that exist in literature— if they aren't commercialized, there is little incentive to complete the design for a robust wearable for other researchers to use long-term.

2.2.2 Education Market

A popular application is collecting physiological data to measure attention and cognitive load in education settings to adapt content for more user-centric learning. For online education systems, only one such system currently exists but it is very much

in its infancy with little information on its effectiveness [37]. Nestor, the system in question, claims that its camera-based system can notify the user if they are being inattentive and trigger quizzes at that point so that students are incentivized to pay attention. However, what is less known is how the machine learning models of individuals is formed, since they claim that "machine learning can use data from social network." This hits a privacy concern since Nestor, although unclear, can potentially be mining user's personal data. What Nestor also doesn't do but is valuable for an online learning system that can quickly adapt to the user is affect classification. In [38], it was shown that a child's displayed affect was correlated to their vocabulary learning but the exact relationship varied by context and the type of interaction the child had. For an online education system, you currently have the ability to control the learning context as displayed on the screen, but it's feasible to imagine a future system being able to tailor a user's environment to create a more suitable learning environment. To get there, we first need to understand more internal metrics of the users to be able to adapt the content appropriately.

In live classrooms, there are a few examples of real-time user sensing systems. EduSense is a perfect example of a system that was successfully used in uncontrolled live classrooms [39]. The platform, created by Carnegie Mellon University, seeks to improve classroom instruction by giving the instructor real-time feedback on their performance and the aggregate classroom response. They achieve this by creating models of engagement in classrooms using two physically non-invasive sensing modalities: acoustic and visual. The data is pushed to a server where real-time classification is performed on the subjects in a classrooms, predicting various activities that relate to engaged classrooms (e.g., hand raised, sit vs. stand, smile detection, student vs. instructor speech, etc.). This data could then be fed via a dashboard to the instructor, giving valuable live statistics of the classroom without signaling any one student.

A similar system is being used in Chinese classrooms to analyze the behavior of students and for attendance keeping [40]. The camera in the classroom is designed to "log six types of behaviors by the students: reading, writing, hand raising, standing up, listening to the teacher, and leaning on the desk. It also records the facial expressions of the students and logs whether they look happy, upset, angry, fearful or disgusted." In regards to privacy, the principal of one of the schools that has the system said, "data collection and application focus on behavior management of the entire class rather than any individuals." Similar to EduSense, the metrics reported to the instructor are an aggregate of the classroom.

2.2.3 Data Privacy

For systems that analyze you pervasively in your own home and other settings, they have to address user privacy concerns to ensure the users are in full control of their own data and nothing about the users' unique identity can be discovered. If proper safeguards are not created, it becomes difficult to persuade wide user adoption and creates vulnerabilities that could be exploited by bad actors. For EduSense and the China-based system, the authors addressed privacy concerns head-on by not transferring the data off-site by using local servers, not storing any video data, and anonymizing the abstracted feature set [39, 40].

Currently, there is a another system that is being integrated into live classrooms that seeks to aid instructors on measuring teaching effectiveness and to find out which students are not engaged. BrainCo, a company that spun out of the Harvard Center for Brain Science, has created a wearable that uses EEG to find patterns in specific regions in a user's brainwaves that relate to attention. This wearable has an externally visible on-board light that has its color mapped to the users attention level. This attention level is also transmitted to the teacher via an application that runs on a local PC. Their device is currently being used for various research applications but is also the subject of a large pilot study in China. As The Wall Street Journal reports, teachers do see a rise in engagement in the classroom but there is a fair amount of skepticism on how effective the science is (i.e., if the observed impact is just a placebo) and if and with whom the data is shared [41]. The increased criticism has caused the device to be retracted from at least one school in China [42].

Chapter 3

System

3.1 Design Considerations

In this section we describe the design considerations that guided the Captivates implementation. These were formed mainly by our experience in creating wearable systems, but also by occupant feedback we obtained through multiple design iterations and visits with manufacturers, ranging from plastic injection to eyeglass assembly. We describe them here to illuminate on the factors that shaped our design.

3.1.1 Form Factor

When initially designing the system, we conceptualized a variety of on-body sensing form factors. We quickly converged to having a sensing system on the face, given the various sensing approaches for on-the-face physiological measurements and the amount of real estate that a device worn on the face would give us. Naturally, the most ideal form factor would be a pair of smart eyeglasses, since this offers us the space to add a comprehensive hardware suite while also giving us the ability to apply sensing anywhere along the footprint of the glasses. The drawback in trying to create a custom pair of glasses is that not all people would be able to wear them, including users that currently wear prescription glasses. We explored the feasibility of creating a clip-on device to existing eyeglasses, but found it to be non-ideal given the amount of hardware we wanted to integrate and the limited space that method offers. An alternative would be to allow for prescription lenses to be inserted into our eyeglass frame, similar to what Focals by North has done [35]. Therefore, we will try to design to allow for this adaptation, but it will not be a focus for this initial system construction.

3.1.2 Scalablity

Given the amount of work that has been put into designing this system, we want it to be able to join and cooperate with a wider network of existing or future sensor systems and actuators. To do this, the system needs to be fitted with a radio transceiver and support multiple radio protocols for easy integration into different network architectures. Another requirement is to have the system be able to join a network with other nodes and have enough resources on-board for it to serve as an active node in the system, allowing for it to mediate between nodes. We would like the system to be capable of joining mesh network architectures for a more dynamic experience as a user traverses through a space. The constraint of meshability is not a strict one, but one that we believe to be advantageous for future Internet-of-Things (IoT) applications in spaces without existing router infrastructures.

3.1.3 Extensibility

We want the system to be extensible for other researchers to use for their own direction without needing to redesign an entirely new pair of glasses. Therefore, we seek two add the functionality for a user to easily program the device without needing to electrically rework any of the system for this will allow for a wider breadth of individuals to be able to fully utilize this device. We also seek to add enough connectivity so that others can create and integrate their own sensing platforms as either an extension to our sensing suite and/or use our existing infrastructure for computation and networking.

3.1.4 Design for Comfort and Signal Robustness

Since this system is intended to be worn during static and dynamic activities (e.g., sitting, walking, etc.), we want to ensure the system is both comfortable to wear and uses sensing modalities that have a higher chance of performing well in most environments. Ultimately, this means using techniques that don't require a significant amount of force or adhesion to couple to the user's skin for sensing, since this won't allow the system to be sustainable for long-term use. Further, electrode-based techniques that require this type of coupling often have reduced signal integrity in long-term, mobile applications. This is due to the dynamic stresses applied to the mechanical bond of the electrode-to-skin contact during movement and on the various internal noise sources, for instance, cardiac activity, ocular movements, eye blinks and muscular activity [43, 15, 44, 14, 45].

3.1.5 Manufacturability

We need to design the system to be easily produced and assembled while balancing the appearance and comfort of the glasses. This requires a design that balances the amount of individual pieces and overall assembly time (e.g., avoid soldering wires, using many fasteners, having fragile components that can break easily if assembled too forcefully, etc.) while not making it difficult to modify for other applications. Our goal is to have a system that is able to be made through processes geared for mass production, allowing us to cheaply produce batches of plastics and assembled printed circuit boards. In addition, we also want the design to be 3D-printable to allow for users to perform any modifications and print their own iteration while not being limited to minimum-order-quantities (MOQ) or tooling costs at an injection molding facility.

3.2 Design Exploration

As part of the MIT Media Lab's Research at Scale Program, we travelled to a variety of manufacturing and assembly facilities in Shenzhen and Seoul during the summer of 2019. The goal of the trip was to see firsthand how products are manufactured and how far we can push design constraints. The capabilities we toured included plastic injection molding, die casting, printed circuit board (PCB) production and assembly, and eyeglass production and assembly.

3.2.1 Plastic Injection Molding

For plastic injection molding, we toured a few facilities in and around Shenzhen, China, that were geared for manufacturing at large quantities (10,000+). Injection molding was of interest to us, since if we wanted to make over 100 pairs of these glasses to use and distribute to other researchers, simply 3D printing the batch would be cost and time prohibitive. Therefore, it would be advantageous for us to design a pair of glasses that can be injection molded so that if we wanted to create a large batch, we can invest in a tool to then make the present batch and any future batches— a tool is the metal mold, usually aluminum for low volumes and steel for high volumes, that the liquid plastic flows into and solidifies in.

We consulted with a variety of injection molding facilities, including ones that specialize in eyeglass manufacturing, on the design of our glasses and on which type of plastic material to use. Interesting enough, we learned TR90 Nylon is a popular

plastic to use for eyeglass frames due to its low density, flexibility, and durability. To demonstrate, one of the manufacturers grabbed an eyeglass frame made out of TR90 and bent it 90-degrees for us to see that once released, it returned back to normal without any deformities in the shape, color, or texture. Another manufacturer told us that he imports his TR90 from Switzerland, since its known to be of the best quality, a fact we could not validate but have heard more than once. From these conversations with the manufacturers, it was clear that TR90 would be the most robust choice as a housing material for our electronics.

3.2.2 Circuit Board Design Iterations

When we first conceived the idea and settled on the form factor, we thought about making a pair of eyeglasses that was flat on all faces for a square-like design. This would allow us to use a traditional rigid circuit board for the entire front without any plastic housing in order to reduce costs and complexity— a prototype of this form factor is shown in Figure 3-1. However, after consulting with a few eyeglass designers in both Shenzhen and Seoul, they strongly recommended us to redesign the frames if we really wanted to make this system as a platform for others to use, especially for measurements in public settings. Their justification was that, although our original design was cheaper to manufacture, it wouldn't be comfortable physically or socially to wear, making it difficult to collect data across a variety of users throughout their natural day.

After much deliberation, we decided to attempt a redesign where instead of making the front rigid, we would design a flexible circuit that runs across the curved face of the glasses, allowing for a traditional eyeglass aesthetic. This method would also reduce the weight of the glasses since FR-4, the traditional core of a rigid PCB, is denser than the TR90 that is required to support the flexible circuit and complete the front housing of the glasses. This method would increase the mechanical and assembly complexity since we would require more piece parts in our build, but this was a trade-off we thought was worth given the recommendations.

3.2.3 Smart Eyewear Design

From our time with a few eyeglass manufacturing facilities, we were able to pickup a few design strategies and pinpoint failure points of current smart eyeglass designs that we can try to solve for early on. Throughout this investigation, they walked us through several of their designs, and we disassembled a variety of glasses.

Figure 3-1: Initial Design of Glasses with Three Rigid PCBs

Hinged Design

There are numerous smart eyeglasses that exist, but the designs we focused on are ones that have electronics running through at least two faces (i.e., the front and either arm). In this type of design, electronics need to run through a hinge, assuming the entire design isn't rigid. This usually manifests itself as a set of flexible wires or a piece of flexible circuitry. A major issue with these designs is that the wires or flexible circuit are being subjected to dynamic bending and will experience fatigue as the glasses are repeatedly opened and closed, leading to an eventual system failure. The most ideal scenario would be to find a hinged electrical connector that can pass signals between circuit board and not require any circuitry to be repeatedly bent. We managed to find only one connector of this type and it was made by Molex (Figure 3-2) but we couldn't find it on any online distribution websites. While touring Shenzhen's electronic shops, we were told to talk to a specific shop owner whose speciality was knowing about and selling all types of connectors; a picture of his stall overflowing with connector reels is shown in Figure 3-3. After consulting with him and having him contact several factories that make Molex connectors, we unfortunately found out that this connector, or anything similar, was no longer in production and hasn't been for a long time.

Given that we couldn't get the connector, our alternatives were to either solder wires between the electronic subsystems or to have a flexible circuit run through the

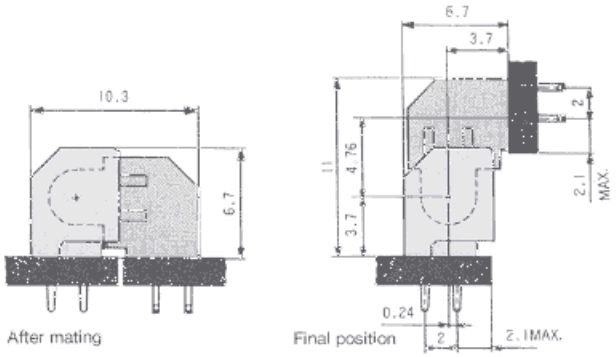

Figure 3-2: Molex Hinged Connector

Figure 3-3: Shenzhen Connector Shop

hinge which carries all the signal and power conductors. Since our system was going to have a variety of sensors, resulting in numerous signals and power lines running between the PCBs, soldering individual wires wasn't a practical option. Therefore, we started to investigate design constraints for flexible circuitry in dynamic bending environments. We first started to disassemble various smart eyeglasses that we could find in the Shenzhen markets and within the factories that make them. We quickly found that most of the hinged designs had two pieces of plastic that were press-fit into each other with a piece of flexible circuit running somewhere between (Figure 3-4). The interesting note was that these designs did not follow the recommended guidelines for the minimum bend radius allowed in a flexible circuit, virtually guaranteeing eventual electrical failure within the hinge.

Figure 3-4: Flexible Circuitry Running Through Hinge of a Common Smart Eyeglass Architecture

To better understand what informs the minimum bending radius of a flexible circuit, we talked to one manufacturer that specializes in these types of systems. Flexible printed circuit boards are comprised of several layers as shown in Figure 3-5. The thickness of a flexible circuit increases depending on the amount of copper electrical layers, polymide thickness, and copper thickness. For bending, the compressive and tensile forces increase the further you are from the neutral axis (i.e., center line) of the bend, resulting in tensile forces that if larger than the yield strength of the material, can crack the copper traces. In addition, apart from tensile forces, there are also shear forces within the flexible stack-up that can separate layers and lead to delamination of layers, particularly in the compressed regions. Therefore, it is important to

keep the thickness of the flexible circuit as small as possible in the bending regions to allow for greater electrical robustness over the life of the system. We've been told by manufacturers that for semi-dynamic bends (i.e., subjected to flexing less than 100 times), the minimum bend radius should be 20 times the thickness while for dynamic bends, the minimum bend radius should be 100 times the thickness.

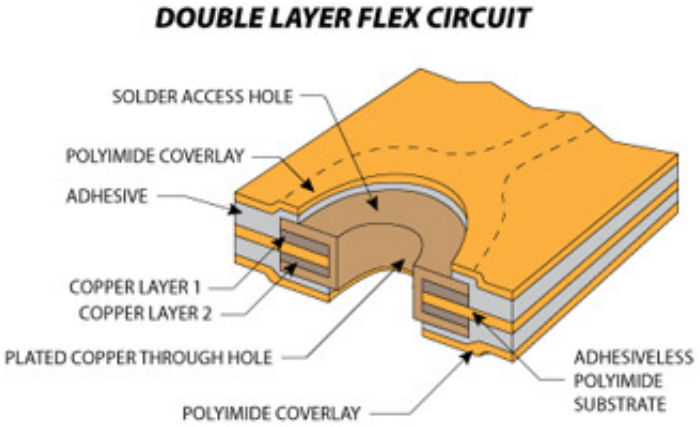

Figure 3-5: Two Layer Flexible PCB Stackup [1]

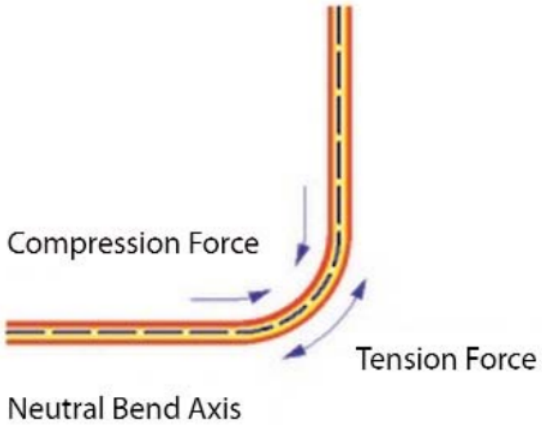

Figure 3-6: Compressive and Tensile Forces in Bent Flexible Circuit [2]

Battery Architecture

Another tip we received from the manufacturers is that battery placement in the glasses is important for weight balance and power capacity considerations. Batteries are usually the most dense component in a pair of smart eyeglasses so their positioning matters due to weight distribution. If placed too far forward, the glasses can slip off

easily and if placed behind the ear, can lead to a less traditional aesthetic due to the battery's large form factor. There are batteries with curved designs that can potentially go around the ear but those designs come at a trade-off of decreased power density and likely require custom manufacturing. Oftentimes, designers put a single battery on only one of the glasses (Figure 3-7) to simplify the design but that biases the weight towards one side and can lead to concentrated stress regions on the nose pad due to the mechanical moment imposed by the battery. The most ideal method would be to place two batteries, one on either side, to better balance the weight but this would lead to a more complicated power distribution architecture and charging circuitry.

Figure 3-7: Battery on One Side of Glasses

3.3 Sensor Selection

Through the design considerations outlined in Section 3.1 and with the valuable knowledge acquired during our design exploration discussed in Section 3.2, we were able to more easily narrow down the sensors for which to incorporate into our system. When selecting, we searched through much of the more recent literature on quantifying cognitive load and attention through physiological measurements, since that's our target research focus. What was quickly apparent was that much of the more recent research has been focused on using electroencephalogram (EEG) on the frontal lobe of the brain to quantify a user's attention. In fact, there have been several companies and researchers that use this technique for a mobile wearable, including BrainCo [46],

Muse [27], and AttentivU [47]. However, as described in Section 3.1.4, electrode-based techniques aren't well suited for long-term use and during mobile activities so we decided to refrain from selecting this type of sensing technique.

We converged to having a non-contact temperature sensor (i.e., thermopile) on the nose since nose temperature has been used before to measure cognitive load [17]. However, since the nose temperature is also influenced by external stimuli (e.g., air flow), a baseline temperature measurement is needed. As a result, a second non-contact temperature sensor is positioned on the temple as a ground truth since the temple surface temperature is a good estimate for internal temperature and should not see rapid fluctuations during internal state changes [19].The second sensor we chose was an eye blink sensor, as cognitive load influences blink rate [13, 22]. To accurately measure blinks, we are using an IR emitter/receiver pair since its relatively low-power compared to camera-based approaches, its also non-contact, and has shown promise in detecting blink rate and intensity [24]. We also include an IMU, which allows us to capture a person's activity level and head pose [48]. It has also been shown that a person's head pose can be correlated with gaze direction which is important to determine what a user is fixated on in a controlled space [26]. Lastly, we adapt the system to be compatible with the VIVE Tracking System so that we can utilize existing hardware infrastructure for 3D location estimation, ideal for crowd applications and monitoring movement in a room. This will also allow us to determine where a user is fixating by using the head pose and 3D localization data. The system also includes on-board user-noticeable LEDs that can subtly alert the wearer to any notifications (e.g., take a break, alertness is dropping, etc.), influence their attention, and, externally-facing for creative use in performances.

3.4 Design

The design of the system consists of three parts: the electrical, mechanical, and firmware. The electrical and mechanical design progressed in parallel since one informed the other. As mentioned in the acknowledgements, this was a joint effort between David Ramsay and I, with both of us collaborating on the initial planning, but him working primarily on the mechanical while I progressed on the electrical and firmware, with the exception of close collaboration on the light pipe design, which will be explained in the mechanical design section. A summary and comparison of our system with other popular smart eyeglass systems is shown in Table 3.1.

Table 3.1: System Specifications and Comparison

Design Target	Captivates	Google Glass Enterprise Edition 2	Focals by North
Weight	35g	46g	72.57g
Battery Life	10 hours	<8 hours (depends on usage)	18 hours (intermittent use)
Battery Recharge Time	3 hours	1 hour (fast charge)	2 hours
Ruggedization	None	Water and dust resistant	Water and dust resistant
Sensing Modalities	•Nose and Temple Temperature •Touch •Blink Rate •Inertial Measurement •3D Location	•Microphone •Touch •Camera •Inertial Measurement	•Microphone •Ambient Light •Inertial Measurement
Actuators	•Eye-facing LEDs •Externally-facing LEDs	•Display Module •Mono Speaker, USB audio, BT audio	•Display Module •Mono Speaker
Price	$100 (manufacturing cost only)	$1000	$599

3.4.1 Electrical System

Circuit Layout

In order to maximize the amount of circuitry that can fit into a traditional eyeglass form factor, we decided to use a system design that uses 3 PCBs, one on either side and a flexible one running across the front and connecting to both side PCBs. The side PCBs would be a standard, FR4-core, circuit board at 0.8mm thickness so as to minimize the overall thickness of the eyeglass housing. We also design for 4/4mil copper spacing to allow for a compact circuit design. A picture of the unpopulated side PCBs is shown in Figure 3-8.

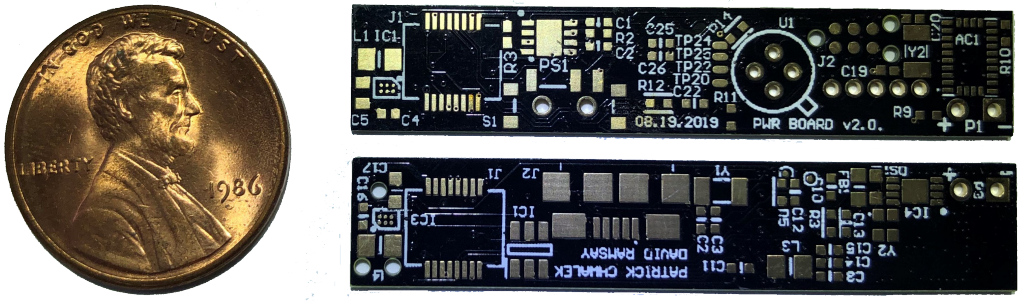

Figure 3-8: Unpopulated Side PCBs

The front flexible circuit was the most novel piece of circuitry that we designed, since the circuit board needed to conform to the contours of the eyeglass's front plastic housing and also needed to be robust enough for the dynamic bending experienced within the hinge. For the former, we designed a 3D rendering of the flex that had the correct curvatures (Figure 3-10) and then used computer-aided design (CAD)

software to unbend it to get the flat footprint (Figure 3-9).

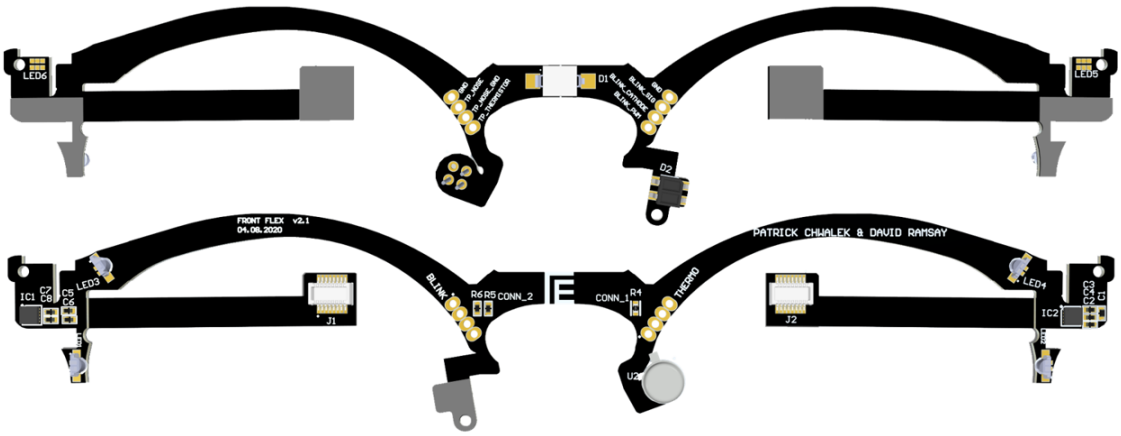

Figure 3-9: Design of Flexible Circuit Board (Top is Front-facing, Bottom is Rear)

As explained in Section 3.2.3, the thinner the flexible circuit is, the tighter bend radius it can robustly tolerate. The issue we had was that we needed a four-conductive-layer flex PCB to be able to fit the amount of power and signal conductors we needed running across the glasses; if we chose to do a 2-layer design, that would dramatically increase the rim thickness of our glasses which would make the overall frame look bulky. With a 4-layer design, the minimum thickness that a manufacturer can guarantee is 0.26mm, which per the 1:20 thickness-to-minimum-bend-radius ratio, would allow for a recommended 5.2mm radius bend, too large for any hinge design we could reasonably integrate.

After touring a flexible PCB manufacturing facility in China and becoming familiar with the process, we came up with a nontraditional solution to increase the robustness of our circuit and allow for a tighter bend radius through the hinge. We settled on having the circuit board be 4-layers across the front of the glasses and then transition to 2-layers going through the hinge, a process we thought would be feasible when we watched a manufacturer assemble each layer of a flexible circuit. This 2-layer portion can be made as thick as 0.1mm, which would allow for a recommended 2mm bend radius. To further improve the robustness, as also shown in Figure 3-10, the bend that we designed in the hinged area is an S-bend that is made by an initial static bend that can tolerate a tighter bend and gives space for a less extreme dynamic bend across the remainder of the 2-layer portion that is activated when the glasses are opened and closed.

One of the most complicated issues with this system was how to incorporate the sensors around the nose into the design. From a design perspective, the easiest

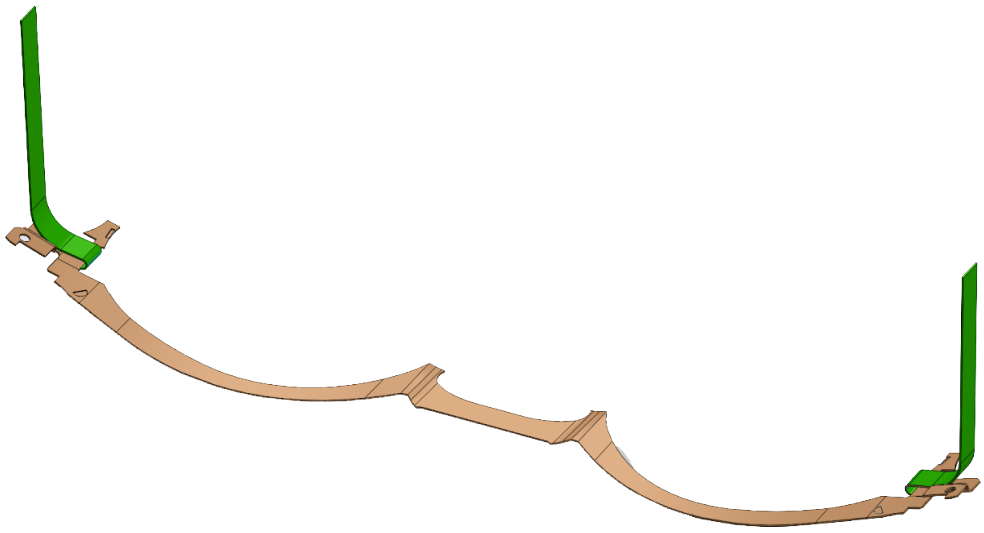

Figure 3-10: Curved Flex Circuit Rendering

approach is to solder the near-nose sensors to the flex circuitry through a set of wires, but this method would significantly lengthen the assembly process. Instead, as we did with the entire flex assembly to match the contours of the glasses, we matched the bends of the flex circuit to that of the mechanical housing around the nose, creating two tails of flex circuitry that included the temperature and blink sensors. During installation, these tails are folded into the groves of the plastic housing, not requiring any additional soldering. The final manufactured flexible circuit is shown in Figure 3-11.

Figure 3-11: Manufactured Flexible Circuit

Processors

A block diagram of the entire system is shown in Figure 3-12. For the main processor, we use a dual-core microcontroller (MCU) that has an ARM Cortex-M4 processor for the user application and an ARM Cortex M0+ processor for the network stack. This processor was chosen because it allows us to treat the MCU as an application processor, not worrying about the overhead of running the network stack. This approach allows for quicker embedded firmware development and increased versatility for applications that don't require the network component. In addition, this processor supports various types of network stacks (i.e., Bluetooth 5, Zigbee 3.0, and OpenThread) so if a different network architecture is required for an application, switching does not require a hardware redesign. Lastly, using this chip allows for a simpler electrical design than using two discrete processors (i.e., an application processor and network processor), since both cores are integrated into a single discrete package.

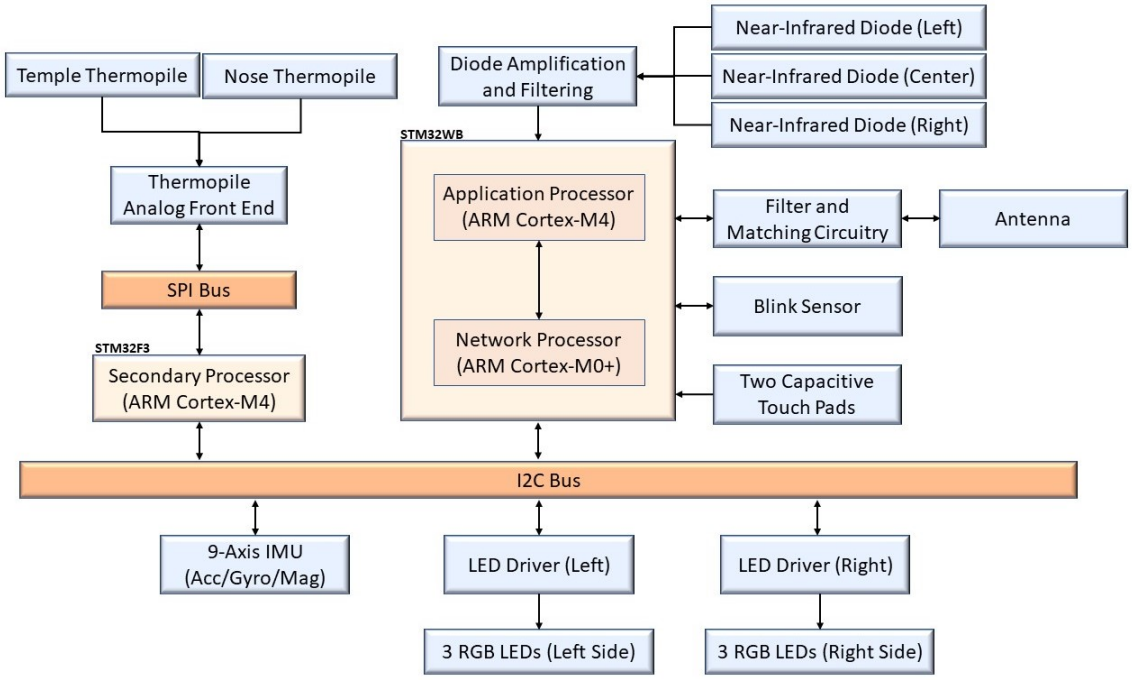

Figure 3-12: System Diagram

The dual-core MCU exists on one of the two-sided PCBs, which are separated and connected by the flexible circuitry. Given that sensors exist on all sides of the glasses and that the long separation between the two-sided PCBs can lead to signal integrity issues, we decided to put a secondary Cortex-M4 MCU on the opposite side

of the glasses. That way, this secondary MCU can manage the sensors on one side of the glasses, apply any preprocessing algorithms to its local sensor streams, compress the data, and send it to the main processor when ready or requested. Although this raises the software complexity, it allows us to simplify the hardware by reducing the amount of conductors going through the flexible circuit.

Power Architecture

Our system utilizes two batteries for better weight distribution, but this approach increases the complexity of our power architecture and charging circuitry. As shown in Figure 3-13, the lithium-polymer batteries and their charge controllers are electrically separated in our system until they are interfaced with our power multiplexer (MUX). The reason for the separation is that since the batteries are physically far apart (i.e., one in each arm) and since we cannot guarantee that they are from the same manufacturing batch, it is unsafe and ill-advised to put them in parallel for charging through one charge controller. For one, if they are not from the same batch, one of the batteries will likely age at a different pace than the other, leading to possible cross charging (i.e., one battery depletes into the other), resulting in a premature system failure. Also, this would constrain the design to always use the same capacity of batteries for both sides– a constraint we didn't want in case there was an application where a larger density battery were to be tethered to one side. Another reason we went with the separated charge controllers is that with one charge controller and two spatially separated batteries, the two charge lines running to the batteries have to be impedance matched to ensure balanced charging, a constraint that would have been difficult to model given the complexity due to multiple connected circuit boards.

Within the multiplexer, the battery with the highest charge is prioritized over the other until discharged to around 0.2V under the idle battery. At that point, the comparator of the MUX triggers a switch and the idle battery is set to active while the active is made idle. If both inputs to the MUX are under 2.8V, the MUX goes into high impedance mode and shuts off the battery inputs and system output completely. This ensures that the batteries aren't over depleted, which can result in damage to the batteries and potentially risking a hazard. The MUX output then powers two 3.3V, low-quiescent, linear regulators, one on each side of the system, and also directly powers the two LED drivers located on the flexible circuit.

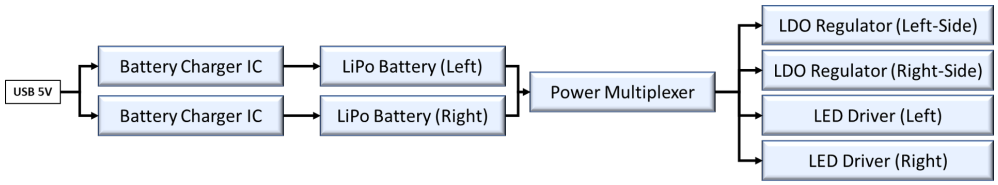

Figure 3-13: Power Architecture

Temperature

The thermopile circuit consists of two analog thermopiles, one faced towards one side of the nose (Figure 3-14) and the other faced towards the user's temple (Figure 3-15). These thermopiles are then connected into an analog front end, which amplifies and filters the signal before being converted by the secondary processor's 12-bit analog ADC. Once a set of readings is collected by the secondary processor, the main processor is alerted over an interrupt line, and after acknowledgement, the data is transmitted over the I2C bus to the main processor.

Figure 3-14: Nose Thermopile

Blink

The blink sensor is an IR emitter/receiver package that showed promising results in [24] on detecting blink events. We experimented with different blink detection locations and found that putting the sensor on the bridge of the nose, facing towards the eye, yielded the best performance. The near-infrared diode that emits light can be

Figure 3-15: Temple Thermopile

modulated through the main application processor to better dynamically adjust the reflected power seen by the receiving diode and enable synchronous detection. This feature is important for environments that have external near-infrared light sources (e.g., sunlight) that can saturate our blink sensor.

3D Localization and Pose Estimation

For the 3D localization, we are using two VIVE base stations [49] as our signaling sources in the environment in order to triangulate the glasses in 3D space. Each base station emits a vertical and horizontal near-infrared sweep, alternating between both sequences and between both base stations. Prior to each sweep, a pair of near-infrared LED flashes occur that signals to the receiver (i.e., the glasses) which base station is sweeping and if the sweep is horizontal or vertical. Ultimately, a set of four sequential sweeps (i.e., 1 horizontal and 1 vertical from each base station) repeats at a rate of 30Hz and is received by a near-infrared receiver. With the receiver being calibrated to know the global coordinates of each base station, it can determine where it is in 3D space by identifying the sweep and comparing the timing between sweeps. VIVE states that a less-than 2mm accuracy and a range of 16 feet can be achieved with their system, but based on preliminary testing we believe our tuned custom system can detect the base station sweeps at a range of at least 30 feet.

The hardware architecture for our receiver consists of three near-IR sensitive diodes, one on each face of the glasses. The search criteria for choosing which type

of diode to use consisted of finding a surface mount component that had a slim form factor but a large enough receptive area that also offered adequate responsivity to capture the momentary sweep. In addition, the field-of-view (FOV) needed to be large enough as to limit any blind spots. We evaluated six types of diodes and found that the VBPW34FASR by Vishay performed the best and conformed to our design constraints. The three Vishay diodes in our circuit are then tied together (i.e., summed) and passed into a transimpedance amplifier to convert the diode-generated current into a voltage. The reason for the summation is that we wanted to treat the system as a point mass that has a large FOV so that as a person turns their head, at least one diode will still be exposed, increasing the reliability of the system. This also allows for a simpler analog front end with the trade-off of decreased accuracy, since the diodes are not spatially co-located but are treated as such in software. Although VIVE can achieve

For head pose, we chose a digital 9-axis IMU that connects to our I2C bus and offers various internally computed metrics, such as activity classification and a quaternion matrix for estimating head pose. Software calibration was done to align the IMU's reference coordinate system to the glasses coordinate system for accurate head pose estimation.

Other Components and Interactivity

Our system offers a few other components for easy usability, interactions, and actuation. In order to easily program our system, we wired the USB micro port to both the battery charging circuitry and to the single-wire programming interface for the main application processor through a modified USB cable and an ST-Link Programmer—this can be done by anyone with just a programmer and a spliced USB cable. In addition, the system has two capacitive touch points on the left arm that are sensitive enough to identify taps and both forward and backward swipe gestures. The sensitivity of the touch points can be adjusted in software to tune for specific applications. Finally, as shown in Figure 3-16, there are six RGB LEDs, one facing towards either eye, two on the top, and two facing forward. The LEDs toward the eye are intended to be used as notification [50] while the other LEDs can be used for crowd applications (e.g., a concert where the LEDs are actuated by the person's physiological signals).

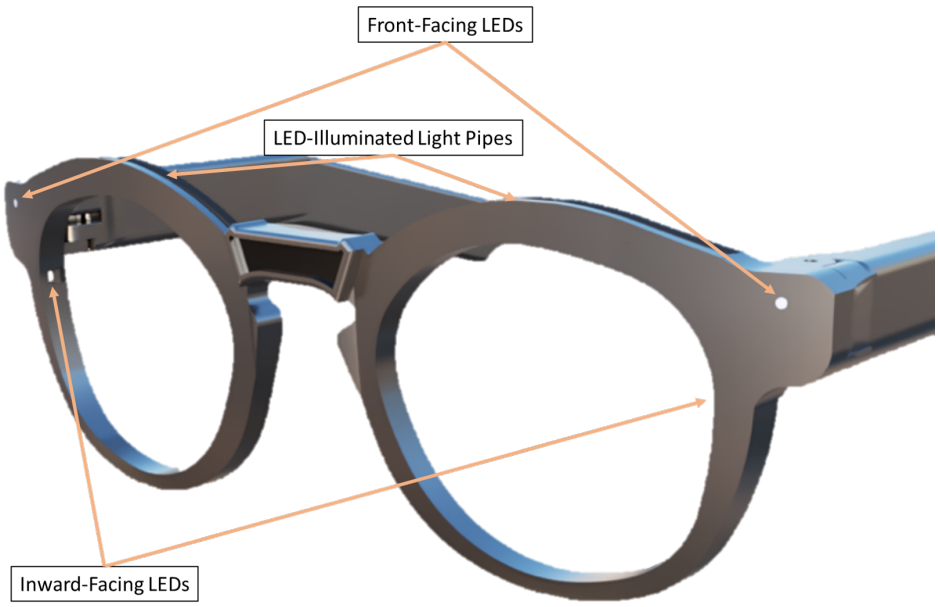

Figure 3-16: LED Locations

3.4.2 Mechanical Design and Assembly

The mechanical design of the glasses has experienced several iterations as we converged on a final form factor. As outlined in Section 3.1.1, we wanted to choose a design that looks less like a research experiment and more like what most people would associate with a traditional pair of glasses. Figure 3-17 shows the front view of our most recent design. What can also be seen is an exposed opening above the bridge of the nose and two exposed openings, one on either side of the glasses (Figure 3-18). These openings are to expose the near-IR diodes to the incident light emitting from the VIVE base stations. Our goal is to cover these openings with near-IR transparent plastic, but for testing we've left them open for easy debugging. These openings can be covered or completely omitted for an application that does not require the 3D localization feature.

Hinged Design

We iterated on various hinge types, including omitting a hinge entirely and creating a rigid design. Eventually, we settled on an approach that is commonly found in smart eyeglasses, where the hinge is a plastic design with the flex circuitry running through the middle of it. Our flex circuitry is preformed (Figure 3-19) and inserted into the front housing, resting on internal raised plastic features (Figure 3-20). The flex circuit is then sealed by a piece of plastic acting as the back plane (Figure 3-21) and then

Figure 3-17: Front View

Figure 3-18: Angled View

the arm is able to be rotated around the flex circuit (Figure 3-22) and inserted into the plastic via a fastener running through the rotating axis of the hinge.

Light Pipes

One of the design features we were most excited about was to have the glasses include a hidden feature that allowed them to transition from normal day use to a more festive environment. To do this, we included a few RGB LEDs, including one on each brow that would light up the entire top ridge of the glasses. The goal was to have the illuminated features blend into the design so that when turned off, it wouldn't be immediately apparent that this feature existed. One way we attempted to do this was by crafting a piece of polymethyl methacrylate (PMMA), also known as acrylic, to form a light pipe for the light from the RGB diode to transmit through. For our

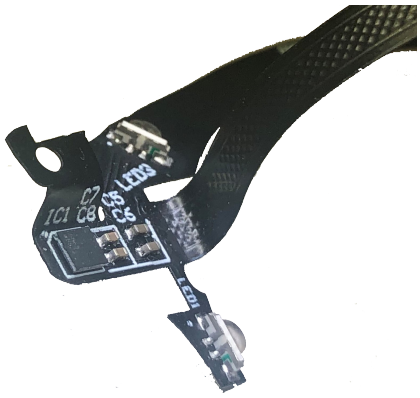

Figure 3-19: Pre-bent Flex Circuitry

design, we needed the light to diffuse through the acrylic and evenly illuminate only one face while internally reflecting on all the others (Figure 3-23). We worked with a manufacturing facility in China to laser cut these PMMA pieces and chrome-coat three of the four sides to maximize the light emmission on the top-face (Figure **??**). Although this method worked, we found it to be cost prohibitive given that a single piece cost us over $50 to custom manufacture.

An alternative approach that we settled on was to use side-glow fiber optic cables. A traditional fiber optic cable is not suitable, since it lacks additional cladding on the cable's surface that allows for greater diffusion of light across the length of the cable. Side-glow fiber optic cables, however, have a cladding, usually clear Teflon, that decreases the critical angle of incident on the internal surface that is required for external emission of light. Figure 3-25 shows the fitted 1.5mm diameter side-glow fiber optic cable illuminated in our design. Apart from the brow lighting, we also have lighting at two points on the front face of the glasses where some traditional glasses have rivets. To increase light visibility, we added commercial-off-the-shelf PMMA light pipes with fresnel lensing on one end so that the emitted light is less directed, offering better angular visibility.

Battery Placement

As we learned from our trip to Shenzhen (3.2.3), strategic battery placement was key to a more comfortable design. We have two, 150mAh, batteries on either side of the glasses, running in series with the side printed circuit board. The circuit boards are designed so that when assembled, the width is uniform across the batteries and across the circuit board assemblies for a streamlined design (Figure 3-26).

Figure 3-20: Inside of Front Housing

3.4.3 Firmware

For a system of this complexity, given all the sensor channels and preprocessing steps, it's warranted to use some low-level operating system to help keep track of all the events, prioritize any events that need more real-time action, and provide the utility for easy debugging. To do this, we decided on using a real-time operating system that already has native support for the processors we are using and has a variety of tools that aid in debugging and in modelling the system's performance.

FreeRTOS

FreeRTOS is a light-weight real-time operating systems that is capable of being run on microcontrollers and small microprocessors which makes it suitable for this project [51]. This operating system is also supported by ST Microelectronics, the manufacturer of our processors, and has several additional tools written by ST Micro for easier system integration. FreeRTOS allows for the encapsulation of individual blocks of code into threads that can be triggered by a variety of sources (e.g., hardware interrupts, timers, other threads, etc.) and has various methods of passing messages across threads (e.g., queues, semaphores, mutexes, etc.). This encapsulation allows you to think of the firmware system as embedded blocks of code that have a deterministic operating behaviour. In addition, FreeRTOS offers various debugging tools

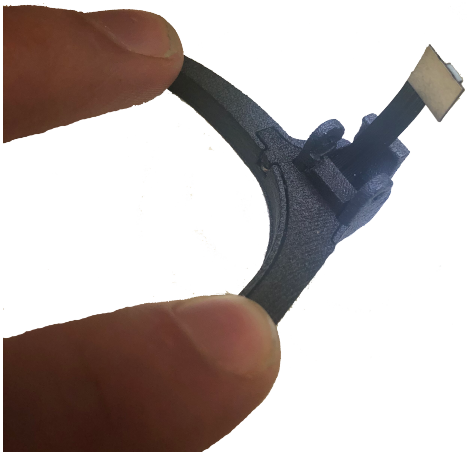

Figure 3-21: Sealed Front Housing with Flex Circuit Integrated

Figure 3-22: Adding a Side Leg

Figure 3-23: PMMA Light Pipe Design for a Single Brow

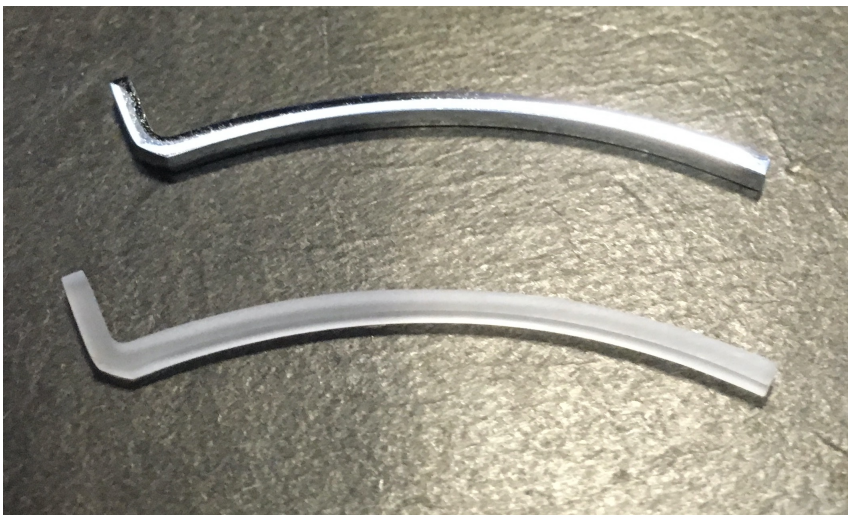

Figure 3-24: PMMA Light Pipe: Chrome-Plated (Top) and Non-Plated (Bottom)

Figure 3-25: Illuminated Brow of Glasses

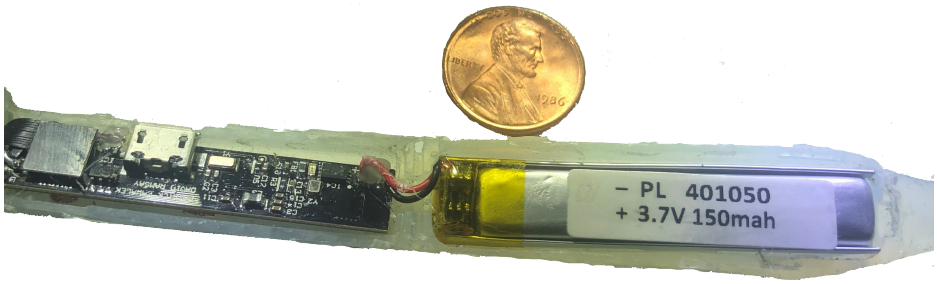

Figure 3-26: Battery Placement in One Arm of the Glasses

that allow you to assess the memory and time utilization of the processor in real-time to model the system and determine if certain threads are causing delays or deadlocks. This allows for more efficient debugging since the problem can be quickly pinpointed to specific logical blocks.

As shown in Figure 3-27, there are a variety of threads in the system, all controlled by a master thread within each processor. The master thread receives queues from either the radio communication thread or inter-processor thread to enable certain sensor channels, actuators, or to send or receive messages. After enabling, each individual sensor channel is triggered by the hardware resource it uses for most efficient operation. For the blink sensing, the direct memory access (DMA) controller is used so that when sampling the blink sensor at 1kHz, the thread only awakens every second (i.e., 1000 samples) to preprocess and pass the data pointer to the master thread where the data is packaged for transmission. For inertial sensing, the digital IMU interrupts the processor when data is ready to be transmitted, triggering the Inertial Sensing Thread to start I2C transmission. The secondary processor also interrupts the main processor once it buffers 10 samples from each thermopile sampled at 10Hz and waits for an acknowledgement before sending the data over I2C.

3D Localization

The 3D localization code was adapted from prior work done in [5]. Every time one of the near-IR diodes is saturated by incidence light (i.e. rising or falling edge), an interrupt is triggered which results in a callback recording the exact timestamp of the event in microseconds. These timestamps are then fed into a thread that classifies if the event was part of a sequence (i.e., two flashes and a sweep) or if the event is noise. If part of a sequence, the code catalogs all the sweeps and calculates an estimate for the 3D location of the device. Since the code requires the global coordinates of the

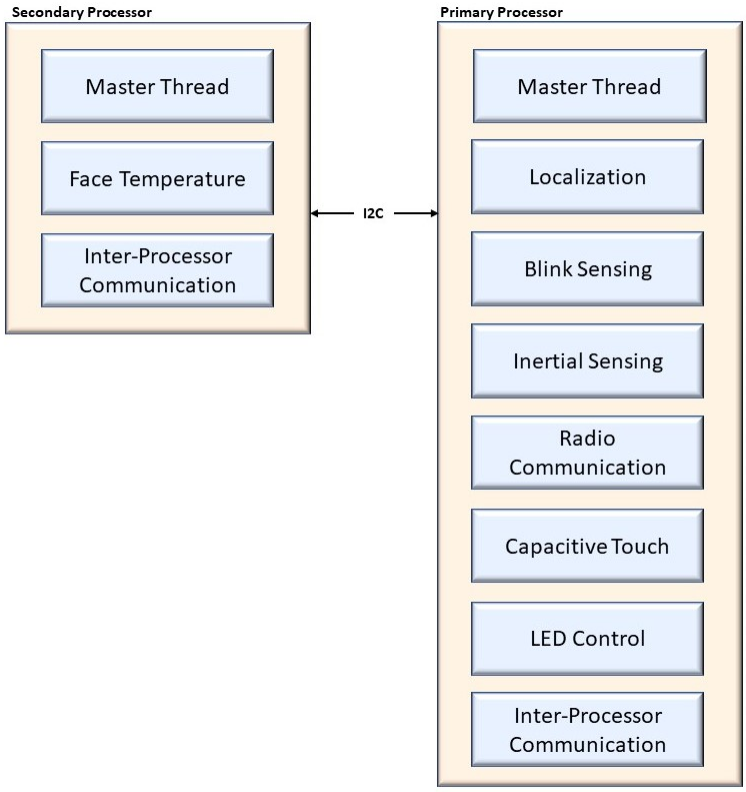

Figure 3-27: FreeRTOS Threads in System

VIVE base stations, they can either be hard coded into the software or updated via the server.

Networking: OpenThread

For networking, there are a few protocols that we could have chosen to use with our network processor. In the end, we wanted to use something that was scalable, supported across a variety of platforms, and could integrate well with server infrastructure and ultimately, the internet. The most ideal candidate was an open-source version of Thread [52] called OpenThread [53]. OpenThread "is an IPv6-based networking protocol designed for low-power Internet of Things devices in an IEEE 802.15.4-2006 wireless mesh network, commonly called a Wireless Personal Area Network (WPAN). Thread is independent of other 802.15.4 mesh networking protocols, such a ZigBee, Z-Wave, and Bluetooth LE" [53]. OpenThread has a variety of features that makes it suitable for our system. One primary advantage is that it's a protocol that is seeing wide adoption across a variety of platforms, and since its being supported by Google, it's integrated into many of their home devices— in theory, our system could use Google Home nodes as hops in our network. It's also scalable to hundreds of nodes, making it ideal for future integration of our glasses into an ecosystem of sensors and actuators. In addition, its IPv6-based so integrating it with other IPv6-based systems is simple and doesn't require a network translator, making it appealing for a "future-proof" implementation. There also exists Contiki, which is another low-power IPV6-based protocol and there is evidence that it has improvements in latency and packet loss rates over OpenThread [54]. However, [54] was just an introductory comparison between both protocols, and more work needs to be done on measuring the power efficiency and performance of both protocols. We chose to stick with OpenThread due to the native support ST offers for our network processor and the active open source community around it.

3.5 Calibration

A few of the subsystems need to be calibrated once the system is assembled and electrically tested. The analog RF circuitry needs to be properly matched to maximize our system efficiency, and the temperature sensors require calibration. In addition, we also have to ensure the tracking system performs optimally across the ranges of a medium-sized room (30 ft x 30 ft).

3.5.1 Radio Tuning

Ground plane size and external plastic casings modify the resonant frequency of an antenna, warranting a proper tuning once a physical system is built in order to minimize losses [55]. To choose the desired transmission channel within our 2.4GHz frequency spectrum, we need to take into consideration the frequencies WiFi commonly transmits at as to avoid any possible noise artifacts. As is shown in Figure 3-28, there are a few channels in the 2.4GHz spectrum that don't overlap with the common Wi-Fi channels of 1, 6, and 11, so for the purpose of this radio tuning, we will choose channel 22 (2450 MHz) as our system's transmission frequency.

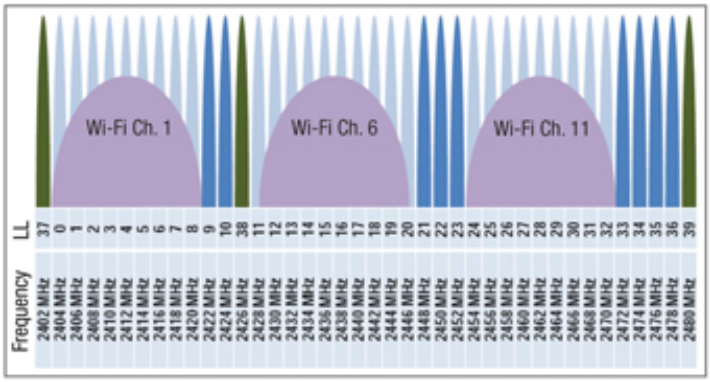

Figure 3-28: 2.4GHz Channels and Corresponding Frequencies with Common Wi-Fi Frequencies Superimposed [3]

The current circuit, as shown in Figure 3-29, includes a PI-matching circuit followed by a low-pass filter (recommended by ST Microelectronics) and then a second PI-matching circuit right before the antenna for additional tuning flexibility. For the start of this tuning, the PI-matching circuit nearest to the antenna is left unpopulated and the inductor is replaced with a $0\,\Omega$ resistor. For testing, the microcontroller is left unpopulated on the PCB and a vector network analyzer (VNA) is attached to the start of the RF transmission line originating from the microcontroller. The entire setup is encased by the eyeglasses plastic housing (Figure 3-30) and each measurement is taken with the device on a person's head for the most real-world setup. The resultant plots from the VNA are shown in Figure 3-31, and as can be seen from reflectance plot, the antenna's resonance is around 2.4 GHz, a bit further from our target of 2.45 GHz.

The Smith chart in Figure 3-31 allows us to view the admittance and impedance of the circuit to tune the PI-matching circuitry. The circles originating from the left-side are of constant conductance, and you can traverse the circle clockwise by

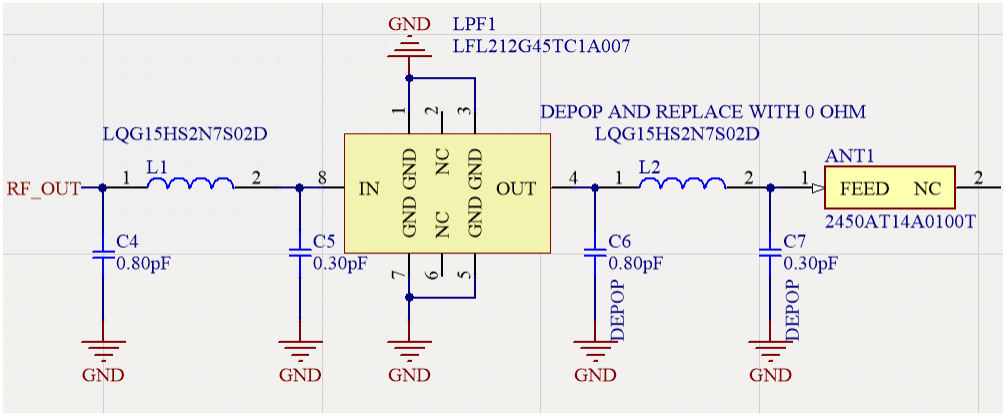

Figure 3-29: Antenna Circuit

Figure 3-30: Antenna Tuning Setup

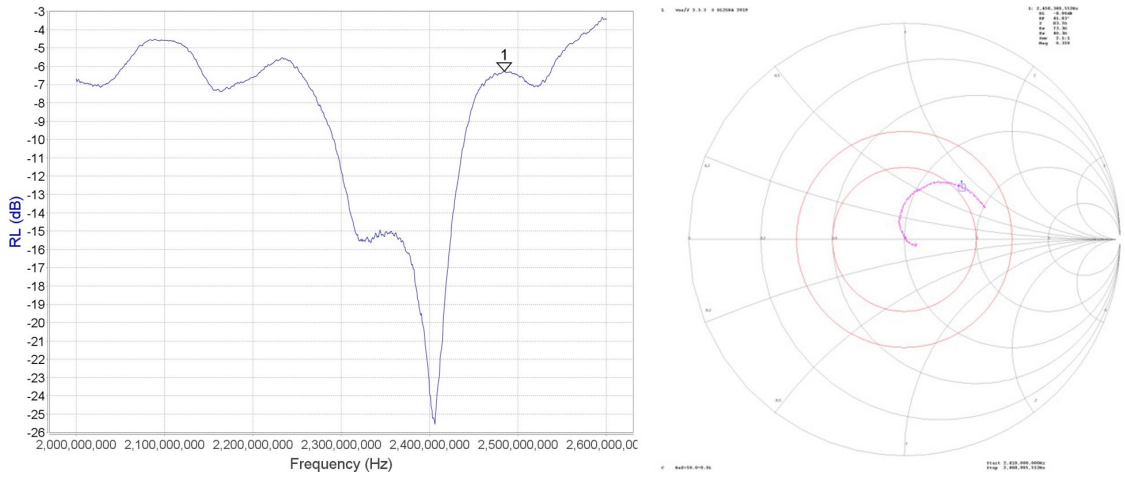

Figure 3-31: RF Reflectance Plot (Left) and Smith Chart (Right) Prior to Tuning

adding shunt capacitance or counter-clockwise by adding shunt inductance. The circles originating from the right-side are of constant resistance. To traverse clockwise, you add series inductance, and to go counter-clockwise, you add series capacitance. Using the traversal techniques, you can tune your circuit to be 50 + 0j at a specific transmission frequency. For our circuit, it first appears that we can increase the performance by adding series capacitance to shift the plot closer to the center (50 Ω impedance). Prior to iterating, there was one small oversight in that the 0 Ω resistor put in place of the inductor has some predicted inherent capacitance and inductance at 2.4GHz which is unideal. To mitigate, the 0 Ω resistor is replaced with a 10pF series capacitor to give us the results shown in Figure 3-32. As you can see, the resonance has shifted near the target frequency of 2.45GHz and the amount reflected has decreased from -8dB to -16dB, leaving us with a more efficient RF circuit. There is no practical reason to try to tune the circuit closer to 2.45GHz resonance, since various environmental conditions (e.g., human hair covering the plastic, ambient humidity differences, etc.) can still shift the resonance slightly, so our current tuning is close enough.

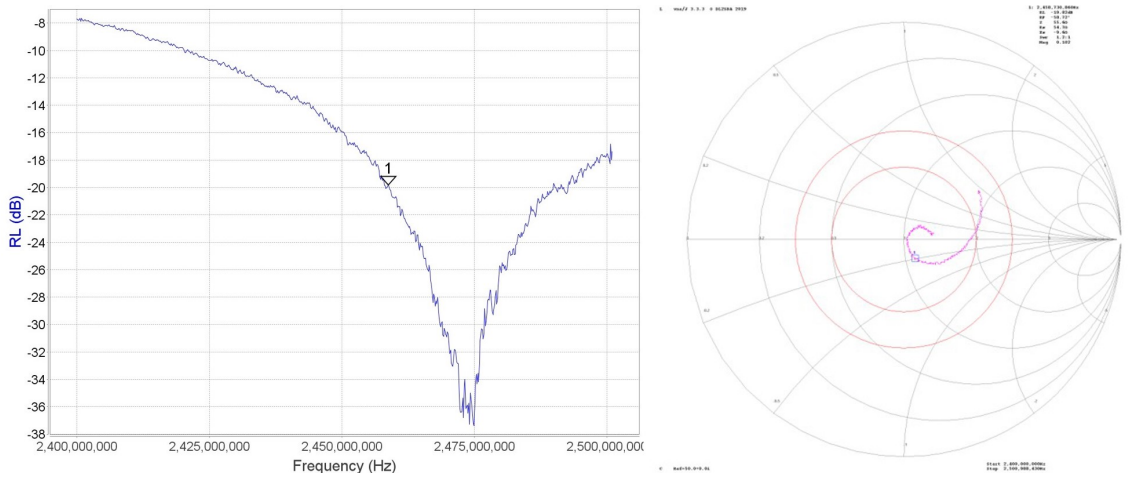

Figure 3-32: RF Reflectance Plot (Left) and Smith Chart (Right) After Tuning

3.5.2 Temperature

For our device, it's important to calibrate the thermopiles because both perform differently based on where they are positioned in our circuit and will experience different resistive losses and noise artifacts. As shown in Figure 3-33, a thermopile is a series of thermocouples with one set of junctions fixed to the sensor housing (reference junction) while the other set (active junction) is coated with an IR absorbant material

and exposed to ambient light, usually blocked by a far-infrared transmssive material to filter the incident light. As a temperature differential forms between the two sets of junctions, a voltage differential is created through a process known as the Seebeck Effect.

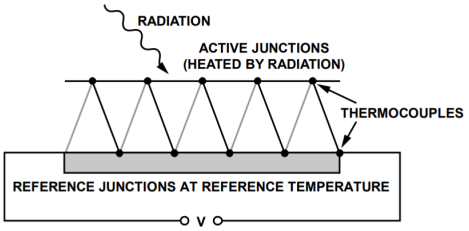

Figure 3-33: Thermopile Architecture [4]

As explained in [56], the total radiation emitted by an object that does not selectively emit or absorb specific wavelengths is defined as

$$P_{obj} = \sigma\varepsilon(T_{object})^4 \tag{3.1}$$

where P_{obj} is the total radiation power emitted by an object, T_{object} is the temperature of the object, σ is the Stefan-Boltzmann constant, and ε is the emissivity of the object. For many objects, the emission factor is between 0.85 to 0.95, where 1 is for an ideal black body. Equation 3.1 is formally known as the Stefan-Boltzmann's Law, and combining it with the seebeck effect, we get

$$V_{TP} = A(T^4_{object} - T^4_{reference}) \tag{3.2}$$

where V_{TP} is the thermopile voltage, T_{object} is the temperature at the active junction, and $T_{reference}$ is the temperature at the reference junction [57]. The constant, A, is a product of the thermal resistance of the thermopile, number of thermocouples within the device, the Seebeck coefficient, net emissivity between the object and the device, Stefan-Botlzmann constant, and field of view (FOV) of the device. In most cases, this constant is approximated for during calibration by varying the temperature of an object and of the thermopile housing, while placing calibrated thermocouples throughout. Rearranging equation 3.2 to get T_{object} as a function of V_{TP}, we have

$$T_{object} = \sqrt[4]{T^4_{reference} + V_{TP}/A} \tag{3.3}$$

However, in [58], it is recommended that we also account for heat flow other than radiation, such as convection and conduction from nearby objects. We modify our

equation as such,

$$T_{object} = \sqrt[4]{T_{reference}^4 + f\{V_{TP}\}/A} \tag{3.4}$$

where

$$f\{V_{TP}\} = (V_{TP} - a_0) + a_1(V_{TP} - a_0))^2 \tag{3.5}$$

and a_0 and a_1 are terms to solve for during the calibration process.

For our system, to estimate each thermopile's reference temperature, $T_{reference}$, each thermopile has a thermistor built in with a manufacturer-provided characterization [59]. To calibrate, we took a peltier module that converts a voltage potential into a temperature differential and using thermal compound, embedded a calibrated thermocouple between the peltier module and a thin piece of laminate material. We then placed a calibrated digital thermopile right next to our analog thermopile as a reference. Finally, using thermal compound again, we attached a thermocouple to the housing of the thermopile to estimate the accuracy of the built-in thermistor. A diagram of the test is shown in Figure 3-37 with a snapshot of the experiment in Figure 3-35.

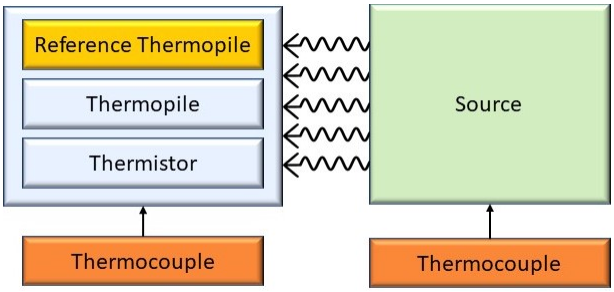

Figure 3-34: Thermopile Calibration Test Setup Diagram

The test was performed over 10 different temperatures for both the nose and temple thermopiles, ranging from 60-115 degrees Fahrenheit, the average temperature ranges expected across the surface of a person's skin. For each iteration, we waited a minimum of 5 minutes to ensure the temperature of the reference object has stabilized. After the experiment, we used non-linear least squares to solve for the constant A given V_{TP}, $T_{reference}$, and T_{object}. This allows for a suitable approximation for temperature within our ranges but assumes constant emissivity. This assumption doesn't hold true given the variety of skin types so a more thorough calibration is warranted where emissivity of the object and $T_{reference}$ are varied. Table 3.2 shows our results from the least squares approximation.

It's important to note that for a commercial system that requires accurate tem-

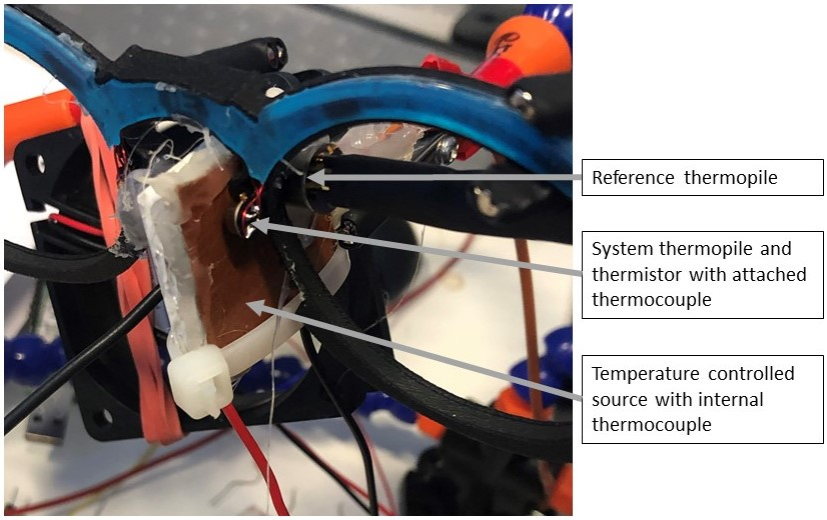

Figure 3-35: Thermopile Nose Calibration Test Setup

Table 3.2: Thermopile Calibration Values

	A	a_0	a_1
Nose Thermopile	7.80e-10	-2.31e-01	3.61e-03
Temple Thermopile	4.21e-10	-3.62e-01	8.31e-02

perature readings of a person's skin, this calibration will need to be done for each pair of glasses due to variability between each thermopile. There exist more expensive digital thermopiles that are factory calibrated but given our main objective is to look at temperature changes, we found it to be more cost and time efficient to choose a cheaper analog thermopile and attempt to generalize the calibration across all the glasses. Going this route also allows us to design an analog front-end circuit that only includes the temperature range naturally seen on the surface of a person's face, further increasing the sensitivity to subtle temperature changes.

3.5.3 Location Tracking System

When first constructing our circuit for the near-infrared tracking system, we followed the sample circuit in [5]. The published circuit consisted of a transimpedance amplifier and a single-stage gain amplifier with a high-pass filter tuned to 338Hz in between (Figure 3-36). Through testing, we realized that the BPV22NF diode that was used in [5] had over 4 times the responsivity as our VBPW32FA diode even though on paper, they were identical. We thought that this was because the BPV22NF had an additional lens, but after shaving it off (Figure 3-36, we still saw a considerable

performance difference. To compensate, our choices were to either increase the transimpedance feedback resistor or increase the gain in the single-stage amplifier, both options which would reduce the bandwidth of our circuit. Given that the VIVE base stations emit a near-infrared sweep, we measured the exposure duration at 30 feet to find that at that distance, the exposure can be modelled as a 125kHz wave. Allowing for some headroom, we settled on optimizing the circuit for a 200kHz bandwidth. Given a transimpedance amplifier gain of 100k, we needed an additional 60dB (i.e., 1000x voltage) on the single-stage opamp to properly saturate the signal at 30 feet. The issue was that our amplifier had a gain bandwidth product (GBW) of 10MHz, which at 60dB, would mean we can only achieve a bandwidth of 10kHz, far from our 200kHz target. The solution would either be to find a different OpAmp, or add a second-stage to spread out the gain. Given that the former was more costly, we decided on adding a second stage and giving both amplifiers a gain of 30dB, resulting in an overall bandwidth of 330 kHz.

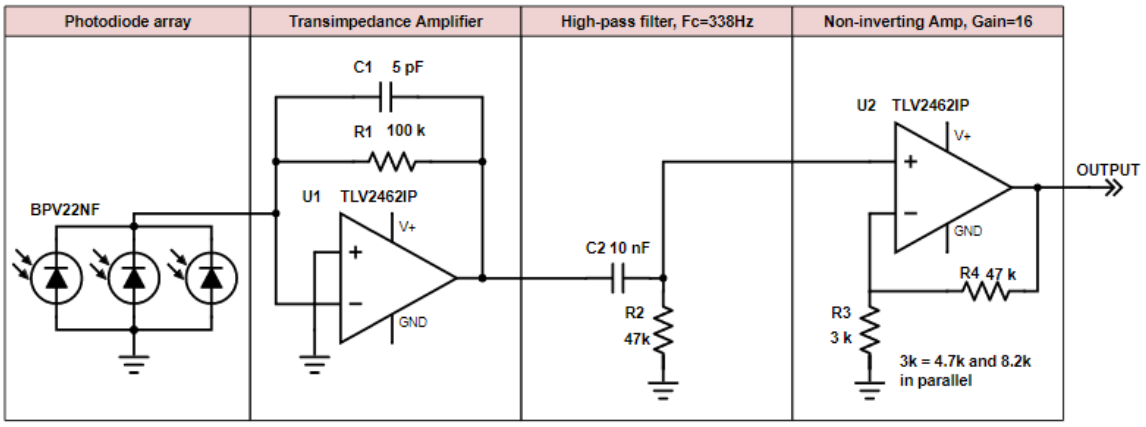

Figure 3-36: Example Localization Circuit [5]

Another issue that we found when tuning our circuit was that we often saw high frequency noise in the amplified diode signal. Through controlling the light exposure on the diodes, we found that the noise was internal to our circuit. After a considerable amount of time of debugging, we narrowed down the issue to originate from the two switching regulators that we were using. Given that our diodes were summed and that a few of the traces were a few inches long prior to entering the transimpedance amplifier, some noise originating from the switching in our regulators was able to cross into the diode signal through capacitive coupling of traces. The only solutions were to either change the switching regulators to linear regulators, increasing power inefficiencies, or adding a transimpedance amplifier at each diode and summing the signal

later in the chain. Given physical space constraints making it difficult to add additional transimpedance amplifiers, we decided to just swap the switching regulators out. The added power inefficiencies were negligible since we are only stepping down the voltage from 3.7-3.8V to 3.3V and we also chose a regulator with a low quiescent power draw. In addition, the LED drivers, the most power consuming components of the system when active, are powered directly from the batteries so switching the regulators has no effect.

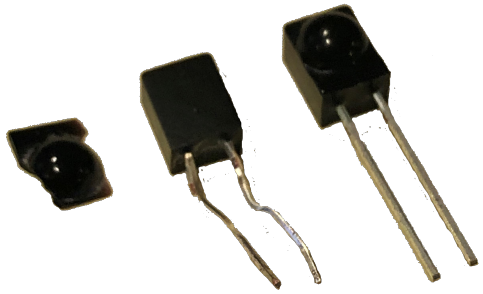

Figure 3-37: Shaved Off Lens of the BPV22NF Diode

3.6 Network

3.6.1 Server

OpenThread has a variety of node types that can exist in a given network [53]. One is a Full Thread Device (FTD) that is classified as a node that is always on and can be used as one of the 32 main routing nodes in the network. Another type is a Minimal Thread Device (MTD), which is a node that always needs a parent to forward traffic through and can go into a sleep mode to conserve power. Within FTDs, there exists a role type of Border Router that is a device that "can forward information between a Thread network and a non-Thread network (for example, Wi-Fi). It also configures a Thread network for external connectivity" [53]. This type of role is suitable for a server, since it has direct access to the thread network and can connect to the internet to outsource the data. Therefore, for indoor applications with internet accessibility, we use a Raspberry Pi equipped with a network node that is configured with the OpenThread Border Router (OTBR) build of the network stack to interact with our system through any internet-connected computer using a dashboard that will be described in Section 3.6.3.

3.6.2 Server Protocol

For the actual server component of the network, our OTBR uses a custom implementation of the Constrained Application Protocol (CoAP) [60]. CoAP is a "specialized web transfer protocol for use with constrained nodes and constrained (e.g., low-power, lossy) networks" [60]. CoAP is ideal because it is "designed to easily interface with HTTP for integration with the Web while meeting specialized requirements such as multicast support, very low overhead, and simplicity for constrained environments" [60]. There exist a few open source implementations so we chose to modify a build of CoAPthon [61] since it was written in Python and was relatively simple to add multicast functionality, a feature that isn't often fully supported in other builds of CoAP.

Each eyeglass node in our network hosts a few server endpoints that other nodes can query to retrieve the node information, such as unique indentifier (UID), node type, sensor status, and server time. Each node can also have its real time clock synchronized by the OTBR and have other settings modified as well (i.e., enable/disable sensor channels, actuate lights, and unicast to OTBR). The OTBR's CoAP server is the main endpoint to where all the eyeglass nodes send their data when logging for an experiment. However, it's easily capable for one eyeglass node to send another few metrics that can be used to synchronize actuations. One example that can be implemented is synchronizing lighting across glasses based on collective physiology measurements to visually identify how an audience is reacting to a performance.

3.6.3 Captivate's Dashboard

The OTBR has also been configured to send and receive encrypted packets over the internet to a graphical user interface (GUI) dashboard that we designed for our system (Figure 3-38). The dashboard was built using PyQt and offers various control and visualization tools [62]. First, it allows you to connect to one of our OTBR nodes and view a list of all the nodes in the network with their IP, node type, short description, and unique identifier shown. The dashboard allows you to select which sensor channels to enable on the glasses and to toggle streaming of those sensor channels. On the right, you can visualize the signals being transmitted from the system in real-time. On the left, you can see the live quaternion values, a list of activities and how confident the system's classification is of each activity being performed (e.g., walking, running, sitting, etc.), and their approximate 3D location in a room if the VIVE system is active. If the inertial sensor channel is activated, there is a secondary

window that appears that shows a block rotating in the same directions as the glasses. The dashboard also allows you to manually toggle and set the color for each LED on a pair of glasses and offers the ability to orient the inertial system in order to properly calibrate the IMU in the reference frame of the glasses.

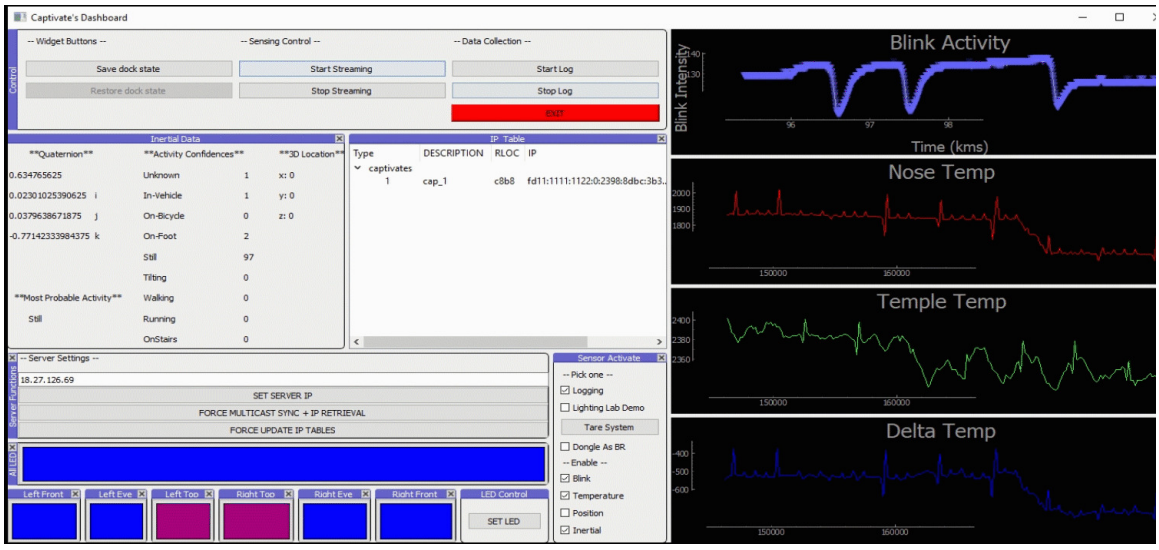

Figure 3-38: Graphical User Interface for our OpenThread Border Router

Chapter 4

System Validation

4.1 Experimental Design

For validating our system, we created a user study that allows us to test the inertial, temperature, and blink sensing components of our circuit. This validation study will also serve as basis to compare our system's performance with the results of past studies discussed in Section 2.1. In conjunction with our system, we will use additional commercial sensing systems that have been used heavily in prior works to predict cognitive states.

4.1.1 Data Collection

A set of 6 participants (3 female and 3 male) were recruited for this study, following IRB protocol #2001000083. Originally, our target was to have 20 participants but due to the COVID-19 pandemic, we had to stop prematurely after our 6th subject. All 6 participants were between the ages of 24 and 32 and were all right handed and fluent in English.

4.1.2 Study Protocol

Stimuli

The study consists of a user sitting alone in a closed room, in front of a computer screen, interacting with an automated test program that we created using PsychoPy [63] and that cycles through a variety of tests (Figure 4-1). Each chosen test has been shown in prior literature to illicit a particular internal state but require different types of use action (i.e., passively watching, physically engaging, and reading com-

prehension). One of the chosen experiments is the Stroop Task [64] where in [17] it was shown that changing the allowed user response time affected "facial temperature patterns, thereby giving a measure of cognitive load." The idea was that a reduced response time requires the user to dedicate more cognitive resources to the task, leading to more working memory usage (i.e., increase in cognitive load). We also incorporated a reading comprehension task, since similar results were shown in [17] for differing levels of reading complexity. The reading comprehension task consisted of the user reading a passage over a fixed length of time and then answering questions in regards to the contents of that passage. The selected reading passage was chosen because it scored a 35.2 on the Flesch reading-ease test [65], indicating it was of college-level difficulty and a reading complexity our subjects likely encounter seeing that they are chosen from the university population. We only presented one passage to the user because we predicted it would cause a response in the thermal signal, but also wanted to see if there would be differing responses in the blink rate and inertial measurements when compared to the Stroop task. In [66], researchers showed that blink rate varies across different reading contexts so, we were interested to see if there was a difference between reading a passage versus quick single-word appearances. To also verify the blink sensor over a longer, more controlled stimuli, we added a 2.5 minute movie clip from a cartoon for the user to watch. The movie clip is different from the reading comprehension and Stroop tasks because it doesn't require any user input.

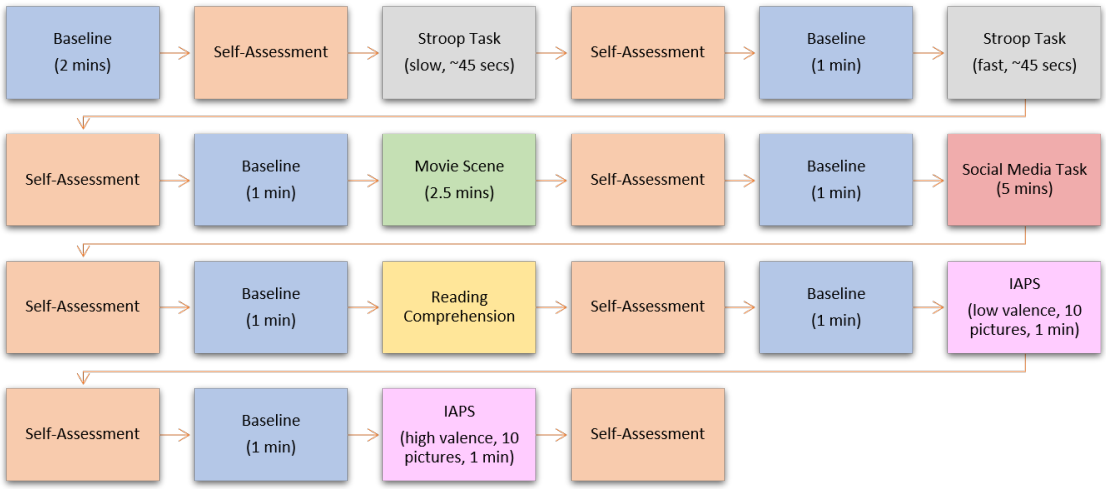

Figure 4-1: Validation Experiment Timeline

For a more valence-controlled stimuli, the International Affective Picture Systems (IAPS) was used [67] for two slideshow sections in the experiment. The IAPS is

a "large set of standardized, emotionally-evocative, internationally accessible, color photographs that includes contents across a wide range of semantic categories." Each photograph has a quantitative measurement of arousal and valence mean on a scale from 1 to 9 with a standard deviation associated with it. For our experiment, we chose a set of 10 pictures that were high valence (greater than 6) and a set of 10 that were low (less than 4). Each set of pictures is shown as a slideshow, 6 seconds for each picture, totalling 1 minute per set. The sets were played sequentially with an intermission in-between to allow the user to return to baseline.

In [68], it was shown that when browsing your Facebook news feed, physiological signals that correlate to high positive valence and high arousal were seen. Given that this is the least controlled of all other tasks that we included, but can potentially offer a significant physiological response, we thought it would be interesting to add it to the experiment. Therefore, we added a section where the user was able to browse their social media for 5 minutes and if they did not have a social media account, they can pass that section of the study.

At the start of the experiment and in-between each section, a baseline video is played to allow the physiological signals to return to a baseline prior to the text task. The baseline videos were of a long nature scene with no dialogue and slow paced, soothing music, with the intention to calm the users. The video was cut into several 1 minute chunks, with the exception of the introductory baseline video being two minutes to properly stabilize the user prior to the experiment. After each task, and prior to the first task, the user is asked a series of self-assessment questions (figures B-1, B-2, and B-3) to measure their affective reaction. For part of this assessment, we use the Self-Assessment Manikin (SAM), which is a "non-verbal pictorial assessment technique that directly measures the pleasure, arousal, and dominance associated with a person's affective reaction" [69]. We also ask the user to self-report questions on their levels of stress, cognitive load, and other internal states [70, 71].

Sensors

The Captivates system was used throughout the entire experiment, with the data being wirelessly transmitted to a server that was also located within the room. We were able to monitor the signals in another office to ensure the system was stable. Prior to each experiment, the optics for the thermal sensors and blink sensor were wiped down with isopropyl alcohol to ensure there was no foreign oil residue absorbing or reflecting the received infrared light and dampening the signals.

In addition to our system, we added three other commercial sensor suites for com-

parison and validation. The first was a camera that was mounted in front of the user with the intention of picking up head dynamics and blink rate to compare with our inertial and blink sensor. We also used a Shimmer [72] device to capture electrodermal activity (EDA) since in [22], EDA rolling averages and frequency power showed promising results in quantifying cognitive load levels. For the Shimmer, we placed two electrodes, one on the index finger and one on the middle finger on the non-dominant hand, as recommended by the Shimmer user guide. Finally, a Zephyr Performance System band [73] was used, since in [74], it was shown "that respiratory behavior generally reflects cognitive processing and that distinct parameters differ in sensitivity: While mentally demanding episodes are clearly marked by faster breathing and higher minute ventilation, respiratory amplitude appears to remain rather stable." These sensor systems will allow for proper validation of our sensors and of the intended internal responses from each individual task.

4.2 Results

Given the limited dataset collected due to the quarantine-imposed interruption of the study, we will focus primarily on qualitative measurements of the system's performance on an individual basis and will attempt to generalize as much as possible. However, given the limited amount of users, we cannot accurately generalize broadly across subjects. The primary goal will be to validate the operation and integrity of the system's sensors and relate the responses to other sensing modalities and cognitive states calculated from the self-assessment.

4.2.1 Self-Assessment

Prior to discussing each individual sensor channel, it is important to analyze the self-assessment results to build an understanding of the internal cognitive states experienced during this experiment. The sample set for these self-assessment results is rather small, so the results will have some error until we are able to increase the sample size. For the scope of this validation, we will focus on quantifying levels of stress, mental effort, happiness, and excitement. In Figure 4-2, we plot the z-normalized responses of the 6 participants— the x-axis is an indication of how strongly they feel that respective state.

As expected, increasing the speed of the Stroop task does considerably yield an increase in mental effort and excitement, but there isn't a considerable increase in

stress. The movie scene was fairly calm across these four states, so we shouldn't see a considerable response on our sensor channels during this task, at least not any due to these cognitive states. The social media task, though, does show a decrease in mental effort required but an increase in happiness compared to the movie task. As for the mental effort during the reading comprehension task, the average rating lies between the slow and fast Stroop tasks but, due to the variance, we are likely to see differing responses across subjects, some stronger than others. Lastly, as expected, the lower valence IAPS did rate significantly lower in happiness and excitement than the higher valence IAPS, further affirming the desired outcome.

4.2.2 Sensor and Experiment Evaluation

In this section, we will review the temperature, blink, and inertial signals from the Captivates System and compare them to the commercial sensors where applicable. We will look at the signal robustness, comment on any edge cases that were experienced, and particular signal artifacts that are of interest.

Thermal Data

Each temperature sensor was sampled at 10Hz and the calibration described in Section 3.5.2 was applied to the data after the experiment. Given that skin temperature is a slow changing signal, a fifth-order low-pass filter was first applied with a cut-off frequency of 0.5Hz on the raw data to remove any movement artifacts or potential electrical noise. One subject's filtered temperature profile for the duration of the experiment is shown in Figure 4-3, with the different parts of the experiment labelled and colored. Its immediately apparent that the temperature signal requires time to stabilize at the start of the experiment, likely a result of the thermopile housing temperature warming up to the incident radiation. As mentioned in Section 3.5.2, this can be compensated for with more accurate calibration that also modulates the housing temperature but that is beyond the scope of this thesis. That being said, as the temple temperature stabilizes, we see a temperature around 310 kelvin, which is slightly higher, but close to, the average max temporal skin temperature of a healthy subject which is 309 kelvin [75]. This indicates that our calibration is close but the tolerance is wider than commercial non-contact thermometers which are often +/- 1-2 kelvin. Also, as expected across all subjects, the nose temperature was below that of the temporal region, but varied in magnitude across individuals and over the duration of the experiment. This should often be the case since the temporal region is

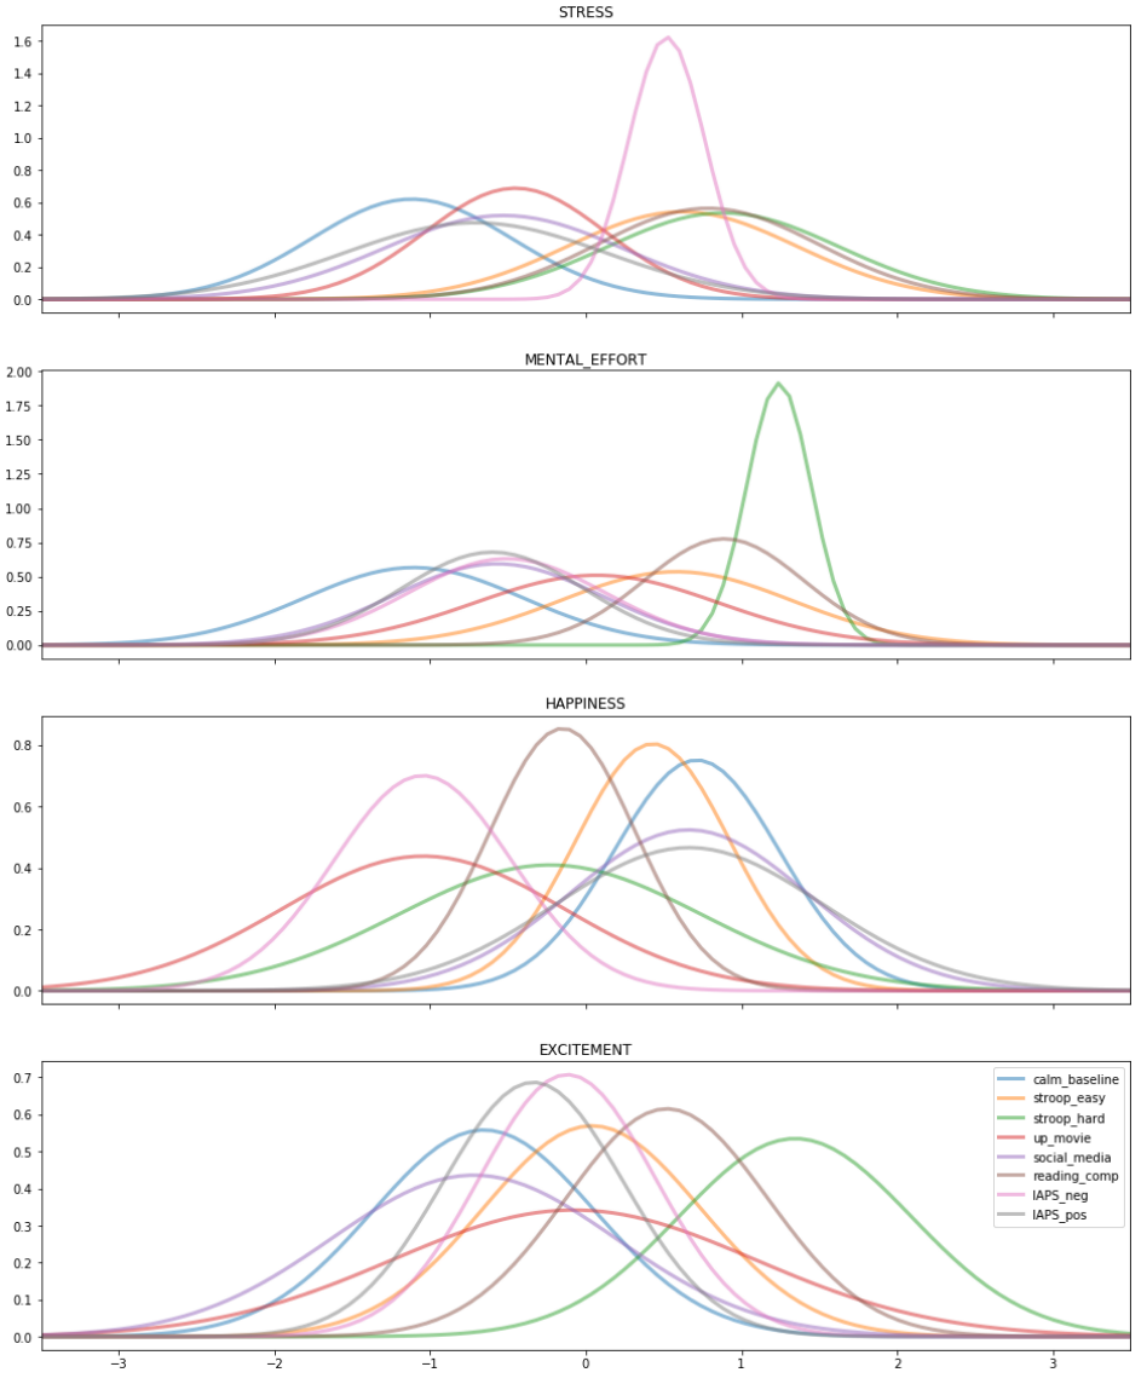

Figure 4-2: Self-Assessment Results (Z-normalized)

a measurement of the body's core temperature [19] and no other region should exceed that unless influenced by an external source e.g., external heat source illuminating the nose.

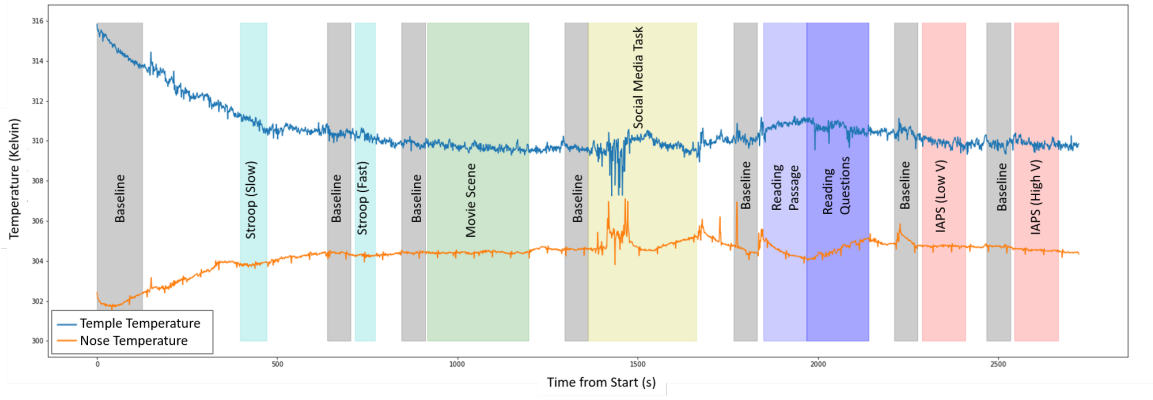

Figure 4-3: Temperature Profile From Both Sensors of a Single Subject Across the Experiment

Given that we are mainly concerned with the temperature difference between the temple region and the nose, we calculate the differential after applying an additional first-order high-pass Butterworth filter to attenuate the initial temperature stabilization. The differential signals across all six subjects are shown in Figure 4-5. Note that the test duration varies slightly for each subject because a few transition windows between tasks that required user input, and also one subject did not have a social media account, so they bypassed that section. We also did not see any considerable noise artifacts besides half-way in the second subject's test where you can see a few seconds of high volatility due to the subject re-positioning the glasses on their face.

Our results show that there do exist changes in the temperature signal across different contexts, but at least from these 6 participants, there is no generalizable trend across individuals. We did not see a significant difference in nose temperature between different speeds of the Stroop tasks for all participants, as was shown in [17]. However, as explained in Section 4-2, there is a difference in perceived mental effort between the two tasks, so one assumption is that the Stroop task wasn't long enough to be able to see a change in the temperature signature. Some fluctuations do exist, up to a 1 kelvin change for one user between the two Stroop tasks, but for most, there wasn't anything significant.

It is apparent, though, that for the reading comprehension task, all subjects saw a rise in their nose temperature, which ranged from 0.2 to 1.5 kelvin, which is inline with expectation given the task should result in greater working memory usage (i.e.,

cognitive load). Further evidence of that is seen with increased skin conductance during reading comprehension, indicating increased cognitive load [22]— an example skin conductance signal is shown in Figure 4-4 for a single subject. We did not see anything statistically relevant with the respiration data, but we surmise that this was likely due to the short duration of our experiment and the rolling window that is used to compute the respiration rate. We also saw a much greater rise in the temperature differential within the reading comprehension task than that of the social media task, a result that was supported by differences in perceived mental effort between the two (Figure 4-2) with the social media task rating less.

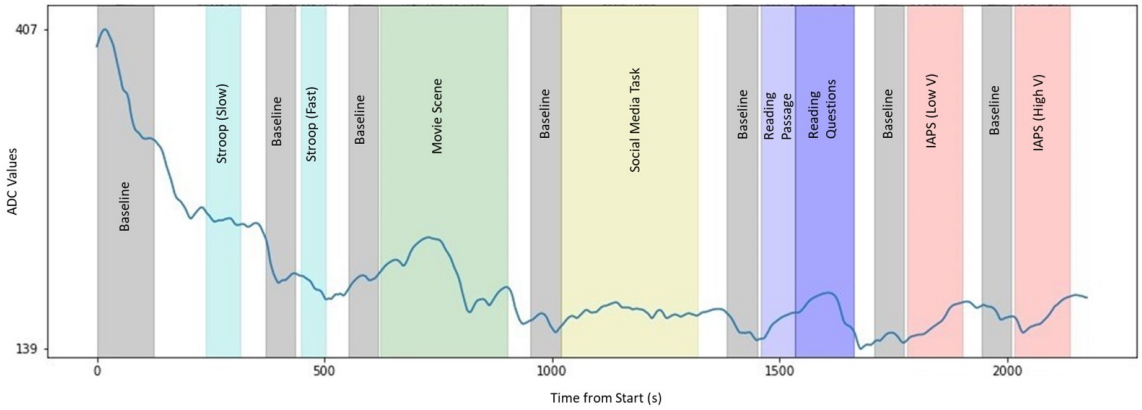

Figure 4-4: Skin Conductance Throughout Experiment for Single Subject

Blink Data

The blink signal through our glasses was sampled at 1 kHz and at 8-bit resolution to reduce the packet size, although 12-bits are achievable. We did not modulate the emitter light, so we are measuring full-power reflectance for this experiment. The lack of modulation may result in a saturated sensor for a few participants but it reduced the complexity of tuning each blink sensor to the individual. We did, however, host the experiment in a room with only LED lighting to limit the amount of environmental near-infrared light other than what is emitted by the blink sensor.

An excerpt of a subject's filtered blink signal is shown in Figure 4-6. For filtering, we passed the raw signal through a second-order Butterworth filter with a cut-off frequency of 0.5Hz to remove any constant light sources. Overlaid with the signal, we processed the camera data using OpenFace [76] and retrieved the binary blink classification for comparison. As can be seen in Figure 4-6, spikes in the signal correlate well with OpenFace blink detection, indicating that there is a sudden change

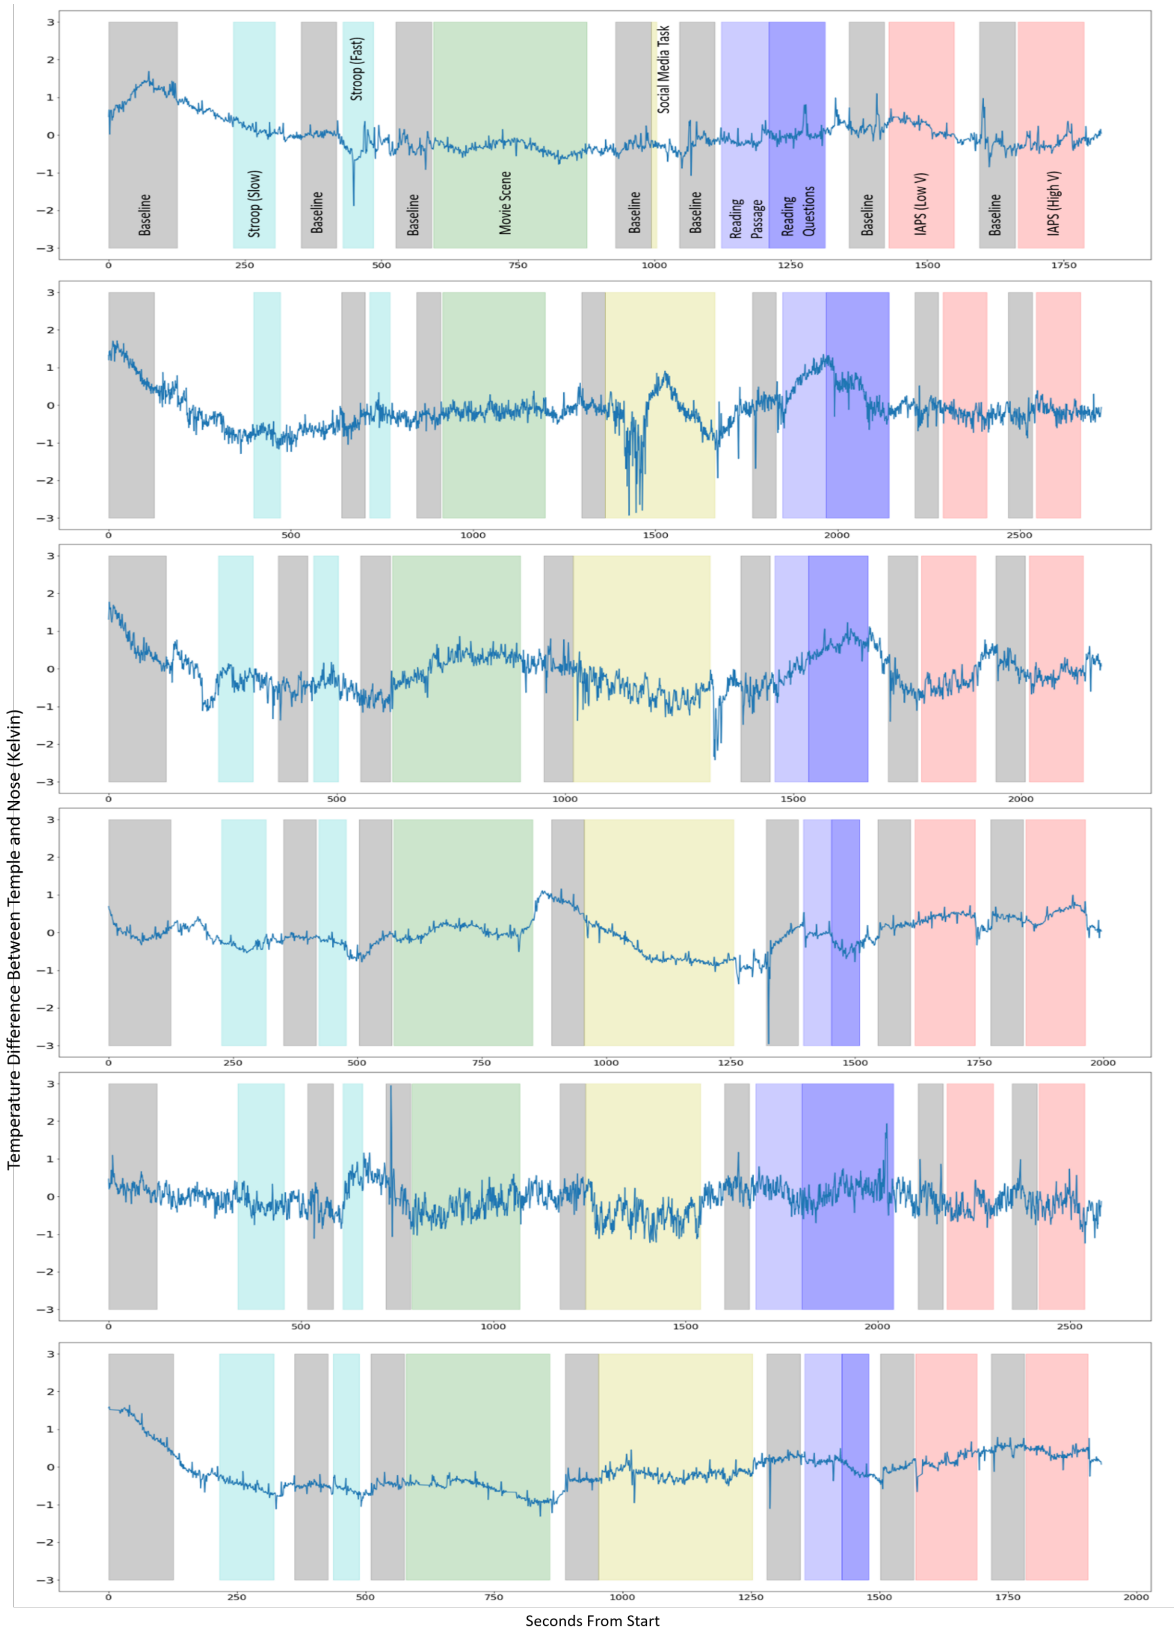

Figure 4-5: Band-passed Temperature Differential between Temple and Nose of 6 Subjects

in reflectance when a blink occurs. However, what is also evident is that there are spikes in the signal that OpenFace does not classify as blinks using the camera data. The average human blink is around 150ms [77], so we assumed that a 30Hz sampling of the camera would be enough, but it seems that the classifier built into OpenFace requires several consecutive positive blink samples to classify the blink. We deduce that these are eye blinks because the lighting in the room is LED-based with no illumination in range of the IR sensors, so there should not be any high-frequency events in the signal besides attenuated electrical noise. For future experimentation, we will increase the camera sampling to validate this assumption but for the purpose of this analysis, we will compare as-is.

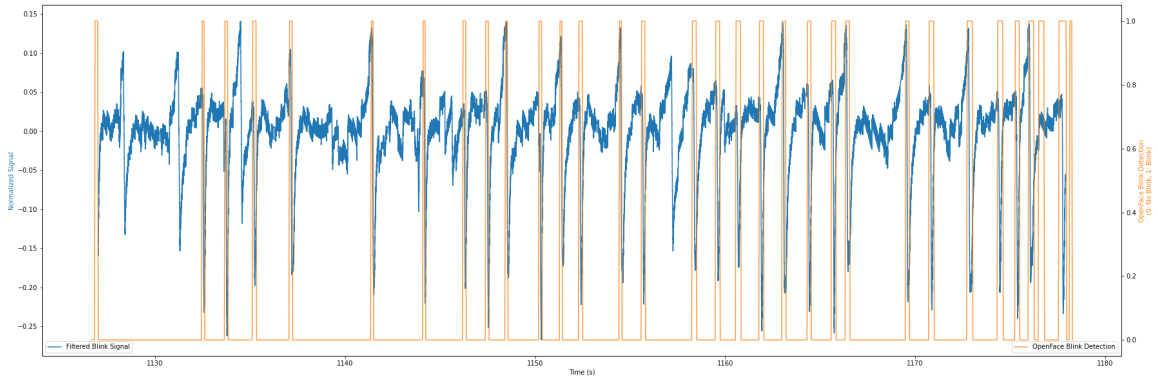

Figure 4-6: Excerpt Blink Signal with OpenFace Blink Classification Overlaid

After the signal is filtered, we apply simple thresholding to see how well the signal correlates with the OpenFace positive blink detections and to see if a static threshold can be used across subjects. As can be seen from Figure 4-7, the OpenFace classification aligns well with our simple thresholding and our system is also able to pick up on other blink events that cannot be validated with OpenFace software but have been witnessed in the raw camera feed data. When attempting to generalize the thresholding technique across users, its quickly apparent that a dynamic thresholding method is warranted, given the varied signal response across individuals. In addition, dynamic modulation of the emitted light is also needed, since for at least one subject, we saw that the diode was saturated and no discernible blinks were detected. This is likely a cause of the varied geometry across faces, and even though we attempted to angle the blink sensor in the most optimal trajectory towards the eye for average face dimensions, it is clearly not suitable for all, so a custom-tailoring of sensor positioning and dynamic signal strength adjustment is warranted. Further, signal robustness can be improved by synchronizing a pulse-width modulated signal with the sampling of

the analog-to-digital converter to collect a sample when the eye is exposed and when it is not. Calculating the differential and filtering allows you to remove any constant light sources in the room (i.e., sunlight) or light modulated at different frequencies. Further validation will need to be done at a higher camera sample rate for a more accurate comparison.

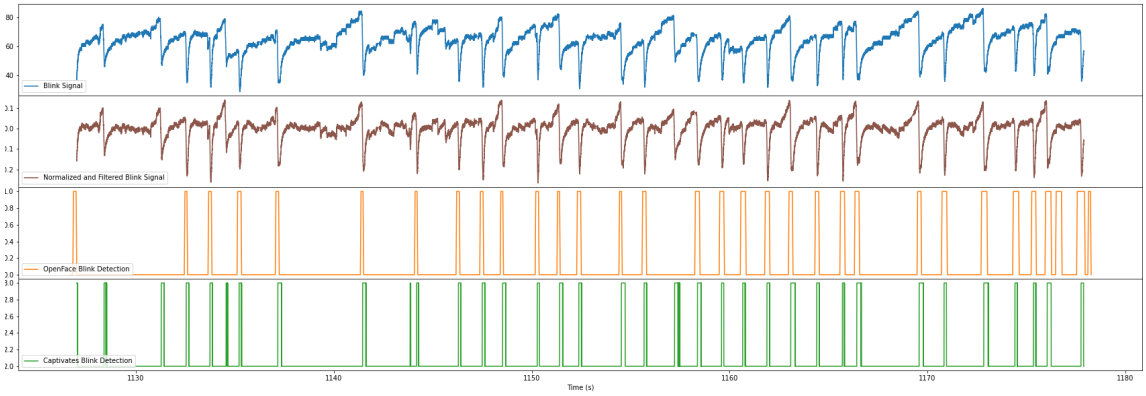

Figure 4-7: Excerpt of Blink Thresholding Compared to OpenFace Blink Detection

Inertial Data

A subject's head pose data calculated from our IMU is shown in Figure 4-8. You can see that when a subject is focusing on a stimulus, there is a reduction in head dynamics when compared to task transitions. The two exceptions we saw to that was when the user looked down to grab their phone for the social media experiment and during the reading comprehension section when the user was actively reading and engaging with the questions being asked. These results, although brief, offer more justification into coupling an inerital sensor with other sensing modalities to look at the physiological differences across contexts.

To further validate the head dynamics as calculated by the IMU, we compared the on-board inertial measurements with ones calculated using the camera feed through the OpenFace software (Figure 4-9) [76]. As you can see, the dynamic movement periods correlated well with what was shown in Figure 4-8. There are differences in magnitude, but that is likely due to jitter in the way OpenFace first calculates landmarks on the face and then tracks their movement, limited by the resolution of the camera feed. As for why the static levels are not identical during low-movement periods, OpenFace tracks head pose in reference to the computer screen, while the IMU tracks head pose in reference to Earth's gravity and magnetic field. If the 3D

tracking system was activated, we would be able to transform between coordinate systems if reference to a particular stimuli was required.

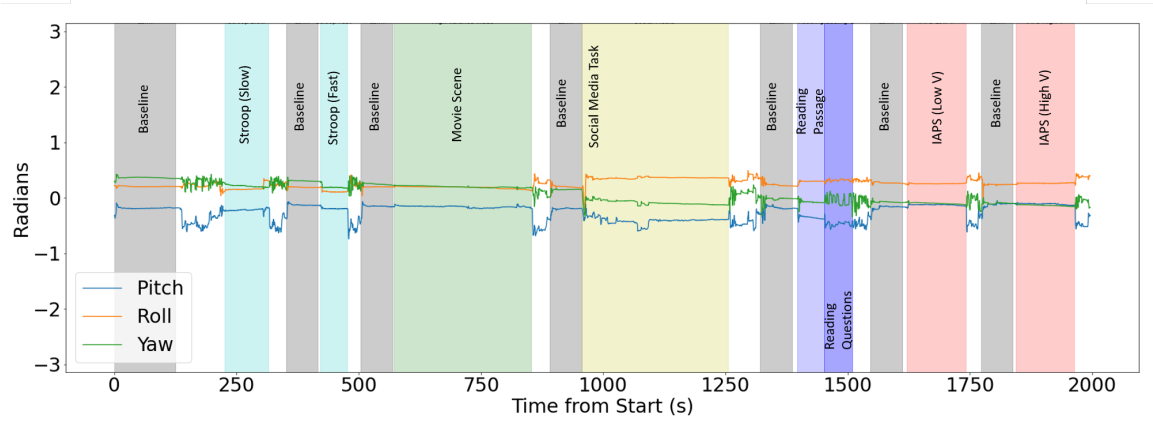

Figure 4-8: Head Pose as calculated by onboard IMU for Single Subject

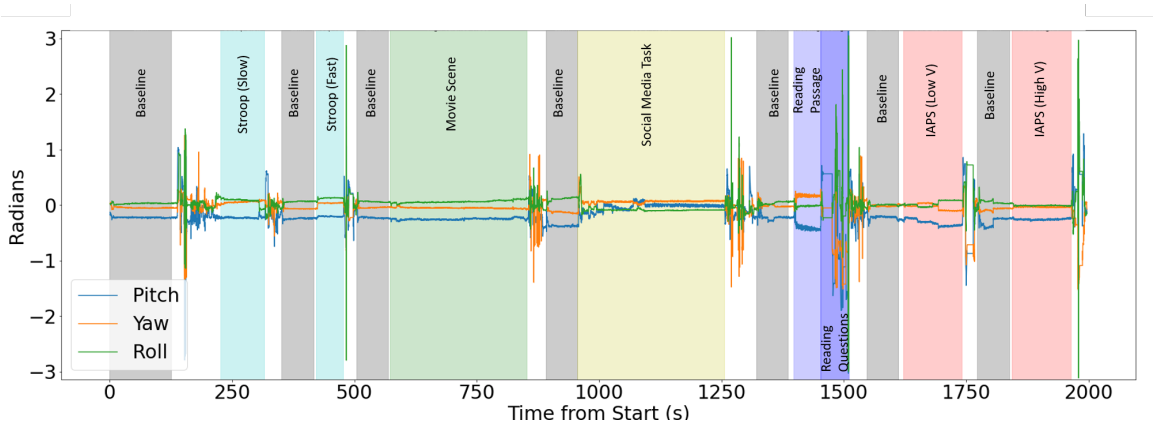

Figure 4-9: Head Pose calculated for Single Subject using Camera

4.3 Discussion

In Section 4.2, we describe our preliminary validation study's data, focusing on the validity of each sensor channel individually and where possible, leveraging self-assessment results and commercial sensing channels to prove how well our observed trends follow that of prior work. Of the data that we collected, we found that for many of our tasks, the signal response varied in magnitude and direction across individuals. Although we don't have enough information to form final conclusions, we have at least enough sensor data to add insight on how to approach predicting user's internal state and how to structure this system for ambulatory measurements.

4.3.1 *Is a generalized model possible?*

Its often repeated in research studies that if there is a strong stimulus in a room, subjects will react to it and we would see correlated trends in physiological responses. But, for every-day life, the stimuli vary across strengths and contexts with the physiological response likely dependent on a person's short-term experience leading up to the point of stimulus exposure— this was the whole motivator on why we wanted to create this system in the first place. We cared less to understand what a person's signal was at a specific moment, but more on the measure of dependence of that response to past moments. We theorize that through a probabilistic approach, we can predict the certainty of a particular cognitive state through historical physiological data, forming a model of human cognitive states. The question is, how generalizable across subjects can this model be?

This preliminary study gives a snapshot to the validity of each sensor in relation to particular emotions and cognitive processes, but more work needs to be done on characterizing the response. However, given what we've seen from the temperature and inertial measurements, there are significant signal differences between different events to where, if given enough data, it's likely that we can measure a person's behaviour and certain cognitive states (e.g., how much they're loading their working memory). But enough differences across subjects exist, particularly in the observed thermal measurements, as this strongly suggests that an individualized modelling approach is the ideal path forward. As was shown in Figure 2.1, the thermal responses across individuals for the exact same sequence of tasks varied in amplitude, and at times, in direction. That's not to say generalized trends don't exist, since we did see significant difference in thermal responses during the social media task and reading comprehension, which were an indicator of cognitive load as supported by the self-assessment results. But even though we saw a hint of correlation, the issue is that there was enough variability throughout the experiment that indicates models will need to be tuned to the user and can perhaps change for individual users over time. More than likely, there will be events that a user won't respond to as strongly, if at all, where as other individuals will. The path forward is to form an individualized model that is initially based on some priors and then updated as a user's data and daily self-report results is made available.

4.3.2 Improvements to Experimental Design

During this round of testing, we discovered a few modifications that can be done to improve the validity of the experiment. One necessary change is to edit each test scenario to be at least 3 minutes long, giving enough time for the user's physiology to respond. We would have liked to see the temperature responses stabilize during each task but the task lengths were too short to allow for that. In addition, we also need to have the user sit with our system for at least 10 minutes at the start of the experiment so that they can acclimate to the environment. This will reduce the dramatic temperature fluctuations seen in some users at the start of the experiment (Figure 4-5). Another modification to the experiment would be to use a reading passage with a lower Flesch reading-ease test score since that would result in a passage that is more difficult to read, thus a greater temperature and blink rate response should result. Although we did see a response for some subjects during the reading comprehension task, using a more complex passage would ensure the response is statistically relevant. Lastly, the frame rate of the video capture needs to be increased from 30Hz to at least 60Hz to decrease the likelihood of not capturing an eye blink, offering a better truth comparison to our custom blink detection system,

Chapter 5

Conclusion and Future Work

In this work, we presented the Captivates system, a wearable eyeglass platform that attempts to fill the sensing gap that existing smart environment infrastructure cannot provide, allowing for the capture of physiological signals that have shown correlation to specific cognitive states. We presented the work we have done on making this platform robust and practical for ambulatory applications. Our design includes systems for measuring face temperature, blink rate, head movement, activity prediction, and estimating 3D location. There are also LED subsystems that can be used as actuators for notifications and interactive experiences. The design is extensible for other researchers to tailor to their own needs without needing to design an entire new wearable system. We have also shown preliminary validation results that show promise in leveraging this system in creating a new type of responsive environment that takes the human element into consideration.

Once the risks of COVID-19 diminish, we plan on continuing our user study to increase the sample size for more conclusive results on the cognitive state-measuring aspect of our system. When the study does resume, we will edit the experiment in accordance to what is described in Section 4.3.2. We are also currently in the process of producing 30 of our systems with the intent of distributing them to users as they go about their daily lives. We will be creating a supplementary, small, data logging device, where all the data from the Captivates system will be streamed and saved; this will allow for users to be completely mobile with our system and not need any static networking infrastructure within their space. Given that much of our work is now confined to our homes, we are interested in seeing how our physiological responses vary across days, coupled with cognitive state self-assessment for accurate modelling. With this at-home study, we also plan to evaluate the long-term usability of our system. With daily self-assessments asking the users to rate their comfort

and elaborate on any specific issues with the the physical fit of the system, we seek to gain a better understanding of the ergonomics of our device and the feasibility of using it for long-term studies, as well as to iterate and improve on the overall design. Throughout these experiments, we also seek to build an understanding of users' comfort when it comes to their physiology constantly being sensed and logged. Ideally, we want to create a system that the users feel they are in control and that they can stop the sensing at any time. However, the only way to understand how to build these types of systems is through user studies where we can ask the individuals who could be the potential consumers. Once our studies are complete, we plan on distributing a subset of our systems to interested researchers for use in similar or diverse applications.

After we are able to collect a large enough data repository, our next effort will be to create a targeted experience through a personalized modelling effort. In the education context alone, its been proven that a personalized learning policy outperforms that of a generic one. In [78], the researchers created a system that modulated engagement during learning based on verbal and nonverbal affective cues of children. They found that when learning vocabulary, a personalized policy leads to a significantly better learning outcome. With our system being used to collect various physiological signals over long periods of time, we are able to form more comprehensive models that better inform learning policies, allowing for an adaptive online education system that's personalized to the learner.

The Captivates system is one part in a larger narrative of creating a more responsive environment. This includes additional sensing nodes placed densely throughout an environment that inform a more accurate description of a space and on what we individually prefer. This thorough sensing will lead to better actuation of environmental resources, such as lighting or heating. Through the Captivates project, we are one step closer in creating a truly tailored responsive environment.

Appendix A

Schematics

A.1 Front PCB Schematics and Bill of Materials

MCU BOARD CONNECTOR

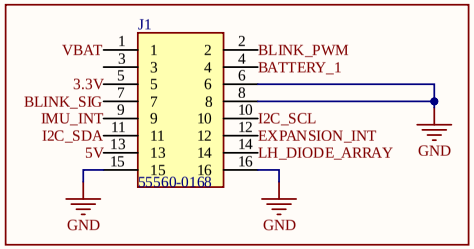

POWER BOARD CONNECTOR

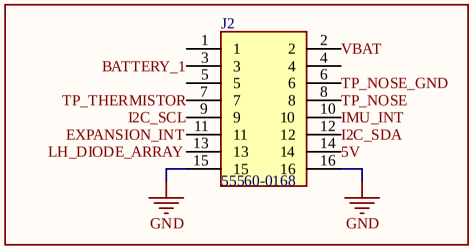

Title	CONNECTORS		
Size: A	Number:		Revision: 2.0
Date:	5/04/2020	Sheet of	
File:	D:\OneDrive - Massachusetts Institute of Technology\Sch...	Drawn By:	PATRICK SWOLEK

THERMOPILE CIRCUIT

THERMOPILE CIRCUIT CONNECTOR

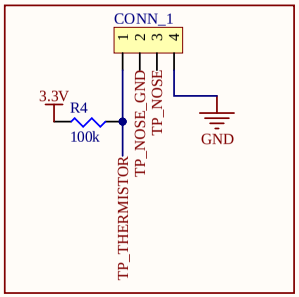

THERMOPILE

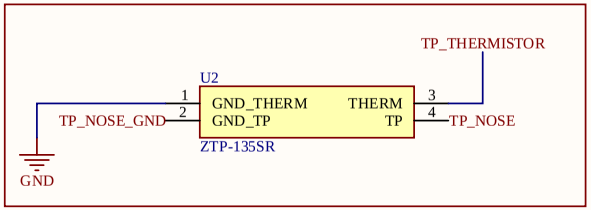

BLINK SENSOR CIRCUIT

BLINK CIRCUIT CONNECTOR

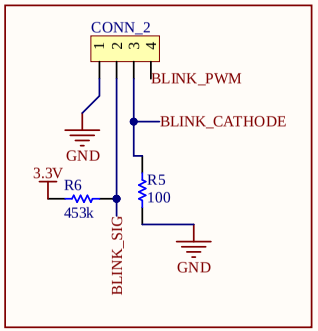

BLINK SENSOR

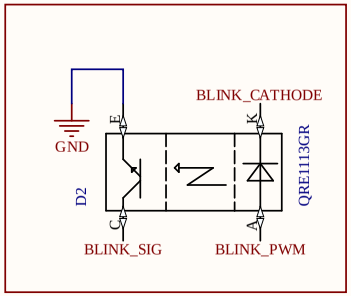

VIVE DIODE

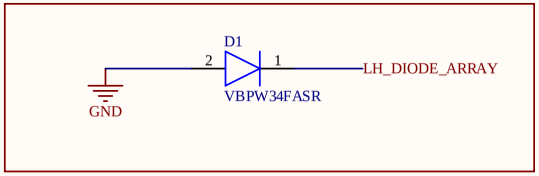

Title: FRONT FLEX SENSORS		
Size: A	Number:	Revision: 2.0
Date: 5/04/2020	Sheet of	
File: D:\OneDrive - Massachusetts Institute of Technology\Sch...	Drawn By: PARUL SHARMA DELWALEK	

LEFT-SIDE CIRCUIT

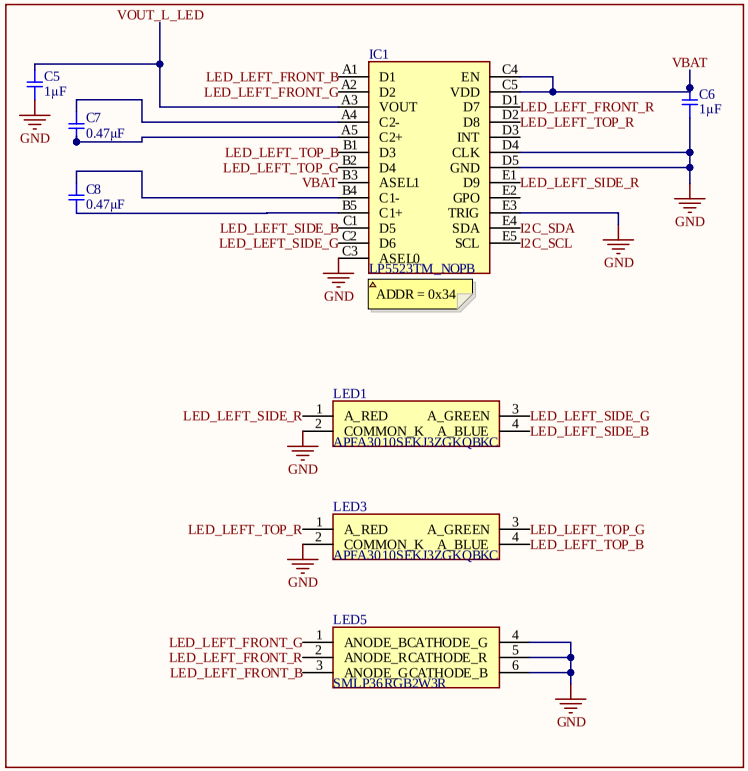

RIGHT-SIDE LED CIRCUIT

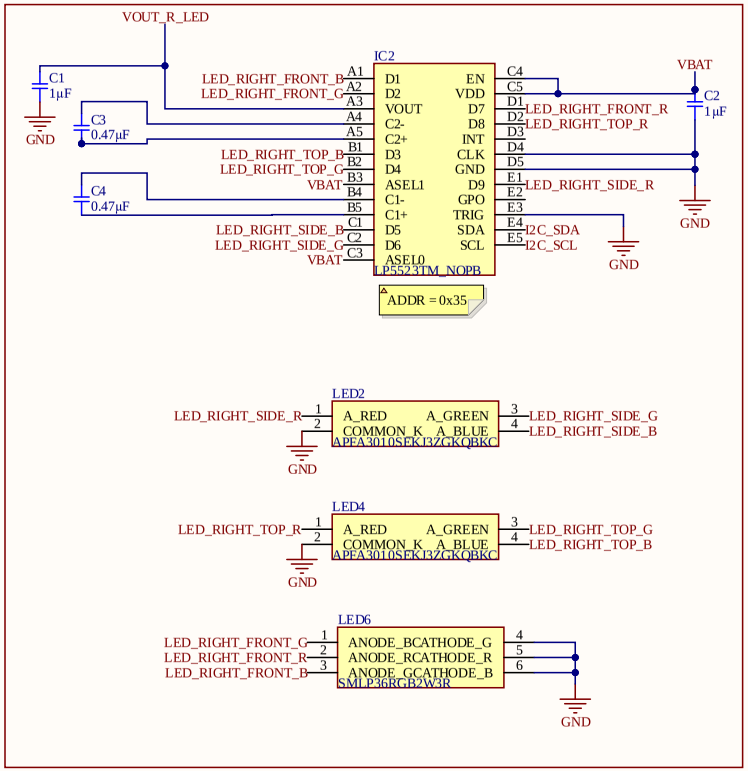

Title: LEDS		
Size: Letter	Number:	Revision: 2.0
Date: 5/04/2020	Sheet of	
File: D:\OneDrive - Massachusetts Institute of Technology\...\Schematic.SchDoc	Drawn By: PATRICK CHWALEK	

Footprint	Name	LibRef	Designator	Description	Quantity
0402R	1µF	GCM155C71A105KE38D	C1, C2, C5, C6	CAP CER 1UF 10V X7S 0402	4
0402R	0.47µF	GCM155C71A105KE38D	C3, C4, C7, C8	CAP CER 0.47UF 10V X7S 0402	4
VBPW34FAS	VBPW34FASR	VBPW34FASR	D1	Diode	1
QRE1113GR	QRE1113GR	QRE1113GR	D2	QRE1113 Series 30 V 50 mA Miniature	1
BGA25C40P5X5_226X226X68	LP5523TM_NOPB	LP5523TM_NOPB	IC1, IC2	Integrated Circuit	2
555600168	55560-0168	55560-0168	J1, J2	Connector	2
KPFA-3010RGBC-11	APFA3010SEKJ3ZGKQBKC	APFA3010SEKJ3ZGKQBK	LED1, LED2, LED3, LED4	LED	4
ROHM-1510-(0604)	SMLP36RGB2W3R	SMLP36RGB2W3R	LED5, LED6	LED	2
SMD-0402-RES	100k	ERJ-2GEJ104X	R4	RES SMD 100K OHM 5% 1/10W 0402	1
0402R	100	Resistor	R5		1
0402R	453k	Resistor	R6	RC0402FR-07453KL	1
ZTP-135SR	ZTP-135SR	ZTP-135SR	U2	Undefined or Miscellaneous	1
					24

A.2 Primary MCU Side PCB Schematics (Left-Side) and Bill of Materials

TOUCH SENSING HEADERS

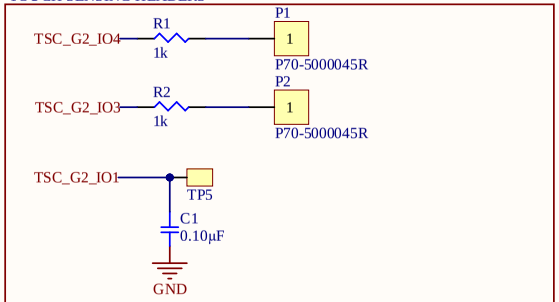

Note: reduce sampling capacitor if less sensitivity and fast acquisition time is needed
Note: do not use tantalium capacitors (per ST)
Note: reduce the capacitor on TSC_G2_I01 to increase sensitivity

FLEX CONNECTOR

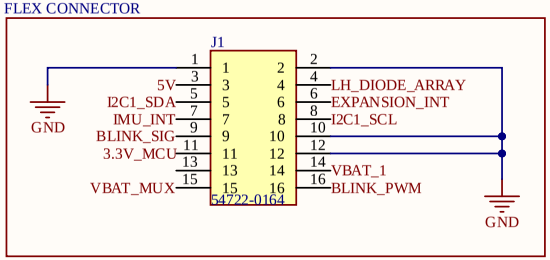

TEST POINTS

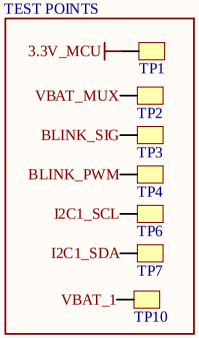

BREAKOUT PADS

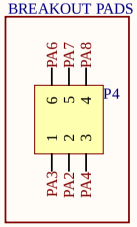

PROGRAMMER USB BREAKOUT

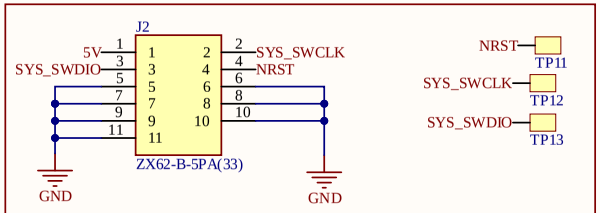

PROGRAMMER TAG BREAKOUT

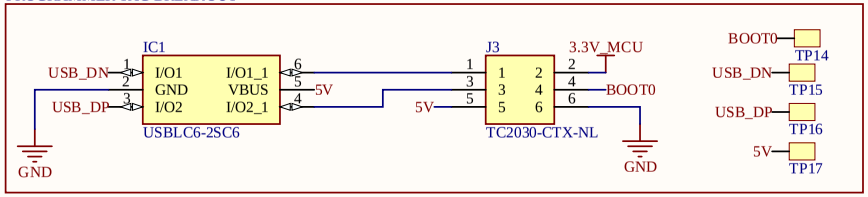

Title: EXTERNAL CONNECTORS		
Size: A	Number:	Revision:
Date: 3/26/2020	Sheet of	
File: D:\OneDrive - Massachusetts Institute of Technology\Sch...	Drawn By: PATRICK CHWALEK	

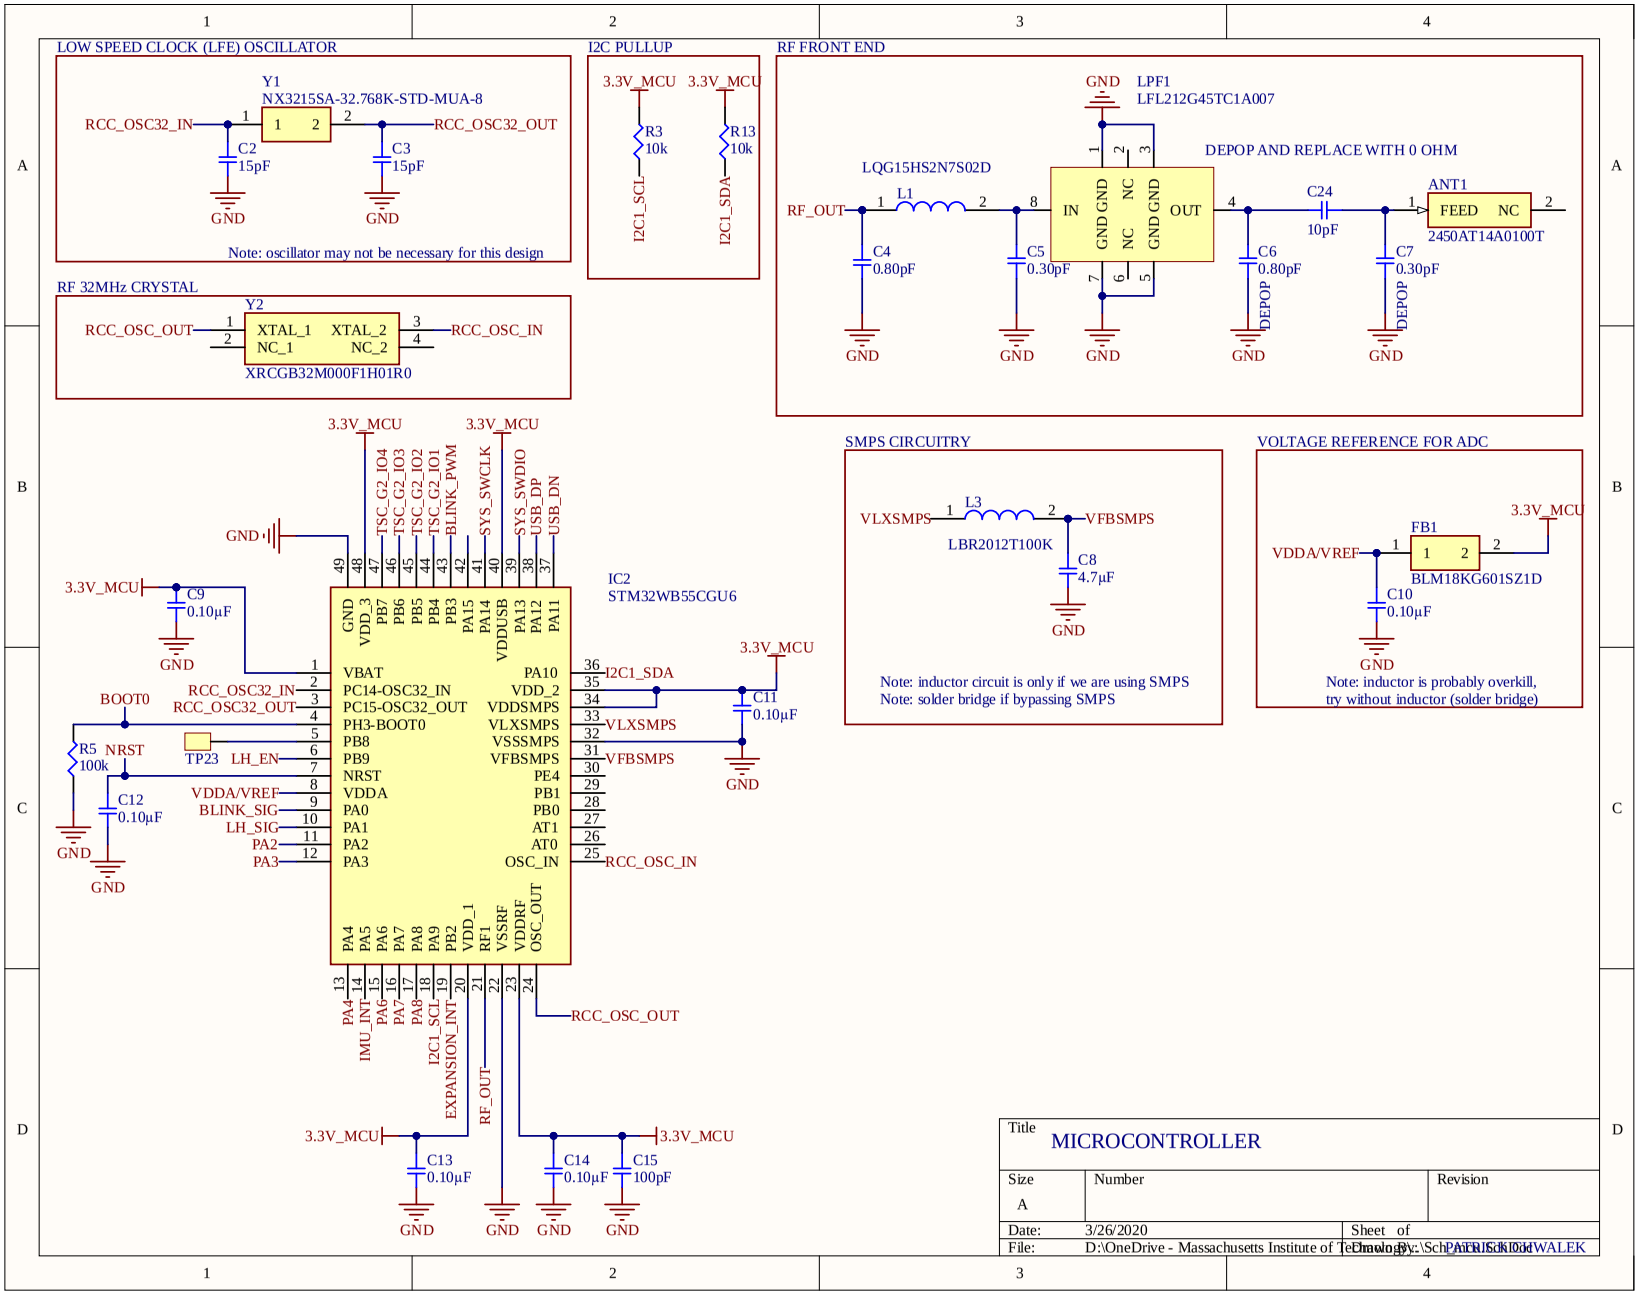
LOW SPEED CLOCK (LFE) OSCILLATOR
Y1
NX3215SA-32.768K-STD-MUA-8
RCC_OSC32_IN
RCC_OSC32_OUT
C2
15pF
C3
15pF
GND
Note: oscillator may not be necessary for this design
I2C PULLUP
3.3V_MCU
R3
10k
R13
10k
I2C1_SCL
I2C1_SDA
RF FRONT END
LPF1
LFL212G45TC1A007
DEPOP AND REPLACE WITH 0 OHM
LQG15HS2N7S02D
L1
RF_OUT
C4
0.80pF
C5
0.30pF
C6
0.80pF
C7
0.30pF
C24
10pF
ANT1
FEED
NC
2450AT14A0100T
DEPOP
RF 32MHz CRYSTAL
Y2
RCC_OSC_OUT
RCC_OSC_IN
XTAL_1
XTAL_2
NC_1
NC_2
XRCGB32M000F1H01R0
IC2
STM32WB55CGU6
VBAT
PC14-OSC32_IN
PC15-OSC32_IN
PH3-BOOT0
PB8
PB9
NRST
VDDA
PA0
PA1
PA2
PA3
PA10
VDD_2
VDDSMPS
VLXSMPS
VSSSMPS
VFBSMPS
PE4
PB1
PB0
AT1
AT0
OSC_IN
OSC_OUT
VDD_1
RF1
VSSRF
VDDRF
C9
0.10µF
C11
0.10µF
C12
0.10µF
C13
0.10µF
C14
0.10µF
C15
100pF
R5
100k
BOOT0
NRST
TP23
LH_EN
VDDA/VREF
BLINK_SIG
LH_SIG
I2C1_SDA
VLXSMPS
VFBSMPS
RCC_OSC_IN
RCC_OSC_OUT
RF_OUT
IMU_INT
EXPANSION_INT
I2C1_SCL
TSC_G2_IO4
TSC_G2_IO3
TSC_G2_IO2
TSC_G2_IO1
BLINK_PWM
SYS_SWCLK
SYS_SWDIO
USB_DP
USB_DN
SMPS CIRCUITRY
L3
LBR2012T100K
C8
4.7µF
Note: inductor circuit is only if we are using SMPS
Note: solder bridge if bypassing SMPS
VOLTAGE REFERENCE FOR ADC
FB1
BLM18KG601SZ1D
C10
0.10µF
Note: inductor is probably overkill,
try without inductor (solder bridge)
Title
MICROCONTROLLER
Size
A
Number
Revision
Date:
3/26/2020
Sheet of
File:
D:\OneDrive - Massachusetts Institute of Technology\...\Schematic1.SchDoc
Drawn By:
PARKER WALEK

POWER REGULATION

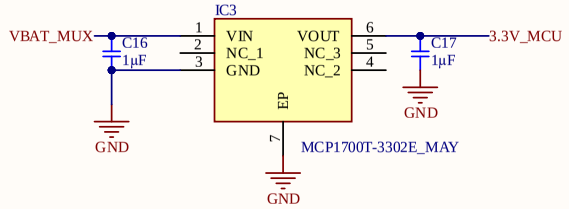

VIVE LIGHTHOUSE POWER

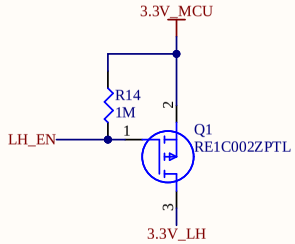

CHARGING CIRCUIT

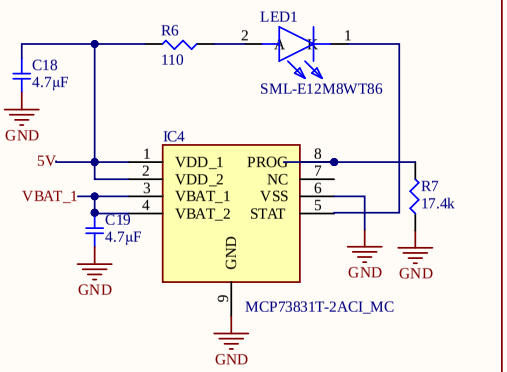

Charge Current : 1000V / 17.4k = 57.47mA

BATTERY INPUT

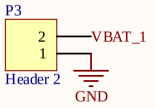

Title: POWER		
Size: A	Number:	Revision:
Date: 3/26/2020	Sheet of	
File: D:\OneDrive - Massachusetts Institute of Technology\Sch...	Drawn By: PAWEŁ CHWALEK	

LIGHTHOUSE DETECTOR

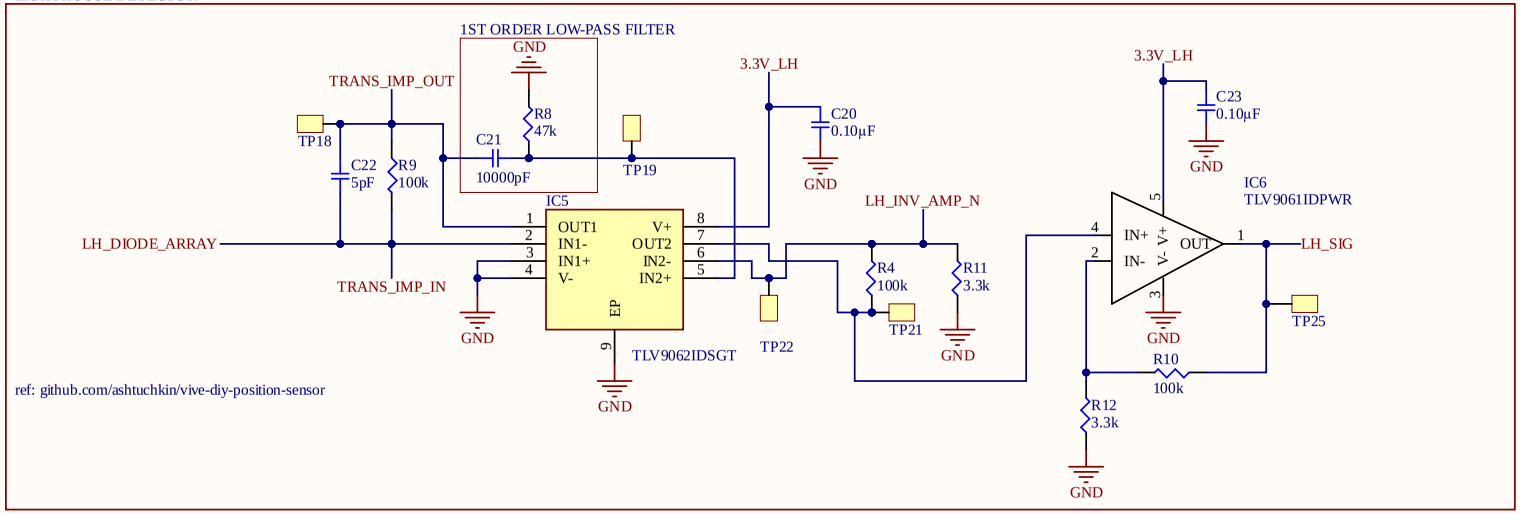

VIVE DIODE

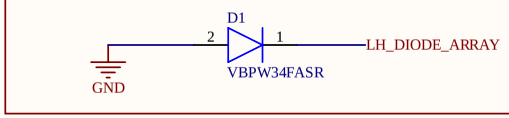

Title: SENSORS		
Size: A	Number:	Revision:
Date: 3/26/2020	Sheet of	
File: D:\OneDrive - Massachusetts Institute of Technology\...	Drawn By: PATRICK CHEVALEK	

Footprint	Name	LibRef	Designator	Description	Quantity
ANTC1608X40N	2450AT14A0100T	2450AT14A0100	ANT1	Antenna	1
SMD-0402C	0.10µF	GCM155R71C10	C1, C9, C10, C11, C12, C13, C14, C20, C23	CAP CER 0.1UF 16V X7R 0402	9
SMD-0402C	15pF	GCM1555C1H15	C2, C3	CAP CER 15PF 50V NP0 0402	2
SMD-0402C	0.80pF	GRM1555C1HR	C4, C6	CAP CER 0.8PF 50V NP0 0402	2
SMD-0402C	0.30pF	GRM1555C1HR	C5, C7	CAP CER 0.3PF 50V NP0 0402	2
SMD-0402C	4.7µF	GRM153R60G4	C8, C18, C19	CAP CER 4.7UF 4V X5R 0402	3
SMD-0402C	100pF	GCM1555C1H10	C15	CAP CER 100PF 50V NP0 0402	1
SMD-0402C	1µF	GCM155C71A10	C16, C17	CAP CER 1UF 10V X7S 0402	2
SMD-0402C	10000pF	GCM155R71H10	C21	CAP CER 10000PF 50V X7R 0402	1
SMD-0402C	5pF	GJM1555C1H5R	C22	CAP CER 5PF 50V NP0 0402	1
SMD-0402C	10pF	GRM1555C1HR	C24	CAP CER 10PF 50V NP0 0402	1
VBPW34FASR	VBPW34FASR	VBPW34FASR	D1	Diode	1
BEADC1608X95N	BLM18KG601SZ1D	BLM18KG601SZ	FB1	Ferrite Bead	1
SOT95P280X145-6N	USBLC6-2SC6	USBLC6-2SC6	IC1	Integrated Circuit	1
QFN50P700X700X65-49N-	STM32WB55CGU6	STM32WB55CG	IC2	Integrated Circuit	1
SON65P200X200X90-7N	MCP1700T-3302E_MAY	MCP1700T-	IC3	Integrated Circuit	1
SON50P300X200X100-9N-	MCP73831T-2ACI_MC	MCP73831T-	IC4	Integrated Circuit	1
SON50P200X200X80-9N	TLV9062IDSGT	TLV9062IDSGT	IC5	Integrated Circuit	1
X2SON	TLV9061IDPWR	TLV9061DPWR	IC6		1
54722-0164	54722-0164	54722-0164	J1	Connector	1
ZX62B5PA33	ZX62-B-5PA(33)	ZX62-B-5PA(33)	J2	Connector	1
TC2030MCPNL10	TC2030-CTX-NL	TC2030-CTX-NL	J3	Connector	1
INDC1005X55N	LQG15HS2N7S02D	LQG15HS2N7S0	L1	Inductor	1
INDC2012X145N	LBR2012T100K	LBR2012T100K	L3	Inductor	1
LEDC1608X46N	SML-E12M8WT86	SML-	LED1	LED	1
LFL212G45TC1A007	LFL212G45TC1A007	LFL212G45TC1	LPF1		1
P70-5000045R	P70-5000045R	P70-5000045R	P1, P2		2
CAPTIVATE_BREAKOUT_P	Captivates_Breakout	Captivates_Brea	P4		1
SOTFL50P160X85-3N	RE1C002ZPTL	RE1C002ZPTL	Q1	MOSFET (P-Channel)	1
SMD-0402-RES	1k	ERJ-2GEJ102X	R1, R2	RES SMD 1K OHM 5% 1/10W 0402	2
SMD-0402-RES	10k	ERJ-2GEJ103X	R3, R13	RES SMD 10K OHM 5% 1/10W 0402	2
SMD-0402-RES	100k	ERJ-2GEJ473X	R4, R5, R9, R10	RES SMD 100K OHM 5% 1/10W 0402	4
SMD-0402-RES	110	ERJ-2RKF8062X	R6	RES SMD 110 OHM 5% 1/10W 0402	1
SMD-0402-RES	17.4k	ERJ-2GEJ103X	R7	RES SMD 17.4K OHM 1% 1/10W 0402	1
SMD-0402-RES	47k	ERJ-2GEJ473X	R8	RES SMD 47K OHM 5% 1/10W 0402	1
SMD-0402-RES	3.3k	ERJ-2GEJ332X	R11, R12	RES SMD 3.3K OHM 5% 1/10W 0402	2
SMD-0402-RES	1M	ERJ-2GEJ105X	R14	RES SMD 1M OHM 5% 1/10W 0402	1
NX3215SA	NX3215SA-32.768K-STD-	NX3215SA-	Y1	Crystal or Oscillator	1
XRCGB32M000F1H01R0	XRCGB32M000F1H01R0	XRCGB32M000	Y2	Crystal or Oscillator	1
					60

A.3 Secondary MCU Side PCB Schematics (Right-Side) and Bill of Materials

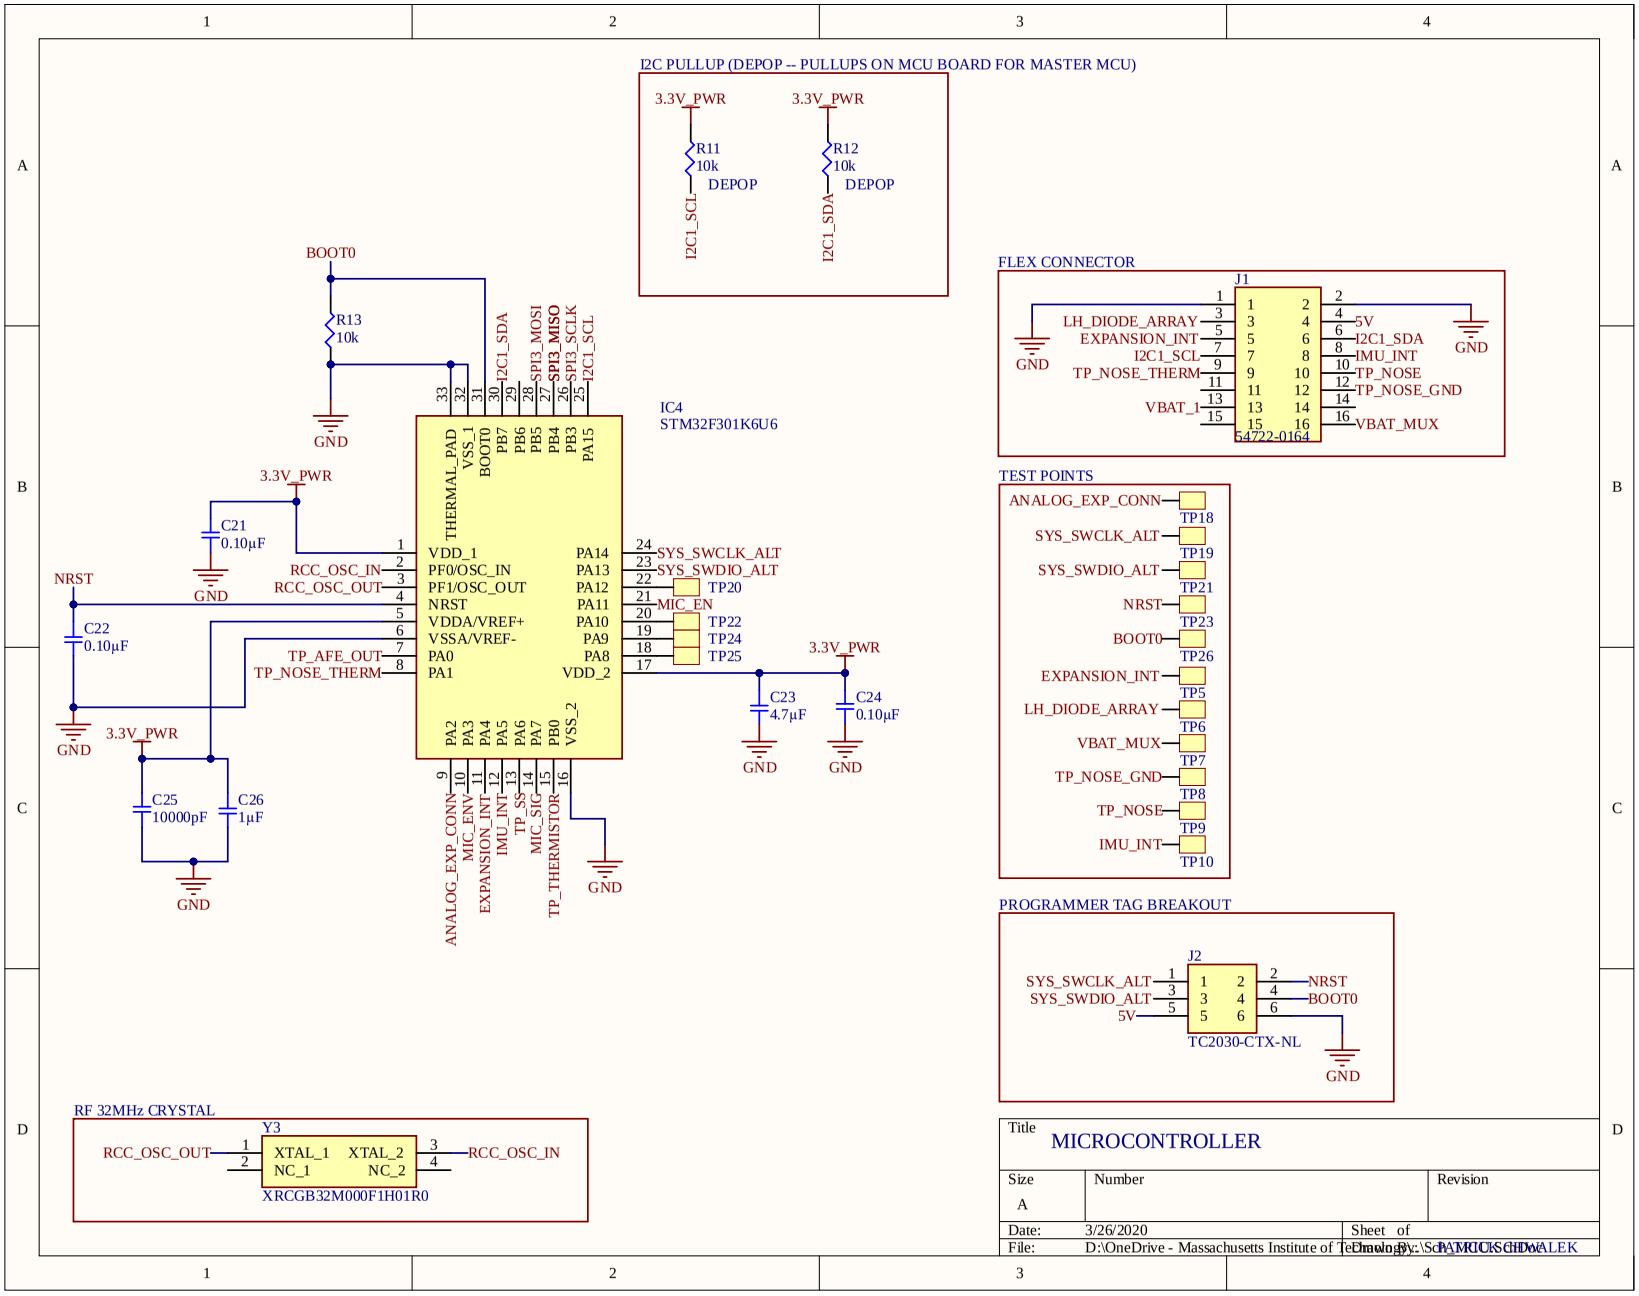
I2C PULLUP (DEPOP -- PULLUPS ON MCU BOARD FOR MASTER MCU)
3.3V_PWR
3.3V_PWR
R11
10k
DEPOP
R12
10k
DEPOP
I2C1_SCL
I2C1_SDA
FLEX CONNECTOR
J1
LH_DIODE_ARRAY
EXPANSION_INT
I2C1_SCL
TP_NOSE_THERM
VBAT_1
5V
I2C1_SDA
IMU_INT
TP_NOSE
TP_NOSE_GND
VBAT_MUX
GND
54722-0164
BOOT0
R13
10k
GND
IC4
STM32F301K6U6
THERMAL_PAD
VSS_1
BOOT0
PB7
PB6
PB5
PB4
PB3
PA15
I2C1_SDA
SPI3_MOSI
SPI3_MISO
SPI3_SCLK
I2C1_SCL
3.3V_PWR
C21
0.10µF
GND
NRST
C22
0.10µF
GND
RCC_OSC_IN
RCC_OSC_OUT
TP_AFE_OUT
TP_NOSE_THERM
VDD_1
PF0/OSC_IN
PF1/OSC_OUT
NRST
VDDA/VREF+
VSSA/VREF-
PA0
PA1
PA14
PA13
PA12
PA11
PA10
PA9
PA8
VDD_2
SYS_SWCLK_ALT
SYS_SWDIO_ALT
TP20
MIC_EN
TP22
TP24
TP25
3.3V_PWR
C23
4.7µF
GND
C24
0.10µF
GND
3.3V_PWR
C25
10000pF
C26
1µF
GND
PA2
PA3
PA4
PA5
PA6
PA7
PB0
VSS_2
ANALOG_EXP_CONN
MIC_ENV
EXPANSION_INT
IMU_INT
TP_SS
MIC_SIG
TP_THERMISTOR
GND
TEST POINTS
ANALOG_EXP_CONN
TP18
SYS_SWCLK_ALT
TP19
SYS_SWDIO_ALT
TP21
NRST
TP23
BOOT0
TP26
EXPANSION_INT
TP5
LH_DIODE_ARRAY
TP6
VBAT_MUX
TP7
TP_NOSE_GND
TP8
TP_NOSE
TP9
IMU_INT
TP10
PROGRAMMER TAG BREAKOUT
J2
SYS_SWCLK_ALT
SYS_SWDIO_ALT
5V
NRST
BOOT0
TC2030-CTX-NL
GND
RF 32MHz CRYSTAL
Y3
RCC_OSC_OUT
XTAL_1
XTAL_2
NC_1
NC_2
RCC_OSC_IN
XRCGB32M000F1H01R0
Title
MICROCONTROLLER
Size
A
Number
Revision
Date:
3/26/2020
Sheet of
File:
D:\OneDrive - Massachusetts Institute of Technology\...\MICRO.SCHDOC
Drawn By:
PATRICK CHWALEK

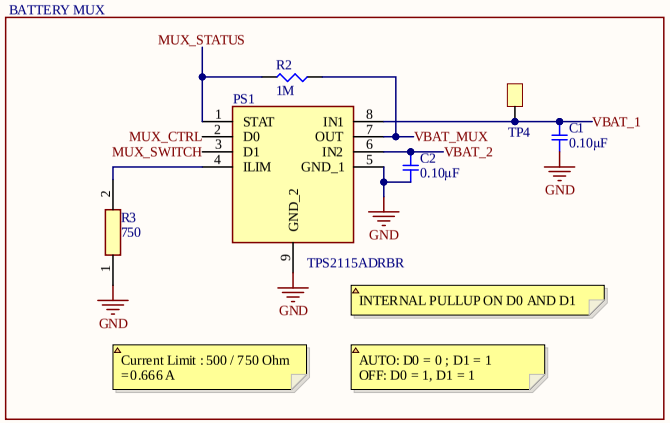

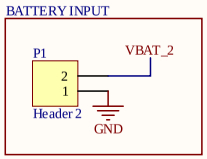

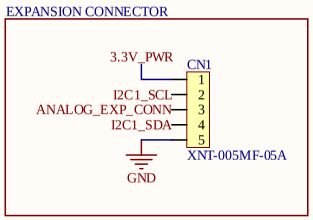

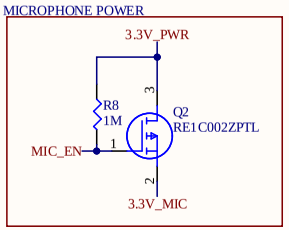

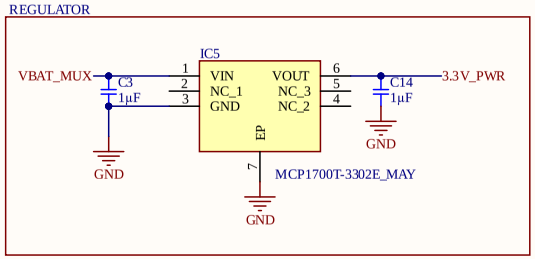

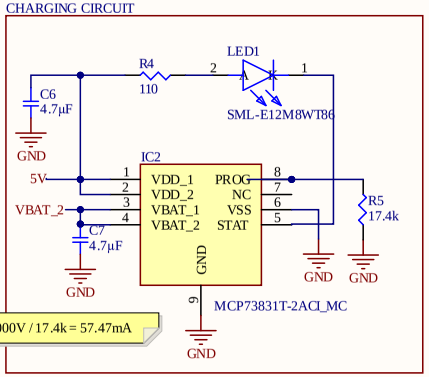

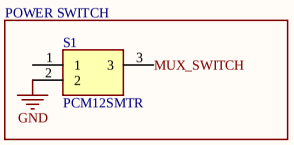

Title: POWER		
Size: Letter	Number:	Revision:
Date: 3/26/2020	Sheet of	
File: D:\OneDrive - Massachusetts Institute of Technology\...\Sch.SchDoc	Drawn By:	

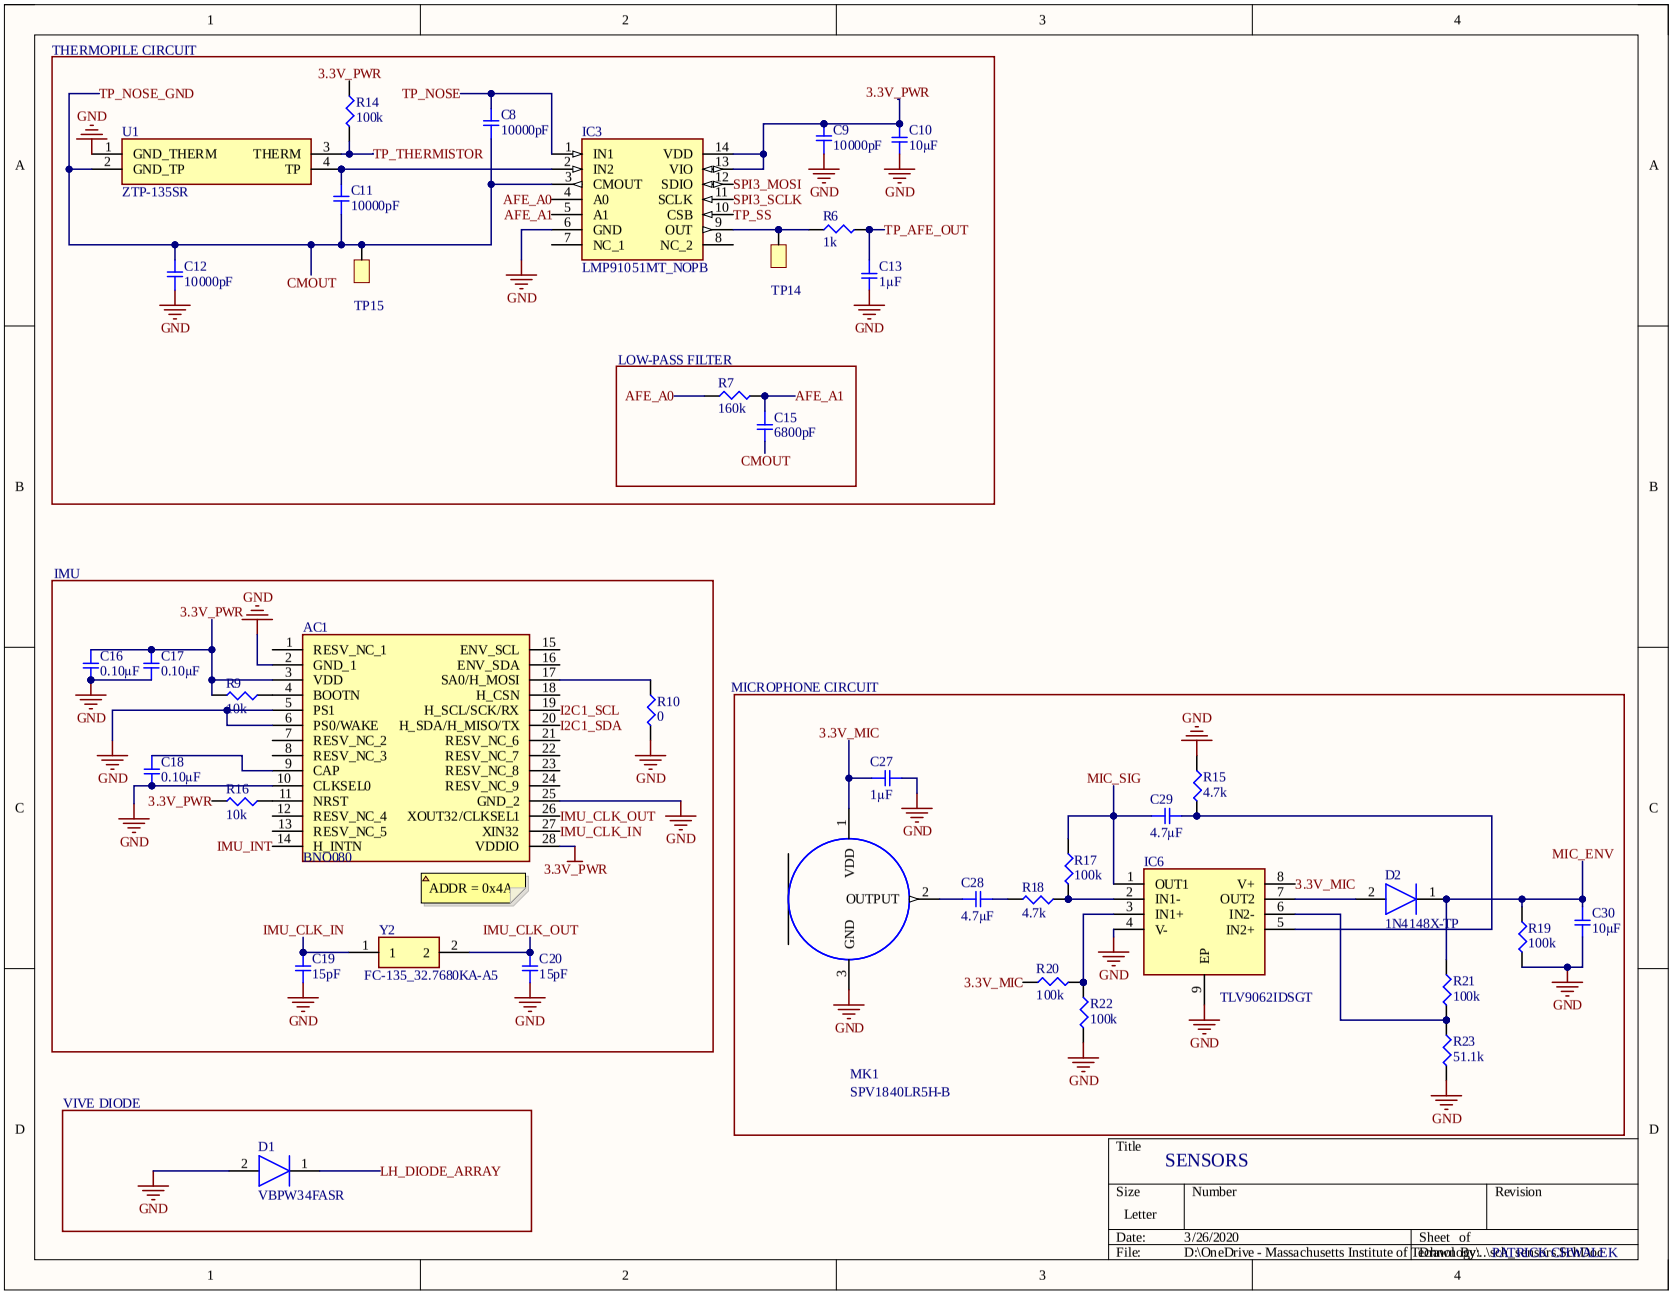
THERMOPILE CIRCUIT
3.3V_PWR
TP_NOSE_GND
TP_NOSE
R14
100k
C8
10000pF
3.3V_PWR
GND
U1
GND_THERM
GND_TP
THERM
TP
TP_THERMISTOR
ZTP-135SR
C11
10000pF
IC3
IN1
IN2
CMOUT
A0
A1
GND
NC_1
VDD
VIO
SDIO
SCLK
CSB
OUT
NC_2
C9
10000pF
C10
10µF
SPI3_MOSI
SPI3_SCLK
TP_SS
AFE_A0
AFE_A1
R6
1k
TP_AFE_OUT
C12
10000pF
CMOUT
TP15
LMP91051MT_NOPB
TP14
C13
1µF
LOW-PASS FILTER
AFE_A0
R7
160k
AFE_A1
C15
6800pF
CMOUT
IMU
GND
3.3V_PWR
AC1
RESV_NC_1
GND_1
VDD
BOOTN
PS1
PS0/WAKE
RESV_NC_2
RESV_NC_3
CAP
CLKSEL0
NRST
RESV_NC_4
RESV_NC_5
H_INTN
ENV_SCL
ENV_SDA
SA0/H_MOSI
H_CSN
H_SCL/SCK/RX
H_SDA/H_MISO/TX
RESV_NC_6
RESV_NC_7
RESV_NC_8
RESV_NC_9
GND_2
XOUT32/CLKSEL1
XIN32
VDDIO
I2C1_SCL
I2C1_SDA
IMU_CLK_OUT
IMU_CLK_IN
C16
0.10µF
C17
0.10µF
R9
10k
C18
0.10µF
R16
10k
IMU_INT
R10
0
BNO080
ADDR = 0x4A
IMU_CLK_IN
Y2
IMU_CLK_OUT
C19
15pF
FC-135_32.7680KA-A5
C20
15pF
MICROPHONE CIRCUIT
3.3V_MIC
C27
1µF
MIC_SIG
C29
4.7µF
R15
4.7k
MIC_ENV
R17
100k
IC6
OUT1
IN1-
IN1+
V-
V+
OUT2
IN2-
IN2+
EP
TLV9062IDSGT
D2
1N4148X-TP
C28
4.7µF
R18
4.7k
R19
100k
C30
10µF
R20
100k
R22
100k
R21
100k
R23
51.1k
MK1
SPV1840LR5H-B
VIVE DIODE
D1
VBPW34FASR
LH_DIODE_ARRAY
Title
SENSORS
Size
Letter
Number
Revision
Date:
3/26/2020
Sheet of
File:
D:\OneDrive - Massachusetts Institute of Technology\...\SENSORS.SchDoc
Drawn By:

Footprint	Name	LibRef	Designator	Description	Quantity
BNO080	BNO080	BNO080	AC1	Accelerometer	1
SOP65P640X120-14N	LMP91051MT_NOPB	LMP91051MT_NOPB	IC3	Integrated Circuit	1
ZTP-135SR	ZTP-135SR	ZTP-135SR	U1	Undefined or Miscellaneous	1
SMD-0402C, SMD-0402-	1µF	GCM155C71A105KE38D	C3, C13, C14, C26, C27	CAP CER 1UF 10V X7S 0402	5
QFN50P500X500X60-33N-	STM32F301K6U6	STM32F301K6U6	IC4	Integrated Circuit	1
SON65P300X300X100-9N-	TPS2115ADRBR	TPS2115ADRBR	PS1	Power Supply	1
54722-0164	54722-0164	54722-0164	J1	Connector	1
SON50P200X200X80-9N	TLV9062IDSGT	TLV9062IDSGT	IC6	Integrated Circuit	1
*MIC_SPV1840LR5H-B	SPV1840LR5H-B	SPV1840LR5H-B	MK1	Mic Mems Analog	1
PCM12SMTR_1	PCM12SMTR	PCM12SMTR	S1	Switch	1
VBPW34FASR	VBPW34FASR	VBPW34FASR	D1	Diode	1
SON50P300X200X100-9N-	MCP73831T-2ACI_MC	MCP73831T-2ACI_MC	IC2	Integrated Circuit	1
SON65P200X200X90-7N	MCP1700T-3302E_MAY	MCP1700T-3302E_MAY	IC5	Integrated Circuit	1
FC135327680KAAG0	FC-135_32.7680KA-A5	FC-135_32.7680KA-A3	Y2	Crystal or Oscillator	1
SMD-0402C	4.7µF	GRM153R60G475ME15D	C6, C7, C23, C28, C29	CAP CER 4.7UF 4V X5R 0402	5
SMD-0402C	10µF	GRJ155R60J106ME11D	C10, C30	CAP CER 10UF 6.3V X5R 0402	2
LEDC1608X46N	SML-E12M8WT86	SML-E12M8WT86	LED1	LED	1
SOTFL50P160X85-3N	RE1C002ZPTL	RE1C002ZPTL	Q2	MOSFET (P-Channel)	1
SMD-0402C	6800pF	GRM155R71E682KA01J	C15	CAP CER 6800PF 25V X7R 0402	1
SMD-0402C,	0.10µF	GCM155R71C104KA55D	C1, C2, C16, C17, C18, C21, C22, C24	CAP CER 0.1UF 16V X7R 0402	8
SODFL1608X70N	1N4148X-TP	1N4148X-TP	D2	Diode	1
SMD-0402-RES	100k	ERJ-2GEJ103X	R14, R17, R19, R20, R21, R22	RES SMD 100K OHM 5% 1/10W 0402	6
SMD-0402-RES	10k	ERJ-2GEJ103X	R9, R11, R12, R13, R16	RES SMD 10K OHM 5% 1/10W 0402	5
SMD-0402C	10000pF	GCM155R71H103KA55D	C8, C9, C11, C12, C25	CAP CER 10000PF 50V X7R 0402	5
RESC1005X40N	750	ERJ-2RKF7500X	R3	RES SMD 750 OHM 1% 1/10W 0402	1
SMD-0402-RES	51.1k	ERJ-2GEJ103X	R23	RES SMD 51.1K OHM 5% 1/10W 0402	1
SMD-0402-RES	110	ERJ-2RKF8062X	R4	RES SMD 110 OHM 5% 1/10W 0402	1
CAPC1005X55N	15pF	GCM1555C1H150JA16D	C19, C20	CAP CER 15PF 50V NP0 0402	2
SMD-0402-RES	4.7k	ERJ-2GEJ103X	R15, R18	RES SMD 4.7K OHM 5% 1/10W 0402	2
SMD-0402-RES	1M	ERJ-2GEJ105X	R2, R8	RES SMD 1M OHM 5% 1/10W 0402	2
SMD-0402-RES	0	ERJ-2GEJ103X	R10	RES SMD 0K OHM 5% 1/10W 0402	1
SMD-0402-RES	17.4k	ERJ-2GEJ103X	R5	RES SMD 17.4K OHM 1% 1/10W 0402	1
SMD-0402-RES	1k	ERJ-2GEJ102X	R6	RES SMD 1K OHM 5% 1/10W 0402	1
SMD-0402-RES	160k	ERJ-2GEJ164X	R7	RES SMD 160K OHM 5% 1/10W 0402	1
XRCGB32M000F1H01R0	XRCGB32M000F1H01R0	XRCGB32M000F1H01R0	Y3	Crystal or Oscillator	1
TC2030MCPNL10	TC2030-CTX-NL	TC2030-CTX-NL	J2	Connector	1
					68

Appendix B

Self-Assessment Survey

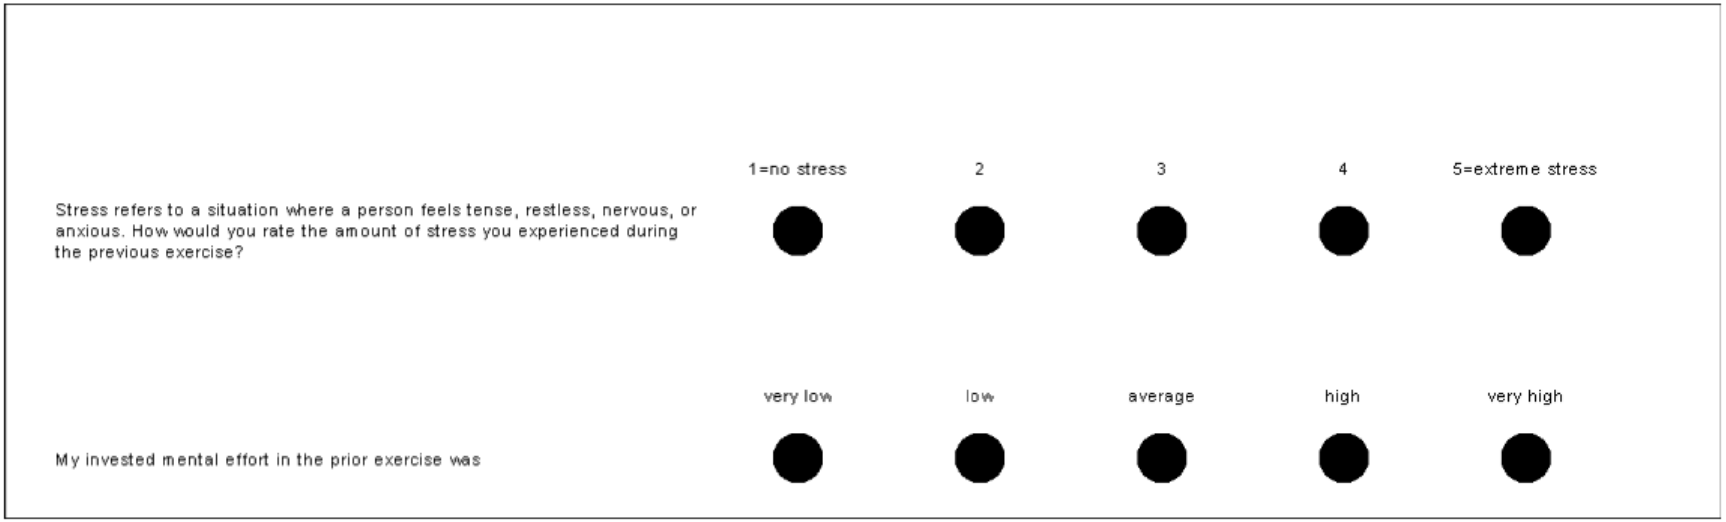

Figure B-1: Self-Assessment Questions (1/3)

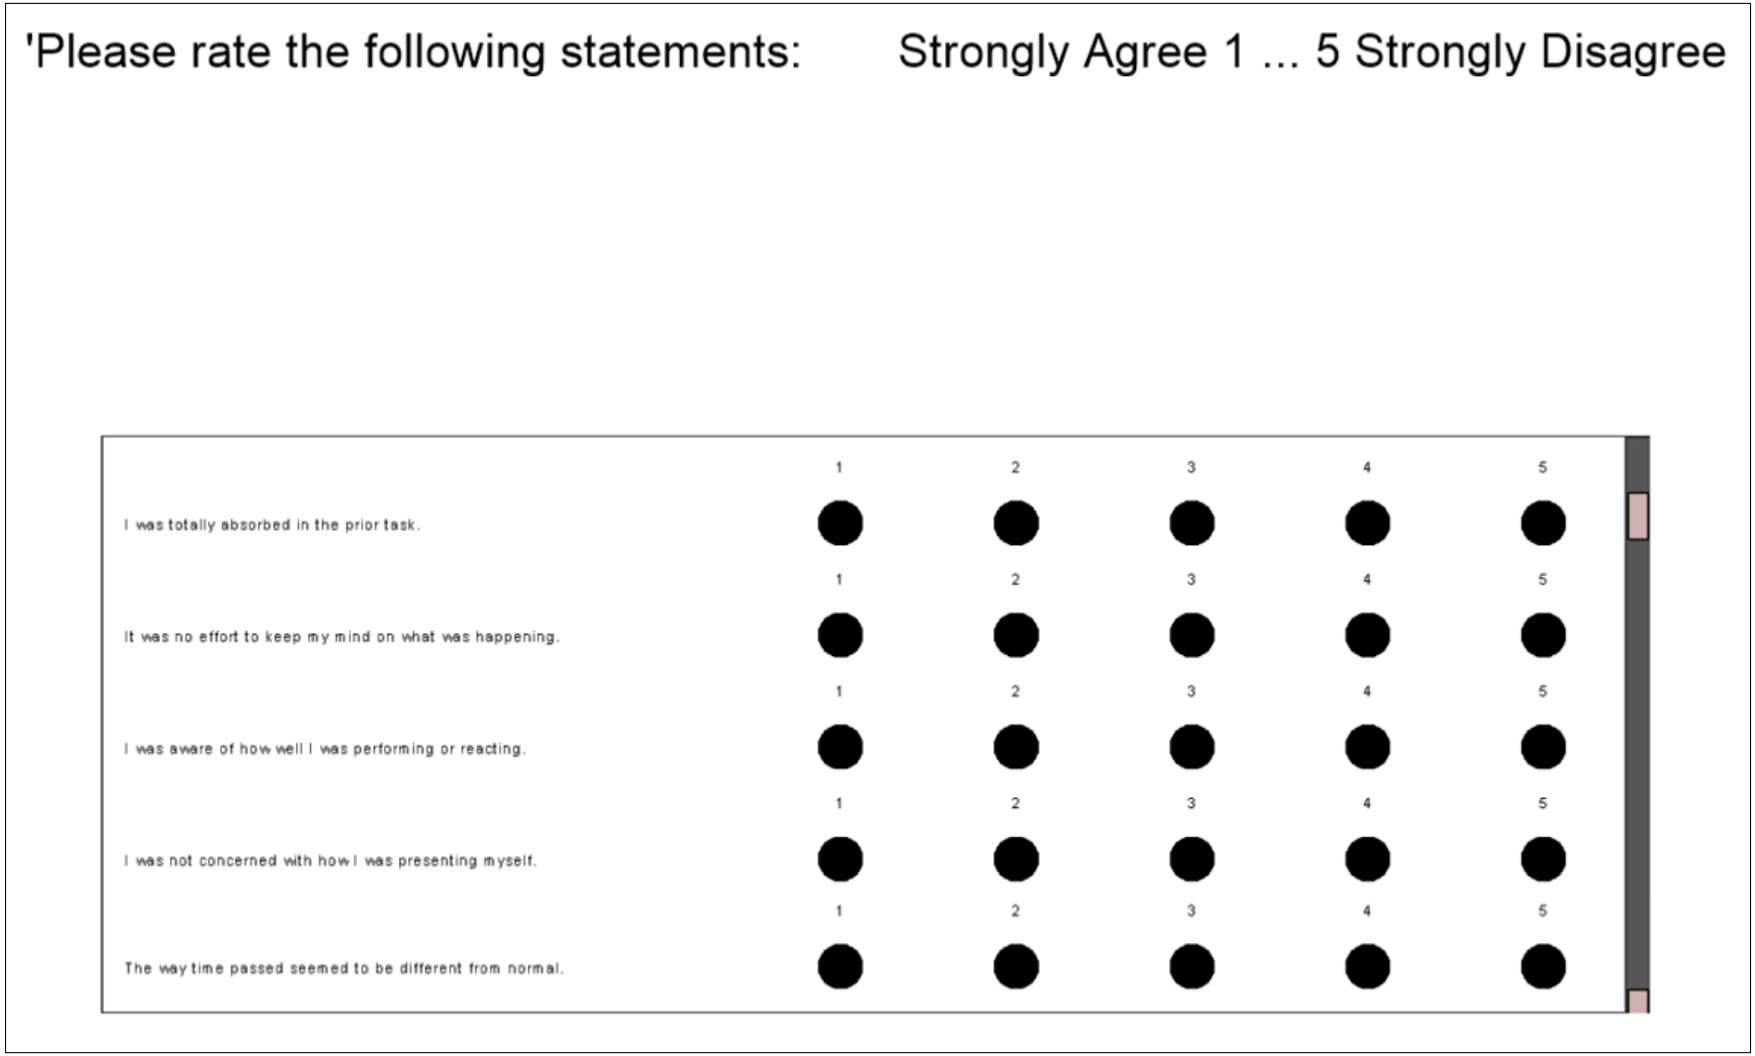

Figure B-2: Self-Assessment Questions (2/3)

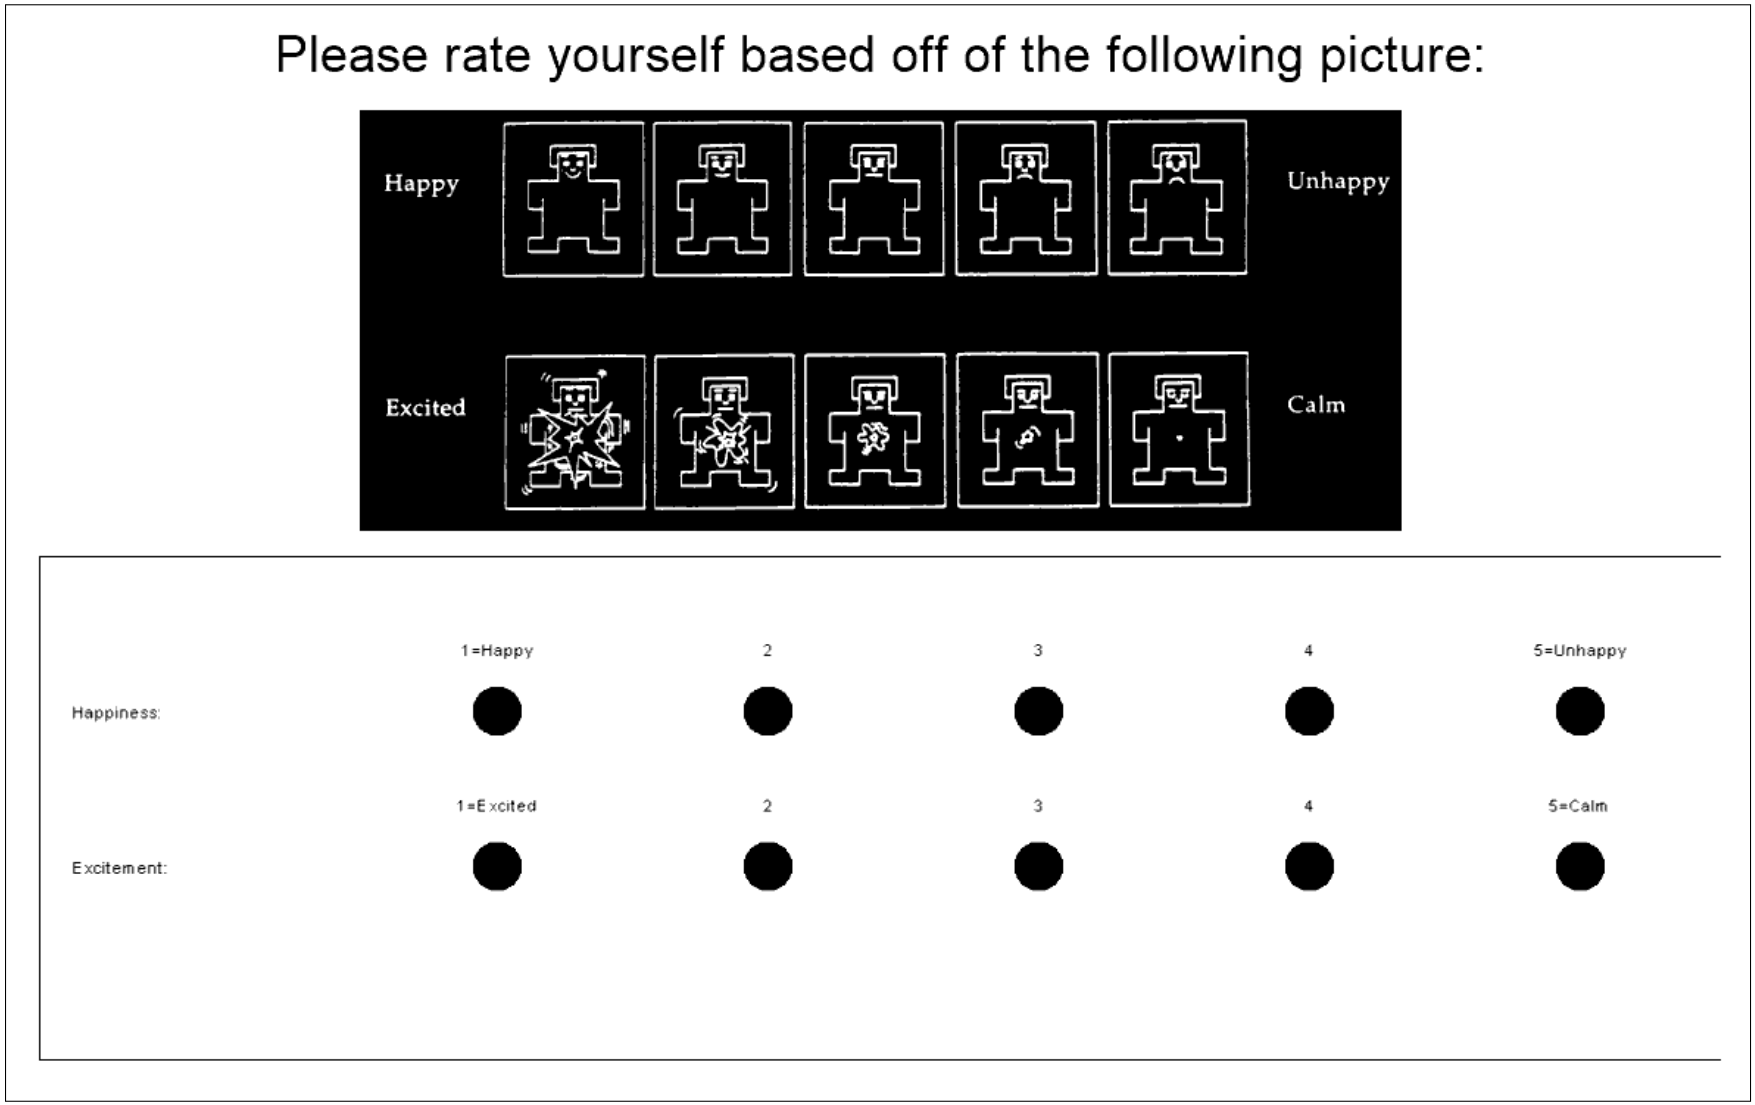

Figure B-3: Self-Assessment Questions (3/3)

Bibliography

[1] "Flex circuit materials, layers and plated through vias," 2020 (accessed May 3, 2020). [Online]. Available: https://www.tech-etch.com/flex/flex_materials.html

[2] M. Finstad, *Basics of Flex Circuit Design*, Minco, 2008.

[3] *Comparing Low-Power Wireless Technologies*, 2020 (accessed May 3, 2020). [Online]. Available: http://www.hotenda.com/media/articles/Comparing-Lowpower-Wireless-Technologies.html

[4] *NDIR Thermopile-Based Gas Sensing Circuit*, Analog Devices, 2020.

[5] A. Shtuchkin, "DIY position tracking using HTC vive's lighthouse," https://github.com/ashtuchkin/vive-diy-position-sensor, 2020.

[6] S. Ioannou, V. Gallese, and A. Merla, "Thermal infrared imaging in psychophysiology: Potentialities and limits," in *Psychophysiology*, 2014.

[7] L. Dubourg, A. R. Silva, C. Fitamen, C. J. A. Moulin, and C. Souchay, "Sensecam: A new tool for memory rehabilitation?" *Revue neurologique*, vol. 172 12, pp. 735–747, 2016.

[8] M. Stikic, C. Berka, D. J. Levendowski, R. F. Rubio, V. Tan, S. Korszen, D. Barba, and D. Wurzer, "Modeling temporal sequences of cognitive state changes based on a combination of eeg-engagement, eeg-workload, and heart rate metrics," *Frontiers in Neuroscience*, vol. 8, 2014.

[9] F. Zennifa, J. Ide, Y. Noguchi, and K. Iramina, "Monitoring of cognitive state on mental retardation child using eeg, ecg and nirs in four years study," *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 6610–6613, 2015.

[10] P. Gerjets, C. Walter, W. Rosenstiel, M. Bogdan, and T. O. Zander, "Cognitive state monitoring and the design of adaptive instruction in digital environments: lessons learned from cognitive workload assessment using a passive brain-computer interface approach," *Frontiers in Neuroscience*, vol. 8, 2014.

[11] S. Lal and A. Craig, "Reproducibility of the spectral components of the electroencephalogram during driver fatigue." *International journal of psychophysiology :*

official journal of the International Organization of Psychophysiology, vol. 55 2, pp. 137–43, 2005.

[12] W. Klimesch, "Eeg alpha and theta oscillations reflect cognitive and memory performance: a review and analysis," *Brain Research Reviews*, vol. 29, pp. 169–195, 1999.

[13] H. Ledger, "The effect cognitive load has on eye blinking," 2013.

[14] I. Daly, N. Nicolaou, S. Nasuto, and K. Warwick, "Automated artifact removal from the electroencephalogram: A comparative study," *Clinical EEG and neuroscience : official journal of the EEG and Clinical Neuroscience Society (ENCS)*, vol. 44, 05 2013.

[15] G. L. Wallstrom, R. E. Kass, A. Miller, J. F. Cohn, and N. A. Fox, "Automatic correction of ocular artifacts in the eeg: a comparison of regression-based and component-based methods." *International journal of psychophysiology : official journal of the International Organization of Psychophysiology*, vol. 53 2, pp. 105–19, 2004.

[16] A. Cömert and J. Hyttinen, "Investigating the possible effect of electrode support structure on motion artifact in wearable bioelectric signal monitoring," *BioMedical Engineering OnLine*, vol. 14, 2015.

[17] Y. Abdelrahman, E. Velloso, T. Dingler, A. Schmidt, and F. Vetere, "Cognitive heat: Exploring the usage of thermal imaging to unobtrusively estimate cognitive load," *IMWUT*, vol. 1, pp. 33:1–33:20, 2017.

[18] T. Mizuno and Y. Kume, "Development of a glasses-like wearable device to measure nasal skin temperature," in *HCI*, 2015.

[19] F. Pompei, "Temporal artery temperature detector," patentus US9 194 749B2.

[20] K. B. Hebbar, J. D. Fortenberry, K. Rogers, R. Merritt, and K. Easley, "Comparison of temporal artery thermometer to standard temperature measurements in pediatric intensive care unit patients." *Pediatric critical care medicine : a journal of the Society of Critical Care Medicine and the World Federation of Pediatric Intensive and Critical Care Societies*, vol. 6 5, pp. 557–61, 2005.

[21] J. A. Stern, L. Walrath, and R. N. Goldstein, "The endogenous eyeblink." *Psychophysiology*, vol. 21 1, pp. 22–33, 1984.

[22] N. Nourbakhsh, Y. Wang, and F. Chen, "Gsr and blink features for cognitive load classification," in *INTERACT*, 2013.

[23] Y. Shi, N. Ruiz, R. Taib, E. H. C. Choi, and F. Chen, "Galvanic skin response (gsr) as an index of cognitive load," in *CHI Extended Abstracts*, 2007.

[24] A. Dementyev and C. Holz, "Dualblink: A wearable device to continuously detect, track, and actuate blinking for alleviating dry eyes and computer vision syndrome," *IMWUT*, vol. 1, pp. 1:1–1:19, 2017.

[25] R. Makepeace and J. Epps, "Automatic task analysis based on head movement," *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 5167–5170, 2015.

[26] R. Stiefelhagen, "Tracking focus of attention in meetings," *Proceedings. Fourth IEEE International Conference on Multimodal Interfaces*, pp. 273–280, 2002.

[27] "Meditation made easy," 2020 (accessed May 3, 2020). [Online]. Available: https://choosemuse.com/

[28] G. D. Flumeri, P. Aricò, G. Borghini, N. Sciaraffa, A. D. Florio, and F. Babiloni, "The dry revolution: Evaluation of three different eeg dry electrode types in terms of signal spectral features, mental states classification and usability," *Sensors (Basel, Switzerland)*, vol. 19, 2019.

[29] P. Sawangjai, S. Hompoonsup, P. Leelaarporn, S. Kongwudhikunakorn, and T. Wilaiprasitporn, "Consumer grade eeg measuring sensors as research tools: A review," *IEEE Sensors Journal*, vol. 20, pp. 3996–4024, 2020.

[30] *Open Source Tools for Neuroscience*, 2020 (accessed May 3, 2020). [Online]. Available: https://openbci.com/

[31] S. Ishimaru, K. Kunze, K. Kise, J. Weppner, A. Dengel, P. Lukowicz, and A. Bulling, "In the blink of an eye: combining head motion and eye blink frequency for activity recognition with google glass," in *AH '14*, 2014.

[32] U. Rehman and S. Cao, "Augmented-reality-based indoor navigation: A comparative analysis of handheld devices versus google glass," *IEEE Transactions on Human-Machine Systems*, vol. 47, pp. 140–151, 2017.

[33] N. J. Wei, B. Dougherty, A. Myers, and S. M. Badawy, "Using google glass in surgical settings: Systematic review," *JMIR mHealth and uHealth*, vol. 6, 2018.

[34] "Your everyday smart glasses." [Online]. Available: https://vueglasses.com/

[35] *Focals by North.*, 2020 (accessed May 3, 2020). [Online]. Available: https://www.bynorth.com/

[36] J. A. Castro-García, A. J. Molina-Cantero, M. Merino-Monge, and I. M. Gómez-González, "An open-source hardware acquisition platform for physiological measurements," *IEEE Sensors Journal*, vol. 19, pp. 11 526–11 534, 2019.

[37] "Nestor: Ai driven courses." [Online]. Available: https://nestor-ai.com/

[38] H. Chen, H. W. Park, X. Zhang, and C. Breazeal, "Impact of interaction context on the student affect-learning relationship in child-robot interaction," *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, 2020.

[39] K. Ahuja, D. Kim, F. Xhakaj, V. Varga, A. Xie, S. Zhang, J. E. Townsend, C. Harrison, A. Ogan, and Y. Agarwal, "Edusense: Practical classroom sensing at scale," *IMWUT*, vol. 3, pp. 71:1–71:26, 2019.

[40] "Facial recognition used to analyze students' classroom behaviors," 2020 (accessed May 3, 2020). [Online]. Available: http://en.people.cn/n3/2018/0519/c90000-9461918.html

[41] S. Hong, "China's efforts to lead the way in ai start in its classrooms," Oct 2019. [Online]. Available: https://www.wsj.com/articles/chinas-efforts-to-lead-the-way-in-ai-start-in-its-classrooms-11571958181

[42] J. Li, "A "brain-reading" headband for students is too much even for chinese parents," Nov 2019. [Online]. Available: https://qz.com/1742279/a-mind-reading-headband-is-facing-backlash-in-china/

[43] G. Gratton, "Dealing with artifacts: The eog contamination of the event-related brain potential," *Behavior Research Methods, Instruments, Computers*, vol. 30, pp. 44–53, 1998.

[44] S. Romero, M. A. Mañanas, and M. J. Barbanoj, "A comparative study of automatic techniques for ocular artifact reduction in spontaneous eeg signals based on clinical target variables: A simulation case," *Computers in biology and medicine*, vol. 38 3, pp. 348–60, 2008.

[45] R. J. Croft and R. J. Barry, "Removal of ocular artifact from the eeg: a review," *Neurophysiologie Clinique/Clinical Neurophysiology*, vol. 30, pp. 5–19, 2000.

[46] *FocusEDU*, 2020 (accessed May 3, 2020). [Online]. Available: https://www.brainco.tech/

[47] N. Kosmyna, C. Morris, U. Sarawgi, and P. Maes, "Attentivu: A biofeedback system for real-time monitoring and improvement of engagement," *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019.

[48] S. Ishimaru, K. Kunze, K. Kise, J. Weppner, A. Dengel, P. Lukowicz, and A. Bulling, "In the blink of an eye - combining head motion and eye blink frequency for activity recognition with google glass," 03 2014, pp. 150–153.

[49] *Base Station*, 2020 (accessed May 3, 2020). [Online]. Available: https://www.vive.com/us/accessory/base-station/

[50] E. Costanza, S. Inverso, E. Pavlov, R. Allen, and P. Maes, "eye-q: eyeglass peripheral display for subtle intimate notifications," 01 2006.

[51] *Quality RTOS Embedded Software*, 2020 (accessed May 3, 2020). [Online]. Available: https://www.freertos.org/

[52] *Thread*. [Online]. Available: https://www.threadgroup.org/

[53] *OpenThread*, 2020 (accessed May 3, 2020). [Online]. Available: https://openthread.io/

[54] C. Ellmer, "Openthread vs. contiki IPv6: An experimental evaluation," 2017.

[55] G. T. Tapan Pattnayak, *Antenna Design and RF Layout Guidelines*, Cypress, 8 2018, rev. H.

[56] *Application Note NDIR-Measurement*, Heimann Sensor, 2013.

[57] H. Karaki and V. Polyzoev, "Demystifying thermopile ir temp sensors," 2020 (accessed May 3, 2020). [Online]. Available: https://www.fierceelectronics.com/components/demystifying-thermopile-ir-temp-sensors-0

[58] *TMP006/B Infrared Thermopile Sensor in Chip-Scale Package*, Texas Instruments, 5 2011, rev. C.

[59] *ZTP-135SR Thermometrics Thermopile IR Sensor*, Amphenol, 8 2018, rev. H.

[60] Z. Shelby, K. Hartke, and C. Bormann, "The constrained application protocol (coap)," *RFC*, vol. 7252, pp. 1–112, 2014.

[61] G. Tanganelli, C. Vallati, and E. Mingozzi, "Coapthon: Easy development of coap-based iot applications with python," *2015 IEEE 2nd World Forum on Internet of Things (WF-IoT)*, pp. 63–68, 2015.

[62] PyQT, "Pyqt reference guide," 2012. [Online]. Available: http://www.riverbankcomputing.com/static/Docs/PyQt4/html/index.html

[63] J. W. Peirce, "Psychopy—psychophysics software in python," *Journal of Neuroscience Methods*, vol. 162, pp. 8 – 13, 2007.

[64] J. Ridley, "Studies of interference in serial verbal reactions," 2001.

[65] J. P. Kincaid, R. P. Fishburne, R. L. Rogers, and B. S. Chissom, "Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel," 1975.

[66] M. A. Sans, G. C. Torradeflot, E. P. Cabré, and M. R. Montiel, "Blink rate and incomplete blinks in six different controlled hard-copy and electronic reading conditions," 2015.

[67] P. J. Lang, M. Bradley, and B. N. Cuthbert, "International affective picture system (iaps): Affective ratings of pictures and instruction manual. technical report a-8," 2008.

[68] M. Mauri, P. Cipresso, A. Balgera, M. Villamira, and G. Riva, "Why is facebook so successful? psychophysiological measures describe a core flow state while using facebook," *Cyberpsychology, behavior and social networking*, vol. 14 12, pp. 723–31, 2011.

[69] M. M. Bradley and P. J. Lang, "Measuring emotion: the self-assessment manikin and the semantic differential." *Journal of behavior therapy and experimental psychiatry*, vol. 25 1, pp. 49–59, 1994.

[70] S. Salminen, A. Kouvonen, A. Koskinen, M. Joensuu, and A. Väänänen, "Is a single item stress measure independently associated with subsequent severe injury: a prospective cohort study of 16,385 forest industry employees," *BMC Public Health*, vol. 14, pp. 543 – 543, 2014.

[71] F. Paas, "Training strategies for attaining transfer of problem-solving skill in statistics: A cognitive-load approach." 1992.

[72] *Shimmer*, 2020 (accessed May 3, 2020). [Online]. Available: http://www.shimmersensing.com/

[73] *Zephyr Performance Systems*, 2020 (accessed May 3, 2020). [Online]. Available: https://www.zephyranywhere.com/

[74] M. Grassmann, E. Vlemincx, A. von Leupoldt, J. M. Mittelstädt, and O. V. den Bergh, "Respiratory changes in response to cognitive load: A systematic review," *Neural Plasticity*, vol. 2016, 2016.

[75] D. K. keung Ng, C. hong Chan, E. Y. tung Chan, K. li Kwok, P. yu Chow, W. fai Lau, and J. C.-S. Ho, "A brief report on the normal range of forehead temperature as determined by noncontact, handheld, infrared thermometer," *American Journal of Infection Control*, vol. 33, pp. 227 – 229, 2005.

[76] T. Baltrusaitis, A. B. Zadeh, Y. C. Lim, and L.-P. Morency, "Openface 2.0: Facial behavior analysis toolkit," *2018 13th IEEE International Conference on Automatic Face Gesture Recognition (FG 2018)*, pp. 59–66, 2018.

[77] Y. Wang, S. S. Toor, R. K. Gautam, and D. B. Henson, "Blink frequency and duration during perimetry and their relationship to test-retest threshold variability." *Investigative ophthalmology visual science*, vol. 52 7, pp. 4546–50, 2011.

[78] H. W. Park, I. Grover, S. Spaulding, L. Gomez, and C. Breazeal, "A model-free affective reinforcement learning approach to personalization of an autonomous social robot companion for early literacy education," in *AAAI*, 2019.