PediaTrack:
A Vital Monitoring System for Pediatric Cancer Patients in Low-Resource Settings

Engineering Sciences 96

Final Report

December 18, 2018

Harvard University
School of Engineering and Applied Sciences

Professor David Mooney & Professor Samir Mitragotri
for the Global Health Initiative
at Dana-Farber/Boston Children's Hospital

Table of Contents

Introduction

Every two minutes, a child in the world is diagnosed with cancer ("Every Two Minutes", 2016). Although this devastating disease affects all countries across the globe, the pediatric cancer burden is not equally distributed: approximately 20% of pediatric cancer patients live in high-income countries (HICs), whereas 80% live in low- and middle-income countries (LMICs).[1] In HICs, survival rates are above 80%, yet, in LMICs, only 20% of children survive[2] – meaning that a staggering 90% of pediatric cancer deaths occur in resource-limited countries (Rodriguez-Galindo et al., 2015). This stark disparity motivates the work of the Global Health Initiative (GHI) at the Dana-Farber Cancer Institute and Boston Children's Hospital, an organization which seeks to combat this reality and improve the lives of all pediatric cancer patients worldwide.

This semester, our team of 20 multidisciplinary Harvard engineering students had the opportunity to work with the GHI in our course, ES96: Engineering Problem Solving and Design, to try to engineer a project that could take some small step in addressing the pediatric cancer survival gap in LMICs. This report details the semester-long process by which our team developed our final project: *PediaTrack*, a vital monitoring system for early detection of sepsis and infection for in-hospital pediatric cancer patients.

This report is divided into three core chapters. First, it documents the Investigate Phase, in which the team researched, identified, and synthesized key issues that contribute to the problem of pediatric

---

[1] In pediatric oncology, countries are distinguished by gross national income (GNI) per capita, and categorized as High Income Countries (HICs) or Low-Middle Income Countries (LMICs). HICs have a GNI per capita greater than or equal to $12,056 while LMICs have a GNI per capita less than $12,056 ("World Bank Country and Lending Groups", n.d.).

[2] The National Cancer Institute defines survivorship as "the health and life of a person with cancer post treatment until the end of life" (NCI Dictionary of Cancer Terms). In practice, this definition is often measured on a five-year scale, where survivorship is considered as disease-free for five-years post-diagnosis without cancer recurrence ("Survival", 2018). For the purposes of this report, "survival outcome" is defined as such.

cancer survival in LMICs. Second, it follows the Ideate Phase, in which the team brainstormed ideas and engineering solutions. Third, it discusses the Create Phase, in which the team created and tested our product, *PediaTrack*. Finally, the report concludes with a discussion of next steps, future directions, and the potential impact that we hope that *PediaTrack* could have in improving the lives of pediatric cancer patients in resource-poor countries.

Chapter I:
Investigate Phase

1.1 | Overview

This chapter summarizes the team's investigation into pediatric cancer in LMICs and the work of our client, the Global Health Initiative (GHI) at the Dana-Farber Cancer Institute and Boston Children's Hospital. Specifically, this chapter describes our client and its mission, values, and existing programs, which served as a framework for our research and analysis. Next, this chapter provides an overview of the various facets of pediatric oncology, including the causes, treatments, and patient outcomes that are relevant to understanding our client's mission, and subsequently analyzes the disparity in pediatric oncology between HICs and LMICs. The chapter then outlines the process by which the team organized and synthesized our investigative findings into underlying themes and key issues contributing to this complex problem. Finally, the chapter concludes with our formal problem statement, where we ultimately conclude that better methods of data collection, communication, and interpretation are crucial to addressing the problem of low pediatric cancer survival rates in LMICs.

1.2 | Our Client: The Global Health Initiative

Our client, the Global Health Initiative (GHI), is an organization dedicated to the mission of improving the care and quality of life of children with cancer and blood disorders worldwide (Irini Albanti, personal communication; "About the Global Health Initiative", 2018). Dually affiliated with the Dana-Farber Cancer Institute and Boston Children's Hospital, the GHI was founded in 2012 on the premise that health is a universal human right, and that, as such, it is the responsibility of institutions such as Dana-Farber and Boston Children's to take action to alleviate the inequalities and disparities in

healthcare outcomes that exist in different areas of the world for pediatric cancer patients (Albanti I., personal communication). In particular, the GHI works to close the gap in survival outcomes for children with cancer in low-resource countries; in the words of the GHI's Director, Irini Albanti, it is the 20% survival rate that drives the GHI "to fight for the opposite numbers" for LMICs (Albanti I., personal communication).

Mission and Values

Our client's mission to combat the pediatric cancer survival gap in LMICs is motivated by a variety of values. First and foremost, the GHI mission unites and aligns the missions of Dana-Farber and Boston Children's Hospital, where Dana-Farber's ultimate goal is the "the eradication of cancer, AIDS, and related diseases and the fear that they engender" ("About the Global Health Initiative", 2018), and Boston Children's strives to work "until every child is well" ("Mission & Values of BCH", n.d.). Additionally, the GHI's efforts stem from a belief in the fundamental importance of knowledge transfer, information sharing, and global interaction to improving global health (Albanti I., personal communication; "Global Health | Armenia", n.d.). Indeed, the GHI believes that through global interactions, Dana-Farber and Boston Children's can become better healthcare providers by becoming more culturally sensitive, compassionate, and understanding of different patterns of disease and healthcare systems (Albanti I., personal communication).

Implementing Their Mission: Current Projects

In practice, the GHI translates their mission and values into action through three pillars: research, education, and capacity building. The GHI conducts research in health services, epidemiology, clinical disparities, and public health. GHI education programming includes the Pediatric Global Oncology and Blood Disorders Education (GLOBE) Program, fellowships, curriculum-development, and on-site training in LMICs (Albanti I., personal communication; "About the Global Health Initiative", 2018).

Finally, GHI implements capacity building through international collaboration, programmatic support, and protocol development at healthcare sites in LMICs.

The GHI brings together healthcare workers from diverse spheres to achieve its mission. Through its international projects at LMIC sites, the GHI unites a multidisciplinary team of nurses, social workers, psychologists, and physicians to work with dealing with childhood cancer and blood disorders (Albanti I., personal communication). Currently, the GHI is involved in global projects or collaborations in Rwanda, Myanmar, Egypt, Armenia, El Salvador, Guatemala, Haiti, and Mexico ("Where We Work", n.d.). GHI programming in these areas is tailored to the specific landscape and needs of each particular location. In Rwanda, for example, the GHI has helped develop national cancer centers and programs, and has set up oncological physician training, whereas in Egypt, the GHI facilitates global fellowship and nursing exchange programs. Prior to the GHI's assistance, the Butaro Cancer Center of Excellence in Rwanda provided low- to no-cost cancer care for pediatric and adult patients; thus, GHI's supplementary assistance came in the form of creating a "1-2 day training program covering the basics of pediatric oncologic care for Rwandan physicians and nurses that has been delivered country-wide," in addition to providing mentor nurses from Boston Children's Hospital every 3 months to teach about chemotherapy treatment and palliative care, among other topics ("Where We Work: Rwanda", n.d.). The Children's Cancer Hospital Egypt (CCHE) also provides free healthcare, and GHI supplements this with a 30-month postdoctoral fellowship providing physicians from North Africa and the Middle East with "clinical assignments, mentorship, research, and live telemedicine lectures" ("Where We Work: Egypt", n.d.).
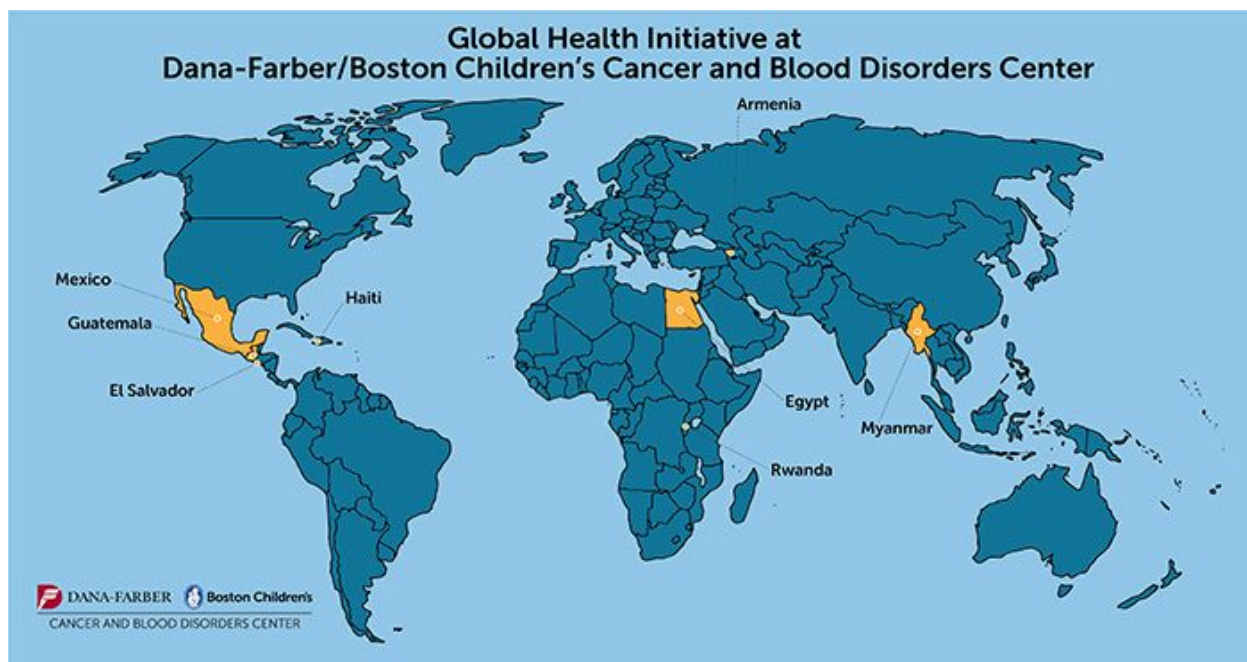
Figure 1.1. Map of the 8 countries in which the GHI is currently involved in projects and programming ("Where We Work", n.d.).

Client Interviews and Site Visits

In addition to performing thorough background research, our team also conducted client site visits and interviews as part of our investigation phase. We first met our client through a presentation and subsequent discussion with GHI Director Dr. Irini Albanti, to gain a big-picture understanding of the pediatric cancer burden in low-resource settings. We also interviewed Persistent Productions, on their documentary, *How I Live*, which follows five pediatric cancer patients from four LMICs, which provided more personal patient anecdotes to help us frame the problem. With this overview in mind, we next interviewed the GHI's GLOBE scholars, who represented several countries in which the GHI conducts its international programming. The scholars this year came from Brazil, Congo, Iraq, and Uganda and each provided an additional perspective into the healthcare process; some were nurses, others physicians, and one was even an oncologist, and with their narratives we were able to begin to understand what pediatric oncology care looks like worldwide. The team was also able to speak with pediatric nursing and

psychosocial healthcare workers at Dana-Farber/Boston Children's who had experience working in LMICs.

This initial round of data-gathering with Dr. Albanti, GLOBE scholars, Dana-Farber/Boston Children's healthcare providers, and Persistent Productions provided the background necessary to help us begin asking the right questions. After gathering qualitative data from these interviews, we went to primary sources and medical literature to gain more quantitative information about pediatric oncology to fill in gaps in our knowledge. The following sections serve as a synthesis of the client and literature research we have collected.

1.3 | Overview of Pediatric Oncology

In order to understand how we could aid the GHI's mission to reduce the survival gap for pediatric cancer patients in LMICs, it is first crucial to have a preliminary understanding of the diverse facets of pediatric oncology itself, before analyzing how these causes, treatments, and outcomes play out in low resource settings. This section thus contains a general overview of pediatric oncology, which sets up the background information necessary for analyzing the disparities in LMICs in following sections.

The Pediatric Cancer Burden

Cancer is the leading disease-related cause of death in children worldwide: each year, 160,000 new cases of pediatric cancer are identified and approximately 90,000 children across the globe die from the disease (Albanti I., personal communication; Ferlay, 2002). Pediatric cancer encompasses a diverse set of disease types, including, most commonly, lymphoblastic leukemia, Wilms tumor, and retinoblastoma, in addition to other types of cancer, each of which presents a different challenge to global health (Bertaut et al., 2015).

Pathology of Cancer

Cancer itself is an abnormal growth of cells that group together to form tumors. This rapid growth often occurs when cells acquire certain genetic mutations ("Risk Factors and Causes of Childhood Cancer", 2016). Irregular expression of certain genes can induce cancerous growth as their regulation of cell division and proliferation becomes uncontrolled (de Leon, M.P., 1994). Cancerous proliferation may also occur if tumor suppressor genes, which express proteins to restrict cancerous growth, are deactivated by a mutation (de Leon, M.P., 1994).

Cancer progresses in four clinically-defined stages. In its early phase, cancer cells grow locally to the primary site – known as stage 1 and 2. In stage 3, the disease state worsens when cancer cells transition to the invasive mesenchymal form which allows detachment and survival in the circulation of the blood (Yilmaz and Christofori, 2009). In stage 4, also known as advanced stage cancer, cancer cells then bind to and proliferate uncontrollably at new metastatic sites, taking up more space and nutrients from functional cells until they obstruct organ function and ultimately cause death (Yilmaz and Christofori, 2009).

Causes of Childhood Cancer

The causes of childhood cancer are still being researched and are not fully understood. In fact, about 75% of cases of cancer have unknown causes while around 15% and 10% are attributed to intrinsic and environmental factors, respectively (Lichtenstein P., 2000).

Intrinsic risk factors of childhood cancer include inherent or pre- and post-natal developmental issues (Spector et al., 2015). For instance, lower- or higher-than-average birth weights are associated with greater risk for acute myeloid leukemia, while risk of hepatoblastoma increases inversely with birth weight (Caughey et al., 2009). Genetic mutations, both inherited and developed in the womb, are also linked to cancers, with higher correlation rates in some; 40% of retinoblastoma cases involve the

inheritance of a mutated *RB1* tumor suppressor gene from germline cells, which predisposes children to the cancer (Kamihara et al., 2017). Over 200 gene variations are currently believed to control tumorigenesis, and more genome-wide association studies and accessible genetic sequencing can provide more information on the frequency of certain mutations and their correlation to different cancers (Johnson, 2018).

External factors can also play a role in the development of childhood cancer, but there have been difficulties in pinpointing the exact effects of exposure during early development ("Childhood Cancers", 2018). Second-hand smoking, ultraviolet radiation from the sun, drinking, and unhealthy eating — all of which are often associated with adult cancers — do not seem to be correlated with childhood cancers, possibly because the effects of these environmental stimuli may only manifest in adulthood ("Why Child Cancer?", 2018). However, cancer risk can increase with exposure to certain infections. For example, kaposi sarcoma incidence rates in sub-Saharan Africa increased tens of times in tandem with the region's HIV/AIDS epidemic (Stefan et al., 2010). In Uganda, kaposi sarcoma now makes up roughly 30% of cancer cases while the incidence of kaposi sarcoma in Australia, Argentina, China, and Iran is negligible relative to other types of pediatric cancer (Bertaut el al., 2015). However, literature concedes that the modifiable risk factors of childhood cancer are not understood well enough to make prevention an effective disease response (Gupta et al, 2015).

Diagnosis

When a child develops cancer, there are three phases to comprehensive care: diagnosis, treatment, and post-treatment. The diagnosis of childhood cancer begins with the recognition and further investigation of clinical symptoms, first by caretakers and then by trained doctors. Many clinical symptoms of pediatric cancer are nonspecific, including increased tendency to bruise, susceptibility to fever, sustained headaches, rapid weight loss, chronic fatigue, paleness, and localized pains or swelling

(Bertaut el al., 2015). These symptoms are often associated with less serious ailments than cancer, but check-ups with specialized doctors help determine when the combination or persistence of these nonspecific signs may indicate cancer (Bertaut el al., 2015).

Healthcare providers may administer medical diagnostic tests to more rigorously determine whether cancer is the underlying morbidity. Blood samples and tissue biopsies by surgery are taken to detect cancer cell biomarkers in pathology labs. For example, doctors analyze peripheral blood cell counts and bone marrow biopsies to determine if children presenting nonspecific symptoms have acute lymphoblastic leukemia; they also analyze stained lymph node biopsies to determine if children have Hodgkin lymphoma (Gupta et al., 2015). Magnetic resonance, ultrasound, and x-ray imaging machines may also be used to identify and understand the spread and size of tumors in the body ("How Cancer is Diagnosed", 2015).

Treatment

Once a patient has been diagnosed with cancer, the next phase involves transitioning them to treatment. Pediatric cancer places a massive emotional burden on the patient and the patient's family, so its optimal treatment is considered holistically, in which a treatment plan is developed to simultaneously control the type and stage of cancer and preserve the patient's quality of life (Marcus, 2012).

Leading pediatric cancer centers, like Dana-Farber and Boston Children's, provide multifaceted support for their patients to protect their physical and mental well-being while simultaneously treating the underlying malady (Psychosocial Oncology & Palliative Care Program). A wide array of experts are involved in a single patient's care to achieve this, including, but not limited to, oncologists, nurses, dieticians, psychologists, social workers, pathologists and palliative care experts (Nursing Team, personal communication). Treatment regimens often involve a multidisciplinary approach, including surgery, chemotherapy, and radiotherapy.

Surgery

If the cancer is contained within a certain part of the body, surgery for targeted removal of the tumor is especially effective (Bertaut et al., 2015). Surgeons usually excise some additional amount of normal, healthy tissue around the tumor to ensure complete cancer removal and prevent relapse (Stanford Health Care, n.d.). There are different types of surgical procedures for different cancers. For example, cryosurgery, which utilizes liquid nitrogen or argon to kill tissue, is used to treat retinoblastoma in children (Memorial Sloan Kettering Cancer Center). Tumor removal is often combined with other treatment options, such as placing intravenous lines or catheters into the body to administer chemotherapy drugs ("Childhood Cancers", n.d.).

Chemotherapy

Chemotherapy is the administration of drugs to stop cancerous growth by killing cancerous cells or stopping cell division (Bertaut et al., 2015). Children's bodies recover faster than adult bodies from chemotherapy and can therefore handle higher doses that kill more cancer cells in shorter periods of time ("Childhood Cancers", 2018). There are over 100 chemotherapy drugs that work differently due to their chemical structure and their relation to other drugs. Chemotherapy drugs can be used in a cocktail of three or more drugs at the same time, as some combinations of cancer drugs have been found to lead to a faster and more successful outcome than the use of one drug alone, but determining the most effective combination can be challenging (Mayer and Janoff, 2007). One group of chemotherapeutics are alkylating agents, which kill cancer cells by damaging their DNA (Bertaut et al., 2015). Cisplatin, a platinum-containing alkylating agent, stops cancer cell growth for neuroblastoma and osteosarcoma in pediatric patients ("Drugs: Chemotherapy Drugs and Other Pharmaceuticals", 2015). Another common class of chemotherapeutics is antimetabolites, which inhibit DNA synthesis in cancer cells to prevent regrowth. Mercaptopurine, a common antimetabolite, is used to treat acute lymphoblastic leukemia and

non-Hodgkin lymphoma in children ("Drugs: Chemotherapy Drugs and Other Pharmaceuticals", 2015). Since antimetabolites only aim to prevent cancer growth, they are often given in combination with other cancer drugs or radiotherapy to both limit growth and reduce the spread of tumors. Many chemotherapy drugs are administered intravenously or orally, which means they affect the whole body and are not targeted specifically for the tumor. This lack of specificity induces negative systemic side effects, including nausea, vomiting, fever, and decrease of white blood cells in a condition known as neutropenia ("Drugs: Chemotherapy Drugs and Other Pharmaceuticals", 2015).

Radiotherapy

Radiotherapy, or radiation therapy, uses high-energy x-rays to target a specific region of cells, causing strand breaks in cancer DNA. External radiation is the most common delivery method for children, where the x-ray source is outside the body and aims its strongest radioactive waves to the targeted site (Thwaites et al.). Treatment regimens usually require a day or two intermission in between to allow healthy cells to recover as much as possible and reduce the impact of side effects, such as nausea, headache, fatigue, and skin changes. Radiation therapy always involves a tradeoff, as it can permanently damage healthy cells, stunting learning, limiting hearing capabilities, and interfering with reproductive development when directed toward the head or neck ("Radiation Therapy in Children", n.d.).

Survival and Long-Term Quality of Life

Though cancer treatment regimens have proven to be effective against cancer, the disease and its harsh treatments can have lasting effects on the well-being of the survivor. A longitudinal study of over 10,000 pediatric cancer survivors treated at 26 hospitals in the United States demonstrated that survivors were at significantly greater risk of negative health outcomes later in life as compared to the healthy population (Oeffinger et al., 2006). On average, pediatric cancer survivors were 54 times as likely to get a major joint replacement, 15 times as likely to experience congestive heart failure, 14 times as likely to be diagnosed with a potentially cancerous growth, 10 times as likely to be diagnosed with

coronary artery disease, and 9 times as likely to go on dialysis (Oeffinger et al., 2006). A longitudinal

study in the Netherlands supported these conclusions, as 14.2% of pediatric cancer survivors experienced

orthopedic disruptions and 11.9% developed second tumors (Geenen et al., 2007). Additionally, specific

cancers were associated with specific late effects, suggesting that the particular needs of a pediatric cancer

patient may extend well into adulthood. For example, bone tumor survivors are at even higher risk of joint

replacement compared to survivors of other cancers, and this is attributed to long-term asymmetric stress

distributions (Oeffinger et al., 2006).

1.4 | Disparities in Pediatric Cancer Outcomes in LMICs

The complexity of pediatric oncology illustrates that, for a child with cancer, there are a number

of hurdles to overcome to reach complete remission – including early diagnosis, successful and effective

treatment, and consistent post-treatment care. However, the reality is that while the treatments and

outcomes as described above are typically standard in HICs, they are not ubiquitous and not widely

available in most LMICs. Consequently, to determine the factors affecting the LMIC survival gap, this

section will draw key comparisons to each element of pediatric oncology care in LMICs to these HIC

standards, in order to illuminate where disparities arise.

Research Limitations in LMIC Pediatric Oncology Data

Before beginning this analysis, it is important to note some of the limitations in our research that

arise from the underdeveloped nature of pediatric oncology care in LMICs. First, pediatric oncology is a

very heterogeneous field, involving the treatment of many different diseases across vastly different patient

profiles and ranging from infancy to adolescence. Additionally, attitudes towards pediatric oncology care

vary widely across and within LMICs, adding to the heterogeneity of the overarching categorization.

Importantly, the information landscape of the healthcare systems across many LMICs is also

incomplete. Most LMICs lack cancer registries, databases which store information on cancer cases

throughout the country, a practice which is common in HICs. In North America and Oceania, 90% of cancer cases are included in cancer registries, whereas only 21% are included in databases throughout Central and South America, 11% in Africa, and 8% in Asia (Rodriguez-Galindo et al, 2013). Additionally, even if registries do exist, LMICs' cancer databases are often incomplete and inaccurate due to limited population data and political and economic instability. The rarity of pediatric cancer further amplifies these issues (Rodriguez-Galindo et al, 2013). This dearth in information is not only detrimental to the countries themselves in understanding pediatric cancer as a national challenge, but also to the global community as a whole in attempting to help grasp and address the divergence between survivability rates in between LMICs and HICs.

These combined factors make pediatric oncology in LMICs so broad and data-sparse that creating a comprehensive review is fundamentally difficult if not impossible. As a result, the investigation phase necessarily relies on narrower evidence, such as personal experiences and limited-in-scope statistics, to give a representative, albeit imperfect, survey of pediatric oncology care in LMICs. Doing so is in line with the body of literature on the subject and is thus treated as a structural limitation in our investigation.

Disparities in Diagnosis

An early, accurate diagnosis can be crucial to a child's likelihood of achieving a survival outcome of over five years. However, in LMICs, one of the greatest challenges for pediatric cancer patients is delayed diagnosis or misdiagnosis, such that by the time a child receives the correct assessment, it is often too late to treat the cancer effectively.

Lack of Cancer Awareness and Education

Lack of awareness and education about cancer in LMICs causes delays in diagnosis both patient- and physician-side. On the patient's side, delayed diagnosis often arises due to a lack of knowledge recognizing that a child's symptoms may be cancer or, at the very least, abnormal enough to seek medical attention. Across the board, the GLOBE Scholars reported that many cancer patients often develop

painless lumps indicative of a solid tumor, but, because it is painless, patients and their families will

assume it harmless and not seek care (GLOBE Scholars, personal communication). Additionally, patients

often confuse their cancer symptoms with more common disease symptoms, such as believing prolonged

fevers and headaches are caused by infection, or mistaking leg pain as growing pains typical in children

(GLOBE Scholars, personal communication). Indeed, the dearth of knowledge about cancer in LMICs is

illustrated by the fact that, as documented in a 2013 study of pediatric Burkitt Lymphoma (BL) patients in

Kenya and Uganda, only a small fraction of patient guardians – 5% in Uganda and 19% in Kenya – knew

that children could even develop cancer (Buckle et al, 2013). Consequently, patients often only seek

medical care when growths becomes painful and thus malignant, or when more serious symptoms of

cancer begin to take hold, by which point the cancer has progressed to an advanced stage (GLOBE

Scholars, personal communication).

 For the physician, lack of awareness of the early cancer signs is equally pivotal. Historically,

public health priorities in LMICs have long focused on highly communicable diseases such as HIV,

malaria, and tuberculosis (Samb B. et al., 2010). In fact, many LMICs, including some of those that the

GHI serves, lack formal pediatric oncology subspecialty training for physicians or nurses (Halbert and

Khaing, 2014). Consequently, if a general practitioner is faced with a child presenting general cancer

symptoms (such as a painless lump), cancer will rarely be the first ailment the practitioner considers, and

the symptoms will often be misdiagnosed as another malady (GLOBE Scholars, personal

communication). Thus, in these low-resource countries, patients with early cancer symptoms will often

not be directly referred to an oncologist or cancer specialist; instead, healthcare workers may refer the

patient to other doctors or specialists, who may then refer the patient to other non-cancer related doctors

or specialists, thereby resulting in a long and inefficient referral process (GLOBE Scholars, personal

communication). Indeed, as seen by patients in Kenya and Uganda in a 2013 study, many patients

experienced significant delays[3] in diagnosis due to having to visit many different healthcare nodes before being referred to a oncology facility (Buckle et al, 2013). This indirect process not only greatly delays the diagnosis, but, by requiring several different appointments and trips on the part of the patient, also induces further financial and personal strain that can often lead the patient to stop seeking care and never receive a diagnosis at all.

Distance and Poor Access to Healthcare Facilities

Even if a patient recognizes abnormal symptoms, poor access to healthcare facilities – including distance and transportation availability – and competing financial or household responsibilities often hinder a child's ability to receive care and diagnosis. For example, in Myanmar, there are only two pediatric oncology and hematology centers; by consequence, patients often have to travel hundreds of miles over several days to be treated (Halbert and Khaing, 2014). These journeys can be difficult to make, and the difficulties can be compounded by the lack of accessible and efficient transportation methods (Persistent Productions, personal communication). Additionally, these trips often require a child's guardian to travel with them, meaning the guardian may lose wages. In a 2013 study of pediatric BL patients in Uganda and Western Kenya, 35% of Kenyan guardians and 52% of Ugandan guardians surveyed said that no one else was available to take the child to the hospital (Buckle et al, 2013). These competing interests take a toll on the family and the community, and these obligations may be perceived as major obstacles in a guardian's decision to seek care for their child (Buckle et al, 2013).

Infrastructure and Resource Limitations

Limitations in resources, technology, and staffing in hospitals also play a significant role in slowing down the clinical diagnosis process. Even if a child has been referred from their local healthcare provider to a bigger, more well-resourced hospital, there still might not be adequate resources to correctly

---

[3] In this study, Buckle et al measured delays from the time of a patient's first symptoms of cancer to the time of diagnosis. Notably, in this sample of pediatric cancer patients, the median "total delay" was 12.1 weeks in Kenyan patients and 12.9 weeks in Ugandan patients (Buckle et al, 2013).

diagnose them (Ngoya, Muhogora, and Pitcher, 2016). Among the countries that the GHI represents, a hematological (blood or biochemical) test is relatively easy to obtain and run. However, investigative imaging tests are more challenging, as they require complex X-ray, CT, and MRI machines that are expensive, need constant maintenance, and require skilled technicians to operate (GLOBE Scholars, personal communication); a 2016 study conducted in Tanzania revealed that there were only 5.7, 0.31, and 0.05 units for general radiology, CT, and MRI per million people respectively, well below the World Health Organization standards of 20 units per million (Ngoya, Muhogora, and Pitcher, 2016). Consequently, many hospitals can only manage a very low capacity for a large volume of patients, if they have diagnostic imaging machines at all. For instance, at Dr. Nono Longombe's hospital in Congo, there is only one diagnostic imaging machine for nearly 2000 patients (GLOBE Scholars, personal communication). Once broken, the shortage of trained technicians on site often results in the machine remaining unusable for long periods of time (GLOBE Scholars, personal communication). This means that a typical test that might take a few days to process in a HIC could take several weeks or even months to process in a LMIC, further delaying diagnosis.

Through a combination of these factors, a majority of pediatric cancer diagnoses in LMICs often occur at the advanced stage. For instance, in HICs, retinoblastoma typically presents while still intraocular; in LMICs, 60–90%of children present with extraocular tumor (Gupta et al, 2015). Indeed, Dr. Nono Longombe reported that advanced stage cancer diagnoses were the most common pediatric cancer diagnoses he encounters at his local site (GLOBE Scholars, personal communication). This late diagnosis results in poor prognosis later on, as early identification of cancer has consistently been correlated with treatment success and increased likelihood of survival (Buckle et al, 2013; Kanyamuhunga et al 2015).

Disparities in Treatment

Although proven effective when carried out correctly and completely, the implementation of multidisciplinary care and treatment regimens in many LMICs is often hindered due to insufficient resources, technology, staffing, and personnel specialization.

Surgery and Chemotherapy

Treatment options are limited in LMICs, depending on the resources and personnel available, but overall, surgery and chemotherapy are relatively common, with surgery attributed with 49% of survival rates and chemotherapy being the most abandoned for financial reasons (Bishr et al, 2018; Ribeiro, Antillon, Pedrosa, and Pui, 2015). The former is the most readily accessible form of treatment, as surgery sometimes does not require a specifically trained oncologist to perform; however, the relative popularity of surgery is countered by its limited efficacy in treating metastasized tumors (Bishr et al, 2018). Chemotherapy is usually available and used in combination with surgery when possible, but chemotherapy drug supplies are not always constant, and administration is delivered by non-specialists as mentioned by a physician-nurse pair from Iraq (GLOBE Scholars Interview, personal communication). For example, Yangon Children's Hospital, a major pediatric cancer center in Myanmar, has access to 17 different types of widely used chemotherapy drugs, but access to asparaginase, a crucial drug used to treat acute lymphoblastic leukemia, is limited (Halbert and Khaing, 2014).

Radiotherapy

In LMICs, radiation treatment strategy for pediatric cancer patients is restricted by a lack of trained personnel and fear of overburdening the limited equipment (Bishr et al, 2018). To illustrate, the International Atomic Energy Agency recommends a ratio of 250,000 people to one megavoltage machine; in LMICs, such as those in Latin America and Africa, this ratio is much higher, reaching 3.56 million and 0.56 million members of the population, respectively (Bishr et al., 2018). As of 2014, 39.5% of LMICs

did not have access to a single radiation therapy facility (Datta, et al). Even if lack of equipment was not

an existing problem, there are a limited number of trained radiation oncologists throughout LMICs. For

example, only 29.3% of Latin American countries have classes or some form of radiation oncology

training, and even fewer African countries (18.5%) do the same (Bishr et al., 2018).

Human Resources: Understaffing of Treatment Teams

Hospitals in LMICs experience a severe lack of trained oncologists, especially those specializing

in pediatrics. Dr. Nono Longombe, for example, is one of only two oncologists servicing the entirety of

the Democratic Republic of the Congo, a nation of nearly 85 million people (Globe Scholars, personal

communication). In contrast, the United States boasts over 20,000 oncologists for a population of 325

million people (Kirkwood M., 2013; US Census Bureau, 2017). Dr. Longombe is responsible for over

2,500 times the number of patients as the average American oncologist, and thus is severely

disadvantaged in amount of direct treatment he and oncologists in other countries under similar

circumstances can provide.

Additionally, pediatric oncology care requires a multidisciplinary team that extends past the

expertise of just general practitioners and oncologists. An ideal team consists of physicians, nurses,

pharmacists, dieticians, psychologists, social workers, and other care experts who all work to treat a

patient's needs (Gupta et al., 2015). However, in LMICs this team often includes only one physician and

one nurse (Dana-Farber Nursing Team, personal communication). Therefore, the role of nurses within

LMICs has evolved to include responsibilities such as monitoring symptoms, taking vitals, and

administering chemotherapy drugs, as well as providing psychological and nutritional support. Their work

is further complicated by the lack of technological resources available to them. For instance, the nursing

team at Boston Children's Hospital emphasized the need for precise anthropometric instruments to help

combat malnutrition and undernutrition, which are common side effects of cancer treatment (Dana-Farber

Nursing Team, personal communication). A high volume of patients, coupled with the multitude of

responsibilities that need to be fulfilled with limited resources and non-specialized training make nurses the overworked nexus of care in LMICs.

Consequently, additional psychosocial teams and nutritionists are integral to pediatric care. This is especially true in LMICs where psychosocial work will not only be an important step in aiding the mental health of the patient and family but will also effectively work towards reducing rates of treatment abandonment as they work against stigma to increase cultural awareness and education of cancer (Gupta et al., 2015). For instance, in 2005, Indonesia introduced an educational video for cancer patients and families which led to a decrease in treatment refusal from 14% to 2% (Mostert et al., 2010). During our interviews at Dana-Farber/Boston Children's, Jorge Fernandez, licensed social worker, emphasized the immense effect a social worker could have through simple tasks, such as calling families for appointment reminders, a task which reduced abandonment rates from 17 to 2% (Dana-Farber Psychosocial Team, personal communication).

Another underrepresented role, often assumed by nurses in LMICs, is that of nutritionists. A report from 2014 discussing nutritional practices and resources in LMICs reported that only 55% of LMICs performed nutritional assessments within the hospital, and only 35% had nutrition education programs for families and patients (Murphy et al., 2014). The main reasons for the dearth of nutritional education are a corresponding lack of specialized personnel, and limited time and financial resources (Murphy et al., 2014). Nurses lack the educational background of a nutritionist to provide full support to a patient. The GHI nursing team has aimed to implement a nutritional support program in Myanmar that will provide dietary educational resources to patients and families as well as increased training for nurses specifically on nutrition (Dana-Farber Nursing Team, personal communication).

Abandonment of Treatment

Beyond the shortages in technological and staffing resources that many LMICs face, abandonment of treatment is yet another significant obstacle in treating pediatric oncology in LMICs

(GLOBE Scholars, personal communication). In a study in western Kenya, researchers found that 54% of pediatric cancer patients ultimately abandoned treatment (Njuguna et al., 2014). Another study on Yangon Children's Hospital in Myanmar recorded a 25.3% abandonment rate for acute lymphocytic leukemia treatment and a 60% abandonment rate for retinoblastoma treatment (Halbert and Khaing, 2018). Abandonment rate varies vastly across particular countries, hospitals, and cancer types, but ultimately the academic literature demonstrates that LMICs consistently have higher treatment abandonment rates than HICs, and this treatment abandonment rate can explain more than a third of the survivability gap between LMICs and HICs (Friedrich et al., 2015).

A primary reason for treatment abandonment is financial strain. Although patients often do not need to cover the cost of medical expenses in their home countries, as either churches or the government would contribute to pay their expenses, finances are still a significant impediment to treatment (Globe Scholars Interview, personal communication). This obstacle mainly arises from transportation costs and the fact that accompanying a child to treatment often precludes a parent's ability to work (Globe Scholars, Dana-Farber Nursing Team, Psychosocial Team, personal communication). Furthermore, due to the ubiquitous overcrowding of hospitals in LMICs, many of these hospitals are unable to accommodate more than one guest with each patient, so any siblings or second parents that travel with the child to treatment must find other means of lodging (Dana-Farber Nursing Team, Psychosocial Team, personal communication).

In several countries, cultural norms and attitudes towards cancer also act as impediments to treatment adherence. For instance, in Iraq, family approval is needed before a child can start cancer treatment, which often leads to some children not obtaining recommended surgery, chemotherapy or radiation treatment (GLOBE Scholars, personal communication). Similarly, in Armenia, it is not common practice to communicate a cancer diagnosis to patients even when they are adults (Dana-Farber Psychosocial Team, personal communication). In some LMICs, cancer is often stigmatized, which can

prevent patients from seeking treatment after receiving a diagnosis (Rodriguez-Galindo, 2013). For

instance, in a 2017 study of cervical and breast cancer in India, stigma related to cancer was a common

contributor to many who choose to postpone seeking treatment; specifically, they found themes related to

three main areas: fear of casual cancer transmission, personal guilt, and inevitability of death or disability

(Nyblade et al, 2017). In the Congo, many patients prefer to leave the hospital and go to a general

practitioner who merely prescribes antibiotics rather than correctly diagnosing or treating the cancer

(GLOBE Scholars, personal communication). Some patients in Uganda and surrounding areas even elect

to see a witch doctor over a licensed oncologist as a result of the culturally ingrained belief that cancer is

a form of witchcraft rather than an illness that can be adequately treated by mainstream medicine

(GLOBE Scholars, personal communication).

Another major contributor to treatment abandonment is low literacy rates in patients,

which leads to breakdowns in communication between the patient and their healthcare workers. LMICs

have, on average, much lower rates of literacy than wealthier countries, which can make communication

between hospital staff and families more difficult ("Literacy of Adults in Developing Countries", 2016).

Consequently, medical diagnoses in general can be confusing for families who are unfamiliar with

medical terminology, which makes it quite difficult for people to understand the purpose, prognosis, and

methods of their children's treatment, leading patients and families to abandon treatment altogether

(Sudore et al, 2009).

1.5 | Research and Information Synthesis

After investigating the myriad challenges that hinder a child's chance of survival in LMICs, the

next task was to crystallize and unify those challenges into a single problem-statement.

Throughout the investigation phase, data collection and organization was oriented towards

understanding the factors at play in each phase of a pediatric cancer patient's care: oncology diagnosis,

treatment, and post treatment. When new research or data was received from interviews, articles, or client-meetings, information was filed according to how it influenced these three phases of care (see Appendix, Figure 1). *Diagnosis* encompassed recognition of symptoms by patients and their families, arrival to primary care, clinical diagnosis, and registration in database, while *treatment* was a two-tier understanding of the administration of treatment and patient adherence to treatment protocol and *post-treatment care* which develops an understanding of outpatient care and management of therapy side effects.

Visual Mapping

After collecting all the data from research, interviews, and client meetings in Week 3, the team synthesized and analyzed the data points collected by creating a variety of visual mappings and process-interaction diagrams. In subgroups, teams crafted information diagrams that categorized the data, processes, and observations we gathered in the investigation phase.

In the first iteration of synthesis, the teams organized the data in one of two ways: relating to time and relating to actors.

One visual mapping compared the relationship between different data points by creating a timeline of the treatment process graphed against the different factors affecting a patient's chance of survival (see Figure 2). Graphing the data's relationship to other points over time visually highlighted how crucial timing is to understanding and improving the care delivered to pediatric patients; specifically, when visually arranged, the clustering of information close to the short-term end of the time axis highlighted the sense that many of the challenges and processes contained in our data were time-dependent – not only related to timely cancer diagnosis, but also to the importance of efficient and immediate response to symptoms, patient referral to healthcare workers and specialists, and more (see Appendix, Figure i.2).

In the second visual mapping, we analyzed the actors involved in the process of delivering care by arranging data points under an egg scheme with categories of "healthcare system," "patient and family," and the "larger community" (see Figure i.3). The distinct actors enveloped in these categories were clarified by a persona exercise carried on in class. By detailing the specific concerns and characteristics of the various personas at play such as patients, families, community leaders, and healthcare staff, we were able to better understand the people involved with pediatric oncology care in LMICs. Creating detailed personas for those involved ensured we would keep the end user in mind when identifying our problem statement and future work. We called this way of visualizing our personas the social-ecological scheme, as it took into account the unique actors, and the environment they generally operated in. While this social-ecological scheme lacked the dimension of time, it did reveal an interesting trend: despite the fact that we tried to separate our information into clear-cut categories of actors and reactors, without fail it was seen that categorized actor-diagrams would consistently contain data concentrated on the borderlines – that is, we found most data points were lying along the intersection of different actors. This revealed the importance of the interface and interactions between various personas/actors as a part of the problem.

Problem-Identification: Delays in Timely Care

Hence, in our next iteration of synthesis, we combined the social-ecological-scheme and timeline into a pie chart with a time dimension, to try and visually understand how these two trends come together, and lead to the problems faced by LMICs' pediatric cancer patients (see Appendix Figure i.4). In this pie chart, most data points laid along the short-term axis, and at the intersection of actors. Indeed, various combinations of our data organization schemes led to the same conclusion: the interface between different actors is where the flow and coordination of care seems to breakdown in LMICs, and causes significant delays in receiving accurate, timely diagnosis and treatment across all phases of a pediatric cancer patient's care. We concluded that this breakdown in timely care is a problem that we wanted to address.

Causes of Time-Dependency and Delays

Having identified time dependency and delays as a key issue, we identified the causes of these delays and categorised them into four main themes. The first is challenges related to with information transfer. This can take on many forms, like transferring patient data within hospitals, getting the preliminary reports from the primary care physician, or informing the patient of follow-up appointments. Many of these communications can be time sensitive and affect patient health outcomes. Secondly, we identified difficulties with patient and doctor education surrounding cancer and awareness of the possibility of cancer. More specifically, doctors may not consider cancer as a possible diagnosis and on the patient side, the severity of the diagnosis may not be immediately internalised. This may result in a late correct diagnosis or delay in action being taken by the patient. Thirdly, and related to the previous point, there may be a late cancer diagnosis due to limited resources that are at disposal to a doctor. For example, if the doctor is seeing many patients, they may not have the time to examine every patient thoroughly or even at all. Likewise, nursing teams may not be able to properly evaluate the patients due to the patient-to-nurse ratio being too high. Finally, there may be challenges with communication between caregivers. Caregivers could struggle to communicate key patient information between themselves, or quickly communicate a diagnosis to the relevant teams, that may require measurements and exams to be done again, for example.

Having recognized these themes, we drafted a set of criteria through which we could examine each theme and assess where our team could best identify and synthesise a nuanced, approachable issue, as we felt that these themes were inherently connected together, and that if we were able to identify how they were connected, then we may have the ability to make the greatest impact. Our criteria consisted of the team's ability to aid in GHI's mission to reduce the outcome gap, identify a specific issue to assist in ideation, have a targeted and measurable impact, implement the solution at scale to help as many people

as possible within reason, and ultimately to implement and feasibility engineer a solution. We included

feasibility in our criteria as a sanity check, rather than considering specific solutions, to keep in mind that

there should exist solutions to the problem that are achievable in the time frame of two months.

Overarching Themes

Through this filtration process, we investigated the problems that surrounded the overarching

theme of data that we observed, and did further research to determine the root causes of the problem. We

finally narrowed down the causes that met our criteria into three main causes of the problem: challenges

in the *collection, communication, and interpretation of data*. We found that, inherently, the issues that

these three categories encompass are inextricably linked, and, as a result, go hand-in-hand in contributing

to the disparities in pediatric cancer patient outcomes in LMICs. In ideation and creation therefore, any

ideas or conclusions we determined, had to be made while keeping in mind the intertwined nature of the

problems faced by LMICs as it pertains to data. This entire process which led us to this tri-fold conclusion

is summed up in Figure 1.2 below.
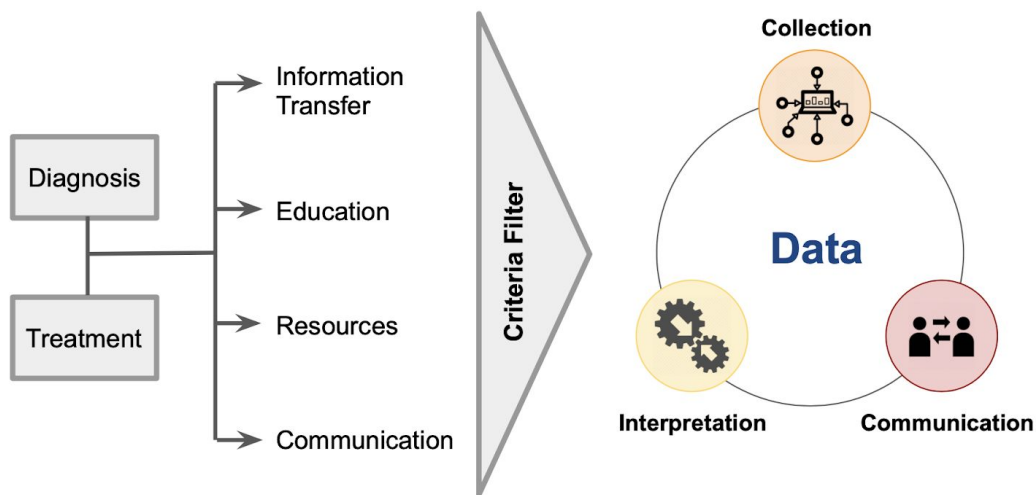
Data collection

In LMICs, data collection in pediatric oncology is a multifaceted issue complicated by the high case to oncologist ratios and high patient to nurse ratios often found in these countries. Research indicates, for instance, that there is a significant shortage of clinical oncologists in 25 countries in Africa (78%) and two countries (11%) in Asia (Matthew, 2018). This is further worsened by the lack of diagnostic infrastructure, such as the low availability of radiotherapy machines and high wait times for the machines that are available. In countries categorized as LMICs, for example, there are about 4,400 megavoltage machines, which account for less than 35% of the world's radiotherapy facilities, while a cancer incidence rate indicates a need of of 9600 units. For example, as seen Figure 1.3 below, "Europe has 17 times as many radiotherapy units as are available in Africa per million inhabitants, while Latin America and the Caribbean region has just one-third of the number of machines available per capita in North America." This means that the countries with the highest case to medical team ratios are also the ones with the most limited medical infrastructure, which renders data collection a particularly challenging task.
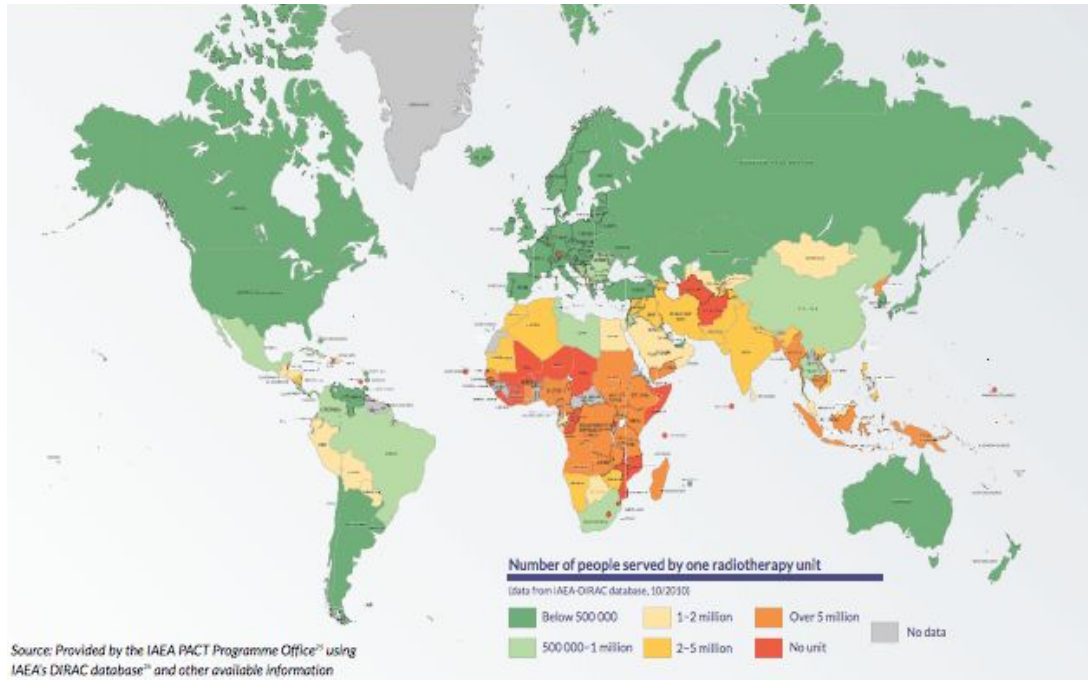
Figure 1.3: Global map of number of persons in a country per radiotherapy machine present (Mathew, 2018)
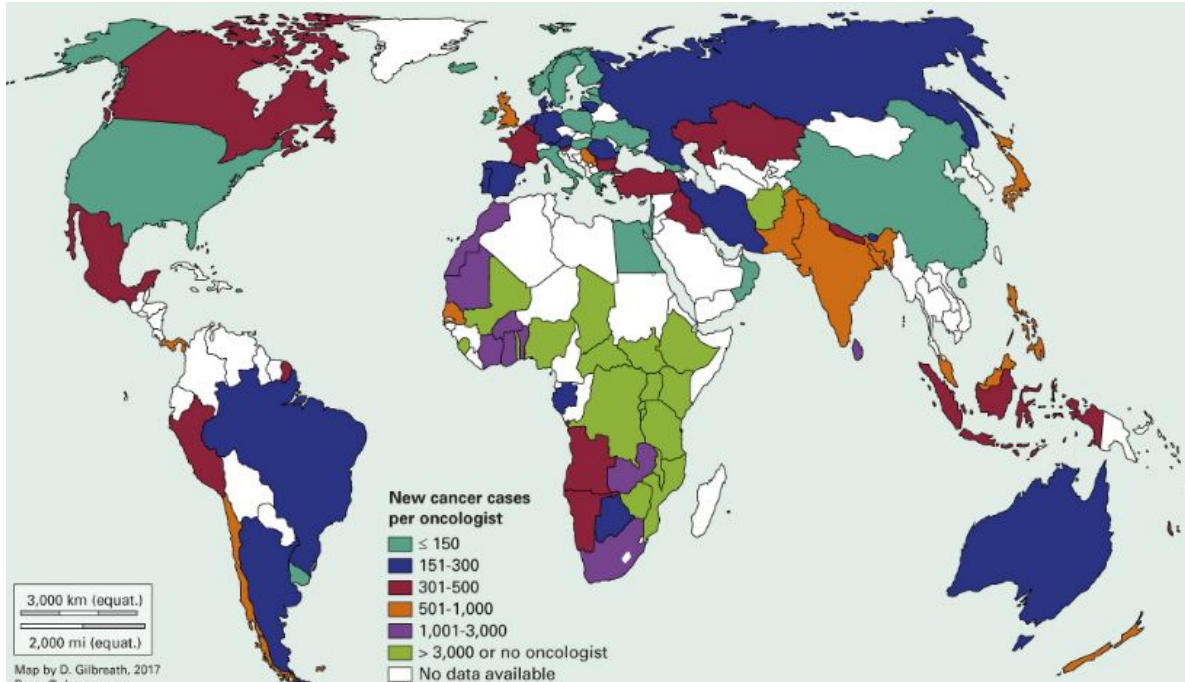
Figure 1.4: Global map of number of cancer cases per oncologist present in a country (Matthew, 2018)

Data Communication

Communication failures within the medical team are the leading root cause of medication errors, wrong-site surgeries, and delays in treatment, according to the Joint Commission on Accreditation of Healthcare Organizations (Hughes, 2008). These problems can be exacerbated in LMICs which often do not have the tools to properly record and communicate patient data. One common method of communicating patient data within a medical team is through electronic medical records. As shown in the following graph, there is a large discrepancy in the use of electronic patient records in high income countries vs. LMICs. In the US, 96% of hospitals use electronic medical records and 88% report clinical benefits as a result of these systems (Patel, 2012).

Use of Electronic Patient Records Worldwide
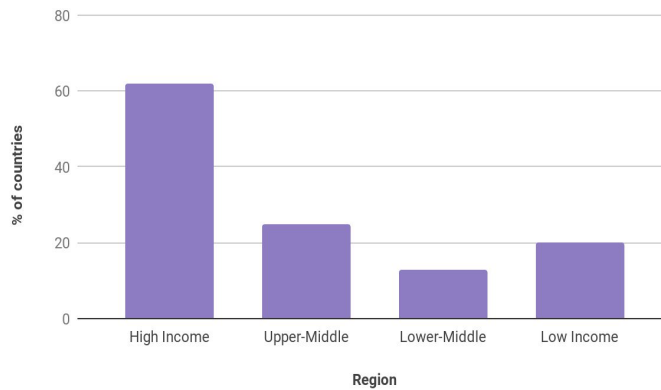
Figure 1.5: Percent of countries using electronic patient records as categorized by their income. (WHO, 2002)

Electronic medical records can be used to create cancer registries, another vital method of data communication globally. Cancer registries help researchers identify causes of cancer and potential future treatments by looking at trends in patient data. As shown in the graph below, many LMICs lack cancer registries with only 11% of cancer cases included in cancer registries in Africa and 8% in Asia, compared to 90% in North America and Oceania. The clear inequality in the tools and methods available in LMICs to record, access, and communicate patient data presents great challenges in improving cancer treatments and outcomes in these regions.

Cancer Registries Worldwide
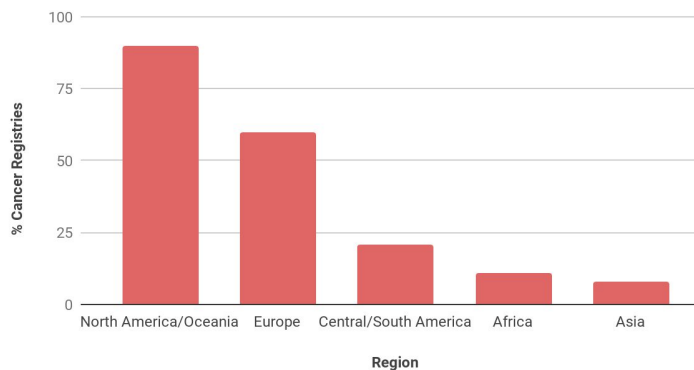
Data Interpretation

Even if challenges in data collection and communication are overcome, data interpretation remains problematic. Pediatric oncology nursing care requires advanced clinical skills and an extensive specialization in pediatric cancer, but this currently only happens in HICs. While nurses play a significant role in the provision of all treatment modalities, often times nurses in LMICs are not formally trained for pediatric oncology. This is a direct result of the unavailability of oncology specialty education, and the lack of opportunities for nurses to specialize due to mandatory nursing rotations throughout the various clinical settings in a hospital (Day et al, 2015). As a consequence, cancer symptoms may be missed, the treatment strategies offered are limited by the expertise and experience of the nurses, and nurses are not fully equipped to deal with the nutritional and physical consequences of receiving treatment.

Problem Statement Development

To bring together our findings into one problem statement, we considered how each classifies as the problem, causes, affected population, or desired outcome. Our problem arose from time: the delays in accurate diagnosis and treatment of LMIC pediatric cancer patients. The cause of this problem is challenges in data collection, communication, and interpretation. Here, our affected population and primary stakeholders in this problem are the actors where that breakdown in data communication and collection occurs: the patients, general practitioners, and cancer specialists in LMICs. Our desired outcome is to aid the mission of our client, GHI, by alleviating the disparity in outcomes in HICs and LMICs through the improvement of cancer survival rates for LMIC pediatric cancer patients. Thus, our investigate phase culminated in the following problem statement:

Problem Statement

Patients, healthcare workers, and cancer specialists in LMICs need better methods and technologies to collect, communicate, and interpret patient data in order to allow for accurate, timely diagnosis and treatment and ultimately increase the likelihood of patient survival.

Affected Stakeholders: patients and families, general physicians, cancer specialists, healthcare workers

Desired Outcome: increased likelihood of patient survival

Problem: delays in timely diagnosis and treatment for pediatric cancer patients

Cause of the Problem: challenges collecting, communicating, and interpreting data

With this problem statement in mind, we moved on to the Ideate phase, where we brainstormed ideas to try and address this need.

Chapter II:

Ideate Phase

2.1 | Chapter Overview

This chapter serves to summarize the team's ideation phase. The purpose of the ideation phase was for the team to systematically converge on a design solution, or a limited set of design solutions, that address our problem statement. Namely, the design solutions must facilitate the collection, communication, or interpretation of patient data, lead to timely diagnosis and treatment of pediatric cancer, and ultimately take some small step to improve the survival rates of children with cancer in LMICs.

The initial brainstorming generated approximately 84 ideas, which were narrowed in the winnowing process by a qualitative group analysis to a set of 15. Through the use of a Pugh chart and a preliminary set of feasibility criteria, those ideas were further reduced to 6 in the reduction stage. Ultimately, a final feasibility analysis, utilizing a more in-depth Pugh chart, as well as a re-iteration through the ideate process, resulted in the choice of one idea to pursue: an in-patient vital monitoring system for early detection of sepsis. The chapter proceeds to describe each process of the ideation phase in detail.

2.2 | The Brainstorming Process

Independent Brainstorming

The team started the Ideate phase with independent brainstorming, to come up with a wide variety of solutions not influenced by others' ideas. To avoid restricting the creativity and novelty of the solutions we came up with, the team initially did not consider general technical or implementation feasibility. However, following idea generation, each team member performed focused literature review on his or her ideas to ensure that each proposal was based on valid research. Through this individual brainstorming, the team came up with approximately 100 total ideas, although many were repeated or similar to one another.

Group Organization

After independent idea generation, the team broke into three groups to sort and organize our ideas. The primary purpose of this organization was to determine the broad categories under which the ideas fell and to identify solution spaces related to our problem statement that current products did not yet address.

The team organized ideas in two main ways, although iterations of each organizational method were employed. In the first organizational method, a grid mapped the intersections of different actors, including doctors, nurses, patients, community members, and the GHI. This framework forced us to think of the use-case of each potential idea. Team members placed sticky notes with brainstormed ideas on the grid to indicate how ideas could facilitate communication

between different actors (e.g. nurse-to-nurse communication). The problem statement encapsulated data collection in addition to data communication, so these ideas were categorized as doctor-to-doctor or patient-to-doctor communication. The grid had obvious gaps in communication related to nurses, communication related to the GHI, and communication related to the patient's community. This organizational scheme demonstrated that these areas could be explored further for potential ideas missed during individual ideation. In the second organizational method, ideas were placed under broad categories, such as education, diagnosis, and post-treatment, which related to specific processes in pediatric oncology care. This framework helped us understand the major factors affecting pediatric oncology care that our solutions targeted and which areas were lacking in solution ideas. After multiple iterations, the primary solution categories were identified as diagnosis, treatment infrastructure and resources, education, and record keeping.

Winnowing Process: 84 to 15 ideas

Ultimately, when duplicates were consolidated, we had 84 distinct ideas. As a group, we went through each idea and subjected it to four general levels of criteria. In order to most efficiently and effectively use the team's energy with a vast number of ideas, this assessment was strictly qualitative. Our first was whether or not the solution was applicable to our problem statement. However, because the brainstorming process took the problem statement into account, only a handful of ideas were eliminated in this way. The second criterion was a surface level feasibility analysis–we wanted to limit our energy to analyzing solutions that we could realistically produce. To assess feasibility, we considered time and resource constraints as well

as whether or not our team had relevant skill sets. Our third level of analysis considered the scale

of the solution. We evaluated whether the solution was too easy or too substantial for the

constraints of the class by considering the amount of time and level of expertise needed to arrive

at a functioning prototype for each idea. The last criterion we subjected our ideas to was the level

of potential impact. Assuming success in engineering a functioning solution that could be

integrated into our target users' workflows, we considered the potential impact and the size of

the affected population. By assessing ideas based on these criteria, the team landed upon 15

potential solutions.

Initial 15 Ideas

The following 15 ideas were the result of our generalized winnowing process. Ultimately,

we went on to evaluate each idea via a thorough Pugh chart process, but initially we conducted

more in-depth research on each idea to understand its respective strengths and weaknesses.

Overall, each idea addressed the problem statement, was relatively feasible within the scope of

the class, had a reasonable scale, and a large potential impact. Below we have provided

descriptions of each of these 15 ideas as a result of our further research into and discussion of

each solution.

Malnutrition Assessment and Intake Monitoring

The Malnutrition Assessment and Intake Monitoring Software would allow all members

of a care team to manually enter relevant nutritional data into the software, such as height,

weight, and mid-upper arm circumference. Using standardized algorithms, the software would

compute a simple nutritional diagnosis with recommendations on combating a patient's

malnutrition. Alerts would be displayed on a patient wearable that would be remotely connected to the software. It would inherently encourage communication between healthcare professionals of an interdisciplinary team while also increasing patient and family awareness and education regarding malnutrition. Because we noticed insufficient and untimely monitoring, identification, and intervention of patients suffering from malnourishment, our software would aim to increase the effectiveness and quickness in making a malnutrition diagnosis.

Diagnosis Web App

The Diagnosis Web App would be an app that flags symptoms that correlate with a cancer diagnosis. General practitioners would be able to input data such as age, sex, symptoms, and anthropometric data into this app. Many primary care physicians do not know about symptoms presented by cancer and therefore are not able to come up with a correct diagnosis. Our app would facilitate cancer diagnosis without assuming that the user has significant awareness of cancer symptoms. Using this tool, general doctors could communicate this data to oncologists earlier in order to increase patients' chance of survival.

Matching Lodging App

The Matching Lodging App is a web-based or mobile app that would match patients' families with host families near health care facilities, eliminating some of the financial and social burden of finding a place to stay for families travelling long distances to seek treatment for their child. The app would assess how far a patient's home is from their treatment facility, and if they are outside a certain radius, the family would be prioritized and matched with a family within

walking distance or an easy transit route to the facility. Impact would be measured by observing the degree to which treatment retention rates increase following the implementation of the system. This idea was inspired by the various lodges converted into communal lodges for cancer patient families in the United States.

Machine Learning Diagnosis Platform

The Machine Learning Diagnosis Platform would ideally reduce the amount of time spent by pathologists to diagnose leukemia, streamlining the process and increasing accuracy. Low and middle income countries have a scarcity of pathologists to conduct timely diagnostic tests, including blood and bone marrow smears used to test for leukemia, the most common of pediatric cancers, which can cause hospitals to outsource these tests, if they take them regularly at all. The solution would have two components: an app where pathologists upload images of blood and bone marrow smears, and a server running in a central location that would analyze images using machine learning and provide a diagnosis along with a certainty rating.

Medical Records App

The Medical Records App addresses the problem of inaccessible medical records caused by disorganized paper systems in many hospitals in LMICs. With this app, a nurse would be able to enter important patient medical information into an electronic form, which would then transfer the data into a searchable database of past medical records. The database could then be backed up to a cloud storage system or a network of external hard drives. The proposed system would be minimally interruptive to nurse workflow since it utilizes technology that they are likely already

familiar with, such as phones or tablets, and simply replaces written data entry rather than

introducing a new procedure (Top Countries/Markets by Smartphone Penetration & Users).

Other electronic medical record softwares have had considerable success in LMICs. For

example, in a case study in three East African countries, hospital staff overwhelmingly agreed

that EMR systems improved patient care and productivity (Tierney et al., 2010).

Wearable Medical Record

The Wearable Medical Record aims to combat the problem that many personal medical

record systems in LMICs fail due to loss of the cards or notebooks they were written on. These

paper record systems prevent access to a patient's previous medical data and thus make

communication of patient data difficult. To combat this, we hoped to implement an

oncology-specific personal medical record system that patients could not easily lose or forget.

Specifically, we envisioned a piece of wearable technology that carries encrypted medical data

that can be accessed via a technology such as USB, radio-frequency identification (RFID), or QR

code. The data would be secure and only decryptable by a trusted physician or the patient herself,

so the patient's information would be kept private. Having access to electronic medical records

like those stored on the Wearable Medical Record has proven benefits in HICs for faster retrieval

of records, improved care, and enhancement of coordination within the care team, so we imagine

that these advantages can be transferred to lower income countries through this device (Gyamfi

et al., 2017).

Communication Platform

The Communication Platform is a web-based or mobile app that connects hospitals and LMICs to teams in HIC's of relevant focus and expertise, where there may be a lack of communication and data exchange. This can be considered an extension of the twinning program. Our app allows for a more expansive approach and more specific matching through an algorithm. The user profile for the solution will be those hospitals which have access to Internet. Ideally this will lead to an exchange of information, which will indirectly improve both knowledge and approaches to pediatric cancer, ultimately leading to an increase in pediatric cancer survival rate.

Anthropometric Data Collection Multi-Tool

The Anthropometric Data Collection Multi-Tool is a device consisting of a hanging-style scale and integrated length board that would allow nurses to easily and accurately take an infant's height and weight. The problem we are trying to address with this device is the lack of material, human, and financial resources that lead to incomplete and sometimes inaccurate data being taken for infantile pediatric cancer patients in LMICs. This device would allow nurses to more accurately assess an infant child's weight. In addition, the tool would streamline the process of anthropometric data collection, letting overworked nurses take more measurements in a shorter amount of time. Improved collection of this data could be used for improved malnutrition tracking for patients, allowing healthcare providers to intervene more quickly and

thus anticipate for and proactively perceive complications and poor outcomes stemming from pediatric malnutrition.

Optical Character Recognition (OCR) Record Digitization

The OCR Digitization is a combination of hardware and software to scan existing records and upload them to a queryable database that can retrieve records instantly via a search of the patient's name or ID number. This idea would likely take the form of a web or phone app or use a small cheap computer with nurses and physicians as the users. This solution requires minimal nurse or doctor workflow intervention and uses keywords to parse the data and sort it into an on-site database. This would improve timely access to the patient data to ensure that long-term treatment plans are adhered to, data trends noticed, and that the correct treatment is administered, ultimately leading to better health outcomes.

Follow-up Appointment Reminder and Tracker

The Follow-up Appointment Reminder and Tracker is a wearable device that both reminds the patient of their appointment date and allows hospitals to monitor patient's appointment attendance. The watch would act as an alarm clock for the patient, while also notifying the hospital when a patient has missed their appointment in order to intervene and encourage treatment adherence. This would help streamline the process of patient tracking by digitizing the system and improving communication between patient and doctor in order to reduce treatment abandonment rates.

Matching Patients to Survivors App

The Matching Patients to Survivors App is a communication platform to match cancer patients with survivors in order to alleviate stigma and aid with patient education of cancer. Survivors would volunteer their time to support cancer patients of similar cultures and language backgrounds and connect digitally on our platform. The goal of this concept is to reduce treatment refusal and abandonment by combating stigma and lack of education.

Data Collection Toy

The Data Collection Toy is a data logger and a children's toy. On the outside the tool is a comforting toy for patients to keep with them throughout their cancer treatment, but the inside of the toy is engineered to measure and store biometric and environmental data such as body temperature, pulse, and respiratory rate. This idea would aim to bridge the gap between medical procedure and child's play, thus allowing for less emotional stress on the child while also giving doctors and nurses an easy to use and intuitive data collection tool. In practice, the child would interact with the toy, an internal sensor mechanism would record specified vitals, and then store the information in a retrievable/accessible form such as USB hardware or Wi-Fi compatibility. We would be able to measure the success of the toy implementation by monitoring both increase in quantity of vitals collection data as well as a reduction of time required on the nurse's side to collect vitals. Additionally, a qualitative measurement of impact would be an increase in quality of experience for the child both mentally and emotionally.

Visual Inpatient Monitor

The Visual Inpatient Monitor is a two-input visual monitor that would allow nurses to quickly assess the status of a patient and the time since their last nurse check-in. The device itself would consist of a button for the patient to request assistance and a button for the nurse to reset the "time since last check in" clock. There would also be a visual output consisting of either a red frowning face (patient button pressed) or green smiling face (button not pressed). Allowing nurses to visualize the status of their patients could allow them to optimize their workflow, combatting the low amount of time nurses can spend with patients due to extremely high nurse to patient ratios.

Chemotherapy Drug Database

The Chemotherapy Drug Database is a platform that would suggest drug prescriptions for a patient's cancer care based on a set of parameters. This platform would aid the low number of oncology specialists and address the inability for doctors to communicate with oncologists. The platform would use an algorithm that takes in the patient's type and stage of cancer and returns possible drug administration plans. As a result, this platform would allow for more accurate chemotherapy treatment plans due to the absence of communication between doctor and oncology specialist.

Cardiac Vital Monitoring Device

The Cardiac Vital Monitoring Device seeks to address the issue of cardiotoxicity, which is the second largest cause of treatment-related deaths (Board et al, 2003). By conducting inpatient measurement of cardiac ejection fraction, this device could collect data to prevent cumulative cardiotoxicity. The device would measure a patient's vitals via Doppler methods, after chemotherapy treatment has been performed. Healthcare workers can observe trends in ejection fractions to judge how future treatment should be managed or whether the patient was suffering from cardiac complications. The intended user for this product would be doctors.

2.3 | Reduction Process of 15 to 6 Ideas
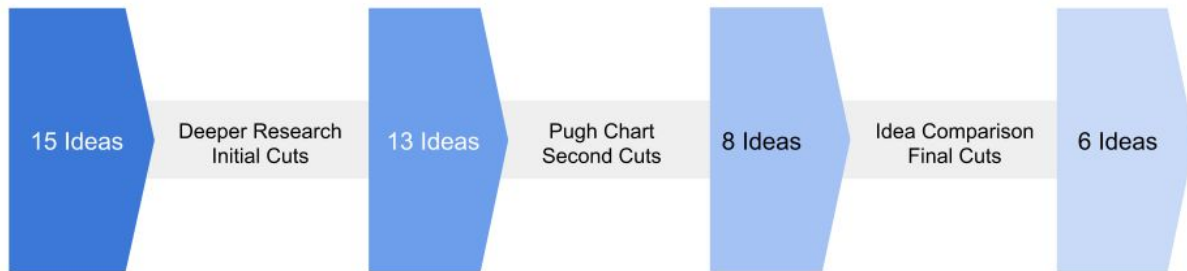
Figure 2.1. Outline of steps in the reduction phase, to narrow from 15 ideas to 6.

Eliminating Assumptions and Deeper Research

To flesh out and narrow down our 15 ideas to a smaller subset of final candidates, we decided to 1) map out and eliminate any assumptions that our solutions might be making, 2) identify holes in our knowledge or research regarding the feasibility of our solution, and 3) fill those holes with deeper research and adjust our ideas accordingly.

We mapped out and eliminated assumptions for each solution using a flowchart, which helped explicitly outline the implicit dependencies required for each solution to be effective. We modeled this process based off of the structure of a mathematical proof, which gradually builds to a conclusion by establishing smaller truths. In other words, by breaking down our solutions into smaller steps, we wanted to determine if there were reasonable grounds to believe whether each step would be achieved. The template followed was as follows:
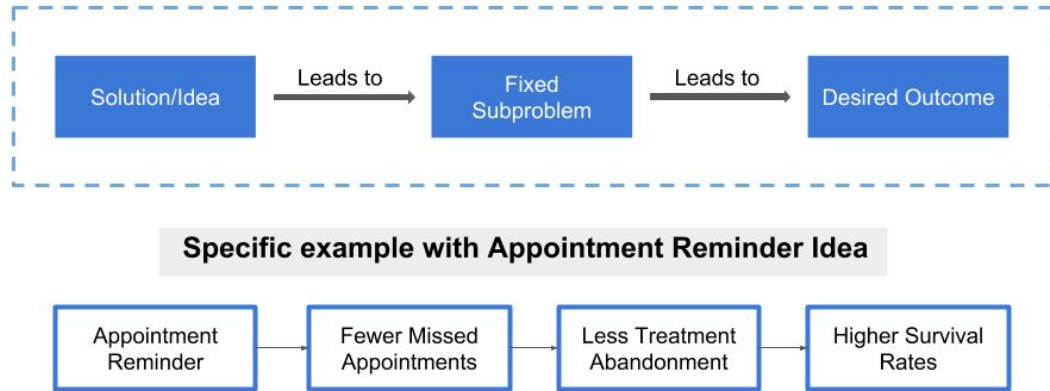
Figure 2.2: Flowchart template used to map out logic behind each potential solution. Each arrow indicates a dependency that is not necessarily true. For example, it is not necessarily the case that a reminder leads to fewer missed appointments, as patients might be aware of the appointment and continue to miss it due to financial strain or other circumstances.

While this flowchart was simplistic, it allowed for an insightful analysis: in particular, it did well to highlight which steps or connections needed more significant support from data. After identifying which connections were more tenuous or data-sparse, we targeted these areas and assigned sub-teams of 3-4 people to thoroughly investigate our areas of concern. If there were areas of concern that we believed could be overcome by shifting the focus of the solution itself, then sub-teams would work to adapt the idea and adjust the dependencies required. For example, given that the data we had did not strongly support our claim that appointment reminders would increase appointment adherence in LMICs, the focus of the Appointment Reminder was completely shifted to become the Appointment Tracker, where it gained the additional feature for cumulatively tracking and sorting patients' missed appointments to try to help hospitals identify patients at risk of treatment abandonment.

If, after further research, there were still significant gaps between the different dependencies required for our solution to work, with no foreseeable way to bridge them, then the idea was eliminated by consensus vote of the entire team. This resulted in two of the 15 ideas being eliminated: the Chemotherapy Drug Database and the Matching Patients With Survivors App, as they were deemed infeasible given the resources, skills, and data available to the team.

Initial 13 Ideas Pugh Chart Elimination

With the remaining 13 ideas, we decided to utilize a simple Pugh chart to assess the relative strengths and weaknesses of each and make further elimination. The first step of making this chart was to establish a list of criteria against which we would measure our ideas, then weigh each criterion according to its relative importance.

In determining our criteria for this initial round of cuts, we decided to use a broad, non-exhaustive list of criteria. Because this analysis was of preliminary ideas that were not yet fully concrete, we excluded highly specific and technical criteria that would depend heavily on the exact specifications of the solution. For example, in this iteration of analysis, we decided to exclude the criteria of technical durability and testable feasibility, as more extensive research and a more finalized idea of our proposed solutions would be necessary to be able to assess these criteria reasonably. This research was something we knew would be a requisite step, but at this stage we were more concerned with narrowing down our ideas by general feasibility, so that our research and idea development efforts could be more focused. The criteria were brainstormed and discussed by the class, and ultimately decided upon by unanimous decision. Our criteria were as follows:

Criteria	Definition
Scope of impact	How large could the impact be if the solution successfully solves the sub-problem it intends to? In assessing the potential size of the impact, both breadth (i.e., reaching a large number of people or locations) and depth (i.e. deeply affecting many facets of one area) of impact were considered.
Measurability of impact	Is the solution measurable in a way that is related to the desired outcome?
Ease of implementation	How difficult would implementation of our solution be? Does it rely on additional infrastructure, technology, or a substantially changed workflow?
Evidence supporting evidence of approach	Do we have reason to believe that the proposed approach would solve the subproblem it attends to? Assessing ideas with this criterion relied heavily on our own research of the challenges faced by similar solutions and of indirect or uncontrollable factors influencing the subproblem.
Alignment with our problem statement	Does our idea align with both our problem statement and the client's core values and mission statement?
Prototype feasibility	How feasible is it to have a working and tested prototype in 2 months?
Cost for implementation	Depending on what the device or solution is, this criterion considered the cost of the physical device or deliverable, as well as the hurdles of obtaining medical approval, such as clinical trials, regulations, or training.

Table 2.1: Table showing the breakdown of each criterion and their definitions analyzed in the Initial Pugh chart.

After establishing our criteria, we gave each criterion a weight. These weights correspond to the relevant importance of each metric with regards to the idea's implementation, the values of GHI, and its ability to impact the pediatric cancer survival gap. For example, we weighted the scope of impact to be a 3 because we wanted our product to be effective in achieving a solution

to the problem. Affecting either a large population of people or a smaller population very effectively or both were significant issues to consider in producing a functional final product. By contrast, we weighted measurability of impact with a 2 because while it was significant, we considered a product that ultimately has a large scope of impact is more important than being able to measure that impact. However, since measuring impact is still very important in assessing the success of a product, we gave it a 2 and not a 1. Weights were valued from 1 to 3, and the final weightings can be seen in Table 2.2.

Once we had decided upon both the metrics we would evaluate our ideas against and their relative weights, we then scored each idea in each category on a scale of -1 to 1, -1 for negative, 0 for neutral, 1 for resoundingly positive. For example, under the ease of implementation category, the anthropometric data collection scored a 1 since the solution does not require additional technical infrastructure, does not demand a drastic change in nurse workflow, and actually streamlines what has the potential to be a time-consuming process. By contrast, a score of -1 was given in cases of shortcomings of the idea in a specific category. For example, computer-aided diagnosis scored a -1 in ease of implementation, since it requires technical infrastructure and change in doctor workflow. Its continued use might also be hindered by the prevalence of infectious disease, and a general practitioner's likelihood to de-prioritize cancer diagnoses. Combined, these hurdles led to the scoring of -1 in this category. Finally, a score of 0 was given if there was not a clear consensus on whether or not the idea would have a resounding positive impact or be lacking in a certain area. Going back to our ease of implementation example, the patient lodging app scored a 0 in this category due to its dependence on deployment location. Deploying in an area with high cell-phone penetration

would result in relatively few barriers to implementation, while deployment in countries with low cell-phone penetration would pose immense barriers for successful implementation.

Applying our criteria also brought holes in our research to our attention; consequently, some criteria were also initially scored as a 0 and flagged if they required further research. For example, the appointment reminder watch was given a 0 in scope of impact since we needed additional data to quantify the number of patients who missed appointments due solely to forgetting or misremembering their appointment date, rather than other competing factors. In this specific case, further research revealed case studies in which appointment reminders did increase patient adherence to appointments, but we were still uncertain if these affects would carry over to LMICs where other barriers such as transportation availability and transportation cost played a more instrumental role in treatment abandonment. For these reasons, the score of 0 was maintained after regrading with additional information.

With these criteria weights in place and our grading scale in mind, the Pugh chart was completed and organized as shown below.

Criteria	Nutrition Tracking Software	Anthropometric Multi-Tool	Appointment Tracking Wearable	Medical Record Wearable	Computer-Aided Leukemia Diagnosis	Digital Medical Record App	Communication Platform	Computer-Aided Referral	Cardiac Monitoring Device	OCR Medical Records	Lodging Matching App	Hospital Triage & Patient Feedback	Toy Vital Monitoring
Scope of impact	1	0	0	1	1	1	0	1	0	1	1	-1	-1
Measurability	0	0	1	0	0	0	-1	0	1	0	0	-1	-1
Ease of implementation and integration	0	1	1	0	0	1	0	1	0	-1	0	-1	-1
Evidence supporting efficacy of approach	1	1	0	1	0	1	0	0	1	1	1	-1	-1
Mission and Values	1	1	1	1	1	1	1	1	1	1	1	1	1
Feasibility	1	1	1	0	0	0	1	0	1	-1	1	1	-1
Cost	1	0	0	0	1	1	1	1	0	-1	1	-1	-1
Total Scores With Weights	12	10	9	7	7	6	6	4	2	-1	-2	-6	-12

Table 2.2: Pugh chart including categories of assessment as well as weighting assigned to each category and scores given to each project within each category of assessment. Initially we eliminated the negatively scored projects (indicated in red).

The completed Pugh chart was used to help inform our decision of which ideas to pursue moving forward. We assessed for clear drop-offs in score first, and then cut ideas which attained

an overall negative Pugh score. The ideas in green were chosen to advance to our next round of ideation, while the ideas in red were dropped.

Pugh Chart Exceptions and Further Elimination

Recognizing that the Pugh chart is an imperfect tool, the final score of each idea was not the only metric we used to eliminate our ideas.

Although it had a higher Pugh chart score than the Oncology Education Platform, the Digitizing Medical Record App was ultimately cut upon comparison to our other medical record idea, the Wearable Medical Record. Since the Wearable Medical Record scored higher, we decided to move forward with that idea and drop the digitization app in an effort to more clearly focus our efforts on what we believed to be the better medical records option.

Additionally, despite scoring well on the Pugh chart, the Cardiac Monitoring Device was eliminated after further research revealed that a majority of cardiac complications from cancer treatment present in long-term survivors, taking on the scale of 3-5 years after treatment to present (Aleman et al., 2014). This knowledge shifted the potential impact of the Cardiac Monitoring Device to a timeframe outside the scope of the 5-year survival rate that the GHI and our team sought to address.

Finally, the Diagnosis Web App was also cut from the list due to a team-wide concern with its high risk of failure. Although the solution could potentially catch the early symptoms of cancer that are often confused with other infectious disease symptoms and recommend the general physician refer the patient to an oncologist, the cases that were actually cancer, compared to those that were actually infectious disease, might be so few that the algorithm could potentially result in many false positives. These false positives might in turn have a significant

adverse effect, in which patients without cancer might have to go through a more complicated referral process that interferes with the timeliness of their care. Without a clear, effective way to address this concern, the idea was eliminated.

Summary of Eliminated Ideas in Reduction Phase

Malnutrition Assessment and Intake Monitoring software

This solution should not be too difficult to create because we could use software templates to code for our specific needs. We also already have access to nutritional questionnaires used in developed countries to assess malnutrition and we can therefore implement these proven assessments in our software. This software should be relatively easy to use because most of the usage would simply be inputting data and parameters that LMICs already currently measure. We plan to design this software to be straightforward and simple to implement. We decided to *keep* this solution.

Diagnosis Web App

The main potential challenges in this solution would be accuracy in diagnosis as well as size of database. The use of this solution would also likely require hospitals to have sufficient access to WiFi. We determined that while this app could be influential with a successful implementation, the ease of utilization and integration was very low. Additionally, the potential feasibility would be minimal because of the web access these hospitals would need as well as the engineering of an app of this caliber. Thus, we decided to *halt development* of the solution.

Matching Lodging App

This solution would only impact a small subset of the patient population that would require accommodation. Additionally, the solution would be challenging to roll out at scale due to local requirements and regulations. Recruiting potential hosts may be challenging, especially given the stigmas surrounding cancer in some of the locations of GHI partner hospitals. Given this, the fact that there is little evidence to show that this would increase survival outcomes, and that it does not relate specifically to our problem statement, we decided to *halt development* of the idea.

Machine Learning Diagnosis Platform

The solution received a relatively high score in the Pugh chart, as we recognized that the target population would be very large if this project were successful, especially as it would be extremely beneficial in medical centers that had limited diagnostic and pathological resources. We also understood that we would need access to large datasets to train the algorithms to accurately diagnose acute lymphoblastic leukemia, though there are significant applicable modern experimental machine learning techniques being implemented. After the intermediate idea analysis phase, we concluded this machine learning tool had enough supporting evidence to prove its feasibility and relevance to our problem statement. Thus, we decided to *keep* this solution.

Medical Records App

We had three potential projects in the area of medical record keeping, and we only wanted to continue developing one in order to present a wider variety of ideas to the client. As a

group, we decided that the strongest of these ideas was the Wearable Medical Record, so we ultimately did not move forward with the app system. We found that the wearable accessory might be easier to implement and maintain, as it does not rely on the use of Wi-Fi or integrated physical hard drives. The ideas were otherwise similar in scope, measurability, and alignment with our values, so this difference in potential ease of integration was the deciding factor that caused us to *halt development* of this solution.

Wearable Medical Record

Of our three potential projects in the field of medical record keeping, we ultimately kept this idea due to its relatively high potential ease of implementation and its uniqueness relative to existing technologies. This idea would not rely on any existing infrastructure beyond a means for scanning the information from the accessory, and this means would likely be as simple as a phone or computer, which are often already available. For example, Brazil, Colombia, and Mexico, countries that all have hospitals with partnerships with the GHI, have smartphone penetrations of 43.1%, 39.8%, and 45.6% respectively (Top Countries/Markets by Smartphone Penetration & Users). The relative independence from existing technical infrastructure of this idea coupled with the low cost of each device and intuitiveness of use makes the accessory one of our overall strongest proposals. Thus we decided to *keep* the solution.

Communication Platform

Assessing the feasibility of the solution, we noted that it scored a value of 5 on our chart. The low score, however, can be explained by the solution's lack of measurability, with it only addressing pediatric cancer indirectly. Additionally, it had mediocre performance in 3 other areas. In efficacy, our reasoning for this was that it may improve knowledge of the cancer but

may not be able to aid split second decisions necessary to save a life, because of time zone difficulties. In implementation, we hypothesized that we may encounter issues with infrastructure and resistance towards a digital culture change. Lastly, our impact scope was limited to a few individuals within a hospital, rather than many millions of patients that could be supported as in other solutions. This idea fell within the top six scores of the Pugh chart. Consequently, we decided to *keep* the solution for the client presentation.

Infant Scale and Length Board

Further research showed that this idea would not be exceedingly difficult to engineer a basic solution, since the individual measurement devices (digital scale and length board) are already commercially available. However, the engineering solution could lend itself to a more elegant design process that considers usability by nurses and broader application to other ages, among other concerns. The cost of this tool would be low, and the solution would be relatively simple to implement, depending only on a simple training on the use of the tool. Its success and impact, however, would be dependent on nurse adoption and the target population. Nurses would likely be open to try this tool if it would streamline data collection process, but we did want to receive feedback from GHI assessing their opinions on nurse usage. We concluded that we should *keep* this idea.

OCR Record Digitization

This solution could be standardized and widespread across hospitals with existing IT infrastructure. However, those that do not have the infrastructure are likely those with the greatest need for this solution. Tracking the impact of the solution is likely to be easy, given that all data should be in theory digitized. However, significant challenges still remain with OCR,

including accuracy and non-English language recognition, impairing the solution's accuracy and usability **(**Holley, 2009**)**. Due to the combined burden of additional infrastructure required and the existing limitations for OCR in recognizing handwritten characters, we decided to *halt development* of this idea.

Appointment Reminder & Tracker

This concept was within our engineering capabilities and aligned very well with our problem statement by focusing on the communication between doctors and patients. In addition, the product had a measurable impact with the ability to track patient attendance and a relatively easy implementation process would prevent it from adding significant burden to the existing workflow. However, the scope of the impact and evidence to support the efficacy of the solution was a concern as the effects of appointment reminders in LMICs remained unclear. For this reason, we re-emphasized the focus of the product towards the patient tracking aspect rather than solely as an appointment reminder, to focus on hospital communication and stronger data collection. With this new focus, we decided to *keep* this idea as a potential solution.

Matching Patients to Survivors app

Although this idea would target the stigma surrounding cancer in LMICs, it would face many challenges in acquiring volunteer survivors, technological access of the users, and measurability of the impact. In addition, this solution did not relate as well to our problem statement, focusing on education and culture rather than communication and data. For these reasons, we decided to *halt development* of this idea.

Data Collection Toy

The feasibility of the was initially determined to be plausible and was included in our intermediate phase of idea analysis. After our Initial Ideation phase with 10-15 ideas, it was concluded that the practicality and usability of a toy to measure vitals such as temperature and heart rate would only lead to further burden on nurses as accurate readings, cooperation from the child, and the reusability of the toy are difficult to achieve. There was also not enough evidence to support the feasibility of the idea or the potential implementation into treatment procedure. Thus, we decided to *halt development* of this solution.

Visual Inpatient Monitor

Although necessary for robust prolonged patient care, patient feedback is not vital to immediate or urgent care. More quantifiable data, such as vital measurements, are favored over potentially unreliable patient feedback. Furthermore, there are already instances of "nurse call" buttons being abused in HICs, meaning the idea could actually inadvertently increase or obfuscate the nurse's workflow. Although we believed the problem of nurse workflow is an important one, we decided that this idea for a solution would not clearly maximize the efficiency of care, and thus we elected to *halt development* of this idea.

Chemotherapy Drug Database

The solution is based upon assumptions about the administration of chemotherapy drugs that are not valid, as it suggests that specialists and doctors are not informed on how to administer drugs properly and in a standardized way. The solution we posed would not be outstandingly accurate, or offer anything new for specialists over what already exists in

electronic or paper form. We also foresaw difficulty incorporating this solution into an existing infrastructure of each hospital and minimally intruding into the workflow for the specialists, pharmacists, and oncologists. Thus, we decided to *halt development* of this solution.

Cardiac Vital Monitoring Device

Further research showed that it was unclear whether this solution would have the intended impact given that cardiac complications reveal themselves late into the treatment process. This device, while having a discernable and measurable impact, may be challenging to implement into nurse workflows, and challenging to physically engineer. For the solution to be successful, each inpatient would be required to own one, increasing the cost-per-patient for the hospitals. For this reason, we decided to *halt the development* of this solution.

2.4 | Top 6 Ideas Pitched to Client

The initial Pugh chart criteria-consideration not only helped eliminate ideas down to six candidates, but also helped us identify how we should improve these remaining six before presenting them to our client. Consequently, throughout the reduction phase, several of these remaining six ideas evolved significantly. The final descriptions of the remaining six, as they were pitched to our client, and their evolution since their inception is described below.

Nutrition Tracking Software

The Nutrition Tracking Software intended to create a method to record and monitor patient nutrition data points. Malnutrition has clear implications for chemotherapeutic drug

clearance, and long-term growth development. Furthermore, research has shown clear correlation

between nutritional status and survival, relapse, and treatment of abandonment rates (Sala et al,

2012). Most developing countries rely primarily on weight, height, and age-based metrics to

diagnose malnutrition, disregarding pertinent metrics such as mid-upper arm circumference

(Bauer, Jürgens, & Frühwald, 2011). Research suggest that there are 5 categories of nutritional

data that correlate with outcome: anthropometric data, biochemical data, medical history,

physical observations and dietary intake (Allessandra, Emanuela, Federico, Ana, & Maselli,

2011). Therefore, the proposed solution is a software that helps record and monitor data points

through a set of key features. Health professionals would insert into the software nutritional

additives they have already available, and different professionals would be able to log-in and

insert live relevant data. The software would output a simple nutritional diagnosis: healthy, risk

of malnourishment, or malnourished. It would then provide a simple step to take to intervene that

incorporates the information fed into the system. Ultimately, this solution seeks to improve

information gathering, and create a summary of patient behavior over time.

Anthropometric Multi-Use Tool

The Anthropometric Multi-Use Tool is a device that combines an infant scale and an

infant length board together in the most efficient way possible. Current methods for monitoring

height and weight in infant cancer patients are unstandardized and inaccurate due to a lack of

human, material, and financial resources. Examples of currently used methods of data collection

include analog bathroom scales that can have up to a 1% error, which can make a significant

difference when weighing the child with a parent, as is done in some LMICs (Nursing Team,

personal communication). Malnutrition is a huge obstacle to cancer treatment and survival,

especially when chemotherapy is involved because of the dependency on accurate body composition analyses. This multi-tool would allow nurses to safely and quickly take more accurate height and weight data, allowing for earlier malnutrition diagnosis and, in turn, opportunities for interventions. As there is no need to re-engineer a scale, we primarily focused on how to modify existing hanging scale technologies to best suit the needs of medical care professionals in LMICs. Though the population of infantile cancer patients is low relative to the number diagnosed–for example, acute lymphoblastic leukemia, which accounts for approximately 25% of all pediatric cases, has an infantile population of 2-5%–the years-of-life saved is undoubtedly considerably beneficial (Reaman et al, 1999).

Wearable Medical Record

The Wearable Medical Record aims to combat the difficulty of transferring and communicating a patient's physical medical data, both within one healthcare facility as well as between local health clinics and larger hospitals. This idea consisted of a combination of the Medical Records App and Wearable Medical Record. We proposed making a personal, portable medical record device that patients would not easily lose or forget. This device would consist of a piece of wearable technology that carries medical data to be accessed via USB, RFID, QR code, or a similar technology. This ideation evolved based on feedback received about storage concerns of an entire medical record and relevance to pediatric oncology. To incorporate this feedback, we decided that the data stored on this wearable should include a file containing only the most relevant information to pediatric oncology care such as diagnosis, symptoms, treatment plan, medications, and allergies. This data file would accessible by a physician or the patients themselves through a scanning process, which would then save the data file in a searchable

database on the hospital's computer or phone. Having immediate access to previous medical data has proven benefits in HICs for faster retrieval of records, improved care plans, more efficient hospital workflow, and enhancement of coordination within the care team (Gyamfi et al., 2017). We imagine that these advantages can be transferred to lower-income countries through this device.

Appointment Tracker and Reminder

The Appointment Tracker and Reminder combines a wearable device with a computer database and interface. In most LMICs, treatment abandonment is a clear obstacle to increasing survival rates in pediatric cancer. We approached this problem similarly to the approach used by Dana Farber Cancer Institute in El Salvador: a hospital maintained a pediatric cancer patient database and used social workers to follow up with patients who missed appointments, significantly reducing treatment abandonment (Piñeros et al, 2018). We envision an interface that tracks the appointment dates of patients and flags patients who have missed their appointments. Each patient will have a wearable wristband that saves their appointment date and hospital identification number, and when they arrive at the hospital for their appointment, medical staff updates the wristband with their next appointment date, and the database is also updated to reflect the new appointment. If the patient misses the appointment, the database will detect that the date has elapsed with no update, flag the patient, and cumulatively sort patients who have missed the most appointments, thus identifying patients at risk for total treatment abandonment. Healthcare providers will then be alerted of these patients, through an automated alert, so that they can follow up with them.

Additionally, we decided to use the concept of the wristband to engineer a solution on the patients' side as well, by proposing to integrate a pre-appointment reminder into the wristband. This would be done by making the wristband a functional watch and incorporating an alarm system that corresponds with the date of appointment. Implementation of pre-appointment reminder systems (usually SMS-based) in both LMICs and HICs has been shown to result in significant increases in treatment adherence (Domek et al, 2016; Schwebel and Larimer, 2018). We theorized that a reminder system would need to be independent of WiFi or network connection since many rural areas in LMICs do not have convenient WiFi or cell access (See prototype in Appendix Figure ii.4).

Machine Learning Diagnosis Tool

The Machine Learning Diagnosis Tool is aimed specifically at diagnosing Leukemia cases in LMICs. Leukemia makes up approximately 25% of pediatric cancer cases and is one of the most challenging to diagnose ("What are the Key Statistics for Childhood Leukemia," 2016). Additionally, the disease progresses quite quickly, so early identification and treatment are vital to survival of the patient. The Machine Learning Diagnosis Tool would seek to supplement the skills of the pathologist by providing an accurate and timely diagnosis of blood and bone marrow smears (Wahhab, 2015). To obtain this diagnosis, a technician would take images of the smears, which are likely already being obtained for a report, and send them via online submission to the machine learning algorithm. The algorithm would be pre-trained on a large set of blood and bone marrow smears for pediatric leukemia cases, and using that "knowledge," could identify cancerous cases. If the accuracy rate of the algorithm was not high enough to guarantee accurate diagnosis, the tool could be used as a winnowing or prioritizing tool for pathologists, and also

serve as a reason to recheck slides if the diagnosis the pathologist reached was different than the diagnosis reached by the tool. Ultimately, the Machine Learning Diagnosis Tool would serve to reduce the time leading to diagnosis of cancerous cases and increase the accuracy of the delivered diagnoses (Wahhab, 2015) (See prototype in Appendix Figures ii5, 6 & 7).

Pediatric Oncology Communication Platform

The Pediatric Oncology Communication Platform is a web-based or mobile app that matches pediatric oncology health professionals in LMICs to those in HICs, facilitating the global transfer of pediatric oncology knowledge. Relatively few healthcare professionals are involved in pediatric cancer care in many LMICs; for example, there are only five practicing pediatric oncologists in Malawi, Uganda, and Botswana combined ("First Comprehensive Pediatric Hematology/Oncology Initiative Launched in Africa"). Consequently, pediatric oncologists and other professionals who treat pediatric cancer patients are isolated, and such a platform could correct this by integrating them with the pediatric oncology network present in HICs. Doing so would be an extension of twinning programs, which establish a long-term partnership between organizations in LMICs and HICs with the mutual goal of increasing pediatric cancer survival rates. The targeted users are medical teams within pediatric oncology wards, whether it's oncologists, pathologists, nurses, and so on. Users would need to complete a detailed form regarding the information they would like to provide or receive, and a back-end algorithm would match LMIC teams with HIC teams based on alignment of form responses. The subsequent communication could occur through the website or externally, depending on user preference. This system however, would need easy access to WiFi (See prototypes in Appendix Figures ii.8, 9 & 10).

2.5 | Client Feedback

Nutrition Tracking Software

First, we presented the Nutrition Tracking Software. The client's main feedback included: 1) There are two main markers of malnutrition, which include arm circumference and head circumference, and both should be incorporated; 2) Constipation, vomiting and zinc should be removed as markers, while MUAC and iron levels should be emphasized. The markers included could be obtained through a focus on published metrics; 3)This solution would present a way to track nutrition over time and provide a longitudinal data set for each patient; 4) It is important to consider whether the program would require internet access, and whether it could be incorporated into existing hospital records.

Anthropometric Multi-use Tool

Our second idea was the Anthropometric Multi-use Tool. The main feedback we received from the client included: 1) The emphasis of size (length, width and circumference) as opposed to weight as a method for assessing malnutrition - the client noted that dosage for chemotherapy is assessed in terms of volume for children over the weight of 8 kg, or approximately 10 months; 2) The size of our prototype was too small for it to be used effectively an throughout a wide range of pediatric cancer cases. The client argued that the scope of our target audience would be greatly reduced as infants only account for 10% of pediatric cancer cases (Albanti and Lehmann, personal communication); 3) This solution could possibly be combined with our nutrition tool idea since longitudinal measurements also provide more useful information that a single reading.

Wearable Medical Record

Our third idea was the Wearable Medical Record. The main feedback we received from the client included: 1) There are difficulties in accessing and reading handwritten notes and doctors are not able to give treatment until they can access a patient's pathology, leading to significant delay; 2) This solution would even be impactful in the US, where we could start a possible pilot program for the technology; 3) Patients in Rwanda and Myanmar have a patient passport with medical identification that they keep securely, as a similar model for our intended solution; 4) There would be consensus among doctors on the data necessary to include in the personal medical record including allergies and chemotherapy plans; 5) The biggest pain point will be people entering the data to update the personal medical record. Since Paul Farmer was able to leverage the power of community workers to deliver HIV medication, they suggested we might use community workers to enter medical data for the personal record; 6) Questions to consider: Is data on medical records transported by the patient? Who would upload the new data? How much data capacity would the Wearable have? Throughout the presentation, the client was verbally enthusiastic about the idea as having a high potential impact.

Appointment Tracker and Reminder

Next, we discussed the Appointment Tracker and Reminder. The client's main feedback included the following: 1) The main reasons for treatment abandonment in LMICs are prohibitive transportation costs, prohibitive care costs, or a perception of treatment inefficacy, so a device that reminds patients when they have appointments will not significantly reduce treatment abandonment rates; 2) An appointment reminder system may be more effective in

reducing treatment abandonment rates in middle- or high-income countries rather than low-income countries, where more significant barriers to treatment exist; 3) In general, data about appointment attendance are important and can be used to efficiently allocate funding and resources; 4) The project's scope is similar to a program in El Salvador which successfully reduced treatment abandonment rates. In the mentioned program, a healthcare facility tracked patient appointments and social workers followed-up with patients to ensure that they got to their appointments.

Machine Learning Diagnosis Platform

Our next idea was the Machine Learning Diagnosis Platform. The client's main feedback included: 1) If this idea was pursued, it would be absolutely necessary to predict cases with the highest accuracy. This accuracy rating should range between 99-100%; 2) Technicians and pathologists would still need to perform the same amount of work, so the client questioned the usefulness of this tool, but did acknowledge that it could be used as a prioritization tool. 3) This idea would be most useful in a central pathology lab within a country, or places with higher patient concentration such as Niger, Sierra Leone, or Ethiopia ("Countries Compared by Health > Physicians", n.d.). The client recommended that if pursued, the idea should be used in a MIC as opposed to a HIC. In general, the client expressed concern with the implementation of Machine Learning due to potential cultural rejection or distrust of the technology's accuracy among the medical field.

Pediatric Oncology Communication Platform

Last, we presented the Pediatric Oncology Communication Platform. The client's main feedback included: 1) That the idea itself, and the problem it was seeking to address, were unclear; 2) It might be difficult to identify, from a design standpoint, who would be actually messaging through the platform, and someone from the hospital would need to be responsible for the communication; 3) Participants would need to be vetted before HIC hospitals would be comfortable providing advice, and there would need to be a longitudinal relationship where parties knew each other, and each other's skills and character; 4) When the focus shifted to treating the platform as an extension of the existing GHI twinning program, the client became more enthusiastic, but retained reservations regarding the potential high density of requests.

2.6 | Feasibility Analysis of Top 6 Ideas

After receiving feedback from our client on our six ideas, we analyzed the feasibility of our ideas with more rigorous criteria a to narrow down to one preliminary idea.

Pugh Chart Analysis

This feasibility analysis took the form of a Pugh chart in order to give a quantitative, and thus more directly comparable, measure of the qualitative features of our solutions. Importantly at this point, our ideas were defined enough to apply more specific, technical criteria in our assessment. Consequently, the criteria we formulated are different from the first Pugh chart analysis, and focus more on concrete details such as "technical sustainability." Additionally, broader criteria from the previous Pugh chart analysis were separated into more specific categories; for example, "cost for implementation" from the previous analysis was broken down into two separate criteria: "affordability" for the end-user and "R&D costs" on the team's and GHI's end. We made these categories distinct because we decided to prioritize the affordability and entrance cost for LMICs over our own costs for developing it, so weighted these categories differently. We formulated the following 12 criteria presented in the table below to evaluate each solution:

Criterion	Description	Explanation of Weight
Aligns with client values	How well does this assist the client's mission? Does it meet the client's values and expectations? Take into account client feedback.	**3**: We weighed this relatively highly because GHI is our client, and we are proposing a solution to address where they see the most need.

Affordability	Will end-users be able to afford this solution? How much would it cost to introduce and integrate this technology?	**2:** We weighed this a 2 because we saw that the solution can be adapted or downscaled to serve specific hospitals in partnership with GHI. As a group, we have greater control to design a solution that is more cost-effective or one that integrates within a specific community.
Ease of adoption and integration	How easily/conveniently will this integrate into the current workflow? How much training/explanation is needed for the doctor and patient for them to accept this solution? Will the user have incentive to adopt this?	**3:** We weighed this relatively highly because ease of adoption/integration is a large determining factor of whether it could be directly helpful to our end-users.
Enhances communication, collection, and interpretation of data	How effectively does it relate to the problem statement?	**1:** We weighed this lowly because we realized we can adapt our problem statement accordingly - if we decided on a solution that instead addresses another aspect of pediatric oncology more directly (e.g. reminding patients of their appointments).
Technical sustainability and maintainability	How long can it last without maintenance and will constant maintenance be necessary?	**3**: We weighed this relatively highly because technical sustainability and maintainability (as opposed to cost) is something more within our locus of control as engineers and should therefore be an important consideration.
Scope of Impact	To what extent will the solution impact pediatric cancer patients? The size of impact can be quantified as breadth*depth, where breadth is the number of affected individuals and depth is the extent to which the solution impacts each individual.	**3**: We weighed this relatively highly because we intend to prioritize the solution we see will affect change in the communities the solution is implemented in, whether it be number of reachable individuals affected, or the temporal extent to which they are affected.
Reliance on Technical Infrastructure	What sort of infrastructure is required for our solution to be integrated fully? Is our solution reliant on computers, smartphones, access to internet, etc.?	**2:** We weighed this relatively lowly because the specific countries and hospitals we implement the solution in can be adjusted if needed.

Directness to desired outcome	In the ideal case scenario where our solution works as intended and is adopted successfully, will the solution result in more accurate measurements or data	**3:** We weighed this relatively highly because we are prioritizing the solutions that would have clear impact on improving situations (e.g. better anthropometric and nutritional data collection).
Engineering feasibility	Can our engineering group create a working prototype before the conclusion of the semester? Do we have the necessary skills or expertise to produce a working product?	**4:** We weighed this highly because we believe that in order to deliver a working solution, it needs to be within the scope of our skills and contain enough modularity for each of us to get involved productively.
Testable feasibility	Is it possible to test the effectiveness of this solution? Can we access or generate preliminary data to experiment with? What is the measurable impact?	**3:** We weighed this relatively highly because we wanted to ensure that we can verify the effectiveness of our design once created.
R&D Costs for implementation	What are the R&D costs for GHI to get our solution deployed on site?	**1:** We weighed this lowly because the cost could be incurred by an actor (e.g. GHI's partners) who can more likely afford the solution.
Market Landscape	Are there existing solutions or startups that are addressing the same problem in the context of LMICs? Are we infringing on any patents?	**1:** We weighed this lowly because although we wanted to be aware of similar solutions already being used in certain GHI countries, we decided that we could likely create a useful, unique solution despite existing competition.

Table 2.3: Breakdown of criteria and weights used to evaluate our six ideas in the Feasibility Analysis.

Each metric was then given a weight between 1 and 4 based on the relative importance of each to the production of our solution. We gave "engineering feasibility" the highest weight because we wanted to create something that worked by the end of the semester to deploy and make an impact with our client GHI. Other categories we weighed higher include "ease of adoption or integration," "technical sustainability," and "relative scope of impact" because we valued how well the solution would be integrated into the LMIC setting and how many people can be positively affected in the long term. A more specific description of our evaluations can be

found in Table 2.3. Another criterion we considered was "risk of failure," or in other words, the likelihood that our product will work as intended, become adopted, and ultimately increase cancer survival rates. However, we decided not to include this as a single category to consider because there were many important components that could be accounted for more specifically by the other criteria we consider (e.g. ease of adoption, technical sustainability, market landscape).

For each criterion, the idea was rated -1, 0, or 1 relative to others. A score of -1 was given if the solution did not favorably meet the criterion or was less clearly fit compared to other solutions, 0 if there was insufficient evidence or were comparable to others, and 1 if we could resoundingly advocate for the solution with adequate evidence.

Criteria	Weight	Communication Platform	Nutrition Tracker	Appointment Tracker and Reminder	ML Diagnosis Platform	Anthropometric Multi-Use Tool	Wearable Medical Record
Aligns with client values	3	0	1	0	-1	0	1
Affordability	2	1	1	0	1	0	0
Ease of user adoption and integration	3	-1	-1	0	-1	1	0
Enhances data communication, collection, and interpretation	1	0	1	0	1	1	1
Technical sustainability and maintenance	3	1	0	0	1	0	0
Scope of Impact	3	-1	1	0	1	-1	1
Reliance on Technical Infrastructure	2	-1	-1	0	-1	1	0
Directness to Desired Outcome	3	-1	0	-1	1	0	1
Engineering Feasibility	4	-1	-1	0	-1	1	0
Testable Feasibility	3	-1	0	1	1	1	1
R&D Costs	1	0	0	1	-1	1	0
Market Landscape	1	-1	1	0	0	-1	0
Total Scores With Weights		-14	1	1	2	10	13

Ideas Eliminated through Feasibility Analysis

When analyzing the results from our Pugh chart, we eliminated ideas based on large scoring drop offs. We saw three main jumps: first, that the communication platform received the lowest score at -14; second, that the nutrition tracker, appointment reminder, and ML diagnosis platform all received similar scores around 1-2; third, that scoring then jumped up to 10-12 with the anthropometric multi-use tool and the wearable medical record. Only the top-tier scores were kept -- that is, the anthropometric multi-use tool and wearable medical record. The specific reasons for eliminating the other ideas are detailed below.

Communication Platform

This idea was unanimously terminated due to its drastically low score (-14). The main problems plaguing this idea include a lack of engineering and testable feasibility, indirectness to our desired outcome of increased survivability rates, and reservations from the GHI on the potential clinical benefits of this platform. Ease of adoption also posed a large challenge, as the communication platform would not easily fit into the already overworked schedules of doctors in LMICs, and there would be little tangible incentive for specialists within HICs to incorporate it into their workflow.

Nutrition Tracking Software

The strongest parts of this idea were that it aligns well with our client's values and our problem statement, as well as feasible and affordable to engineer and use. However, ease of adoption may be a problem, as consistent data entries would be necessary and would likely add

to the already overloaded workflow of medical staff, and there may not be enough technological infrastructure to support a hospital-wide software. Additionally, since the solution intends for nurses to take more patient nutritional measurements than is currently the norm in many LMICs, there were concerns that this solution may require an educational component stressing the importance of these measurements in order for it to be readily adopted. Finally, since the idea proposed to simultaneously store patient nutrition data and recommend steps to improve nutrition, the engineering feasibility of its proposed recommendation algorithm would be a large challenge, due to its need to be clinically accurate enough that all recommended treatments or nutritional supplements -- and all their potential interactions -- would be guaranteed safe for a patient. These challenges contributed to the idea's low Pugh Chart score (1) and were the basis for its elimination.

Appointment Reminder & Tracking Software

Initially the appointment reminder tracking software seemed promising, with its simplistic technological design scoring higher than both the machine learning and communication platform in terms of engineering feasibility and ease of adoption. However, we found that the product was less direct about helping improve pediatric cancer survival rates, as a reminder of an appointment would not be able to surmount the many of the primary obstacles to appointment adherence, such as financial strain or lack of transportation. Additionally, considering that some LMICs have initiated effective social worker intervention practices for appointment reminding and tracking, this product would not necessarily be making a novel impact (Moulik et al, 2016). For these reasons, the idea received a relatively low Pugh Chart score (1) and was eliminated.

ML Diagnosis Platform

This idea received a low Pugh Chart score (1) primarily due to the combination of the score of -1 in the "ease of adoption and integration" category and the score of -1 in the "aligns with client values and expectations." Based on the feedback from our client, it became apparent that doctors and communities in LMICs may be unlikely to embrace the significant cultural shift of having a machine performing a task that typically requires a high level of medical training and expertise, or trust the diagnoses from the potentially foreign technology. Thus, though it performed well in other criteria, this fatal shortcoming ultimately eliminated the idea.

Elimination of Final Two Ideas and New Idea Generation

Following this Pugh Chart feasibility analysis, only the anthropometric multi-use tool and the wearable medical record were kept for further consideration. Ultimately, however, through rigorous research and analysis into both ideas, the team decided not to move forward with either the wearable medical record or the multi-use tool, and instead decided to pursue an additional iteration of ideation.

Justification for Eliminating the Anthropometric Multi-Use Tool

Although this idea scored well on the Pugh chart, a critical concern with this idea is it's relatively small scope of impact. Since less than 10% of pediatric cancer is in infants, this instrument would only be beneficial to a very small proportion of cancer patients (Albanti and Lehmann, personal communication). Indeed, as we learned from the nutrition presentation by

members of the GHI, pediatric cancer patients also only have anthropometric data points taken every couple months at most, after which the infant patients may outgrow the tool. Thus, this instrument would only provide a few data points for this small subsection of the population. Ultimately, this -1 in scope of impact was also determined to be a "kill criterion" that eliminated the idea from consideration.

Justification for Eliminating the Electronic Medical Record

In order to more fully address the feasibility of the wearable medical record idea, we needed to pin down the particular technicals specifications we would choose in order to ensure that a version of the idea still could meet our design goals, namely improving hospital/health-center workflow while also being user friendly and intuitive. The technologies we considered were RFID/NFC, QR codes, and solid-state storage (USB drive, SD card). However, after analyzing each of the design specification options (as seen in Table 2.5), we realized that none would be able to meet enough of the design goals to become a valuable product with a scope that matched our goal. Creating a device specifically tailored to an individual hospital and its record system would have been possible, but would vastly reduce the scope of the project to an undesirably low level. On the other hand, creating a device that would have relatively wide-use and integrate well at both the health centers and hospital level of care would be almost impossible, given the technical limitations as well as infrastructure limitations of low-income settings. Most specifically, the device would require internet connection (to integrate with a web-based EMR system), or the device would substantially increase workflow, as clinicians would have to constantly update data. Either of these qualifiers did not match up

with our original design goals, as they vastly limited the scope of the impact or the success of the

initial idea.

	RFID + Web	RFID Volatile	RFID - stable non-editable	QR - local	QR - web	USB Stick/SD - custom file	USB Stick/SD card - standard file
Data Storage	100 bytes + unlimited	100 bytes w/reader	64 KB	~1000 characters	Unlimited	Unlimited	Unlimited
Access (Phone)	Yes	Yes	Yes	Yes	Yes	Need adapter	Need adapter
Access (Computer)	Yes w/reader	Yes w/reader	Yes w/reader	Yes w/camera	Yes w/camera	Yes	Yes
Updatable	Yes	Yes	No	Y/N (need to reprint code)	Yes	Yes	Yes
Need software package	Yes	Yes	Yes	Yes	Yes	Yes	No
Replacement to current records system?	Yes	No	No	No	Yes	No	No
Adds additional workflow?	No	**Yes**	**Yes**	**Yes**	No	**Yes**	**Yes**
Needs Internet/LAN	**Yes**	No	No	No	**Yes**	No	No

Ideate Iteration: Vital Monitoring for Sepsis Detection

After eliminating the wearable medical record, we decided that we needed to re-iterate through the ideate process and re-examine ideas previously discarded, to see if we had prematurely closed out any viable, more promising solutions. Consequently, we focused in on our initial 15 ideas and revisited the reasoning for eliminating each one. In particular, we re-visited the cardiac vital monitoring idea.

Originally, the cardiac monitor idea had been cut because the cardiac complications it tried to track only manifested long-term, past the window of pediatric care and into adulthood (see Section 2.4). However, when we returned to this idea and researched the importance of cardiac vitals in the literature in more depth, we found an immense amount of data on the importance vital signs for a different application: diagnosing pediatric cancer patients with sepsis. Thus, we decided to re-invent the cardiac monitor in a new form: as the in-patient vital monitoring system, targeted towards the early recognition of sepsis.

Final Idea Motivation: The Problem of Sepsis

Infection and sepsis are one of the largest causes of mortality in pediatric cancer patients. In one case study conducted in Brazil, of the 220 pediatric cancer patients admitted to the oncology ward, the prevalence of infection and sepsis was 73% and mortality due to sepsis was 40% (Costa et al. 2017).
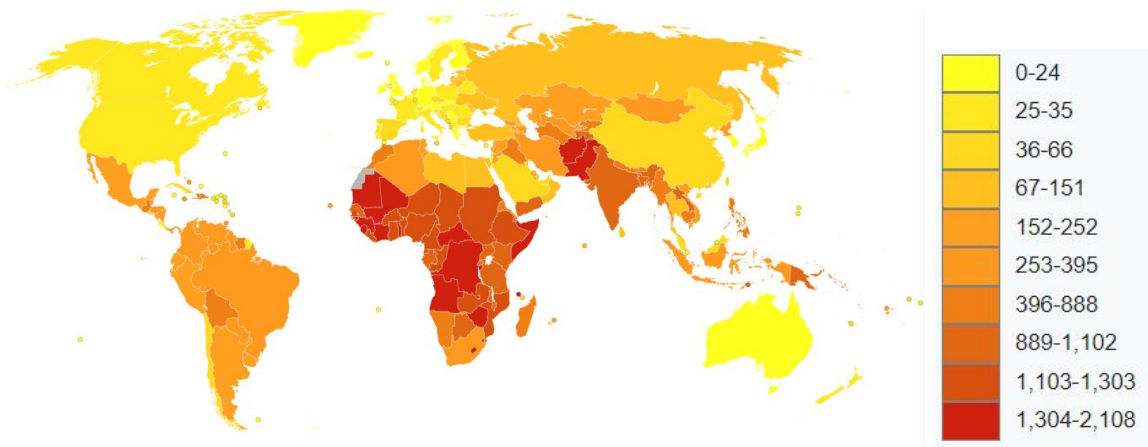
Figure 2.3: Pediatric Sepsis Deaths per Million Persons: Global map showing the number of sepsis deaths in each country per million persons (Zanotti-Cavazzoni, 2007)

Sepsis is caused by weak immune response to infection. In a healthy individual, the immune system serves to kill harmful bacteria and microbes that invade the body. During sepsis, however, the immune system is overworked and unable to combat the infection, and the byproducts that it produces are released into the blood, causing inflammation throughout the body (Recknagel et al. 2012). Pediatric cancer patients are particularly vulnerable to infection and sepsis because, as a result of harsh treatments such as chemotherapy, the child's healthy cells are destroyed and the immune system is compromised, leaving the body susceptible to bacteria. Indeed, sepsis is even more of a concern in LMICs than it is in HICs, as care providers and patients often have less access to sanitary tools and devices necessary to prevent infection.

Sepsis presents in a patient as a whole-body response to the infection. By consequence, it manifests through several recognizable physical symptoms. Specifically, sepsis is clinically diagnosed based on evidence of systemic inflammatory response syndrome (SIRS), defined by abnormal body temperature, elevated heart rate, and high respiratory rate (Kenzaka, et al., 2012).

Although an abnormal temperature is the most typical response to sepsis, if any two of these three vital signs are abnormal, then a patient is suspected to to have sepsis (Kenzaka, et al., 2012).

For detecting and treating sepsis, time is of the essence. The graphic below illustrates the progression of bacterial growth and sepsis for a neutropenic[4] patient, whose white blood cells are low (Stefan et al., 2014). As illustrated, about 6 hours after exposure to bacteria, early signs of sepsis may begin with a fever. About 8 to 10 hours after initial exposure, sepsis itself can begin to present, in which blood flow to vital organs becomes restricted. If left unaddressed, just 2 hours later, the patient can experience severe sepsis, in which organs or organ systems begin to fail, resulting in complications such as cardiovascular system dysfunction or acute respiratory distress syndrome (Fiser et al. 2005). Finally, two more hours later, and just 12 hours after exposure to infection, the patient might enter septic shock and, ultimately, death. Consequently, the first 6 hours of clinical symptoms are deemed the 'critical period' in the treatment of sepsis; indeed, each hour of delay is estimated to decrease the patient's likelihood of survival by 8 percent (Kumar et al., 2006). Thus, to treat sepsis effectively, early intervention is crucial: in one large study of septic shock, 80% of patients survived if treatment was administered within 1 hour (Kumar et al., 2006).

---

[4] A neutropenic patient is one that has low white blood cell count. This condition is typically due to chemotherapy treatment ("What is Neutropenia and How Is It Treated?" 2018).

Figure 2.4: Progression of bacterial growth and infection in a septic patient. Bacterial growth when uninhibited follows an exponential curve; by 6 hours, clinical symptoms begin to manifest in fever, which is then followed by sepsis, severe sepsis, and septic shock on a close to two-hour interval (Stefan et al. 2014). The left is a visualization of the rapid timeline of sepsis development.

However, early identification and intervention in sepsis is made particularly difficult in LMICs due to a low nurse-to-patient ratio and inadequate amount of vital monitoring for pediatric cancer patients. Although the optimal nurse-to-patient ratio in HICs is 1:4, in most LMIC hospitals, these ratios are as high as 1:20 (Friedrich et al, 2014). As a result, nurses are unable to consistently take patient vitals, and often only record vitals 2-3 times a day (Weenk et. al. 2017). This inhibits early detection of sepsis, as, in the worst case scenario, a patient could have already developed septic shock in this 12-hour window, in which symptoms could go unnoticed. Indeed, pediatric patients in particular often experience more rapid vital sign deterioration, which further underscores the need for more frequent vital monitoring in LMICs (Cootes N. 2010). However, as it stands, vital monitoring products available in HICs are expensive and require complex technology and infrastructure often not available in low-income settings, making this technology inaccessible to the pediatric patients in LMICs that need it.

2.7 | Final Project Idea: Vital Monitoring System for Early Recognition of Sepsis

To address this need, we propose to develop a simple, cheap in-patient vital monitoring system that will help understaffed nurses in LMICs with early detection of sepsis and infection for in-patients. Our product will allow nurses to monitor patient vitals that are critical to early detection of sepsis or infection -- primarily heart rate, respiratory rate, and temperature -- and alert nurses of any abnormalities in their patients' vitals. By collecting these three vital signs, nurses will be provided with information tailored to one of the leading causes of death for pediatric cancer patients. In particular, by monitoring respiratory rate, our product will monitor respiratory rate adds important clinical value, because respiratory rate is a highly informative indicator of physiological state and is the most sensitive vital sign marker of clinical deterioration, particularly amongst children (Schein et al 1990, Goldhill et al 1999, Ridley 2005, Cretikos et al 2008,& Braun 1990). However, in both HICs and LMICs, respiratory rate is often considered the 'forgotten vital,' as nurses often fail to record respiratory rate due to poor knowledge, subjective assessments about whether to assess respirations, and a lack of time (Elliott, 2016).

The product will be made up of two parts: an on-patient device that measures the vital signals, and a central 'hub' monitor. The on-patient device will measure heart rate, respiratory rate, and temperature and transmit that data to the 'hub', where nurses will be able to visualize which patients (if any) are in alarm state and need attention. Additionally, this hub will allow nurses to pull up longitudinal data for a patient of concern, to provide a more comprehensive picture of any peculiar trends.

To determine thresholds for normal or abnormal vitals, this device will draw from PEWS, or Pediatric Early Warning Scores, which is a standard system used in many HIC hospitals, such as Boston Children's, to assess the physiological stability of a child based on their vital signs. Specifically, the trigger-based PEWS system is a threshold-based scoring method that assesses a child relative to a normal or abnormal threshold (Duncan, Hutchison & Parshuram, 2006).

In essence, our device, which we have called PediaTrack, is catered towards the collection of vitals that are significant to the pediatric population the GHI works with, automizes nurse tasks while ensuring that immediate intervention is accessible, and collects longitudinal data for more informed care.

2.8 | Feasibility Analysis of PediaTrack, Vital Monitoring System

At the end of our Ideate phase, PediaTrack stood out as the idea with the strongest

potential to pursue. The following section provides an in-depth analysis of the feasibility of

PediaTrack and our reasons for selecting it, based on our criteria.

Alignment with Client Values

By allowing for the collection of vital signs pertinent to pediatric oncology patients, the

device can increase the ability to collect, record, and interpret relevant patient data without

overburdening the overworked nurses in LMICs. These benefits directly address GHI's mission

to "improve the lives and enhance the care of children with cancer and blood disorders

worldwide" (The Global Health Initiative). This solution also aligns with the client's values of

enhancing communication and capacity building at hospitals they serve in LMICs, as the vital

monitoring system would help build better programs to record, store, and communicate patient

data. For these reasons, the wearable medical record scored a 1 in this category.

Affordability

The main components of this technology that would require an additional cost would be

the physical on-patient device and the hub component. While the on-patient device will have to

be provided for each patient, the cost per unit will be low when mass produced. The components

of the medical device are not expected to exceed **$300,** which is significantly less than the vital

monitoring systems found in high income countries that are on the scale of hundreds to

thousands of dollars per patient. The hub component will add an additional cost component. However, since a single hub can communicate with multiple medical devices, this will be a one-time implementation cost. The potential displays that we would explore cost around $50 (Micro Electronics, Inc. n.d.). The software component that allows for the medical team to interact with the patient data being collected, including a database to store the patient's data locally would require almost no additional cost. Since implementation of the system and software component is cheap, yet not costless, the idea received a 0 within the affordability category.

Ease of Adoption and Integration for Users

The vital monitoring system presents minimal additional workload for hospital staff. While it introduces a new component to the daily nurse workflow that would require an education process and adaptation, it does so whilst automating vital collection for them. Unlike other ideas that received a -1 in this category, such as the communication platform or the nutrition tracking software, which either introduce completely new steps to physician's workflow or modify multiple steps of their workflow, the wearable vital monitoring would only modify a single component of the workflow. While we recognize that there would initially be an addition to workflow and a timely education process for all healthcare professionals, eventually this solution would optimize hospital workflow and even reduce the burden on nurses, which would increase incentive for users to adopt and integrate this into their healthcare facility. Additionally, once adopted, this solution is easily sustainable over time as the technological implementation in terms of the wearable and software would not need to be changed. Hence, the idea received a 0 within the category.

Enhances Communication, Collection, & Interpretation of Data

By collecting patient vitals, communicating them to nurses through a central hub, and interpreting the data by sending off an alarm when recorded data exceeds defined thresholds the vital collection device directly aligns with the objectives of our problem statement. For these reasons, the Vital Monitoring System received a score of 1 in this category.

Technical Sustainability

The technical sustainability of the Vital Monitoring System was ranked at a 0. The implementation requires low upkeep. The medical device will need to be charged, and the hub will need to be powered. Besides that, the device itself is intended to be disposable. There is the possibility of individual devices being broken and losing connection to the central hub. As a response to the this concern, the device itself will allow for rudimentary alarms, such that a nurse can be alerted to a patient even if his/her device is no longer communicating with the hub. Thus, we determined that the long-term technical concerns of this device are nontrivial yet surmountable.

Scope of Impact

For scope of impact, the Vital Monitoring System scored a 1. Not only does the solution increase the availability of accurate, more frequent patient but does so in a manner that decreases the work burden on nurses. By focusing on the three vitals- respiratory rate, heart rate, and temperature- the device collects data that can be indicative of sepsis which affects 40% of pediatric oncology patients. Given that data suggests that each hour of delay decreases the

patient's likelihood of survival by 8%, the ability to continuously monitor patient vitals and be alerted to abnormal vitals is likely to improve survival rates. Furthermore, this solution could be implemented in every pediatric cancer facility in our target communities.

Reliance on Tech Infrastructure

We rated the reliance on tech infrastructure of the Wearable Medical Record as a 1. The Vital Monitoring would require no additional infrastructure in a hospital. We would provide the display needed for the central hub component of the product. Since the data would be communicated through RF technology, this product would not require any network connectivity in the base case. Compared to the other ideas, the Wearable Medical Record received a 1 in this category because the pure software solutions, such as the Nutrition Tracking Software or the communication platform, require access to wifi in *addition* to a phone or computer.

Directness to Desired Outcome

We determined that the Vital Monitoring System would have a resoundingly direct effect on achieving the outcome we are seeking, thus receiving a score of 1 in this category. Since we envision the device being implemented in a hospital setting, the patients will be in a location with the most immediate access to treatment. A successful version of this device would surely provide quick, easy access to relevant patient information. Having access to this sort of frequent vital patient data will likely decrease the amount of time required to identify physical deterioration and sepsis within a patient and the quality of care both in LMICs like Uganda and also in higher income countries. For instance, a research study shows that 76% of children who

had all five vital signs measured in the emergency room received fluid resuscitation in the first hour after onset of sepsis as opposed to 61% of those who had an incomplete set of vital signs. Similarly, twenty percent of children who had all five vital signs measured received antibiotics in the first hour as opposed to 9% in children who had fewer vital signs measured (Hébert et al., 2017).

Engineering Feasibility

When rating the engineering feasibility of the idea , we not only considered whether we would be able to build a working prototype in two months, but also whether this solution would leverage the skill set of the entire class and could be split up amongst our team. The device is more challenging to engineer compared to the anthropometric multi-use scale, which received a 1 in this category. However, the Wearable Medical Record did not receive a -1 because it is one of the most modular solutions we have proposed. There are distinct mechanical, electrical, software, and bioengineering parts to this solution that could be divided easily among the class. This would also leverage the varying expertise of the entire class, whereas the more software heavy solutions would only target a subsection. Although this device is more challenging to build because there are multiple parts, depending on the technology chosen, the device could actually be fairly straightforward to construct. For these reasons, it received a 0 in this category.

Testable Feasibility

For testable feasibility, the vital monitoring device scored a 1. Once we have a working system, we will be able to test the functionality of the program, by running data transfer trials

between the device and the hub in order to debug problems, and the accuracy of the data collected by comparing it to positive controls. For the physical device itself, we anticipate conducting durability tests to determine the performance parameters of the hardware. Examples of tests could include shock and waterproof testing. Most importantly, GHI has offered to run a pilot test of both our hardware and software to allow for real-time clinical feedback on our solution, specifically whether it reduces time to access patient data and improves communication of patient data within a healthcare team. This will give us a much more relevant evaluation of our product than an attempt at a simulated hospital workflow.

R&D Costs

We gave a neutral rating of 0 to the research and development costs associated with our final idea. We determined that the approval process for this device would not be as costly as one for a diagnosis or medical treatment device; however, there may be some approval cost as a result of the collecting handling of potentially sensitive patient data. Phone apps and paper cards containing personal medical information do already exist and are not subject to any sort of stringent approval process, but we were somewhat concerned that the potential interaction of the device with a larger patient information database might raise cost-incurring legal concerns.

Market Landscape

The vital monitoring system received a neutral score of 0 in the market landscape category. Vital monitoring systems that record heart rate, respiratory rate, and temperature on pediatric patients do exist, so there is the possibility of competing with other similar products.

Both the EQ02 and BioRadio vital monitoring systems can be used on children and measure heart rate, respiratory rate, temperature, and oxygen saturation (McCarthy et al.). However, McCarthy et al. also note that most vital monitoring systems are designed for adults so there are relatively few solutions for pediatric care. Additionally, the battery life of the BioRadio is only 8 hours, so it is not suitable for continuous use in a pediatric cancer ward; meanwhile, the EQ02 is most likely not a low cost solution since the price is accessible only after contact with a vendor (McCarthy et al.). Therefore, there is room for the market for a low-cost solution designed to monitor vitals in a pediatric ward setting.

Chapter III:
Create

3.1 | Chapter Overview

This chapter serves to summarize the team's create phase, the purpose of which was to produce a working prototype of our vital monitoring system. We describe the preliminary technical specifications and designs considerations that significantly influenced our final product design, and then navigate through the different modalities of our product: the sensors, the housing for the sensors, the central hub for nurses, and the manner in which data was communicated between the device worn by the patient and the central hub that nurses are intended to interact with. For each of the modalities we detail the design considerations specific to the modalities, the testing that was completed to validate design, and further steps that will be taken to continue to refine the product.

Target Design Considerations & Technical Specifications

In the first part of the create phase, we established several design considerations for our device to guide our technical specifications. Some of the considerations we decided on were direct results of the constraints imposed upon our design by the inherent limitations to working in a low income hospital setting. Paramount amongst these was cost. We knew we needed to keep our cost as low as possible so as not to be a strain on the limited funds that these hospitals have to purchase new equipment. Additionally, due to the overburdened workload of many nurses in LMICs, we need our device to be minimally obstructive to workflow; thus, it needed to

be intuitive and easy to use. Also, due to the lack in robust technical infrastructure in several of our target hospitals, our device had to be able to function without the use of existing WiFi or computers.

In addition to the specific constraints of a low-income setting, we also outlined the general design constraints that arise from the actual intended function of our device. Specifically, we needed our device to reliably communicate across a hospital ward, so that all on-patient devices can transmit their data to the central hub for the ward. Furthermore, our sensors needed to collect data that is accurate enough to be clinically useful to hospital staff. If our data were too noisy, alarm fatigue could begin present issues for nurse care, and if our device measured too many false negative alarm states, important physiological warning signs could be missed, meaning our device would not be serving its purpose. Finally, we wanted to ensure that every child using our device is afforded maximal mobility and the best quality of life possible, so we set out to design a vital monitoring system that would allow each patient to be able to walk about as they please, and would be easy to wear for long periods of time without any discomfort.

Low-Income Setting Considerations	Low Cost
	Minimally Obstructive to Workflow
	Intuitive to Use
	Not Reliant on WiFi
	Not Require Existing Computers
General Considerations	Can Communicate Across Ward
	Accurate Data Collection
	Child Mobility

To set more quantitative guidelines for our design process, we next established the initial technical specifications for our product. With respect to cost, we set an upper limit of $300 per device, ensuring that our device would be as affordable and accessible as possible for low budget hospitals in LMICs. We also required our device to have a minimum lifetime of one year, so that the devices did not need to be replaced often enough to impose a significant economic burden on a hospital using the system. One way we intended to address this issue of longevity was by making our device water resistant. Since the device does make direct contact with human skin, particularly at the underarm, sweat could pose trouble for our sensors if they are not sufficiently water resistant. Specifically, we wanted to achieve IPX4 resistance, which requires our sensors to be able to withstand water splashes from any direction. Moreover, in order to prevent nurse workflow obstruction, we wanted the battery for the device to only require recharging once every 24 hours or more. This corresponds roughly to a given nurse needing to change the battery once every other shift, since shifts are typically 10-12 hours (Rogers, 2004), which we considered to be an acceptable workload addition. Considering this minimum battery life constraint, we set a maximum battery recharge time of one day so that a new set of batteries could be fully charged in the same amount of time it takes the old set to discharge. All the targeted technical specifications for the overall device can be found in Table 3.2.

Criteria	Requirements
Cost	< $300 per device
Lifetime	> 1 year
Battery Life	> 1 day
Recharge Time	< 1 day
Safety	Meet electrical & chemical safety standards
Water Resistance	Sweat resistant

Table 3.2: Table of general technical specifications for the vital monitoring device.

System Overview

After setting our technical goals, we began the creation phase by designing how the two-device system will interact with one another. Figure 3.0 outlines an overall block diagram for the whole system. First, each PediaTrack on-patient device will have a few main components: sensors to gather the temperature, HR, and RR data, a microprocessor to gather and analyze the information, an LED indicator to denote patient status, and a battery to power the device. However, as noted before, we also decided that nurses and clinicians would want a place to see this data in one central location to more efficiently analyze the patients in their ward and to best fit within their workflow, so each patient device needed a way to communicate module to send and receive information. The central location where we would send this information we named the hub. There, nurses would be able to see the status of each of their patients and the alarm status of their patients as well as their vital numbers. Furthermore, the hub would not only shows the most recent data points, but longitudinal data, which is important in giving clinicians a better context for their patient. Lastly, a key feature of this system is that the hub can

communicate with many patient devices and is very scalable, therefore reducing infrastructure

costs and providing a more unified place for patient information within a pediatric cancer ward.
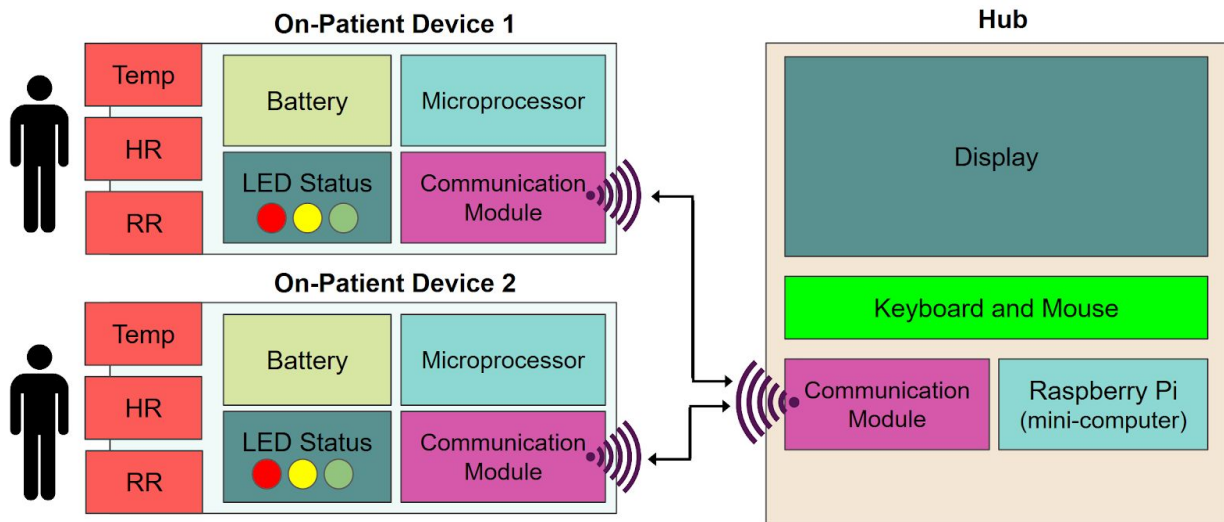
Figure 3.0 Block diagram of the Pediatrack system, with two on-patient devices and the Hub

Initial Prototype

Our initial prototype for the on-patient device was a wearable chest band with the heart

rate and temperature sensors embedded in the fabric. The band would measure respiratory rate

by being comprised of resistive fabric, such that it would register chest and abdominal stretching

from breathing as a voltage difference that could then be used to calculate respiratory rate per

minute. The band would then have heart rate and temperature sensors embedded in the band. We

looked at existing architectures to draw inspiration for our design, including the ADI respiratory

belt transducer ("Respiratory Belt Transducer," n.d.). In the create process, we built a prototype

which demonstrated the way we would secure our technology to the patient and the elastic band

that measured breathing (Figure 3.1).

The initial design had a two-part belt made of inextensible nylon webbing and elastic fabric. Most of the belt is made of the webbing and a small part located at the center of the chest (Figure 3.1) is the elastic fabric. The belt has shoulder straps to ensure proper placement of the belt. To make the model, we used super glue to bond the webbing and the flexible fabric. We used safety pins to attach the shoulder straps. To make sure it would fit, we selected one of the team members as our user and measured his chest and shoulders before cutting the fabric webbing. For adjustability, we used standard plastic belt adjusters that are commonly used in backpack straps and bicycle helmets.
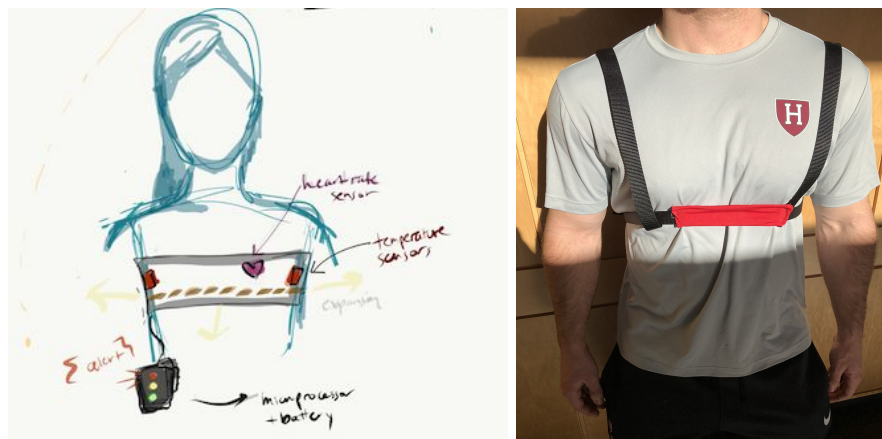
Figure 3.1: Initial design was a respiratory band intended to be placed on the upper torso of a patient

After creating this prototype, we sought client feedback on this on-patient design by attending the GHI's weekly meeting on November 13 and speaking with three physicians on the team. Their feedback emphasized that the resistive band would be challenging for a child to wear, as it might make them feel weighed down and restricted, and suggested that we redesign the on-patient prototype to prioritize patient comfort and mobility. Based on this feedback, we ideated possible designs to mount sensors and measure respiratory rate effectively on the patient without using a resistive band. Ultimately, we decided that a sticker-based device would allow

the most freedom of movement, that would incorporate our sensors for heart rate (HR),

respiratory rate (RR), and temperature sensor into a disposable adhesive.

3.2 | Sensors

Introduction

The purpose of the sensor components was to retrieve the data for respiratory rate, heart rate, and temperature from the patient's body in a manner that allows for continuous monitoring of their vitals and early recognition of abnormal vital signs. Doing so required selecting the proper sensors, determining their optimal location to prevent impediment to the child's mobility, ensuring accurate measurements, and processing the sensor-received data to standardized output values, which would be easily packaged to send to the hub as heart rate, respiratory rate, and temperature measurements.

Design

After eliminating our initial resistive band prototype based on feedback that it would be obstructive to the child, the team first had to design a method for sensing heart rate, respiratory rate, and temperature in a minimally invasive way.

Ultimately, the team decided to use an echocardiogram (ECG) to measure heart rate (HR) and respiratory rate (RR) and temperature sensor for body temperature. This design was based on literature research demonstrating that ECG signals can be used to calculate both HR and RR, as the QRS peaks in an ECG occur once every heartbeat, but also change in amplitude with one's breathing (Charlton, Bonnici,Tarasenko,Clifton,Beale, Watkinson, 2016). The ECG was particularly compelling for this reason, as determining respiratory rate would not require any obstructive physical attachment; rather, all that would be required would be adhering three

simple electrode stickers on the patient, which aligned with the team's overall aim to create a sticker-based device.

ECG Circuit and Design

After determining that an ECG would serve as our HR and RR sensor, we then needed to build our circuit. While ECG circuits exist commercially, these circuits are often pricy and overly complex: for example, even simple single-lead ECGs are up to $20 dollars commercially ('Single Lead Heart Rate Monitor, n.d.'). Instead, the team opted to build our own ECG circuit, as this was estimated to be much cheaper, since many of the necessary components come for under a dollar each.

Before building our circuit, we determined that the final build would need to follow technical specifications:

Specification	Threshold	Explanation
Cost	< $20-30	Simplest and cheapest ECG heart rate monitors are around this price range ('Single Lead Heart Rate Monitor, n.d.')
Accuracy	+/- 7.5 beats/minute +/- 3/4 breaths/minute	ECGs can have up to 100% accuracy; natural heart rate variability is 15 bpm from difference between inhalation + exhalation (Shaffer & Ginsberg, 2017)
Range	25 - 240 bpm	Physiological range and specification from similar solution
Precision	+/- 0.5 bpm +/- 0.5 breath/min	Clinically, heart rate is reported to the nearest bpm, and respiratory rate is reported to the nearest breath/min.

Operating Temperature	No more than 110° F	1st degree burns begin at 110° F
Power	< 20mW	Minimum: $4mA \times 3V = 12\ mW$ Maximum: $4mA \times 5V = 20\ mW$
Current	< 10 mA	10-20 mA is an average range for "let-go" current, and 100 mA is enough to cause ventricular fibrillation (Fish & Geddes, 2009)

Table 3.3: Technical specifications and notes for ECG circuit.

Our ECG circuit design took inspiration from an ECG circuit that was built in ES152: Circuits, Devices, and Transduction, and designed by Teaching Fellow Evan Smith. This circuit includes a 3-lead ECG with a safety block (1), an amplifier (2), a high pass filter (3), a notch filter (4), and a low pass filter (5).
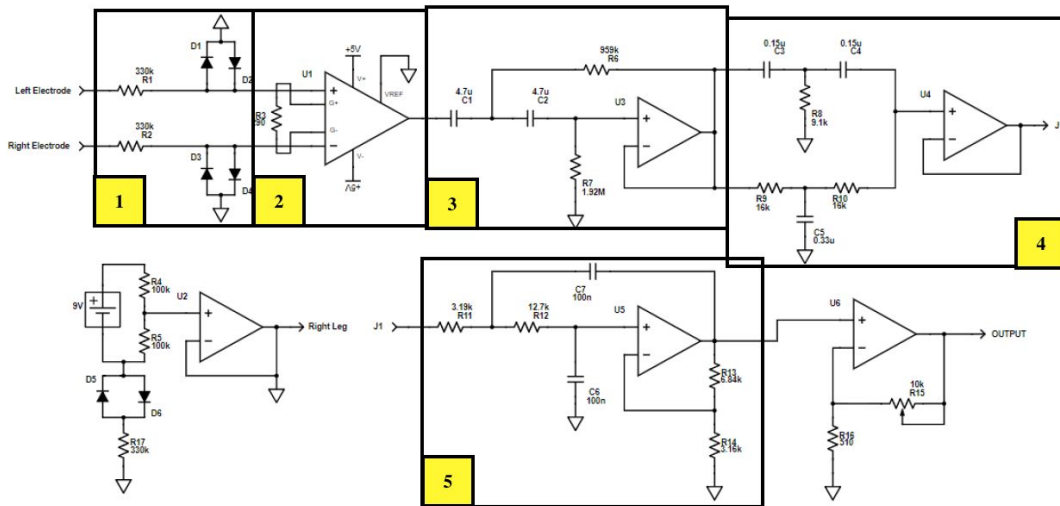
Figure 3.2: ES 152 ECG build with specific blocks highlighted: safety block (1), amplifier (2) high pass filter (3), notch filter (4), and low pass filter (5)

The original ES 152 ECG circuit design was useful inspiration, but ultimately its various filters would have likely required a large circuit board. To cut down the complexity of the circuit, the final design of our ECG circuit only used a protection circuit, a difference amplifier, and a

high pass filter–essentially, the first three pieces of the ES 152 build. Each of the two leads feed

into one of the ECG inputs on the difference amplifier, denoted as ECG-L and ECG-R. These

each pass through the pair of over-current- and over-voltage-protecting diodes in box 1 of Figure

3.5. The difference amplifier takes the difference between the voltages at the two ECG leads to

find the unfiltered ECG signal, which is then altered with a high pass filter with corner frequency

of approximately 550 Hz, as calculated in below.

$$fc = \frac{1}{2\pi RC}, \text{ where } R = 290 \text{ k}\Omega \text{ and } C = 1\mu F$$

$$fc = \frac{1}{2\pi*(290*10^3)*(1*10^{-9})}$$

$$fc = 550 \ Hz$$

The third ECG lead serves as our ground. The output of this block is then fed into an

Arduino Pro Mini, loaded with the script necessary to analyze the readings for heart and

respiratory rate.
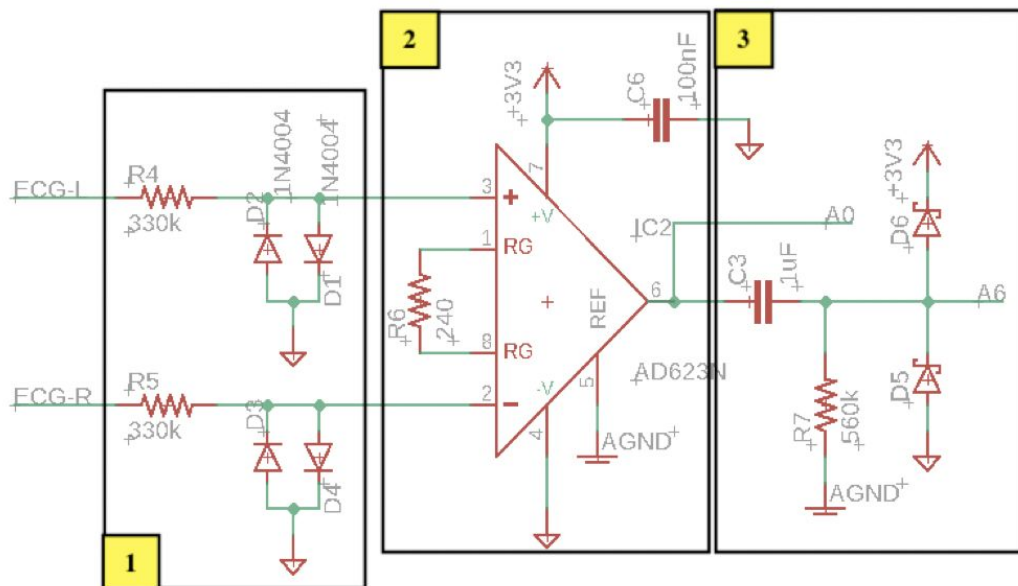
Figure 3.3: Schematic of ES 96 final build for ECG circuit as drawn on Eagle PCB design software.

Another design consideration was that the device must be portable; thus, the entire circuit must be powered by a portable source, such as a battery. The ES 152 circuit design was created for a non-rechargeable 9V battery. However, this was unideal for our purposes, as these batteries are single-use, and often cost around $3-4 (Walmart, n.d.), therefore having an increased cost. Thus, we instead opted for a rechargeable battery to avoid a recurring cost and large waste. As such, the final design for this ECG assumes a rechargeable 3.7V battery, with the construction of a pseudo-ground to support a single-supply circuit design.

Optimizing ECG Lead-Placement

Once we had built our circuit and were receiving a clear QRS-signal, we ran several preliminary tests to determine the optimal placement of our ECG electrodes on the patient's chest. We recorded signal response at various locations on the human body under changing positions - both still and in motion. We initially chose 5 different placements to test on both sides of the body (Figure 3.6): middle of the chest (Location I), the left clavicle (Location II), the right clavicle (Location II'), under the left arm (Location III), and under the right arm (Location III'). Placement was tested under the arms to see whether we could get an accurate ECG signal in a location close to the temperature sensor.
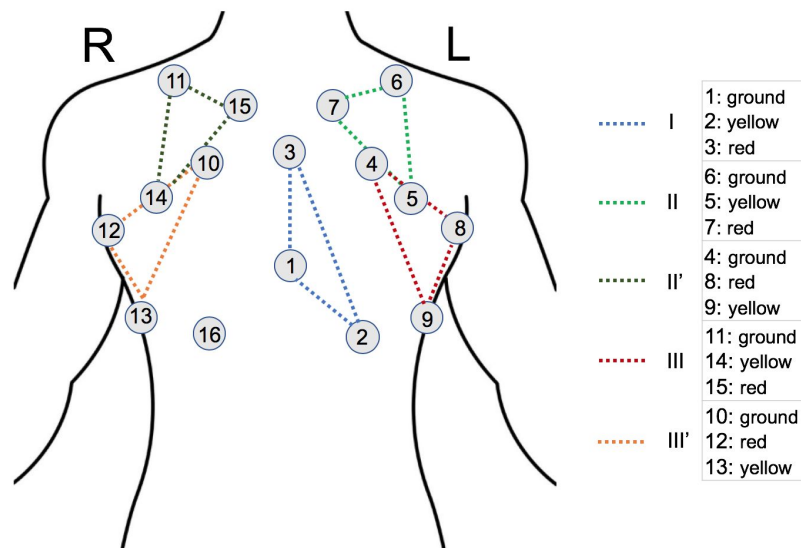
Figure 3.4: Diagram of tested ECG electrode placement locations. Note that 'red' is positive lead, 'yellow' is the negative lead.

We attached a subject to the ECG circuit with the electrodes placed in each subsequent configuration and recorded the signal quality as well as the maximum and minimum amplitude peaks when the subject was resting, coughing, and making various movements of the chest and

arms. After initial analysis, we eliminated placement at the right clavicle as well as under the right arm as these locations resulted in nearly no detection of an ECG signal. We proceeded with more rigorous trials at the remaining locations at the middle of the chest and on the left side of the body. The conclusion to test placements on the left side of the body is supported by the understanding that proximity to the heart improves the accuracy of the acquired signal.

As the subject's ECG signal was recorded in each of the electrode placements, we used an oscilloscope to determine the maximum and minimum voltage amplitudes while the subject was resting and then qualitatively analyzed the impact of movement on the noise of the signal in each of the locations. As shown in Table 3.4, for the signal recorded at the middle of the chest, the maximum change in voltage was 1.66 V and the minimum change in voltage was 1.12 V. At the left clavicle, the maximum change in voltage obtained was 1.34 V and the minimum was 1.04 V. Finally, under the left arm, the maximum change in voltage observed was 0.720 V and the minimum change was 0.560 V. As seen in the images below, the signal obtained at the middle of the chest had the least amount of noise and the largest detectable amplitude peaks, which is important in determining heart rate and respiratory rate. The signal from the left clavicle was corrupted by quite a bit of noise and did not have as large of an amplitude. And the signal obtained from under the left arm was not detectable with much smaller amplitude peaks.
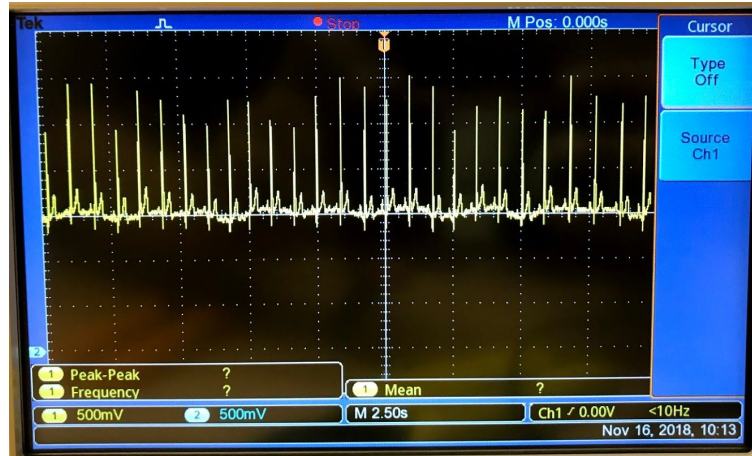
Figure 3.5: ECG signal recorded at the middle of the chest (Location I). Note that the signal is not very noisy and has large detectable amplitude peaks.

Figure 3.6: ECG signal recorded at the left clavicle (Location II). Note that the signal is quite noisy and has smaller amplitude peaks.
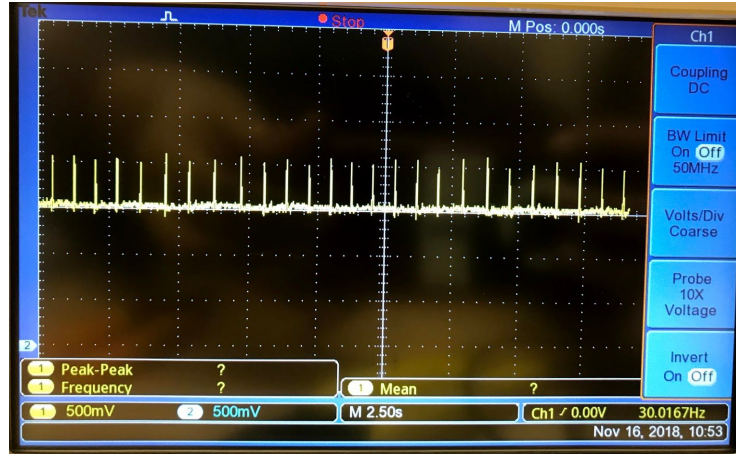
Figure 3.7: ECG signal recorded under the left arm (Location III). Note that the signal is extremely noisy and a proper heartbeat including the P or T-wave is not detectable.

Placement Configuration	Electrode Leads	Max. ΔV (peak to peak amplitude)	Min. ΔV (peak to peak amplitude)
I (middle chest)	1: ground 2: negative 3: positive	1.66 V	1.12 V
II (left clavicle)	6: ground 5: negative 7: positive	1.34 V	1.04 V
III (under left arm)	4: ground 8: negative 9: positive	0.720 V	0.560 V
II ' (right clavicle)	11: ground 14: negative 15: positive	no signal	no signal
III ' (under right arm)	10: ground 12: negative 13: positive	0.320 V	0.200 V

Table 3.4: The maximum and minimum voltages recorded, for all five placement configurations, when the subject remained seated. Location I resulted in the strongest signal with the maximum voltage peaks.

We also qualitatively examined the impact of various movements on the signal obtained at each of the three locations. In three separate trials, the subject coughed, moved forwards and backwards, and then moved their arms up and down. We then observed the signal recording to see how stable the ECG was even under each of these manipulations. As seen in Figures 3.10-3.11, with samples of the signal obtained while the subject moved their arms, at the middle of the chest, there was a slight increase in noise, the signal remained fairly stable. At the left clavicle, the signal was even more corrupted with noise and the peaks became less detectable. The signal obtained under the left arm became extremely noisy with a great deal of shift, making the detection of each distinct heartbeat very difficult.
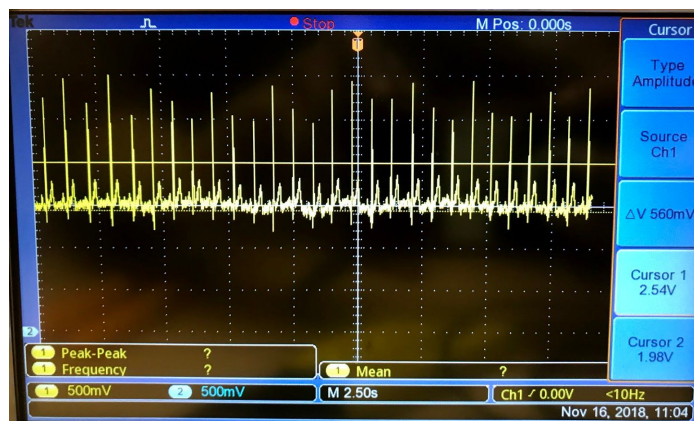
Figure 3.8: ECG signal recorded at the middle of the chest (Location I) while subject moved their arms up and down. Note the relative stability of the signal despite the movement.
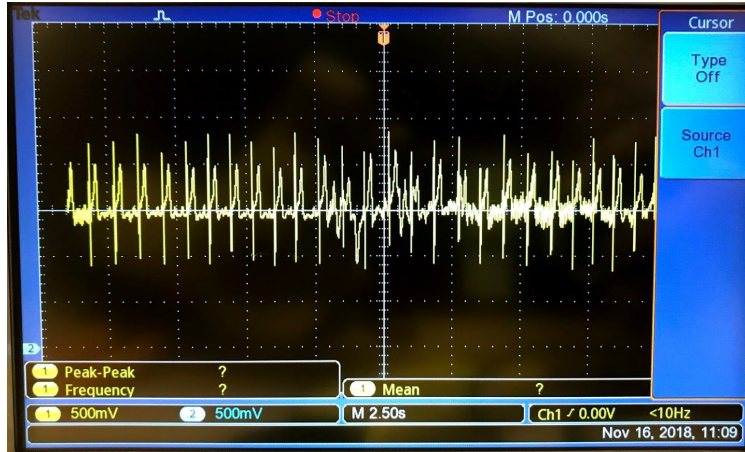
Figure 3.9: ECG signal recorded at the left clavicle (Location II) while subject moved their arms. The signal becomes fairly noisy and the peaks more difficult to detect.
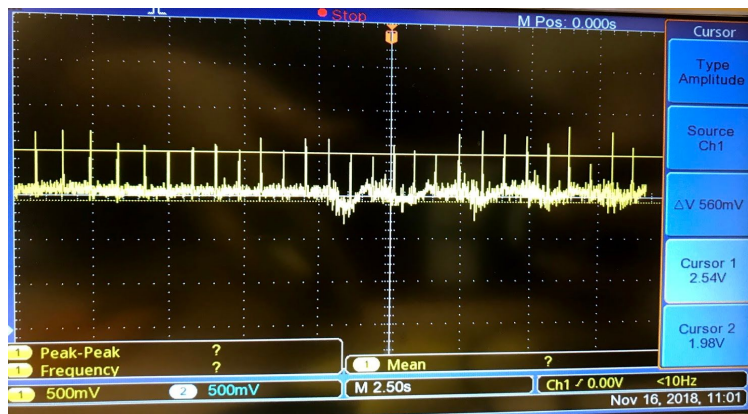
Figure 3.10: ECG signal recorded under the left arm (Location III) while subject moved their arms. The signal is corrupted by extreme noise with variable shift in amplitude size of the QRS peaks.

From our quantitative analysis of the amplitudes and qualitative analysis of the stability of the ECG signal at each location, we determined that the optimal placement for our ECG electrodes would be Location I, or the middle of the chest. At this placement, the ECG signal is clear, with easy detection of the QRS peaks for heart rate and respiratory rate determination, as well as relatively stable under various simulated movements of the subject. Further testing of the accuracy of obtaining heart rate and respiratory rate from the ECG was conducted at this location.

Heart Rate Algorithm

In order to obtain heart rate from our ECG, we had to code an algorithm in Arduino. To interpret the raw data, we designed a simple three step algorithm. The first step of the program had to determine QRS peaks so that we could uniquely identify each heartbeat. We accomplished this by creating a dynamic voltage threshold that changes based on the average height of the last ten peaks. Before collecting data, the algorithm undergoes a short initialization setup. In this period, the code identifies one peak so that it can set a first threshold. The peak is determined by identifying a maximum point over a period of a few seconds. Once the initial threshold has been set, only peaks that are at higher voltages than the threshold will be flagged as QRS complexes and averaged into the threshold moving forward.

The algorithm then collects data, identifying QRS peaks, over a period of sixty seconds. Once sixty seconds of data has been collected, the patient's heart rate is determined by taking the number of beats detected, which is equal to the number of detected QRS complexes, over the interval over which they were collected. This process reduces random, physiologically insignificant fluctuations in heart rate such that the resulting value best reflects the patient's actual heart rate at the time of measurement.

The algorithms we ultimately chose went through a series of design changes in response to different constraints and considerations. One of the foremost limitations that Arduino imposed on our algorithm was that each Arduino Uno only had a very small amount of RAM. So, to avoid the issue of insufficient storage, we intended to use FRAM chips, which add additional storage for us to utilize. We ultimately did not pursue this design as it added cost and because we were able to implement an accurate detection program without relying on the extra storage on the chip.

Additionally, we first had not used dynamic thresholding and had instead opted to manually set a

voltage threshold for determining QRS complexes, but we worried that this would not produce

reliable results. Thus we moved to the dynamic system. Another system we have considered for

QRS detection is a dynamic threshold that would detect slopes rather than absolute voltage

values. This could ideally avoid noise registering as heartbeats, and early tests of this version of

our peak detection on sample data in MATLAB show that it is remarkably accurate when well

implemented, so moving forward we may want to look further into putting such a system onto an

Arduino and comparing it with our current algorithm.
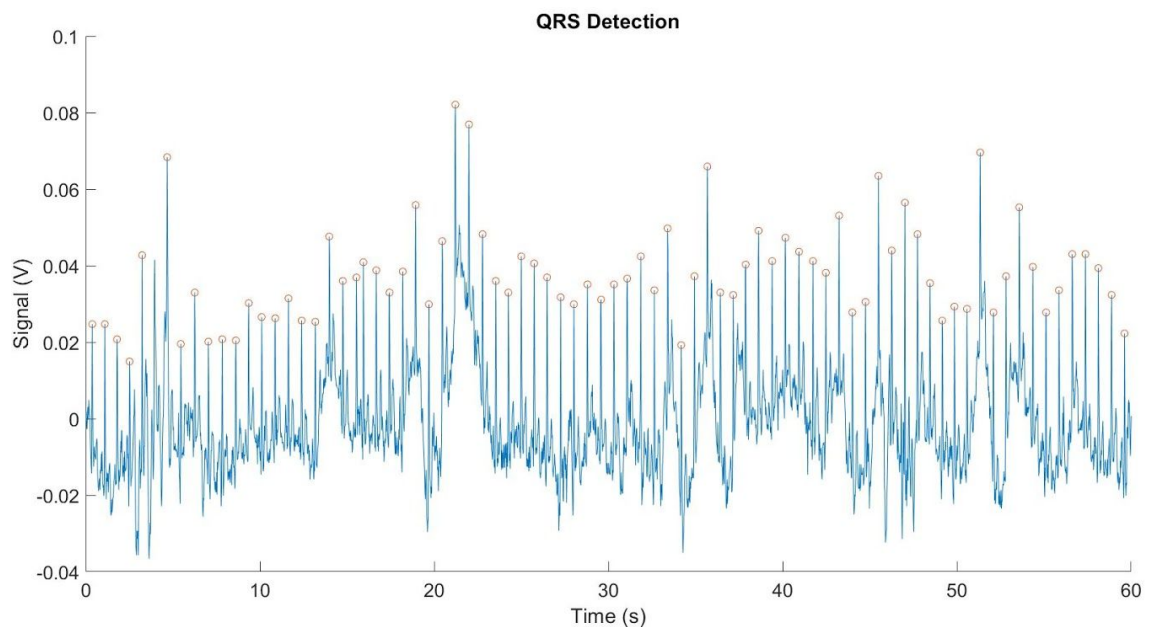
Figure 3.13: Test slope thresholding algorithm run on sample data in MATLAB. The ecg signal is shown in blue. Detected QRS complexes are shown as orange circles. We observe only one false negative and no false positives over the sixty second interval of the highly noisy signal.

In order to test for the accuracy of our ECG and algorithms, we tested one live subject by

hooking them up to the prototype while simultaneously testing their heart rate manually and

comparing the results. Each trial was 60 seconds and each subject was tested over 15 trials. The

manual testing consisted of measuring heart rate with our fingers placed at the neck of the subject and counting the number of beats that occured per minute. The prototype testing consisted of placing the leads of our prototype onto the appropriate parts of the body and reading the algorithm's output displayed on the screen.We then analyzed these recordings comparing the differences between the two for each trial in order to calculate the average difference between the two as well as the p-value to determine there was any statistical significance between these two methods of testing and their results.

Figure 3.11. Representative graph for one subject of manual heart rate (blue) as compared to sensor-determined heart rate (red) over 15 trials of 60 seconds total.

Our results found that the heart rate remained relatively constant over the period of 15 trials, which is what you would expect from heart rate. There should be no significant variation in heartbeat over the span of 30 minutes assuming that there is no large change in activity. We did find, however, an average difference between recordings of 9.3% and a p-value of under

0.05, indicating a significant statistical difference between the methods of recording. We believe that this difference might be due to an inaccurate manual HR reading vs an inaccurate sensor reading. Given that we are untrained medical professionals, and that we used a manual method of recording as our control, there could have been room for error - for example, we may have miscounted the number of beats or missed a few within the trials. Indeed, we manually measured heart rate by placing two fingers at the neck of our subject and counting the number of beats that we felt per minute. In fact, in our graph, our manual HR recording is far less consistent with the more steady sensor HR recording, supporting the notion that our manual HR recording was more erroneous than the manual one, contributing to the statistically significant difference in recording values. We believe that consulting a medical professional on how to more accurately record manual HR will go far in improving our control manual HR recording. Once we have a more stable manual heart rate, we will be able to more thoroughly test the accuracy of our built sensor.

Respiratory Rate Algorithm

The algorithm that assesses respiratory rate from our recorded ECG signal is largely divided into two stages: extraction of the respiratory signal and estimation of respiratory rate. To extract respiratory signal, we needed to identify the location and amplitude of the R waves, which we have already done when estimating heart rate. We simply use the location and amplitude of each R peak to impose a sinusoidal over the collected R-waves. We apply a smoothing function to the sinusoidal such that the number of cycles of the sinusoidal over a period of 60 seconds was then equated to respiratory rate. To determine the number of cycles within a given time frame, we compared each data point comprising the smoothened sine wave

to a point falling before and after it. When a point was associated to an amplitude larger than the point falling before and the point falling after itself it was recorded as a breath.
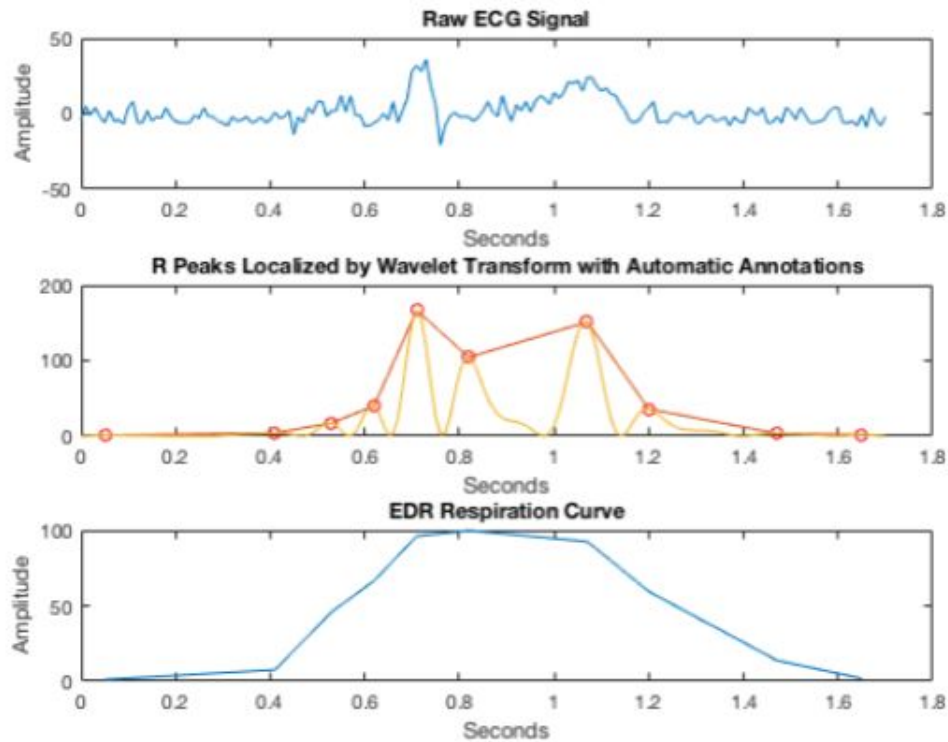
Figure 3.12: (1) The timestamp and amplitude of R-wave peaks of the raw ECG signal are identified and recorded, (2) a sine curve is then imposed over the compiled R-wave peaks, (3) once the sine curve is smoothened the number of cycle/minute equals respiratory rate.

In order to test for the accuracy of our respiratory algorithm, we tested our subject by attaching them to the prototype while simultaneously testing their heart rate manually and comparing the results. Each trial was 60 seconds, and our subject was tested over 15 trials. The manual testing consisted of measuring the difference in temperature under the nose of our subject with our fingers indicating a breath. The prototype testing consisted of placing the leads of our prototype onto the appropriate parts of the body and reading the algorithm output displayed on the screen. We then analyzed these recordings comparing the differences between the two for

each trial in order to calculate the average difference between the two as well as the p-value to determine there was any statistical significance between these two methods of testing and their results.

Figure 3.13: Representative graph for one subject of manual respiratory rate (blue) as compared to sensor-determined respiratory rate (red) over 15 trials of 60 seconds total.

We found an average difference between recordings of 21.7% and a p-value of under 0.05, once again indicating a significant statistical difference between the methods of recording. In this case, and as evidenced by the graph above, while our manual respiratory rate remained relatively constant (expected in respiratory rate over the course of 30 minutes with no significant change in activity level), the sensor recording was far more volatile. Indeed, we believe this might be due to a technological difficulties with our respiratory rate algorithm, which will be commented on in future steps.

Temperature Sensor

Circuitry

For our temperature sensor, we wanted a contact sensor (rather than ambient temperature) with medical grade accuracy and a small profile that would be comfortable for patient use. We also needed something with a data output type that would be easily integratable into our final sensor design. After searching for commercially available temperature sensors, we settled on the MAX30205 Human Body Temperature Sensor by Maxim Integrated. The full list of required technical specifications, as well as the specifications of our chosen sensor, are listed below.

Qualification	Intended Specs	Sensor Specs
Price	< $5	$1.60/unit
Size	< 5mm x 5mm	3 mm x 3mm
Accuracy	35.8ºC -39ºC: ±0.2º	35.8°C - 37°C: ±0.2º 37°C - 39°C: ±0.1º
Operating Supply Current	< 10mA	6µA
Supply Voltage	work on 3.3V	2.7V - 3.3V
Communication Interface	I2C or analog	I2C

Table 3.5: Temperature sensor technical specifications

This device converts the temperature measurements to digital form using a high-resolution, sigma-delta, analog-to-digital converter (ADC), and when soldered to the final auxiliary PCB, the accuracy of our sensor meets clinical thermometry specification of the ASTM E1112.

The device comes in an 8-pin TDFN package, meaning that we needed a TDFN to 12-pin converter. This conversion was achieved using a temperature "buddy-board", allowing us to interface directly with our PCB.

The algorithm we use to obtain an actual temperature value starts by obtaining a raw 16-Bit reading from an output pin of the sensor. This reading can then be converted to Centigrade via a constant representing the LSB (Least Significant Bit). Our algorithm then repeats this process, generating a new reading every 100ms for 15 seconds. We can then average these 150 readings directly on the patient-device, and then output our final temperature reading to the hub.

Testing Protocol

After having built the circuitry for the temperature sensor, we tested the accuracy and sensitivity of the temperature sensor against a medical grade thermometer. We conducted these tests on two separate locations, the underarm and the upper center chest region, to see which location gave the optimum and most accurate results. A series of temperature readings was taken in both locations with both the temperature sensor and the standard thermometer.

We established a set protocol to standardize the testing and allow for the data to be accurately compared. For the underarm location we first took 15 measurements using our thermometer. When taking measurements from the sensor, however, we applied the sensor under our subject's arm and then waited for the readings to stabilize on the display before recording the displayed reading as our temperature reading for that test. We used a similar system to test the accuracy of the chest region, but when taking measurements from this area we used a cloth and tape to both secure and insulate the temperature sensor and the thermometer, ensuring that

ambient temperature was not affecting the reading. The data from these tests can be seen in the graphs below.
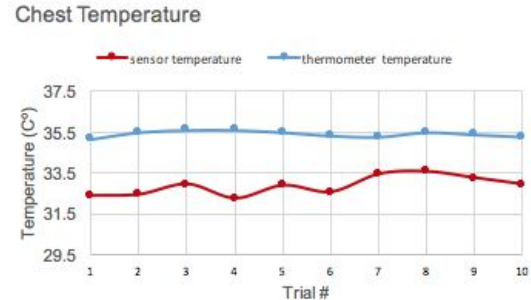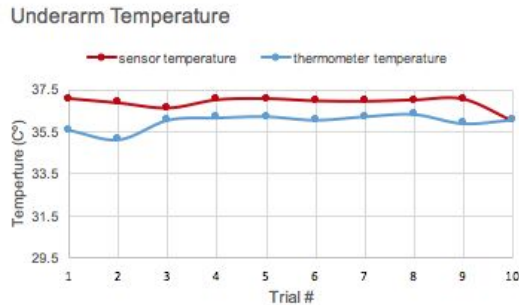
Figure 3.14(a) - Underarm Temperature   *Figure 3.14(b) - Chest Temperature*
In the figures above, red represents the temperature obtained from the MAX30205 sensor and blue represents the temperature from the thermometer.

Upon the completion of these tests we found that our sensor's temperature readings under the arm were very similar to those of the thermometer whereas the readings at the chest region were not so closely related. The data depicted in the graphs above confirm the correlation between the underarm temperature versus the lack thereof for the chest temperatures. Because of these tests we concluded that the underarm was the best and most accurate place for our sensor to take temperature readings.

Data and Results

After running the preceding tests, we then analyzed the data to determine if our results were statistically significant. We first ran an F-test to determine whether we should run a T-test assuming equal or unequal variances. From here, we ran our T-tests, and the raw data from our experiments (visualized in box-and-whisker format), as well as our P-values and variances, can be found below.

Calibrated Sensor Testing (Trial 1)	Calibrated Sensor Testing (Trial 2)
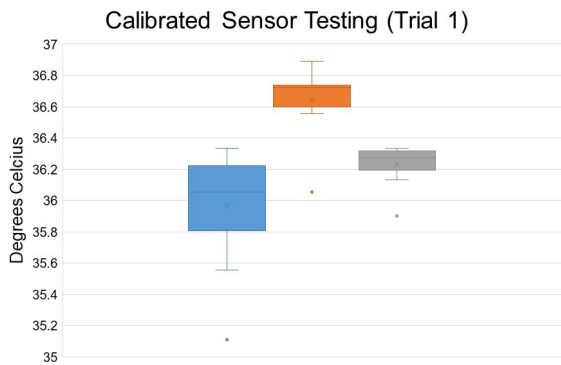
P-value = 0.06   P-value = 0.25

Sensor variance = 0.02   Sensor variance = 0.05

Figure 3.15(a) - Calibrated sensor test, trial 1   *Figure 3.15(b) - Calibrated sensor test, trial 2*

In each of the above figures, blue represents the standard thermometer underarm temperature, orange represents the standard thermometer oral temperature, and the grey represents our underarm temperature sensor.

For the above results, our null-hypothesis was the two data sets (Underarm thermometer temperature and underarm sensor temperature) were statistically different. However, since our p-value in both tests was greater than our set p-value threshold ($p < 0.05$), we could reject our null hypotheses and conclude that there is no statistically significant difference between our two data sets. This means that our sensor is generating the accurate values that we would expect.

Conclusion

In conclusion, the sensor components as a whole met most of its goals of collecting accurate data, processing that data into useful information, and preparing the data for transmission. All sensors gave accurate output for relatively low cost, and heart rate and temperature were able to be fairly accurately derived from these data points. Specifically, the temperature sensor demonstrated no significant statistical difference from a medically approved thermometer. There were some statistically significant discrepancies between manual and sensor-gathered heart and respiratory rates, but these are at least partially due to limitations in our ability to accurately take manual measurements for these quantities.

Next Steps

Going forward, there are a few changes to be made that could improve the function of the sensing portion of our device.

First, we need to address the saturation problem that has been a ongoing issue in the ECG circuitry. The saturation issue occurs when the offset of the ECG signal extends beyond the limits of our power supply (0 and 3.3V), and often depends on the placement of the ECG and subject movement. If the drift of the signal remains completely within the limits, there is no saturation. If part (or all) of the signal extends beyond the limits, then that part (or all) of the signal is lost[5]. If only part of the signal is lost, then often our heart rate algorithm still functions effectively; however, because we often lose the precision of the peak amplitude, this issue affects the respiratory rate algorithm. And unfortunately, when the signal saturates completely, we do not have a useful ECG signal. To account for this problem, we implemented a potentiometer that can be used to adjust the voltage reference to account for the saturation problem and bring the ECG signal back within range of the power supply limits. This was a successful work-around, but requires manual adjustment when the signal saturates. Going forward, we would use a digital potentiometer that would allow us to dynamically change that without the need for manual intervention, completely removing this as an issue in the device.

Secondly, we should improve our respiratory rate algorithm, as this measurement has the highest average difference from manual measurements. Solving the ECG saturation issue would help address this, as it would lead to more oscillatory peak amplitudes, therefore allowing our algorithm to better detect respiratory rate. Additionally, the respiratory rate algorithm itself could

---

[5] This issue is not present on commercial ECG monitors due to the wide voltage range they operate on (usually greater than 12 Volts).

be improved to make sure that each breath flagged is, rather than just a local max, is true the max over the time period of a breath (reducing false positive breath detections). Lastly, we should also test our heart rate algorithm more rigorously, and determine if our algorithm works better using a slope rather than value based threshold.

PCB Design and Implementation

Introduction

Following the selection, evaluation and validation of the sensors the device would be using, the team required a method by which to integrate them all in a compact manner and in such a way as to allow for a 'plug-and-play' user hardware interface. Additionally, since the above tests were conducted on prototyping test beds, such as breadboards and regulated benchtop power supplies, we needed this solution to incorporate a mobile regulated battery power supply. Additionally, the solution needed to incorporate the LEDs that would give a patient and caregiver information about the alarm state.

Subsequently, the team decided to use a custom-made printed circuit board (PCB) to achieve this. A custom PCB allows for a bespoke circuit design that is easier to debug than alternatives (such as a solderable protoboard) in a compact design. This would also ultimately allow more flexibility for further amendment as well, given the design considerations specified below.

Component Selection and Schematic Capture

When making these considerations, we had to keep in mind that we had limited time to get a functioning prototype. As such we used as many off-the-shelf components for non-bespoke functionalities as we could. This would mean that we save time in testing, manufacturing and validating the PCB, as we would not need to assemble and test as many subsystems. Additionally, as we only had time for one PCB revision, we wanted to make the design as flexible as possible to allow us to integrate new functionalities as fast and seamlessly as possible.

Power Management

We decided to continue using the Adafruit batteries combined with their LiPo battery charger. This prevented us from having to integrate battery charge and safety circuity on the board. The charge controller's output was connected to the Arduino's built-in 3V3 regulator, further removing the need for a regulator circuit.
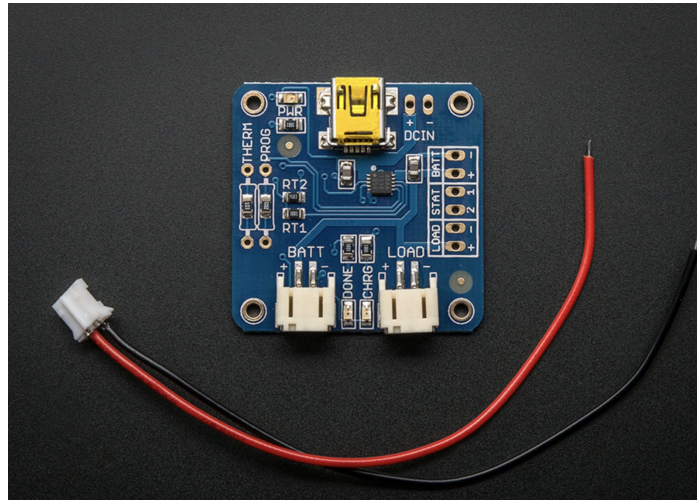
Figure 3.16: Adafruit's Battery USB charger (USB LiIon/LiPoly charger, n.d.)

Safety and User Interfaces

In addition to the Adafruit charge controller for battery safety, we needed to keep in mind that we are interfacing with the human body. As such, we needed to have current and voltage limiters – in the form of a resistor and diodes respectively – on each lead that connects to the human body. As demonstrated below, a series 330kΩ resistor is a current limiter and the diodes will short current to ground if voltage exceeds their 0.6V forward voltages. The diodes are through-hole to allow for high currents to flow before burning out.
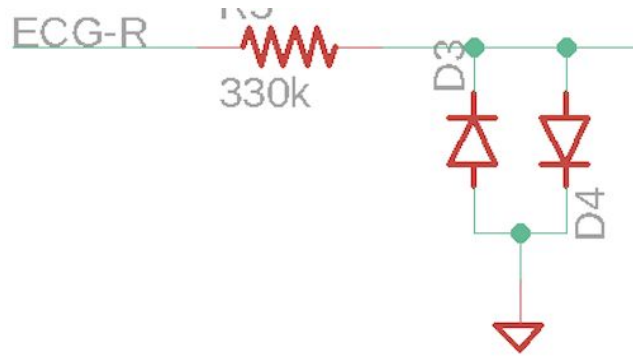
Figure 3.17: ECG lead safety circuitry.[6]

In order to indicate alarm states to the user and to allow us to determine the state of the

program while debugging, we included three programmable LEDs and one power indicator LED.

While the current draw is on the order of 10mA which reduces battery lifetime, the only LED on

consistently is the power indicator LED, which is required for the user to establish correct

operation. The others are programmable and can be turned off from the microcontroller. These

components are surface mounted devices (SMD) to reduce footprint size.
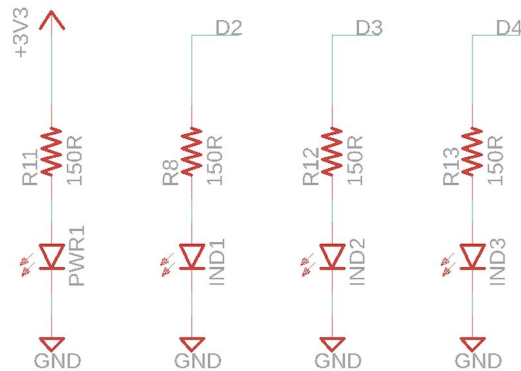
Figure 2.18: Indicator LEDs schematic

---

[6] Design courtesy of Evan Smith and ES152.
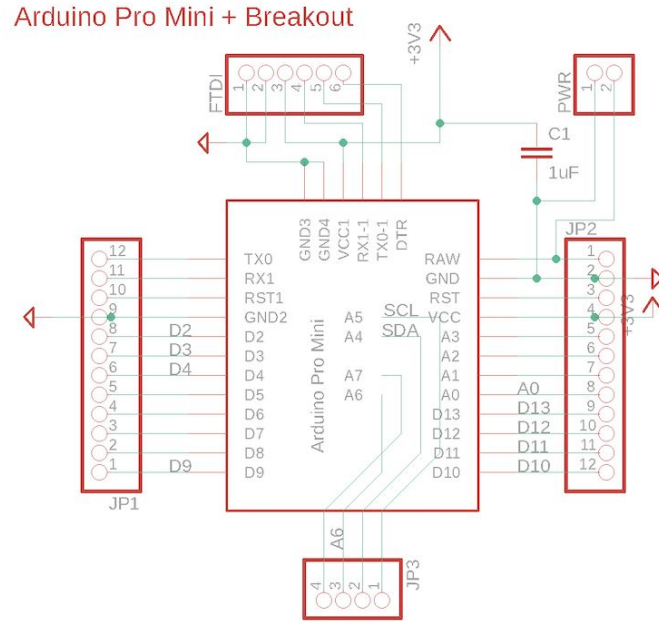
Figure 2.19: Arduino Pro Mini schematic

We chose to use Sparkfun's Arduino Pro Mini due to its compact size and high performance specifications. This would be mounted onto the PCB by three sets of 2.54mm female header pins that connect to the 2.54mm male header pins soldered downwards from the Arduino. This way, we could easily swap out microcontrollers should one break or get damaged.

Keeping in line with the expandability specification, we broke out all of the arduino pins to 2.54 mm male header pins that would allow us to easily connect the Arduino to extra components should the need arise. Likewise, we broke out the unregulated and 3V3 regulated power pins for debugging.

In order to program the Arduino Pro Mini, we need to use an off-board FTDI chip to translate the Rx/Tx serial port on the arduino to USB to interface with the Arduino IDE. These pins were likewise broken out and positioned to allow for easy connection to Sparkfun's FTDI connector.

All wires in the schematic that have labels are used on the PCB. The others are unused.
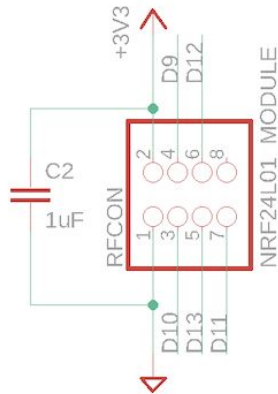
Figure 3.20: Radio Module Connector Schematic

In order to further ease prototyping, the RF radio connector was designed to match the
pinout of the RF radio module, thus allowing for the module to sit neatly on top of the PCB.
Again, this prevents the need for us to assemble a sensitive radio module on the PCB proper and
allows us to swap the modules should one malfunction. C2 is bypassing the 3V3 supply bus to
improve power stability. The connector is a 2.54mm female 2x8 header, thus also providing
mechanical stability as well as a good electrical contact.
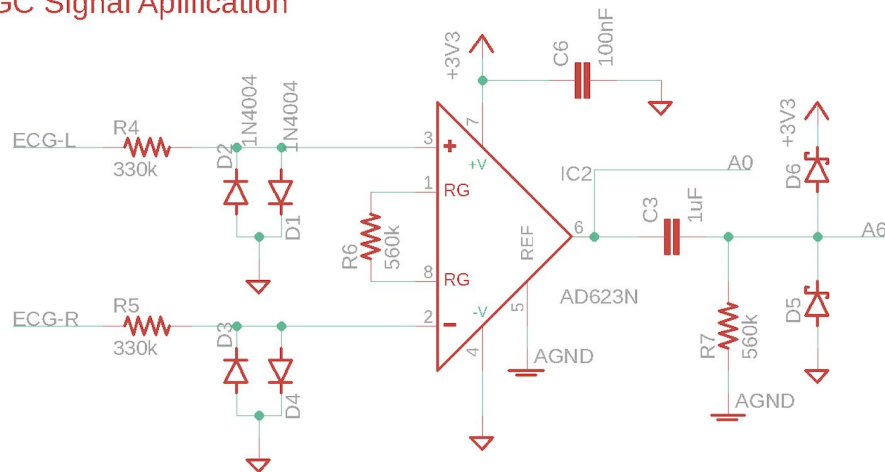
Figure 3.21: AD623N Schematic

The rationale behind using this topology has been covered earlier. The diodes and AD623N are all through-hole to allow for easier reworking of the component topology should the need arise. The capacitors and resistors are SMD to minimise the size of this circuit. C6 is bypassing the 3V3 supply bus to improve power stability. D5 and D6 are schottky diodes to protect the arduino pins from exceeding their specified 0V – 3V3 range.

AGND refers to the pseudo-ground reference supply that is required for single-supply operation. As we are using a battery to power the circuit, in order to process negative voltage analog signals, we use AGND as a zero-reference for all analog circuitry to ensure that the voltages at the microcontroller inputs stay within the 0V – 3V3 range.
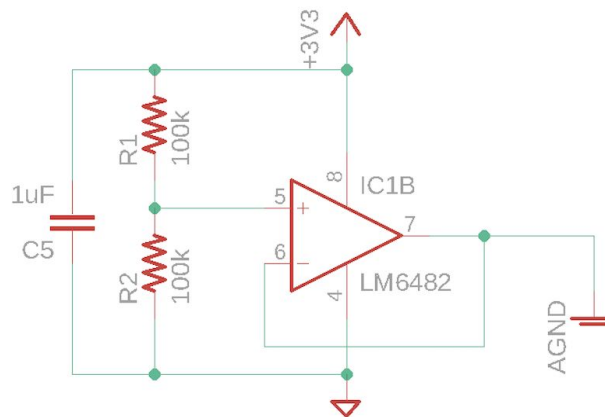
Figure 3.22: 1.15V Pseudo-Ground Schematic

Finally, to improve the user experience when interacting with the device, we decided to use a familiar 3.5mm jack ("headphone connector") to connect the ECG leads to the device.
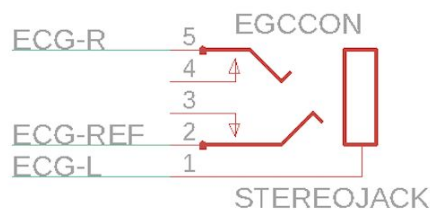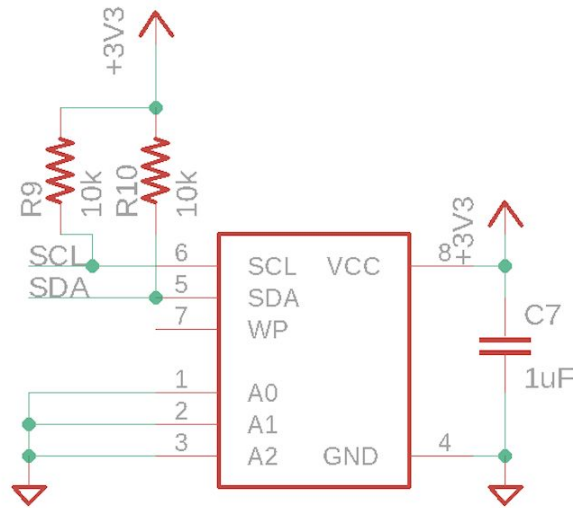
Figure 3.24: Ferromagnetic Random Access Memory Chip (FRAM) Schematic

Before the sensor algorithms were optimised, there was a need for extra off-board

memory in the form of an I2C FRAM chip. Ultimately it was not used. It is worth noting that if

we did indeed need to use it, in future revisions of the board we would simply choose to use a

different microcontroller rather than adding this component. In our case, we were limited by size

constraints on the PCB, and the Arduino Pro Micro fit within those constraints.
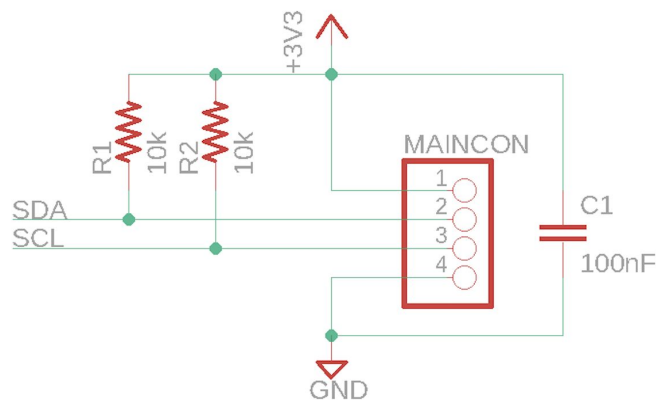
Figure 3.25: Temperature Sensor Buddy Board Connection. JST-PH Schematic

We decided to use a seperate board for the chosen temperature sensor. Due to its small size, the second board – called the Temperature Sensor Buddy Board – would break out the small pads on the temperature sensor and allow us to easily interface the sensor with the main board. We chose to use the smaller 2mm JST-PH connector family to interface the two boards as shown in the schematic. Capacitors on both ends bypass the 3V3 supplies. We chose the JST-PH because of its ubiquity, ability hand assemble and compact size that would not be too bulky on a patient.

Component Placement

Throughout this process, we used Eagle version 9.2.2 from Autodesk, as it provided a wide feature set that was easily available on the educational license. We decided to use a two-layer design to ease PCB design without exceeding the PCB budget. When placing components, we needed to keep several things in mind to minimise the size of the PCB while still maintaining flexibility in the design, ideal component placement for optimal functionality and ease of hand assembly. We considered:

- All off-board connectors go on the edge of the PCB
- Analog components and traces are separated from digital ones to improve noise response
- Analog signal trace lengths are minimised to reduce noise
- Four 3mm diameter mounting holes are needed for mechanical mounting of the the main board. Four 2mm mounting holes are needed on the Temperature Sensor Buddy Board.

- Bypass capacitors are placed next to the power connections of their respective component. Additional bypass capacitors are added in sparsely populated areas of the board

- 25mil thick (1mil = 0.001 inches) traces for power to provide higher currents with minimal resistive losses

- 16mil thick analog and digital signal traces for reduced noise and reduced ground coupling.

- Arduino breakout headers are positioned close to and in line with the Arduino interface headers to ease pin number recognition on the prototype

- I2C 10kΩ pullup resistor pair located next to the Arduino I2C pins and on the Temperature Sensor Buddy Board to reduce noise

- RF radio module connector is located in such a way as to ensure that the radio module stays within the dimensional bounds of the PCB

- LEDs positioned on the edge to allow for good visibility.

In addition to the above considerations, we also made the top and bottom layers of copper ground planes with buried vias connecting the two at regular intervals to reduce noise and ground coupling. For the silkscreen – the white layer of text covering the solder mask – we placed component numbers to allow for easy component recognition when hand assembling the PCB. Additionally, we placed the PCB name and version number on the back solder mask, as well as the name of the designer to ensure accountability and help people know who to ask in

case of questions. A white box is placed on top to allow for easy numbering of individual boards for our inventory and rework tracking system.
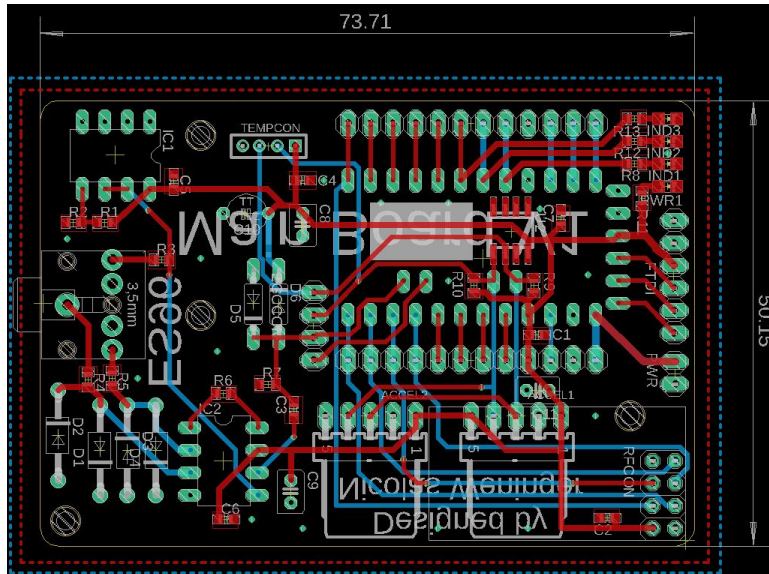
Figure 3.26: MainV1 Board Layout View. Units in mm.
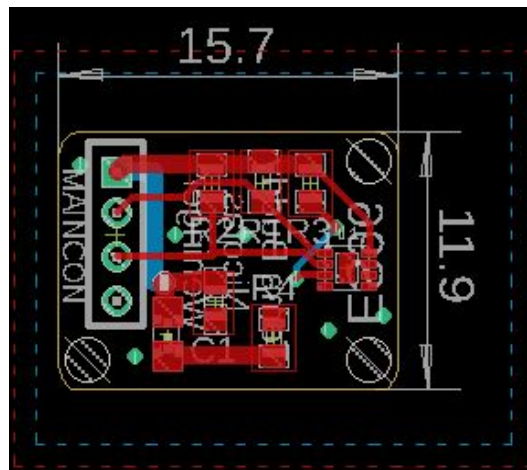
Figure 3.27: Temperature Sensor Buddy Board Layout View. Units in mm.

Manufacturing

Both PCBs were manufactured by a 3rd party manufacturer, as we did not have access to the requisite tools in-house. The order was placed just before thanksgiving break so that they

would arrive the week following the break. Before the order was placed, we preformed design rule checks on the schematic and board files of both PCBs to ensure that we were in compliance with the tolerances stated by the manufacturer.

Assembly and population of the boards with components was done in-house using a combination of equipment available the SEAS EE lab in Maxwell-Dworkin and Paul Horowitz and Jim McArthur's electronics lab in Cruft. Leaded solder paste and a fine tipped soldering iron was used to solder the SMD components. Following each SMD component being placed, we verified that the connection was as intended by checking for the resistance of the component. All capacitors had infinite resistance and all resistors had a resistance within ±1% of the marked value, as expected.

Through-hole components on the main board were soldered using leaded solder as well, as this makes it easier to conduct reworks. On the Temperature Sensor Buddy Board, unleaded solder was used, due to safety concerns with skin contact.
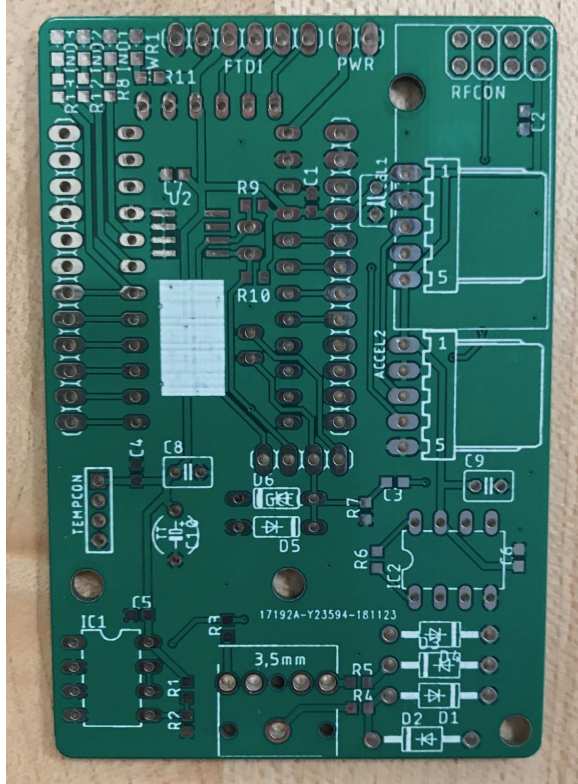
Figure 3.28: Unpopulated MainV1 PCB front
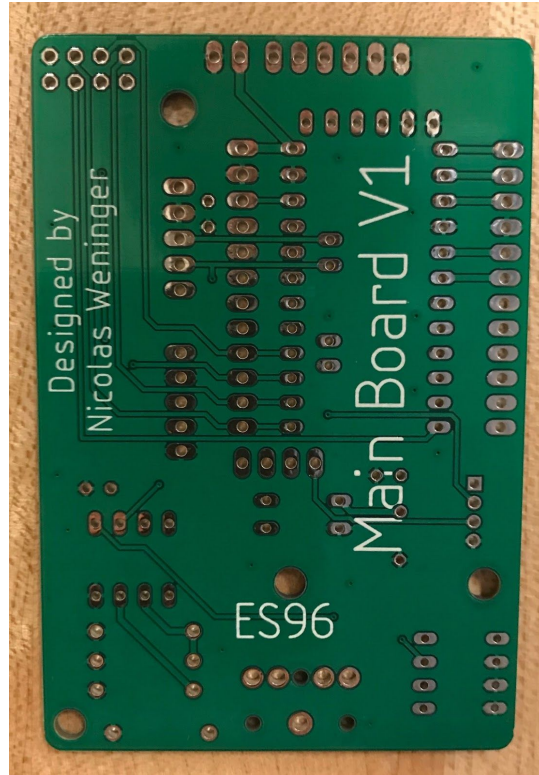
Figure 3.29: Unpopulated MainV1 PCB rear

The Temperature Sensor Buddy Board proved to be too challenging to assemble, as the footprint for the temperature sensor is designed for machine assembly rather than hand assembly. We were thus unfortunately unable to use the Temperature Sensor Buddy Board as planned. However, we did devise an alternate temporary solution, as detailed in the following section.
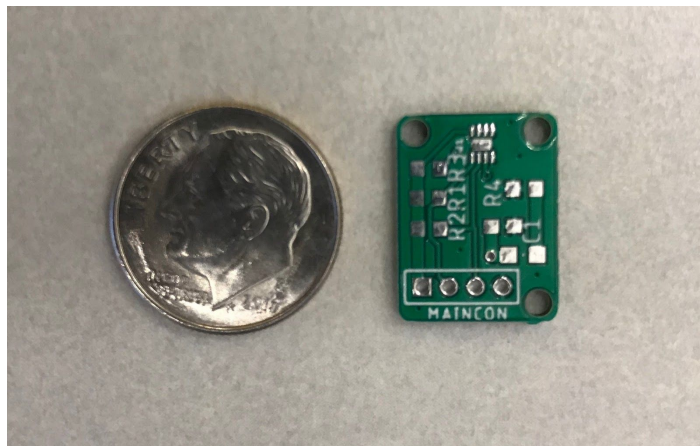
Figure 3.30: The unpopulated Temperature Sensor Buddy Board front, as compared to a nickel

Figure 3.31: The unpopulated Temperature Sensor Buddy Board rear

Required Reworks

Due to oversights in our design validation process and subsequent workarounds to these problems, a few components on the Main PCB ECG circuitry needed to be reworked, as discussed in an earlier section:

- Added a 10kΩ potentiometer to the former location of a now unused accelerometer connection. Connected 3V3 and GND

- Pin 3 of the op-amp was connected to the potentiometer wiper connection

- Pin 6 of the AD623N was connected to Arduino pin A0

- Pins 1 and 2 of the op-amp were shorted together

- R3 was not placed

- The ECG reference lead connection was connected to Pin 1 of the op-amp

- R7 of the high-pass filter was connected to pseudo-ground, not ground, via a different 560kΩ resistor.
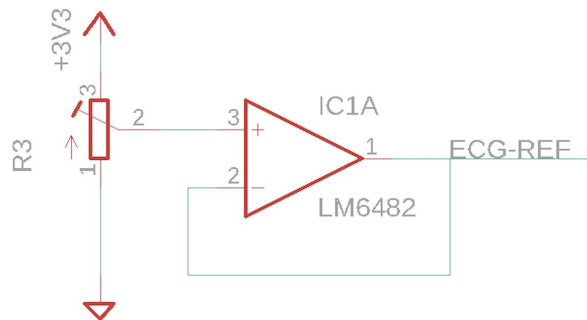
Figure 3.32: Reworked ECG Adjustment Schematic

In order to overcome the challenges faced with the Temperature Sensor Buddy Board in such a way to let us interface with the main PCB as desired, we decided to use a small protoboard combined with a breakout board for the temperature sensor that we had previously assembled and validated. This would allow us to wire the sensor to a 2.54mm connector that would be wired to the male JST-PH connector for the main PCB, as well as including the external circuitry required for proper operation of the sensor.
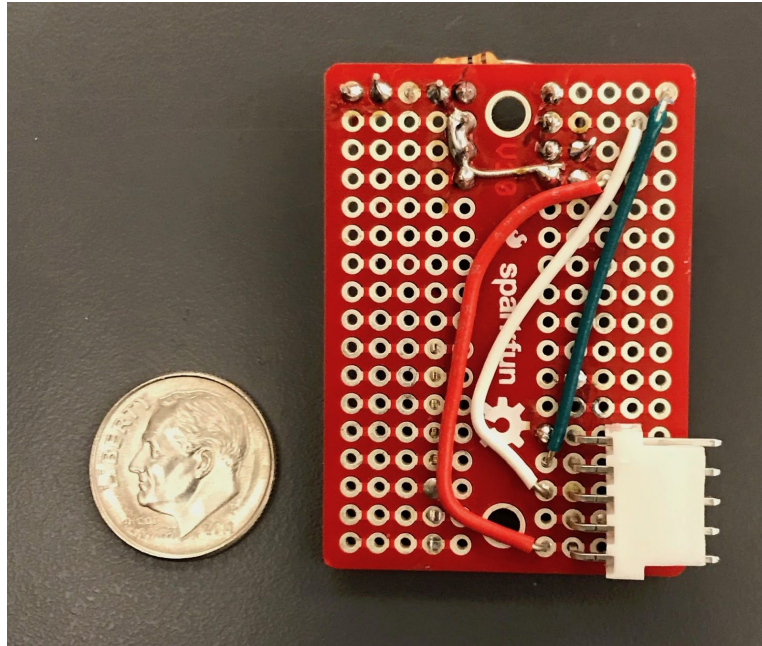
Figure 3.33: Temporary Temperature Sensor Buddy Board as compared to a nickel
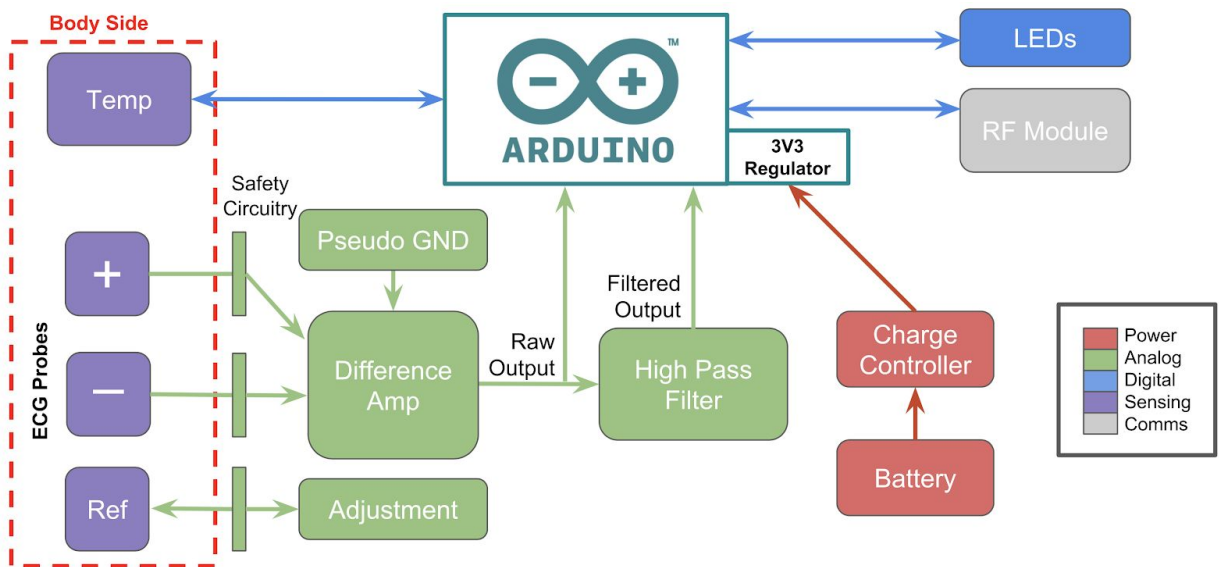
Figure 3.34: Final MainV1 PCB block diagram
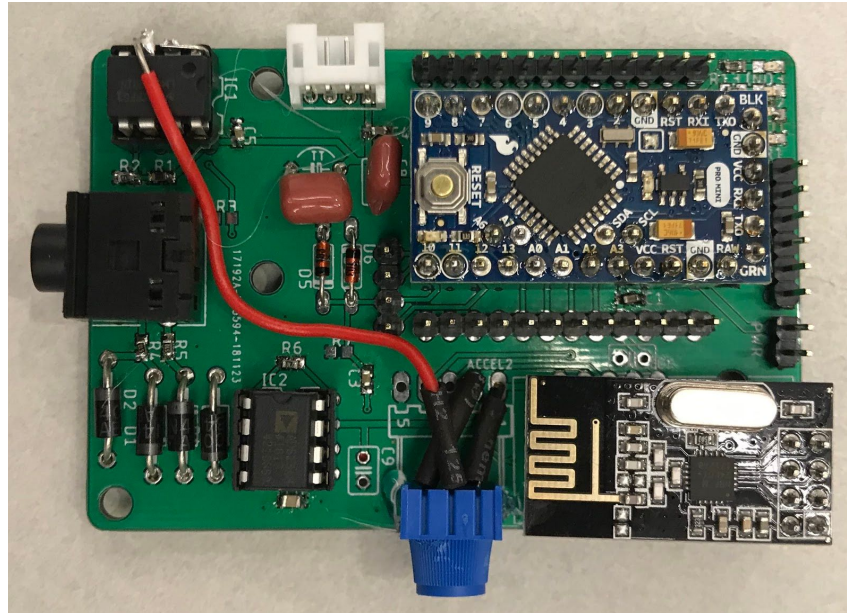
Figure 3.35: Populated MainV1 PCB front
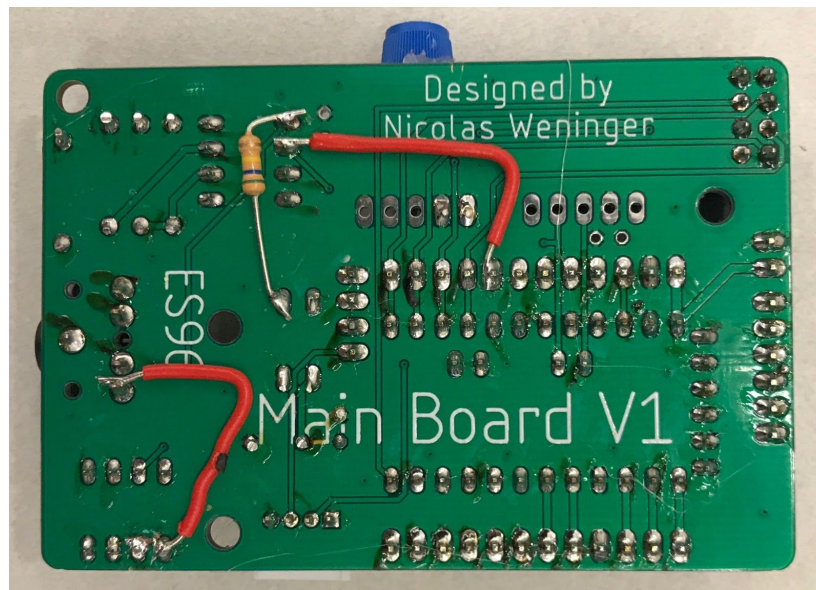
Figure 3.36: Populated Main V1 PCB rear

PCB Bringup and Validation

Once the PCBs were assembled, we conducted tests to see whether all subsystems were working correctly. The steps and results are outlined below. Between every step, power is turned off and on again to prevent issues that might arise with hotplugging. All but one subsystem

144

worked – the ECG analog processing. Unfortunately, we did not have time to fix this. Discussion is continued in the next steps chapter.

Instructions	Subsystem Under Test	Result	Observations
With nothing plugged into PCB, power up.	PCB assembly		No shorts and issues with power delivery.
Plug in Arduino	Arduino		3V3 regulator and electrical pin connections work.
Plug in ECG leads.	ECG analog processing		Components work but output is saturated. Likely issue with high-pass filter.
Turn on pins D2-4	LED Indicator LEDs		All LEDs illuminate
Plug in RF module	RF Module		Rx/Tx with hub confirmed.
Plug in temperature sensor	Temperature Sensor Buddy Board		Data from Buddy Board successfully read via I2C.

Table 3.6: Results of PCB bringup and validation

We also needed to validate whether the V1 prototype indeed allowed for 1 full day of operation on the chosen 2000mAh battery. We first conducted a preliminary theoretical lifetime analysis, looking at the current draw of the major components in the full circuit. This analysis indicated that the current battery would meet the specified 1 day lifetime.

Measurements of the "as built" hardware indicated that the current draw was higher than the preliminary analysis indicated. This is likely due to passive components such as voltage dividers and filters, and differences in specific component current draw within manufacturer stated tolerances. The measured current draw still comfortably meets the 1 day lifetime requirement.

Component	Average Current Draw (mA)	Comment
LEDs	20.00	Only 2 power LEDs on continuously
AD623N	0.55	
Arduino Pro Mini	4.40	No modifications to the Arduino. In active mode.
RF Module	12.30	
LM6482	1.20	
MAX30205	0.93	
Adafruit Charge Controller	2.00	MCP73833
TOTAL	39.38	
Expected Lifetime (h)	**50.79**	

Table 3.7: Theoretical device lifetime analysis

Battery Capacity (mAh)	2000
Max. Current Draw (mA)	55
Average Current Draw (mA)	41.5
Minimum Lifetime (hours)	**36.4**
Average Lifetime (hours)	**48.2**

Table 3.8: Measured current draw lifetime analysis

3.3 | Housing

Objective

The goal of the housing is to provide structural support and encasement of the vital monitoring device. The housing holds the circuitry and affixes sensors to intended locations of the patient's body. The housing must be sturdily constructed so as to ensure the safety of the components and the longevity of the device.

Design

We had many considerations when differentiating and ultimately choosing how to house the electronics of the device and to secure sensors. These criteria include breathability, wear time, and reliability, all of which contribute to our overall goal of creating something comfortable that a child can wear without feeling restricted. It's a burden for the child and his/her family to go through the fight against cancer; we want to provide as much support as possible by making the in-hospital processes more comfortable. We achieved this goal by evaluating different options, as follows (* indicating chosen approaches on the final design proposed in this report):

1. Chest band: elastic band that would stretch around the patient's torso with the circuit board, as well as sensors embedded

2. Armband: elastic band wrapped around the patient's upper arm with a pocket for the circuit board and ECG wires leading to electrodes on the chest

3. Clip-on case*: rectangular case enclosing the circuit board with ports for leads to sensors

4. Medical adhesive*: tape for adhering sensor to skin

5. Temperature sensor case: rectangular piece fitting around the sensor that reduces direct contact of non-essential parts of the temperature sensor with skin

Criteria	Threshold	Explanation
Dimensions Weight	85 mm x 60 mm x 42 mm 58g	Fit PCB and battery and charger
Durability	No cracks with drop from 5ft	Approximation from medical staff dropping device
Fluid resistance	Submersion of 1 ft depth - water, ethanol or saline solution	Approximation of device being soaked from sweat or cleaning with ethanol or mistakenly dropped in patient-cleansing basin
Heat protection	33 degrees Celsius	Typical skin temperature is 32.5 degrees Celsius. Increase comfort by minimizing difference in temperature.
Patient interface	2N minimum contact force 30 degree angle of inclination from vertical	Approximation from 180 pound person imposing a sheer force for 0.001 seconds when walking into patient at 5mph.

Table 3.9: Technical Specifications for Case

Criteria	Threshold	Explanation
Dimensions	1 mm x 3.5 mm x 3.5 mm	Fit temperature sensor
Comfort	Breathable (with RET score <= 13) Biocompatible	Standard for "good breathability" rating for fabrics Minimal skin tissue response after peeling adhesive
Durability	12 hours without replacement Surface energy <= 25 dyne/cm	Long enough to last nurse shift Sticks better if surface energy less than that of skin
Fluid resistance	Immersion with water or saline solution for 12 hours	Account for sweat on or cleansed patient skin, lasting nurse shift

Electrochemical safety	Conducts <= 5mA of electricity	US safety standard

Table 3.10: Technical Specifications for Adhesive

We used the feedback in the meeting to think about how we could prioritize patient comfort. We ideated possible solutions that allowed us to mount sensors on the patient without using a restrictive chest band. The sensor team had to select a different method of measuring respiration rate other than a respiratory transducer chest band, which we accommodated. We considered the following criteria in a Pugh chart that helped us identify what form factor we wanted to pursue: length of wear, breathability, range of motion, battery heating, integration to child's image, physical feel, reliability, safety, ease of integration of temp sensor (Table 3.11).

	Details, Clarification	Chest Adhesion	Armband	Clip
length of wear	How long can the format be worn without intervention from hospital staff; i.e. access to skin/clothes under attachment, need to reposition device, need to replace the batter	1	2	3
breathability	To what extent does the format allow air flow for perspiration	1	2	3
range of motion	To what extent does the format allow for child's natural range of motion	2	2	2
battery heating	To what extent does the format protect the child from battery heat	1	2	3
integration to child's image	How would it emotionally affects patient's percepion of self/appearance	2	2	3
physical feel	What kind of bulkiness or irritation would it cause?	1	2	3
reliability, safety	Will nurses be able to consistently use it on each patient and obtain good results?	2	2	2
ease of integration of temp sensor	How convenient will sensor placement or measurement be with this central case design?	1	3	1
		11	17	20

Table 3.11: Pugh chart outlining criteria considered when evaluating next design beyond chest band.

Our three options for mounting the electronics onto the patient included directly adhering to the chest, mounting in an armband, and in a case that would be clipped on the waist. We

decided to explore rough prototypes of both the armband and clip-on case, but ultimately, the waist-mounted box scored slightly better than the armband because it would be less visible and obtrusive to the patient wearing it.
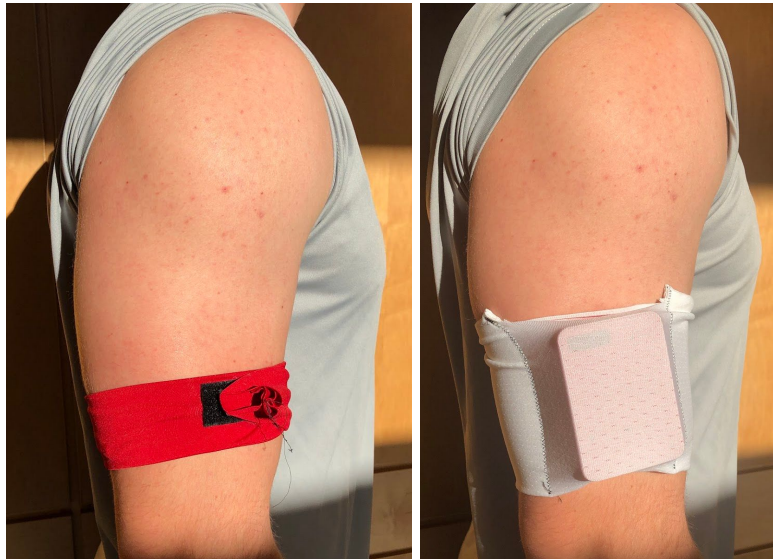
Fabrication: Armband

We came to the conclusion that it would be better if the technical elements of the device were housed in a centralized and safe location that could be secured to the patient. As mentioned before, through the Pugh chart, we explored two options for securing it on the patient, in the sleeve of an armband and attached via a clip, the former of which will be discussed in this section.

The armband would allow the patient to have normal range of motion while ensuring that the device and its elements are protected. For the first prototype, we cut breathable and flexible athletic material ("Dri Fit Active Wear Fabric") into a long narrow strip approximately 6 inches wide and 14 inches long. After folding the material in half, we sewed the open ends together, leaving a partial opening in the center for the device to be slipped into. In order to secure the piece around the arm, we used fabric glue to attach velcro to both ends.

While evaluating the first prototype, we determined that it was too narrow to comfortably fit the device case, and the adjustable velcro piece was not strong enough to feel secure and comfortable on the arm (Figure 3.37). For the second prototype iteration, we cut a similar athletic fabric piece twice as wide to allow for the entire device case to fit. We then selected a stronger velcro and sewed it onto the band instead of using fabric glue for a more secure fit (Fig. 3.38).

This more recent prototype proved to be a consistent and comfortable option for carrying the device case. We decided not to market the armband as the primary method of attaching the

device to the patient because the clip would present a more versatile use of diverse areas of placement. However we hypothesized that it would be beneficial to have more than one option for a medical care provider to choose from, in the event that a patient's treatment or personal preference dictates clothing choices that are incompatible with the device case and clip.

*Figure 3.37: Armband Prototype 1*   *Fig 3.38 Armband Prototype 2*

Design and Fabrication: On-patient unit

The sensors team established that we would need to house a custom-printed circuit board, battery, and power distribution module on the patient-side unit. We also have requirements to mount two sensors: an electrocardiogram (3 leads) and a temperature sensor. We decided that in order to prioritize patient comfort, we should only adhere the ECG leads and temperature sensor on the patient while leaving the electronics in an external unit not adhered to the patient. This would lower the size and weight of things that we would have to directly mount near the upper chest of the patient. We decided to package the electronics in an enclosed case. We propose that this case could be in a patient's pocket, clipped waist-side, or stored in an arm band. In general,

creating a small box with a clip would keep the electronics safe and allow the device to be placed where it is most comfortable for the patient.

Our final design, the clip-on electronics housing, went through a number of iterations before arriving to the final product. A photo of the variations can be seen below in Figure 3.45.
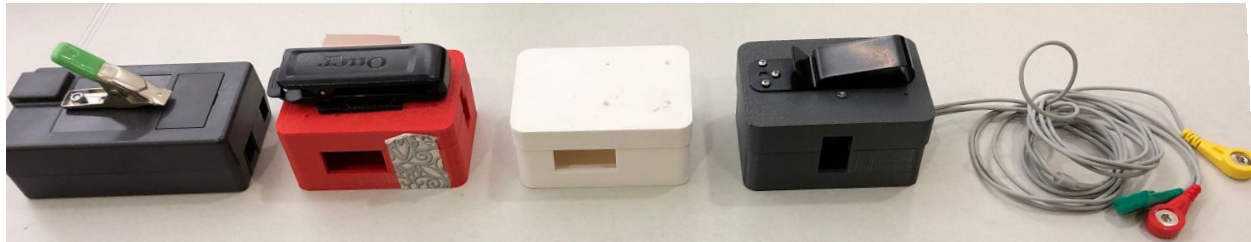
Figure 3.39: Design evolution

From left to right, the prototypes in the figure are numbered 1 through 4. To minimize the form factor of our system and maximize the user comfort, we housed the majority of electronics in a box, with a spring clip for attachment purposes (ideally their pants, gown, or pocket). This led to our first prototype on the left, prototype 1. Prototype 1 was made from commercial, off-the-shelf components we had available to us. Before the dimensions of the electronics components were frozen, we approximated the size of the electronics to be the same as an Arduino Uno. So, we modified an existing Arduino Uno case as a first prototype to get user experience feedback. We mounted a spring loaded clip to the Arduino case in order to make it easy to clip the case on a waistband or pocket. The process for mounting the clip included drilling holes into the clip and to the case and joining them with screws and nuts.

This model was used to get feedback from members of the team. They found that it was bulky and that the clip was not the right shape and stiffness to clip onto the patient. In our later steps, we designed the case in SolidWorks and 3D printed it in house. This enabled us to make changes quickly and have prototypes design and created in as little as 6 hours. In prototype 2, we

designed the case to match the dimensions of the printed circuit board and accommodate the necessary ports into the case. These ports included an auxiliary jack, charging port, window for LED and pin out access, port for a tuning the potentiometer, and port for the temperature sensor cable. The case is made of two parts that clasp together. In prototypes 2-4 we refined the dimensions of the parts so they mated together more tightly. In prototype 4, the sensors team needed to increase the size of the battery by a factor of 4, so we modified the case to accommodate these requirements. The figure shows several iterations of clips used in the design. We ordered a variety of belt clips and tested them to determine which one was easiest to mount to the case and most usable.

On-patient Unit Tests

Durability

In designing our device, we wanted to ensure that it would be durable against day to day drops. An example of this would be the device not breaking or cracking open if dropped from a certain height (such as a patient bed). This is an important feature to ensure compliances with the base case use of our device, outlined in our technical specifications (Table 3.9). The goal for our tests was to determine whether there was any separation of the lid from the body or fracture of the material when dropped from various heights.

Parameters of the experiment include various heights, as well as different orientations from which it was dropped. The impact surfaces that we alternated between were the large face, small face, medium face, and corner. Each impact surface was tested 3 times at 0.6 m (2 ft), 0.9

m (3 ft), and 1.5 m (5 ft). We selected an ideal height of 1.5m, because It is an approximation of the height at which a nurse will handle the device.
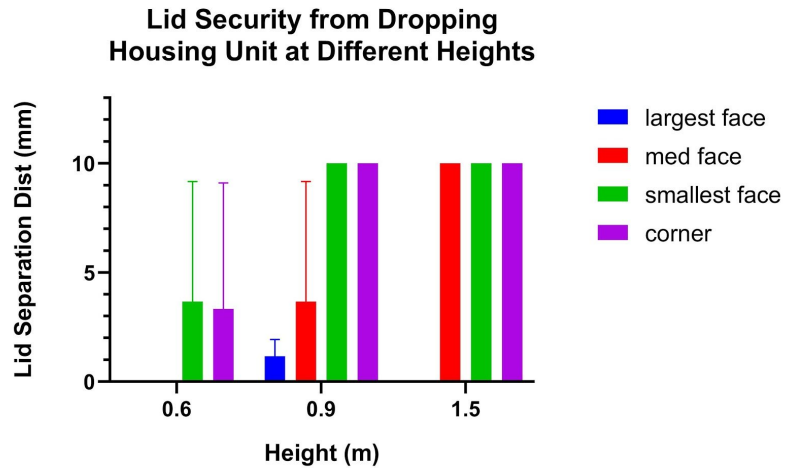
Figure 3.40: Measurement of distance between lid and rest of case when the closed case is dropped from various heights, where 10mm separation indicates full opening of the unit.

As shown in the Figure 3.46, in general, no drop at a height of 0.6 m caused the lid to completely separate from the frame. Increasing to 0.9 m, however, we saw that when dropped on the corner or smallest face, the lid did completely separate.It is apparent that a 1.5 m height drop will cause the box to separate entirely. Considering that a conservative approach would use the corner orientation as the standard, this suggests we need a tighter locking mechanism, whether via screws and nuts, a bolt, or a clipping mechanism.

Clip Weight Stretch Test

We conducted a clip stretch test to determine how much weight the different clips would be able to withstand when attached to the device case and clipped onto different materials. The two clips tested were a plastic clip and a metal clip, labeled otter clip and metal clip respectively.

This was an important test to make sure that the device housing unit would have a strong and reliable clip that was securely attached to the patient.

To build the testing apparatus, we attached equal size squares of canvas and athletic fabric to a wooden backboard. The two fabrics were chosen for their widely differing properties- the canvas was less flexible and more representative of a potential hospital gown while the athletic fabric resembled more casual, softer fabrics used for daily wear. We attached the device to the clip being tested and measured the distance each fabric stretched and deformed as well as how secure the clip was while carrying a specified mass. For every fabric and location we tested each clip with seven different masses up to 70 grams (more than the weight of the PCB and battery that would need to fit into the unit), and we conducted three iterations for each.
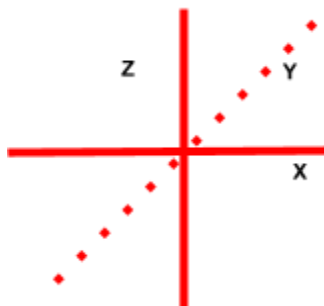
*Figure 3.41: Set up for testing clip performance as
well as the effect of different clips on the device case.*

Effect of Clip on Material Forward Stretch
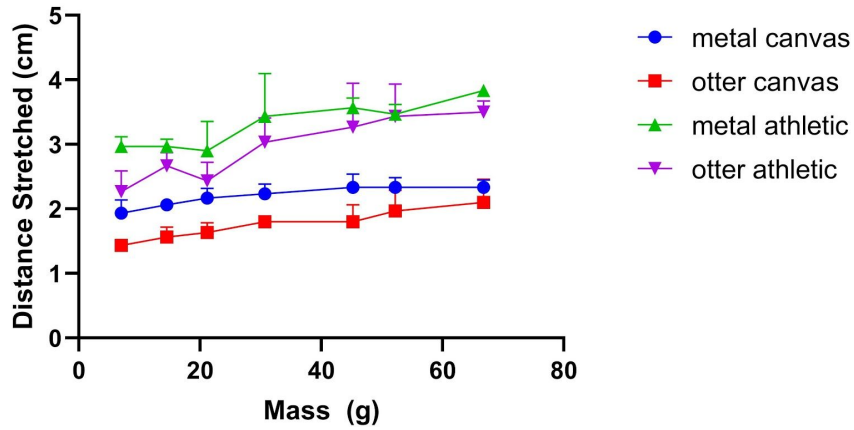
Fig 3.42: Measurement of y direction change- the distance between the hanging device and the wood backboard

With these results, we ran a T-test to determine if there was a statistically significant difference between how far the material stretched for each clip; ultimately, given a fabric, we found that there was no statistically significant difference (p>0.05) between how far the material was stretched and the type of clip (Figure 3.48).

Effect of Clip on Material Downward Stretch
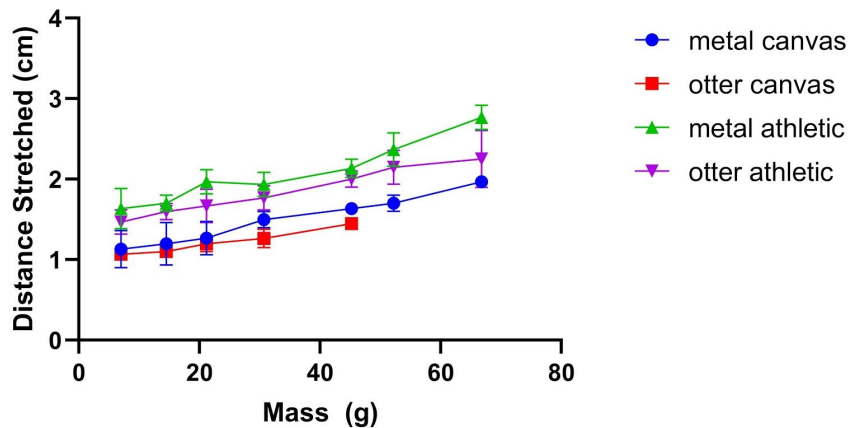
However, during the inverted trials, where we hung the device upside down at the bottom of the test pieces of fabric to imitate different ways the device could be clipped onto the patient, we measured how far the fabric is stretched downwards and we found that the "otter" plastic clip failed at 43-grams weight by completely slipping off the canvas fabric. The metal clip succeeded in holding all levels of mass loads without any visible deformation of the clip or fabric material.

These results along with the clip movement stability test suggested the metal clip as a good candidate to use on our final case prototype. In conclusion, we would not suggest using the device clipped upside down but both clips would hold for both stretchy and stiff fabrics, with more of a displacement on stretchier fabrics.

Clip Movement Stability

The purpose of this test was to analyze the security of the clip on the fabric. While previous test allowed us to determine the durability of the clip itself, this test aims to determine the level of functionality of the clip when attached to a patient's clothing and subjected to typical human movement. We would then be able to choose an optimal clip type that keeps the device case securely attached to the patient while they conduct normal day to day movements.

	X-axis	Y-axis	Z-axis
Otter athletic bottom			
Otter athletic top		x	x
Otter canvas bottom			

Otter canvas top		x	x
Metal athletic bottom	x	x	x
Metal athletic top	x	x	x
Metal canvas bottom	x	x	
Metal canvas top	x	x	x

Table 3.12: Results for clip staying secured (marked by x) while undergoing movement in three directions. Note that the X, Y, and Z axes are labeled in Figure 3.47, which depicts the set up for this test.

Using the same testing interface and positions as the Clip Weight Test, we tested the extent to which the clip-on device stayed affixed to the fabric. With the approximate device weight of 58g secured inside the device case, we moved the system side to side in the x, y, and z directions to simulate patient movement and possible device handling by the healthcare providers. We found that the otter clip failed (completely fell off of the fabrics) in all three directions for both fabrics when in the inverted orientation, as well as the z direction when on the top of the athletic test fabric (Table 3.12). Although the metal clip did fail in the z direction when in the inverted location on the canvas, it was significantly more secure than the plastic clip. Using this information as well as the previous clip weight stretch, we were able to conclude that metal clip would be a more consistent and reliable choice for our device.

Fabrication: Medical Adhesive

The goal of the medical adhesive is to keep sensors in place on the skin for accurate measurements. We chose four commercial products from 3M, selecting on the criteria listed, to consider and test for our final proposed device (Table 3.13).

	Breathability	Water Resistance	Wear Time (Days)	Backing
2477P	✓	✓	3	Elastomer double-sided
9834	✓	✓	14	Polyurethane
9917	✓	**X**	14	Nonwoven double-sided
9907HTW	✓	✓	14	Nonwoven

Table 3.13: Comparison of different properties between commercial 3M medical adhesive products.

The defining considerations made for these products include: breathability, to allow sweat to pass through; water resistance, to prevent reduced stickiness from fluids; wear time, which we want at least 12 hours for to match nurses' work shifts, and backing of the tape, for the least irritating against skin. We ultimately choose the 9907HTW for the temperature sensor based on two tests, as discussed below.

Comfort and Biocompatibility Test

We wanted to minimize skin response to the adhesive and thus evaluated each adhesive through a comfort and biocompatibility test. Four individuals wore the four different adhesives in two different locations, on the forearm (except for one individual whose samples were adhered to the back, due to sports involvement reasons) and under the arm (the intended location of the temperature sensor), for 12 hours. They then subjectively scored each adhesive out of 4 based on how red their skin tissue turned after peeling, as well as recorded how long it took for the skin to return to normal color (Figure 3.44).
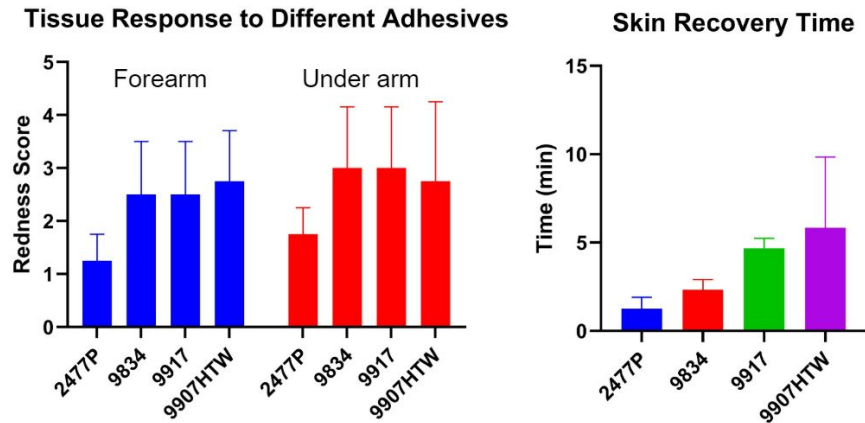
Figure 3.44: Comparison of tissue response scores between different adhesives on the forearm and under the arm and skin color recovery times on the forearm.

The adhesives 9834, 9917, and 9907HTW are not significantly different under analysis of variance (ANOVA) in redness scores and have similar skin recovery times (within 5 minutes of each other), suggesting adhesive 2477P as a potential candidate. However, 2477P was the least sticky, falling after 3-5 hours instead of lasting the entire 12 hours as did the other adhesives. Another design consideration was the pain of peeling off the adhesives for the child. Adhesives 9834 and 9917 were reported to be the most painful to peel off, which eliminated them from consideration. These results leave 9907HTW as the most viable adhesive to use for the temperature sensor.

Breathability Test

In the breathability test, we measured how much fluid was able to evaporate through the adhesive material over time. Here we had cups of 50g saline solution, to imitate sweat, with different adhesives covering the opening (Figure 3.50).
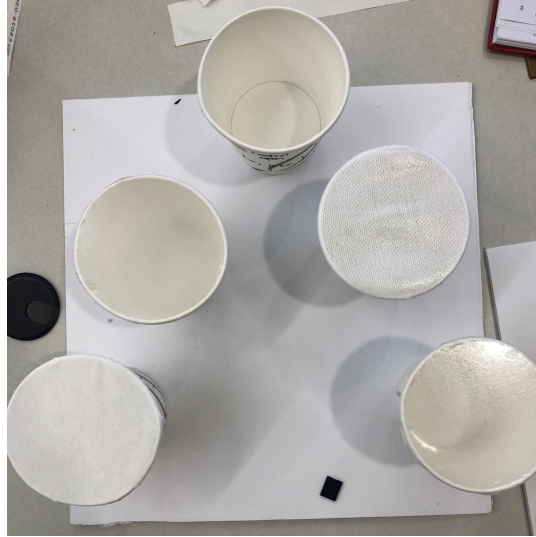
Figure 3.50: Overhead view of samples, materials from left to right: 9907HTW, 9834, control (no lid), 9917, and 2477P.

Cups were weighed after different timepoints up to 48 hours to determine how much solution

was still remaining with evaporation, as demonstrated by the graph in Figure 3.51.
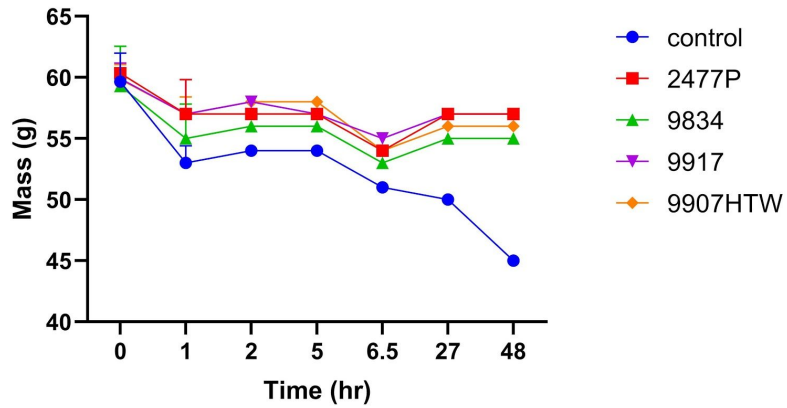
Figure 3.45: Line graph depicting change in mass of HBSS in each sample over time as a result of moisture vapor transmission through the material.

%-Change in Mass after 48 Hours Evaporation with Different Adhesives
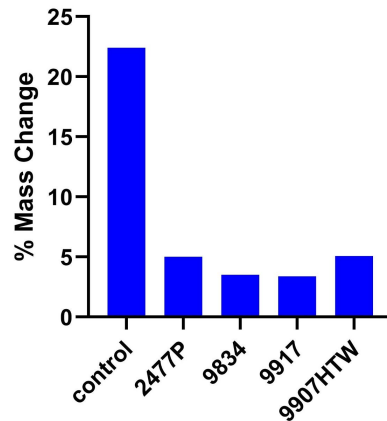
Fig 3.46: Bar graph comparing percentage difference in mass of samples after 48 hours due to solution evaporation between different adhesive materials.

The control sample had no lid and affirmed that evaporation was taking place, with the most percentage of mass change over 48 hours, shown in Figure 3.52. Our collected data also suggests that 9907HTW allowed higher or comparative vaporization rates and therefore more breathability. With these results we have decided that 9907HTW is the most viable adhesive to use. In the future, we can consider more products of similar properties that have even less redness response from the skin.

Fabrication: Temperature sensor case

The sensor that the sensors team is using needs to be placed directly onto the human skin to achieve accurate readings. This use case, however, exposes the sensor to a difficult environment. It would be in danger of corrosion via human sweating and would also be subject

to repeated sheer forces of repetitive adhesive application and removal. The team concluded that the construction of a temperature sensor housing would be beneficial to this use case, whereby the sensor would be waterproofed and resistant to wear and tear from the adhesive.

Due to the small form factor, we decided to 3D print the model with PVC, as we can obtain a high amount of precision with this technique. Additionally, the sensor would be built with 100% infill to ensure good strength of the housing. The model was designed on Solidworks and then printed in the Active Learning Labs. Two models were designed and built, as shown in Figure 3.53 and Figure 3.54.

Figure 3.47: Photo frame design

Figure 3.48: Roof design

In both cases, these models were not only designed to be a tight enough fit for the temperature sensor to ensure security but also loose enough fit so as to insert and remove the sensor without too much difficulty. In order to achieve this balance, a number of different models with different tolerances on the dimensions were printed. The models were then tested, by inserting the sensors and qualitatively determining fit. The variation of each model with the best result was chosen.

Although a final prototype for the housing has not yet been decided on, each of our models exhibit different advantages. The photo frame design is extremely simple and easy to

manufacture. It is easy for the user to insert the sensor, by simply clipping the sensor into the frame. The roof design is advantageous as it only exposes the bare minimum of the sensor, to the skin, in order to achieve accurate readings whilst minimizing sweat exposure. In this case, the sensor is slid into the housing. Ultimately, both sensors achieve our design needs to some extent depending on which model is chosen, but moving forward, further design iterations will be needed in order to reach an optimal solution.

Conclusion

These design considerations of different components helped us arrive to our final product, which can be seen below in the figures below.
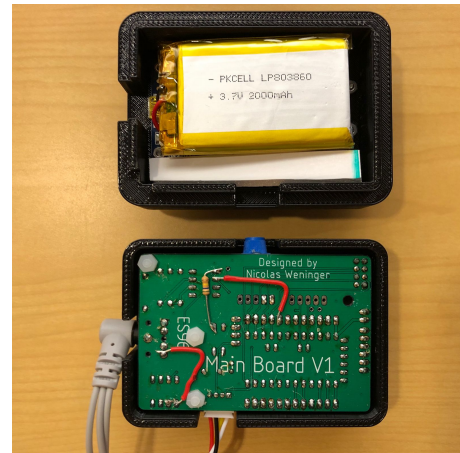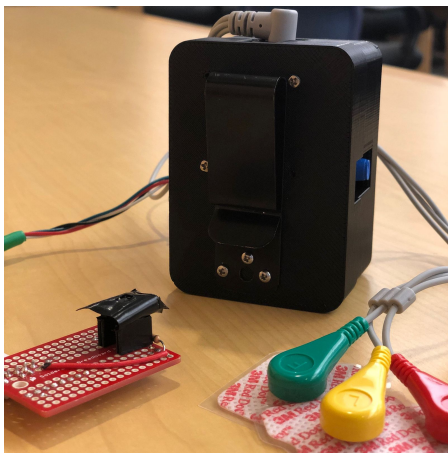
Figure 3.49(a): Final product view with PCB and ECG electrode connections.

Figure 3.49(b): Lid and body, top view
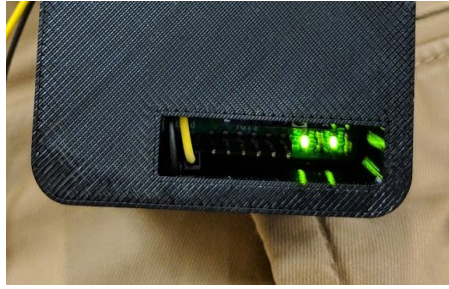
Figure 3.49(c): LED Alarm cutout

Figure 3.49(d): Housing Unit Clippes onto Individual

The circuitry housing is made of ABS plastic, with a final weight of 267g and dimensions of 85 mm x 60 mm x 42 mm. Figure 3.58 shows the final product clipped to a patient.

In our next steps, the box can be made smaller, so it would be less obtrusive for the patient. The box can be made smaller by redesigning the printed circuit board and minimizing battery size by focusing on minimizing the power consumption of the device. In a product-ready version of this device, we would manufacture the case via injection molding, which could ensure that the case stays closed during normal use and when dropped from up to five feet by using a simple locking mechanism. Additionally, we should make the case water resistant, which would require sealing the sensor and battery ports as well as covering the LED display hole with a clear acrylic.

3.4 | Data Communication

Introduction

After having decided on the project the team would pursue in the create phase, we realized that we would need a dedicated sub-team to determine the optimal means of transferring

the patient data we were planning to collect with our on-patient devices to the proposed central

device with a display that the healthcare staff could easily monitor to see when a patient's vitals

move out of the healthy range (the ***hub***).

Main Goals of Data Communication Sub-Team

This team would be in charge of determining and designing the optimal means of data

transfer from patient to hub. We made these decisions based on the following criteria:

- No requirements of existing hospital infrastructure.

- Minimal addition to nurse workflow.

- Minimizing the alarm fatigue that can result from false-positives.

- Minimizing the number of points of weakness in the communication system.

- Allowing for quick updates of the patient database to prevent negative trends from
  being missed.

- Allowing for new patients to be easily integrated into the system, and for the easy
  removal of patients who have been discharged.

- Minimizing the processing power requirement and battery drain on the on-patient
  device.

Design - Hardware

Design Considerations

In choosing a technology to transmit our data, we quickly eliminated a wired data

communication solution, due to the added infrastructure, greater need for system care, and

general desire to reduce the number of components that nurses had to interface with. We settled

on a wireless solution, and debated among three main wireless technologies; Bluetooth, local WiFi, and radio frequency (RF) transmission. Our analysis is outlined in the following table, where green indicates that the device would satisfy that condition well, yellow indicated the device would meet the specifications but with reservations, and red indicated that the technology poorly matched our needs in that category:

	Transmission technology	Range	Cost	Infrastructure Complexity	Coding and library complexity	Battery Usage
	Bluetooth					
	2.4/5 GHz WiFi					
	Radio Frequency (RF)					

Table 3.14: Table showing analysis of possible wireless options for the communication device

Our specific thresholds were:

- Excellent signal transfer (that is; no dropped packets) over a range of at least 50 feet

- Minimal cost per device, ideally less than $5 per device

- Minimal infrastructure complexity, specifically on the hardware components

- Existence of pre-existing software protocols, implementations, and usable libraries

- Minimal battery drain of at least less than ~50mA

Following this analysis, we decided to pursue RF communication as our means of data transfer, and all further design considerations were made with this decision in mind.

We decided to begin our RF testing by using two Arduinos connected to RF chips to interface with each other, since the Arduino IDE has a simple Serial Monitor built in to it, by which we could easily print the data being sent. The system would be the basis for having many on-patient devices, which we henceforth refer to as the **_child device_**, and one central hub device, henceforth called the **_parent device_**. Our plan was to transition to using a Raspberry Pi 3 B+ as the parent device, since the hub would already have a Pi acting as a mini-computer, and the Pi has the same input/output (I/O) pins that the Arduino has. Furthermore, having the parent RF chip directly on the Pi would improve the processing time, and reduce the coding complexity as the Pi can be directly interfaced with Python script, which the Hub team was already using to code the user interface.

With respect to the hardware, we chose the NRF24L01+ RF chip because of its easy I/O interface. The chips were sold at $1.20 per device and came with breakout pins soldered on to the board, allowing for easier testing. For our initial prototype, we used simple jumper cables to connect the I/O pins of the NRF24L01+ to the corresponding I/O pins of the Arduino board, with the intention that for a final design, it would incorporate a breakout board connection, allowing for easy replacement of any malfunctioning parts.

Despite our initial intentions to transition from using an Arduino as the parent device to having the RF chip directly interface with the Pi, due to many hours of debugging issues, we decided to have an Arduino connected via serial USB to the Pi, and this would have the RF chip attached since we were successful in getting the Arduinos to communicate with the RF chips. We were unable to determine the nature of this problem, and thus this workaround was determined to be the best course of action, to still obtain a working prototype. The added demand on the Pi was

insignificant, as we implemented a Python script which interfaced with the Serial Monitor of the Arduino IDE, parsed the data that was printed and then inserted it into the database that the Hub team created.

Our initial prototyping was done using Arduino Unos, but once we had successfully created, tested and debugged all the software protocols, we transitioned to using an Arduino Micro board for the child device implementation, as this could execute the same functions that we needed and occupied a smaller form factor. The parent device maintained the same set-up with the Raspberry Pi 3 B+ connected via serial USB to an Arduino Uno which had the RF chip attached.

<u>Final Hardware Components</u>

- 1 Raspberry Pi 3 B+

- 1 Arduino Uno

- 1 Arduino Micro per child device

- Jumper cables to connect I/O pins

<u>Initial Test Results</u>

We carried out range testing on the NRF24L01+ to determine how well it satisfied our range requirements and to see what obstacles we encountered as it pertains to physical obstructions. The following table shows the results of the test, where green indicates no dropped data packets, yellow indicates frequent dropped data packets, and red indicates complete signal loss. This data was obtained by having one Arduino powering an RF chip which continually tried to send an iterating counter located in a central location, while a second Arduino connected to a laptop as a Serial Monitor was moved around and monitored for any missed counters.

Obstacles	Number of Obstacles	Distance	Low Power	High Power
Open air	-	10 feet	Green	Green
Open air	-	50 feet	Green	Green
Concrete wall	One wall	5 feet	Green	Green
Concrete wall	Two walls	15 feet	Green	Green
Concrete wall	Three walls	50 feet	Red	Green
Glass	Glass door	10 feet	Green	Green
Glass	Two glass walls	15 feet	Yellow	Green
Glass	Two glass walls	20 feet	Yellow	Green
Glass	Two glass walls	25 feet	Yellow	Yellow
Glass	Glass wall & glass door	25 feet	Red	Yellow
One floor above	Directly above hub	40 feet	Yellow	Yellow
One floor above	Not above hub	50 feet	Red	Red

Table 3.15: Table showing the results of range testing on NRF24L01+ RF chip

The data led us to the following conclusions:

- Increasing the power supply to the RF chip increased the signal transmission strength. This however was decided against, due to the significantly increased battery drain of the high-power mode

- The signal transmission through open air is excellent, and ideal for our system

- Glass quickly becomes an obstacle to signal transmission over short distances. We determined that this was not an obstacle to choosing the NRF24L01+ however, as we did not forsee large glass walls or doors being prevalent within a single pediatric cancer ward.

- The signal was transmitted well through concrete, except over long distances. Similar to our conclusion for glass, we determined that this was not an obstacle to choosing the chip, as concrete would not be prevalent within a single pediatric cancer ward.

- The signal transmission through multiple floors of a building was poor. We determined that this would not be an obstacle to choosing the chip as most pediatric cancer wards should only be one floor, and if there happens to be multiple floors, the system would simply require one parent device per floor.

<u>Design - Software Protocols</u>

<u>Design Considerations</u>

After having chosen the NRF24L01+ chip for our project, we then analyzed the libraries available for this chip, and decided to use the RF24 library provided by TMRH20 as an open source library, due to its availability both as a Python library and on the Arduino IDE library interface. We began by studying online documentary outlining the hardware and software protocols for using the chip, to determine how best to set up the infrastructure for having one central device interfacing with many devices, as this was integral to the proper functioning of our system. To begin, we analyzed the set-up protocols of the chip.

```
void setup() {

    // setup serial communications for basic program display
    Serial.begin(9600);
    Serial.println("[*][*][*] Beginning nRF24L01+ ack-payload child device program [*][*][*]");

    // ----------------------------- RADIO SETUP CONFIGURATION AND SETTINGS -----------------------------//
    radio.begin();

    // set power level of the radio
    radio.setPALevel(RF24_PA_LOW);

    // set RF datarate
    radio.setDataRate(RF24_250KBPS);

    // set radio channel to use - ensure it matches the target host
    radio.setChannel(0x76);

    // open a reading pipe on the chosen address for selected node
    radio.openReadingPipe(1, nodeAddress);

    // enable ack payload - slave reply with data using this feature
    radio.enableAckPayload();

    // preload the payload with initial data - sent after an incoming message is read
    radio.writeAckPayload(1, &sensorData.asbytes, sizeof(sensorData.asbytes));

    // print radio config details to console
    printf_begin();
    radio.printDetails();

    // start listening on radio
    radio.startListening();
    // -----------------------------------------------------------------------------------//
}
```

Figure 3.50: Set-up code used for NRF24L01+ on child device (Arduino)

The chip had some variables which were modifiable, including:

- Power level (with preset MIN, LOW, MED, HIGH and MAX thresholds) which alters the strength of the RF signal the chip produces and thus the distance it can travel without signal loss. It is important to note that this also affects battery life.

- Data Rate, or the speed of data transmission. Similarly, this improves the accuracy of data transfer, reducing dropped packets, but also increases the battery drain.

- Radio channel, the frequency channel at which the RF chip would operate, which can range from 0 to 127. For chips to communicate, they must operate on the same channel, and 76 was chosen as a less commonly used channel to minimize any interference.

In addition to these variables, the chip functions on a pipe structure, that is; each chip can at any point in time have one reading pipe and one writing pipe opened, independently of each other. These pipes are simple addresses which serve as a gate; a chip will only open communication channels with another chip if they both are seeking communication on the same pipe address. A chip which is seeking to send a message, must have the same writing pipe address open as the reading pipe of a chip looking for a message. Using the openReadingPipe and openWritingPipe commands built into the RF24 library, it is possible to alternate between open pipes, a feature we determined to be critical for the functioning of the system we had envisioned as the parent device needed to be able to communicate with all the child devices independently.

Finally, to understand how the chip sends its information, we read the online protocol of the payload structure. We determined that there was a payload limit of 32 bytes in a single data

transmission, which meant that we would have to pre-determine the order and structure of the data we sent since it would be difficult to also send identifiers for the data.

Figure 3.51: Enhanced ShockBurst packet structure for NRF24L01+. (Nordic Semiconductors NRF24L01P Product Specifications)

Figure 3.52: Packet Control Field structure for NRF24L01+. (Nordic Semiconductors (NRF24L01P Product Specifications)

The NRF24L01+ uses an Enhanced ShockBurst packet structure, which is an upgrade over most existing RF chip data structures, as it contains the Packet Control Field (PCF) which allows for the variation of payloads, and sending of unique packet IDs.

Design concept

Handshaking: We determined that the dynamic nature of the data structure was essential for the functions we needed, specifically the packet ID that is sent with each payload, allowing the parent device to quickly recognize if a data set received is new or whether it is a retransmission. We deduced that this would make the acknowledgement process of the data transfer easier, and reduce the instance of dropped packets, as if there is an unsuccessful transmission, the child device will save the data to be sent later, and the parent device will be able to send a

confirmation back when it has received a new payload. The handshake protocol became essential not just for the default communication acknowledgement, but also the de/commission protocol, as it was critically important to know if the set/reset process had occurred successfully.

De/Commissioning protocol: We determined that it would be logistically difficult to ship every child device with a unique identifier which remains fixed for the duration of the device, as this adds burden to the nurses of having to keep track of each device and assign them sequentially. For this reason, we chose to create our own protocol which would assign a child device a unique ID only when a nurse needs to use it, and this would allow the parent device to sort the patient data accurately. This system would be significantly more modular, allowing the nurses to pick up the nearest device and assign it to a patient as opposed to having to search for the next sequential device to be assigned. We also determined that an essential part of this protocol would have to be the wiping of a device when a patient is discharged from care on the ward, for both patient data confidentiality.

Age and alarm state data structure: We determined that the parent device would have to consistently send the age to the child device, since it was considerably more difficult to send during the commissioning due to the instantaneous and sporadic nature of the commissioning protocol. Since the default communication would always involve sending the patient's age and alarm state (whether the child device should be showing an alarm LED), we chose to make a custom union of structs, to minimize the RAM usage of the Arduino and improve the speed of data transfer due to only using the minimum number of required bytes. We also made the decision to send the age in months, to avoid having to send floating point values for the number of years old a patient is, thus making calculations more accurate.

```
// struct union to collect/send age & alarm state
typedef union writeArrayT {
  struct {
    unsigned long age;
    bool alarmstate;
  } asStruct;
  byte asbytes[5];
} writeArrayT;

// global write array
writeArrayT writeArray;
```

Figure 3.53: Sample code showing age and alarm state initialization

Timestamp and vitals data structure: Similar to our conclusions regarding the data structure of

the age and alarm state, we determined that to minimize RAM usage and maximize data transfer

speed, we needed to make a custom array for the sensor data on both the parent and child

devices.

```
// struct union to collect/send sensor data
typedef union sensorarray{
  struct {
    unsigned long stime;
    float HR;
    float RR;
    float temp;
  } asvalues;
  byte asbytes[16];
} sensorarray;

// global sensor data array
sensorarray sensorData;
```

Figure 3.54: Sample code showing timestamp and sensor data initialization

Details of Protocols:

For both the parent and child devices, we determined that their normal operations would

be based on the same protocols, albeit with different implementations applicable to either side.

The flow diagram below shows the general operational layout of the two devices.
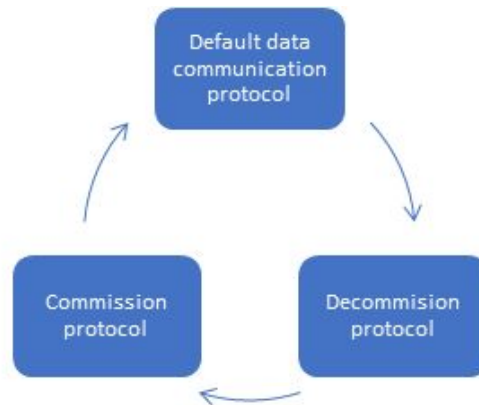
FIgure 3.55: Flow diagram showing the protocol relationships for both parent and child devices

Parent Device

The parent device needed to be able to execute the same basic functions as the child, that

is, default communication, commission and decommission. The default communication

procedure is somewhat different due to the parent device having to obtain data from multiple

devices. The de/commission procedure is different as the parent device has to send the

information for the child device to interpret as activating the de/commission protocol.

Default Data Communication

The parent device will have a pre-built array of all the possible pipe addresses (listed in

order as hexadecimal characters from 0x01 to 0xff, allowing for a maximum of 255 devices to be

commissioned at any point in time, a number we deduced to be reasonable, but also upgradeable if we saw the need to scale up). The parent device will have a second parallel Boolean array, which holds Boolean values that indicate whether there is a child device with the same pipe index in the Boolean array corresponding to the pipe array.

The parent device will iterate through the Boolean array, and for every pipe address that it detects an existing child device at (indicated by a TRUE Boolean), it will open that reading and writing pipe and attempt to communicate with the child device to obtain data. This is simplified in Figure 3.65 below. If the parent device does not successfully connect to the child device, it will scan again for all devices every 10 seconds, an interval which we determined to be suitable for the parent device as it should run often enough that it does not miss any data packets from the child device. This procedure allows us to minimize the run-time of our data collection and thus lets us increase the speed at which nurses will receive an alert if a patient's vitals are out of the ordinary. It also does not add to nurse burden, since the child device will still send data only every 10 minutes, reducing the risk of alarm fatigue.

Furthermore, once the parent device has received the data successfully from the child device, it saves the information and executes its alarm threshold algorithms, from which it finds an alarm state (whether the on-patient device should be displaying an alarm). It then packages the patient's age in months and alarm state and writes it to the child device.

Device Decommission

When the nurse clicks the decommission button after choosing which patient's device to decommission, the parent device checks its database to find the pipe address of that device, and opens that address as its reading and writing pipes. It then writes a unique decommission

identifier, a compound of an age of 0 and alarm state of 1, which we chose to not conflict with our target population age of a few months to adolescence. The child device would interpret this as an overwrite command and then revert to the zero pipe as its pipe address. Finally, the parent device will open the zero pipe and search for a success marker from the child; once this is obtained, the parent device changes the Boolean in the Boolean array for the pipe that the device previously occupied to FALSE, thus that pipe is no longer scanned for during default communication, and is free for recommissioning.

Device Commission

When the nurse clicks the commission button after entering the patient's details, the database will scan the Boolean list to see which the lowest available pipe address is. The parent device then opens the 'zero pipe' 0x00 to look for a child device to commission and it attempts to send the pipe address to be assigned. The parent then opens the reading and writing pipe for that new device that it attempted to assign and looks for a response from the newly assigned device with a success identifier (0xF1). If it obtains this identifier, it changes the corresponding Boolean in the Boolean array to TRUE, thus adding that device to the list of pipes that will be scanned for in the default operations, and sending the pipe address to the hub database so the device is permanently assigned to that patient until decommissioning. If the assignment was unsuccessful, it will repeat it for 15 seconds (an interval we determined to be suitable for connection, as if it takes longer that 15 seconds then the device might be malfunctioning and should be switched out or charged), and if after this time the rewrite is still unsuccessful, the parent device will abort the commissioning procedure and return an error message.

```
// defining pipe number array length
#define node_no 20

// Defining all pipes
char pipeAddresses[node_no] = {0x01,0x02,0x03,0x04,0x05,0x06,0x07,0x08,0x09,0xA,0xB,0xC,0xD,0xE,0xF,0x10,0x11,0x12,0x13};

//Defining bool array of available pipes
bool pipeAvailable[node_no] = {1,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0};

// Defining zero pipe
char zeroPipe = 0x00;
```

Figure 3.56: Sample Arduino code showing pipe and boolean array declaration for parent device

Note that from Figure 3.65, we can see that the only pipe currently assigned is pipe 1, and the rest are pre-set as unassigned. The example is made with 20 preset pipes, but in the final implementation there will be 255. Due to unusual difficulties in achieving successful communication with 0x01, we decided to always leave that as 'occupied', and never assign it to a child device, thus eliminating the complication. If a new child device is commissioned, the parent protocol will detect that the next free pipe is pipe address 2, or 0x02, and send this to the child device so that it can save it as its new pipe. It will then change the second Boolean value to 1, and the device will scan for data from 0x02 in its default communication until 0x02 is decommissioned. Finally, Figure 3.65 also shows the initialization of the zero pipe that the de/commissioning protocol will use.

Child Device:

We determined that there would be three core functions on our child device as explained in Figure 3.67. The child device's protocols differ from the parent device's, as it has to decode the information sent by the parent, interpret commission/decommission instructions, and send data by default.

Default Data Communication

Once the device is activated, it will continually collect data from the sensor processing algorithms, package it along with a timestamp, write it into the payload and then attempt to send the data to the parent for a preset amount of time. This would happen once every 10 minutes to minimize battery drain and alarm fatigue for the nurses. In response, the child device expects to receive an array from the parent consisting of the patient's age and the alarm state upon every successful connection. The child device continually checks the alarm state to determine if it should be signaling an alarm for the nurses.

Device Decommission

This will only be activated by a nurse-controlled press of a button on the parent device and is designed to wipe the child device when the patient is discharged. This will allow the nurse to select which patient's device they wish to disconnect by their name or hospital ID number, and the database on the parent device will determine which child device that corresponds to. The parent will then send a unique identifier to the desired child device, which the child device will interpret as a signal to cease all sensor functions and erase previous data, resetting its reading and writing pipes to the 'zero pipe' 0x00, a predefined value.

Device Commission

This will only be activated by a nurse-controlled press of a button on the parent device and is designed to activate a child device by sending it the next available pipe address, which it saves as its new reading and writing pipe, and the parent device will register that the connection has been made. This will be achieved by the parent device scanning for a child device on the

predefined 'zero pipe', which will be the address of only a decommissioned device waiting to be assigned. After this, the child device opens the new reading and writing pipe, and writes a success identifier to the parent device, to confirm that the handshake has been accepted and the device has been successfully assigned. Once the commission procedure is successfully executed, the device will carry out its sensor and data communication functions as outlined in the default communication, until it is again decommissioned.

Design Iteration and Client Feedback

When the idea was pitched to the client, it was well received, and thus we continued pursuing RF as our means of communication. We did add a few features to our system based on feedback from other subteams and the client, and these included:

- Printing of the data in an easier format for the Hub team's python script to parse.

- Minimizing the RAM requirement on the child device due to RAM limitations on the Arduino Micro by writing certain permanent values to the EE-PROM memory (permanent memory).

- Ensuring that for the child device, the commission and decommission code were gated by if-loops so that they only ran when needed, to optimize RAM usage and battery life.

- Ensuring that for the parent device, the commission and decommission code were also gated by if-loops that would only be triggered when a button is clicked on the HTML script written by the Hub team.

There were a few design considerations raised by the other subteams which we decided too complex to implement in the given time frame, but useful advice for future iterations:

- Implementing a system which would recognize if the child device was unable to send its data set, and hold on to the data set with a unique timestamp and ID until it is successfully sent. This also included adjusting the parent device to be able to parse the unique timestamp and ID and successfully insert the data set into the correct slot in the database.

- Using a custom library for the RF chip based on its baseline protocols, thus minimizing memory use on the Arduino

- Data storage and transmission encryption would be necessary before commercial applications of the product, due to patient confidentiality concerns.

Final Communication Device Prototype
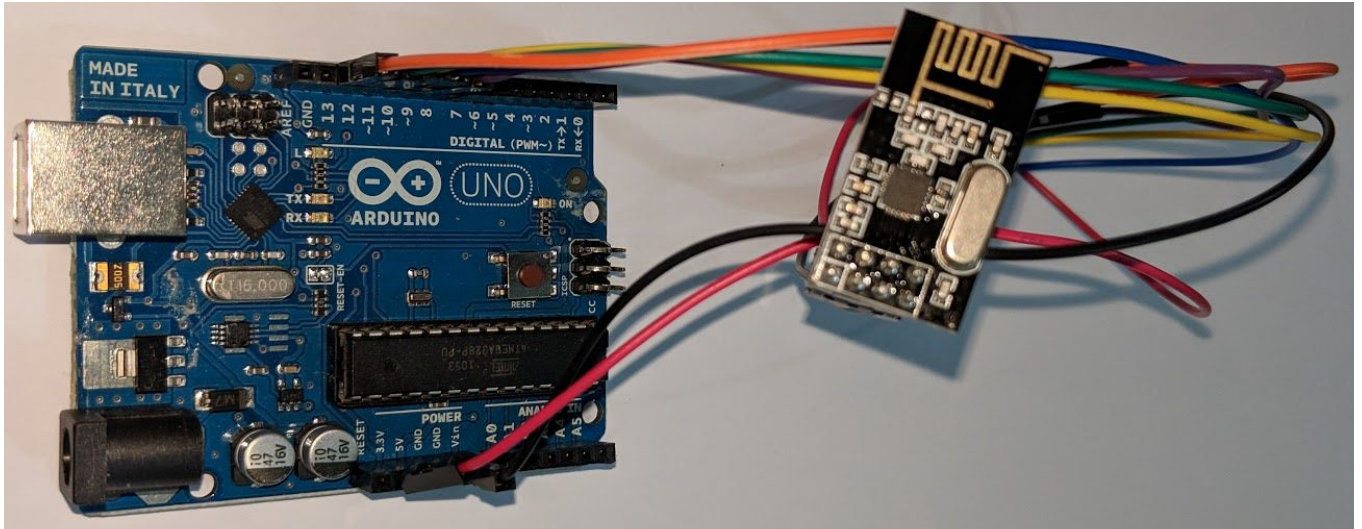
Figure 3.57: Picture of child device debugging prototype with Arduino Uno connected to NRF24L01+
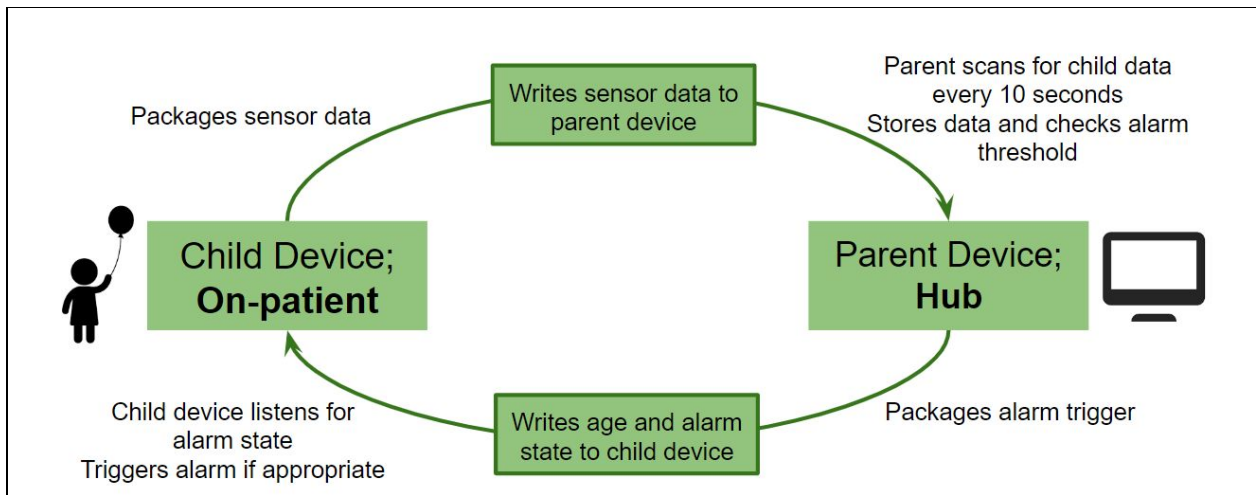
FIgure 3.58: Cyclic flow chart showing communication procedure between parent and child devices
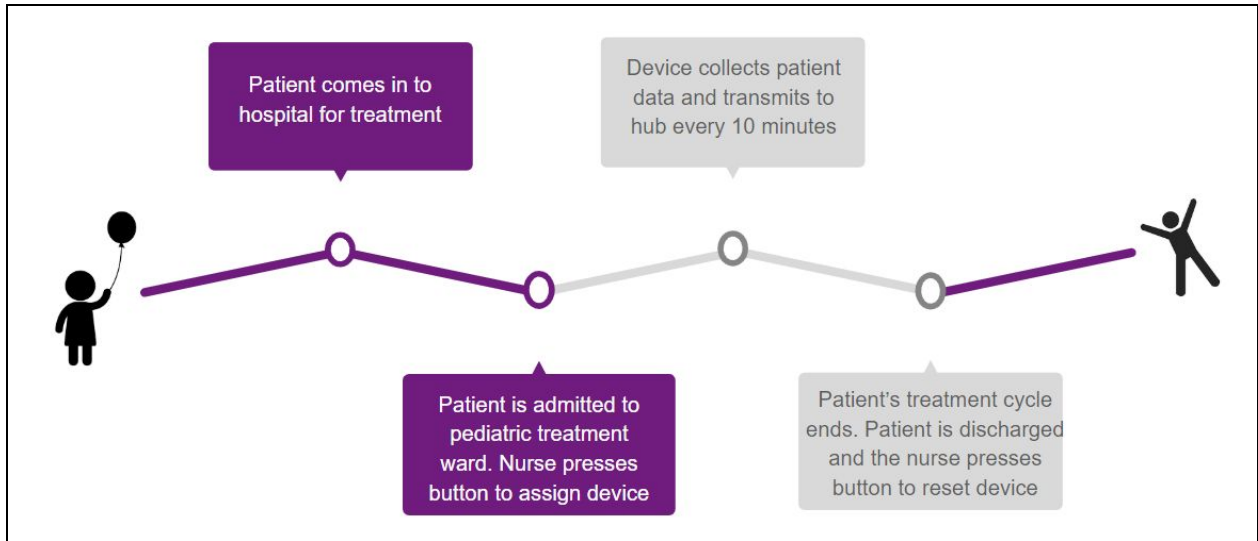
Figure 3.59: Timeline showing hospital workflow incorporating PediaTrack
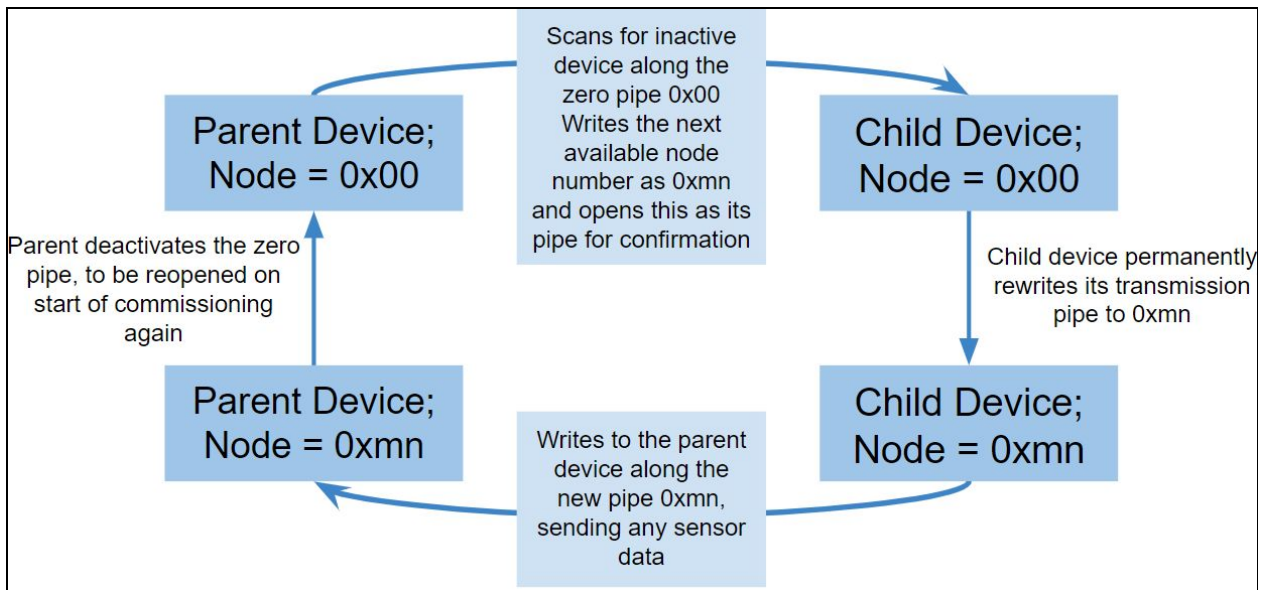
Figure 3.60: Cyclic flow chart showing the interaction between parent and child devices during commissioning
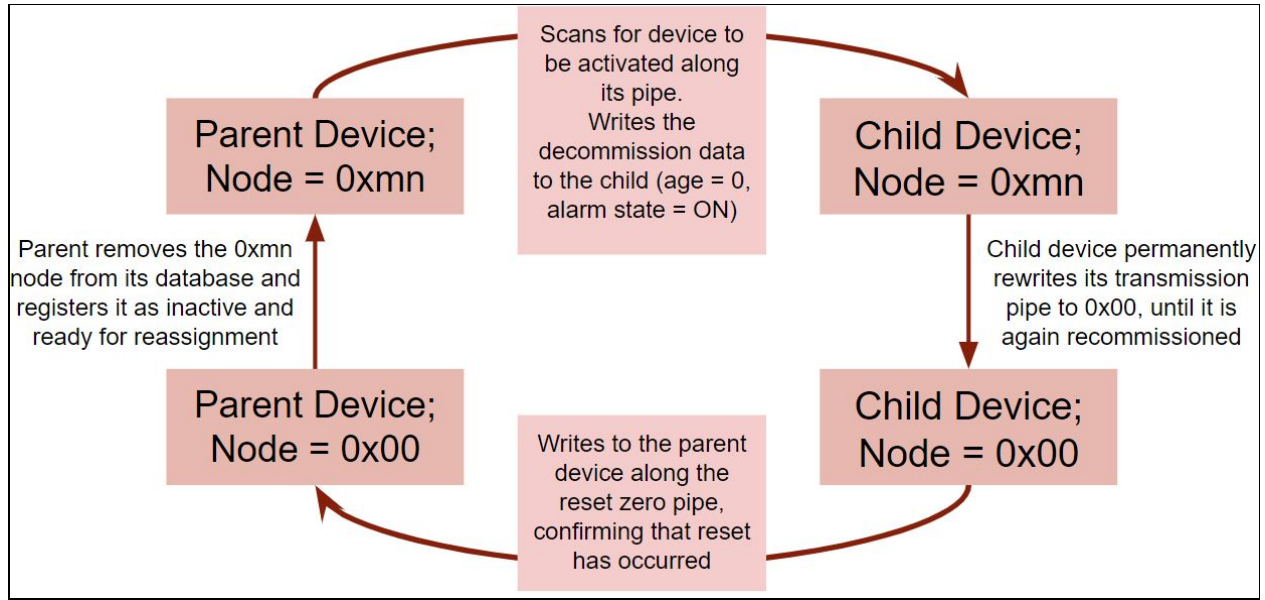
Figure 3.61: Cyclic flow chart showing the interaction between parent and child devices during decommissioning

Conclusion

We set out to satisfy the design criteria outlined in the introduction of this section, and through this first implementation, we succeeded in creating an initial prototype of the data communication system. In choosing RF, we were taking a small risk due to its higher difficulty of implementation, but through significant research into the protocols behind the NRF24L01+ module we chose, we were able to design a solution that:

- Has no requirements of existing hospital infrastructure such as WiFi, computers, or bedside power supplies.

- Keeps nurse workflow burden to a minimum, and potentially improving patient check-in time due to the constant data transmission of our system.

- Would keep alarm fatigue to the minimum level by regulating the time intervals in which data would be collected and sent.

- Uses no externally wired data transmission components, reducing points of weakness.

- Allows for quick updates of the patient database to prevent negative trends from being missed through the time intervals chosen for data sending and collecting.

- Allows for new patients to be easily integrated into the system, and for the easy removal of patients who have been discharged through the detailed commission and decommission protocols.

- Minimizes the processing power requirement and battery drain on the on-patient device.

Next Steps

Moving forward, we acknowledge that there are a few design specifications that should be pursued to make the communication procedure more effective, and make the best use of the data that we are transmitting. Primary among these will be the debugging of the Raspberry Pi direct interface with the NRF24L01+, as this will improve the run-time of our data collection and analysis on the hub. Another large undertaking that would be necessary moving forward would be to create our own library for the RF chip, as the existing library we are using has many unused functions and abstracts away a lot of the workings of the chip. If we were to build our own library, we would be able to eliminate some of the abstraction, and only keep the functions that we actually need. Customizing the functions for our system would also minimize the run-time on both the child and parent devices, and optimize battery life on the child device.

Another large change in the current protocol would be the addition of data counters to the information that the child device sends. This would allow for the child device to stockpile any data-sets that it did not successfully send and send them sequentially when it gets a connection to the hub, with a last-in-first-out (LIFO) system, that is; the most recent data point will always be sent first, followed by the next most recent until all the previously unsent data is sent. The LIFO system would work best, as in the unlikely instance of a single connection alone being achieved over a long period of time, the parent device will still have the most up-to-date information to display, and will fill in the longitudinal data when it gets a period of better connectivity. Having such a system would also allow for the parent device to recover, should the power go out on the the hub for a prolonged period of time. The timestamp currently sent with the sensor data is a

good foundation for the data counter structure, so the groundwork is already laid for such an addition to the protocol.

Finally, one last step would be to investigate the process by which we can encrypt the data storage, and also encrypt the RF transmission, as the system uses open-bandwidth RF, and someone who follows the same protocol we used for the chips would be able to intercept the communication and read the otherwise confidential patient data.

3.5 | Hub

Introduction

The purpose of the hub is to provide a centralized location where health professionals in a pediatric cancer ward can view and monitor patient vital data collected by the sensors, integrated by the on-patient microprocessor, and transmitted via radio frequency. To increase its usefulness, the hub alerts health professionals of a patient's immediate status based on the latest vital readings and allows health professionals to view patient vital data longitudinally. Therefore, the hub is a centralized alert and monitoring system, with the promise of helping detect the rapid onset of sepsis and alleviating nurse workload in wards with high patient-to-nurse ratios.

Design Considerations

There were two overarching design considerations for the hub. Since the hub is the main user-facing component of the PediaTrack vital monitoring system, it must be intuitive for its users: nurses and physicians. Additionally, it must make use of cost-effective hardware and software components to be deployed in low-income settings.

The hub team identified three ways to achieve an intuitive user interface for health professionals. First, the interface needed to be as *simple* as possible to streamline integration into nurse and physician workflows. This meant that the number of pages and the number of ways that a health professional could interact with the hub was minimized during the design process. In the final prototype, there were only three unique pages across the entire application and fewer than twelve editable fields per patient. Second, the hub needed to be accommodating of medical practice norms which would also streamline integration into pre-existing workflows. Several

pieces of patient information, such as medical ID, diagnosis, and date of birth, were included in the final implementation because health professionals regularly require access to this information. Third, the hub needed to display vital data in such a way that it could convey useful information to health professionals without overloading them. The representation of too many data points could make it more difficult for health professionals to tease out important patient information and act effectively. By emphasizing these design considerations, the team was able to make a user-friendly final prototype.

To help make PediaTrack a cost-effective vital monitoring solution, the hub needed to be built with cheap hardware and software components. On the hardware side, the main way to cut costs was to use a bare-bones central processing unit. We opted for a Raspberry Pi 3 Model B+ which is known for its ease of use and can be bought online for $35. Even with the purchase of external modules, such as a monitor, keyboard, and mouse, we believed this solution to be hundreds of dollars cheaper than purchasing a computer. Additionally, since the central processing unit would only be running a single program, the Raspberry Pi's reduced RAM, storage space, and processing power did not pose a problem. On the software side, there were enough open-source options to build the application free of charge. We opted for an application built around Flask, an open-source web framework, rather than using commercial products popular for Raspberry Pi application development like Xojo, which has a $100/year licensing fee (Xojo: Store).

The overall design considerations of making an intuitive hub interface and using cost-effective hardware and software components informed the platform we used to build the hub, downstream design decisions during the iteration process, and the data we collected to

evaluate the design. Namely, since the use-case was the central design consideration, we sought qualitative feedback from doctors rather than evaluating the application with less relevant quantitative metrics.

Design Concept

The hub application consists of three front-end pages. The main page is the default page and the one that health professionals will interact with most often. It consists of containers for each patient connected to the PediaTrack system. The containers will change color based on the patient's alarm state and display the patient's most important unique identification information. Additionally, each container contains two links that will redirect the user to patient-specific pages. The first page is a form where health professionals can view and edit patient information, such as location, diagnosis, and vital thresholds. The second page shows a patient's longitudinal vital data with graphs.

On the back-end, the application continuously updates by reading in new vital data sent from the on-patient sensors. The new vital data are time-stamped and stored in a database where they can readily be accessed to update the front-end display. Any changes made to patient information, including vital thresholds, are stored in a separate database from the vital records. Beyond listening for and storing patient information, the back-end also performs a basic analysis of a patient's vital thresholds and latest vital data to determine an alarm state, which indicates whether or not a patient requires immediate medical attention. The alarm state is indicated by the color of the patient's container on the main page.

The software is built on the Raspbian Stretch operating system and the Flask web microframework, and the hardware is built around the Raspberry Pi 3 Model B+ and additional hardware modules necessary to display and interact with the hub application.

Software Components

This section lists, with a granular level of detail, the software components we used to build the hub. All of the Python packages were installed using Python's preferred installer program (pip3), and pip3 itself was installed using Debian Linux's advanced package tool (APT).

Operating System

Software Component	Purpose	Installation
Raspbian Stretch	Raspberry Pi's Linux-based operating system which houses the rest of the software components and provides a graphical user interface.	Downloaded New Out of the Box Software (NOOBS) from Raspberry Pi's website and copied it onto a microSD card. NOOBS contains the Raspbian Stretch OS.

Table 3.16. The software components necessary for the Raspberry Pi's operating system.

Web Application Dependencies

Software Component	Purpose	Installation
Chromium	Browser where the application is displayed.	Pre-downloaded with Raspbian Stretch.
Python 3	Computer language for back-end processing.	Pre-downloaded with Raspbian Stretch.
pip3	Python package manager.	Installed with APT.
Flask	Web microframework that combines Python, HTML, CSS, and JavaScript files into	Installed with pip3.

	a web application.	
Jinja2	Software that allows Python to template an HTML page.	Pre-downloaded with Flask.
Bootstrap 4	Framework for building well-formatted web applications from HTML, CSS, and JavaScript files.	Installed from Bootstrap's website.
Chart.js	Tool that uses JavaScript to build graphs and charts.	Downloaded from Github.

Table 3.17. The software components necessary for the web application component of the hub.

Database Dependencies

Software Component	Purpose	Installation
SQLite3	Database to store patient information.	Installed with APT.
SQLAlchemy	Python toolkit to easily interact with SQL databases.	Installed with pip3.
Flask-SQLAlchemy	Flask extension that supports SQLAlchemy.	Installed with pip3.

Table 3.18. The software components necessary to support the database.

Interaction with the Arduino

Software Component	Purpose	Installation
Arduino IDE	Software necessary to read in packaged vital data from the hub-side Arduino.	Downloaded directly from Arduino's website.
PyZMQ	Python binding to asynchronous messaging software that allows communication between simultaneously running applications.	Installed with pip3.
PySerial	Python toolkit to interact with serial ports on a computer.	Installed with pip3.

Hardware Components

This section lists the hardware components necessary for the hub. Since most of the components did not need to be specified during the prototyping process, many of the components listed are generic. Additionally, power supplies and cords to connect the various components are not listed though necessary.

Hardware Component	Purpose
Raspberry Pi 3 Model B+	Central processing unit on which all of the software runs.
MicroSD Card (8 GB)	Provides memory for the Raspberry Pi.
PCF8523 Real Time Clock Module	Allows the Raspberry Pi to retain the correct time when not connected to Wi-Fi.
Monitor	Allows the web application to be viewed on a screen.
Keyboard & Mouse	Allow the health professional to interact with the web application.

Table 3.20. The main hardware components necessary for the implementation of the hub.

First Prototype

Application Structure and Function

After outlining the necessary software and hardware components, we began to prototype the hub. The initial prototype was a self-contained web application with three functional pages: a page displaying all patients on a main screen with an alarm state assigned to those in unhealthy conditions, a page with each patient's individual information and editable vital threshold fields, and a page with the patient's vital data over a longitudinal period. Although time was arbitrarily

designated with numbers, this period was meant to represent 24 hours in the final

implementation which is an appropriate time scale to detect the onset of sepsis. To achieve this

basic functionality, we implemented the files in the following file structure, and this file structure

stayed the same in the final implementation of the hub application.
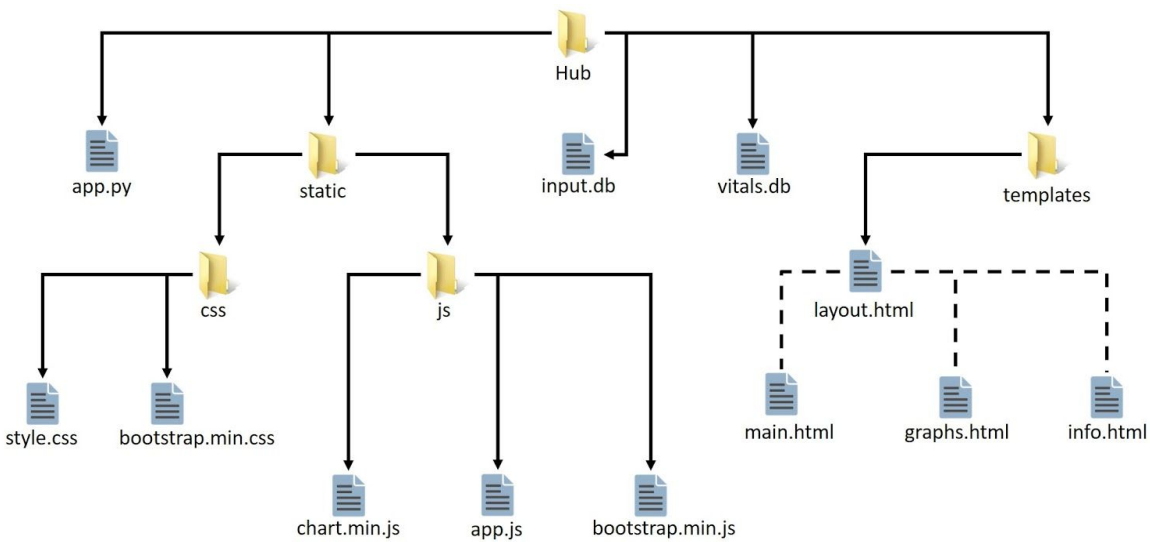
Figure 3.62. The file structure of the first prototype of the hub.

The application file, app.py, defines the routes supported by the application and specifies

the application's response when routes are accessed. On the main route, the page queries the

patient information database, called input.db, passes each patient's identification number and

name to main.html, and renders main.html in the browser. The main route is called when the

application is first opened in the browser or the main page is accessed from either the patient

information or patient vital data pages via the "return to main" button.

On the patient information route, if the request method is "get," the page queries a single

patient's information from the patient information database, passes this information to info.html,

and renders info.html. The patient information route's "get" method is called when the user

clicks on the "edit" button of a patient's container on the main page. The specific container that

is clicked also passes the patient identification number to the route which indicates which patient's information should be queried. On the patient information route, if the request method is "post," the page reads the form input fields and updates the patient information database if fields have been changed. The patient information route's "post" method is called when the user clicks the "set" button on the patient information page.

On the patient vital data route, the page queries a single patient's vital data from the vital.db database, passes this information to graphs.html, and renders graphs.html. The graphs are made with JavaScript functions from the chart.min.js file. The patient vital data route is called when the user clicks on the "data" button of a patient's container on the main page. The specific container clicked dictates which individual patient's data will be queried and displayed.

Finally, there is a back-end updating mechanism hidden from the user. When the application is loaded in the browser, it calls a JavaScript function in app.js. This function calls the update route in app.py which, in turn, queries the vital data and vital thresholds of each patient, determines each patient's alarm state, updates the database if alarm states have changed, and returns the patient identification numbers and corresponding alarm states to the function in app.js. The function then updates the color of each patient's container based on his or her alarm state, with white indicating normal vitals and red indicating abnormal vitals. If any vital is below the lower threshold or above the upper threshold for that vital, the container will become red, indicating that the patient requires medical attention. The function in app.js then calls itself to restart the update process.

The styles.css, bootstrap.min.css, and bootstrap.min.js files are referenced by the html pages to stylize the content in a consistent, clean way.

User Experience

When the application is opened, the user starts on the main page. Here, the user sees a separate container for each patient, and the patients' containers are arranged regularly in a grid with up to six columns. Within each container, the patient's identification number and name are displayed, as well as one button called "edit" and another button called "data." Some of the containers are red, indicating these patients have abnormal vitals, and some of the containers are white, indicating these patients have normal vitals. If the data in the databases are changed by a separate Python script while the page is open, the user will see some containers change color as patient alarm states change. This simulates live-updating data in the final working prototype.
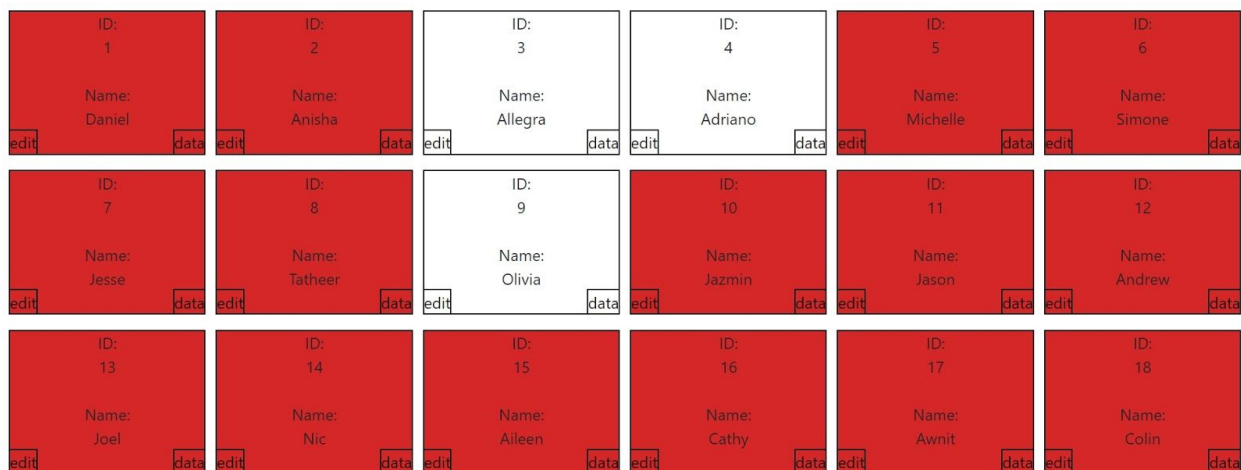
Figure 3.63. The first prototype's main page with simulated patients.

From here, the user can click the "edit" button in any patient's container to reach his or her patient information page. On the patient information page, the user sees pre-filled form fields for the patient's identification number and name as well as the lower and upper thresholds for respiratory rate, heart rate, and temperature. The user can change any field except for the patient identification number and click "set" to save those changes in the database and redirect back to

198

the main page. This allows health professionals to change the vital thresholds on a patient-by-patient basis. The next time the user opens the form page for this patient, the fields will reflect any changes made by the user. The user can also get back to the main page without making changes by pressing the "return to main" button.

Figure 3.64. The first prototype's patient information page.

Again on the main page, the user can press the "data" button to view the longitudinal vital data of a patient. The page shows three graphs, one for each vital sign. Because these vital data are simulated in the first prototype, time is arbitrarily represented with floats. This page does not update in real time but takes a snapshot of the patient's vital data as represented by the database when the "data" button is pressed. To return to the main page, the user can once again click the "return to main" button.
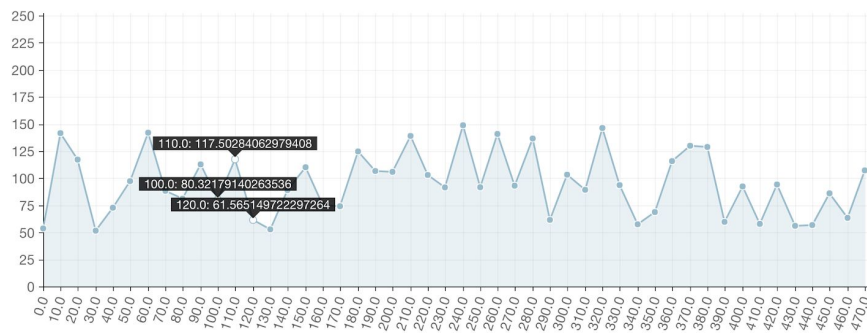
Figure 3.65 The heart rate graph on the patient vital data page from the first prototype. In this iteration, users could not cleanly interact with a single data point.

Limitations

The first prototype of the hub was a stand-alone application. In order to function, the vital sign database needed to be pre-populated with simulated data, and live-updating could only occur when new data was simulated. There was no mechanism for the hub to update the database based on data points sent from the sensors as integration with the communication team had not yet happened. Integration with the communication team posed a technical challenge that would need to be overcome by the implementation of the final prototype.

Additionally, certain details were deferred in this first prototype, such as time-stamping each new vital sign measurement before saving it in the database, smoothing out the users' interactions with individual data points on the graphs, and compacting the patient information page to reduce the need to scroll extensively. However, it reasonably demonstrated the hub's intended functionality and use-case, so we sought feedback from medical professionals about how we could improve on this iteration.

Design Feedback and Iteration

After the first iteration, we sought feedback from Dr. Chen and Dr. Rizwan at Boston Children's Hospital. Dr. Chen is a cardiologist and the director of the Cardiovascular Health for Cancer Survivors Program. Dr. Rizwan used to be a general physician in Pakistan and is now one of Dr. Chen's colleagues at Boston Children's Hospital.

This meeting highlighted the importance of having unique identifying information for patients, such as date of birth and a medical ID number, which health professionals often need to access to provide quality care to their patients. Although we did include an ID for each patient, we did not have a field for health professionals to enter the medical identification number for each patient as issued by a separate healthcare system. They also suggested that the location of the patient in the ward be indicated on the main page as this would allow health professionals to act more efficiently if a patient's vitals became unhealthy. Finally, they advised that the hub system's patient identification number (i.e. not the medical identification number) be made less prominent or removed altogether from the user's experience. Dr. Chen and Dr. Rizwan also recommended that we include another graph for each vital–heart rate, respiratory rate, and temperature–which would indicate average vital trends over the past few days. These graphs would especially be helpful in providing information about the patient's normal vitals on a daily basis and any abnormalities that subsist over a longer period of time. The graphical representation of vital signs over 24 hours, as we previously envisioned, is more typical of an intensive care unit setting and is unlikely to be useful for patients who are in a stable condition.

We sought additional feedback for our application from Dr. Albanti and Dr. Lehmann from the Global Health Initiative. They stressed that age is a determining factor that affects

healthy vital thresholds and that the hub could incorporate age-based thresholds. Additionally, they advised that displaying a separate intermediary state for patients who were near but not beyond their vital thresholds would be very useful in the prioritization of patient care. This would make our vital alarm state similar to a threshold-based PEWS system, which is a method of detecting early warning signs in patients and used for detecting sepsis and infection.

Final Prototype

There were two main improvements in the final prototype compared to the initial prototype. First, the design implemented recommendations by medical professionals to make the hub more useful in a care setting. Second, the hub was integrated with the communication system so that it could update databases and the application pages with live vital data.

The following specific changes were made to the hub application based on recommendations made by doctors at Boston Children's Hospital and the Global Health Initiative:

1. On the main page, the containers display a patient's first and last name, location, date of birth, and medical identification number. Additionally, the device identification number is shown in the top left corner so that it is less prominent.

2. The containers turn yellow when a patient is close to their vital thresholds but still within healthy bounds. If the difference between a vital and the nearest vital threshold is less than 10% of the difference between the upper and lower thresholds, the patient's container will turn yellow.

3. The patient information database was expanded to include values for patient surname, location, date of birth, medical identification number, and diagnosis.

4. The hub has pre-set age-based thresholds for each vital based on literature values (see the figure below). The literature vital thresholds are used to determine patients' risk for sepsis. When a patient is added to the system, these vital thresholds are set based on their age though a health professional can still edit the thresholds on a patient-by-patient basis afterwards.

5. The patient vital data page includes two graphs for each vital, one indicating the trend over the last 24-hours and the other indicating the daily averages for up to five days.

Age	Heart Rate Upper Threshold (bpm)	Heart Rate Lower Threshold (bpm)	Respiratory Rate Upper Threshold (bpm)	Respiratory Rate Lower Threshold (bpm)	Temperature Upper Threshold (C)	Temperature Lower Threshold (C)
0-1 month	180	100	60	30	38.5	36
1 month - 1 year	180	90	45	25	38.5	36
1-5 years	140	70*	30	20	38.5	36
5-12 years	130	60*	24	16	38.5	36
12-18 years	110	50*	20	14	38.5	36

Table 3.21. Normal vital ranges specified by age group. The heart rate and temperature thresholds are from Mathias et al. and the respiratory rate thresholds are from similar age ranges as given in Bradshaw et al. The asterisked values are extrapolations from the data in Mathias et al. because literature values could not be found.

Additionally, the 24-hour vital graphs indicate the upper and lower vital thresholds with horizontal red and blue lines respectively, and a gray container was added to the main page to indicate what each field corresponds to. These changes are represented in the sample pages for the final prototype shown in the figures below.

Figure 3.66. The main page in the final prototype with additional patient information and a yellow color to indicate an intermediate state of risk.

Figure 3.67. The patient information page in the final prototype with new fields: medical ID, date of birth, location, diagnosis, and age (calculated from date of birth).

Jane

Last 24 Hours | Last 5 day average

Heart Rate
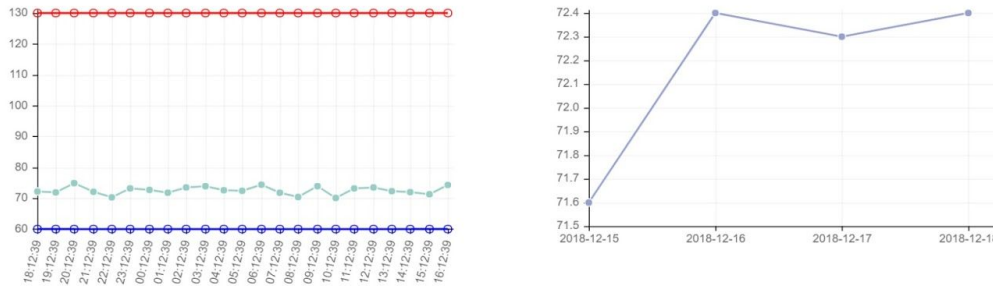
Figure 3.68 . The 24-hour and daily average heart rate graphs in the final prototype vital data page.

In order to display live data coming from the PediaTrack sensors, the hub needed to have some additional functionality. A separate process, running in a separate terminal, runs a script called serialreader.py. This script establishes a connection with the serial port where information is being transmitted from the hub-side Arduino. It reads in a complete string from the Arduino, consisting of the device identification number, time delay between the measurement recording and its transmission, heart rate, respiratory rate, and temperature. Then, it sends this string to app.py, which parses the string and stores the new vital information in the vital database with a time-stamp. This occurs within the update route specified in the first prototype and allows the hub to integrate fully with the rest of the PediaTrack system.

Conclusion

There are a couple of additional features that can be added to the hub. Most importantly, the hub needs to have pages that link it to the commissioning and decommissioning procedure

that will be specified by the communication component of the PediaTrack system. This will allow users to add and remove patients based on who is in the ward, a feature that was deferred for this prototype.

Based on feedback from the GHI, the hub could eventually include an electronic medical record system, which could build off of the limited patient information page. In addition, the main page could be adapted for each hospital to make the patient location feature more specific. This would include integrating the hospital layout into the hub design, and, instead of displaying a grid of patient bed numbers, it would display the actual bed layout. Finally, since the goal of PediaTrack is implementation across LMICs, the hub will need multi-language functionality. This could take the form of adding a multi-language API such as Google Translate or a more simple option to translate the few words on the page by switching the html file based on language. But either option would require more research into cost and usability to determine the best option.

The hub was able to achieve the desired functionality that worked with the rest of the PediaTrack system. We were able to receive and display live patient data as well as assign an alarm state to help with sepsis recognition for the medical professional. Finally, we were able to design a hub with hardware and software dependencies that are at a low cost while maintaining efficiency and simple usability for medical professionals.

3.6 | Final Prototype

We arrived at our final prototype in three major stages. The first of which was a series of disconnected sensors on breadboards, which allowed us to quickly construct functional prototype circuits, though we knew these would not be what we ultimately implemented (Figure 3.69). We then moved toward a solderable breadboard design, which allowed us to consolidate all our circuits onto one device (Figure 3.70). This device has all the functionality of our final prototype, but is not as compact or specific to our design, as there are a number of open, unused pins. Finally, we arrived at our final PCB-based prototype (Figure 3.71). This includes all the sensing, processing, communication, and powering functionality that we defined as necessary for our device. It is as compact as possible, and easiest to mass produce of the prototypes. Ultimately this is essentially what would be included in a commercial version of our device.
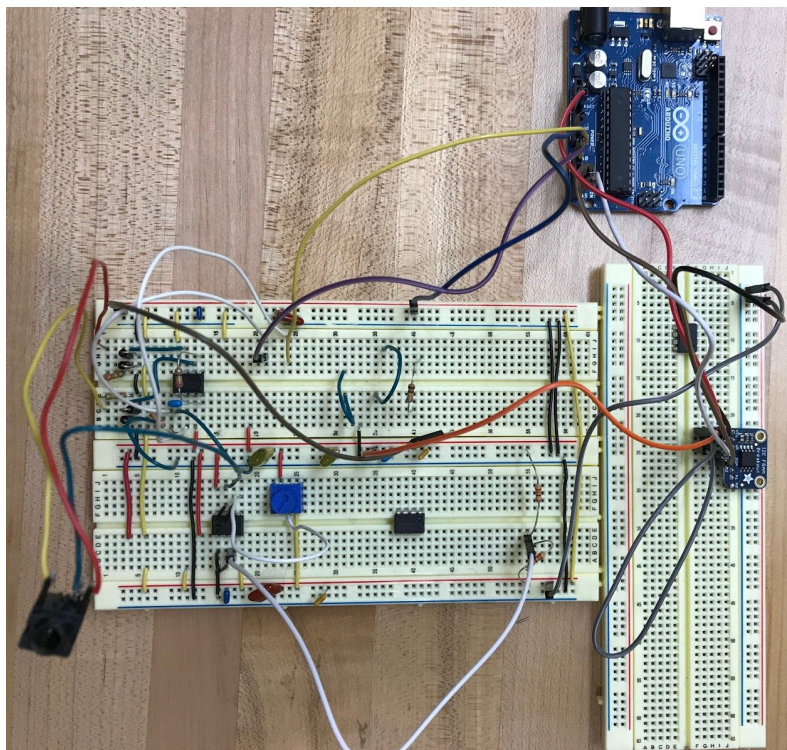
Figure 3.69: Breadboard prototype of ECG lead circuitry.
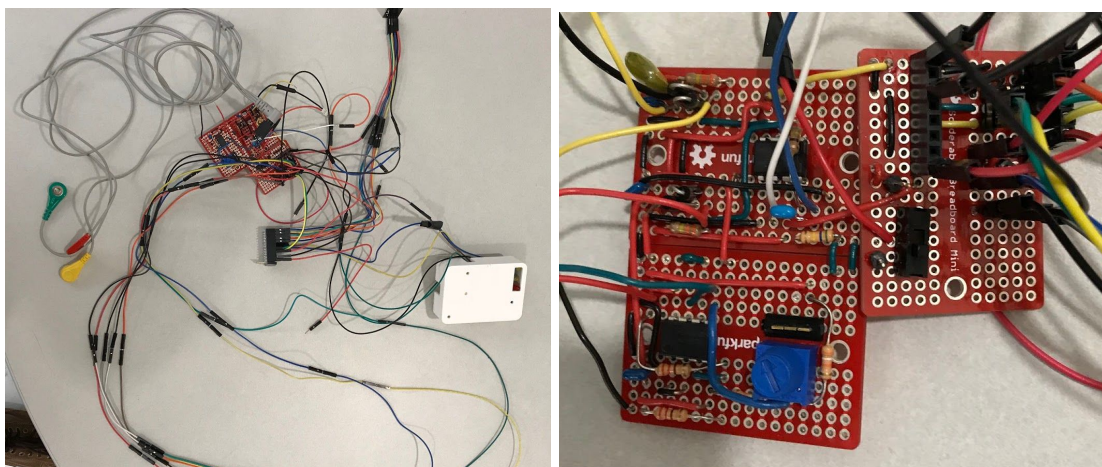
Figure 3.70: (left) Solderable breadboard prototype with attached ECG leads, RF chip, and temperature sensor. (right) Zoom in of the solderable breadboard itself.
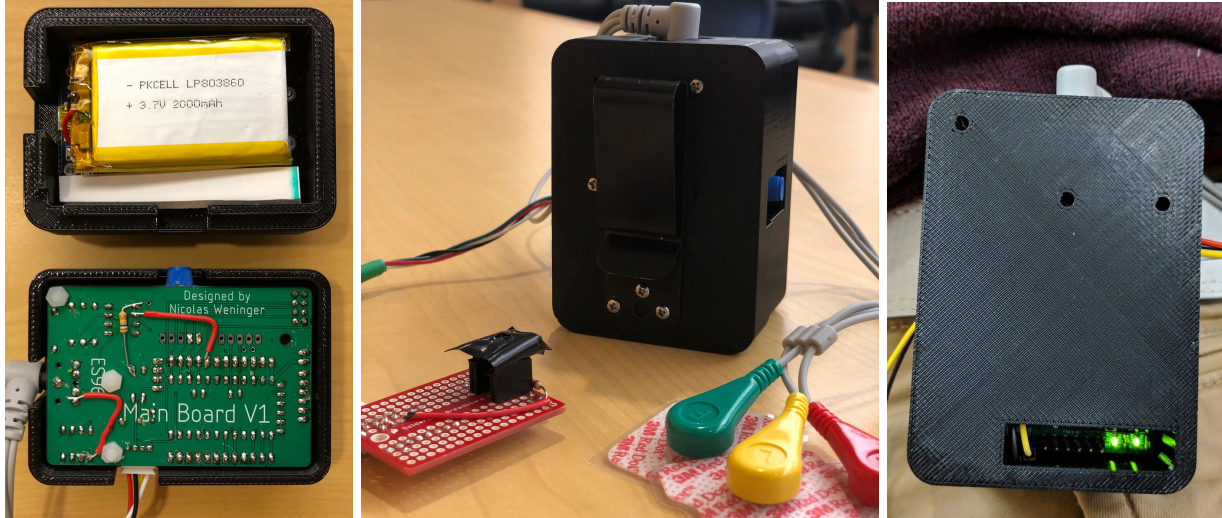
Figure 3.71: (Left) Inside of final prototype with PCB, battery, and charge controller inserted with ECG leads and temperature sensor connected. (Center) Final assembled prototype with leads and temperature sensor attached. (Right) LED alarm displaying a green light for healthy vitals.

With our final prototype constructed, we needed to assess whether or not the technical specifications we defined had been met and whether we had satisfied our initial design considerations. In the end, all of the design considerations were satisfied and most of the technical specifications were met, though some require more extensive testing for true verification.

Specification	Satisfied / Unsatisfied
Inexpensive (<$300)	Satisfied ($55 cost of goods)
Allow child mobility	Satisfied
Communicate across a ward	Satisfied (see table 3.15)
No WiFi reliance	Satisfied (RF communication)
No computer reliance	Satisfied
Lifetime > 1 year	Needs verification
Battery life > 1 day	Satisfied (40 hours)
Recharge time < 1 day	Satisfied (4 hours)

Electrically and chemically safe	Needs Verification
Sweat resistant	Needs Verification

Table 3.22: Table assessing the validation of the original specifications we determined for our design

The specifications that still need verification, lifetime, safety, and sweat resistance, are all feasible for us to achieve and verify, but not within the time constraints of this course. We are unable to state with certainty the lifetime of our device until we have been able to test it in real-life use cases over a the lifetime of the device, which should take at least a year. The safety and sweat resistance specifications require external verifications and could be tested with specific equipment that was not available to us during the development of our prototype; however, we see no reason our device could not pass these tests given more time.

As a demonstration of the PediaTrack system and its response to alarming vitals, we used a heat pad to simulate a fever. Our subject had ECG and temperature leads attached appropriately and was assessed to have standard vital signs as seen in Figure 3.71(a) before applying the heat pad to the temperature sensor reading. Once the temperature sensor measured the influx of heat from the pad, the warning indicator on the on patient device lit up, as seen in Figure 3.71(b). Then , within ten seconds, the alarm-raising temperature data was transmitted to the hub. The hub received the alarm state and adjusted the subject's indicator box to turn red, signalling a possible detection of sepsis (Figure 3.71(c)). The subject's longitudinal data was then accessed from the homepage for visualization (Figure 3.71(d)).
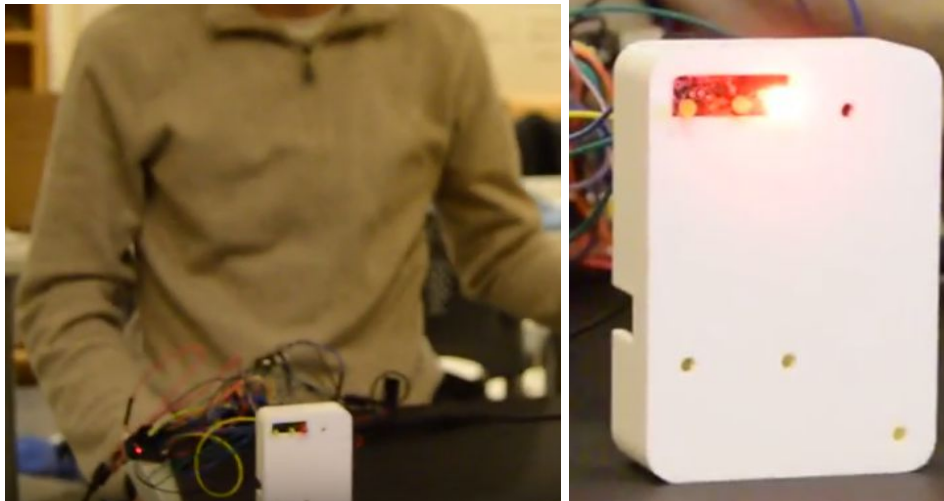
Figure 3.71(a): (Above left) Subject with PediaTrack vital monitoring system attached and exhibiting normal vitals.
Figure 3.71(b): (Above right) Subject exhibits simulated elevated temperature using a heating pad. In response a red indicator LED lights up to alert bystanders to potential sepsis.

Figure 3.71(c): The hub receives elevated temperature from the on-patient device. In response, the indicator box for

the subject changes color to red.
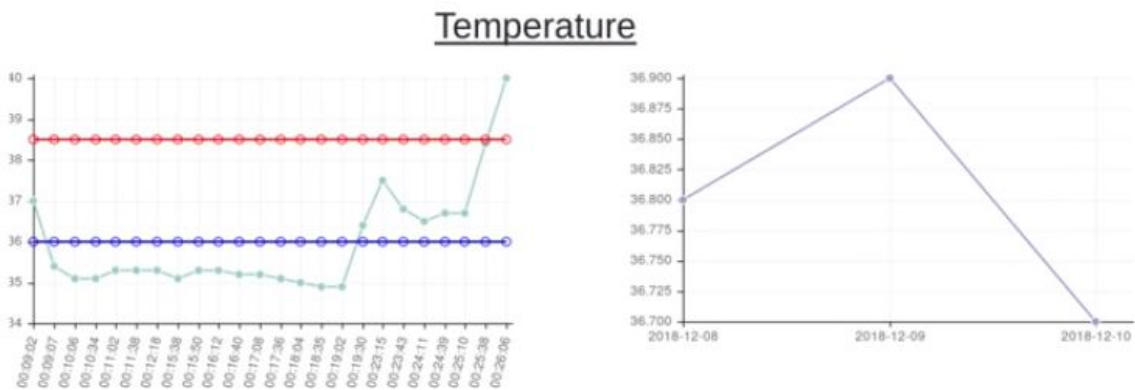
Figure 3.71(d): Longitudinal temperature data plotted over the past 24 hours (left) and 3 days (right)

3.7 | Future Improvements

Going forward, there are many features and additions to the system that would vastly increase its usability and successful integration into a hospital setting, in addition to the important refinements and additional testing we've listed previously.

Although a major upside of the hub is that all the data is in one central location, the one downside is that this data is tied to the hub itself and cannot be viewed by a nurse or clinician at somewhere that isn't the central ward location. Providing an app interface in which healthcare providers could see the same hub data on a smartphone or tablet (and therefore when at the patient's bedside) would bring the data closer to the patient, and improve workflow of when nurses and doctors are checking on patients. An app interface could also be used by clinicians or parents to input behavioral data, another thing the GHI presented would be helpful in analyzing the status of a patient and their risk for sepsis. Additionally, the system could, in future iterations, integrate with existing medical record systems that might exist in some of these locations, offering more fluid combination with current practices.

Finally, this idea is not specific to pediatric cancer; thus, PediaTrack could be extended to various other use cases in different diseases and conditions in the medical field. Although it has been a great specific use case to for our focus (as outlined in our background information), sepsis is a problem in many hospital settings outside of pediatric and cancer wards, and our system could be used in many cases to help save lives by providing early warning for sepsis.

Concluding Remarks

The PediaTrack system offers unique value in a vital monitoring system as it targeted particular towards pediatric oncology patients. While most vital monitoring systems are designed for adults and incompatible with children, the application of PediaTrack is an easy-to-apply, easy-to-wear monitoring system designed for children. It is also targeted to monitoring vitals that are essential to the recognition of sepsis. Particularly, it allows for continuous monitoring of body temperature, which is the most important vital in determining if a patient has an infection which could become septic, and yet not recorded by most vital monitoring systems. To illustrate, among seven commercial vital monitoring systems suitable for children, only two monitored temperature (McCarthy et al.). By continuously monitoring patients and sending alarms when abnormal vitals are detected, Pediatrack will allow nurses to effectively prioritize care and by displaying the data collected in a simple to use interface it will provide nurses with access to longitudinal data and make more informed decisions towards treatment. Finally, and most importantly, PediaTrack is a cost-effective, infrastructure-independent solution which can be implemented across low-resource settings where vital monitoring is not the norm.

Ultimately, by resting at the intersection of the collection, communication, and interpretation of patient data, PediaTrack is our class's small step forward to closing the gap in survival rates for pediatric cancer patients in LMICs.

Works Cited

Duncan, H., Hutchison, J., & Parshuram, C. S. (2006). The Pediatric Early Warning
        System score: a severity of illness score to predict urgent medical need in
        hospitalized children. Journal of critical care, 21(3), 271-278.

Friedrich, P., Ortiz, R., Fuentes, S., Gamboa, Y., Ah Chu-Sanchez, M. S., Arambú, I. C.,
        & Central American Association of Pediatric Hematologists and Oncologists
        (AHOPCA). (2014). Barriers to effective treatment of pediatric solid tumors in
        middle-income countries: Can we make sense of the spectrum of nonbiologic
        factors that influence outcomes?. Cancer, 120(1), 112-125.

Holley, R. (2009). How good can it get? Analysing and improving OCR accuracy in large
        scale historic newspaper digitisation programs. D-Lib Magazine, 15(3/4).

McCarthy, S., O'Connor, Y., Thompson, M., O'Sullivan, T., & Ryan, D. (2015). Wearable
        Vital Sign Sensors and Their Potential within Low and Middle Income Countries,
        11.

Day, S., Challinor, J., Hollis, R., Abramovitz, L., Hanaratri, Y., & Punjwani, R. (2015).
        Paediatric oncology nursing care in low-and middle-income countries: a need for
        baseline standards. Cancer Control, 2015, 111-116.

Mathias, B., Mira, J., & Larson, S. D. (2016). Pediatric Sepsis. Current Opinion in
        Pediatrics, 28(3), 380–387. https://doi.org/10.1097/MOP.0000000000000337

Xojo: Store. (n.d.). Retrieved December 18, 2018, from https://www.xojo.com/store/

Nordic Semiconductors NRF24L01P Product Specifications:
        https://www.nordicsemi.com/-/media/DocLib/Other/Product_Spec/nRF24L01PPS
        v10.pdf?la=en

Roy Moulik, N., Kulkarni, K. P., & Kumar, A. (2016). Tracking children with acute
        lymphoblastic leukemia who abandoned therapy: Experience, challenges, parental
        perspectives, and impact of treatment subsidies and intensified counseling.
        Pediatric Hematology and Oncology, 33(5), 327–337.
        doi:10.1080/08880018.2016.1212956

Mathew, A. (2018). Global Survey of Clinical Oncology Workforce. Journal of Global
        Oncology, 4, 1-12.

Hughes, R. G. (2008). Nurses at the "sharp end" of patient care.

Jamoom, E., Patel, V., King, J., & Furukawa, M. (2012, August). National perceptions of
        ehr adoption: Barriers, impacts, and federal policies. National conference on
        health statistics.)

World Health Organization. (2002). MHealth : New Horizons for Health through Mobile
        Technologies:Second Global Survey on eHealth.

"Pediatric Cancer and Blood Disorders Around the World." Dana-Farber/Boston
        Children's Cancer and Blood Disorders Center.

Respiratory Belt Transducer. (n.d.). Retrieved December 19, 2018, from
        https://www.adinstruments.com/products/respiratory-belt-transducer

Holley, R. (2009). How Good Can It Get?: Analysing and Improving OCR Accuracy in
        Large Scale Historic Newspaper Digitisation Programs. D-Lib Magazine, 15(3/4).

Wal-Mart Stores, Inc. (n.d.). 9V Batteries. Retrieved December 18, 2018, from
        https://www.walmart.com/c/kp/9v-batteries

Weenk, M., van Goor, H., Frietman, B., Engelen, L. J., van Laarhoven, C. J., Smit, J.,
        Bredie, S. J., … van de Belt, T. H. (2017). Continuous Monitoring of Vital Signs
        Using Wearable Devices on the General Ward: Pilot Study. JMIR mHealth and
        uHealth, 5(7), e91. doi:10.2196/mhealth.7208

Cootes N. (2010). Managing the unwell child. London journal of primary care, 3(1),
        19-26.

Sterling, S. A., Miller, W. R., Pryor, J., Puskarich, M. A., & Jones, A. E. (2015). The
        Impact of Timing of Antibiotics on Outcomes in Severe Sepsis and Septic Shock:
        A Systematic Review and Meta-Analysis. Critical care medicine, 43(9), 1907-15.

Micro Electronics, Inc. (n.d.). AOC E2070SWHN 19.5. Retrieved December 18, 2018,
        from
        https://www.microcenter.com/product/482457/e2070swhn-195-hd-60hz-vga-hdmi
        -led-monitor

Bradshaw, C., Goodman, I., Rosenberg, R., Bandera, C., Fierman, A., & Rudy, B. (2016).
        Implementation of an Inpatient Pediatric Sepsis Identification Pathway.
        Pediatrics, 137(3), e20144082. https://doi.org/10.1542/peds.2014-4082

Global Health Initiative. (n.d.). Where We Work: Rwanda. Retrieved December 17, 2018,
        from
        http://www.danafarberbostonchildrens.org/why-choose-us/global-health/where-w
        e-work/rwanda.aspx

Global Health Initiative. (n.d.). Where We Work: Egypt. Retrieved December 17, 2018,
        from
        http://www.danafarberbostonchildrens.org/why-choose-us/global-health/where-w
        e-work/egypt.aspx

About the Global Health Initiative. (2018). Retrieved October 1, 2018, from
        https://www.dana-farber.org/global-health-initiative/

Bishr, Mai K. et al. (2018). Radiation Therapy Availability in Africa and Latin America:
        Two Models of Low and Middle Income Countries. *International Journal of*
        Radiation Oncology, Biology, Physics, 102(3), 490-498.

Buckle, G., Collins, J., Sumba, P., Nakalema, B., Omenah, D., Stiffler, K., . . .
        Moormann, A. (2013). Factors influencing time to diagnosis and initiation of
        treatment of endemic Burkitt Lymphoma among children in Uganda and western
        Kenya: A cross-sectional survey. Infectious Agents and Cancer, 8(1), 36.

Brandon, K., & Gordon, A. (2018, September 19). Dana Farber Nurse Interviews
        [Personal interview].

Longombe, N. A., Nalbandyan, A., Perez, M., Zia, N., & Kandasamy, S. (2018,
        September 18). GLOBE Scholars Interview [Personal interview].

Shea, M., & Rogers, M. (2018, September 10). Persistent Productions Interview
        [Personal interview].

Albanti, I. (2018, September 10). [Personal interview].

Fernandez, J. (2018, September 19). Dana Farber Psychosocial Team [Personal
        Interview].

Datta, Samiei, & Bodis. (2014). Radiation Therapy Infrastructure and Human Resources

in Low- and Middle-Income Countries: Present Status and Projections for 2020. *International Journal of Radiation Oncology, Biology, Physics, 89*(3), 448-457.

de Leon M. P. (1994) Oncogenes and Tumor Suppressor Genes. In: Familial and Hereditary Tumors. Recent Results in Cancer Research, vol 136. Springer, Berlin, Heidelberg

Every Two Minutes a Child is Diagnosed with Cancer. (n.d.). Retrieved October 1, 2018, from https://www.stbaldricks.org/press-releases/every-two-minutes-a-child-is-diagnosed-with-cancer/

Food and Agriculture Organization of the United Nations (FAO), International Fund for Agricultural Development (IFAD), the United Nations Children's Fund (UNICEF), World Food Programme (WFP), & World Health Organization (WHO). (2018). The State of Food Security and Nutrition in the World 2018. Building Climate Resilience for Food Security and Nutrition.

Geenen MM, Cardous-Ubbink MC, Kremer LCM, et al. Medical Assessment of Adverse Health Outcomes in Long-term Survivors of Childhood Cancer. JAMA. 2007;297(24):2705–2715.

Gupta, S., Howard, S., Hunger, S., Antillon, F., Metzger, M., Israels, T., Harif, M., & Rodriguez-Galindo, C. Treating Childhood Cancer in Low- and Middle-Income Countries. (2015, November 1). In (pp. 121-146).

Halbert, Jay, & Khaing, Aye. (2014). Overview of pediatric oncology and hematology in Myanmar. South Asian Journal of Cancer, 3(1), 78-82.

How Cancer Is Diagnosed - National Cancer Institute. (2015). Retrieved October 1, 2018, from https://www.cancer.gov/about-cancer/diagnosis-staging/diagnosis

Kanyamuhunga, A., Tuyisenge, L., Stefan, D. (2015, Aug. 31). Treating childhood cancer in Rwanda: the nephroblastoma example. Retrived from https://www.ncbi.nlm.nih.gov/pmc/articlesPMC4646289

Lichtenstein, P., Holm, N. V., Verkasalo, P. K., Iliadou, A., Kaprio, J., Koskenvuo, M., Hemminki, K. (2000, July 13). Environmental and heritable factors in the causation of cancer--analyses of cohorts of twins from Sweden, Denmark, and Finland. Retrieved from https://www.ncbi.nlm.nih.gov/pubmed/10891514

Literacy of Adults in Developing Countries: New Data from a Skills Survey | Center For Global Development. (2016, August 2016). Retrieved October 1, 2018, from https://www.cgdev.org/blog/literacy-adults-developing-countries-new-data-skills-survey

Mayer, L. D., & Janoff, A. S. (2007). Optimizing Combination Chemotherapy by Controlling Drug Ratios. *Molecular Interventions,7*(4), 216-223. doi:10.1124/mi.7.4.8

Murphy, A.J., Mosby, T.T., Rogers, P.C., Cohen, J., & Ladas, E.J. (2014). An international survey of nutritional practices in low- and middle-income countries: A report from the International Society of Pediatric Oncology (SIOP) PODC nutrition working group. *European Journal of Clinical Nutrition*, 68(12), 1341-1345.

Mostert, S., Sitaresmi, Gundy, Janes, Sutaryo, & Veerman. (2010). Comparing childhood

leukaemia treatment before and after the introduction of a parental education programme in Indonesia. Archives of Disease in Childhood, 95(1), 20-25.

NCI Dictionary of Cancer Terms. (2011, February 2). [nciAppModulePage]. Retrieved October 1, 2018, from https://www.cancer.gov/publications/dictionaries/cancer-terms

Njuguna F., Mostert S., Slot A., Langat S., Skiles J., Sitaresmi M. N., van de Ven P. M. , Musimbi J., Muliro H., Vreeman R. C., Kaspers G. J. (2014). Abandonment of childhood cancer treatment in Western Kenya. Archives of Disease in Childhood, 99(7), 609-14.

Pediatric Cancer and Blood Disorders Around the World. Retrieved September 16, 2018, From http://www.danafarberbostonchildrens.org/why-choose-us/global-health-initiative / global-health-statistics.aspx

Risk Factors and Causes of Childhood Cancer. (n.d.). Retrieved October 1, 2018, from https://www.cancer.org/cancer/cancer-in-children/risk-factors-and-causes.html

Rodriguez-Galindo, C., Friedrich, P., Morrissey, L., & Frazier, L. (2013). Global challenges in pediatric oncology: Current Opinion in Pediatrics, 25(1), 3–15.

Samb, B., Desai, N., Nishtar, S., Mendis, S., Bekedam, H., Wright, A., Etienne, C. (2010). Chronic Diseases: Chronic Diseases and Development 4: Prevention and management of chronic disease: A litmus test for health-systems strengthening in low-income and middle-income countries. The Lancet,376(9754), 1785-97.

Memorial Sloan Kettering Cancer Center. (n.d.). Retrieved from https://www.mskcc.org/

Spector, Pankratz, & Marcotte. (2015). Genetic and Nongenetic Risk Factors for Childhood Cancer. The Pediatric Clinics of North America, 62(1), 11-25.

Sudore, Landefeld, Pérez-Stable, Bibbins-Domingo, Williams, & Schillinger. (2009). Unraveling the relationship between literacy, language proficiency, and patient–physician communication. Patient Education and Counseling, 75(3), 398-402.

Survival | Cancer Trends Progress Report. (n.d.). Retrieved October 1, 2018, from https://progressreport.cancer.gov/after/survival

Thwaites, D. I., & Malicki, J. (2011). Physics and technology in ESTRO and in Radiotherapy and Oncology: Past, present and into the 4th dimension. Radiotherapy and Oncology, 100(3), 327-332. doi:10.1016/j.radonc.2011.09.014

Yilmaz, M. and Christofori. (2009). G. EMT, the Cytoskeleton, and Cancer Cell Invasion. Cancer Metastasis Rev 28, 15–33.

Ferlay, J. & Boyle, P. (2005). Cancer incidence, mortality and prevalence worldwide. *IARCPress*, 16(3), 481-488.

Mission & Values of BCH. Retrieved October 1, 2018, from http://www.childrenshospital.org/about-us/mission-and-values

Where We Work. Retrieved October 1, 2018, from http://www.danafarberbostonchildrens.org/why-choose-us/global-health/where-w e-work.aspx

Global Health | Armenia. Retrieved October 1, 2018, from http://www.danafarberbostonchildrens.org/why-choose-us/global-health/where-w e-work/armenia.aspx

Oeffinger, K. C., Mertens, A. C., Sklar, C. A., Kawashima, T., Hudson, M. M., Meadows,
  A. T., Robison, L. L. (2006). Chronic Health Conditions in Adult Survivors of
  Childhood Cancer. *New England Journal of Medicine*,355(15), 1572-1582.

Friedrich, P., Lam, C. G., Itriago, E., Perez, R., Ribeiro, R. C., & Arora, R. S. (2015).
  Magnitude of treatment abandonment in childhood cancer. *PLOS ONE*, *10*(9),
  e0135230.

Bertaut T., Bray F., DeSantis C., Fedewa S., Forman D., Islami F., Kramer J., Lancaster
  A., Lortet-Tieulent J., McMikel A. Miller K., O'Brien M., Parkin M., Piercy A.,
  Sangrajrang S., Sankaranarayanan R., Sauvaget C., Simpson S., Smith R., PhD;
  Steliarova-Foucher E., Sullivan R., Wagner D., & Elizabeth Ward E. (2015).
  Global Cancer Facts & Figures. *American Cancer Society*.

Psychosocial oncology & palliative care program - dana-farber cancer institute (n.d.).
  Retrieved October 1, 2018, from
  https://www.dana-farber.org/adult-psychosocial-oncology/

Caughey, R. W. and Michels, K. B. (2009), Birth weight and childhood leukemia: A
  meta-analysis and review of the current evidence. *International Journal of
  Cancer*, 124: 2658-2670.

Kamihara, J., Bourdeaut, F., Foulkes, W. D., Molenaar, J. J., Mossé, Y. P., Nakagawara,
  A., Brodeur, G. M. (2017). Retinoblastoma and Neuroblastoma Predisposition
  and Surveillance. Clinical Cancer Research,23(13).

Marcus J. Psychosocial Issues in Pediatric Oncology. *The Ochsner Journal*.
  2012;12(3):211-215.

Types of Surgery for Cancer Treatment. (n.d.). Stanford Health Care. Retrieved from
  https://stanfordhealthcare.org/medical-treatments/c/cancer-surgery/types.html

Radiation Therapy in Children. (n.d.). CureSearch for Children's Cancer. Retrieved from
  https://curesearch.org/Radiation-Therapy-in-Children

Johnson, A. V. (2018). An Update: Genetic Mutations and Childhood Cancers. *The
  Journal for Nurse Practitioners,14*(4), 230-237.

AJ, M., TT, M., PC, R., & J, C. (2014). An international survey of nutritional practices in
  low - and middle-income countries: a report from the International Society of
  Pediatric Oncology (SIOP) PODC Nutrition Working Group. *Eur J Clin Nutr.
  68(12):1341-5.* Retrieved from
  https://www.researchgate.net/profile/Alexia_Murphy_Alford/publication/2635508
  22_An_international_survey_of_nutritional_practices_in_low-and_middle-incom
  e_countries_A_report_from_the_International_Society_of_Pediatric_Oncology_S
  IOP_PODC_Nutrition_Working_Group/links/00b7d53b3d80a54066000000.pdf?
  origin=publication_detail

Albanti, I., & Lehmann, L. (2018, October 17). GHI Client Interview [Personal
  interview].

Aleman, B., Moser, E., Nuver, J., Suter, T., Maraldo, M., Specht, L., Vrieling, C., &
  Darby, S. (2014). Cardiovascular disease after cancer therapy. *EJC Suppl, 12*(1),
  18-28. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4250533/

Allessandra, S., Emanuela, R., Federico, A., Ana, M., & Maselli, T. (2011, July 05).
  Nutritional status at diagnosis is related to clinical outcomes in children and
  adolescents with cancer: A perspective from Central America. Retrieved from

https://ac-els-cdn-com.ezp-prod1.hul.harvard.edu/S095980491100400X/1-s2.0-S0
95980491100400X-main.pdf?_tid=1de35eda-4207-47fb-bf6a-91b78f4a2b2a&acd
nat=1539558906_d369edddaa178a1db39e297201cc92fa

Bacheldor, Beth. (2014). RFID News Roundup. *RFID Journal*, 4,
https://www.rfidjournal.com/articles/pdf?12062

Bauer, J., Jürgens, H., & Frühwald, M. C. (2011, March). Retrieved from
https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3065754/

Hewitt, M., Weiner, S. L., & Simone, J. V. (2003). *Late Effects of Childhood Cancer*.
National Academies Press (US). Retrieved from
https://www.ncbi.nlm.nih.gov/books/NBK221735/

Brandon, K., & Gordon, A. (2018, September 19). Dana Farber Nurse Interviews
[Personal interview].

Domek, Contreras-Roldan, O'leary, Bull, Furniss, Kempe, & Asturias. (2016). SMS text
message reminders to improve infant vaccination coverage in Guatemala: A pilot
randomized controlled trial. *Vaccine, 34*(21), 2437-2443.
https://www-sciencedirect-com.ezp-prod1.hul.harvard.edu/science/article/pii/S02
64410X16300615

Escalante, H. J., Montes-Y-Gómez, M., González, J. A., Gómez-Gil, P., Altamirano, L.,
Reyes, C. A., . . . Rosales, A. (2012). Acute leukemia classification by ensemble
particle swarm model selection. *Artificial Intelligence in Medicine, 55*(3),
163-175. doi:10.1016/j.artmed.2012.03.005
https://www-clinicalkey-com.ezp-prod1.hul.harvard.edu/#!/content/playContent/1
-s2.0-S0933365712000450?returnurl=null&referrer=null&scrollTo=%23fn0005

First Comprehensive Pediatric Hematology/Oncology Initiative Launched in Africa - The
ASCO Post. (n.d.). Retrieved October 17, 2018.
http://www.ascopost.com/issues/july-10-2017/first-comprehensive-pediatric-hema
tologyoncology-initiative-launched-in-africa/

The Global Health Initiative (GHI) at Dana-Farber/Boston Children's. (n.d.). Retreived
from
http://www.danafarberbostonchildrens.org/why-choose-us/global-health.aspx

Gyamfi, A., Mensah, K., Oduro, G., Donkor, P., & Mock, C. Barriers and facilitators to
Electronic Medical Records usage in the Emergency Centre at Komfo Anokye
Teaching Hospital, Kumasi-Ghana. *AFJEM, 7*(4), 177-182. Retrieved from
https://www.afjem.org/article/S2211-419X(16)30109-4/abstract

Improved Diagnostics & Patient Outcomes. (n.d.). Retrieved from
https://www.healthit.gov/topic/health-it-basics/improved-diagnostics-patient-outc
omes

Longombe, N. A., Nalbandyan, A., Perez, M., Zia, N., & Kandasamy, S. (2018,
September 18). GLOBE Scholars Interview [Personal interview].

Piñeros, M., Frech, S., Frazier, L., Laversanne, M., Barnoya, J., Garrido, C., & Morgan,
D. (2018). Advancing Reliable Data for Cancer Control in the Central America
Four Region. *Journal of Global Oncology, 4*(4), 1-11.
http://ascopubs.org/doi/10.1200/JGO.2016.008227

Poissant, L., Pereira, J., Tamblyn, R., & Kawasumi, Y. (2005). The Impact of Electronic

Health Records on Time Efficiency of Physicians and Nurses: A Systematic Review. *Journal of the American Medical Informatics Association, 12*(5), 505-516.

Reaman, G., Sposto, R., Sensel, M., Lange, B., Feusner, J., Heerema, N., Leonard, M., Holmes, E., Sather, H., Pendergrass, T., Johnstone, H., & Uckun F. (1999). Treatment outcome and prognostic factors for infants with acute lymphoblastic leukemia treated on two consecutive trials of the Children's Cancer Group. *J Clin Oncol, 17*(2), 445-455. Retrieved from https://www.ncbi.nlm.nih.gov/pubmed/10080584

Sala, A., Rossi, E., Antillon, F., Molina, A. L., de Maselli, T., Bonilla, M., … Barr, R. (2012). Nutritional status at diagnosis is related to clinical outcomes in children and adolescents with cancer: A perspective from Central America. *European Journal of Cancer*, *48*(2), 243–252. https://doi.org/10.1016/j.ejca.2011.06.006

Schwebel, & Larimer. (2018). Using text message reminders in health care services: A narrative literature review. *Internet Interventions, 13*, 82-104. https://www.sciencedirect.com/science/article/pii/S2214782918300022

Solution. (n.d.). Retrieved October 21, 2018, from http://emmunifyv3.weebly.com/solution.html

Hébert, A., Boucher, M., Guimont, C., & Weiss, M. (2017). Effect of measuring vital signs on recognition and treatment of septic children. *Paediatrics & Child Health,22*(1), 13-16. doi:10.1093/pch/pxw003

Spitzer, J. (2018, June 18). *Walmart patent proposes storing patient records on the blockchain*. Retrieved from https://www.beckershospitalreview.com/healthcare-information-technology/walmart-patent-would-store-patient-records-on-blockchain.html

Tierney, W., Achieng, M., Baker, E., Bell, A., Biondich, P., Braitstein P., Kayiwa, D., & Yiannoutsos, C. (2010). Experience implementing electronic health records inwah three East African countries. *Medinfo 2010 - Proceedings of the 13th World Congress on Medical Informatics, 160, 371-375.* https://utexas.influuent.utsystem.edu/en/publications/experience-implementing-electronic-health-records-in-three-east-a

Top Countries/Markets by Smartphone Penetration & Users. (n.d.). Retrieved October 21, 2018, from https://newzoo.com/insights/rankings/top-50-countries-by-smartphone-penetration-and-users/

Wahhab, H. T. (2015). Classification of Acute Leukemia Using Image Processing and Machine Learning Techniques. Retrieved from http://studentsrepo.um.edu.my/5925/1/Thesis.pdf

Wearable devices for EMR system. (2018, February 8). Retrieved October 21, 2018, from https://www.vinfotech.com/blog/healthcare/wearable-device-integration-with-emr-system

What are the Key Statistics for Childhood Leukemia. (2016, February 3). Retrieved October 21, 2018, from https://www.cancer.org/cancer/leukemia-in-children/about/key-statistics.html

Rogers, A. E. (2004). *Keeping patients safe: Transforming the work environment of*

nurses. Retrieved December 18, 2018.

SparkFun Single Lead Heart Rate Monitor - AD8232. (n.d.). Retrieved December 16, 2018, from https://www.sparkfun.com/products/12650

Shaffer, F., & Ginsberg, J. P. (2017). An Overview of Heart Rate Variability Metrics and Norms. Frontiers in Public Health,5(258). doi:10.3389/fpubh.2017.00258

Fish, R. M., & Geddes, L. A. (2009). Conduction of electrical current to and through the human body: a review. Eplasty, 9, e44.

USB LiIon/LiPoly charger. (n.d.). Retrieved December 16, 2018, from https://www.adafruit.com/product/259

Stefan, Daniela Cristina, Rodriguez-Galindo, Carlos (Eds.). (2014). *Pediatric Hematology-Oncology in Countries with Limited Resources*. Springer-Verlag New York.

Kumar, Anand, et al. "Duration of Hypotension before Initiation of Effective Antimicrobial Therapy Is the Critical Determinant of Survival in Human Septic Shock." Critical Care Medicine, vol. 34, no. 6, June 2006, pp. 1589–96. PubMed.

What is Neutropenia and How Is It Treated? (2018, April 9). Retrieved December 19, 2018, from https://blog.dana-farber.org/insight/2018/04/what-is-neutropenia-and-how-is-it-treated/

Appendix I: Investigate Phase Figures
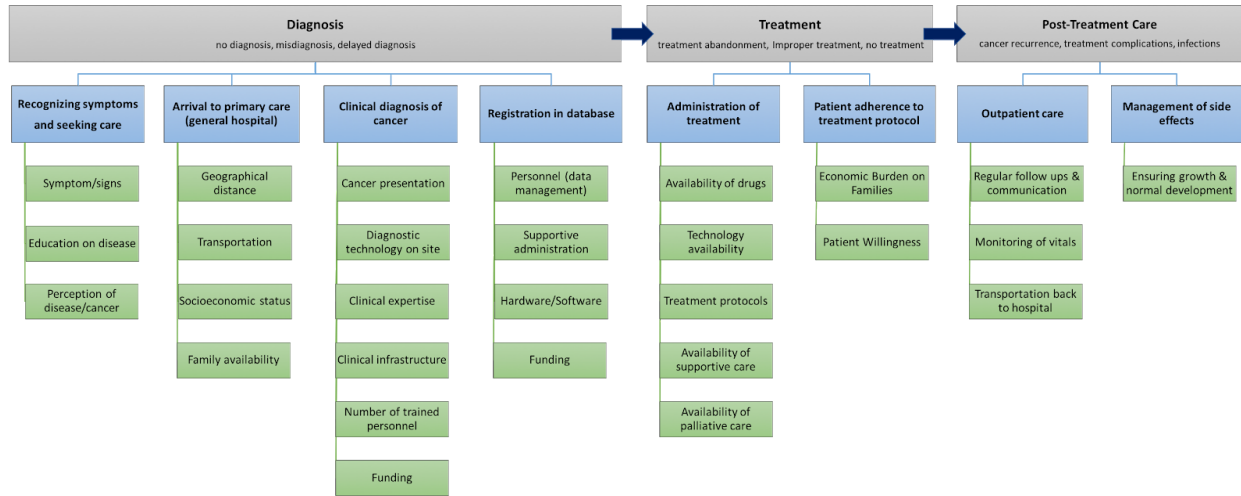
Figure i.1: Flowchart organizing data collection by factors influencing stages of pediatric cancer patient's care - diagnosis, treatment, and post-treatment care.
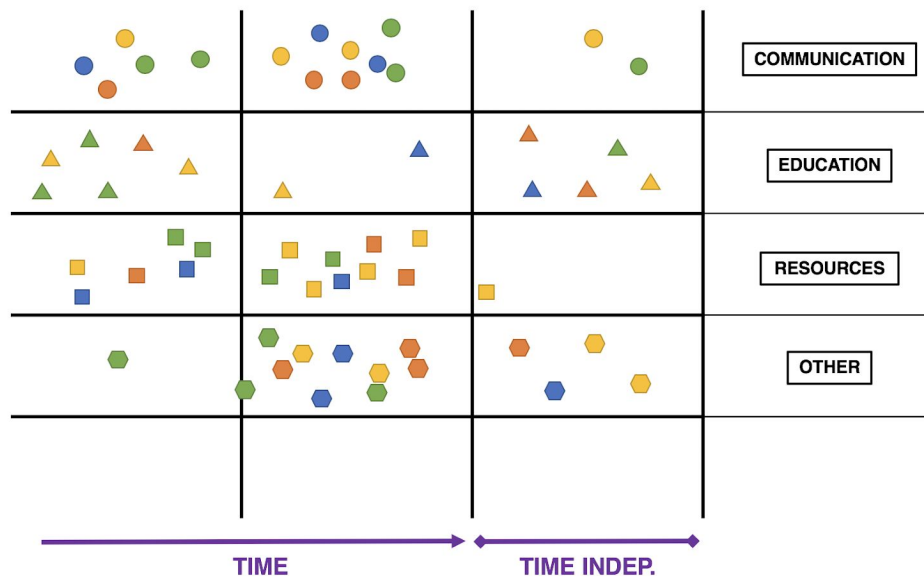
Figure i.2: First iteration of synthesis organizing data by graphing factors affecting patient survival against timeline of a patient's treatment highlighting their temporal relationship.
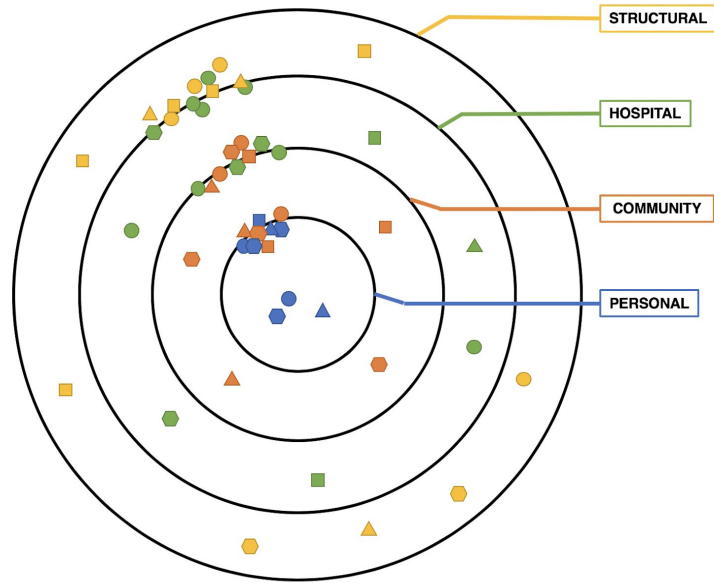
Figure i.3: The second iteration of visual mapping analyzing the actors involved in patient care organized as an egg model with categories of healthcare system, patient and family, and the larger community.
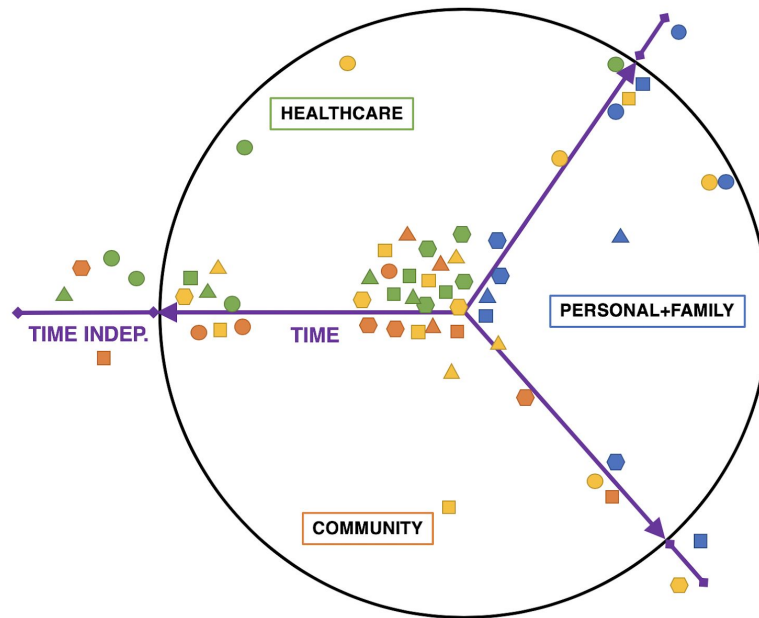
Figure i.4: Temporal pie chart organizing data by combining time dependency and actor-reactor mapping highlighting that patient care breaks down in a short term scale at the interface between actors.

Figure i.5: Visual representation of data collected analyzing actors (depicted as the inscribed rectangles) and their actions over time (depicted clockwise starting from the bottom left-hand corner).
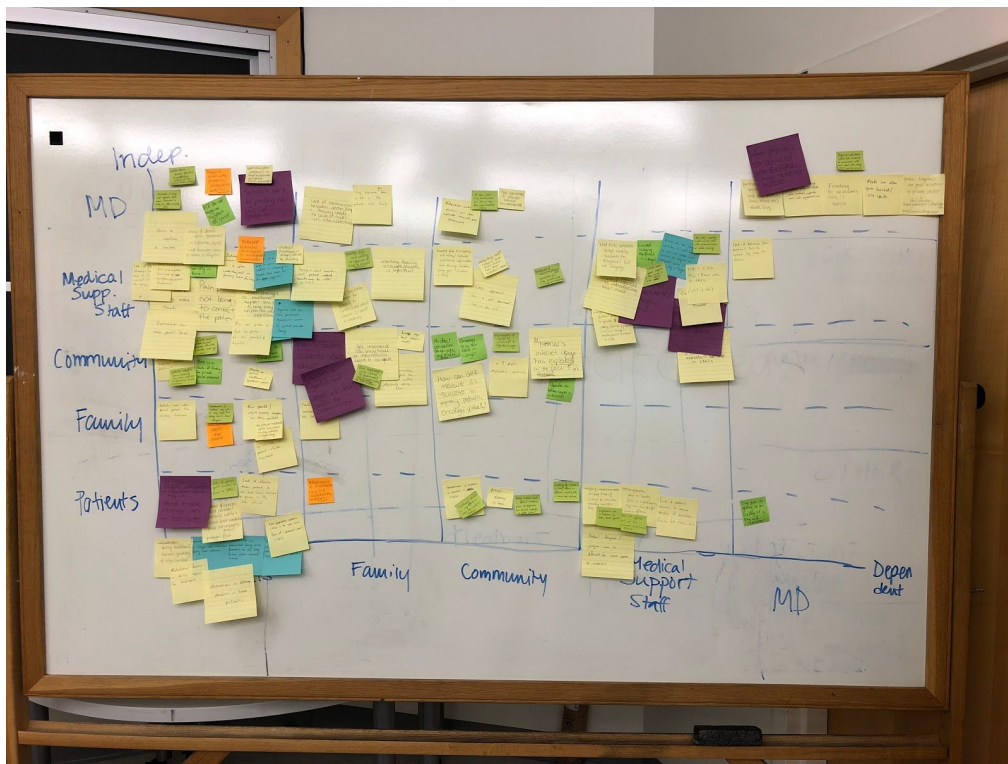
Figure i.6: Image of another model of visual data representation; here, a matrix is formed with the actors (MD, medical supplementary staff, community, family, and patients) across the x- and y-axes.
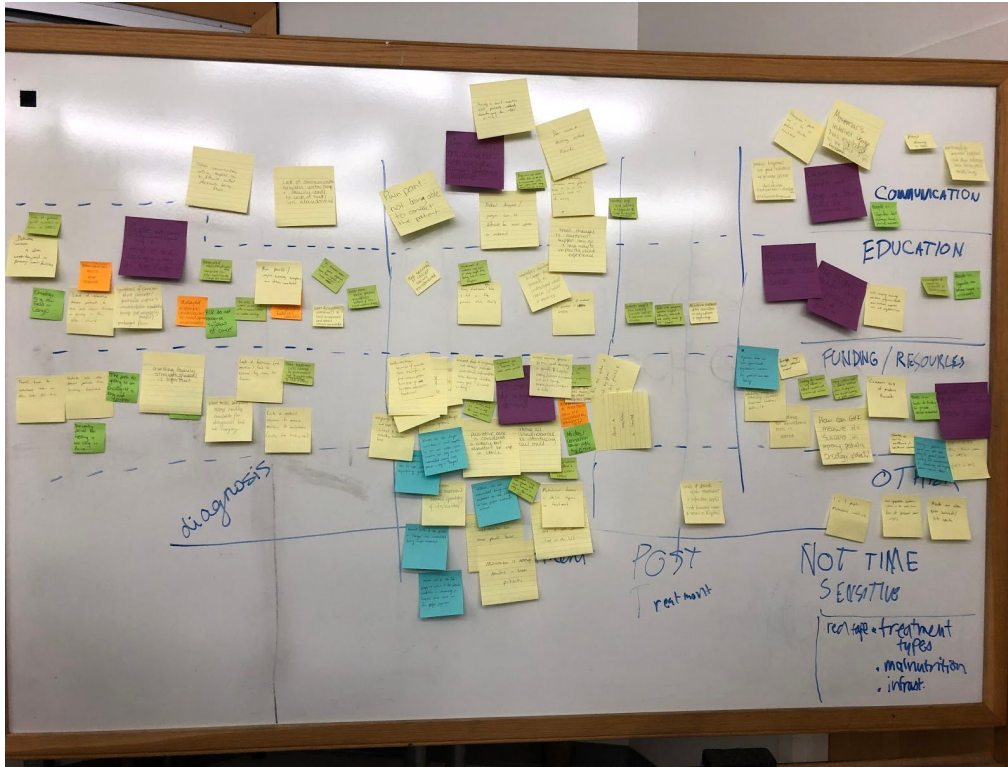
Figure i.7: A visual representation of our data with time (relative to healthcare process, including diagnosis, treatment, post-treatment, and not time-sensitive issues) against themes (communication, education, funding/research, other).
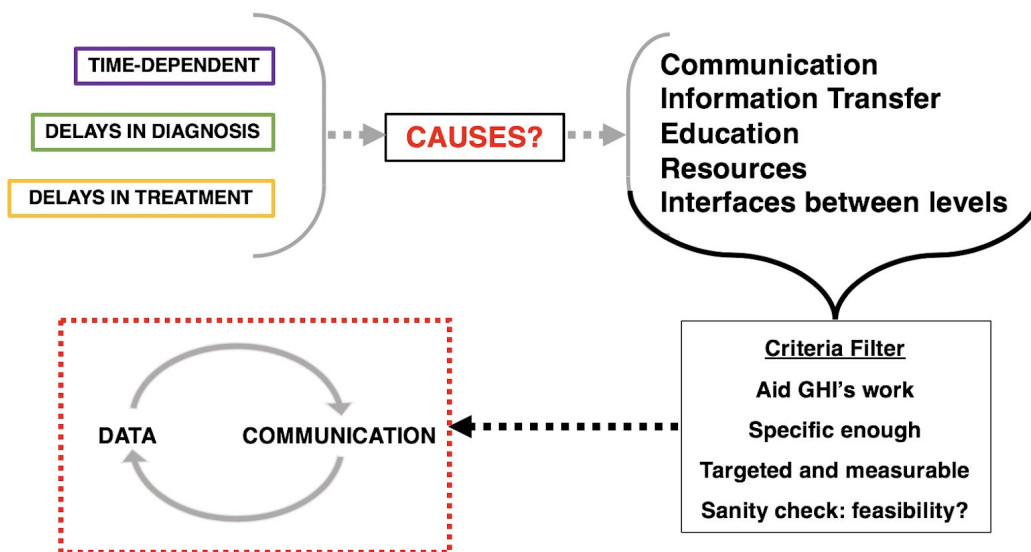
Figure i.8: Overall synthesis workflow and criteria filtering process leading to the essence of the problem: collection and communication of data. Causes of the problems identified in the visual mapping exercise were subjected to criteria to determine the most nuanced, approachable issue.

Figure i.9: Criteria Questions

- Does solving this problem meet client (GHI) needs and aid GHI's mission?
- Is this problem specific?
- Is this problem impactful? Is it important to a lot of people?
- Can we work in a space that is scalable, as opposed to a limited area?
- Do we have data to back up our description of this problem?
- Can we possibly work in this area, access data, and impact in our short timeline?

Appendix II: Ideate Initial Prototypes

Figure ii.1: Prototype 1: Nutrition Software - "Nutritrack" Software
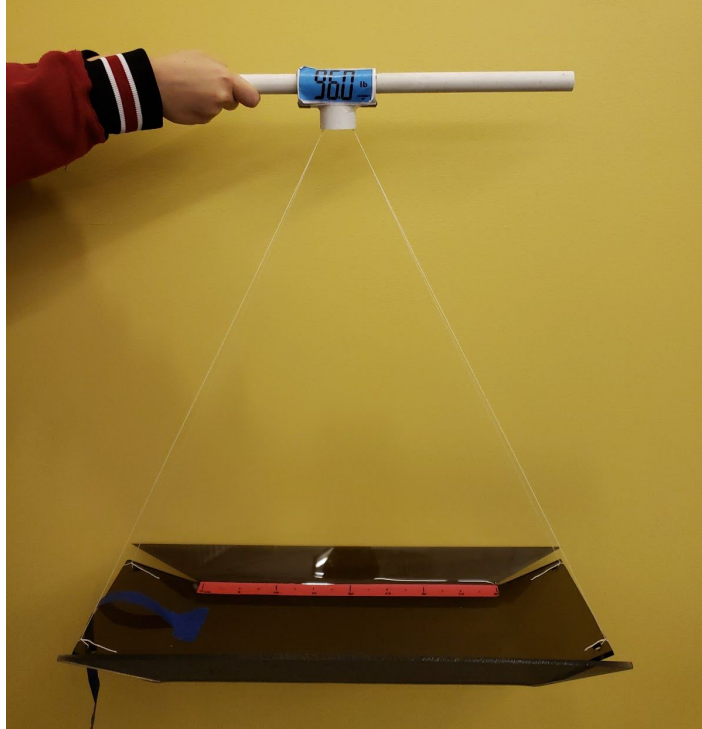
Figure ii.2: Prototype 2: Infant Anthropometric Multi-Tool

Figure ii.3: Prototype 3: Wearable QR code medical records

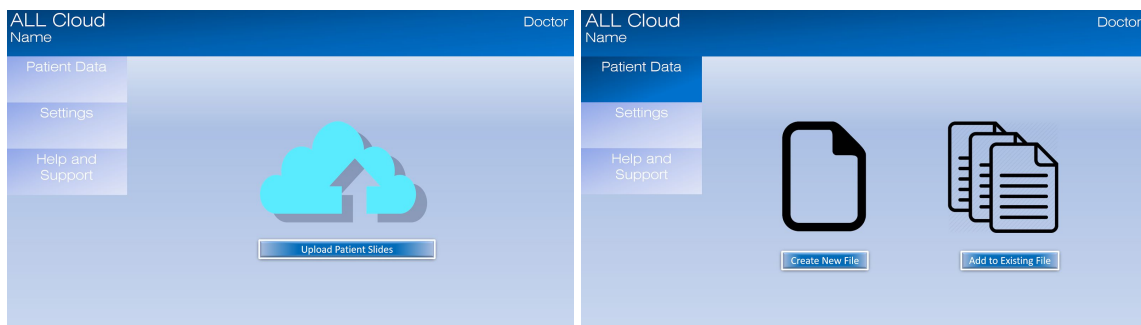*Figure ii.4: Prototype 4: Appointment Tracker and Reminder.*

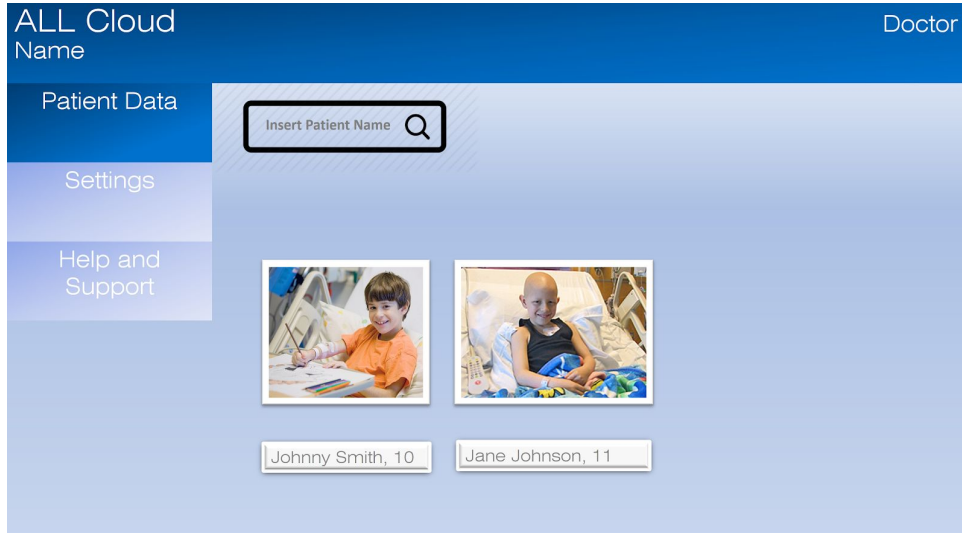*Figure ii.5: Prototype 5: ML Diagnosis Data Upload Page*

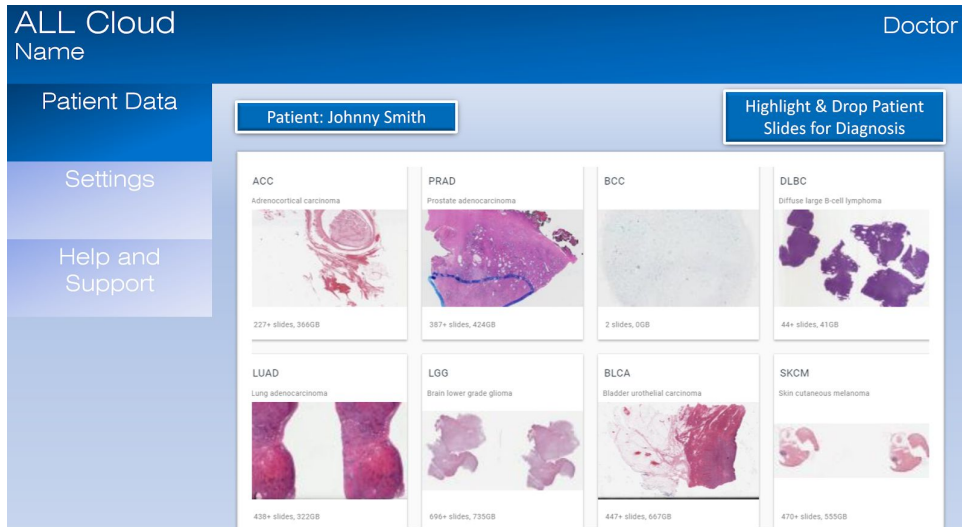*Figure ii.6: Prototype 5: ML Diagnosis Patient Listing Page*

Figure ii.7: Prototype 5: ML Diagnosis

Figure ii.9: Prototype 6: Communication Platform Registration Page
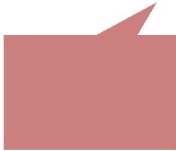
Figure ii.10: Prototype 6: Communication Platform Chat Page

Appendix III: Sensor Testing Protocols

Temperature Sensor

Objective

After having build circuitry for the temperature sensor, we will be testing the accuracy and sensitivity of the temperature sensor to medical grade thermometer readings. We are testing two locations: the underarm and the upper center chest region against each other and against a standard thermometer.

Procedure

For underarm:

1. Wipe underarm with dry towel before and between measurements
2. Use thermometer to measure body temperature and record.
3. Apply the temperature sensor to underarm region. Hold sensor in place until readings have stabilized on the display
4. Repeat for at least 15 measurements.
5. Save the temperature readings

For chest:

6. Use cloth/medical tape to insulate the temperature sensor, and ensure that ambient temperature is not affecting the reading
7. Apply sensor to mid-upper chest. Be sure to secure with tape. Do not use your hand to hold the temperature sensor onto your chest.
8. Repeat for at least 15 measurements
9. Save the temperature readings

Respiratory Rate

We are currently working with two models to derive respiratory rate from ECG.

1. The first is based of the Pan-Tompkins R wave detection algorithm. Matlab code can be found:
2. The second model operates with a slope based threshold of R peak detection. Matlab code can be found:

Both models will be tested first in a simulation, where an existing, publicly available data set is used to test the accuracy of algorithms being used. Data set can be found at http://peterhcharlton.github.io/RRest/bidmc_dataset.html

Secondly, we will use live data collected from our ECG circuit against our algorithms. The code used to transfer ECG data from an oscilloscope through an Arduino Uno to personal computers can be found here:
https://docs.google.com/document/d/1eHdjEWkSZUmghC8nrPvQEIgXHcT6EakaWmy0-uZ1Lds/edit

Live Data for Heart and Respiratory Rates

Heart Rate

We need to test our heart rate monitoring algorithm for two things: QRS detection and heart rate calculation. These two parameters determine the effectiveness of the algorithm.

Procedure

1. Modify the algorithm to also write data to a CSV file which we will later analyze in MATLAB to determine true heart rate and peaks. This should be at the same sampling frequency as is used in the arduino algorithm.
2. Code to determine the location of QRS complexes, not just the number over each interval.
3. Run the algorithm on a subject for five minutes to determine an average heart rate over that interval.
4. Stop recording.
5. Use the accurate algorithms available in MATLAB to determine the locations of R waves.
6. Compare this to the arduino-based algorithmically determined R wave locations. Determine percent of R waves accurately identified.
7. Calculate heart rate in MATLAB. Compare to arduino results. Determine percent error.
8. Have the subject elevate her heart rate through exercise. This will simulate a patient with potentially septic heart conditions.
9. Repeat steps 3-7 for this scenario.
10. Repeat steps 3-9 for ideally two more subjects to obtain a variety of different heart data.

Respiratory Rate

Assumes RR code uses same peak detection as in HR code
Given that RR interval detection uses the same code as is used in the HR code, we only need to test the respiratory algorithm for the rate it determines, not its peak detection.

Procedure

1. Modify the algorithm to also write data to a CSV file which we will later analyze in MATLAB to determine true heart rate and peaks. This should be at the same sampling frequency as is used in the arduino algorithm.
2. When running the algorithm, count the subject's breaths.
3. Run the algorithm on a subject for five minutes to determine average respiratory rate over the interval.
4. Stop recording.
5. Compare the arduino-derived respiratory rate to the observed respiratory rate.
6. Repeat steps 2-5 but have the subject talk during the recording. This will determine if we can still determine an accurate respiratory rate while the child is holding a conversation.
7. Have the subject perform somewhat vigorous exercise to elevate his respiratory rate.
8. Repeat steps 2-5 with the elevated respiration.
9. Repeat steps 2-8 with two more subjects.

ECG Location Optimization

Objective

We will be using ECG data to derive respiratory rate and heart rate. Therefore, we aim to optimize the location and type of electrodes used for ECG data collection. We will be comparing the relative signal strength- determined through peak amplitude and noise level- determine by deviations caused by coughing and moving from different placements under three conditions: normal breathing, fast breathing, and deep breathing.

Procedure
Lead placement on right side of chest.

1. Place electrodes in all positions V1-V6 depicted on the diagram
2. Breathe normally for 1 minute after the ECG reading has calibrated. Take a picture of the   resulting ECG graph
3. Breathe fast for 1 minute. Take a picture of the resulting ECG graph
4. Breathe deeply for 1 minute. Take a picture of the resulting ECG graph
5. Place electrodes in positions V1-V3. Repeat steps 2-4
6. Place electrodes in positions V4-V6. Repeat steps 2-4.

Lead placement on the left side of the chest

1. Mirror electrode placement on the left side of the chest. Repeat steps 2-6

Box Placement of ECG Leads

1. Place electrodes as depicted in the diagram below .Repeat steps 2-4.
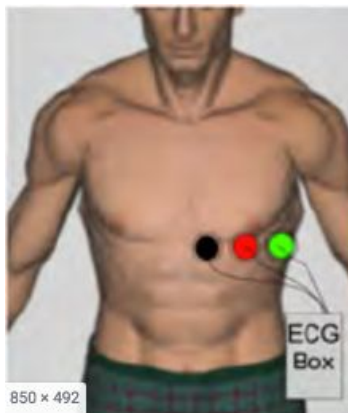
Appendix IV: Housing Test Protocols

If a procedure was completed and adequate data obtained, then it is noted along with the date on which the procedure was completed.

Casing Test Protocol Outline

Durability
- Compress with 700N
    - Safe approx of patient weight = 160 lb (16 year old boy)
- Drop from 5 ft
    - Average hospital bed = 1.777 ft x safety factor of 2 → 3.5ft
    - Average height of male (upper limit) doctor shoulders = 5 ft

Fluid-resistant / cleaning capability
- Submerge in tap water depth of 1 ft for 24 hours
- Submerge in salt water or HBSS (pH 4-5.5) depth of 1 ft for 24 hours
- Submerge in ethanol depth of 1 ft for 24 hours

Heat protection
- Measure temperature on surface of the case with running electronics inside → need to be less than 33 degrees Celsius

Interface with patient
- Binary evaluation of whether certain movements (arm, jumping, etc) will cause slippage of casing
- Effect of weight, angle of drooping

Adhesive Test Protocol Outline

Comfort

- Breathable woven design
  - Resistance of Evaporation (RET) score <= 13
- Biocompatible
  - Whether there are any adverse effects after wearing it for the intended extended period of time
  - Acute Systemic Toxicity Test (US Pharma Convention Class VI reactivity assessment)

Durability

- Should maintain contact with the skin for at least 12 hours without needing replacement (matching with typical nurse shifts)
- Testing protocol ideas
  - Physically wear the adhesives for >12 hours and see how many fall often (binary evaluation)
    - Combined into biocompatibility test
  - Surface adhesion energies? With adhesion tests using skin substitute?

Water resistance

- Remains adhered to "skin" after soaked/immersed in water or saline solution of pH 5.5 for 12 hours
  - Addressed in biocompatibility/durability test

Electrochemical safety

- <5 mAmp current passing through leads
  - Should comply since we're using commercial medical adhesives…

Case Testing Protocols

Durability Test

Goal: To determine whether case can withstand day-to-day compression or drop forces. This would use n=3 replicates per material.

Materials:

- 1-3x Assembled case (locked, with replacement battery and PCB inside)
  - PVA
  - ABS
- Hydraulic Press? (I am not sure what the machine is called in the ALLab)
- Meter Stick

Compression Procedure (for a given replicate of a material)

1. Place assembled and secured case inside of the hydraulic press
2. Calibrate machine to apply 500N of force (baseline needed, approximating patient weight of 160 lb of 16 year old boy) for 5 seconds
3. Inspect case for damage
   a. Locking mechanism still works
   b. Frame contains no fractures or cracks
   c. Case material is free from indentations or scratches > 1mm
4. Repeat process with increasing 50N of force until either
   a. Case is damaged (see Step 3)
   b. Depression of lid by more than 5 mm
   c. 1000N reached
5. Intended output: bar chart (with error bar) comparing maximum withstood force between materials

Drop Procedure - COMPLETED 12/5/18

1. Place assembled and secured case inside of the hydraulic press
2. Ensure *pre-drop* functionality of temperature sensors, ECG sensors, and PCB
3. Release case with three different orientations to ground from 2 ft
4. Inspect case for damage
   a. Locking mechanism still works
   b. Frame contains no fractures or cracks
   c. Surface of case is free from indentations and inconsistencies
5. Test *post-drop* functionality of temperature sensors, ECG sensors, and PCB
6. Repeat procedure on given replicate 3 times
7. Repeat Step 3-Step 6 for a drop height of 3 ft and 5 ft
8. Intended output: bar chart comparing number of drops that maintained functioning components between materials

Fluid Resistance Test

Goal: To determine change in case material properties with exposure to different fluids.

Materials

- 3 empty cases of each material
  - PVA
  - ABS
- 3 8-cm depth plastic containers to hold cases
- Tap water (control case)
- HBSS solution
- 70% ethanol

Duration: 12 hours

Procedure

1. Fill tap water, salt water, and ethanol into their respective containers
2. Place one closed case in each container of fluid. A weight may need to be added inside to prevent floating.
3. Wait 12 hours.
4. Inspect case for damage
   a. Locking mechanism still works
   b. No moisture inside of the case
   c. Case material has no obvious changes in flexibility or bending properties
5. Intended output: table with binary evaluation whether each material passed each fluid test

Heat Protection Test

Overview + Motivation: Our device is to be clipped on to patient's clothing and have the electrodes stuck onto the patient's body at various places. The driving question would be: "Does the skin experience an increase in temperature above a certain threshold, whereby the patient would experience discomfort?" Our device may produce heat from the PCB, battery, power management unit (charger) (which are all enclosed in the case), and temperature sensor. Patients will be wearing both the case and temperature sensor probe in close proximity with the skin, so it is important that our device does not produce excessive heat which may cause skin irritation or burns.

Goal: Monitor surface temperature of case under operation. This would use n=3 replicates for per material.

Materials
- 3 empty cases of each material
    - PVA
    - ABS
- External temperature sensing probe (not the one part of device)
- Arduino, temperature reading software
- Medical adhesive tape (MAT) 9907HTW

Duration: 12 hours (Patients will wear device for a minimum of 12 hours without taking it off. Pads are designed to give 3 days of wear, a conservative estimate would be to take the temperature measurements of the device's components for 3 days, but seeing as time is limited, the testing period will be kept to 12 hours)

Procedures
1. Record temperature of case surfaces before running.
2. Begin operation of the softwares to continuously measure ECG and temperature.
3. Secure the external temperature sensing probes with MAT on 1) the main device case and 2) the temperature sensor case.
4. Detect and record temperature from each probe continuously for the duration of the experiment.
5. Intended output: line plot over time for increasing temperature of material and include/indicate a pain threshold value

If Time Allows
Verify comfort in a blinded way with an actual user to assess if the temperatures reached ever reach a level of discomfort (from the patient's perspective).

Temperature Sensing Integrity Test

Overview + Motivation: There is a need to test how much heat the temperature sensor produces within its housing under load. Should it heat the housing excessively, there will be heat transferred from the housing to the skin, which could distort the amount of heat that the temperature measures and thus distort the accuracy of our sensor. Question to be answered: "Is the temperature measurement accuracy compromised when the sensor is placed within the housing?"

Goal: Determine the effect on the temperature reading from the heating up of temperature sensor by comparing the measurements of the temperature sensor with housing versus those of the sensor with no housing.

Duration: 2 hours

Procedures

1. Place a temperature sensor with no housing on a point far from the temperature sensor housing, place a temperature sensing probe near the temperature sensing housing and place a temperature sensor probe on the casing of the housing.
2. Operate the temperature sensor under maximum load for two hours (while software records the measurements of temperature probes).
3. Intended output: line graphs over time for temperature readings on the two sensors and analyze by calculating percent error relative to the control case

Clip Weight-Stretch Test - COMPLETED 12/3/18

Goal: To determine how much dragging the weight of the case material has on the device interface with patient attire. Testing the angle at which the clothing sits on the body - with respect to the vertical - versus the angle of the clothing at the site where the device is clipped onto the patient. Total of 3 (case material) x 3 (clip types) x 5 (clothing material) test conditions with 3 replicates each, for 135 tests.

Materials

- Assembled case
  - PVA
  - ABS
- Different clips
  - Metal
  - Otter plastic
- Ruler
- Various clothing materials
  - Khakis/stiff canvas material
  - Nylon athletic material

Procedures

1. Weigh and record the empty case.
2. The resting angle for a given fabric at the chosen point of contact will be measured from the vertical when no device is clipped to the patient.
3. The device will then be attached to the patient at the chosen point of contact.
4. The angle will then be measured again at the point of contact and be recorded.
5. Repeat removal and attachment of case for three replicates.
6. Intended output: spreadsheet recording various combinations (case + clip + material + angle) and the severity of its effect on the device's interaction with the patient… does functionality, comfort, and mobility increase or decrease?

Clip Movement Stability Test - COMPLETED 12/5/18

Overview + Motivation: To interface correctly with the patient, the case should remain clipped onto the patient and remain unbroken under daily stress-inducing activities, such being brushed against by a person passing in the passage.

Goal: Comparing the ability of the device to remain clipped to the patient when no force is applied versus when the force is applied, mimicking casual contact.

Materials
- Different clips
  - Metal
  - Otter plastic
- Various clothing materials
  - Khakis/stiff canvas material
  - Nylon athletic material
- Simulated activities
  - Vertical shaking (z-direction)
  - Horizontal shaking (x-direction)
  - Forward shaking (y-direction)

Procedure
1. Clip device on to given clothing material.
2. Apply shaking force from all 3 dimensions - x, y and z. Whether the device slips off or not for each of these experiments will then be recorded.
3. Repeat 3 times for each different area where the device can be clipped.

Adhesive Testing Protocols

Biocompatibility + Durability Test - COMPLETED 11/28/18

Goal: Compare biocompatibility and sustained adhesion of different adhesives. This would use n = 2x4 replicates, assuming same physiological response on left vs right side with four test subjects, for each of two locations and four adhesives, totaling 64 test samples.

Materials (all 3M medical adhesives)

1. 9917 (double-sided)
2. 9907HTW
3. 9834
4. 2477P (double-sided)

Duration: 12 hours

Test Conditions

- Under arm (approx place for temperature sensor) vs wrist -- 2 adhesives on each test subject
- Shower vs no shower -- 2 test subjects for each

Procedures

1. Cut out 4 rectangular pieces (3cm x 3cm) of each adhesive, per person.
2. For double-sided adhesives, add a thin layer of gauze to the backing.
3. Place adhesives in a rectangle formation, as shown in Figure 1, under the arm with 5mm in between adhesives. Repeat this rectangle formation on your wrist.
4. Repeat for other arm in both locations.
5. Leave the adhesive on for 12 hours. Two people cannot expose adhesives to water for the duration of the experiment, while two should immerse them for 15 minutes or "take a shower."
6. Count how many of the adhesive patches remain in contact with the skin. Evaluate what degree of redness (scale of 1-5) for each patch.
   a. Need to peel consistently between test subjects and locations.
7. Intended output: bar chart comparing how many of the adhesives remained stuck on skin + bar chart comparing skin responses

Breathability Tests - COMPLETED 11/28/18

Goal: Measure rate of moisture vapour transmission through the different adhesive materials.

Materials
1. 9917 (double-sided)
2. 9907HTW
3. 9834
4. 2477P (double-sided)
5. Paper towel
6. Plastic cup with diameter x inch
7. HBSS solution

Duration: 12 hours

Procedures
1. Cut out 6 square pieces of each adhesive (or paper towel as paper towel) with side length 10 cm.
2. Weigh cup.
3. Fill the cup with 50 g of saline solution.
4. Fasten each piece of adhesive to a cup, sealing it completely.
5. After 12 hours, weigh cup and subtract from initial mass value to determine how much moisture has passed through the fabric.
6. Intended output: bar chart with error bars for how much fluid has passed through

Electrochemical Safety Test

Goal

Asses potential for corrosion, as well as potential impacts if corrosion were to occur. Measure charge and discharge rate of batteries at different temperatures, and evaluate lifetime.

Materials

Battery

Battery specification sheet
Multimeter
Battery health measurement tool

Test Conditions

For battery health measurement, specifically in regard to battery lifetime, tests must be run at different temperatures.

For corrosive rates, tests can be run at one temperature, or a variety of different temperatures.

Duration

24 hours for corrosive
Battery health will be dependant on procedure of measurement tool

Procedures

For specifics, follow procedure provided by measurement tool
General Procedure:
Connect battery to test probes
Select data type to record
Transfer data to computer

Appendix V: Main Board and Temperature Buddy Board Full Schematics