


HEARTBEAT:

DESIGN AND DEVELOPMENT OF A HEADPHONE-MOUNTED INFRARED HEART RATE MONITOR

Orion Buske, Chris Neils, Mike Regnier
Department of Bioengineering, University of Washington



ABSTRACT:

A headphone-mounted photoplethysmographic (PPG) heart rate monitor was developed in order to provide a new method for continuous personal health monitoring. A pair of headphones was augmented with infrared LED/phototransistor sensors and was able to measure the light absorbance through the concha of the ear in order to isolate a PPG waveform. A receiver circuit was designed to amplify and filter the resulting signal, which was then digitally analyzed in a Labview user interface to identify the heartbeats and display the results. The device was tested on twelve subjects across eight experiments, for which the signals from two headphone monitors, a pulse oximeter, and an electrocardiogram control were simultaneously measured. Across all experiments, the headphone monitors performed as well or better than the pulse oximeter, with an average accuracy of around 85%, ranging from around 88% during no motion to 78% during heavy respiration. There was significant variation between the accuracies of the two headphone monitors across experiments and subjects. However, the accuracies of the headphone monitors were very similar to that of the pulse oximeter for both of these dimensions. Suggested areas of future work include improvement of the receiver circuit filtering, headphone fit, and performance of the user interface.

TABLE OF CONTENTS

1	INTRODUCTION	2
1.1	PROJECT DEFINITION	2
1.2	SIGNIFICANCE	3
1.3	SOCIAL, ETHICAL, ECONOMIC ISSUES	5
1.4	TECHNICAL BACKGROUND	6
1.5	TECHNICAL CHALLENGES	11
2	DESIGN	13
2.1	DESIGN AND RESEARCH PLAN PROPOSED IN BIOEN 481	13
2.2	MATERIALS AND METHODS	14
2.3	DESIGN OF EXPERIMENTS	21
3	RESULTS	24
3.1	TIMELINE.....	24
3.2	DETAILED CHRONOLOGICAL DESIGN NARRATIVE	24
3.3	FINAL SYSTEM.....	33
3.4	EXPERIMENTAL RESULTS.....	39
3.5	ANALYSIS AND CONCLUSIONS.....	43
3.6	SUGGESTIONS FOR FUTURE WORK.....	45
	ACKNOWLEDGEMENTS	48
	REFERENCES	49
	APPENDIX	52
A1.	DATA ANALYSIS SCRIPTS.....	52

1 INTRODUCTION

1.1 PROJECT DEFINITION

The purpose of this project was to design, develop, and test a headphone-mounted, light-based heart rate monitor. A standard headphone was augmented with a near-infrared sensor able to detect a photoplethysmographic (PPG) signal through the concha of the ear. A hardware circuit was developed to receive and filter this PPG signal in the analog domain in order to extract the pulsatile component that corresponds to the heartbeats. An accompanying Labview user interface was programmed to accept the analog PPG signal, perform limited digital processing and time-based analyses, and display the current heart rate, the recent heartbeat history, and a heartbeat indicator. Multiple iterations of design and testing were performed on the headphone sensor device, the supporting sensor circuitry, and the computer user interface.

The device and circuit system was then tested, in duplicate, across 12 subjects of various height (69 ± 5 in), weight (158 ± 30 lbs), sex (58% male), and fitness level (3.1 ± 0.7), although the ages of the subjects were very similar (21.2 ± 1.1 years old). For each subject, the accuracies of two headphone monitors and a standard pulse oximeter were measured with respect to an electrocardiogram (ECG) in eight different tests. These tests assessed the performance of the devices under conditions of varying motion, heart rate, and blood pressure. The data were then statistically analyzed to determine the effect of these variables on the accuracy of each of the devices.

1.2 SIGNIFICANCE

Overweight and obesity accounted for \$51 to \$78 billion in national costs in 1998 alone, over 60% of Americans are classified as either overweight or obese, and heart disease is still the primary cause of death in the US.^{1,2} Good diet and exercise are the best methods for avoiding these health problems, but even though the average American watches over four hours of television *per day* according to a report by A.C. Nielsen Co., a 2006 Gallup poll found the average American only gets three hours of moderate exercise *per week*. This can hopefully be improved by encouraging people to take personal responsibility for their health and by creating new technologies to promote this through simplicity and convenience. This project developed out of the idea of combining listening to music with physiological monitoring to produce a pair of headphones capable of non-invasively monitoring the vitals of the wearer. By making heart rate monitoring more accessible and enjoyable, it would provide the consumer with an ability to measure the health benefit of the exercise they get, both during the workout and in the long term.

The need for a new method of non-invasive physiological monitoring is present in many different settings. For instance, easy vital monitoring during sports exercise or for personal fitness management, clinical settings where patient cooperation with current monitoring practices is low, or settings where finger-clip, chest-band, or wrist-watch monitors are not ideal (e.g. during physical activity where these devices may be restrictive). Personal physical monitoring devices and non-invasive medical diagnostics are becoming more pervasive, and with well over 110 million iPods sold to date³ (other estimates predict closer to 200 million) there are a lot of

¹ Finkelstein, Fiebelkorn, and Wang, "National Medical Spending Attributable To Overweight And Obesity: How Much, And Who's Paying?"

² Kung et al., "Deaths: final data for 2005."

³ "Apple's "The beat goes on" special event."

headphones being worn. These two developments can be brought together with a consumer headphone-mounted heart rate monitor for use with a mobile music player. Further, with the most recent release of the iPhone 3.0 SDK, developers can interface hardware devices directly with the iPhone, providing the potential for the iPhone to take readings from the device, store data, and display relevant information.

In addition to this consumer market, there are clinical applications for comfortable, continuous heart rate monitors as well. Current cardiac rehabilitation practice involves patients attending supervised, in-clinic exercise while being monitored. Many of these patients need to be monitored continuously with an electrocardiogram (ECG), but for lower risk patients this monitoring is not always necessary.⁴ A headphone heart rate monitor would therefore be beneficial to provide a less obstructive means to continuously monitor patient vitals during exercise. Also, for long-term rehabilitation, tele-home care (telephone calls and home visits), or personal fitness plans, health improvements could be tracked over time by recording the individual's vitals continuously throughout each day. Real-time monitoring of heart rate would be useful for ensuring that athletes remain within a reasonable heart rate zone during intense exercise. Current athletic devices for real-time monitoring tend to be mobility-restrictive (e.g. require chest straps or the wearer to press the fingers of one hand onto electrical contacts on a watch on the other hand), although a finger slip (Intervent) and ring sensor (LifeSpan) have gone into commercial production to try to avoid this problem. The reviews for both devices are mixed, however, because their performances dramatically diminish during periods of motion. Motion artifacts are a critical problem faced by PPG sensors and developing a device that is relatively

⁴ Fletcher et al., "Exercise Standards : A Statement for Healthcare Professionals From the American Heart Association."

robust to motion would allow PPG-based devices to be effectively used during exercise.⁵ Such a motion-resistant headphone heart rate monitor would increase the average consumer's ability to personally monitor his or her health or provide athletes with a tool for more efficient, safer training.

1.3 SOCIAL, ETHICAL, ECONOMIC ISSUES

The main issues concerning this device are acceptance, privacy, accuracy, and cost. The impact of this device directly corresponds to its degree of acceptance and the effect it has on the exercise regiment of the owner. It is possible that, no matter how convenient fitness monitors get, most people will not get a healthy amount of exercise. As far as privacy, if the device were further improved to utilize Bluetooth wireless communication between the earpiece and a receiver, encryption may need to be utilized to prevent eavesdropping. Because of the low sensitivity of the information being transmitted and the short transmission range, such concerns are likely minimal. The accuracy of the device is an important consideration. If it is not sufficiently accurate in a variety of situations, it will have low appeal to consumers. If its accuracy is very high, it has the potential to be able to be used in clinical situations as a substitute for standard heart rate monitors. The cost of the device will have a very large impact on its popularity. The device was very inexpensive to build and could likely be mass-produced for the same price as a standard pair of sport headphones.

⁵ Grajales and Nicolaescu, "Wearable multisensor heart rate monitor."

1.4 TECHNICAL BACKGROUND

1.4.1 Theory

In order to identify a heartbeat, the heart rate monitor uses a method known as photoplethysmography (PPG). This method measures the absorption of light by a tissue to determine properties of that tissue. There are two main forms of PPG, transmittance and reflectance, which differ in the relative orientation of the light emitter, light receiver, and tissue being investigated.⁶ In transmittance PPG, the emitter, usually a light-emitting diode (LED), is directed into one side of the tissue and the receiver is aligned directly on the opposite side of the tissue and oriented to face the LED. The receiver then measures the intensity of light transmitted through the tissue. In contrast, reflectance PPG works by placing both the emitter and receiver on the same side of the tissue, with both components directed into the tissue. In this case, the receiver measures the intensity of light reflected by the tissue. In both cases, changes in the relative absorption of the tissue directly result in changes in the measured light intensity at the receiver. Each time the subject's heart beats, the perfusion of blood in the tissue rises and then falls, and the absorbance of light in the infrared (IR) spectrum increases and then decreases accordingly because blood exhibits strong absorbance in the IR spectrum. This enables heartbeats to be detected by measuring the pulsations in the intensity of light at the receiver. There is some constant absorbance by the skin, muscle, tendons, bone, and average blood volume, but variation in the absorbance is mainly due to the pulsatile blood flow through the tissue. The variable, or AC, component of the signal that corresponds to the heartbeats can then be isolated from the static, or DC, component of the signal that corresponds to the background tissue with a high-pass filter.

⁶ "Principles of Pulse Oximetry Technology."

In situations where sample rate, computing power, or delay are significantly limited or quantization error is an issue, it may be necessary to perform signal filtering in the analog domain prior to digital conversion and processing. Such analog filters fundamentally take advantage of the frequency response of RC, RL, and LC circuits to generate high-, low-, or band-pass filters which operate on continuous, analog signals in real-time. These filters, at the most basic level, are of two types, passive and active. Passive filters consist solely of unpowered circuit components, usually resistors, capacitors, and/or inductors, and provide a theoretically linear frequency response over any frequency range, but provide no amplification. Active filters incorporate powered components such as operation amplifiers into the circuit to provide amplification and can provide sharper frequency cutoffs and more powerful filters, but they are limited by the performance of the active components in the filter.

1.4.2 Review of literature

Since the invention of the electrocardiogram (ECG), technology for continuous heart rate monitoring has been improving. With the discovery of photoplethysmography (PPG), another source of physiological data was provided that could be recorded from almost any skin area.⁷ Decades later, almost all heart rate monitors still utilize either one or the other of these two fundamentally different data sources.

For decades, development of comfortable and portable heart rate monitors has been explored because of the value of such a device for exercise physiology and training.⁸ Heart rate and heart rate variability have been shown to provide good indicators of exercise intensity and fitness

⁷ Challoner and Ramsay, "A photoelectric plethysmograph for the measurement of cutaneous blood flow."

⁸ Seaward et al., "The Precision and Accuracy of a portable heart rate monitor."

level, respectively, which can help inform healthy, effective exercise behavior.⁹ There are a variety of heart rate monitors currently on the market, and most utilize an ECG signal in one way or another. Many consist of a chest strap that measures an ECG signal and wirelessly transmits it to a nearby watch-style device.¹⁰ Most newer strap-less versions acquire an ECG signal by requiring the wearer to touch metal contacts on a wristwatch with the opposite hand, but this is only useful for querying the current heart rate, not continuous recording. PPG sensors, on the other hand, provide continuous, direct monitoring of blood flow, which may help explain the growing interest in them as continuous health monitoring continues to gain popularity. Other reasons for the increasing research include the affordability, simplicity, and reliability of PPG devices.¹¹

As PPG has become better explored, additional applications have been uncovered. One additional benefit of using PPG over ECG is the presence of both heart and respiratory information in the signal. This provides the ability for an effective PPG-based device to gather both heart rate and respiratory rate information from the same probe,^{12,13} and has been found to provide signals of equal or superior quality to traditional ECG for heart rate and transthoracic impedance plethysmography for respiratory rate.¹⁴ Further, by adding another PPG sensor at a different wavelength, the proportion of oxy-hemoglobin to hemoglobin in the blood can be estimated, providing a measure of blood oxygenation. Thus, it is much simpler to add blood oxygenation measurement to a PPG-based heart rate monitor than one that is ECG-based. PPG

⁹ Achten and Jeukendrup, "Heart Rate Monitoring: Applications and Limitations."

¹⁰ Christensen and Kushner, "Commercial Program and Product Review Heart Rate Monitors."

¹¹ Allen, "Photoplethysmography and its application in clinical physiological measurement."

¹² K. Nakajima, T. Tamura, and H. Miike, "Monitoring of heart and respiratory rates by photoplethysmography using a digital filtering technique."

¹³ Shelley, "Photoplethysmography: Beyond the Calculation of Arterial Oxygen Saturation and Heart Rate."

¹⁴ Johansson, Öberg, and Sedin, "Monitoring of Heart and Respiratory Rates in Newborn Infants Using a New Photoplethysmographic Technique."

signals are able to be derived from a variety of locations on the body, the most typical being finger¹⁵, earlobe¹⁶, and forehead¹⁷. Recent studies have indicated that the ear is at least as good, if not better in some ways for PPG than the traditional finger pulp.¹⁸ PPG signals have also been able to be obtained without any contact with the subject's skin, making them very useful for burn victims or where the subject needs to be mechanically isolated.^{19,20} PPG also provides an option for monitoring patients during magnetic resonance imaging (MRI) investigation, where the metal required to implement ECG-based monitors cannot be used. With PPG-based monitors, a fiber-optic sensor can be attached to the patient and connect to a processing device out of the range of the MRI.²¹

PPG-based heart rate detection presents a number of difficult challenges as well, including obfuscation by venous pulsations, motion resistance, and insufficient blood perfusion to the area being investigated. There have been many methods, both in hardware and software, developed to overcome each of these problems. One common method of reducing the effect of motion artifacts is to use measurements from accelerometers built into the device to mathematically compensate for the motion, usually with the aid of an adaptive filter.^{22,23,24,25,26} A variant of this method uses additional phototransistors to detect the frequency of the motion artifacts, which can

¹⁵ Sherebrin and Sherebrin, "Frequency analysis of the peripheral pulse wave detected in the finger with a photoplethysmograph."

¹⁶ Stern, "Ear Lobe Photoplethysmography."

¹⁷ Casati et al., "Forehead Reflectance Oximetry: A Clinical Comparison with Conventional Digit Sensors during Laparotomic and Laparoscopic Abdominal Surgery."

¹⁸ Awad et al., "Analysis of the Ear Pulse Oximeter Waveform."

¹⁹ Cheang and Smith, "An Overview of Non-contact Photoplethysmography."

²⁰ Verkruysse, Svaasand, and Nelson, "Remote plethysmographic imaging using ambient light."

²¹ Lindberg, Ugnell, and Öberg, "Monitoring of respiratory and heart rates using a fibre-optic sensor."

²² Foo and Wilson, "A computational system to optimise noise rejection in photoplethysmography signals during motion or poor perfusion states."

²³ Celka et al., "Motion Resistant Earphone Located Infrared based Heart Rate Measurement Device."

²⁴ Renevey et al., "Wrist-located pulse detection using IR signals, activity and nonlinear artifact cancellation."

²⁵ Foo, "Motion Artefact Reduction of the Photoplethysmographic Signal in Pulse Transit Time Measurement."

²⁶ Relente and Sison, "Characterization and adaptive filtering of motion artifacts in pulse oximetry using accelerometers."

then be subtracted from the frequency spectrum of the signal to identify the real heart rate.²⁷ One novel approach also measures light absorption at a wavelength where water dominates the absorption to compensate for motion.²⁸ One software-based approach analyzes the signal data using algorithms that are specifically designed to be more robust to motion effects.²⁹ The motion sensitivity of PPG sensors is still an issue, however, and that was one of the main technical challenges of this project, especially if the device were to be used during exercise.

While this project was underway, two papers were discovered which describe research that bears striking similarity to this project. These papers concern the research presented in 2004 by Celke et al.³⁰ and the work presented in 2007 by Brodersen et al.³¹. Both research papers cover the development of in-ear photoplethysmographic devices to continuously monitor heart rate. The similarity between this project and the “PulseEar” developed by Celke et al. is remarkable, as their device is even built upon a headphone scaffold, although apparently not with the intent of playing music. The PulseEar is patented and under commercial development. Despite these concurrent advancements in the field, however, work was continued on this project as planned.

Recent developments in e-textiles,³² wireless sensors,³³ and wearable devices^{34,35} promise to significantly increase the pervasiveness of personal health monitoring in the coming years. E-textiles and clothing-embedded sensors have become increasingly popular in recent years as a

²⁷ Lei Wang, Lo, and Guang-Zhong Yang, “Multichannel Reflective PPG Earpiece Sensor With Passive Motion Cancellation.”

²⁸ Debreczeny and Baker, “Pulse oximetry motion artifact rejection using near infrared absorption by water.”

²⁹ Yong-Sheng Yan and Yuan-Ting Zhang, “An Efficient Motion-Resistant Method for Wearable Pulse Oximeter.”

³⁰ Celka et al., “Motion Resistant Earphone Located Infrared based Heart Rate Measurement Device.”

³¹ Brodersen et al., “In-Ear Acquisition of Vital Signs Discloses New Chances for Preventive Continuous Cardiovascular Monitoring.”

³² Teng and Zhang, “M-health: trends in wearable medical devices.”

³³ Gay and Leijdekkers, “A Health Monitoring System Using Smart Phones and Wearable Sensors.”

³⁴ Asada et al., “Mobile monitoring with wearable photoplethysmographic biosensors.”

³⁵ Hung, Zhang, and Tai, “Wearable medical devices for tele-home healthcare.”

way of continuously monitoring vitals through a wearable device, such as Nike+iPod³⁶, NuMetrex sports bra³⁷, and recent research into “smart” fabrics woven of carbon nanotube-coated cotton³⁸. Current research in mobile health technology has used wireless ECG sensors and smart phones to continuously provide personalized heart monitoring for high-risk cardiac patients.³⁹ As an extension of these efforts, the device developed in this project is another step toward convenient health monitoring and a healthier world.

1.5 TECHNICAL CHALLENGES

At the outset of this project, there were two known technical issues that needed to be addressed. Most predominant was the issue of motion resistance and motion artifacts in the PPG signal. As previously described, PPG sensors are notorious for motion artifacts, and much work has been put into devising motion resistant sensors and motion compensation algorithms. This is likely one reason why ECG sensors dominate the current consumer heart rate monitor market, but significant research efforts are currently underway to improve PPG devices to broaden their application. For this project, this challenge was met through targeted device and filter design, with the attempt to try to reduce the motion of the device as much as possible, as well as eliminate motion artifacts from the captured signal.

Another challenge was the issue of extracting a viable PPG signal through the highly cartilaginous concha. As previously described, transmittance PPG signals have been taken through the ear lobe and reflectance PPG signals around the perimeter of the ear, but no research

³⁶ “Nike + iPod.”

³⁷ “Strapless Heart Rate Monitor Clothes.”

³⁸ Shim et al., “Smart Electronic Yarns and Wearable Fabrics for Human Biomonitoring made by Carbon Nanotube Coating with Polyelectrolytes.”

³⁹ Gay and Leijdekkers, “A Health Monitoring System Using Smart Phones and Wearable Sensors.”

was found which explored transmittance PPG through anywhere other than the ear lobe. Further, the signal to noise ratio was so poor in initial trials that digital signal processing was not a reasonable approach. A major focus of the project was thus on the challenge of acquiring an adequate signal and effectively filtering it in the analog domain.

2 DESIGN

2.1 DESIGN AND RESEARCH PLAN PROPOSED IN BIOEN 481

As originally proposed, the project plan was to design, build, and test a headphone-based physiological monitor and supporting software. This monitor was to consist of a set of sensors, infrared-light, red-light, and potentially electrocardiogram (ECG), embedded into a pair of stereo headphones. This device was to accept a stereo audio input signal through a stereo headphone, or tip-ring-sleeve (TRS), jack and output a data signal through a stereo microphone, also TRS, jack. The TRS microphone jack was known to provide a limited amount of power, and it was intended for the device to be run directly from the computer power through this connection. The device would thus connect to both the headphone and microphone jacks on a personal computer, play music received through the headphone jack, and output measured physiological data through the microphone jack.

Supporting analysis software and a user interface were then to be designed for the personal computer to receive the physiological data from the microphone jack, perform any necessary de-multiplexing, filtering, and analysis, and then display the measured vitals in a descriptive, user-friendly way. The physiological data outputted by the device was intended to initially consist of just the single filtered signal received by the infrared-light sensor. Then, as sensors with additionally monitoring capabilities (red-light, ECG) were added to the device, sensor signals would be time-multiplexed before transmission to the computer, and then de-multiplexed by the computer software.

The development of the device was intended to be performed in a series of stages. Each stage consisted of adding a new sensor to the device, updating the hardware and software to handle the new signal, testing the device for motion resistance, accuracy, and ergonomics, and then refining the device, hardware, and software to correct any problems found. First, an infrared (IR) sensor was to be added to measure heart rate, then a red sensor to measure blood oxygenation when combined with the IR sensor, and finally, if time permitted, an electrocardiogram (ECG) or some other sensor would be added to test for an additional physiological signal such a brain activity or awake state.

As it turned out, the project never advanced past the first development threshold to be able to add an additional sensor. The device was never considered sufficiently resistant to motion to be able to justify complicating it by adding additional sensors. Thus, the focus of the project was restricted to designing and constructing the best headphone-mounted IR heart rate monitor possible. This left more time for trial and error, experimentation, and focus on designing a system which might be able to be used while the wearer is in motion, rather than giving up on motion and adding additional sensors whenever possible. Both routes are reasonable, but the former was taken for this project.

2.2 MATERIALS AND METHODS

2.2.1 Device construction

In the creation of the measurement device, a pair of Sony stereo headphones (model MDR-J10) was fitted with an IR LED (940 nm, 30°, 10 mW) and IR phototransistor (QSC112QT). Coated, stranded, 24-gauge speaker wiring was first soldered to each terminal of both the LED and

phototransistor. The exposed wiring was then covered with heat-shrink tubing to prevent contact with other wires or skin. The speaker of the headphone was then removed and a 7/64" drill bit used to widen the cylindrical support for the audio vent out the back of the headphone. After expanding the support, a small slot in the edge of the headphone was cut out for the phototransistor wiring. The phototransistor was then placed in the drilled hole, the wires routed out the slot, and the speaker snapped back into place, holding the phototransistor firmly in place. The LED was then hot-glued to the subject-side of the ear clip, past the peak, and angled such that the IR light was directed toward the base of the ear. The wires for the LED and phototransistor were cut to a length of approximately five feet and bundled and secured to the base of the headphone by wrapping a rubber band around the collection of wires until secure. This helped ensure that tugs on the wiring would not propagate to the individual components and would rather move the entire device as a whole.

2.2.2 Device testing

For the device testing, an IR sensor was embedded, by the above procedure, in both ears of a single pair of headphones. One was placed in each ear of the subject, and the voltage of the LED was set to 0.02 V below the lowest voltage at which either of the photoreceivers saturated. As a control, Kendall Q-Trace 5400 Ag/AgCl Resting ECG electrodes were attached to the top of both wrists of the subject, and a ground electrode was placed just above the subject's left hip. These three electrodes were connected to an A-M Systems, Inc Differential Amplifier, (Model 3000). The significant settings were the high-pass and low-pass filter frequencies and the gain, which were set to 0.1 Hz, 100 Hz, and 500, respectively. The gain was initially set to 1000, but after saturating during the measurements on one subject, it was reduced to 500. To provide a comparison with existing commercial heart rate monitors, a Nellcore N-200 Pulse Oximeter was

attached to the ring finger of the subject's non-dominant hand. In order to record when the device identified a heartbeat (and not the pulse waveform), the voltage across the speaker output was recorded. Because the device beeps whenever it identifies a heartbeat, this provided a clear signal from the pulse oximeter. The voltage signal from each of these four devices was simultaneously recorded at 50 ms per sample (software-timed) through a NI-DAQ (NI USB-6009).

The NI-DAQ measurements were recorded with a Labview VI that exported the measurements on the four channels to a data file for each subject and test scenario at 50ms per sample. These data files were analyzed at a later date with a set of MATLAB scripts (see 2.2.3 Data filtering and heartbeat identification and Appendix A1) that processed the data and determined the accuracy of all devices with respect to the ECG for every 30-second period in every file. In order to measure how the devices performed under varying blood pressure, the subjects were reclined and inclined on a tilt table. Blood pressure measurements were taken every minute with an Omron IntelliSense Automatic Digital Blood Pressure Monitor (Model HWM-737).

Three complications arose during testing. First, the blood pressure measurements were accidentally taken from the right arm of six of the twelve participants. Second, for six of the participants, the blood pressure cuff was on the same arm as the pulse oximeter, invalidating the pulse oximeter readings during this experiment because of insufficient blood flow whenever the blood pressure cuff tightened to make a measurement. To resolve these issues, the corresponding invalid blood pressure measurements and pulse oximeter readings were excluded from further analyses. Third, the sample rate ended up being insufficient to accurately record the R-wave of

the ECG waveform, and additional processing was performed to reconstruct the waveform as well as possible.

2.2.3 Data filtering and heartbeat identification

After collecting data from 96 trials (8 experiments with 12 subjects), MATLAB scripts were written to filter this data. The specifications for these scripts are explained in detail in Appendix A1, but a summary of the analytic methods will be included here. The waveforms for the pulse oximeter, headphone monitors, and ECG were filtered and then thresholded to identify heartbeats. The pulse oximeter signal was initially low-pass filtered with a 0.5 second (10 data point) moving average. The headphone monitor signals were initially high-pass filtered by subtracting a 0.4 second (8 data point) moving average. The ECG signal was band-pass filtered by subtracting a 0.5 second (10 data point) moving average and then low-pass filtered by applying a 2 second (20 data point) moving average.

Two additional filtering steps, failure identification and signal flipping, were performed just on the ECG signal. ECG failure was characterized by abnormally low (signal flat-lines) or high (signal spikes or saturates as electrodes detach and are reattached) standard deviations in the waveform. In an effort to remove these segments from the analyses, the ECG signal was then scaled by a 2 second (20 data point) moving standard deviation. Sections of the ECG signal with standard deviations below 0.01 or above 1.0 were discarded. Because the left and right ECG leads were occasionally reversed, the ECG signal orientation had to be standardized. Comparing the standard deviations of the positive subset of the waveform with that of the negative subset yielded an effective determiner. If the R-waves pointed downward, the standard deviation of the positive subset was much less than the negative subset, and the signal was inverted.

Finally, the signals were thresholded and heartbeats were identified as the midpoint of each above-zero segment. For the ECG signal, an additional post-processing step was performed to merge together pairs of peaks that correspond to the R and S waves of the same complex. If two peaks appeared within 6 data points, the second peak was ignored. This cutoff is reasonable because at 50 ms per sample, a peak pair separated by less than 6 data points corresponds to a heart rate above 250 bpm.

2.2.4 Data segmentation and heartbeat classification

After skipping the first 10 seconds of recording (to allow all devices to initialize and settle), the data from each trial was divided into 600-data-point samples (30 seconds at 50 ms per sample). For each sample, the accuracy of each device was calculated with reference to the ECG by classifying each device-registered heartbeat as either a true positive (device heartbeat with corresponding ECG heartbeat), false positive (device heartbeat with no corresponding ECG heartbeat), or false negative (ECG heartbeat with no corresponding device heartbeat). The concept of a true negative is not clearly defined for this application and its definition and incorporation would not provide any meaningful data so it was not included in the analyses.

Heartbeats were classified with a sliding window approach described in pseudocode in Appendix A1: `analyzedata.m`, to allow for variable device latency. To summarize this algorithm: for a set of sequential ECG heartbeats, h_1 , h_2 , h_3 , and h_4 , the first device heartbeat in the range $(h_1, h_3]$ was classified as a true positive (matching h_1). If there is no device heartbeat in $(h_1, h_3]$, record a false negative (missing match to h_1). Any remaining device heartbeats in the range $(h_1, h_2]$ were classified as false positives. This algorithm provides substantial leniency in the device latency,

but an adjacent false positive and false negative have the potential to be incorrectly classified as a true positive if they are near enough in time.

2.2.5 Accuracy analysis

Once the accuracy was computed for each device for each 30-second sample in the experiment, samples were filtered out in the following ways. If there were too few samples (under half the expected count, based upon the experiment length), all subject data for that experiment was discarded because the sampling rate was inadequate during that experiment. Else, if there were too many accuracy values, any excess data points were removed. If, on the other hand, there were the correct number of data points, the last sample for each device was discarded because it may have edge effects or be incomplete.

Additionally, a threshold was optionally imposed to censor periods of device failure in the following way. If the sample accuracy was under 50% for any of the devices being compared, that sample was thrown out because of a presumed device error. The corresponding samples in all other devices were also thrown out to keep the number of data points consistent and avoid biasing the results. This 50% threshold was an estimate based upon preliminary accuracy measurements for several trials, where most device errors resulted in sample accuracies around 30%, but if there were no errors the accuracies were at least 80%. Thus, 50% seemed to be a reasonable threshold to distinguish between samples contained device failures and those that did not.

2.2.6 *Statistical analyses*

Statistical analyses were performed on the data to compare the accuracies of the pulse oximeter and the two headphone monitors with respect to the ECG for a variety of conditions. The accuracies of the two headphone monitors were compared for each experiment over all subjects with paired t-tests. Similarly, the worse of the two headphone monitors was compared to the pulse oximeter with paired t-tests. An ANOVA was used to determine if the angle of the tilt table had a significant impact on accuracy for each of the devices. A two-way ANOVA was used to determine if accuracy was significantly affected by the experiment or the device, and another two-way ANOVA was used to assess the impact of the subject on the device accuracies. Finally, a linear regression was used to estimate the degree to which device accuracy was correlated with the time into the sampling.

2.2.7 *Costs*

A summary of the main expenditures incurred over the course of this project are summarized in Table 2.1, below.

Table 2.1: Summary of costs incurred for project expenses.

TRS Plugs	3.99
Modeling clay	10.95
Sony headphones x 2	25.98
50ft stranded wire	5.49
Ear buds	9.99
Hot glue gun	19.99
Glue sticks	4.49
IR LED x 6	3.60
IR receiver x 6	3.60
Resistors x 20	2.00
	\$90.08

2.3 DESIGN OF EXPERIMENTS

In order to assess the performance of the headphone device, an ECG signal was employed as a control signal. A commercially available pulse oximeter was included in the study for comparison as an example of current optical heart rate monitor technology.

Potentially influential independent variables that needed to be explored or accounted for included: 1) blood pressure, 2) motion, 3) heart rate, 4) headphone device variation, 5) subject (height, weight, sex, ear structure), and 6) variation with time. A regimen of tests was devised in order to examine the effect of each of the above variables on the accuracy of the headphone heart rate monitor and the pulse oximeter. In order to quantify variability in blood pressure, a blood pressure monitor was added for the relevant experiments. Accelerometers would have been necessary to precisely quantify the variability in motion, but instead a more qualitative measure was used. Twelve (12) subjects each performed a 50-minute regimen in which data was collected from 2 ear sensors (one in each ear), a pulse oximeter, an ECG, and optionally a blood pressure cuff. Voltage measurements for each device were obtained every 50 ms (software-timed) through a single NI-DAQ. The following is the testing regimen and devices recorded from for each test:

Table 2.2: Summary of testing regimen. Duration, conditions, and recording devices for each of eight experiments.

#	Duration (min)	Conditions	Ear monitor?	Pulse oximeter?	ECG?	BP monitor?
1	5	Sitting still	X	X	X	
2	10	General activity*	X	X	X	
3	3	Rapping fingers on table*	X	X	X	
	5	Tilt table (0 degrees)				
4	5	Tilt table (-45 degrees)*	X	X	X	X
5	5	Tilt table (0 degrees)*	X	X	X	X
6	5	Tilt table (+45 degrees)*	X	X	X	X
7	variable	Seated, holding breath	X	X	X	X
	5	Running up and down stairs				
8	5	Seated, cool-down*	X	X	X	

* Eating, drinking, seated motion, and talking were allowed and encouraged.

- Test 1 provides a control for heart rate and motion for all devices. For most other experiments, general motion was not restricted or quantified to provide a better estimate of actual device accuracy.
- Test 2 provides a qualitative, instead of quantitative, measure of the accuracy of the device under reasonable operating conditions. Because of the immobility of the power supplies and data recording and the lack of any available in-place exercise equipment, this test provided the most reasonable measure of general activity. Although not individually quantifiable, the effect averaged over all participants would provide insight into how well the device might work under actual operating conditions.
- Test 3 provides a measure of how well the pulse oximeter responds to extreme motion.
- The subject was allowed to rest for five minutes on a horizontal tilt table to let the subject's blood pressure stabilize.
- Test 4 tests the devices under decreased blood pressure.
- Test 5 provides a blood pressure control for tests 4 and 6 and an intermediate data point.
- Test 6 tests the devices under increased blood pressure.
- Test 7 tests the devices under increased thoracic pressure, decreased heart rate, and no respiration.
- Test 8 tests the devices under increased heart rate and severe respiration.

The testing protocol was established through a compromise between the competing factors of available participants, reasonable experiment length per participant, and the need for sufficient data to provide statistical power to analytic tests. In order to balance these factors, most experiments were performed for 5 minutes per subject, and the data were divided into 30-second

samples. This provides 10 samples per subject per experiment. With 12 subjects, this provided a reasonable number of samples for most statistical tests.

More technically, the variability in the accuracy data was unknown before any formal tests were performed. In preliminary trials, the standard deviations tended to have values around 10%.

Using this estimate, to identify a difference in accuracy of 5% between two devices with a 95% confidence level (with a two-tailed, paired t-test), at least 18 sample pairs were required. To identify a difference on one subject for one test, the test should thus have been at least 9 minutes long. On the other hand, if the data were compiled across subjects with 5-minute tests, only 2 subjects were really required for such a test. The testing protocol was thus extremely generous and should have provided enough data for sufficient statistical power.

3 RESULTS

3.1 TIMELINE

The final timeline of this project was:

Investigate TRS connection	06/15 - 07/23
Preliminary device design	07/03 – 07/23
Filter design	08/06 – 08/18
Filter refinement	08/20 – 10/18
Data acquisition and UI	10/18 – 02/08
Filter redesign from scratch	01/03 – 01/20
Filter analysis, testing	02/08 – 02/22
Device redesign	02/08 – 02/22
Experiment design	02/22 – 03/17
Final circuit fix	03/17 – 03/18
Subject signup	03/14 – 03/17
Device testing	03/17 – 03/20
Analysis scripting	03/20 – 03/24
Draft writeup	03/21 – 03/29

3.2 DETAILED CHRONOLOGICAL DESIGN NARRATIVE

3.2.1 *The switch from TRS to DAQ*

In the initial project proposal, it was intended for the device to output its signal to the computer through a TRS microphone jack, and to draw power from this jack as well. Even though it is unlikely that any non-prototype version of this device would use a TRS output, this specification was decided because of its novelty and convenience, as the computer is already necessary for processing and display. After experimentation, however, there developed two problems which resulted in eventually switching to another method of data acquisition.

The first of these problems is that although the TRS microphone jack does provide a some power, enough to power small, inexpensive microphones, this is nowhere near sufficient to

power the IR light-emitting diode (LED) needed for the emitter component of the first sensor.

When the power channel was loaded with just an IR LED, the voltage difference provided by the jack (originally 3.2 V) dropped to 1.1V and the current was only 0.8 mA, far below the standard 20 mA required to power an LED. Although the emitter component was unable to be powered by the microphone jack, the initial receiver circuit was able to be powered by this source.

Unfortunately, as soon as active components were added to the receiver circuit, the microphone jack was unable to provide sufficient power. A power supply was used instead to power both the emitter and receiver components of the IR sensor in the interest of maintaining focus on the project as a whole.

The second problem encountered with the TRS jack output was the way in which a signal needed to be presented for the signal channel to transmit it to the computer. The voltage on the signal channel was actually pushed high by the computer, necessitating a different approach to set the voltage on the channel. The receiver circuit was redesigned several times to try to accommodate this but eventually the TRS output was given up in favor of utilizing a National Instruments Data Acquisition Device (NI-DAQ). The NI-DAQ additionally provides analog to digital (A2D) conversion and comes with convenient Labview modules to incorporate the measured signal directly into a Labview Virtual Instrument (VI). This switch helped focus the project on building a working device system rather than trying to develop the interface with the TRS output.

3.2.2 Getting an initial signal

The first component built was a prototype headphone scaffold onto which a transmittance PPG IR sensor could be affixed to measure heart rate through the ear. Since most in-ear headphones

make contact with the cavum conchæ (see Figure 3.1, adapted from Bisaccia et al.⁴⁰), it made sense to have either the LED emitter or phototransistor receiver at this contact point and the other aligned on the other side of the ear, but the concha is thin and cartilaginous, so it was unclear whether or not a PPG signal would be able to be captured through this tissue. A mold of the experimenter's concha was made with modeling clay, baked, and hollowed to provide a scaffold for a speaker and sensors. Because the phototransistor is light-sensitive, it was placed in the cavum conchæ, allowing the mold to shield environmental light, and leaving the emitter to be aligned around the back of the ear to measure light transmittance. In order for the LED and phototransistor to be stably aligned and pressed against both sides of the ear, a sheet metal was bent into a U shape with one end attached to the concha mold and the LED mounted on the other. The sheet metal was, however, just flexible enough to allow the device to be fitted onto the ear. Wires were soldered to the LED and phototransistor to connect them, respectively, to the power supply and a breadboard on which an initial receiver circuit was built. The completed prototype is shown in Figure 3.2.




Figure 3.1: Anatomy of the outer ear with the cavum concha (and other features) labeled. (Adapted from Bisaccia et al.)

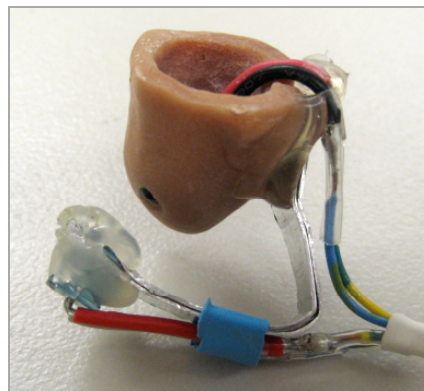


Figure 3.2: Completed ear monitor design prototype.

⁴⁰ Bisaccia et al., "The surgical correction of protuberant ears."

This initial circuit, shown in Figure 3.3, both converted the IR phototransistor current signal to a voltage signal, as well as provided passive high-pass filtering ($f_c = 0.16$ Hz). The resulting signal, however, was under 20mV peak-to-peak and was extremely sensitive to motion. Still, this rudimentary prototype was able to detect a PPG signal through the cartilaginous concha. The next stage of development focused on improving the circuitry to filter and amplify the PPG signal.




Figure 3.3: Initial phototransistor receiver circuit and high-pass filter.




Figure 3.4: 300 mV peak-to-peak PPG-based signal after initial amplification and band-pass filtering.

3.2.3 Improving the circuitry

First, an amplifier was added after the initial receiver-filter circuit to boost the signal to around 300 mV peak-to-peak. Next, an active, second-order, Sallen-Key Tschebyscheff (3dB passband ripple) low-pass filter ($f_c = 3.0$ Hz) was designed, constructed, and added in series after the amplifier.⁴¹ The circuit was then able to output an amplified, band-pass filtered PPG signal as shown in Figure 3.4. Interestingly, the low-pass filter was so aggressive that the normal characteristics of the PPG waveform were lost in the signal, which instead had been essentially reduced to a sinusoid.

⁴¹ Mancini, *Op Amps for Everyone: Design Reference (SLOD006B)*.

The next iteration of circuitry improvement resulted in trying new combinations and configurations of active high- and low-pass filters. First, the initial passive high-pass filter was replaced with an active high-pass filter of the same form as the low-pass filter mentioned previously, but the signal quality was not as good, so the old configuration was restored. The low-pass filter was then updated with a higher cutoff frequency ($f_c = 4.0$ Hz) to reduce the attenuation of reasonable 3 Hz (180 bpm) input signals. An additional high-pass filter ($f_c = 0.5$ Hz) was added to remove the offset generated by the low-pass filter, the signal was amplified again, and then passed through another low-pass filter ($f_c = 3.5$ Hz).

With two amplification stages, two second-order low-pass filters, one second-order high-pass filter, and one first-order high-pass filter, the signal was very clean, but a new problem presented itself. The symptom was a very large second peak of the end of every beat, resulting in what appeared to be two heartbeats for every one of the subject's. The cause seemed to be non-uniform gain in the pass-band. Because the filters had sharp frequency cutoffs at the cost of 3 dB ripples in the pass-band, these ripples seemed to be amplifying the post-heartbeat venous pulsations. If the subject heart rate was around 60 bpm, (1 Hz), the filtering amplified the venous pulsations and resulted in an output signal that was a 2 Hz sinusoid, where the second of each pair of peaks was slightly smaller in amplitude. This issue was addressed by changing filter component values to lower the cutoff frequency for the low-pass filters (to shift the pass-band ripple below where it would be a problem), switching the initial passive high-pass circuit and amplifier with a standard current to voltage converter, and removing the final low-pass filter. These improvements reduced the complexity of the circuit, improved the response time, and resulted in the semi-final circuit configuration shown in Figure 3.5.




Figure 3.5: Intermediate circuit configuration consisting of C2V conversion, low-pass, high-pass, and amplification stages.

In an effort to explore whether or not reducing the pass-band ripple (at the cost of softening the frequency cutoff) would improve the filtering behavior, the circuit was redesigned from scratch as a pair of fourth-order Chebyshev (0.5 dB pass-band ripple) filters, high-pass first and low-pass second. This new circuit did not suffer from the same half-beat amplification problem mentioned previously, but the settling time made it unacceptable for use. Where the original circuit had been designed in order to keep the settling time after a significant disruption (such as removing and replacing the device in the ear) to within a few seconds, the newly designed circuit took around 20 seconds to recover. Thus, the old circuit was restored and work on the user interface was started.

3.2.4 User interface design

The user interface (UI) was written as a Labview Virtual Instrument (VI), which provides convenient support for the NI-DAQ being used to record the IR receiver circuit's output signal. A VI was built to record the filtered PPG signal, identify and display the heartbeats, and report an approximation of the current heart rate. One complication faced in this design was the calculation of the time between heartbeats (necessary to compute the heart rate) from the number of data points between peaks in the signal. It ended up being difficult to know how long had actually passed between these two peaks because it was unknown when, in real-time, each data

point was actually measured. This is because the data acquisition was software-timed, resulting in incredibly large variability in the time between consecutive samples. Even if a sampling period of 50 ms was specified, it was not uncommon to see 100 or 200 ms between a pair of samples. Although hardware timing was considered, it would make the desired real-time nature of the UI extremely difficult to create. One approach to solving this would be to query the system clock at the same time every data point is acquired, but this would make the VI too slow to effectively process the data in real-time. Another approach, and the approach eventually taken, was to query the system time occasionally, after a set number of data points, and compute the average sample rate over that interval. The number of data points used was the sampling period, so if the sampling period is 50 ms, the system time will be requested every 50 samples, or 2.5 second. If the sample rate is higher, the relative error in sampling periods is likely higher, so sampling the time more often seemed reasonable. While on a small time scale there is a potential for large errors in the reported heart rate, on the scale of heartbeats it is unlikely that the time between beats differed dramatically from that which can be estimated from this estimated sampling rate. After getting the VI working sufficiently well, the device itself was revisited and redesigned.

3.2.5 Device redesign

The device was recreated on two different headphone scaffolds. The first was a pair of generic ear buds. The speaker was removed, holes drilled for the phototransistor and wires, but there was no natural way to place the LED around the back of the ear. A metal U clip similar to the one fashioned for the first device prototype was tried, but because the ear buds did not fit the ear as well as the mold, the U clip was not strong enough to hold the device against the ear. An over-the-ear clip was added but the amount of tension required to hold the LED precisely aligned with

the phototransistor was unachievable with the available materials. The weight of the new device was too great to hold up without an over-the-ear clip, but there was no clear way to firmly position the LED opposite the phototransistor. Because of the sensitivity to motion, a reasonable design needed to hold the LED and phototransistor in line, and needed to hold the device onto the ear to prevent the tissue from moving between the emitter and receiver.

During experimentation with LED position, it was discovered that if the LED power was increased, the LED could be placed almost anywhere around the ear and still provide illumination for the phototransistor to measure a PPG signal. Thus, the next design placed the LED out of the way at the top of the ear, but in such a place that motion would be restricted. A pair of Sony over-the-ear stereo headphones (MDR-J10) were purchased and outfitted with the IR sensor as shown later in Figure 3.6. This change in the arrangement of the sensor dramatically improved the motion resistance of the device. Previous to the change, the emitter and receiver needed to be precisely lined up. If either moved, the transmitted light intensity changed dramatically, burying the signal. A signal was able to be detected with extremely low LED current (2 mA) in this configuration, but any motion or adjustment destroyed the signal. With the alternate orientation, the LED current needed to be much higher (15 mA), but the signal could be detected despite moderate motion. The new orientation was adopted because improved motion resistance was favored over low power consumption.

3.2.6 Finalizing the circuit

In response to failures in initial testing on volunteers, the circuitry was redesigned and finalized. Prior this redesign, the circuit did little amplification of the signal prior to cascaded filtering. An instrumentation amplifier (INA 126) was added after the current to voltage converted but prior to

any filtering to amplify the signal ($G = 85$). The low-pass and high-pass filters were then swapped in order to remove the substantial offset created by the instrumentation amplifier. This new circuit exhibited superior noise filtering and motion resistance and was used for the duration of the project.

3.2.7 Testing

In order to assess the accuracy, reliability, and robustness of the heart rate monitor system, two versions of the device were created. Data was recorded from an electrocardiogram (ECG) standard, pulse oximeter, and two headphone heart rate monitors (one in each ear) simultaneously over a one-hour test. The first problem encountered had to do with the fact that the UI was designed to analyze one IR sensor measurements in real-time, not two. When the Labview VI was adjusted to perform the same analyses on each data stream, the computation load was too great and the lowest sampling period achievable was increased over two-fold to around 100 ms per sample. This sampling period is near the bottom limit of acceptable sampling rates for the headphones, but is completely unacceptable for the ECG and pulse oximetry data being simultaneously recorded through the NI-DAQ. Thus, although the UI still provides a real-time display for data from one IR sensor, it was abandoned for testing in favor of a simple VI that collected data from all four measurement channels and wrote the measurements to a data file. After testing was completed, a series of MATLAB scripts were written to analyze this data, identify heart beats in all signals, and measure the accuracy of the headphone devices and pulse oximeter with respect to the ECG signal.

During the analysis, it became apparent that the 50 ms sampling period was insufficient to accurately record the ECG signal. In fact, the R-wave of the heartbeat was completely missed

every few heartbeats, making it difficult to correctly extract the heartbeats from this control signal. Through informal inspection, it appeared that only around 75% of the ECG heartbeats could be correctly identified. This somewhat limited the ability to effectively measure the device accuracies, as the control could not be entirely trusted. In fact, better results were able to be reported by using the pulse oximeter as a control signal. This will be further discussed in the Analysis and Conclusions.

3.3 FINAL SYSTEM

The project resulted in the construction of a working heart rate measurement system, consisting of a headphone-mounted measurement device, a receiver circuit, and an accompanying Labview user interface.

3.3.1 *The device*




Figure 3.6: Photos of the completed device, consisting of a pair of headphone scaffolds augmented with IR PPG sensors to continuously monitor heart rate.

The final device (Figure 3.6) was created by embedding an IR sensor into a Sony stereo headphone scaffold. The IR phototransistor was embedded in the earpiece and the IR LED was affixed to the ear clip and pointed into the base of the ear. This illuminates the entire ear with IR light and allows the phototransistor to pick up a PPG signal almost regardless of where it is

pointed, so long as it is stable with respect to the ear. Because the LED is attached near where the ear clip makes contact with the top of the ear, it is not subject to the same severity of motion effects as if it were attached to the long end of the flexible ear clip. Additionally, the weight of the device helps hold the LED in position behind the ear, keeping the IR illumination of the ear relatively constant.

3.3.2 The receiver circuit




Figure 3.7: Completed receiver circuit, consisting of C2V conversion, amplification, high-pass filtering, low-pass filtering, and amplification stages (from left to right).

The final receiver circuit, shown in

Figure 3.7 above, accepts the current signal provided by the IR phototransistor and converts it to an amplified, filtered voltage signal. It processes the signal in five stages. The first stage converts the current signal to a voltage signal.⁴² The second stage takes the raw voltage signal and amplifies it ($G = 85$) with an instrumentation amplifier to allow the signal to be effectively filtered. The third stage is a second-order Chebyshev high-pass filter ($f_c = 0.40$ Hz) to remove the DC-offset imposed by the instrumentation amplifier. The fourth stage is a second-order Chebyshev low-pass filter ($f_c = 1.26$ Hz). The final stage amplifies ($G = 8$) the filtered signal again before it is read by the NI-DAQ. When modeled in MATLAB, the two cascaded filters of this circuit (ignoring gain components) have the frequency response shown in Figure 3.8.a and

⁴² Horowitz and Hill, *Art of Electronics*.

the impulse response shown in Figure 3.8.b. The settling time of these filters is around 1.7 seconds, as shown in Figure 3.8.b. From the INA126 documentation, the instrumentation amplifier should add less than 0.2 seconds to the settling time, making the overall circuit settling time around 2.0 seconds.

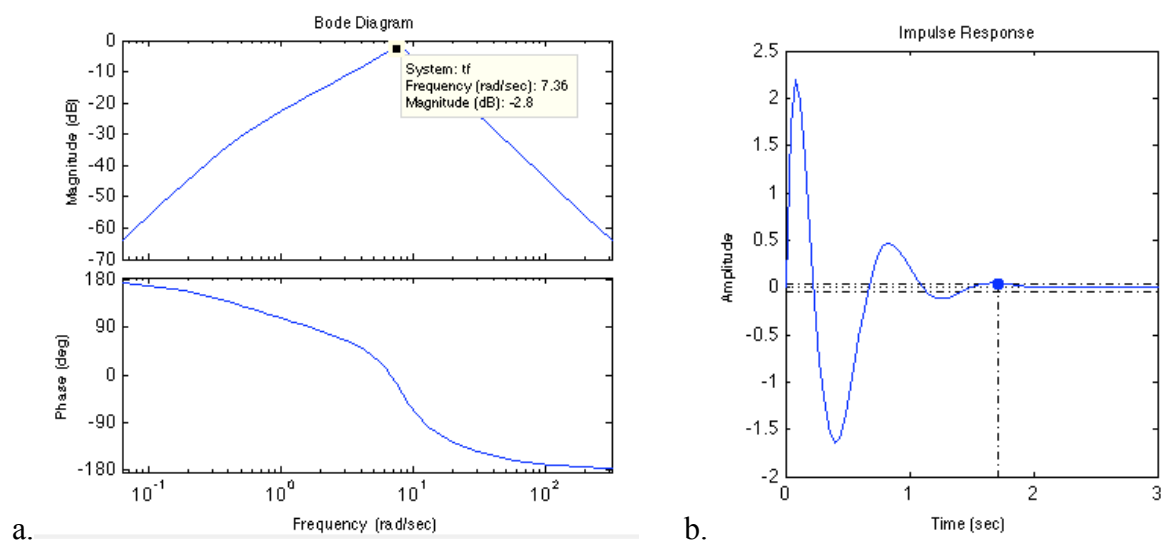


Figure 3.8: a) Magnitude and phase plot for cascaded filters in completed circuit (assuming unity-gain). Peak in frequency response corresponds to 1.2 Hz. b) Impulse response and settling time of filter components.

3.3.3 The user interface

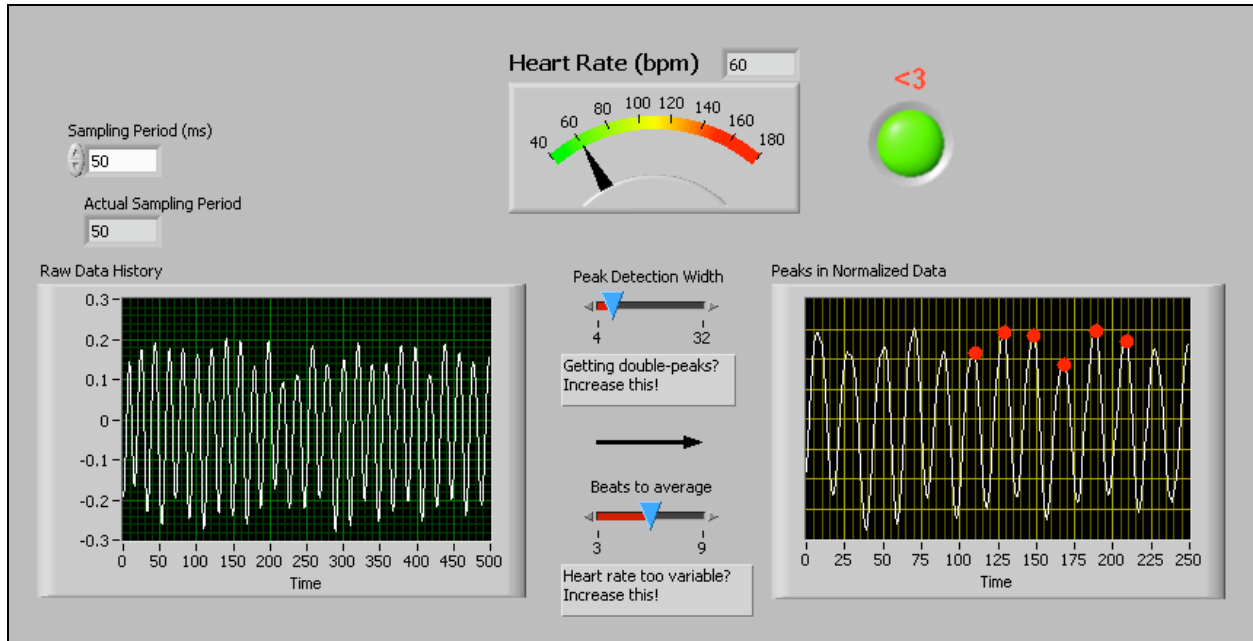


Figure 3.9: Screenshot of completed Labview user interface during device recording.

The user interface (UI) was written as a Labview Virtual Instrument (VI), which provides convenient support for the NI-DAQ being used to record the IR receiver circuit's output signal. The completed user interface, shown in Figure 3.9, displays in real time the PPG signal read by the sensor, a normalized and heartbeat-annotated version of the signal, an indicator of whether a heart-beat is currently in progress, and a meter that displays the subject's current heart rate. The three user-specifiable parameters are:

Sampling interval (default 50 ms): This is the rate at which the VI is requested to acquire voltage measurements from the NI-DAQ. An adjacent label displays the actual interval between measurements taken, which takes into account hardware and software delay. If the measurements are software-timed (as they are in this implementation), the computer cannot acquire samples at an interval shorter than

roughly 30 ms, but 50 ms is more than sufficient for measuring a headphone signal.

Peak detection width (default 8): This is the number of data points over which a peak is identified. More precisely, a quadratic curve fitting is used to fit a parabola over this many data points, and the coefficients from this fitting are used to decide whether or not a peak is present. This is user-specifiable because if the input signal is not clean enough, it is possible for the peak detector to report false peaks in the signal noise. By increasing the peak detection width, it acts as a low-pass filter and can help ensure that only actual peaks are reported.

Number of heartbeats to average (default 6): This is the number of the most recent heartbeats that will be included in the computation of the current heart rate. By increasing this number, the reported heart rate will be an average over a large number of samples and will tend to be more stable and less prone to rapid changes. By decreasing, the reported heart rate will be a better reflection of the current heart rate, but will be more variable and prone to error.

The VI operates by acquiring samples from the NI-DAQ at the specified rate and saving these data points in an array for further processing. The length of this history array is internally set to 250 data points. This should be sufficient for most applications, since it corresponds to over 12 seconds with a 50 ms sampling period, and the maximum number of heartbeats to average is 9. The data in this window is then normalized and then run through the Wavelet Analysis

Multiscale Peak Detection VI provided in the Labview Advanced Signal Processing Toolkit. This VI was chosen for peak detection because of its superior performance over the standard peak detection VI. After all the peaks in the history array are identified, the user-specified number of most recent peaks (heartbeats) is observed to calculate the heart rate. In order to calculate the heart rate, the difference between the locations of sequential pairs of heartbeats is computed and the median period is used to approximate the current heart rate. Mean and weighted averages were also tried, but the median average provided the most robustness and best results. The period is measured as a difference in data point indices, which is then converted into a heart rate by multiplying by the current sample interval and scaling appropriately.

Prior to testing, it was apparent that this UI would not be easily extendable to process two headphone monitor signals, a pulse oximeter signal, and an ECG simultaneously with any adequate sample rate so a simple skeleton VI was constructed to dump the NI-DAQ data directly into log files for later analysis. The default sample period of 50 ms was accidentally left unchanged, which resulted in significant problems because a higher sample rate is required to accurately record an ECG signal. This is further discussed in 3.5.5 Other analyses.

After all testing was performed, the UI was again revisited and adjusted to take advantage of several point-by-point VIs built into Labview to improve the efficiency of the windowing processes. This revised UI is able to achieve sampling rates of around 20 ms when recording from a single headphone monitor.

3.4 EXPERIMENTAL RESULTS

After cleaning the data as described in the Methods section, the data were first pooled by experiment and device. Averaged across experiments, the means of the right ear monitor, left ear monitor, and pulse oximeter were $85.9\% \pm 3.0^*$, $84.2\% \pm 6.4$, and $81.6\% \pm 9.8$ without applying the 50% thresholding described in the Methods (2.2.5 Accuracy analysis) and $86.7\% \pm 2.8$, $86.3\% \pm 4.2$, and $86.3\% \pm 4.7$ with said thresholding.

3.4.1 *Between headphone monitors*

The two ear monitors were then compared. Pooled across all experiments, the accuracy of the right ear monitor ($86.2\% \pm 8.2$) was significantly higher than the left ear monitor ($85.7\% \pm 9.1$) ($p = 0.02$, $N = 690$), with an average paired difference of $0.6\% \pm 7.0$. When separated by experiment, the right ear performed significantly better in experiments 6 ($p = 0.04$, $N = 154$) and 8 ($p < 0.01$, $N = 79$).

3.4.2 *Comparison with pulse oximeter*

The worse (left) ear monitor was then compared to the pulse oximeter, pooling data across subjects. Pooled across the relevant experiments (1, 2, 8), there was no significant ($p = 0.16$, $N = 370$) difference between the accuracy of the two devices. When the experiments were analyzed independently, the pulse oximeter was found to perform significantly better than the left ear monitor for experiment 2 ($p = 0.03$, $N = 203$), averaging an accuracy $1.5\% \pm 1.1$ higher. For experiment 2, there was, however, no significant difference between the right ear monitor and the pulse oximeter ($p = 0.37$).

* This notation was adopted to clarify that both the mean and standard deviation are in units of percent and the standard deviation is not a percent of the mean.

3.4.3 Effect of experiment

The mean accuracy of each device for each experiment is summarized in Figure 3.10.




Figure 3.10: Mean accuracy (+/- one standard deviation) of each device for each experiment, pooled across all subjects. Bar plot generated with “barweb” (reference 43).

The effect of the experiment on device accuracy was analyzed in three ways.

1. The pooled data for each experiment was analyzed with a separate ANOVA for each device, which determined that the accuracy significantly varied across experiments for all three devices ($p \ll 0.01$ for all devices).
2. A two-way ANOVA over experiments and devices confirmed a significant variation between experiments ($p < 0.01$), but revealed that there was no significant difference in the way the devices varied with experiment ($p = 0.91$). In fact, 84% of the total variation in the data was due to the variation between experiments.
3. Two-sample t-tests were used to compare the performance of each device for each experiment with the control for that device. As summarized in Table 3.1, at least one of the

⁴³ Ajiboye, *barweb*.

devices performed significantly worse in experiments 2, 3, 5, and 8. None of the devices ever performed significantly better in an experiment than in the control.

Table 3.1: Average device accuracy for each experiment. Performance that was significantly ($\alpha = 0.05$) worse than the control is shaded. No device performed significantly better in an experiment than in the control.

Experiment	Right ear	Left ear	Oximeter
Control (1)	0.884	0.891	0.896
2	0.851	0.844	0.859
3	0.858	0.863	0.814
4	0.890	0.896	0.909
5	0.873	0.848	0.900
6	0.880	0.874	0.896
7	0.890	0.908	0.857
8	0.807	0.777	0.776

3.4.4 *Effect of blood pressure*

An ANOVA of the data from the tilt table experiments (4, 5, 6) revealed that only the left ear sensor had any significant change in accuracy between the experiments ($p < 0.01$, $N = 243$). The right ear sensor and pulse oximeter remained relatively constant ($p = 0.33$, $N = 239$ and $p = 0.87$, $N = 129$, respectively). Upon further investigation, the accuracy of the left ear sensor was lower during the horizontal state than either the declined or inclined states (85% versus 89% and 87%, respectively).

3.4.5 *Effect of subject*

The variation in device accuracy across subjects is summarized in Figure 3.11. Although not shown, the data has similar form and variation when restricted to just the control experiment. A two-way ANOVA over subjects and devices showed that the accuracy of the devices varied significantly between subjects ($p < 0.01$), but also that this variation was similar between the

devices ($p = 0.67$). The variation between devices accounted for less than 1% of the total data variation, as opposed to 76% explained by the variation between subjects.




Figure 3.11: Mean accuracy (+/- one standard deviation) of each device for each subject, pooled across all experiments. Bar plot generated with “barweb” (reference 44).

3.4.6 Effect of time

The data for experiments 1 and 2 were analyzed to explore the effect of the recording time on device accuracy. A two-way ANOVA across the three devices and the sample time showed a significant effect from each of these two factors ($p = 0.01$, $p < 0.01$, respectively). Linear regression analyses, however, revealed that there was no significant correlation between time and accuracy for any of the devices. This is especially true for the control (experiment 1), but all three devices tended to improve slightly over the course of experiment 2. For each experiment

⁴⁴ Ibid.

and device, the estimated slope and corresponding 95% confidence interval (CI) are shown in Table 3.2.

Table 3.2: Estimated slope and 95% confidence interval for the correlation between time and accuracy for each device.

	<i>Experiment 1</i>			<i>Experiment 2</i>		
	Slope	95% CI		Slope	95% CI	
Right ear	-0.004	-0.017	0.010	0.003	-0.001	0.007
Left ear	0.002	-0.013	0.017	0.004	0.000	0.008
Pulse oximeter	-0.010	-0.034	0.014	0.004	-0.002	0.009

3.5 ANALYSIS AND CONCLUSIONS

3.5.1 Overall device performance

Across all tests, the headphone heart rate monitors developed were quite accurate (typically over 85%) and as accurate or more accurate than the pulse oximeter. The better of the two constructed monitors was significantly better than the pulse oximeter, and the worse performed similarly or slightly worse. This variation between devices was likely due to two reasons. First, there is some variation in the construction of the two devices, most notably the angle and position of the LED on the ear clip. This manifested itself in one of the headphones (the worse one) consistently falling out of the ear of the subject more frequently than the other. This was especially true for the horizontal tilt table and hard breathing experiments where the left ear monitor performed significantly worse than the right due to more frequently being dislodged. Second, there is substantial variation in the actual component values in the two duplicate receiver circuits. The result of this, in combination with the LED variation, was a two-fold difference in the peak-to-peak amplitude of the signals outputted by the two circuits at the same LED current. Because the LED voltages were controlled with the same power supply and could not be individually tuned, one device's IR sensor was always closer to the saturating level than the other. Aside from the

two previously-mentioned experiments in which one headphone fell out more frequently than the other, however, the two headphone monitors seemed to perform very similarly.

3.5.2 Experiment effects

As expected, the devices performed significantly worse during general activity than during the control, but surprisingly, these drops were less than 5%. The ear monitors performed very well under a variety of circumstances although the accuracy was drastically lower just after the subject had exercised. This drop in performance was similarly matched by the pulse oximeter, however, and was likely due to the increase in respiration effects in the PPG signal.

The tilt table experiments (4, 5, and 6) showed no relationship between blood pressure and accuracy.

3.5.3 Subject effects

The performance of the ear monitors was significantly affected by the subject being monitored. One of three potential causes is that differences in ear shape and size could affect the ear monitor accuracy by varying how the headphone stability and the extent to which the IR receiver is exposed. This did not seem to be the main reason for the variation in accuracy, however, because there was not a significant difference between the accuracies of the ear monitors and the pulse oximeter for any subject. Thus, if the ear monitor performed worse on a given subject, it was likely that the pulse oximeter performed worse on that subject as well. Thus, the second potential cause is variation in PPG or ECG signal between subjects. If the ECG control performed poorly, all devices were assessed as having performed poorly, as discussed in 3.6 Suggestions for future work. If there were significant variation in the strength or waveform of the blood flow between

subjects, this would have affected both the ear monitors and the pulse oximeter, since both are PPG-based. A final potential cause is variation in the behavior of the subjects during the experiments. The accuracy of all measurement devices was likely lower for subjects who moved, talked, ate, or drank more than other subjects. This cannot be verified, however, since no record of the subjects' movements exists.

3.5.4 Time effects

The length of time since starting recording did not seem to have a significant impact on the accuracy of any of the devices. During the 10-minute experiment 2, there was a slight increase in accuracy over time, but this was likely due to a decrease in the average activity of the subjects over time rather than any change in the performance of the device.

3.5.5 Other analyses

During data analysis it became apparent that the sampling period of 50 ms was insufficient to consistently capture the R-wave of the ECG signal. Since the R-wave tends to be 30 Hz or less, the sampling period should be 16 ms or less (60 Hz), rather than the 20 Hz used. The result of this was occasionally missing R-waves completely, thus making ECG analysis very difficult. Since many real heartbeats were not identified in the ECG signal, this incorrectly classified many true positives identified by the devices as false positives instead, reducing their apparent accuracy. The real accuracies of the devices are thus likely higher than the reported values.

3.6 SUGGESTIONS FOR FUTURE WORK

Should this project be continued, there are a large number of areas in which improvements could be made. First of all, the receiver circuit could be redesigned to provide better filtering,

automatically adjust the LED power to prevent receiver saturation, and automatically adjust the amplifier gain to normalize the circuit output. Due to a flaw in an operation amplifier (op amp), it was discovered that if the signal were pushed onto the negative power channel in a particular way, the power supply effectively filtered the signal, improving the device performance substantially, including making it more motion resistant. Even though this behavior could not be replicated with a functional op amp, it shows that with superior filters, the device may be able to work well during exercise.

A second area of improvement is in the design of the headphone device itself. The receiver outlet could be drilled at a location that would point more toward the ear, thus helping to reduce the impacts of environmental light and motion. A significant factor in the current location choice was convenience. Another improvement would be standardizing the placement of the LED along the ear clip and embedding it. The placement variation may have contributed to differences in device accuracy, and embedding the LED in the ear clip would allow it to be centered and stabilized for improved comfort and performance. Similarly, using a different headphone scaffold could improve the comfort and stability of the monitor, which would result in improved performance. Newer and more expensive headphones fit tighter into the ear and would help prevent dislocation, a major problem with one of the ear monitors. Another area of improvement in the device is in the wires connecting the IR sensor to the receiver circuit. Currently, four 24-gauge, coated wires connect to each ear's headphone, and the weight and torque can dislodge the devices or cause them to move. By using fewer, thinner, shielded wires or a wireless (e.g. Bluetooth) link, the ear monitors would be less prone to falling out of the ear and would be more comfortable. If the receiver circuit were printed onto a microchip, the entire system would be

less bulky and more comfortable, although then additional power issues would present themselves.

Another area of improvement in this project would be to improve the computer software supporting the system. The digital filter and peak detection could be redesigned to perform better, and the existing code could be optimized to increase the possible sample rate. This is critical in order to ensure that the ECG control signal is properly recorded and the R-wave is not missed because of an insufficient sample rate. Further, the user interface should be updated to work for two headphone monitors at once, as in its current state only one is supported. Another improvement to the user interface would be to allow the user to specify a target heart rate and for the system to play music with a compatible tempo through the headphones. Perhaps more interesting, the system could be improved to automatically select music from a playlist with a tempo which matches the user's current heart rate. Thus, the music selection would automatically be tailored to the user's current physiological condition.

ACKNOWLEDGEMENTS

I would like to thank Mike Regnier and Chris Neils for their understanding, support, and direction through the duration of this project, and the Department of Bioengineering and the University of Washington for providing the equipment and resources that enabled this research to be performed. I am grateful for the technical aid provided by Yung-Chun Chen and Trevor Fowler, and the emotional support by Krista Hiatt and Linda Josephson. This research could not have been completed without the selfless dedication of time provided by the test volunteers.

REFERENCES

- Achten, J., and A.E. Jeukendrup. "Heart Rate Monitoring: Applications and Limitations." *Sports Medicine* 33 (2003): 517-538.
- Ajiboye, Bolu. *barweb (BARgraph With Error Bars)*. MATLAB Central. <http://www.mathworks.com/matlabcentral/fileexchange/10803>.
- Allen, John. "Photoplethysmography and its application in clinical physiological measurement." *Physiological Measurement* 28, no. 3 (2007): R1-R39.
- "Apple's "The beat goes on" special event." *Steve Jobs live*. <http://www.engadget.com/2007/09/05/steve-jobs-live-apples-the-beat-goes-on-special-event/17>.
- Asada, H.H., P. Shaltis, A. Reisner, Sokwoo Rhee, and R.C. Hutchinson. "Mobile monitoring with wearable photoplethysmographic biosensors." *Engineering in Medicine and Biology Magazine, IEEE* 22, no. 3 (2003): 28-40.
- Awad, Aymen, Robert Stout, M. Ghobashy, Hoda Rezkanna, David Silverman, and Kirk Shelley. "Analysis of the Ear Pulse Oximeter Waveform." *Journal of Clinical Monitoring and Computing* 20, no. 3 (June 1, 2006): 175-184. doi:10.1007/s10877-006-9018-z.
- Bisaccia, E, A Lugo, B Johnson, and D Scarborough. "The surgical correction of protuberant ears." *Skin Therapy Letter* 10, no. 2 (2005): 7-9.
- Brodersen, Olaf, D. Römhild, D. Starke, A. Steinke, and S. Vogel. "In-Ear Acquisition of Vital Signs Discloses New Chances for Preventive Continuous Cardiovascular Monitoring." In *4th International Workshop on Wearable and Implantable Body Sensor Networks*, 13:189-194. Aachen University, Germany: Springer Berlin Heidelberg, 2007. http://dx.doi.org/10.1007/978-3-540-70994-7_33.
- Casati, Andrea, Grazia Squicciarini, Marco Baciarello, Marta Putzu, Alessandra Salvadori, and Guido Fanelli. "Forehead Reflectance Oximetry: A Clinical Comparison with Conventional Digit Sensors during Laparotomic and Laparoscopic Abdominal Surgery." *Journal of Clinical Monitoring and Computing* 21, no. 5 (October 1, 2007): 271-276. doi:10.1007/s10877-007-9084-x.
- Celka, P, C Verjus, R Vetter, P Renevey, and V Neuman. "Motion Resistant Earphone Located Infrared based Heart Rate Measurement Device." In *Biomedical Engineering*. Innsbruck, Austria, 2004.
- Challoner, A V J, and C A Ramsay. "A photoelectric plethysmograph for the measurement of cutaneous blood flow." *Physics in Medicine and Biology* 19, no. 3 (1974): 317-328.
- Cheang, Y S, and P R Smith. "An Overview of Non-contact Photoplethysmography." *Electronic systems and control division research* (2003): 57-59.
- Christensen, Kelli, and Robert Kushner. "Commercial Program and Product Review Heart Rate Monitors." *Obesity Management* 2, no. 6 (2006): 242-243.
- Debreczeny, Martin, and Clark Baker. "Pulse oximetry motion artifact rejection using near infrared absorption by water." <http://www.patentstorm.us/patents/7277741.html>.
- Finkelstein, Eric A., Ian C. Fiebelkorn, and Guijing Wang. "National Medical Spending Attributable To Overweight And Obesity: How Much, And Who's Paying?." *Health Aff* (May 14, 2003): hlthaff.w3.219.
- Fletcher, Gerald F, Gary Balady, Victor F Froelicher, L. Howard Hartley, William L Haskell, and Michael L Pollock. "Exercise Standards : A Statement for Healthcare Professionals From the American Heart Association." *Circulation* 91, no. 2 (1995): 580-615.

- Foo, Jong, and Stephen Wilson. "A computational system to optimise noise rejection in photoplethysmography signals during motion or poor perfusion states." *Medical and Biological Engineering and Computing* 44, no. 1 (March 27, 2006): 140-145. doi:10.1007/s11517-005-0008-y.
- Foo, JY A ;Wilson. "Motion Artefact Reduction of the Photoplethysmographic Signal in Pulse Transit Time Measurement." *Australasian Physical & Engineering Sciences in Medicine* 27, no. 4 (December 2004): 9 p.
- Gay, Valérie, and Peter Leijdekkers. "A Health Monitoring System Using Smart Phones and Wearable Sensors." *International Journal of ARM* 8, no. 2 (June 2007): 29-36.
- Grajales, L., and I.V. Nicolaescu . "Wearable multisensor heart rate monitor." In *Wearable and Implantable Body Sensor Networks, 2006. BSN 2006. International Workshop on*, 4 pp.-157, 2006.
- Horowitz, Paul, and Winfield Hill. *Art of Electronics*. Cambridge University Press, 1980.
- Hung, K, Y T Zhang, and B Tai. "Wearable medical devices for tele-home healthcare." *Engineering in Medicine and Biology Society* 2 (2004): 5384-5387.
- Johansson, Anders, Per Åke Öberg, and Gunnar Sedin. "Monitoring of Heart and Respiratory Rates in Newborn Infants Using a New Photoplethysmographic Technique." *Journal of Clinical Monitoring and Computing* 15, no. 7 (December 1, 1999): 461-467. doi:10.1023/A:1009912831366.
- K. Nakajima, T. Tamura, and H. Miike. "Monitoring of heart and respiratory rates by photoplethysmography using a digital filtering technique." *Medical engineering & physics* 18, no. 5 (July 1, 1996): 365-372.
- Kung, HC, DL Hoyert, J Xu, and SL Murphy. "Deaths: final data for 2005." *Natl Vital Stat Rep* 54, no. 10 (April 24, 2008): 1-120.
- Lei Wang, B.P.L. Lo, and Guang-Zhong Yang. "Multichannel Reflective PPG Earpiece Sensor With Passive Motion Cancellation." *Biomedical Circuits and Systems, IEEE Transactions on* 1, no. 4 (2007): 235-241.
- Lindberg, L., H. Ugnell, and P. Öberg. "Monitoring of respiratory and heart rates using a fibre-optic sensor." *Medical and Biological Engineering and Computing* 30, no. 5 (1992): 533-537. doi:10.1007/BF02457833.
- Mancini, Ron. *Op Amps for Everyone: Design Reference (SLOD006B)*. 2nd ed. Newnes, 2003. <http://www.ti.com/litv/pdf/slod006b>.
- "Nike + iPod." Promotional site. *Apple*. <http://www.apple.com/ipod/nike/run.html>.
- "Principles of Pulse Oximetry Technology." *Pulse Oximetry*, September 10, 2002. <http://www.oximeter.org/pulseox/principles.htm>.
- Relente, A.R., and L.G. Sison. "Characterization and adaptive filtering of motion artifacts in pulse oximetry using accelerometers." In *[Engineering in Medicine and Biology, 2002. 24th Annual Conference and the Annual Fall Meeting of the Biomedical Engineering Society] EMBS/BMES Conference, 2002. Proceedings of the Second Joint*, 2:1769-1770 vol.2, 2002.
- Renevey, P., R. Vetter, J. Krauss, P. Celka, and Y. Depeursinge. "Wrist-located pulse detection using IR signals, activity and nonlinear artifact cancellation." In *Engineering in Medicine and Biology Society, 2001. Proceedings of the 23rd Annual International Conference of the IEEE*, 3:3030-3033 vol.3, 2001.

- Seaward, B., R. Sleamaker, T. McAuliffe, and J. Clapp. "The Precision and Accuracy of a portable heart rate monitor." *Biomedical Instrumentation & Technology* 24, no. 1 (1990): 37-41.
- Shelley, Kirk H. "Photoplethysmography: Beyond the Calculation of Arterial Oxygen Saturation and Heart Rate." *Anesth Analg* 105, no. 6S_Suppl (2007): S31-36.
doi:10.1213/01.ane.0000269512.82836.c9.
- Sherebrin, M.H., and R.Z. Sherebrin. "Frequency analysis of the peripheral pulse wave detected in the finger with a photoplethysmograph." *Biomedical Engineering, IEEE Transactions on* 37, no. 3 (1990): 313-317.
- Shim, Bong Sup, Wei Chen, Chris Doty, Chuanlai Xu, and Nicholas A Kotov. "Smart Electronic Yarns and Wearable Fabrics for Human Biomonitoring made by Carbon Nanotube Coating with Polyelectrolytes." *Nano Letters* 8, no. 12 (2008): 4151-4157.
- Stern, Robert. "Ear Lobe Photoplethysmography." *Psychophysiology* 11, no. 1 (1974): 73-75.
- "Strapless Heart Rate Monitor Clothes." Promotional site. *NuMetrex*.
<http://www.numetrex.com/>.
- Teng, XF, and YT Zhang. "M-health: trends in wearable medical devices." *Zhongguo Yi Liao Qi Xie Za Zhi* 30, no. 5 (2006): 330-340.
- Verkruysse, Wim, Lars O. Svaasand, and J. S. Nelson. "Remote plethysmographic imaging using ambient light." *Opt. Express* 16, no. 26 (2008): 21434-21445.
- Yong-Sheng Yan, and Yuan-Ting Zhang. "An Efficient Motion-Resistant Method for Wearable Pulse Oximeter." *Information Technology in Biomedicine, IEEE Transactions on* 12, no. 3 (2008): 399-405.

APPENDIX

A1. DATA ANALYSIS SCRIPTS

analyzedata.m: This function accepts a *data file and the length of each sample (in seconds) analyzed from the data. The input data is an NxM array with a row for each time point. Each row must be of the form: [<index>, <right ear>, <left ear>, <pulse oximeter>, <ecg>]. The input data is then divided into chunks of length specified by the chunk length parameter. For each chunk, all heartbeats are found in each of the four signals, and each heartbeat found in the ear and pulse oximeter signal is classified according to the identified ECG heartbeats. The algorithm run to classify these heartbeats can be summarized in pseudocode as follows:*

```

Keep track of the current heartbeat for each device
For each ECG heartbeat identified
    Let the index of the current and next two ECG heartbeats be  $t_i$ ,  $t_{i+1}$ , and  $t_{i+2}$ , respectively
    For each device
        For each heartbeat before  $t_i$ 
            Add one to the count of false positives
            Advance to the next device heartbeat
        End
        If there are any heartbeats identified between  $t_i$  (inclusive) and  $t_{i+2}$  (exclusive)
            Add one to the count of true positives
            Advance to the next device heartbeat
        Else
            Add one to the count of false negatives
        End
    End
End

```

This greedy algorithm allows a reasonable amount of delay (two measured heartbeats) between when the ECG detects a heartbeat and any other device, and correctly classifies each device heartbeat as either a true positive (device heartbeat corresponds to ECG heartbeat), false positive (device heartbeat with no corresponding ECG heartbeat), or false negative (no device heartbeat within window to correspond to ECG heartbeat). The concept of a true negative is not clearly defined for this application because it is unclear how the absence of a heartbeat could be correctly identified.

findheartbeats.m: This function accepts a string which specifies the type of data ('ecg', 'pulseox', or 'ear') and a vector of data points, and returns three vectors, heartbeats, filtereddata, and sigerror.

heartbeats: The heartbeats vector contains a list of indices which correspond to the center of the heart beat signal for the specified data type. For ECG data, the reported indices correspond to the centers of the QRS complexes. For pulse oximetry data, the reported indices correspond to the indices at which the pulse

oximeter emitted a beep, signaling a recorded heartbeat. For the ear sensors, the reported indices correspond to the peaks in the sinusoidal PPG signal.

filtereddata: The *filterreddata* output is a vector of the same number of data points as the input data vector, and contains the filtered version of that vector used to identify the heartbeats in the signal.

sigerror: The *sigerror* vector is used by the ECG data type and contains a Boolean value for each element in the input data vector. If the value is 0, *then the corresponding input data point was determined to be part of a valid ECG signal. If the value is 1, then the ECG signal was not valid at this input and if there was a peak at this data point, it was not reported. This is useful if the ECG signal is used as a control, as it allows the receiving code to ignore all signals for the set of data points when the control signal was not reasonable. This essentially provides a report of any events which were detected in the ECG signal. Such events were identified when the running standard deviation of the ECG waveform exceeded a reasonable threshold.*

findpeaks.m: This function accepts a vector of binary data and returns an array of *indices which correspond to the center of above-zero regions of the data.* If a segment of ones is an even number of data points in length, it returns the lower index.