Sociometric Badges: Wearable Technology for Measuring Human Behavior

by

Daniel Olguín Olguín

M.S. in Electronic Systems Engineering,
Tecnológico de Monterrey Campus Monterrey (Mexico, 2005)
B.S. in Electronics and Communications Engineering,
Tecnológico de Monterrey Campus Monterrey (Mexico, 2003)

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

Master in Media Arts and Sciences

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2007

Author_____

Program in Media Arts and Sciences
May 11, 2007

Certified by_____

Alex P. Pentland
Toshiba Professor of Media Arts and Sciences
Program in Media Arts and Sciences
Thesis Supervisor

Accepted by_____

Andrew B. Lippman
Chair, Departmental Committee on Graduate Students
Program in Media Arts and Sciences

Sociometric Badges: Wearable Technology for Measuring Human Behavior

by

Daniel Olguín Olguín

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
on May 11, 2007, in partial fulfillment of the
requirements for the degree of
Master in Media Arts and Sciences

Abstract

We present the design, implementation and deployment of a wearable computing research platform for measuring and analyzing human behavior in a variety of settings and applications. We propose the use of wearable sociometric badges capable of automatically measuring the amount of face-to-face interaction, conversational time, physical proximity to other people, and physical activity levels using social signals derived from vocal features, body motion, and relative location to capture individual and collective patterns of behavior. Our goal is to be able to understand how patterns of behavior shape individuals and organizations. We attempt to use on-body sensors in large groups of people for extended periods of time in naturalistic settings for the purpose of identifying, measuring, and quantifying social interactions, information flow, and organizational dynamics. We deployed this research platform in a group of 22 employees working in a real organization over a period of one month. Using these automatic measurements we were able to predict employees' self-assessment of productivity, job satisfaction, and their own perception of group interaction quality. An initial exploratory data analysis indicates that it is possible to automatically capture patterns of behavior using this wearable platform.

Thesis Supervisor: Alex P. Pentland
Title: Toshiba Professor of Media Arts and Sciences, Program in Media Arts and Sciences

Sociometric Badges: Wearable Technology for Measuring Human Behavior

by

Daniel Olguín Olguín

The following people served as readers for this thesis:

Thesis Reader_____

Joseph A. Paradiso
Associate Professor of Media Arts and Sciences
Program in Media Arts and Sciences

Thesis Reader_____

David P. Reed
Adjunct Professor of Electrical Engineering
Program in Media Arts and Sciences

Acknowledgments

I would first like to thank my advisor Alex (Sandy) Pentland for encouraging me to pursue new directions and providing me with his guidance and support over the last two years. I look forward to working with him on many new exciting projects. I would next like to thank my thesis readers: Joseph A. Paradiso and David P. Reed for their invaluable feedback and comments that made of this thesis a better work.

Next in the list are the members of the *Responsive Environments* group who were there to answer many of my questions during the design process of the sociometric badges: Mark Feldmeier for his invaluable help designing and testing the electronic circuitry, Joshua Lifton for sharing his work and knowledge with me, and Mat Laibowitz for his thoughtful recommendations.

Thanks also to the *Human Dynamics* group members with whom I have worked over the last two years: Jonathan Gips, Wen Dong, Juan Carlos Barahona, and the "Badge Team": Benjamin Waber, Taemie Kim, Akshay Mohan, and Koji Ara for the hard work that made this project possible. Thanks to Mary Heckbert who played an important role making sure that I had acces to all the resources that I needed.

Special thanks the MIT Media Lab sponsors: Hitachi, Cisco, Motorola and other members of the TTT consortium that supported this project. Thanks to Analog Devices for their generous donation. Thanks to Daniel Oster for making the first experiment possible.

I would also like to express my gratitute to all the UROPS who participated in this project and who made it a really fun experience: Nathan Davis, Agustya Mehta, Alice Yeh, and Christopher Moh.

Finally, I would like to thank my parents and sister: Daniel Olguin Salinas, Amparo Olguin Olmos, and Marisol Olguin Olguin for their endless support and encouragement to pursue my dreams.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The study of human behavior has always intrigued social scientists interested in enhancing organizational effectiveness and individual well-being in the workplace. Organizational behavior is a multidisciplinary field that seeks knowledge of behavior in organizational settings by systematically studying individual, group, and organizational processes. Some of the questions it tries to answer are: How can goals be set to enhance people's job performance? How may jobs be designed so as to enhance employees' feelings of satisfaction? Under what conditions do individuals make better decisions than groups? What can be done to improve the quality of organizational communication? How can leaders enhance the effectiveness of their teams? [35]

According to Bowditch [23], some of the behavioral skills necessary for effective managerial performance are: communicating with peers, subordinates, and bosses; obtaining and sharing information; running meetings; allocating resources to different groups; and handling conflict within or between teams. Understanding of management processes, often referred to as organizational behavior, requires the study of people, groups, and their interactions in organizations. Advanced information and manufacturing technologies such as computer networks and mobile communication devices are revolutionizing the ways in which organizations operate and people perform their jobs.

In this thesis we propose the use of wearable electronic badges to automatically collect

and analyze unconscious social signals and context information, with the purpose of studying individual and group behavior in organizational settings. We set the foundations for developing the technology that will enable social scientists to automatically measure individual and collective patterns of behavior, predict human behavior from unconscious social signals, identify social affinity among individuals working in the same team, and enhance social interactions by providing feedback to the users of this platform.

1.1 Synopsis

This thesis provides a detailed description of the design process and implementation of a wearable computing research platform for measuring and analyzing human behavior in a variety of settings and applications. We present initial exploratory results that show that it is possible to measure several aspects of human behavior using this platform. The thesis is organized as follows: Chapter 1 describes the research problem, chapter 2 gives an overview of the background and previous work related to the research topic, chapter 3 presents a study on activity recognition accuracy across common locations for wearable sensors, chapter 4 describes the proposed hardware platform, chapter 5 presents experimental results and exploratory data analysis, and finally chapter 6 presents the conclusions and future work.

1.2 Problem Statement

Social interaction can take on several forms, face-to-face interaction being one of the most frequent and important of them [41]. As a result of the digital revolution, people can now communicate with others over long distances using mobile telephones, e-mail, instant messaging clients, video conferencing, and other forms of digital media. However, little has been done to exploit the fact that people already carry with them wearable electronic devices such as mobile phones, PDAs, MP3 players, digital watches, radio frequency ID (RFID) badges, and the like. These devices have the necessary computational power to process and

analyze information about their users' behavior. These body-worn sensor networks mean that we can potentially know who talks to whom, and even how they talk to each other.

Even though some devices already incorporate sensors capable of capturing context information, we believe that there is no single platform capable of measuring a wide range of variables such as the amount of face-to-face interaction, non-linguistic social signals, location, physical proximity to other people, and context information to facilitate the study of human behavior. We propose a new research platform composed of a wearable electronic badge and a mobile phone, capable of measuring human behavior by capturing and analyzing the aforementioned social signals for a variety of applications.

1.2.1 Research Problem

How can we automatically collect and analyze social signals that are characteristic of human behavior in large groups of people over extended periods of time?

Standard methods to measure and evaluate human behavior, such as surveys, often suffer from subjectivity and memory effects. In [65] Pentland envisioned a device that could accurately and continuously track the behavior of hundreds of humans at the same time, recording even the finest scale behaviors with great accuracy. Such a device would replace expensive and unreliable human observations with automated, computer-mediated ones.

In this thesis we attempt to create such a device. We believe that the automatic discovery and characterization of face-to-face communication and social interaction would allow us to gather interaction data from large groups of people. This could potentially remove two of the current limitations in the individual analysis of human behavior: the number of people that can be surveyed and the frequency with which they can be surveyed.

It has been demonstrated that the effectiveness of similar automatic methods of measuring and predicting human behavior can be evaluated following a social psychologist perspective: the test for the ability to automatically measure social signals would be the ability to

predict outcomes from *thin slice* observations of human interactions [28]. To demonstrate the effectiveness of our tool, a series of experiments in collaboration with other researchers has been planned and one has been carried out. For details on these experiments refer to chapter 5.

1.2.2 Research Hypothesis

Some characteristics of human behavior can be automatically measured and analyzed by using a wearable electronic badge capable of capturing social signals derived from speech, body motion, proximity to other people, relative location, and face-to-face interaction.

1.2.3 Research Assumptions

Every person participating in our experiments will wear a badge around their neck every day while they are at work. The badge will record different speech features, without recording the content of conversations in order to maintain privacy, measure the amount of body motion, record relative location using fixed base stations as reference points, capture face-to-face interactions with other participants wearing the badge, and estimate proximity to other participants via periodic Bluetooth scans.

1.2.4 Research Methods

- While various hardware platforms have been developed in the past, we are unaware of any commercial device capable of measuring the social signals of interest and making them available for research purposes. Therfore, a new small form factor wearable electronic badge will be designed and manufactured.
- Several tests will be performed to assess the functionality of the hardware.
- Several user studies will be carried out in real organizations to gather data and analyze individual and group behavior.

- Participants will be asked to answer a short survey at the end of the day that will allow us to correlate our automatic measurements with their own perceptions.

1.3 Motivation

We believe there is a need for automatic tools to measure individual and group behavior in the social sciences. Researchers often rely on surveys and human observers to study human behavior. This is often expensive and time consuming. Therefore, we have decided to build an experimental research platform that could prove useful and effective in the assessment of human behavior. We predict there will be a large demand for such tools in different research communities and later on in the consumer market for self-monitoring devices.

In [61] we presented the design of the first prototype of a *"Star Trek"*-like wearable communicator badge as part of the research platform proposed in this thesis. Three hundred of these badges have been manufactured to date. These badges will enable people working in large organizations to communicate, find information, and interact in more efficient and intelligent ways. The badge should be able to perform speech analysis using a microphone and state-of-the-art micro-power electronics. It will be used as a push-to-talk system and will be capable of playing audio messages and reminders through a speaker. An accelerometer will allow us to study how people move and behave throughout the day: Are they walking to a meeting? Are they talking to someone? Are they sitting in front of their computers? An infrared (IR) sensor will be used to capture face-to-face interactions and study social networks. A 2.4 GHz radio transceiver will send and receive information from base stations distributed along a specific area, as well as determine proximity to other badges. A Bluetooth module will enable it to interface with cell phones, PDAs, portable computers, and other Bluetooth-enabled sensors and devices.

1.4 Scope of Research

We are developing a technology to study human behavior and enhance communication between humans by using physical information from speech, body motion, proximity sensing, and location, among others. We would like to bring this technology to real organizations with the long-term goal of redefining managerial practices in order to increase productivity, efficiency, and facilitate the flow of information. During the time frame of this thesis work we performed several feasibility studies in real organizations to recognize potential improvements, enhancements, and problems that the system might present in terms of functionality, performance, user acceptance, and privacy concerns.

We designed, implemented, and deployed the necessary hardware and software infrastructure for collecting and analyzing data from several individuals over extended periods of time (at least one month) in different settings. The purpose of these studies was to validate the functionality and performance of the system.

The main contributions of this thesis work are: 1) The design and implementation of a wearable sociometric badge. 2) The search for an optimal configuration of sensors' placement on the human body to perform daily activity recognition and capture social signals. 3) The experimental deployment of this research platform. 4) An initial exploratory data analysis to assess the functionality of this platform.

1.5 Proposed Approach

As we begin to deploy hundreds and even thousands of wearable sensors on regular workers, hospital patients, and the general population, the question shifts more toward a balance between what information can be gained and their broad, immediate user acceptance. In [62] we studied human activity recognition using wearable sensors in different parts of the body. We found that the best global activity classification performance was achieved when using three accelerometers placed on the chest, hip and wrist. However, we also found that it is possible to obtain similar results using only two accelerometers placed on the chest and hip.

We are convinced that a minimal system formed by the combination of an electronic badge worn around the neck and a mobile phone carried by the user is the best type of platform that would allow us to study human behavior in organizational settings over extended periods of time.

We propose the use of a new badge that has a small form factor, is comfortable to wear over long periods of time (at least eight hours a day), and has a long battery life. To achieve this, the badge was designed for very low power wake-up directly from sensor stimuli. In addition to some of the main features offered by previous badge platforms, the sociometric badge is also capable of:

- Measuring physical activity levels and recognizing common daily human activities (sitting, standing, walking, and running) in real time using a single 3-axis accelerometer and unsupervised learning methods similar to the one presented in [38]. This range of activities could be expanded by adding a mobile phone containing an accelerometer [6].
- Detecting when the user is participating in a conversation and extracting speech features in real time to measure non-linguistic social signals and identify the social context while ignoring the words themselves in order to assuage privacy concerns [67].
- Interacting with the user through voice commands, a push-button, and a tri-directional switch to find resources and information in a timely fashion.
- Communicating with radio base stations in the 2.4 GHz radio band for sending and receiving information to and from different badges. The base stations could either be other badges placed at fixed locations or compatible radio base stations, such as the *Plug* sensor network developed in the Responsive Environments group at the MIT Media Laboratory [46].
- Measuring the received signal strength (RSSI) from other radio transceivers to perform indoor localization using triangulation algorithms.
- Communicating with Bluetooth enabled cell phones, PDAs, and other devices to study the user's behavior and detect people in close proximity [40].
- Capturing face-to-face interactions using an IR sensor [26].

We believe our proposed approach to capture social signals and measure human behavior

has several advantages over existing methods such as direct observation by humans, the use of pervasive cameras to videotape social interactions, or the use of surveys. Direct observation of humans by humans is expensive and limited to a few people per observer, and observers do not always agree. Deploying pervasive cameras is extremely expensive and their range of measurement is constrained to a particular place. The use of surveys is subjective, inaccurate and time consuming. In contrast, the ability to automatically capture not only visible characteristics of human behavior such as face-to-face interactions, relative location, and motion but also the underlying psychological processes that occur during social interactions in hundreds of people, at the same time and with a single unobtrusive tool, represents a great advantage.

1.6 Sample Application Scenarios

We have envisioned several applications scenarios for the sociometric badges:

Sensible Organizations. Organizations will become truly *sensible* when they start deploying hundreds or thousands of wireless environmental and wearable sensors capable of monitoring human behavior, extracting meaningful information, and providing managers with group performance metrics and employees with self-performance evaluations and recommendations [15]. *Sensible Organizations* is a new concept of social sensor network technologies that will help improve organizational practices. A sociometric badge could potentially measure, analyze, and reveal organizational dynamics by closely looking at interactions and social behavior among employees of an organization [77]. Companies using this research platform could have a better understanding of how they work and how they can improve their daily routines in order to increase productivity, innovation, and job satisfaction.

Personal sales coach. Recent experiments in the *Human Dynamics* group at the MIT Media Laboratory have shown that it is possible to measure how persuasive a person is being when talking to others, how interested a person is in a conversation, how much

attention a person is paying to someone, and how effective someone is at negotiating by measuring different voice features and body motion [73]. The sociometric badges could be used to track individual and global sales performance in retail stores and give advice on how to interact with clients more effectively. Sales representatives today lack sufficient feedback on their selling capability. Top performing sales representatives usually have good communication skills, exhibit a lot of enthusiasm and energy, and form a personal bond with their clients. The sociometric badge could be used as a personal sales coach device allowing sales representatives and managers to reflect on their own performance and improve their sales skills.

Health care. The sociometric badge could also be used in a variety of health care monitoring applications, such as depression state tracking, elder care, triage in the emergency room, weight loss monitoring, and others. Sung et al. [74] showed that non-invasive behavioral measures such as voice features and body motion are correlated with depression state and can be used to classify emotional state and track the effects of treatment over time. The sociometric badge could be used as a self-monitoring device to extract speech features, body motion, and social interaction information and alert the user or their family members of early symptoms of depression. A device with a similar form factor and capabilities as the sociometric badge could be used to monitor daily activities of the elderly. Capable of detecting falls and behavioral changes, the badge could automatically alert a family member or a doctor of a potentially dangerous situation. One last example in the health care domain is in the area of weight loss monitoring. Using the badge and a mobile phone could be useful for keeping track of energy expenditure, food intake, and daily routines.

Hospitals. The badge could also be used as a push-to-talk system in hospitals. Doctors and nurses need to be in continuous communication and perform tasks that require the use of both hands. If they wore a badge that knew the location and context of each person inside the building, doctors and nurses could be easily accessible at all times. They could potentially obtain information in the right place and at the right moment. One example of a commercially available push-to-talk system is the

Vocera communications system, already widely accepted and used in hospitals [72]. The technology used in the sociometric badge could potentially augment the current functionalities of such systems.

Children and Learning. Classroom behavior and social interactions of children could also be studied using this wearable platform. How much attention are students paying to the teacher? How much does each student participate in a class? How do children socialize? Using an automatic tool could help behavioral scientists answer questions like these.

Retail stores. Often times customers need to find someone capable of answering their questions or helping them find a product in retail stores. One could potentially know when employees are available to talk to customers or when they are busy already talking to someone if they wore a sociometric badge [40].

In this thesis we will focus on collecting data from individuals participating in different studies with the purpose of giving them feedback on their overall behavioral patterns, such as the amount of time they spend in conversations, the number of people with which they interact face-to-face everyday, and how much they tend to move relative to their peers. We compare individual patterns of behavior with the group average in order to maintain privacy.

Chapter 2

Background and Previous Work

2.1 Electronic Badges

Wearable badges are common devices that employees wear in large organizations to identify themselves to others or to gain access to certain locations or information. The *Active Badge* developed at Xerox PARC in 1992 was one of the first attempts to augment inanimate name tags with electronics. Containing only a small microprocessor and an IR transmitter, this badge could broadcast the identity of its wearer and trigger automatic doors, automatic telephone forwarding, and computer displays [78].

More complex badge platforms have been developed after the Active Badge, several of them at the MIT Media Laboratory and some others in different research labs. In 1996, the *Thinking Tags* [22], developed at the MIT Media Laboratory, were the first computationally augmented name tags that were capable of displaying how much two people at a conference or meeting had in common simply by lighting LEDs of different colors depending on how many questions the participants had previously answered the same. Two years later they evolved into the *Meme Tags* [21], allowing conference participants to electronically share brief ideas or opinions through a large LCD screen. This system further evolved into the *nTAG System*, a commercial solution to improve, measure, and automate meetings and events [7].

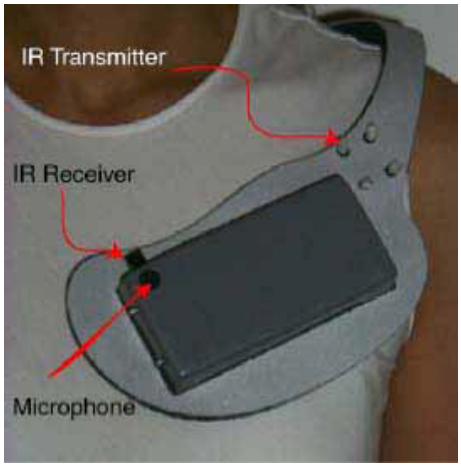

Figure 2-1: The *SocioMeter*.

The *wearable sensor badge* developed at Philips Research Labs in 1999 [31] was capable of detecting simple pre-ambulatory activities using an accelerometer. The *iBadge* [64] was designed to be worn by children to capture interactions with teachers and common classroom objects. It was run by an Atmel ATMEGA microcontroller, a Digital Signal Processor (DSP), and contained a 2-axis accelerometer, a light sensor, a pressure sensor, a temperature sensor, a microphone, and a Bluetooth module. It weighed 65 grams and the battery lifetime was on average 4 hours and 30 minutes.

A more recent example is the *SocioMeter*, a wearable sensor package designed at the MIT Media Laboratory to measure face-to-face interactions between people with an IR transceiver, a microphone, and two accelerometers [27]. This badge was used to learn social interactions from sensory data and model the structure and dynamics of social networks. However, due to its size and weight, users reported feeling somewhat uncomfortable while wearing it. Figure 2-1 shows a picture of the *SocioMeter*. A more detailed description is given in section 2.3.

The *UbER Badge* [43] is the most recent research platform developed at the MIT Media Laboratory for facilitating interaction in large groups of people. This badge has both radio frequency (RF) and IR communication, a 5 x 9 LED display capable of presenting graphics and scrolling text that users in the vicinity can read, an on board microphone for 12-bit audio sampling, a 12-bit audio output available at a headphone jack, a pager motor

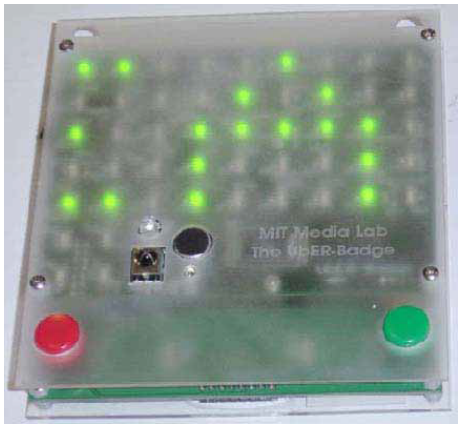

Figure 2-2: The *UbER* badge.

for vibratory feedback, 3 inboard processors, capacity for up to 256 MB of flash memory, provisions for connecting LCD displays, and connectors that allow a variety of different sensors to be integrated. It measures 11 x 12 cm and weights 170 grams. Its average current consumption is 100 mA and batteries last for about 15 hours of continuous usage. Figure 2-2 shows a picture of the *UbER* badge.

The best known commercially available badge system is the 802.11-based *Vocera Communications System* [10]. With it, users interact through wearable badges that can be clipped to coat pockets, worn as pendants, or carried in holsters. The system architecture centers on a server that maintains voice dialing phrases, badge session identifiers, e-mail addresses, telephone numbers, and names. The Vocera badge provides a voice-controlled user interface and enables instant, hands-free conversations among people throughout the workplace. It weighs 53.9 grams and measures 10.6 x 3.5 x 1.5 cm. The standard battery (Li-ion 600 mAh) lasts for 2 hours of talk time and 44 hours of standby time, and it is rechargeable through a single bay or 8-bay charger. This platform is already used by thousands of people in hospitals, retail stores, and service organizations [72].

2.2 Context Awareness

Context-aware computing is an enabling technology for adapting different functions in computer-aided devices to new situations. The development of context awareness for mobile devices requires recognition and extraction of implicit context information from the usage situations and environment of a device. Context information is provided for applications and services, which adapt their appearance and functions accordingly. The main challenge in sensor-based context-aware computing for mobile devices is how to define and carry out context recognition from sensor signals to facilitate the use of context information in mobile applications [49].

According to Lehtonen [44], *context* can be defined as an abstract description of a real-world situation and *context-aware application* as an application that can change its behavior according to the situation. Different variables can be used to discriminate context and they can be anything from physical to social, for example temperature, light intensity, or the number of people in the vicinity. After the variables have been identified it is necessary to find the right sensors that can measure the variables with sufficient accuracy and minimal cost. The next task is to find algorithms that recognize the context with maximum certainty.

Different approaches to make environments and devices more intelligent have been proposed. These include proactive computing, pervasive computing, tangible user interfaces, affective computing [68], and wearable computing. In all these approaches the system should be able to use information about the user and the situation in order to make the interaction between the user and the system more fluent.

Cakmakci *et al.* [24] propose a statistical approach, which could scale down to microcontrollers with scarce resources, to model simple contexts based on raw sensor data. They characterize each context with samples of sensor data taken from each context to estimate probability densities corresponding to contexts. Once they have an estimate for the density characterizing a context, they proceed to compute the probability of being in a certain context given sensor data.

Mayrhofer *et al.* [56] have proposed an architecture to recognize and predict user context

by utilizing multiple heterogeneous sensors. They stated that context recognition and prediction should be embedded in mobile devices with limited resources. In addition, they believe that learning and adaptation should happen on-line without explicit training phases, and that user intervention should be kept to a minimum with non-obtrusive interaction paradigms.

The area of context-aware computing is relevant to the wearable computing platform presented in this thesis since the variables captured by the sociometric badges (amount of speech, proximity, body movement, and face-to-face interaction) are directly related to the user's context. Knowing when a person is speaking tells us something about their availability and interruptability. Knowing when a person is sitting, standing or walking is also useful to determine their daily activity patterns. The same can be said about proximity to other people and face-to-face interaction time. Combining these variables and studying the correlations with other objective measurements could prove useful in predicting future behavior and performance metrics in the case of organizational applications.

2.3 Socially-Aware Systems

Pentland has proposed that *social signals* are particularly powerful for analyzing and predicting human behavior [66, 67]. He found that after a few minutes of analyzing social signals within conversations it was possible to predict who would exchange business cards at a meeting; which couples would exchange phone numbers at a bar; who would come ahead in a negotiation; who was a connector within a workgroup; and a range of subjective judgments.

In [66] Pentland developed texture-like measures for four types of non-linguistic social signals: *activity level* (how much time a person takes in a conversation), *engagement* (how involved a person is in a conversation), variation in *prosodic emphasis*, and *mirroring* (occurring when one person unconsciously copies another participant's prosody and gesture).

The *Human Dynamics* research group at the MIT Media Laboratory has developed several

socially aware platforms to measure different aspects of social context. We describe some of these platforms next:

The SocioMeter

Choudhury [26] performed the first successful experiments in learning communication patterns by equipping a group of people with *SocioMeters*. She explored methods for computationally modeling the underlying social network structure and integrated speech processing methods to demonstrate that it was possible to extract information about people's patterns of communication. Our sociometric badges were inspired by the *SocioMeter* concept and we try to take further steps in the directions proposed by Choudhury: comparing face-to-face interaction with other communication channels, adding value for the users by providing feedback, and improving the form factor of the device.

VibeFones

The VibeFone application is mobile social software that uses location, proximity and tone of voice to gain a sophisticated understanding of people's social lives by mining their face-to-face and phone interactions. It was used in several applications such as automatic characterization of social and workplace interactions, a courtesy reminder for phone conversations, and a personal trainer for dating encounters. This application augments traditional means of gathering social interaction data (surveys or ethnographic studies) and speech data. The mobile phone platform is highly conducive to collecting long-term continuous data and sampling the user for training labels [48].

Social Motion

In this application two types of sensors were used: proximity sensors and motion sensors. By proximity sensors, Gips [34] refers to RF scanners capable of detecting other devices within a fixed proximity and IR transceivers capable of detecting when two similar devices

are facing each other. Gips showed that these sensing modalities can be incorporated into current mobile devices without disrupting the usage patterns and form factors to which people have grown accustomed. He also demonstrated that these sensor signals contained the information necessary to infer the underlying social structure of groups of people about which no information is known *a priori* by identifying groups of friends taking part at a career fair, team membership of students participating in a treasure hunt game, and company affiliation of visitors attending a conference [34].

2.4 Human Activity Recognition

Previous work in human activity recognition using accelerometers has shown that it is possible to classify in real time several postures and activities. In [29], the authors developed a two-layer model that combined a Gaussian mixture model with first-order Markov models to classify a range of activities including sitting, walking, biking, and riding the subway. A single three-axis accelerometer placed on the torso was used. Similarly, Van Laerhoven augmented a pair of pants with an accelerometer to register what the wearer was doing and anticipate their behavior [75].

In [47] the authors combined data from accelerometers and microphones placed in different body locations to classify activities performed in a wood shop with 84.4% accuracy. They modeled most of the activities using single Gaussian hidden Markov models (HMMs). The number of hidden states to model each activity was selected through visual inspection.

Kern and Schiele investigated the acquisition of context information from audio and acceleration sensors [39]. They developed a model that uses personal and social interruptability of the user to decide both whether or not to notify the user and to decide which notification modality to use. Motion classification was achieved using two-state ergodic HMMs to describe each class (sitting, standing, walking, walking upstairs, walking downstairs, and running), achieving 91.9% accuracy.

Several algorithms to classify twenty different physical activities from data acquired using

five biaxial accelerometers were evaluated in [17]. Activity recognition was performed using decision tables, instance-based learning, C4.5 decision tree, and naïve Bayes classifiers with an overall recognition rate of 84% (decision trees having the best performance). They showed that acceleration of the dominant wrist is more useful in discriminating their set of activities than that of the non-dominant wrist and that acceleration of the hip is the second best location for activity recognition.

The authors of [51] used decision trees to classify six different activities with a single accelerometer, placed on six different body positions commonly used for wearing electronic devices. They achieved accuracies ranging from 16.7% to 92.8% depending on the position of the accelerometer and the features used.

In chapter 3 we present a study on activity recognition accuracy across common locations for wearable sensors. We find the classification accuracy of a single accelerometer placed in three different parts of the body and evaluate whether there would be a significant improvement in recognition accuracy if multiple accelerometers were added. We wish to determine if a single accelerometer placed on the chest could yield good results when classifying common daily activities, as well as if adding a second accelerometer on the hip (which could be incorporated into a mobile phone) would improve the results.

Chapter 3

Human Activity Recognition Using Accelerometers

The development of miniature sensors that can be unobtrusively attached to the body or can be part of clothing items, such as sensing elements embedded in the fabric of garments, have opened countless possibilities for monitoring individuals over extended periods of time [20]. Being able to automatically recognize human motion patterns using unobtrusive wearable sensors can be useful in monitoring the elderly in their homes and keeping track of their daily activities and behavioral changes. This could lead to a better understanding of numerous medical conditions and treatments. Other applications of human activity recognition range from context aware computing to physical training, physical rehabilitation, and military applications such as intelligent outfit design for soldiers.

In this chapter we study different configurations of accelerometer placement to classify human activities that are frequent in at least one of these application areas. Modeling human activities using hidden Markov models (HMMs) trained with different features derived from accelerometer signals has proven effective [17,29,45,47,51,58]. HMMs are capable of modeling important properties of gestures such as time variance and repetition. They also handle noise due to sensors and imperfect training data by providing a probabilistic framework [47].

We chose to model eight different activities using HMMs and continuous Gaussian observation vectors based on the previous work on activity recognition described in section 2.4. We compare the activity classification accuracy of four different configurations of accelerometer placement on the human body, and determine the location of the accelerometer that gives the largest amount of information about each activity, as well as the accelerometer configuration that results in the best activity recognition accuracy. We find the number of hidden states that best models each activity by achieving the lowest test error using K-fold cross-validation.

Three subjects were asked to wear the set of accelerometers and perform the following sequence of activities: sit down, run, duck, walk, stand, crawl, lay down (on the chest), and perform small hand movements (while standing). The data collection process was repeated three times for each subject. 90 seconds of data were collected for each activity and labeled according to the start and stop times of each activity performed. The data was divided into nine datasets (three datasets per subject) and an HMM was trained for each activity using eight datasets and tested on the ninth dataset. Each sequence in the test dataset was classified according to the HMM that resulted in the maximum likelihood. This process was repeated nine times, each time using different training and test sets to obtain the K-fold cross-validation classification accuracy and test errors.

In the following sections we describe the data collection process and the accelerometer calibration procedure that we followed. We also show how to find the number of hidden states that best models each activity. Our results show that it is possible to improve the classification accuracy using different configurations of accelerometer placement and modeling each activity with a different number of hidden states.

3.1 Data Collection and Processing

Three wireless accelerometers (*MITes: MIT Environmental Sensors*) [59] were placed on different parts of the body: right wrist (A1), left hip (A2), and chest (A3). We selected these particular locations on the basis of lengthy discussions with potential user groups

concerning their general acceptance of wearable sensors, and because existing devices that are already worn on those parts of the body (i.e. watches, mobile phones, and electronic badges) may incorporate accelerometers in the near future.

Each subject was asked to perform eight different activities in a predefined order. The start time and stop time of each activity was recorded and the data was labeled according to these timestamps. This process was repeated three times for each subject giving a total of nine datasets (three datasets per subject). The order of the activities performed by each subject was: 1) sitting, 2) running, 3) ducking, 4) walking, 5) standing, 6) crawling, 7) lying down, 8) small hand movements. The data collection process was repeated three times for each subject. 90 seconds of data were collected for each activity and labeled according to the start and stop times of each activity performed. The data were divided into nine datasets, each of them containing 80 observation sequences of one-second duration. An HMM was trained for each activity using eight datasets (640 observations) and tested on the ninth dataset (80 observations).

The number of samples that we obtained after collecting each dataset was different for each accelerometer. Therefore, it was necessary to perform a linear interpolation to resample the data and get the same number of samples from each accelerometer. We resampled each accelerometer at $f_s' = 50$ Hz. Figure 3-1 shows dataset 1 (subject 1) after resampling each accelerometer. Blue rectangles enclose the acceleration data from the three axes corresponding to the different activities performed by each subject in the order previously specified.

3.2 Accelerometer Calibration

The range of values for the signals recorded by each accelerometer in the three axes varied from 0 to 399. This range corresponds to acceleration values between $-2g$ and $+2g$. To normalize the signals a calibration procedure is necessary to obtain the absolute value of gravity, $|\vec{g}|$, and the zero gravity point $\vec{g}_0 = (g_{x_0}, g_{y_0}, g_{z_0})$.

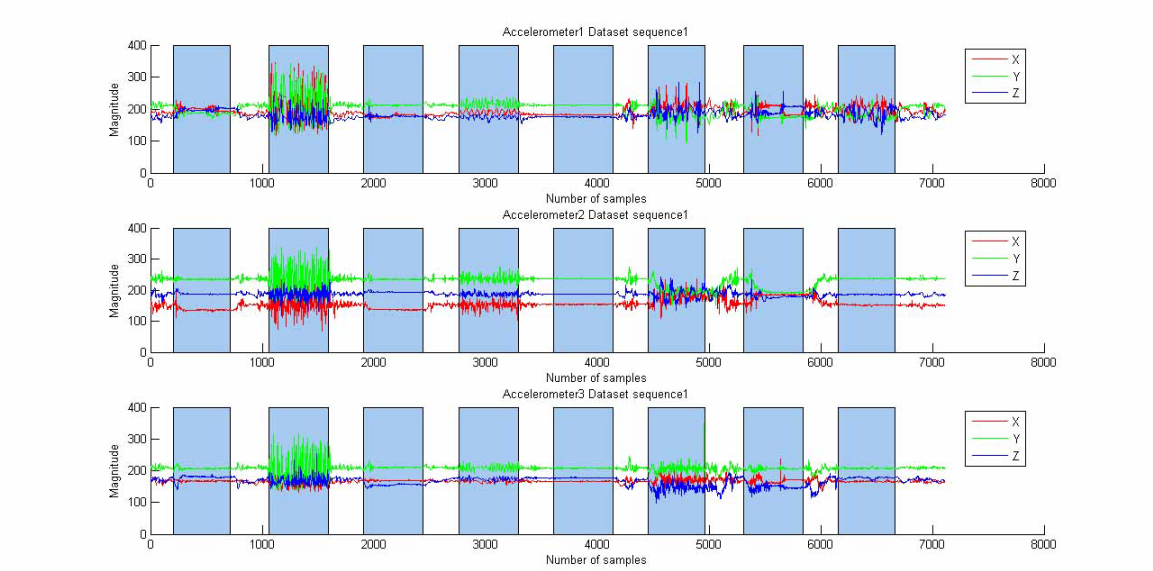

Figure 3-1: Acceleration data recorded from the three different accelerometers: right wrist (top), left hip (middle), chest (bottom). Dataset 1, (subject 1) after resampling.

To obtain these values the three accelerometers were placed together and slowly rotated in all directions. The Euclidean distance d_i from the i^{th} sample point $\vec{a}_i = (a_{x_i}, a_{y_i}, a_{z_i})$ to the zero gravity point is:

$$d_i = \sqrt{(a_{x_i} - g_{x_0})^2 + (a_{y_i} - g_{y_0})^2 + (a_{z_i} - g_{z_0})^2}. \tag{3.1}$$

The zero gravity point is obtained by minimizing the variance of this distance

$$\vec{g}_0 = \arg\min_{\vec{a}} \sigma_d^2, \tag{3.2}$$

where

$$\sigma_d^2 = \frac{1}{n}\sum_{i=1}^{n}(d_i - \mu_d)^2, \tag{3.3}$$

and

$$\mu_d = \frac{1}{n}\sum_{i=1}^{n} d_i, \tag{3.4}$$

which corresponds to the absolute value of gravity, $|\vec{g}|$

$$|\vec{g}| = \mu_d. \tag{3.5}$$

The data recorded from each accelerometer was then normalized as follows:

$$\vec{a}_i' = \frac{\vec{a}_i - \vec{g}_0}{|\vec{g}|} \tag{3.6}$$

3.3 Feature Extraction

Each dataset is composed of raw acceleration data in the x, y, and z axes from each of the three accelerometers. The mean and variance in acceleration for each axis was calculated over 200-millisecond time slices and one-second-long observation vectors were formed by grouping five time-slices. These features have been successfully used in [19] and [39]. Other commonly used acceleration features that could also be explored include correlation coefficients, fast Fourier transform coefficients, energy, entropy, root mean square value, among others.

3.4 Human Activity Modeling Using HMMs

Hidden Markov models are typically used for speech recognition and many other applications of machine learning and pattern recognition. An HMM is a stochastic signal model that represents the probability distributions over a sequence of observations. The HMM assumes that the observation at time t was generated by some process whose state S_t is hidden from the observer. It also assumes that the state of this hidden process satisfies the Markov property; that is, given the value of the previous state S_{t-1}, the current state S_t is independent of all the states prior to $t-1$ [33].

Following [69], an HMM is characterized by:

1) N, the number of states in the model. Although the states are hidden, there is often some physical significance attached to the states of the model. In this application, we model the different body postures of each activity with different hidden states. We denote the individual states as $S = \{S_1, S_2, ..., S_N\}$ and the state at time t as q_t.

2) M, the number of distinct observation symbols per state. The observation symbols correspond to the physical output of the system being modeled, which in this case are formed by the features extracted from the sensor readings. We denote the individual symbols as $V = \{v_1, v_2, ..., v_M\}$.

3) The state transition probability distribution $\mathbf{A} = \{a_{ij}\}$, where

$$a_{ij} = P[q_{t+1} = S_j | q_t = S_i], 1 \leq i, j \leq N. \tag{3.7}$$

4) The observation symbol probability distribution in state j, $\mathbf{B} = \{b_j(k)\}$, where

$$b_j(k) = P[v_k \text{ at } t | q_t = S_j], \quad 1 \leq j \leq N$$
$$1 \leq k \leq M. \tag{3.8}$$

5) The initial state distribution $\pi = \{\pi_i\}$ where

$$\pi_i = P[q_1 = S_i], 1 \leq i \leq N. \tag{3.9}$$

We can specify the complete set of parameters for an HMM as $\lambda = (\mathbf{A}, \mathbf{B}, \pi)$. Given a number of observations, there are three problems that we might want to solve [11]:

1) Given a model λ, we would like to evaluate the probability of any given observation sequence $O = \{O_1 O_2 ... O_T\}$, namely $P(O|\lambda)$.

2) Given a model λ, and the observation sequence O, we would like to find out the state sequence $Q = \{q_1 q_2 ... q_T\}$, which has the highest probability of generating O, namely, we want to find Q^* that maximizes $P(Q|O, \lambda)$.

3) Given a training set of observation sequences, $\mathcal{X} = \{O^l\}_L$, we would like to learn the model that maximizes the probability of generating $\mathcal{X}$, namely, we want to find λ^* that maximizes $P(\mathcal{X}|\lambda)$.

In this work, we are interested in solving problems (1) and (3). Given a training set of labeled observation sequences (features extracted from the acceleration readings in the x, y, and z axes from three accelerometers placed on different parts of the body), corresponding to each of the activities that we want to classify, we first want to estimate the model parameters λ for each activity (problem 3). Given a new set of observations we would like to classify each sequence according to the model that gives the maximum likelihood for that particular sequence (problem 1).

Figure (3-2) shows how we model each observation sequence as a 5-state left-right HMM with continuous Gaussian observation vectors and 2 hidden states. Each observation vector is formed by combining the mean and variance in the x, y, and z axes from each accelerometer.

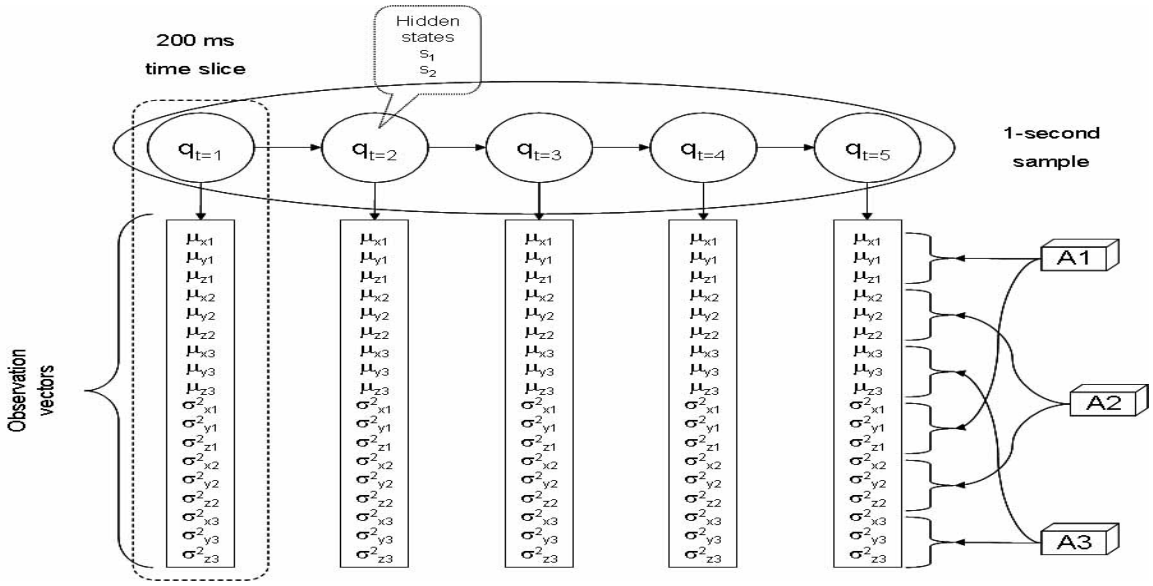

Figure 3-2: HMM representation of an observation sequence using the three accelerometers.

An HMM was trained for each class $(\lambda_1, \lambda_2, ..., \lambda_C)$, where λ_c indicates the learned HMM model for class c, and $C = 8$ is the total number of classes, using the labeled data from 8

datasets as training dataset $\mathcal{T}_k$. The ninth data set was used as the validation dataset $\mathcal{V}_k$.

The HMM toolbox for Matlab developed by [60] was used to train and test the different models. The log-likelihood of each model was calculated for each observation sequence in the ninth dataset, and each observation sequence $O^l = \{O_1^l O_2^l ... O_T^l\}$ (with $T = 5$ time slices) in the validation dataset $\mathcal{V}_k = \{O^l\}_L$ was classified according to the model that gave the maximum likelihood.

The probability of O^l given a model λ_c is obtained by summing the joint probability of O and Q over all possible state sequences q:

$$P(O^l|\lambda_c) = \sum_{\forall Q} P(O^l|Q, \lambda_c)P(Q|\lambda_c), \tag{3.10}$$

$$P(O^l|\lambda_c) = \sum_{i=1}^{N} \alpha_T(i). \tag{3.11}$$

Equation (3.11) has been written in terms of the *forward variable* $\alpha_t(i)$, defined as the probability of observing the partial sequence $\{O_1 O_2 ... O_t\}$ until time t and being in state S_i at time t, given the model λ_c,

$$\alpha_t(i) \equiv P(O_1 O_2 ... O_t, q_t = S_i|\lambda_c) \tag{3.12}$$

The likelihood that an observation sequence belongs to a particular model is calculated as follows:

$$\ell(\lambda_c) = \prod_{l=1}^{L} P(O^l|\lambda_c), \forall c \tag{3.13}$$

$$\mathcal{L}(\lambda_c) = \sum_{l=1}^{L} \log P(O^l|\lambda_c), \forall c \tag{3.14}$$

Combining equations (3.11) and (3.14):

$$\mathcal{L}(\lambda_c) = \sum_{l=1}^{L} \log \sum_{i=1}^{N} \alpha_T(i), \forall c \tag{3.15}$$

The final classification is obtained as

$$\hat{G}(O^l) = \arg\max_c \mathcal{L}(\lambda_c). \tag{3.16}$$

The classifier $\hat{G}$ takes values in the class set $G = \{1, 2, ...C\}$.

This process was repeated $K = 9$ times using k-fold cross-validation. The following sections describe the average cross-validation test error and classification accuracy per class for the 4 different possible configurations of accelerometer placement shown in table 3.1.

	Configuration			
	C1	C2	C3	C4
Right wrist	YES	YES		YES
Left hip	YES		YES	YES
Chest		YES	YES	YES

Table 3.1: Accelerometer configurations.

3.5 K-fold Cross-validation

The entire dataset $\mathcal{X}$ was divided into $K = 9$ equal-sized parts, $\mathcal{X}_k$, $k = 1, ..., K$. To generate each pair of training and testing sets we kept one of the K parts out as the validation set $\mathcal{V}_k$, and combined the remaining $K - 1$ parts to form the training set $\mathcal{T}_k$. Doing this K times, each time leaving one of the other K parts out, we obtained the K pairs:

$$\mathcal{V}_1 = \mathcal{X}_1 \qquad \mathcal{T}_1 = \mathcal{X}_2 \cup \mathcal{X}_3 \cup \cdots \mathcal{X}_K$$
$$\mathcal{V}_2 = \mathcal{X}_2 \qquad \mathcal{T}_2 = \mathcal{X}_1 \cup \mathcal{X}_3 \cup \cdots \mathcal{X}_K$$

$$\mathcal{V}_K = \mathcal{X}_K \qquad \mathcal{T}_K = \mathcal{X}_1 \cup \mathcal{X}_2 \cup \cdots \mathcal{X}_{K-1}$$

3.6 Confusion Matrix

Our error calculation is based on the $C \times C$ confusion matrix. With $C = 8$ different classes (corresponding to each of the activities that we want to be able to discern).

	Predicted Class	
True class	**Yes**	**No**
Yes	TP: True Pos	FN: False Neg
No	FP: False Pos	TN: True Neg

Table 3.2: Confusion matrix.

The class confusion matrix is a $C \times C$ matrix such that its entry (i, j) contains the number of instances that belong to class i but are assigned to class j. Ideally, all off-diagonals should be 0, for the case of no misclassification.

We define the accuracy per class c as the proportion of instances correctly classified in each class as

$$\text{accuracy}_c = \frac{|\text{TP}|_c}{|\text{TP}|_c + |\text{FN}|_c} \tag{3.17}$$

Under 0/1 loss we define the error rate for each class c as:

$$\text{error}_c = \frac{|\text{FP}|_c + |\text{FN}|_c}{|\text{FP}|_c + |\text{FN}|_c + |\text{TP}|_c}. \tag{3.18}$$

3.7 Results

Figure 3-3 shows a comparison of classification accuracy when using a single accelerometer for activity classification. We are able to discern activities such as sitting (66.05% accuracy) and walking (65.68% accuracy) using only accelerometer A1 (right wrist). Accelerometer A2 (left hip) plays the most important role when classifying activities such as running (97.78% accuracy), crawling (69.26% accuracy), and lying down (87.04% accuracy). Accelerometer A3 (chest) is important for classifying activities such as ducking (75.8% accuracy) and standing (77.78% accuracy).

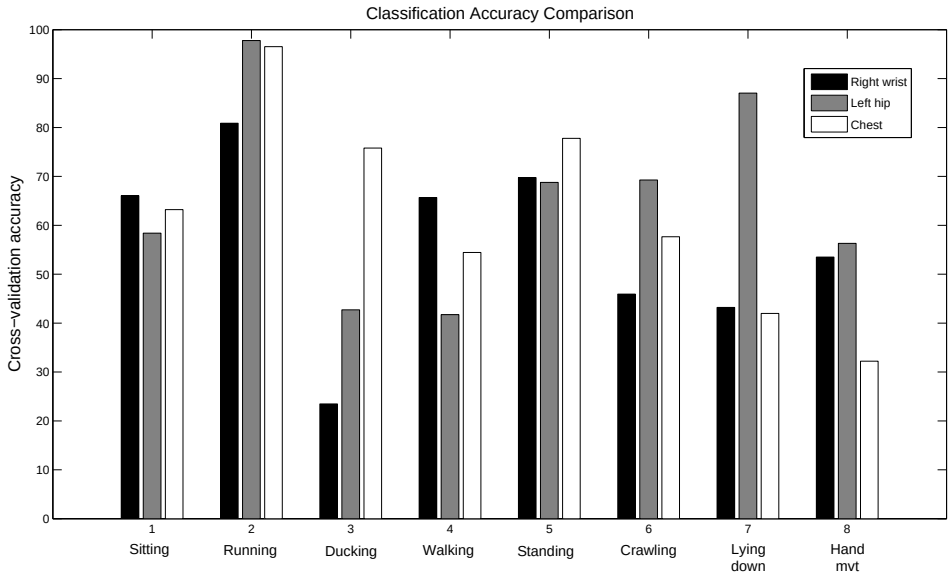

Figure 3-3: K-fold cross validation accuracy comparison using a single accelerometer.

Figures (3-4) to (3-7) show the average cross-validation test error per class using the different configurations listed in table 3.1. Each box has lines at the lower quartile, median, and upper quartile values. Lines extending from each end of the box show the extent of the rest of the data. Outliers are data points beyond the ends of the whiskers.

3.7.1 Configuration C1

This configuration of accelerometer placement achieves a low test error on activities such as running, standing, and performing small hand movements. The action of ducking is usually confused with sitting since it might involve the action of sitting during the process of ducking. The action of sitting gets also confused with ducking. Walking is sometimes confused with standing or running since different subjects might walk at different speeds, and usually standing is followed by walking. Crawling is confused with lying down and vice versa, since both activities are not mutually exclusive.

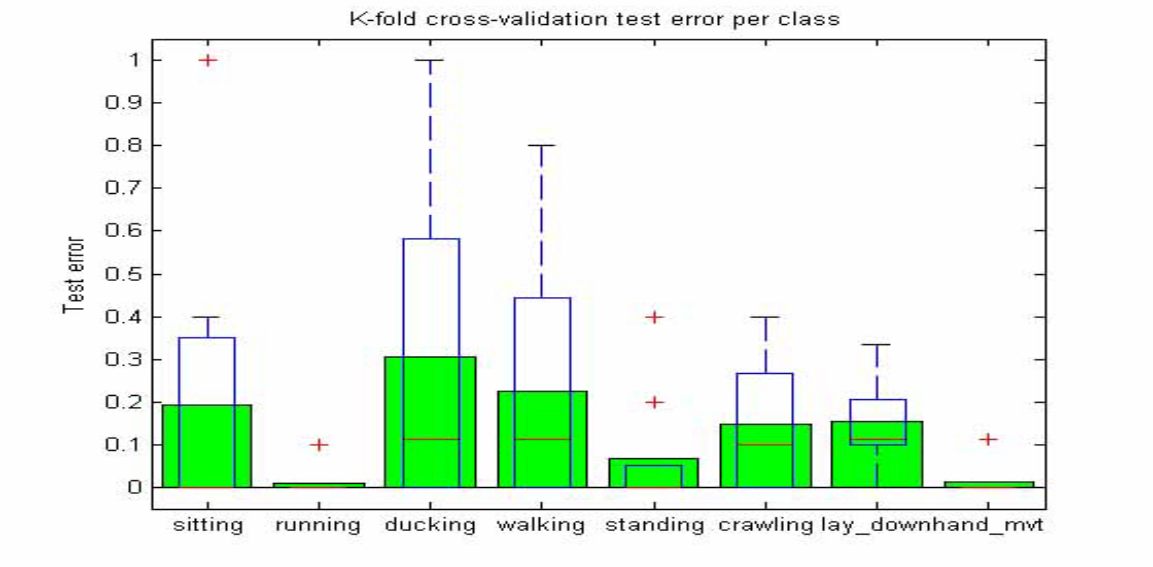

Figure 3-4: K-fold cross validation test error per class (configuration C1).

Classified as⇒	Sitting	Running	Ducking	Walking	Standing	Crawling	Lying down	Hand mvt
Sitting	76	0	9	1	0	0	0	0
Running	0	86	0	1	0	0	0	0
Ducking	23	0	61	1	0	0	0	0
Walking	0	2	0	73	11	0	0	0
Standing	0	1	0	6	77	0	0	2
Crawling	0	0	0	3	0	72	10	1
Lying down	1	0	0	3	0	10	69	2
Hand mvt	0	0	0	0	0	1	0	81

Table 3.3: K-fold cross validation confusion matrix for configuration C1.

3.7.2 Configuration C2

This configuration achieves a low test error rate for the activities of sitting, running, and standing. Ducking is still mainly confused with sitting, and walking is confused with standing. Crawling is confused with lying down, and vice versa.

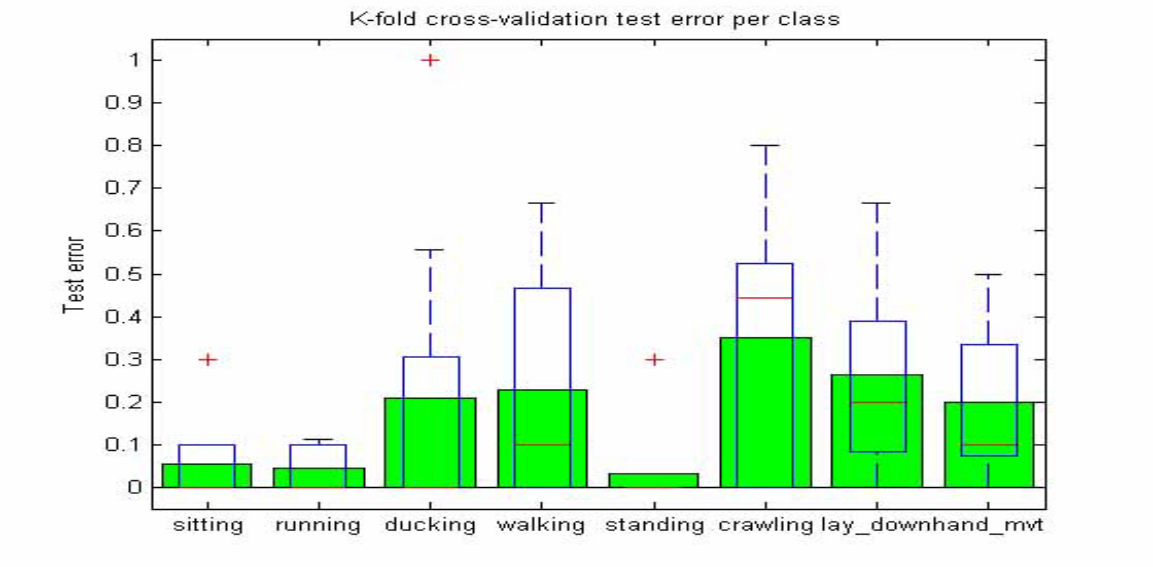

Figure 3-5: K-fold cross validation test error per class (configuration C2).

Classified as⇒	Sitting	Running	Ducking	Walking	Standing	Crawling	Lying down	Hand mvt
Sitting	81	0	3	1	0	0	1	0
Running	0	83	0	1	0	2	1	0
Ducking	11	0	58	2	3	2	0	9
Walking	0	1	0	67	18	0	0	0
Standing	3	0	0	1	81	0	0	1
Crawling	0	0	1	0	0	59	24	1
Lying down	0	0	2	0	1	11	65	6
Hand mvt	6	0	1	0	1	2	6	66

Table 3.4: K-fold cross validation confusion matrix for configuration C2.

3.7.3 Configuration C3

This configuration achieves low test error on classifying the activities of sitting, running, and standing. We can see that the lack of an accelerometer on the wrist causes an increase in the test error when trying to classify walking and small hand movements. The activities of crawling and lying down are still confused with each other.

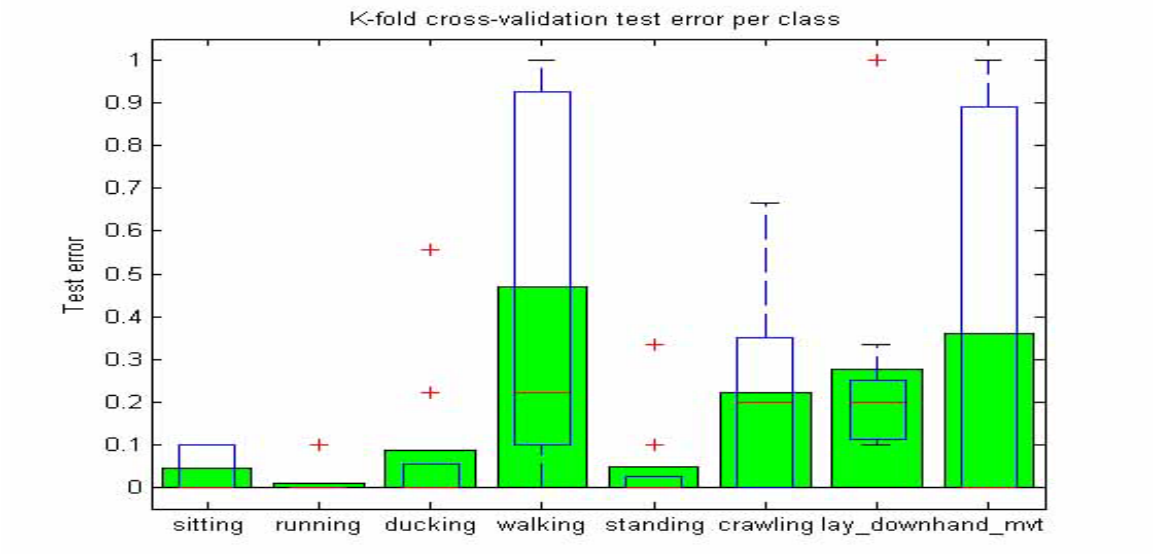

Figure 3-6: K-fold cross validation test error per class (configuration C3).

Classified as⇒	Sitting	Running	Ducking	Walking	Standing	Crawling	Lying down	Hand mvt
Sitting	82	0	1	0	0	0	0	3
Running	0	87	0	0	0	0	0	0
Ducking	5	0	79	0	0	0	0	1
Walking	1	1	0	53	28	0	0	3
Standing	1	1	0	0	84	0	0	0
Crawling	0	0	0	0	0	63	22	1
Lying down	0	0	0	0	1	8	73	3
Hand mvt	0	1	1	0	36	0	0	44

Table 3.5: K-fold cross validation confusion matrix for configuration C3.

3.7.4 Configuration C4

We can observe that using the three accelerometers in configuration C4 achieves a lower test error in most of the activities compared to other configurations. It achieves the best results when classifying activities such as sitting, running, standing, and small hand movements. Crawling and lying down are still confused for each other but in fewer instances.

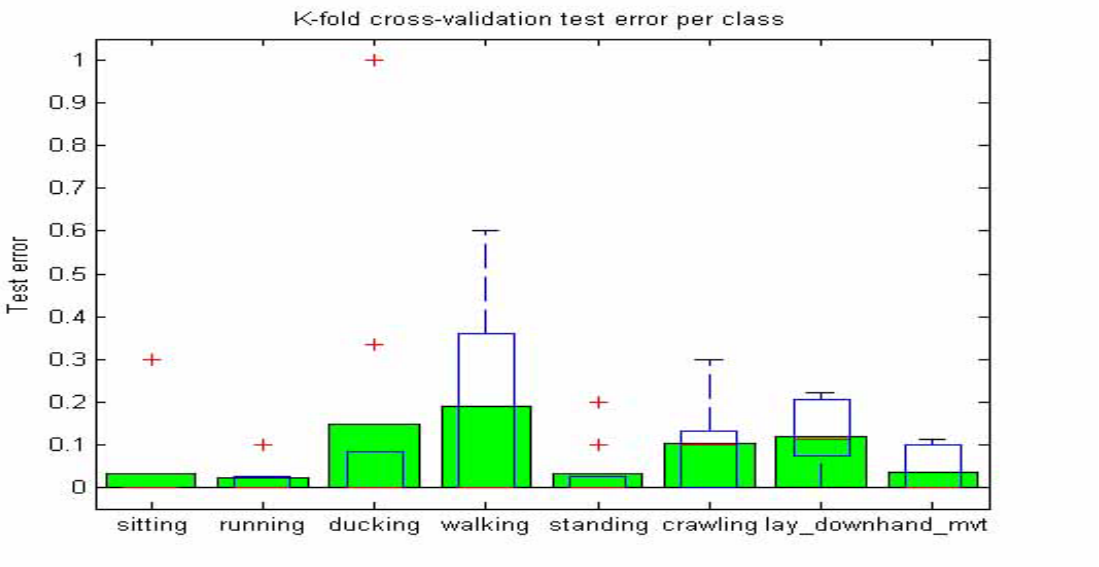

Figure 3-7: K-fold cross validation test error per class (configuration C4).

Classified as⇒	Sitting	Running	Ducking	Walking	Standing	Crawling	Lying down	Hand mvt
Sitting	84	0	0	2	0	0	0	0
Running	0	86	0	1	0	0	0	0
Ducking	14	0	70	1	0	0	0	0
Walking	0	0	0	73	13	0	0	0
Standing	0	0	0	1	82	0	0	3
Crawling	0	0	0	0	0	79	6	1
Lying down	1	0	0	1	0	8	74	1
Hand mvt	0	1	0	0	0	0	0	81

Table 3.6: K-fold cross validation confusion matrix for configuration C4.

3.8 Average Classification Accuracy

Figure 3-8 shows the average classification accuracy per activity when combinations of two and three accelerometers were placed in the four different configurations described in table 3.1. Table 3.7 shows a comparison between the classification accuracy obtained when a single accelerometer (A1 to A3) was used, and the classification accuracy obtained when multiple accelerometers were used (C1 to C4).

Our results show that it is possible to recognize some of the most common activities using

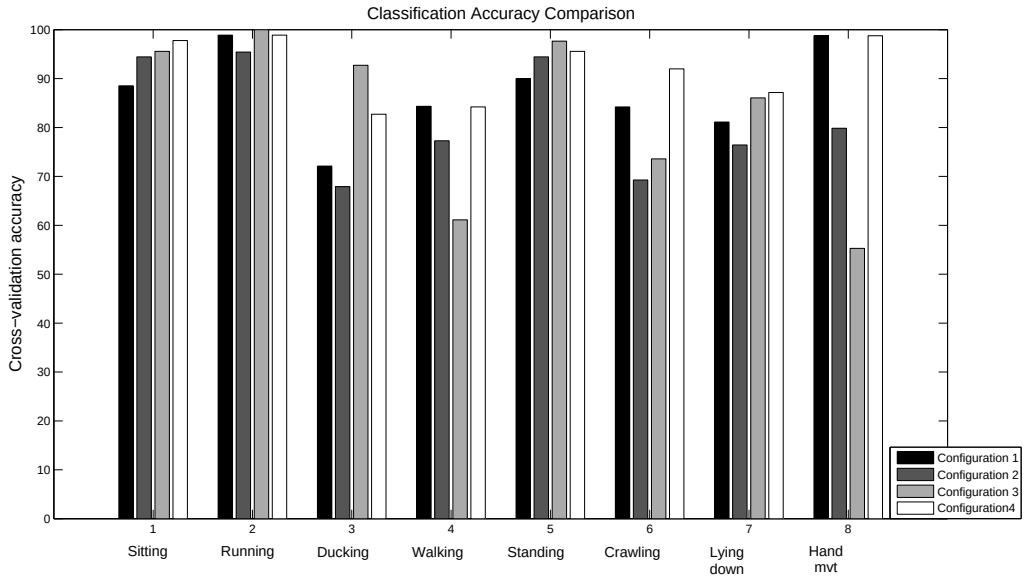

Figure 3-8: K-fold cross validation accuracy comparison for the four different configurations of accelerometer placement.

a single accelerometer on the chest (A3) with 62.45% average accuracy. Adding a second accelerometer on the hip or the wrist improved our classification accuracy by approximately 20%. Adding a third accelerometer improved the global classification accuracy by an additional 10%.

3.9 Finding the Optimum Number of Hidden States

The results presented so far were obtained by modeling each activity with an HMM having 2 hidden states. However, we think that classification results could be improved by modeling each activity with a different number of hidden states. In some cases, 2 hidden states might not be enough for capturing the different stages of a particular activity, especially when the activity involves different movements and body positions.

Figure (3-9) shows the cross-validation training and test errors as a function of model complexity (number of hidden states) for the eight classes using configuration C4 of accelerometer placement. We can observe how the training error decreases as the number of

Class	A1	A2	A3	C1	C2	C3	C4
Sitting	58.40%	**66.05%**	63.21%	88.52%	94.44%	95.56%	**97.78%**
Running	80.86%	**97.78%**	96.54%	98.89%	95.43%	**100.0%**	98.89%
Ducking	23.46%	42.72%	**75.80%**	72.10%	67.90%	**92.72%**	82.72%
Walking	**65.68%**	41.73%	54.44%	**84.32%**	77.28%	61.11%	84.20%
Standing	69.75%	68.77%	**77.78%**	90.00%	94.44%	**97.65%**	95.56%
Crawling	45.93%	**69.26%**	57.65%	84.20%	69.26%	73.58%	**91.98%**
Lying down	43.21%	**87.04%**	41.98%	81.11%	76.42%	86.05%	**87.16%**
Hand mvt	**56.30%**	53.51%	32.22%	**98.77%**	79.85%	55.28%	**98.77%**
Global	55.45%	65.86%	62.45%	87.24%	81.88%	82.74%	92.13%

Table 3.7: Classification accuracy comparison.

hidden states for modeling each activity is increased, while the test error decreases up to a minimum value (at Q_{min}), and then it starts to increase due to possible overfitting of the models to the training data.

Table 3.8 shows the cross-validation accuracy when modeling each activity with 2 hidden-state HMMs and when using the number of hidden states, Q_{min}, that gives the minimum test error for configuration C4 of accelerometer placement. We can see that there was an improvement in classification accuracy in four of the activities (walking, standing, lying down, and small hand movements). There was no change in classification accuracy for running and ducking, and there was a decrease in classification accuracy for sitting and crawling. Based on these results, we choose the number of hidden states, Q_{opt}, that best models each of the activities studied in this chapter.

Class	$Q = 2$	Q_{min}	Var	Q_{opt}
Sitting	97.78%	93.33%	-4.45%	2
Running	98.89%	98.89%	0%	2
Ducking	82.72%	82.72%	0%	2
Walking	84.20%	88.64%	+4.44%	4
Standing	95.56%	96.67%	+1.11%	4
Crawling	91.98%	86.29%	-5.69%	2
Lying down	87.16%	87.26%	+0.1%	4
Hand mvt	98.77%	99.01%	+0.24%	5

Table 3.8: Classification accuracy variation and optimal number of hidden states.

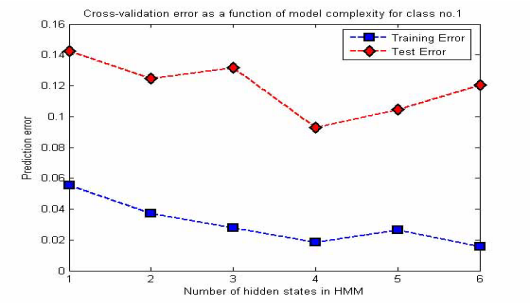

(a) Class 1 (sitting) $Q_{min} = 4$, $Q_{opt} = 2$

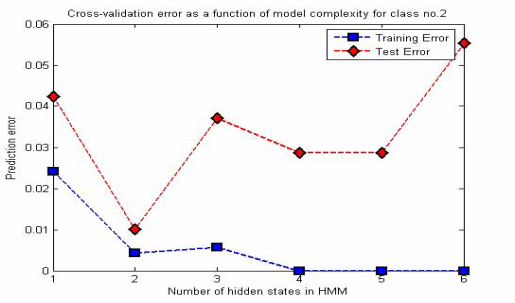

(b) Class 2 (running) $Q_{min} = 2$, $Q_{opt} = 2$

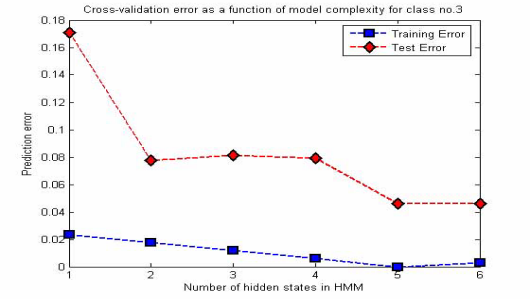

(c) Class 3 (ducking) $Q_{min} = 5$, $Q_{opt} = 2$

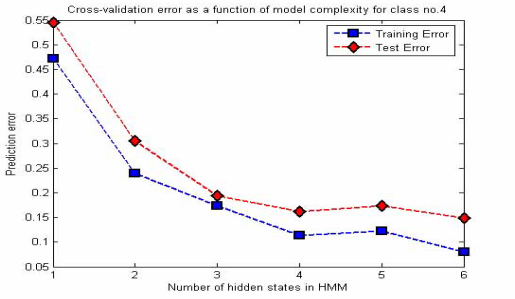

(d) Class 4 (walking) $Q_{min} = 4$, $Q_{opt} = 4$

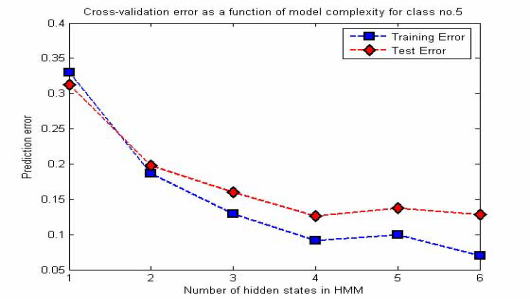

(e) Class 5 (standing) $Q_{min} = 4$, $Q_{opt} = 4$

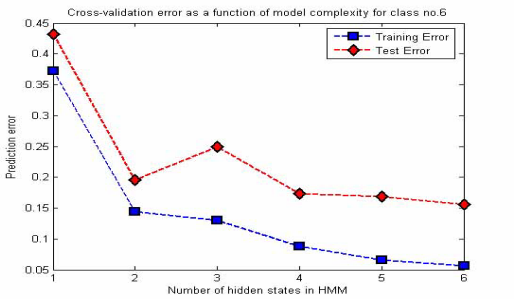

(f) Class 6 (crawling) $Q_{min} = 6$, $Q_{opt} = 2$

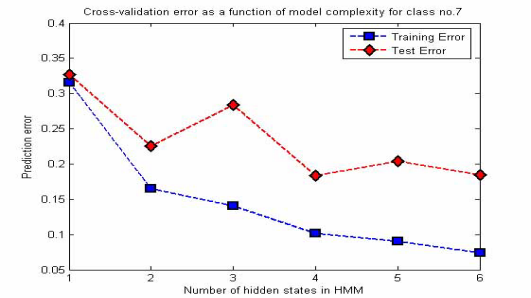

(g) Class 7 (lying down) $Q_{min} = 4$, $Q_{opt} = 4$

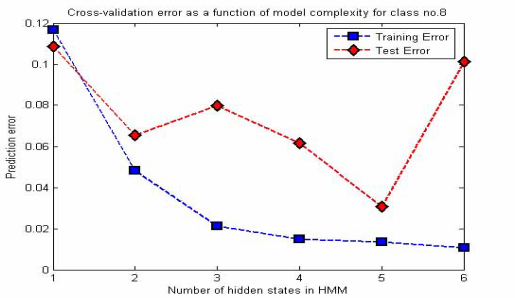

(h) Class 8 (hand mvt) $Q_{min} = 5$, $Q_{opt} = 5$

Figure 3-9: Cross-validation training and test errors as a function of model complexity (number of hidden states) for each activity using configuration C4 of accelerometer placement.

3.10 Remarks

In this chapter we studied human activity classification using wearable accelerometers and modeled eight different activities (sitting, running, ducking, walking, standing, crawling, lying down, and performing small hand movements) using hidden Markov models with two hidden states and continuous Gaussian observations. We compared the performance of four different configurations of accelerometer placement: (C1) right wrist and left hip, (C2) right wrist and chest, (C3) hip and chest, and (C4) right wrist, left hip, and chest; in terms of activity classification accuracy and average test error using K-fold cross-validation. We found that the best global classification performance was achieved when using configuration C4 (92.13%), although it might be possible to obtain similar results using only two accelerometers following configuration C1 (87.24%).

We showed that the global classification accuracy that can be achieved using a single accelerometer is around 60% and determined the activities that are best classified with each accelerometer placed on a different part of the body. Increasing the number of hidden states for modeling activities such as sitting, running, ducking, and crawling did not increase the classification accuracy (and in some cases it even decreased), so modeling these activities with two hidden-state HMMs seems to be appropriate. On the other hand, activities such as standing, crawling, lying down and small hand movements can be better modeled with a higher number of hidden states. We are convinced that a minimal system formed by a wearable badge and a mobile phone can achieve fairly good results in daily activity recognition (80%).

Chapter 4

Hardware Design

4.1 Initial Requirements

When designing portable electronic devices the functionality must be clearly defined, focused, and bounded [36]. The primary function of the sociometric badges is to automatically collect social signals derived from speech, motion, proximity, and face-to-face interaction over extended periods of time. As a secondary function, the user should be able to interact with the device by issuing voice commands or pressing a push-button. During the planning phase of the project we evaluated the feasibility of designing and manufacturing several hundred sociometric badges capable of measuring the aforementioned variables. Figure 4-1 shows a strawman drawing by Joseph Paradiso during a one-hour brainstorming session with several students from the *Responsive Environments* group and visiting scientists from Hitachi Japan.

During the brainstorming session the basic requirements of the system were specified:

- Each sociometric badge should be able to detect whether the user is speaking or not by using a set of four analog band-pass filters capable of extracting different speech features (such as amplitude, energy, and entropy) in each of the four frequency bands and estimate voice pitch based on these features.

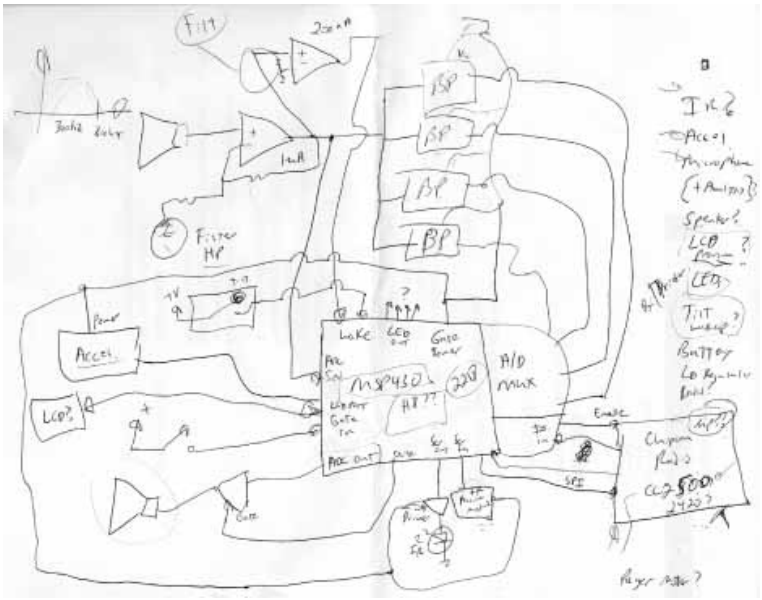

Figure 4-1: Strawman drawing.

- Detect when two badges are facing each other in order to measure the user's face-to-face interaction time.
- Detect when the user is moving or stationary in order to automatically switch between different power consumption modes.
- Communicate with base stations in the 2.4 GHz frequency band in order to transfer data.
- Store data onboard and transfer it wirelessly when the user is in close proximity to base stations.
- Connect with mobile phones via Bluetooth.
- Play audio messages through a speaker.
- Provide visual feedback to the user using either a liquid crystal display (LCD) or light emitting diodes (LEDs).
- Enable the user to use the badge as a push-to-talk system and interact with it using voice commands.

- Have a long battery life so that the user doesn't need to recharge it for at least one day.
- It must be small and comfortable to wear over extended periods of time.

Another constraint was the cost of the badges. Section 4.18 details the resources that were available for this project. The goal was to build several hundred badges with an original estimated cost of $150 per badge. The final cost was $250 per badge including battery, charger, memory card, and custom enclosure. Three hundred sociometric badges were manufactured in total.

In the following sections we describe the choice of electronic components used in the sociometric badges.

4.2 Functional Description

All the circuits were designed to work with a 3.3 V single supply. A voltage reference equal to one half of the supply voltage is required. Therefore, a micro-power bias voltage generator using an ultra low power, rail to rail output operational amplifier (OP281) with a supply current of 4 μA per amplifier is used.

The output of an omni-directional micro electro mechanical systems (MEMS) microphone (SPM0103-NE3) is connected to the non-inverting input of a micro-power single-stage non-inverting operational amplifier (AD8542), with a high-pass filtering cut-off frequency of 85 Hz realized via its feedback loop. The amplified microphone signal is then applied to an array of micro-power single-op-amp Sallen-Key band-pass filters that divide the speech frequency spectrum into four octaves: f_1 from 85 to 222 Hz, f_2 from 222 to 583 Hz, f_3 from 583 to 1527 Hz, and f_4 from 1527 to 4000 Hz. These frequency bands were selected in order to achieve a constant quality factor $Q = \sqrt{2}/(4-K)$ in each of the filters (where $K = 2.58$ is the gain of each filter). A diode-capacitor peak detector is used after each band-pass filter to obtain the spectral envelope in each frequency band. These four spectral envelopes are

applied to the analog-to-digital converter (ADC) on the microcontroller, and are used to extract different speech features and segment audio into speaking and non-speaking regions.

A 3-axis MEMS accelerometer (ADXL330) is used to detect when a person is moving and to distinguish several common activities such as sitting, standing, walking, or running. An IR transceiver module (TFDU4300) is used to detect when two people are facing each other. All sensor electronics are powered up only when a significant sound is detected (indicating a possible conversation) or the accelerometer detects significant motion. A bridged-output audio power amplifier (SSM2211) drives an electromagnetic speaker on the badge to play back messages and reminders.

The main processing unit is an ARM microcontroller (AT91SAM7S256). A 2.4GHz radio transceiver (CC2500) and a class 2.0 Bluetooth module (BR-46AR) have been incorporated for enabling wireless communications with fixed base stations and other Bluetooth-enabled devices. A microSD memory card socket has been included for storing data when the user is out of range of a fixed point or when the badge is used as a self contained sensor package. The badge is powered by a 950 mAh lithium-polymer battery rechargeable through USB. In addition, data can also be transferred through the USB port. Figure 4-2 shows a block diagram of the badge.

4.3 Form Factor

Figure 4-4 shows the front 4-4(a) and back 4-4(b) views of the first prototype's assembled printed circuit board (PCB). Figure 4-5 shows the front 4-5(a) and back 4-5(b) views of the final prototype's assembled PCB. The dimensions of the printed circuit board (PCB) are 4.2 x 8 cm. Two design iterations were made with a few changes between the first prototype and the final version. The 2.4 GHz radio transceiver's daughter board was moved from the back of the PCB to the front in order to have both radio antennas facing out when the badge is worn and to avoid blocking the radio signals with the wearer's body. The microSD card socket was moved to the back of the PCB in the final prototype to allow for the relocation

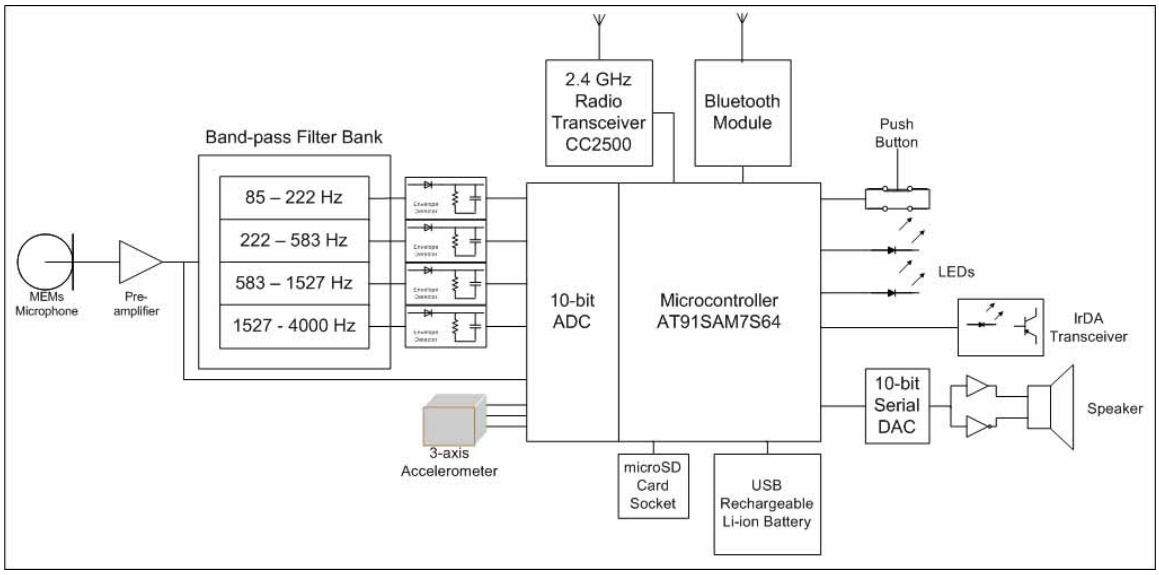

Figure 4-2: Block diagram.

of the radio board. The microphone was placed on a daughter board perpendicular to the badge board in order to orient it towards the wearer.

The placement of each sensor in the sociometric badges was chosen to optimally capture the signals listed in section 4.1. The shape of the badge is similar to some commercially available products that people already wear around the neck using a lanyard [3, 5, 10].

The dimensions of the badge inside the plastic enclosure (as shown in figure 4-6(a)) are 4.5 x 10 x 2 cm and the total weight including the battery is 110 grams. Figure 4-6(b) shows the author wearing the final prototype.

The enclosure was designed and manufactured by Toolless Plastic Technologies [9]. It is formed by three plastic layers that snap together. Figure 4-7 shows the three layers of the enclosure.

4.4 Voltage Regulator

Both lithium batteries evaluated in section 4.16 maintain an average voltage of 3.7 V during most of their discharge cycle, going from a maximum voltage of 4.2 V when fully charged to

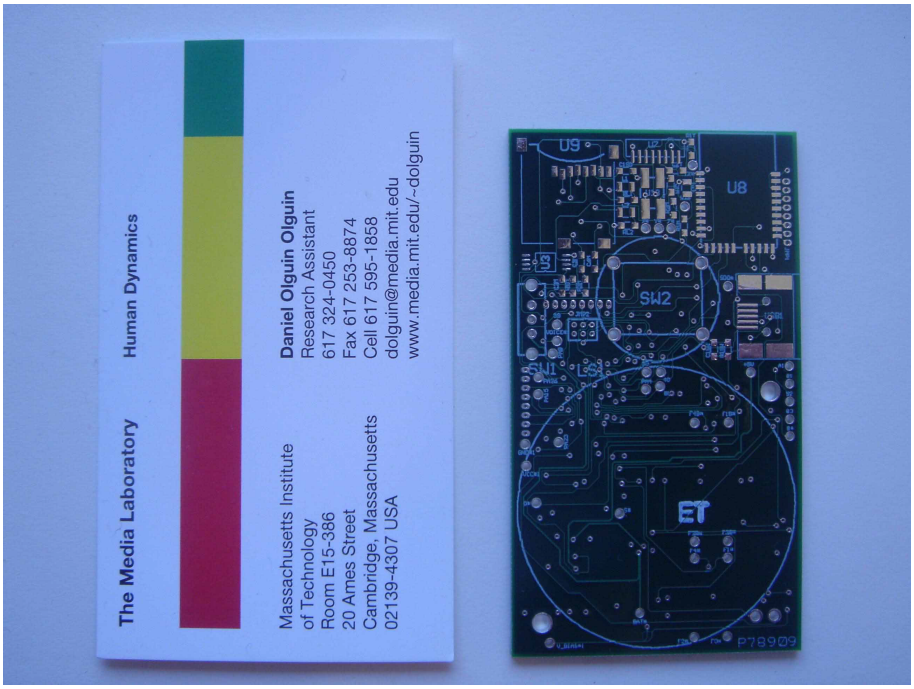

Figure 4-3: Size comparison between unassembled PCB and business card.

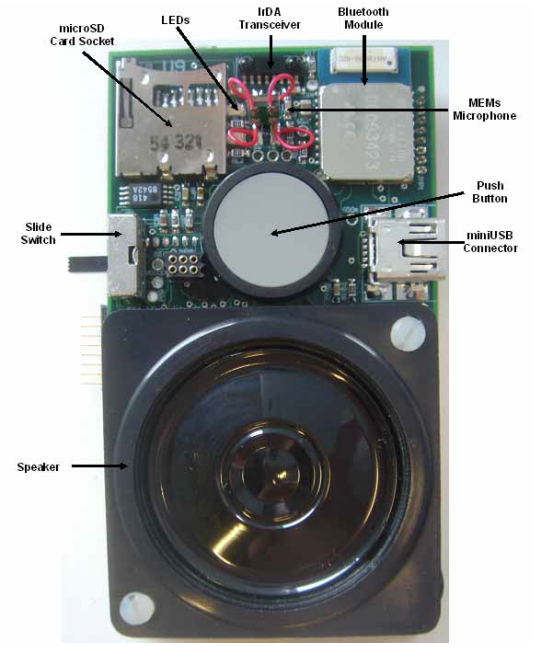

(a) PCB front view.

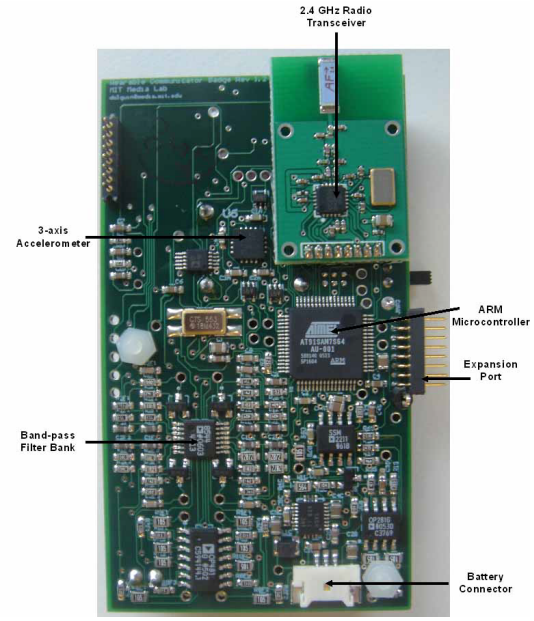

(b) PCB back view.

Figure 4-4: First prototype's assembled PCB.

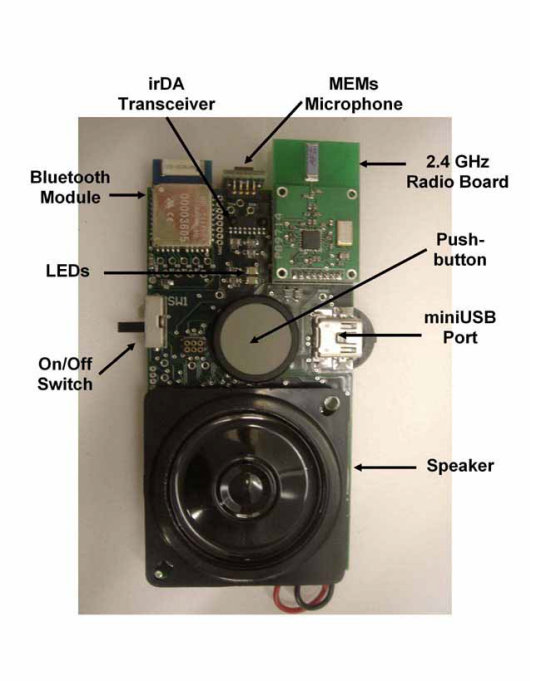

(a) Final PCB front view.

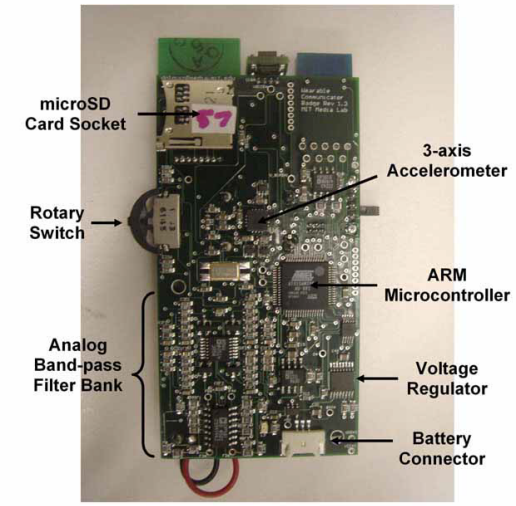

(b) Final PCB back view.

Figure 4-5: Final prototype's assembled PCB.

(a) Wearable sociometric badge.

(b) The author wearing a sociometric badge.

Figure 4-6: Final prototype inside enclosure.

Figure 4-7: Middle, back and front layers of the badge enclosure.

a recommended minimum voltage of 3 V. However, all the electronic circuits in the badge require a direct current (DC) power supply that can maintain a fixed voltage while supplying enough current. Therfore, a voltage regulator designed to automatically adjust the amount of current flowing (so as to maintain a constant output voltage) is required. Table 4.1 shows the different voltage regulators that were evaluated. Based on the maximum peak current that the speaker could draw and a tradeoff between a low dropout voltage that will maximize the battery life and the output voltage noise the MAX1793 low-dropout, low-I_Q, 1 A linear regulator by Maxim Integrated Products was selected.

Part Number	Supply Voltage Range (V)	Maximum Output Current (mA)	Regulated Voltage (V)	Dropout Voltage (mV)	Supply Current (μA)	Shutdown Supply Current (μA)	Output Noise (μVrms)	Size (mm)	Cost (US$)	Output Type
MAX1793	2.5 - 5.5	1000	1.25 - 5.0	210 @ 1A	125	0.1	115	4.9x4.3x1.1	$1.50	16pin TSSOP
MAX8869	2.7 - 5.5	1000	0.8 - 5	200 @ 1A	500	0.1	150	4.9x4.3x 1.1	$1.65	16pin TSSOP
MAX1792	2.5 - 5.5	500	1.25 - 5	130 @ 500mA	80	0.1	115	2.93x2.95x1.1	$1.15	8pin μMAX
MAX604	2.7 - 11.5	500	3.15 - 3.45	240 @ 200mA	15	2	250	5x1.75x4	$1.68	8Pin SO
MAX1806	2.25 - 5.5	500	0.8 - 4.5	175 @ 500mA	210	0.02	300	2.95x4.78x0.94	$1.40	8pin μMAX
TI REG103	2.1 - 15	1000	3.3	115 @ 500mA	1000	0.5	33	6.55x3.55x1.8	$2.50	8pin SOIC
TI TPS76833	3.65 - 10	1000	3.3	350 @ 1A	85	1	55	4.5x5.1x 1.2	$0.90	8pin SOIC
ADP3338	2.75 - 8	1000	3.3	190 @ 1A	30,000	0	95	6.5x3.7x1.7	$1.40	3pin SOT-223
ADP1716	2.7 - 5.5	1000	3.3	259 @ 500mA	650	0	450	3.2x3.2x0.95	$0.96	8pin MSOP
LT3028	1.8 - 20	500	1.8 - 20	320 @ 500mA	30	< 0.1	20	5x3x0.75	$3.78	16pin TSSOP

Table 4.1: Linear voltage regulators comparison.

The MAX1793 low-dropout linear regulator (LDO) operates from +2.5V to +5.5V and delivers a guaranteed 1A load current with a low 210 mV dropout. The high-accuracy (±1%) output voltage is preset at an internally trimmed voltage of 3.3 V. An internal PMOS pass transistor allows a low 125 μA supply current, making the MAX1793 ideal for battery operated portable equipment. Other features include a built-in reset output, low-power shutdown, and short-circuit and thermal overload protection. The MAX1793 is available in a 1.5W, 16-pin power TSSOP package, which is 30% smaller than a SOT223 package and only 1.1 mm high [52]. Figure C-3 shows the schematic diagram of the voltage regulator and battery charger.

4.5 Biasing Circuit

A voltage reference equal to one half of the supply voltage is required. Figure C-2 shows the design of a micro-power bias voltage generator using the OP281 ultra low power, rail to rail output operational amplifier by Analog Devices with a supply current of 4 μA per amplifier. The two 1 MΩ resistors generate the reference voltage while drawing only 1.65 μA of current from a 3.3 V supply. A capacitor connected from the inverting terminal to the output of the op amp provides compensation to allow for a bypass capacitor to be connected at the reference output. This bypass capacitor helps to establish an alternate current (AC) ground for the reference output. The entire reference generator draws less than 5 μA from a 3.3 V supply source.

The OP281 and OP481 are dual and quad ultra-low power, single-supply amplifiers featuring rail-to-rail outputs. Each operates from supplies as low as 2.0 V and are specified at +3 V and +5 V single supply as well as ±5 V dual supplies. Fabricated on Analog Devices' CBCMOS process, the OP281/OP481 features a precision bipolar input and an output that swings to within millivolts of the supplies and continues to sink or source current all the way to the supplies. The output's ability to swing rail-to-rail and not increase supply current, when the output is driven to a supply voltage, enables the OP281/OP481 to be used as comparators in very low power systems. This is enhanced by their fast saturation

recovery time. Propagation delays are 250 ms. The OP281 dual amplifier is available in 8-lead SOIC surface-mount and TSSOP packages. The OP481 quad amplifier is available in narrow 14-lead SOIC and TSSOP packages [12].

4.6 USB Battery Charger

Present-day lithium chemistries are typically 4.1V to 4.2V when the cells are charged to their maximum-rated capacity. Typical prismatic lithium-ion (Li+) and lithium-polymer (Li-Poly) have capacities of 600 mA to 1900 mA. The preferred charge profile for both Li+ and Li-Poly cells is to start the charge with a constant charge current until the cell voltage reaches the rated voltage. When this occurs, the charger then regulates the voltage across the cell. These two regulation states are called constant-current (CC) and constant-voltage (CV) charging. Therefore, this type of charger is usually referred to as a CCCV charger. When the CCCV charger is in CV mode, the current into the cell begins to drop. For typical charge rates, between 0.5C and 1.5C, the transition between CC and CV mode occurs when the cell has accepted about 80% to 90% of its full-charge capacity. Once the charger is in CV charging mode, it monitors the cell current. When it reaches a low threshold (milliamps or tens of milliamps), the charger terminates charge [55].

The USB interface specification includes the ability to power devices. One way to use USB power is to charge batteries. Since many portable devices, like MP3 players and PDAs, exchange information with PCs, device convenience is significantly enhanced if battery charging and data exchange take place simultaneously over one cable. Combining USB and battery-powered functionality gives rise to a whole range of "untethered" devices, such as removable web cameras, that operate whether or not they are connected to a PC. In many cases, it is no longer necessary to include an awkward AC adapter. All USB host devices, like PCs and notebooks,can source at least 500 mA, or five "unit loads" per USB socket. In USB terminology, "one unit load" is 100 mA. Self-powered USB hubs can also supply five unit loads. Bus-powered USB hubs are guaranteed to supply only 1 unit load (100 mA). According to the USB specification the minimum available voltage from a USB

host or powered hub at the peripheral end of the cable is 4.5V, while the minimum voltage from a USB bus-powered hub is 4.35V. These voltages allow very little headroom when charging lithium batteries, which typically require 4.2V, making charger dropout extremely important. All devices that plug into a USB port must start out drawing no more than 100 mA. After communicating with the host, the device can determine if it can take the full 500 mA. Power is taken from pins 1 (+5V) and 5 (GND) on the series mini-B receptacle. Once connected, all USB devices must identify themselves to the host. This is called "enumeration". In the identification process, the host determines the power needs of the USB devices and gives, or withholds the confirmation for the device to increase its load from 100 mA maximum to 500 mA maximum [54].

The USB battery charger selected for the badge is the MAX1555 by Maxim Integrated Products. It charges a single-cell lithium-ion (Li+) battery from both USB and AC to USB adapter sources. It operates with no external FETs or diodes, and accepts operating input voltages up to 7V. On-chip thermal limiting simplifies PC board layout and allows optimum charging rate without the thermal limits imposed by worst-case battery and input voltage. When the MAX1555 thermal limits are reached, the charger does not shut down, but progressively reduces charging current. The MAX1555 features a CHG output to indicate charging status. The CHG output signal is connected to a red LED that turns on when the battery is charging, and off when the battery is fully charged.

With USB connected, but without DC power, charge current is set to 100 mA. This allows charging from both powered and unpowered USB hubs with no port communication required. However, the circuit was designed so that the charging current is set to 280 mA by connecting both pins 1 and 4 to the USB (+5V) pin. No input-blocking diodes are required to prevent battery drain. The MAX1555 is available in 5-pin thin SOT23 packages and operates over a -40°C to +85°C range [53]. This circuit charges the battery whenever the device is docked to USB or plugged into the AC to USB adapter. At the same time the system load is always connected to the battery, in this case through a simple linear regulator (MAX1793), which can supply up to 1 A. If the system continuously draws that amount of current while the battery is charging at 280 mA from USB, the battery will still

Figure 4-8: Badge's side view: miniUSB port, tri-directional switch, and microSD card socket.

discharge since the load current exceeds the charge current. However, the peak load current occurs only for a fraction of the total operating time, so as long as the average load current is less than charging current, the battery will still charge. Figure 4-8 shows the miniUSB port in the badge.

4.7 Microphone Daughterboard

A microphone daughter board containing a MEMS microphone and a pre-amplifier circuit was designed and placed perpendicular to the main badge board, with the microphone facing the wearer. Figure 4-9 shows a picture of the microphone daughter board.

The amplified "Mini" SiSonic SPM0103NE3 [42] microphone from Knowles Acoustics was selected because it represents a low-cost, high-performance alternative to traditional electret microphones (ECMs). Manufactured from poly-silicon, this microphone takes advantage of semi-conductor manufacturing processes that result in a robust microphone that yields a high degree of repeatability, stable acoustic performance, and flexibility for future design enhancements [4]. Table 4.2 shows a comparison of the different MEMS and electret microphones that were considered.

Part Number	Supply Voltage Range (V)	Output Impedance (Ω)	Signal to Noise Ratio (dB)	Directivity	Current Consumption (mA)	Sensitivity dB)	Maximum Input Sound Level (dB)	Size (mm)	Cost (US$)	Design Type
Knowles Acoustics **SPM0103NE3**	1.5 - 5.5	< 100	59	omni	0.100 - 0.350	-22 @1kHz	115	1.5x4.72x3.76	$4.25	MEMs
Knowles Acoustics SPM0103ND3	1.5 - 5.5	< 100	59	omni	0.100 - 0.350	-22 @1kHz	115	1.5x4.72x3.76	$4.73	MEMs
Knowles Acoustics SPM0102NE3	1.5 - 5.5	< 100	59	omni	0.100 - 0.250	-42 @1kHz	115	1.5x4.72x3.76	$6.50	MEMs
ICC MEU 627PD-02-704	0 - 10	< 100	> 35	uni	0.5	-22 @1kHz	NA	6x2.7	<$10.00	Electret
Panasonic WM-G10BX	0 - 10	< 2.2k	> 58	omni	< 0.5	-44 ± 3	NA	4x1	$2.52	Electret
Panasonic WM-G10BT	0 - 10	< 2.2k	> 58	omni	< 0.5	-46 ± 3	NA	4x1	$3.00	Electret
Panasonic WM-E13U	0 - 10	< 2.2k	> 58	omni	< 0.5	-43 ± 4	NA	6x1.3	$3.00	Electret
Panasonic WM-55A103	0 - 10	< 680	> 60	uni	< 0.5	-47 ± 4	NA	9.7x5.0	$2.52	Electret
Panasonic WM-G15C	0 - 10	< 2.2k	> 60	omni	< 0.5	-44 ± 3	NA	4.0x1.5	$2.52	Electret

Table 4.2: Microphones comparison.

Figure C-4 shows the schematic diagram of the microphone and it's amplifier circuit. The gain of the microphone's internal non-inverting operational amplifier (op-amp) is defined in the manufacturer's datasheet [42] as:

$$G = 1 + \frac{R_1}{R_2 + R_3} \tag{4.1}$$

with internal resistors $R_1 = 22$ kΩ and $R_2 = 2.4$ kΩ and external resistor $R_3 = R_{1M} = 1.6$ kΩ, $G = 6.5$.

The high-pass-filter corner frequency is:

$$f_{HP_M} = \frac{1}{2\pi(R_2 + R_3)C_1} \tag{4.2}$$

with $C_1 = C_{4M} = 0.47$ μF, $f_{HP} = 84.65$ Hz.

The microphone's pre-amplified output is then fed into an AD8542 non-inverting micro-power op-amp (U3A) with a gain of:

$$G_M = 1 + \frac{R_{4M}}{R_{3M}} = 1 + \frac{75}{4.02} = 19.65 \tag{4.3}$$

and low-pass cutoff frequency of:

$$f_{LP_M} = \frac{1}{2\pi R_{4M} C_{2M}} = \frac{1}{(2\pi)(75 \times 10^3)(470 \times 10^{-12})} = 4515\text{Hz}. \tag{4.4}$$

Finally, the amplified signal (*Voice**) is connected to the microcontroller's analog-to-digital converter (ADC) and to the input of the band-pass filter bank described in the following section.

In telephony, the usable voice frequency band ranges from approximately 300 Hz to 3400 Hz. The bandwidth allocated for a single voice-frequency transmission channel is usually 4 kHz, including guard bands, allowing a sample rate of 8 kHz to be used as the basis of the pulse code modulation system. The voiced speech of a typical adult male will have a

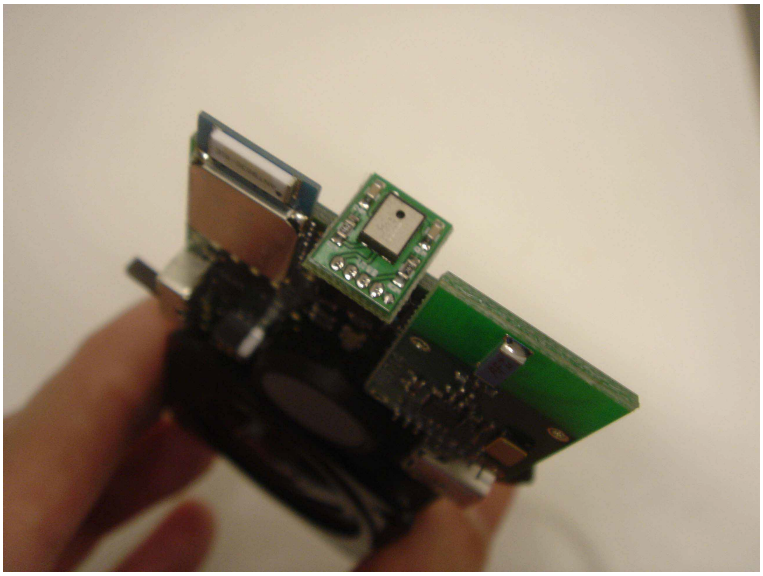

Figure 4-9: Sociometric badge's microphone daughter board.

fundamental frequency from 85 to 155 Hz, and that of a typical adult female from 165 to 255 Hz. Thus, the fundamental frequency of most speech falls below the bottom of the "voice frequency" band as defined above. However, enough of the harmonic series will be present for the missing fundamental to create the impression of hearing the fundamental tone [16].

4.8 Band-pass Filter Bank

The speech frequency spectrum (85–4000 Hz) was divided into four frequency bands: f_1 from 85 to 222 Hz, f_2 from 222 to 583 Hz, f_3 from 583 to 1527 Hz, and f_4 from 1527 to 4000 Hz in order to achieve a constant quality factor $Q = 1$ and constant gain $K = 2.58$ in each of the filters.

Four unity-Q single-op-amp Sallen-Key band-pass filters were designed following the procedure described in appendix A. By setting $C = 22$ nF, $R_A = 1$ kΩ, and the four frequency bands shown in table 4.3, the value of R and R_B for each of the filters was determined. Table 4.3 shows the low-pass (f_L), high-pass (f_H), and center (f_0) frequencies, as well as the bandwidth (BW), quality factor (Q), and RC characteristics of each of the filters.

Filter	f_L	f_H	f_0	BW	Q	RC	R
1	85	222.63	137.56	137.63	0.9995	0.001636199	74372.67
2	222.63	583.10	360.30	360.47	0.9995	0.000624706	28395.74
3	583.10	1527.21	943.67	944.12	0.9995	0.000238515	10841.59
4	1527.21	4000	2471.61	2472.79	0.9995	9.10658E-05	4139.35

Table 4.3: Band-pass filter bank parameters.

A diode-capacitor peak detector with time-constant $\tau = 1/f_0$ is used after each band-pass filter to obtain the spectral envelope in each frequency band. These four spectral envelopes are applied to A/D inputs on the microcontroller and are used to extract different speech features and segment audio into speaking and non-speaking regions. Figure C-5 shows the schematic diagram of the four band-pass filters and peak detectors.

4.9 2.4 GHz Radio Transceiver Daughterboard

A radio transceiver daughter board designed by Mathew Laibowitz in the *Responsive Environments* group at the MIT Media Laboratory was incorporated into the badge. The radio board is based on the 2.4 GHz Chipcon CC2500 transceiver, a low cost single chip 2.4 GHz transceiver designed for very low power wireless applications. It is integrated with a highly configurable baseband modem that supports various modulation formats and has a configurable data rate up to 500 kbps. It's key features are [25]:

- Frequency range: 2400-2483.5 MHz.
- High sensitivity (101 dBm at 10 kbps, 1% packet error rate).
- Programmable data rate up to 500 kbps.
- Low current consumption (13.3 mA in RX, 250 kbps, input 30 dB above sensitivity limit).
- Programmable output power up to +1 dBm.
- Very few external components. Completely on-chip frequency synthesizer, no external filters or RF switch needed.

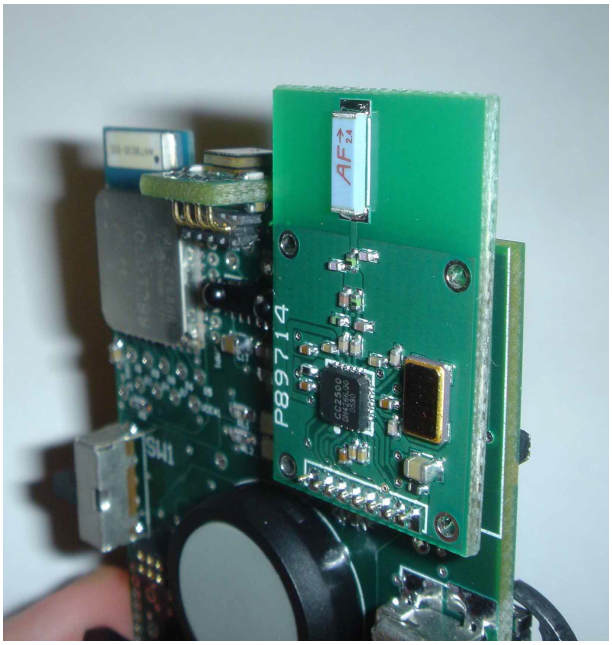

Figure 4-10: CC2500 radio daughter board mounted on the sociometric badge's main board.

- Ideal for multi-channel operation.
- Suitable for frequency hopping systems due to a fast settling frequency synthesizer.
- Separate 64-byte RX and TX data FIFOs.
- Efficient SPI interface. All registers can be programmed with one "burst" transfer.
- Digital RSSI output.
- Wake-on-radio functionality for automatic low-power RX polling.

Figure 4-10 shows the CC2500 radio board attached to the sociometric badge.

4.10 Bluetooth Module

Bluetooth technology enables short-range wireless voice and data communications. This wireless technology is especially appropriate for cable replacement and for use with portable devices in pervasive computing applications. At the baseband level, when two devices

establish a Bluetooth link, one acts in the role of *master* and the other in the role of *slave*. The master device determines the frequency hopping pattern and the phase for the hopping sequence (based upon its clock). All slaves communicating with a given master hop together in unison with the master. The master role generally is assumed by the device that initiates the communication. The master-slave relationship is necessary in Bluetooth low level communications but in general devices operate as peers. For a detailed description of the Bluetooth specification refer to [57].

Each sociometric badge has been provided with a Bluetooth module that enables it to communicate with mobile phones and other Bluetooth-enabled devices. Several commercial modules were considered and the *BlueRadios* BR-C46AR class 2 module was chosen because of its small size, low price and ease of use. The *BlueRadios* serial radio modems can be configured, commanded, and controlled through simple ASCII strings over the Bluetooth RF link or directly through the hardware serial UART. Its main features are [18]:

- Dedicated PCM voice channel for audio applications.
- UART baud rate speeds: 1200bps up to 921.6Kbps, and customizable.
- 10 meter (33 feet) range.
- Software adjustable transmitter power from short to long range.
- Low power consumption (50mA TX, 40mA RX, 1.4mA idle mode, and 30uA deep sleep).
- Small-form factor SMT radio modem.
- Operating temperature range: -40 to $+85°$C.
- Secure and robust communication link

Figure C-9 shows the schematic diagram of the Bluetooth module and its connection with the microcontroller's UART.

4.11 Infrared Transceiver

Radio frequency is not the only form of wireless communication. Infrared technology is used in devices such as notebook computers, PDAs, and electronic remote controls. Infrared wireless communication makes use of the invisible spectrum of light just beyond red in the visible spectrum. In particular, one standard method for IR communication is specified by the Infrared Data Association (IrDA) [2]. This method is commonly used with mobile phones and notebook and handheld computers. Infrared largely requires line-of-sight paths while RF can penetrate many objects.

The TFDU4300 IR transceiver by Vishay Semiconductors was selected for this application. It is a low profile (2.5 mm) IR transceiver module that contains a PIN photodiode, an IR emitter, and a low-power control IC in a single package. It operates from 2.4 V to 5.5 V and covers an extended IrDA low power range of close to 1 m, or a typical remote control range of 12 m. It has a low supply current of < 12 mA and a power shutdown mode current of $< 5\mu$A [76]. Therefore, it is ideal for battery operated applications. Figure C-8 shows the schematic diagram for this IR transceiver.

4.12 Speaker

Each sociometric badge contains a square plastic micro-miniature low profile speaker that measures 40 mm x 40 mm and weights only 9 grams (S150SLA-R5W by International Components) [37]. Its maximum rated power is 1 W and its frequency response goes from 500 to 10,000 Hz. It is driven by a low distortion audio power amplifier (SSM2211 by Analog Devices). The SSM2211 is a high performance audio amplifier that delivers 1 W rms of low distortion audio power into a bridge-connected 8 Ω speaker load. Battery life is extended by using shutdown mode, which typically reduces quiescent current drain to 100 nA [13]. The SSM2211 is a low distortion speaker amplifier that can run from a 1.7 V to 5.5 V supply. It consists of a rail-to-rail input and a differential output that can be driven within 400 mV of either supply rail while supplying output current of 350 mA.

Figure C-12 shows the schematic diagram of the speaker driven by the SSM2211. Pins 4 and 3 are the inverting and non-inverting terminals to the internal op-amp A1. Pin 2 should be connected to pin 3 for use in single-supply applications. The output of A1 appears at pin 5. A second op-amp, A2, is configured with a fixed gain of $A_V = -1$ and produces an inverted replica of pin 5 at pin 8. The SSM2211 outputs at pins 5 and 8 produce a bridged configuration output to which the speaker is connected. The SSM can be configured for gain much like a standard op-amp. The gain from the audio input to the speaker is

$$A_V = 2 \times \frac{R_{2S}}{R_{1S}} = 2 \times \frac{43}{20} = 2.15 \tag{4.5}$$

The 2 x factor comes from the fact that pin 8 has the opposite polarity from pin 5, providing twice the voltage swing to the speaker from the bridged output configuration. C3S is a coupling capacitor which creates a high-pass filter with a corner frequency of

$$f_{H_S} = \frac{1}{2\pi R_{1S} C_{3S}} = \frac{1}{2\pi(20,000)(1 \times 10^{-6})} = 8\text{Hz}. \tag{4.6}$$

C4S is a coupling capacitor that prevents DC current from flowing through the speaker. This capacitor forms a low-pass filter with R2S and the corner frequency is

$$f_{L_S} = \frac{1}{2\pi R_{2S} C_{4S}} = \frac{1}{2\pi(43,000)(150 \times 10^{-12})} = 24,675\text{Hz}. \tag{4.7}$$

4.13 3-axis Accelerometer

A 3-axis accelerometer (ADXL330 by Analog Devices) is used to measure body motion and energy expenditure. The ADXL330 is a small, thin, low power accelerometer with signal conditioned voltage outputs on a single integrated circuit. It measures acceleration with a minimum full-scale range of ±3 g [14]. Figure C-7 shows the schematic diagram and capacitor selection ($C_x = C_y = C_z = 0.22\mu$F) to set the signal bandwidth to 25 Hz in each of the three axes. Table 4.4 shows the characteristics of different accelerometers. We chose

the ADXL330 because it was the least expensive 3-axis accelerometer available at the time and its bandwidth is high enough to measure human acceleration.

Part Number	Supply Voltage Range (V)	Current per Axis (mA)	Number of Axes	Range	Sensitivity	Bandwidth (kHz)	Turn on Time (ms)	Size (mm)	Cost per Axis (US$)	Output Type
ADXL330	1.8 3.6	0.06	3	±3 g	300 mV/g	0.55	1	4x4x1.45	$1.82	Analog
ADXL320	2.4 6	0.25	2	±5 g	174 mV/g	2.5	20	4x4x1.45	$1.88	Analog
ADXL321	2.4 6	0.25	2	±18 g	57 mV/g	2.5	20	4x4x1.45	$1.88	Analog
ADXL322	2.4 6	0.25	2	±2 g	420 mV/g	2.5	20	4x4x1.45	$1.88	Analog
ADXL323	1.8 5.25	0.09	2	±3 g	300 mV/g	1.6	1	4x4x1.45	$1.75	Analog
AXDL311	2.4 5.25	0.2	2	±2 g	174 mV/g	6	4	5x5x2	$2.12	Analog
ADXL202	3 5.25	0.3	2	±2 g	12.5 %/g	6	1-5	5x5x2	$2.13	PWM
ADXL210	3 5.25	0.3	2	±10 g	4.0 %/g	6	1-5	5x5x2	$4.25	PWM
ADXL213	3 6	0.35	2	±1.2 g	30.0 %/g	2.5	20	5x5x2	$4.85	PWM
ADXL204	3 6	0.25	2	±1.7 g	620 mV/g	2.5	20	5x5x2	$6.00	Analog
ADXL203	3 6	0.35	2	±1.7 g	1000 mV/g	2.5	20	5x5x2	$6.00	Analog
ADXL103	3 6	0.7	1	±1.7 g	1000 mV/g	2.5	20	5x5x2	$7.75	Analog
SD 1210	5	7	1	±5 g	800 mV/g	0.4	N/A	9x9x3	$120.00	Analog
ST LIS2L0xAL	2.4 5.25	0.43	2	±2 g	1 V/g @ 5V	2	N/A	5x5x1.6	$7.25	Analog

Table 4.4: Accelerometers comparison.

4.14 AT91SAM7S256 ARM Thumb Microcontroller

The AT91SAM7S256 is a low pin-count Flash microcontroller based on the 32-bit ARM7TDMI RISC processor. It features 256K bytes of embedded high-speed Flash with sector lock capabilities, a security bit, and 64K bytes of SRAM [1]. It was chosen as the main processing unit for the sociometric badge because of its extensive peripheral set that includes a USB 2.0 Full Speed Device Port, USARTs, SPI, SSC, TWI and an 8-channel 10-bit ADC. The design of the *Lug* board [8] by Mark Feldmeier in the *Responsive Environments* group at the MIT Media Laboratory was used as a starting reference. Figure C-1 shows the schematic diagram of the microcontroller and its peripherals. For details about the firmware used with the sociometric badge refer to appendix B. Appendix D describes the peripherals configuration.

4.15 MicroSD Card Socket

The badge was provided with a microSD card socket that allows it to write and read data to and from a microSD memory card. Secure Digital (SD) cards are flash-based memory cards based on a nine-pin interface (clock, command, four data lines, and three power lines) designed to operate in a low voltage range. The SD card interface allows for easy integration into any design, regardless of the microprocessor used. It also offers an alternate communication protocol based on the serial peripheral interface (SPI) standard [70]. Currently available microSD cards include up to 2-GB of data storage. Figure C-10 shows the schematic diagram of the microSD card socket and its connection to the microcontroller's SPI bus. Figure 4-11(a) shows a picture of the microSD card before being inserted into the socket, and figure 4-11(b) shows a picture of the card after being inserted.

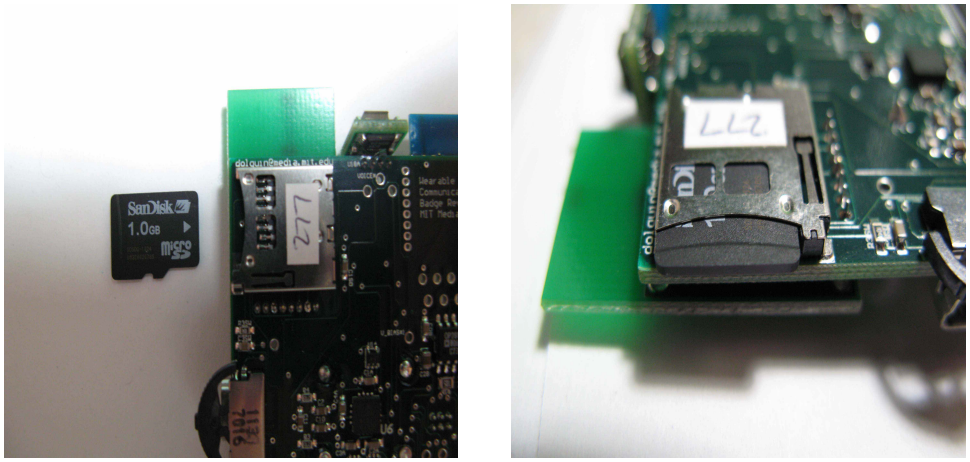

(a) MicroSD card before being inserted. (b) MicroSD card after being inserted.

Figure 4-11: Final prototype inside enclosure.

4.16 Battery

Two batteries by Ultralife Batteries were considered and tested. Table 4.5 shows the characteristics of the 930 mAh UBC001 lithium-polymer battery used with the badge. Figure 4-12 shows the dimensions of this battery, and figure 4-13 shows the charge/discharge characteristics provided by the manufacturer. Table 4.6 shows the characteristics of the 1.7 Ah UBP001 lithium-ion battery originally considered. Figure 4-14 shows the dimensions of this battery, and figure 4-15 shows the charge/discharge characteristics provided by the manufacturer. The lithium-polymer battery was preferred over the lithium-ion one due to it's significantly smaller size and lighter weight.

4.17 Current Consumption Estimation

Table 4.7 shows a current consumption estimation for the main components in the sociometric badge. The average current for each component can be calculated as

$$I_{avg} = D \times I_{max} + (1 - D) \times I_{shutdown}, \tag{4.8}$$

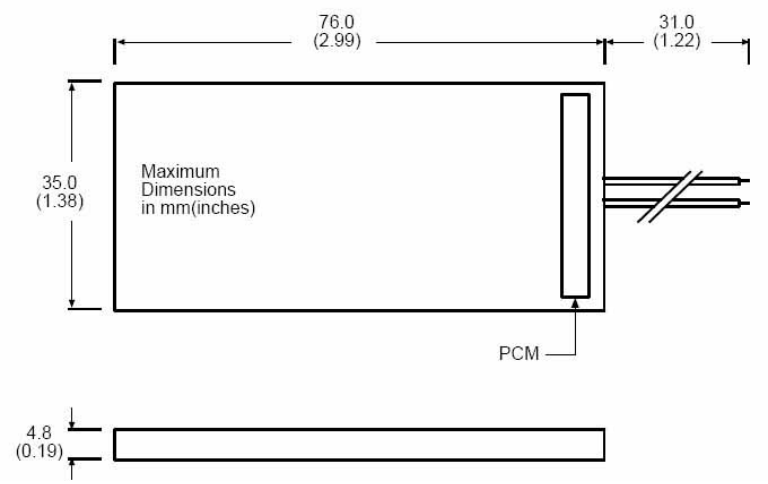

Figure 4-12: UBC001 lithium-polymer battery dimensions.

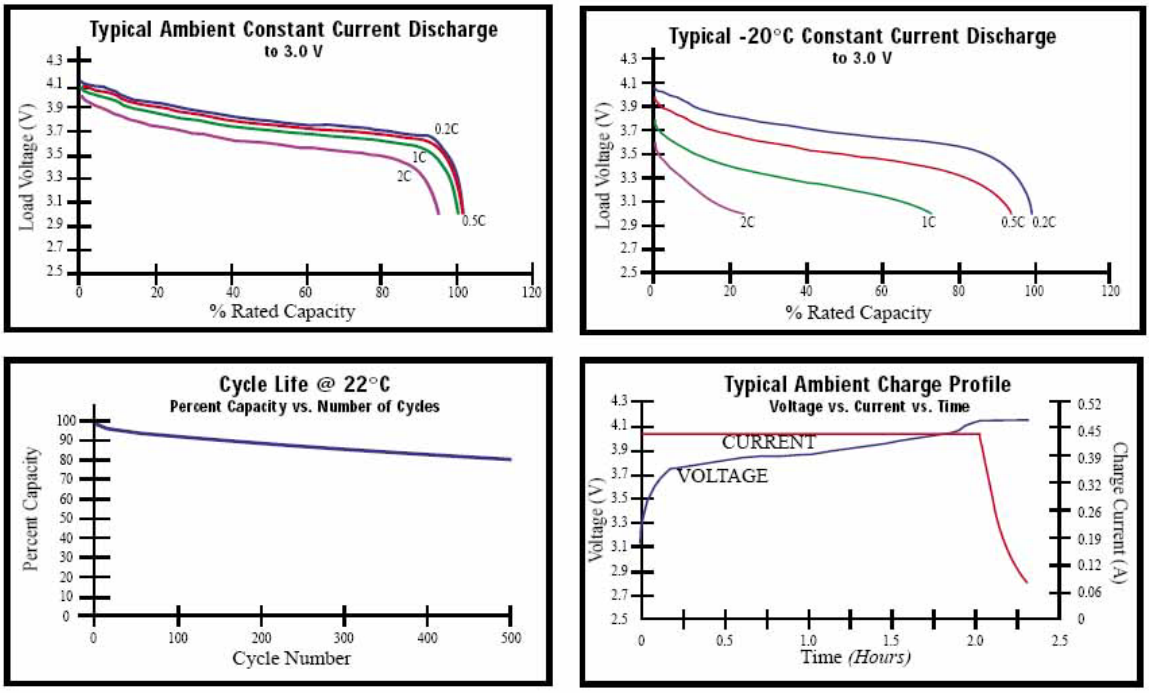

Figure 4-13: UBC001 lithium-polymer battery manufacturer charts.

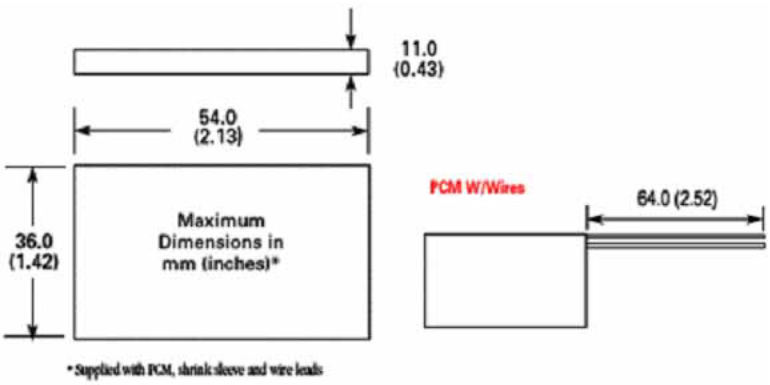

Figure 4-14: UBP001 lithium-ion battery dimensions.

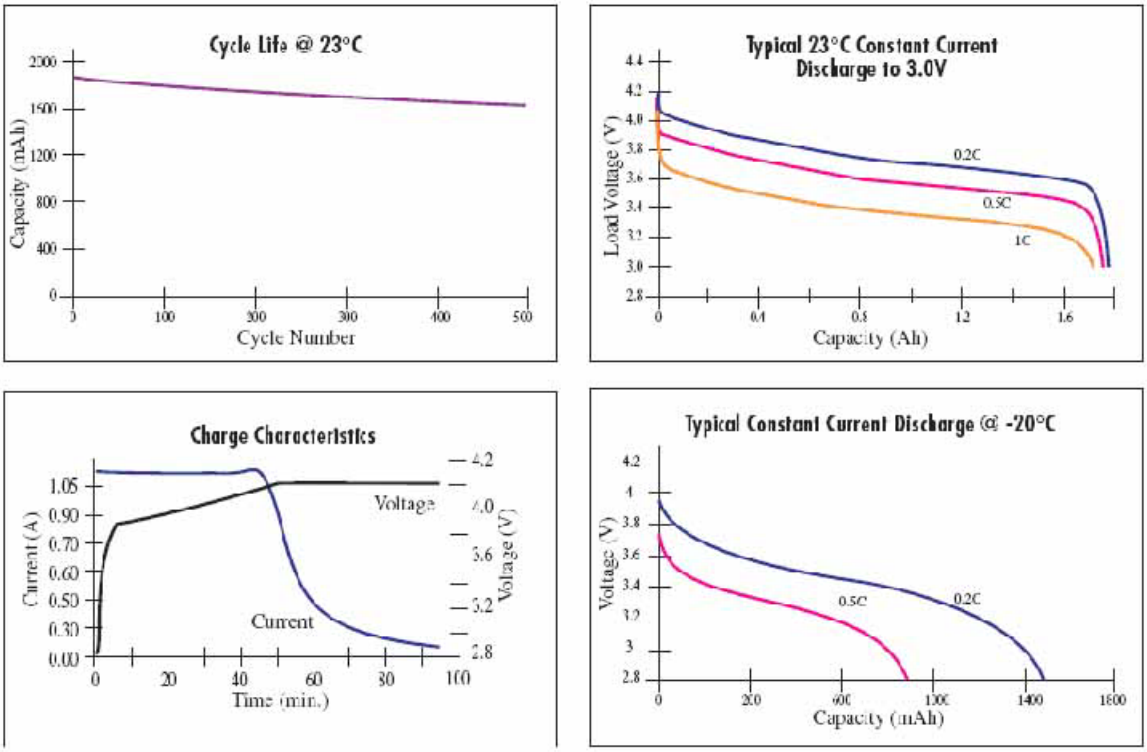

Figure 4-15: UBP001 lithium-ion battery manufacturer charts.

Attribute	Description
Voltage Range:	3 to 4.2 V
Average Voltage:	3.7 V
Nominal Capacity:	930 mAh @ C/5 Rate @ 23°C
Max. Discharge:	2 C continuous
Pulse Capacity:	Varies according to pulse characteristics, temperature, cell history and the application
Energy:	3.4 Wh
Energy Density:	156 Wh/kg, 296 Wh/I
Weight:	22 Grams
Cycle Life:	> 300 cycles @ C/5 to 80% of initial capacity
Memory:	No Memory Effect
Operating Temp.:	−20°C to 60°C
Storage Temp.:	−20°C to 60°C
Self-Discharge:	< 10% per Month
Exterior/Housing:	Laminated Foil
Terminals/Connector:	30 AWG Wire: Red (+), Black (-)
Safety:	Material Safety Datasheet - MSDS014.
Transportation:	Excepted from Regulations
Protection Circuit Module:	Over Voltage Limit: 4.285 +/- 0.025 Volts
Under Voltage Limit:	2.3 +/- 0.07 Volts
Over Current Protection:	2.0 A
Max. Quiescent Drain:	< 6μA
Charging:	Maximum charge rate at C/2 to 4.2 Volts in a temperature range of 0° to 45°C. Hold at 4.2 Volts until current declines to C/10.

Table 4.5: UBC001 lithium-polymer battery specifications.

where D is the duty cycle or percentage of time each component is active, I_{max} is the maximum current each peripheral can draw at any point in time, and $I_{shutdown}$ is the current each component draws in shutdown mode.

Considering that the Bluetooth module, the RF transceiver, and the LEDs would be active only 10% of the time (duty cycle = 0.1), and the speaker only 2% of the time (duty cycle = 0.02), the battery life using the average current would be:

$$\text{Battery life} = \frac{\text{Capacity}}{\text{Total average current}} = \frac{930\text{mAh}}{33.615\text{mA}} = 27.66 \text{ hours.} \tag{4.9}$$

Attribute	**Description**
Voltage Range:	3 to 4.2 V
Average Voltage:	3.7 V
Nominal Capacity:	1.7 Ah @ C/5 Rate @ 23°C
Max. Discharge:	C Rate continuous
Pulse Capacity:	Varies according to pulse characteristics, temperature, cell history and the application.
Energy:	6.5 Wh
Energy Density:	176 Wh/kg, 382 Wh/I
Weight:	41 Grams
Cycle Life:	$>$ 500 cycles @ C/5 to 80% of initial capacity
Memory:	No Memory Effect
Operating Temp.:	-20°C to 60°C
Storage Temp.:	-20°C to 45°C
Self-Discharge:	$<$ 10% per Month
Exterior/Housing:	PVC Shrink Wrap
Terminals/Connector:	24 AWG Wire: Red(+), Black (-)
Safety:	Material Safety Datasheet - MSDS092.
Transportation:	Excepted from Regulations
Protection Circuit Module:	Over Voltage Limit: 4.25 +/- 0.03 Volts
Under Voltage Limit:	2.3 +/- 0.03 Volts
Over Current Protection:	1.8 A - 3.0 A @ Room Temp.
Max. Quiescent Drain:	$< 100\mu$A
Charging:	Maximum charge rate at C to 4.2 Volts in a temperature range of 0° to 45°C. Hold at 4.2 Volts until current declines to C/20.

Table 4.6: UBP001 lithium-ion battery specifications.

If all components were in shutdown mode the battery life would be:

$$\text{Battery life} = \frac{\text{Capacity}}{\text{Total shutdown current}} = \frac{930\text{mAh}}{0.8335\text{mA}} = 46 \text{ days and } 11.76 \text{ hours.} \quad (4.10)$$

Device	Function	Shutdown current (mA)	Maximum current (mA)	Duty cycle	Avg current (mA)
ADXL330KCPZ	Accelerometer	0.35	0.35	1	0.35
OP281	Op-amps (1)	0.003	0.005	1	0.005
OP481	Op-amps (4)	0.012	0.02	1	0.02
AD8542	Op-amps (1)	0.08	0.09	1	0.09
AD5432	DAC	0.0004	0.005	1	0.005
MAX1793	Voltage regulator	0.016	0.2	1	0.2
SSM2211	Audio power amplifier	0.0001	9.5	1	9.5
AT91SAM7S256	Microcontroller	0.08	10	1	10
BR-C29A	Bluetooth module	0.03	50	0.1	5.027
SPM0103NE3	Microphone	0.1	0.35	1	0.35
TFDU4300	IR Transceiver	0.005	0.12	1	0.12
CC2500	RF Transceiver	0.157	21.5	0.1	2.291
MLS25070-2100-19	Speaker	0	282.84	0.02	5.657
HSMM-C170	LEDs	0	66	0.1	6.6
	Total	**0.8335**	**440.98**		**33.615**

Table 4.7: Current consumption estimation.

(a) Badge with lithium-polymer battery.

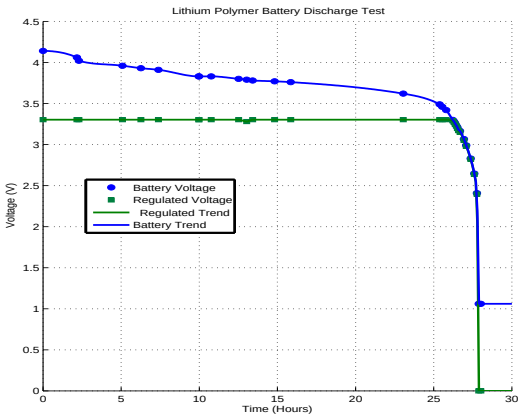

(b) Battery discharge test.

Figure 4-16(b) shows the discharge graph for the lithium-polymer battery and the regulated voltage on a badge running the firmware used in the experiment described in chapter 5.

4.18 Resources

Budget

- $77,593 USD budget provided by the MIT Media Laboratory for prototyping and manufacturing a set of wearable electronic badges for the *Sensible Organizations* research collaboration with Hitachi, Co (MIT Media Lab Sponsor).
- $1,500 USD worth of components donated by Analog Devices, Inc.
- 5 Windows Mobile Smart-phones provided by Motorola, Inc.

Costs

- The cost of raw materials and assembly for manufacturing a set of 100 badges is approximately $250 USD per badge. Appendix E contains the detailed bill of materials.
- $1,500 USD personal desktop computer used for electronic design and as a programming station.

- $300 USD CrossStudio for the ARM integrated development environment (IDE) for building, testing, and developing ARM applications.
- $195 USD CrossConnect for the ARM JTAG to USB programming cable.

4.19 Remarks

Several tests were performed to assess the functionality of the design. We present a list of design goals that were met and of possible improvements that could be made:

- It is possible to detect when the user is in a conversation using the microphone and analog band-pass filter bank, however a unidirectional microphone might be a better option if one only wants to detect the wearer's voice.
- The IR transceiver works well up to one meter and within an angle of $\pm 15^{\circ}$. Multiple IR transceivers could be added to cover a wider angle.
- The 3-axis accelerometer was useful to detect when the users were active. The use of a piezoelectric vibration sensor was also considered, however, its size was a major constraint and therefore it was not included. A gyroscope could also be included to measure the exact orientation of the badge.
- The CC2500 2.4 GHz transceiver is functional and it allows us to transfer data between badges and base stations containing a similar radio transceiver. The radio board could be replaced with a ZigBee radio such as the Chipcon CC2430 or a System on a Chip (SoC) such as the Chipcon CC2511 (these options were not available at the time the badge was designed).
- The microSD card socket was functional when storing data onboard and transferring it via USB to a computer.
- The Bluetooth module is functional and several applications using mobile phones are being developed [40].

- The speaker is the peripheral that consumes the most power and the largest component as well. It could be excluded in designs that do not require the push-to-talk functionality.

- It was decided not to include an LCD on the badge because the user would not be able to easily read the display. Based on previous experience only people facing the user are able to read the display [43]. Two LEDs indicate when the badge is powered on and when the battery is charging.

- The battery lasts approximately one day using the current power management scheme. The battery life can be extended by implementing power saving modes on the microcontroller to put other peripherals into power saving mode when they are not being used.

- Users reported feeling comfortable when wearing the badge for a period of one month (eight hours a day).

Chapter 5

Experimental Results

We deployed the research platform described in chapter 4 for a period of one month (20 working days) in the marketing division of a bank in Germany that consisted of 22 employees distributed into four teams. Each employee was instructed to wear the badge every day from the moment they arrived at work until they left the office. The employee pool has exactly the same number of men as women, but all of the managers are men. The division contains four functional teams consisting of either three or four employees. Each of these teams is overseen by a manager, who is in turn supervised by a mid-level manager. These mid-level managers are responsible for two teams, and they report directly to the division manager. The division's organizational chart is shown in figure 5-1. We treated the mid- and division-level managers as a single team in the analysis.

The bank division itself also has a very interesting physical layout. The division is split across two floors and some teams are co-located in a single room while others have employees from multiple teams in them. In fact, one of the reasons this division took such an interest in the experiment was to determine precisely what effect this physical organization had on the interactions that occur within the division.

The objective of the experiment was to use data collected using our wearable sociometric badges to correlate temporal changes in social interaction patterns (including amount of

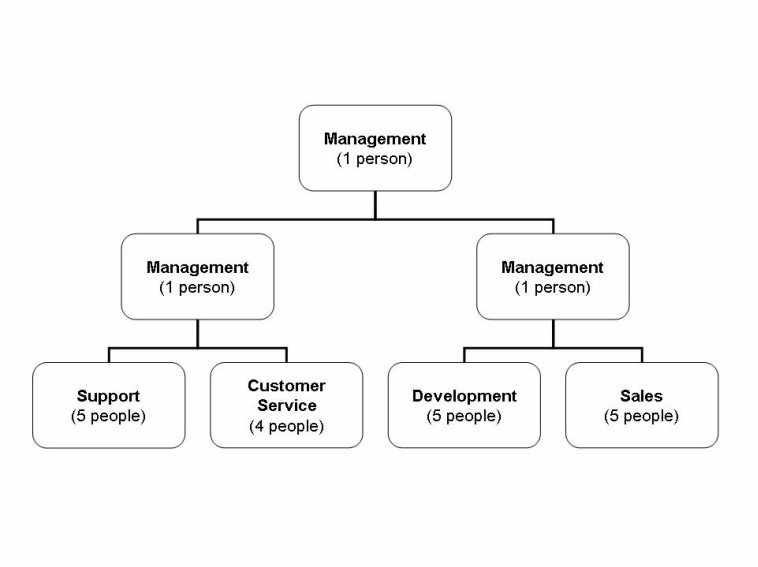

Figure 5-1: Organizational chart.

face-to-face interaction, conversational time, physical proximity to other people, and physical activity levels) with performance of individual actors and groups. We obtained e-mail logs as well as self-reported individual and group performance data as part of a case study on the impact of electronic communications on the business performance of teams [63]. This data gave us a very detailed picture of the inner operations of the division.

5.1 Experimental Procedure

Two kinds of badges were deployed in this experiment. The first were the badges worn by individuals. These badges logged IR receptions (containing the transmitting badge's ID) every time they were facing other badges, Bluetooth devices' IDs, raw and bandpass filtered audio sampled at 8 kHz and averaged over 64 samples so that the raw speech signal could not be reconstructed, and motion data using the accelerometer. Each badge transmitted its own ID via IR every two seconds and was detectable over Bluetooth every ten seconds. The badges are only able to record a limited number of Bluetooth IDs every five seconds, so not every device that is within range is logged. We also had the blue LED blink once a minute as a cue to denote that the badge was operational. The second type of badge

Figure 5-2: Base station placement.

we had were base stations deployed throughout the bank to roughly track the location of interaction events as well as subjects. Figure 5-2 shows a base station placed on a wall. Fourteen of these base stations were distributed across two floors of the bank's building and were continually discoverable over Bluetooth. A central computer was used for data collection and was placed in the division's conference room, where employees could easily retrieve their badges when they arrived and plug them into a USB hub before they left for the day. This operation allowed data to be automatically transferred via the badge's USB port and uploaded to a server in our laboratory once a day while at the same time allowed the battery in the badge to be recharged.

At the end of each day employees were asked to respond to an online survey that included the following questions:

- Q1. What was your level of productivity today?
- Q2. What was your level of job satisfaction today?
- Q3. How much work did you do today?
- Q4. What was the quality of your group interaction today?

Each question could be answered according to the following scale: (1 = very high) (2 = high) (3 = average) (4 = low) (5 = very low). Each person had to enter their badge number when they answered the survey. The following exploratory analysis on face-to-face interaction, proximity, speech, and e-mail was done in collaboration with Benjamin Waber, Taemie Kim, and Akshay Mohan in the *Human Dynamics* group at the MIT Media Laboratory.

5.2 Face-to-Face Interaction

Methodology

IR can be used as a proxy for the detection of interactions between people. In order for another badge to be detected through IR, the badges must have a direct line of sight to each other with a distance of less than one meter. Therefore the badge wearers must be standing roughly facing each other ($\pm 15^{\circ}$). This range is also a limitation, since if both wearers are turned slightly to the side but are engaged in a conversation, detection will not occur. Similarly, if there is an occlusion between the badges, say a computer monitor or a hand, detection will not occur. We can assume that if there are a number of detections between individuals in a short period of time they are indeed involved in face-to-face interaction.

Results

Analysis of IR data yielded interesting results on interaction within teams. Different rooms exhibited different communication patterns in this regard ($p < 0.05$) but the difference across teams was striking ($p < 0.01$). It appeared that there were approximately three types of teams: those with high intra-team communication (one team), those with a moderate amount (one team), and many with low communication (three teams).

Interestingly, teams that had more members tended to have more intra-team interaction, noting that these results were normalized by the number of team members. The hypothesis suggested by this finding is that individuals have to achieve a certain level of connectedness that is a function of team size, so that those on lager teams have more intra-team communication. There was a clear dichotomy between teams with higher communication and

those with lower communication. The former were teams that were project and product intensive, which naturally require intra-team collaboration, while the latter was involved in individual tasks, support roles, and control roles. These tasks intrinsically require a looser social structure since workers have little need for intra-team communication to accomplish their goals.

We were also able to identify features of face-to-face interaction that are significantly correlated with the survey data. In particular, the number of unique people that an individual interacted with over the course of the day was moderately correlated with question Q1 responses ($r = 0.19$, $p = 0.01$), while the total amount of time spent in face-to-face communication was correlated with the responses for questions Q1, Q2, and Q3 ($r = 0.27$, 0.22, 0.31, $p = 0.0008$, 0.006, 0.0001, respectively).

These findings reinforce the intuition that productivity and overall happiness is positively related to the amount of face to face interaction that one engages in. We can also explain this correlation by noting that when individuals are happy and productive, they tend to engage with a larger number of people. Conversely, when people are unhappy and unproductive they tend to withdraw. Looking at causality in the other direction, it may be that interacting with many individuals simply makes one feel happier and more productive.

5.3 Proximity and Location Using Bluetooth

Methodology

The badge can detect other Bluetooth devices in close proximity in an omni-directional fashion. In the past, this functionality has been used to identify location, behavioral patterns, and social ties [30]. We take a similar approach to using Bluetooth technology. In our experiment Bluetooth was able to determine approximate location from base stations and other mobile badges within a 10-m distance. If a person is detected in a 10-m distance, this does not necessarily mean that they interacted with each other. However we can posit that the person was in a close range, easily reachable for face-to-face interaction.

Initially we hypothesized that Bluetooth detections could be used to recognize office level locations and conversational groups. However, the large range of the Bluetooth receivers made this task extremely difficult, limiting the resolution of our data. Using the 2.4GHz radio receiver of the badge, we corrected this problem in the next version by exploiting the received signal strength from other badges.

The limitations of our current Bluetooth data have caused us to take a different approach to the analysis. Since closer devices are detected more often, we used a probabilistic method to determine what other badges were in the area. We chose to say that a badge was detected for a 5-minute time-frame only if the badge was detected for more than 30% of the time-frame, which accounts for the limited Bluetooth detection rate. We defined a person as easily reachable when a person was detected for over 30 minutes a day.

Results

The number of people easily accessible was significantly different between the people on different floors. People on the second floor had on average 4.08 people easily reachable per day while people on the third floor had 15.3 ($p < 0.01$). This same trend held for inter- and intra-team communication (both $p < 0.01$) and for inter- and intra-room communication ($p < 0.001$ and $p < 0.01$ respectively). We can posit that people on the third floor were more stationary, making them more available to others on their team and in their room. This might be a result of the type of work people are engaged in since most people on the third floor had a job description that required them to be present in their office, resulting in high Bluetooth detection rates. Our notion of accessibility, however, fails to capture the fact that mobile individuals would likely come into contact with more people than stationary individuals. Thus, it may be more appropriate to say that the people on the third floor are more reachable for planned interactions, while those on the second floor are more reachable for serendipitous interactions.

5.4 Physical Activity

Methodology

The 3-axis accelerometer signals were sampled at $f_s = 250$ Hz. This sampling frequency should be able to capture the range of human movement and could be as low as 30 Hz since 99% of the acceleration power during daily human activities is contained below 15 Hz [50]. The range of values for the accelerometer signals varies between $-3g$ and $+3g$, where $g = 9.81$ m/s^2 is the gravitational acceleration. To normalize the signals, the calibration procedure described in section 3.2 was followed.

The acceleration Signal Vector Magnitude (SVM) provides a measure of the degree of movement intensity that includes the effect of signal variations in the three axes of acceleration [38]. The SVM is calculated on the normalized i^{th} acceleration sample as follows:

$$\text{SVM}_i = |\vec{a}'_i| = \sqrt{a'^2_{x_i} + a'^2_{y_i} + a'^2_{z_i}} \tag{5.1}$$

To distinguish between periods of activity and rest the average SVM was calculated over one-minute segments:

$$\text{SVM}(k) = \frac{1}{f_s T} \sum_{i=1+f_s T(k-1)}^{f_s T k} \text{SVM}_i \tag{5.2}$$

where $T = 60$ is the time segment (in seconds) over which the average SVM is calculated, and $k = 1...K$ is the number of minutes each person was wearing the badge during the day. When the badge is static and not being worn $\text{SVM}(k) \leq 1$, since only the component of gravitational acceleration is detectable. Individual daily activity level was defined as $\text{SVM}_d = 1/K \sum_{k=1}^{K} \text{SVM}(k)$, where K is the number of minutes that the person was wearing the badge, and d is the date.

Results

We say that a person is in a *high* activity level when their activity level is one standard deviation above the mean value of everyone wearing the badge at the same time. A person is in a *low* activity level if their activity level is one standard deviation below that mean value, and they are in a *regular* activity level if their activity level is within one standard deviation of that mean value. Figure 5-3(a) shows the daily activity level graph for a person who was highly energetic throughout the duration of the experiment, and figure 5-3(b) shows the daily activity level graph for a person who was almost always below the average activity level.

Even though no significant difference was found in the activity level patterns of employees grouped by team or by floor, we found that the percentage of time a person is in a *low* activity level during the day is negatively correlated ($r = -0.202$, $p = 0.01$) with that person's individual perception of productivity (Q1). A low activity level could mean the person is sitting in their office and not moving too much or walking around. If the percentage of time a person is sitting working in their office increases, so does their individual perception of productivity.

5.5 Speech Analysis

Methodology

Objective social signaling measures based on non-linguistic vocal attributes to determine social context have been developed within our research group [67]. We take a similar approach here to characterize the interaction between individuals and determine the percentage of time that an individual is engaged in a conversation.

The audio signal was sampled at $f_s = 8000$ Hz, then it was passed through an array of four band-pass filters and peak detectors to obtain the spectral envelopes for each frequency band. These five values derived from the raw audio signal and the filter envelopes were averaged over 64 samples (8 milliseconds). The averaging ensured that one cannot determine

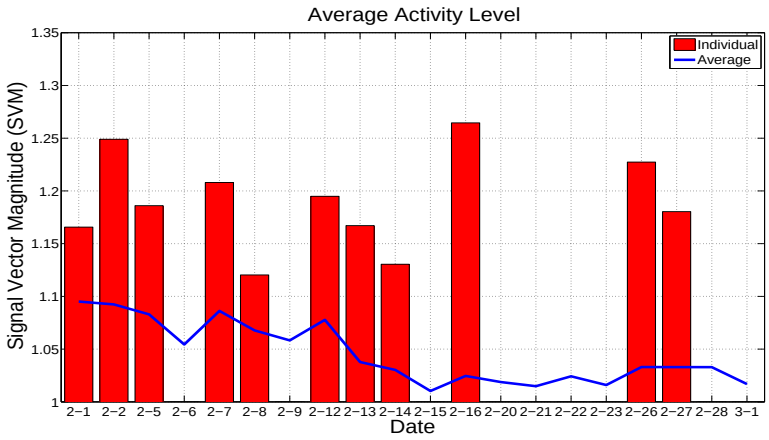

(a) Highly active person

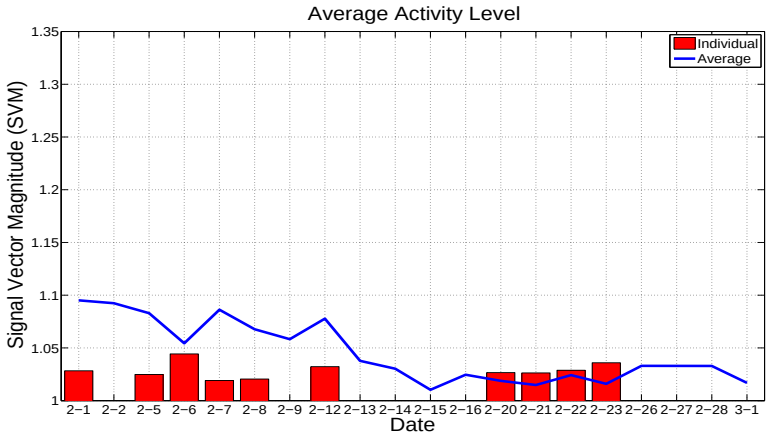

(b) Highly passive person

Figure 5-3: Comparison of daily activity levels for a highly active person and a highly passive person.

the content of the conversation or identify the speaker from the data. Offline analysis of the stored data was done to determine time segments when an individual was in a conversation. The audio was divided into frames and variation in amplitude calculated for each frame, where a frame here means 32 milliseconds of audio with an overlap of 16 milliseconds, corresponding to four and two samples of the averaged signal, respectively. This variation in amplitude for each frame was used to determine whether the individual was involved in a conversation during the given frame using a decision tree designed by Akshay Mohan.

Results

We were able to determine the percentage of total time when the individual was engaged in a conversation with respect to the total time duration for which the badge was worn. In our analysis we distinguish between conversation time and speaking time: conversation time being when any of the individuals involved in a conversation was speaking. No significant difference was found in speaking patterns of employees grouped by team or by floor. However, the daily percentage of speaking time was correlated with the number of different people with whom employees interacted every day based on IR detection ($r = 0.23$, $p = 0.004$).

5.6 E-mail

Methodology

E-mail has been frequently used to measure social ties between individuals [71]. Not only is it easy to measure, but in the modern workplace employees are interacting with each other more and more frequently through e-mail. This data is also easily quantifiable, since we know exactly who sent an e-mail to whom and when. Since it only captures digital interactions, it was unclear whether this accurately represented "real world" interactions.

In general, large scale unidirectional e-mails have little value when analyzing one-on-one interaction. Therefore we only considered reciprocated e- mails when examining relationships between individuals.

Results

In the case of e-mail communication, we found that the amount of e-mail exchanged between individuals varied by floor ($p < 0.0005$), with people on the third floor sending considerably less e-mail than individuals on the second floor. This same trend held for inter and intra-team communication ($p < 0.001$). This is a very interesting result since in contrast to proximity, these results did not hold at the room level. That is, when individuals were grouped into rooms there was no significant difference in the amount of e-mail communication. However, when looking at e-mail communication within a room, there was a large difference between rooms ($p < 0.0001$). One room had over double the average e-mail communication of all other teams.

We can view these results as implying that individuals on the second floor have tasks that require high e-mail communication, while those on the third floor require a lower level of e-mail to work effectively. Although individuals tended not to interact too often with in-room others, there was a significant exception. We could infer that this room had a more asynchronous work flow, where individuals needed to work on critical tasks in real time and exchanged trivial information only when free time was available.

When we look at how people communicate within their teams, however, we get another facet of the overall picture. Managers communicated more with their team than regular employees ($p = 0.01$), exchanging on average 7 more e-mails with members of their team. This may be because managers must coordinate the various members of the team.

5.7 General Findings

We first determined whether a relationship existed between each of the different measurements that can be automatically obtained with the badge on a daily basis (table 5.1) and each of the four survey questions described in section 5.1 to determine the degree of probability that a linear relationship existed between them. Table 5.2 shows the significant correlation coefficients that were found.

x_1	Percentage of time in each activity level
x_2	Percentage of speaking time
x_3	Number of people with whom a person interacted face-to-face
x_4	Total face-to-face interaction time
x_5	Average face-to-face interaction time per person met
x_6	Number of easily reachable people in a 10-m radius
x_7	Number of e-mails sent

Table 5.1: Automatically measured variables.

	Q1	Q2	Q3
x_1	-0.20 (0.01)		
x_3	0.19 (0.01)		
x_4	0.27 (0.0008)	0.22 (0.006)	0.31 (0.0001)
x_5	0.22 (0.005)	0.20 (0.01)	0.24 (0.002)
x_6		0.20 (0.01)	

Table 5.2: Correlation values (significance value p).

A multilinear regression was fit to model questions Q1, Q2, and Q3 using the variables that were significantly correlated with each of them. Question Q1 (perception of productivity level) was modeled as a function of x_1, x_3, x_4, and x_5, obtaining $r^2 = 0.11$ ($p = 0.001$) and regression coefficients $b_0 = 2.5384$, $b_1 = -0.0067$, $b_3 = 0.0155$, $b_4 = 0.0006$, $b_5 = -0.0004$. Question Q2 (job satisfaction) was modeled as a function of x_4, x_5, and x_6, obtaining $r^2 = 0.09$ ($p < 0.005$) and regression coefficients $b_0 = 2.4962$, $b_4 = 0.0005$, $b_5 = 0.0006$, $b_6 = 0.1535$. Finally, Question Q3 (perception of work load) was modeled as a function of x_4, and x_5, obtaining $r^2 = 0.1$ ($p = 0.0002$) and regression coefficients $b_0 = 2.5904$, $b_4 = 0.0012$, $b_5 = -0.0038$. These results imply that while we can account for a respectable percentage of the variation in responses, we must go to a finer level of analysis to find the important factors underlying this data.

Over the whole month, proximity from being in the same room, floor, or team had a high

negative correlation with the number of e-mails exchanged between people ($r = -0.55$, $p < 0.01$). This has powerful implications for previous work that has used e-mail communication as a proxy for the social network of an organization, since these two are in fact negatively related.

We can attribute this to several factors. First, if you are in close proximity to another individual, it makes more sense to interact with them in the real world rather than send them an e-mail. Second, proximity information also picks up on informal relations, while in this particular organization e-mail is used mainly for business purposes. This is because if you spend a lot of time with someone you are more likely to be their friend and therefore less likely to send an e-mail to them. This result points towards the necessity of having face-to-face interaction information in order to have a full view of the social network.

If we group people by floor, we see that people on the second floor had lower Bluetooth counts and higher e-mail counts than people on the second floor (mean = 4.08, 15.08), $p < 0.01$ and (mean = 96.4, 37.7), $p < 0.001$, respectively. Hence we can posit that people on the third floor are more stationary, staying in their office most of the time, allowing for predictable face-to-face communication and mitigating the need for e-mail. On the other hand, people on the second floor are more mobile, often out of their office, requiring them to use more asynchronous communication channels such as e-mail. This result suggests that while serendipity is important, it causes an individual to be less available for face-to-face communication, thus increasing e-mail exchanges. It would be interesting, however, to examine which way the causality actually goes. With only the e-mail count, we would have seen a continuous pattern in the individuals. However Bluetooth detection allows us to see another dimension of how people interact with each other, verifying the limits of using only e-mail as a measure of social interaction.

Supporting the high negative correlations between proximity and e-mail activity, we discovered that betweenness in the daily IR social network had a moderately negative correlation with daily e-mail activity ($r = -0.19$, $p < 0.05$). This is an intuitive finding since we would expect these socially powerful individuals to be able to call up information through face-to-face interaction rather than through asynchronous e-mail. Their central position in

the network also ensures that they will be privy to a large amount of information, further decreasing the need for e-mail communication.

We also found interesting correlations between an actor's betweenness in the daily IR social network and their daily physical activity. Average activity level had a moderate negative correlation with betweenness ($r = -0.2$, $p = 0.01$), and percent of time in a high activity level mirrored this relationship with betweenness ($r = -0.18$, $p < 0.05$). Not surprisingly, percent of time in a low activity level also showed a moderately strong correlation with betweenness ($r = 0.25$, $p < 0.005$).

These results are almost counterintuitive. They imply that those individuals who are the most critical to the network over short periods of time are in fact the least active. One explanation for this is that what makes an actor central is not that they are seeking out other actors but that other actors are physically coming to them. This would occur over the course of a day and not a month because who holds critical information may be different from one day to the next. Since central individuals are in a position of power, they would most likely be sitting or in a comfortable physical position due to their status, while others coming to them would be standing and gesturing so that the more powerful other will be encouraged to impart the desired information. We will need to examine the annotated meeting data over the month in order to gain a better understanding of this phenomenon, although it appears that causality in the other direction is unlikely.

Chapter 6

Conclusions and Future Work

We have presented the design, implementation, and deployment of a wearable computing research platform for measuring and analyzing human behavior in a variety of settings. We demonstrated the use of sociometric badges capable of measuring social signals derived from vocal features, body motion, relative location, proximity, and face-to-face interaction to capture individual and collective patterns of behavior. We deployed this system in a group of 22 employees in a real organization over a period of one month.

Using these automatic measurements, we predicted employees' self-assessment of productivity, job satisfaction, and their own perception of group interaction quality. Our exploratory analysis of face-to-face interaction time and physical proximity to other people indicated that people in different rooms and teams exhibit different patterns of communication behavior ($p < 0.01$). We found interesting results involving different communication channels and behavior patterns. An important finding was that physical proximity to other people was strongly negatively correlated with e-mail communication ($r = -0.55$, $p < 0.01$). We believe that the results from the exploratory data analysis done in this thesis need to be further validated by designing and executing several tests to prove that several aspects of human behavior can be automatically measured using sociometric badges.

6.1 Future Work

Some elements that were left for future work since they were beyond the time frame of this thesis are:

- An in-depth analysis of the band-pass filters' characteristics needs to be done in order to be able to calculate a variety of speech features (such as energy and entropy) based on the spectral envelopes.
- A robust algorithm that accurately detects face-to-face interaction and that does not only depend on the IR transceiver's line of sight must be formulated.
- The necessary algorithms to perform activity and context recognition in real time must still be implemented.
- A triangulation algorithm needs to be implemented in order to perform user localization in real time using the received signal strength indicator (RSSI) values currently captured by the badge's radio transceiver.
- A file system to read and write data in a FAT16 format to the microSD card would enable the user to transfer data without the need of a USB cable.
- Transferring voice over Bluetooth needs to be implemented in order to enable the push-to-talk functionality.
- The use of different power saving modes in the microcontroller and peripherals should increase the battery life.
- More user studies need to be done in order to determine how comfortable the badge is and if further modifications need to be made.

All sensing and processing units included in the sociometric badges will be available in smaller sizes in the near future. Even though the current form factor was accepted by most of the people participating in our studies and they didn't report feeling uncomfortable

when wearing the badge, a truly unobtrusive version should be extremely light weight and almost imperceptible to the user. The technology to build such a device already exists. The badges that were manufactured during the course of this thesis work will be used for research purposes and to further validate the measurements that can be made automatically.

6.2 Future Experiments

Team Performance. (May 2007) Each spring all first-year students in a Master of Public Policy program spend two weeks in and exhaustive study of a particular policy issue. Through readings and briefings by experts on the subject in question, each team of students develops and presents a professional analysis of the policy problem. This year, a group of 140 students participating in this exercise will wear sociometric badges. We expect to be able to identify behavioral patterns of high and low performing teams, as well as team formation processes.

Using Technology to Improve Call Centers. (May - June 2007) During a period of four weeks we will deploy our system in a call center facility to monitor a group of 30 people. The goal is to identify potential problems, user acceptance, and provide feedback to the users.

SCALE: Spontaneous Collaboration Assistant and Linking Engine. (June 2007) A company specializing in solving the problems of human performance in today's complex socio-technical systems will conduct a small pilot study of the SCALE system. They will provide a small focus group of senior executives and staff with the SCALE software and our wearable electronic badges to study physical interactions and communications. The long-term goals of this project are to integrate:

- Text mining of digital documents and communications to construct an organizational model of members' areas of knowledge/work and existing collaboration relationships,
- Data mining of physical interactions between people to augment the model of organizational knowledge and collaboration,

- Social network analysis to identify new productivity-enhancing opportunities for spontaneous collaborations between members of the organization.

This company is a leader in the design of organizations, user-centered technology and training systems that make individuals and teams more effective.

Appendix A

Band-pass Filter Design

A.1 The Band-pass Response

Following [32], the standard form of all second-order band-pass functions is $H(j\omega) = H_{0BP}H_{PB}(j\omega)$, where H_{0BP} is called the *resonance gain*, and

$$H_{BP}(j\omega) = \frac{(j\omega/\omega_0)/Q}{1-(\omega/\omega_0)^2+(j\omega/\omega_0)/Q}. \tag{A.1}$$

To express selectivity quantitatively, the concept of *bandwidth* is introduced:

$$\text{BW} = \omega_H - \omega_L, \tag{A.2}$$

where ω_L and ω_H are the -3-dB frequencies, that is, the frequencies at which the response is 3 dB below its maximum. One can prove that

$$\omega_0 = \sqrt{\omega_L\omega_H}. \tag{A.3}$$

The resonance frequency ω_0 os the *geometric mean* of ω_L and ω_H, indicating that on a logarithmic scale ω_0 appears halfway between ω_0 and ω_H. The narrower the bandwidth,

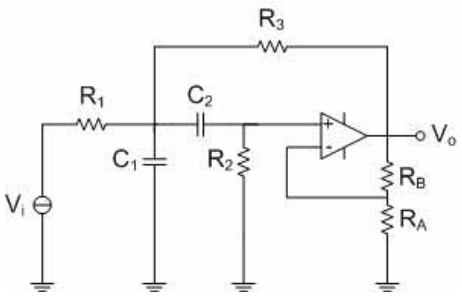

Figure A-1: Band-pass KRC filter.

the more selective the filter. However, selectivity depends also on ω_0, therefore a prosper measure of selectivity is the ratio ω_0/BW. This value is called the *quality factor* of a filter:

$$Q = \frac{\omega_0}{\mathrm{BW}}. \tag{A.4}$$

A.2 Band-pass KRC Filters

The circuit of figure A-1 consists of an R-C stage followed by a C-R stage to synthesize a band-pass block, and a gain block to provide positive feedback via R_3. This feedback is designed to bolster the response near $\omega/\omega_0 = 1$.

The transfer function of the filter is given by $V_o/V_i = H_{0BP}H_{BP}$, where H_{BP} is given in equation A.1, and

$$\omega_0 = \frac{\sqrt{1 + R_1/R_3}}{\sqrt{R_1C_1R_2C_2}}, \tag{A.5}$$

$$H_{0BP} = \frac{K}{1 + (1-K)R_1/R_3 + (1 + C_1/C_2)R_1/R_2}, \tag{A.6}$$

$$Q = \frac{\sqrt{1 + R_1/R_3}}{[1 + (1-K)R_1/R_3]\sqrt{R_2C_2/R_1C_1} + \sqrt{R_1C_2/R_2C_1} + \sqrt{R_1C_1/R_2C_2}}. \tag{A.7}$$

One can vary R_1 to tune ω_0 and R_B to adjust Q. When $Q > \sqrt{2}/3$, a convenient choice is $R_1 = R_2 = R_3 = R$ and $C_1 = C_2 = C$, in which case the filter expressions reduce to:

$$\omega_0 = \frac{\sqrt{2}}{RC}, \tag{A.8}$$

$$H_{0BP} = \frac{K}{4-K}, \tag{A.9}$$

$$Q = \frac{\sqrt{2}}{4-K}. \tag{A.10}$$

The corresponding design equations are:

$$RC = \sqrt{2}/\omega_0, \tag{A.11}$$

$$K = 4 - \frac{\sqrt{2}}{Q}, \tag{A.12}$$

$$R_B = (K-1)R_A. \tag{A.13}$$

Appendix B

Microcontroller's Firmware

Several people contributed to the development of the AT91SAM7S microcontroller's firmware. Joshua Lifton from the MIT Media Laboratory *Responsive Environments* group developed the base code for using the microcontroller and its peripherals. Akshay Mohan, Benjamin Waber, Taemie Kim, and Christopher Moh from the *Human Dynamics* group created the necessary firmware to program the sociometric badge with the main data collection application used in the experiment described in chapter 5. This firmware can be downloaded from the MIT Media Laboratory's CVS repository:

http://mlforge.media.mit.edu/projects/communicator/

Appendix C

Schematic Diagrams and PCB Layout

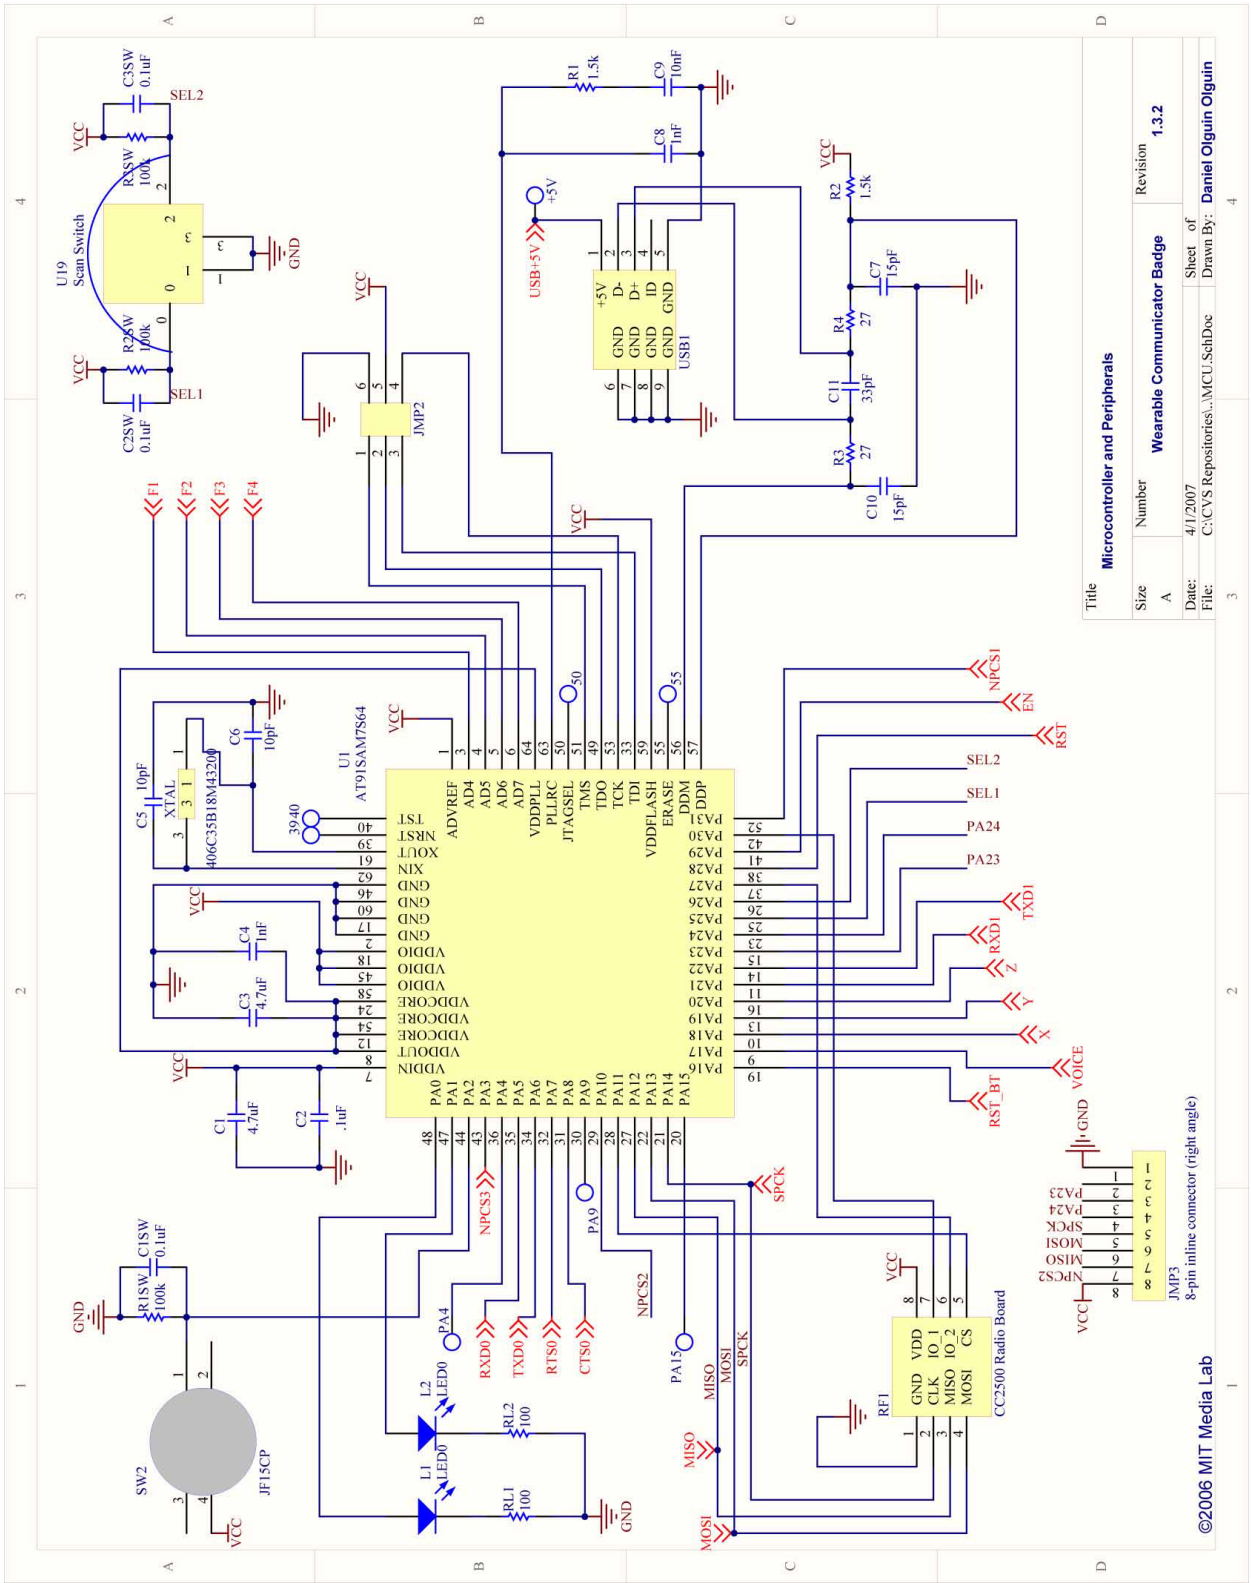

Figure C-1: Microcontroller and peripherals.

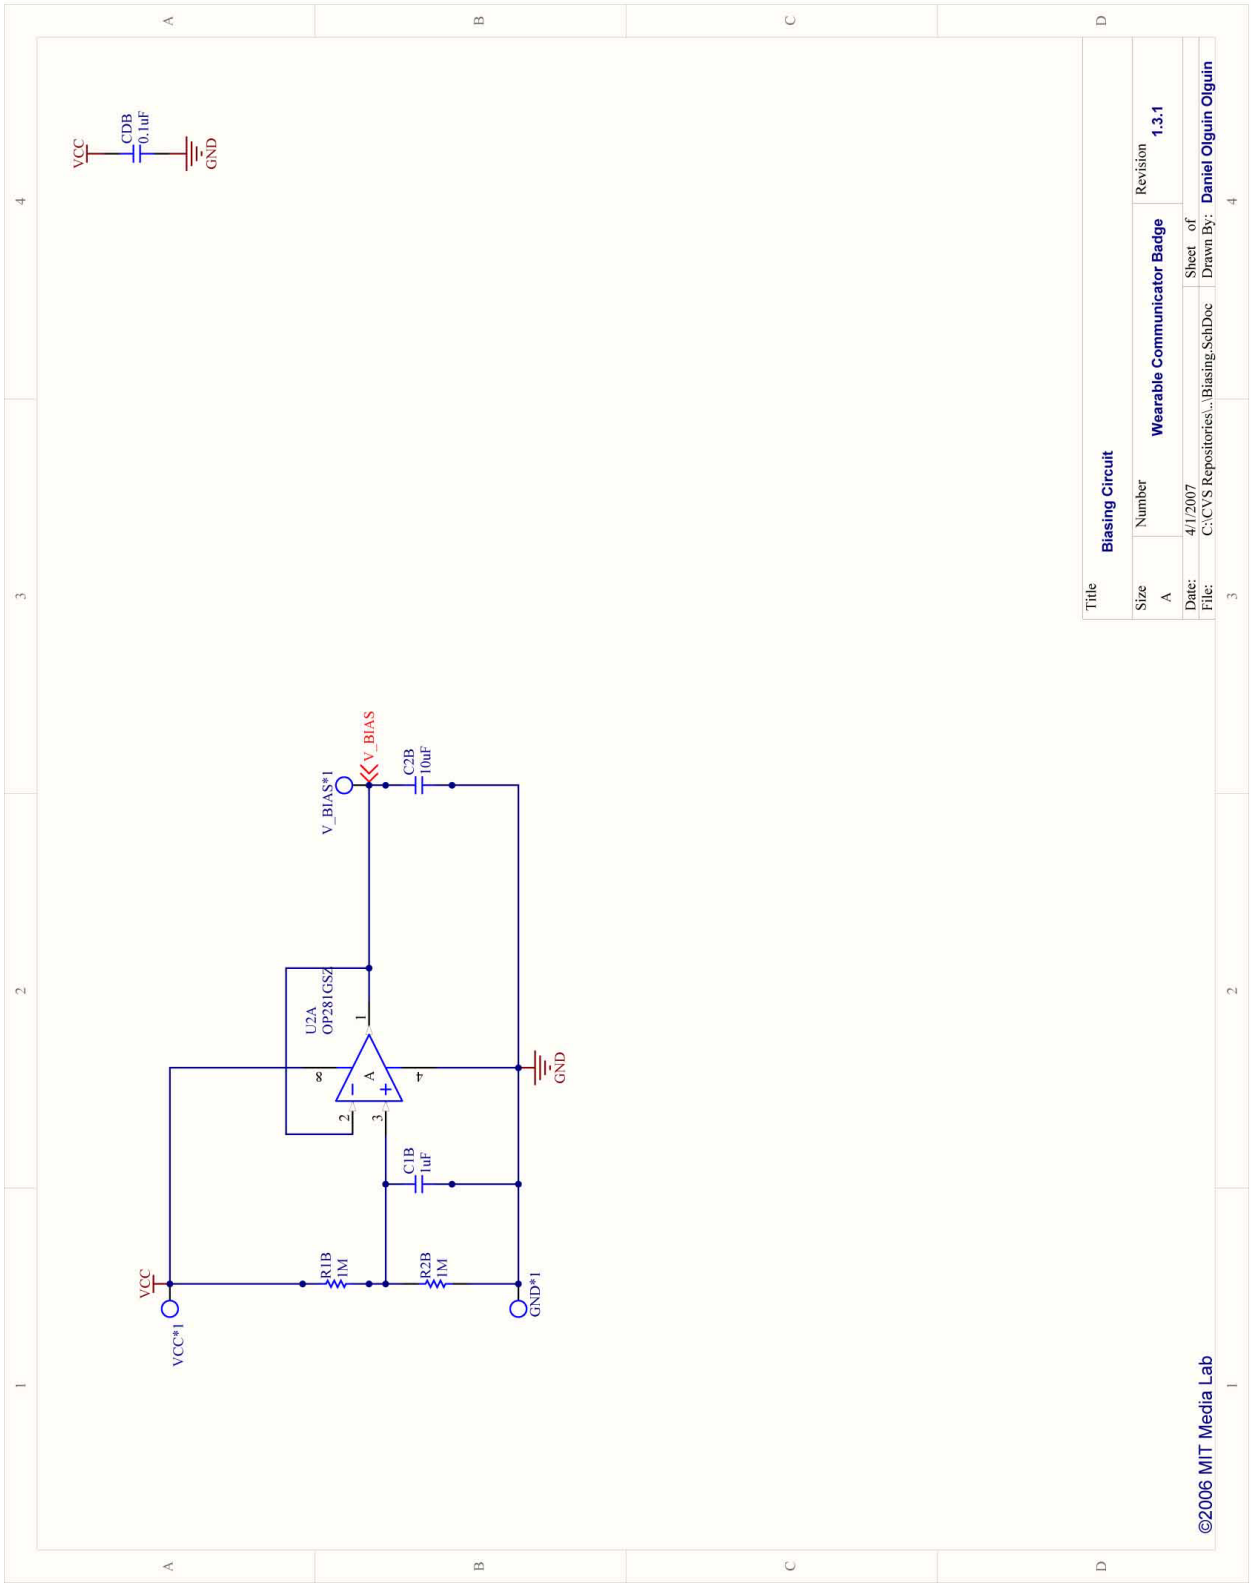

Figure C-2: Biasing circuit.

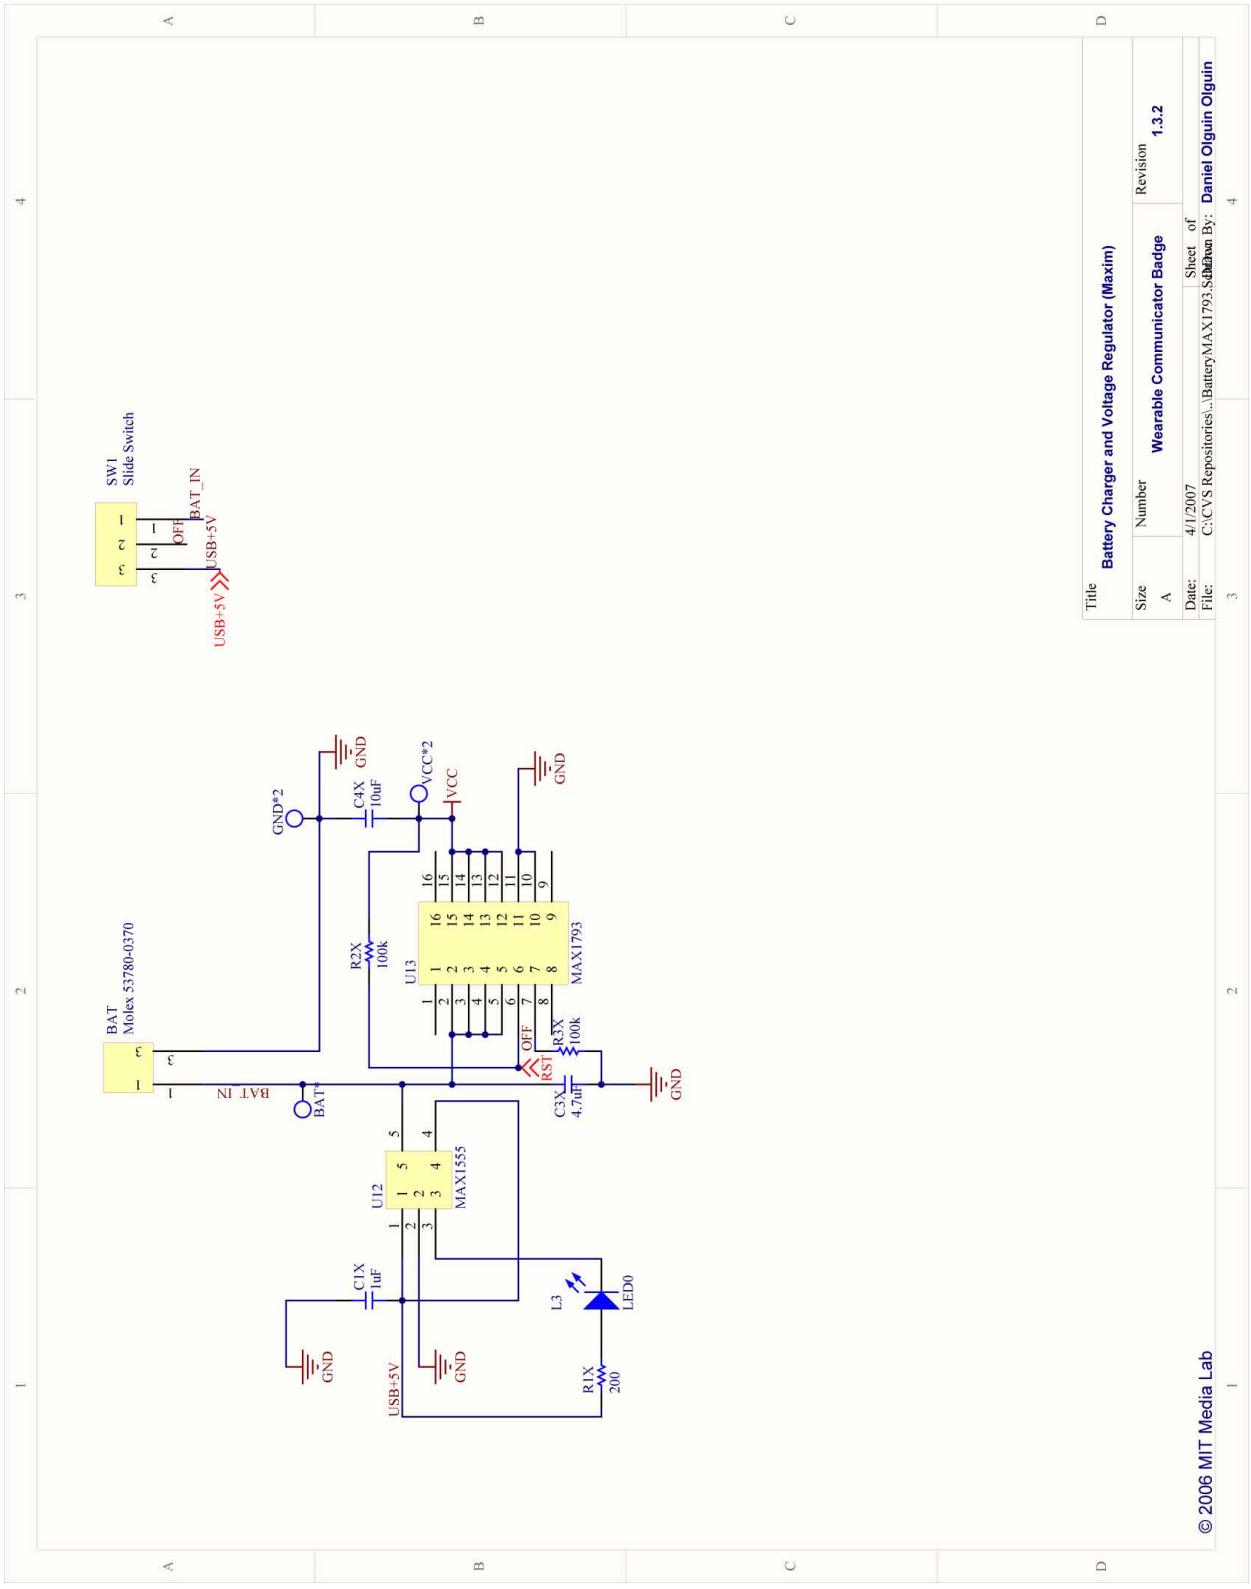

Figure C-3: Voltage regulator and battery charger.

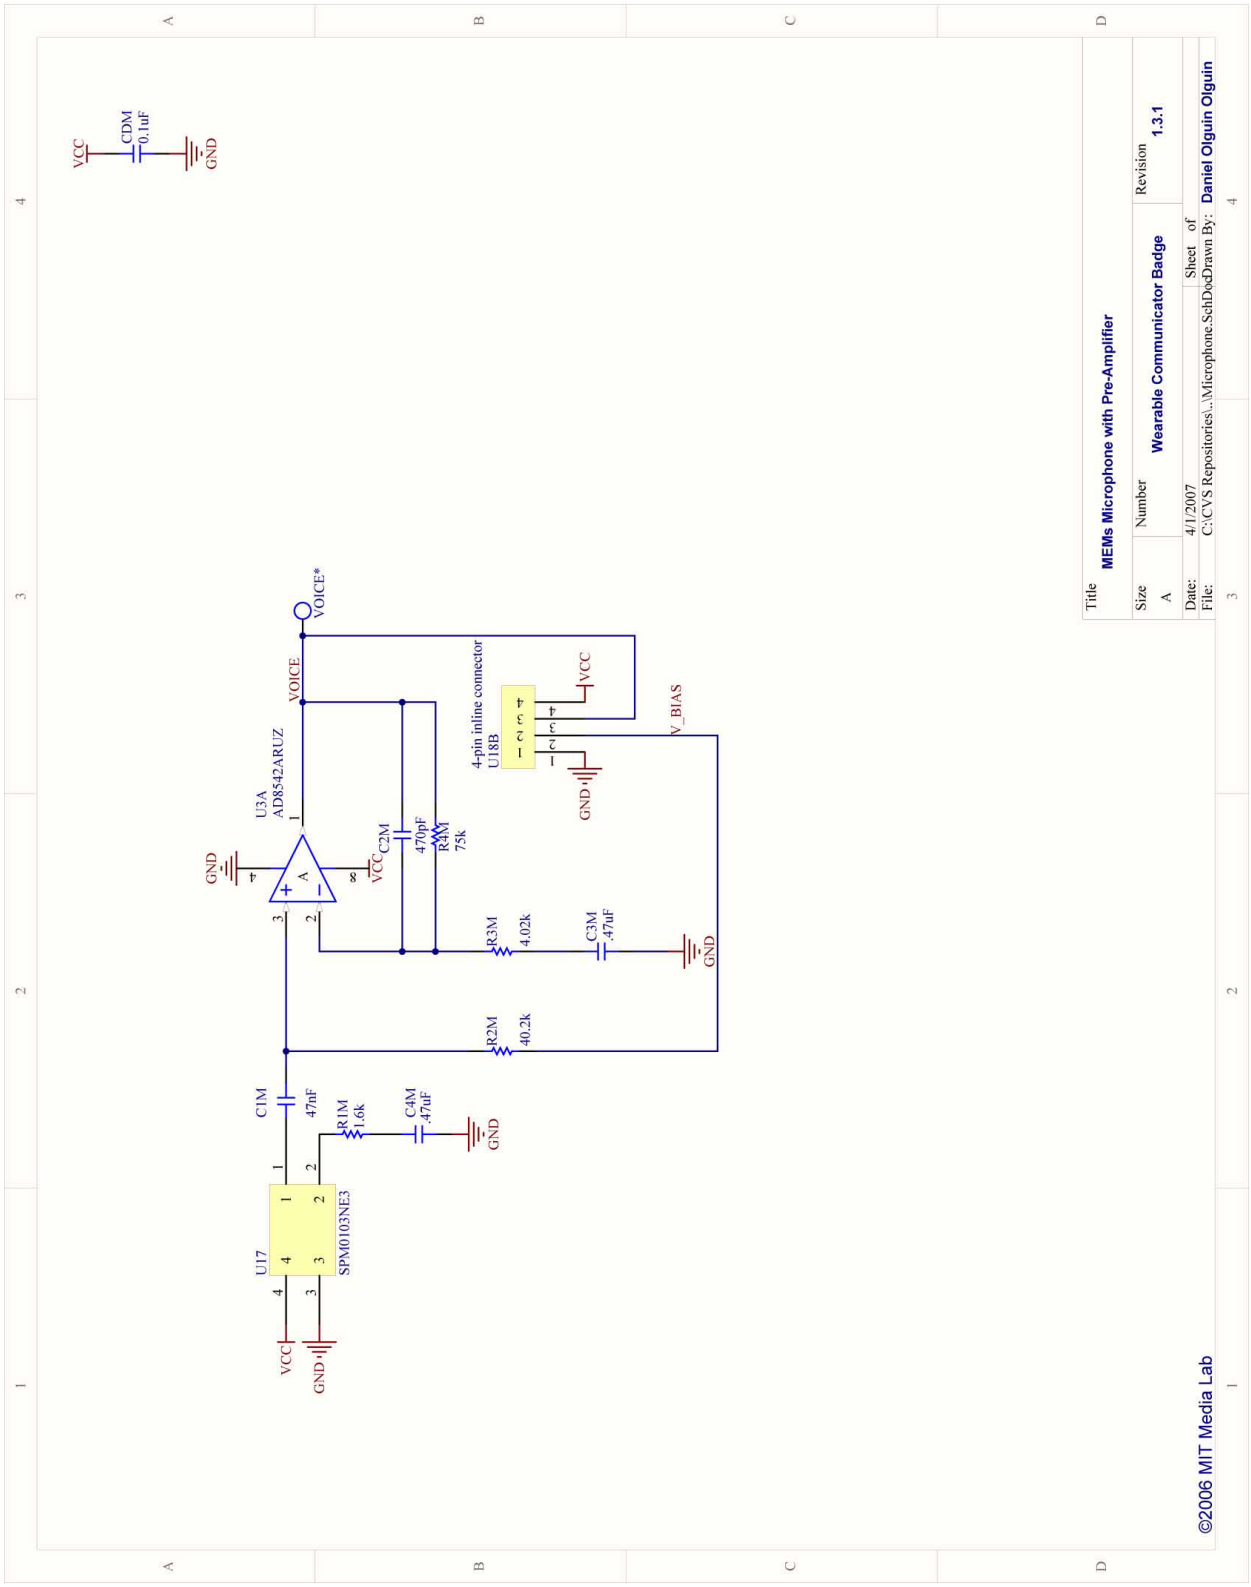

Figure C-4: Microphone with amplifier.

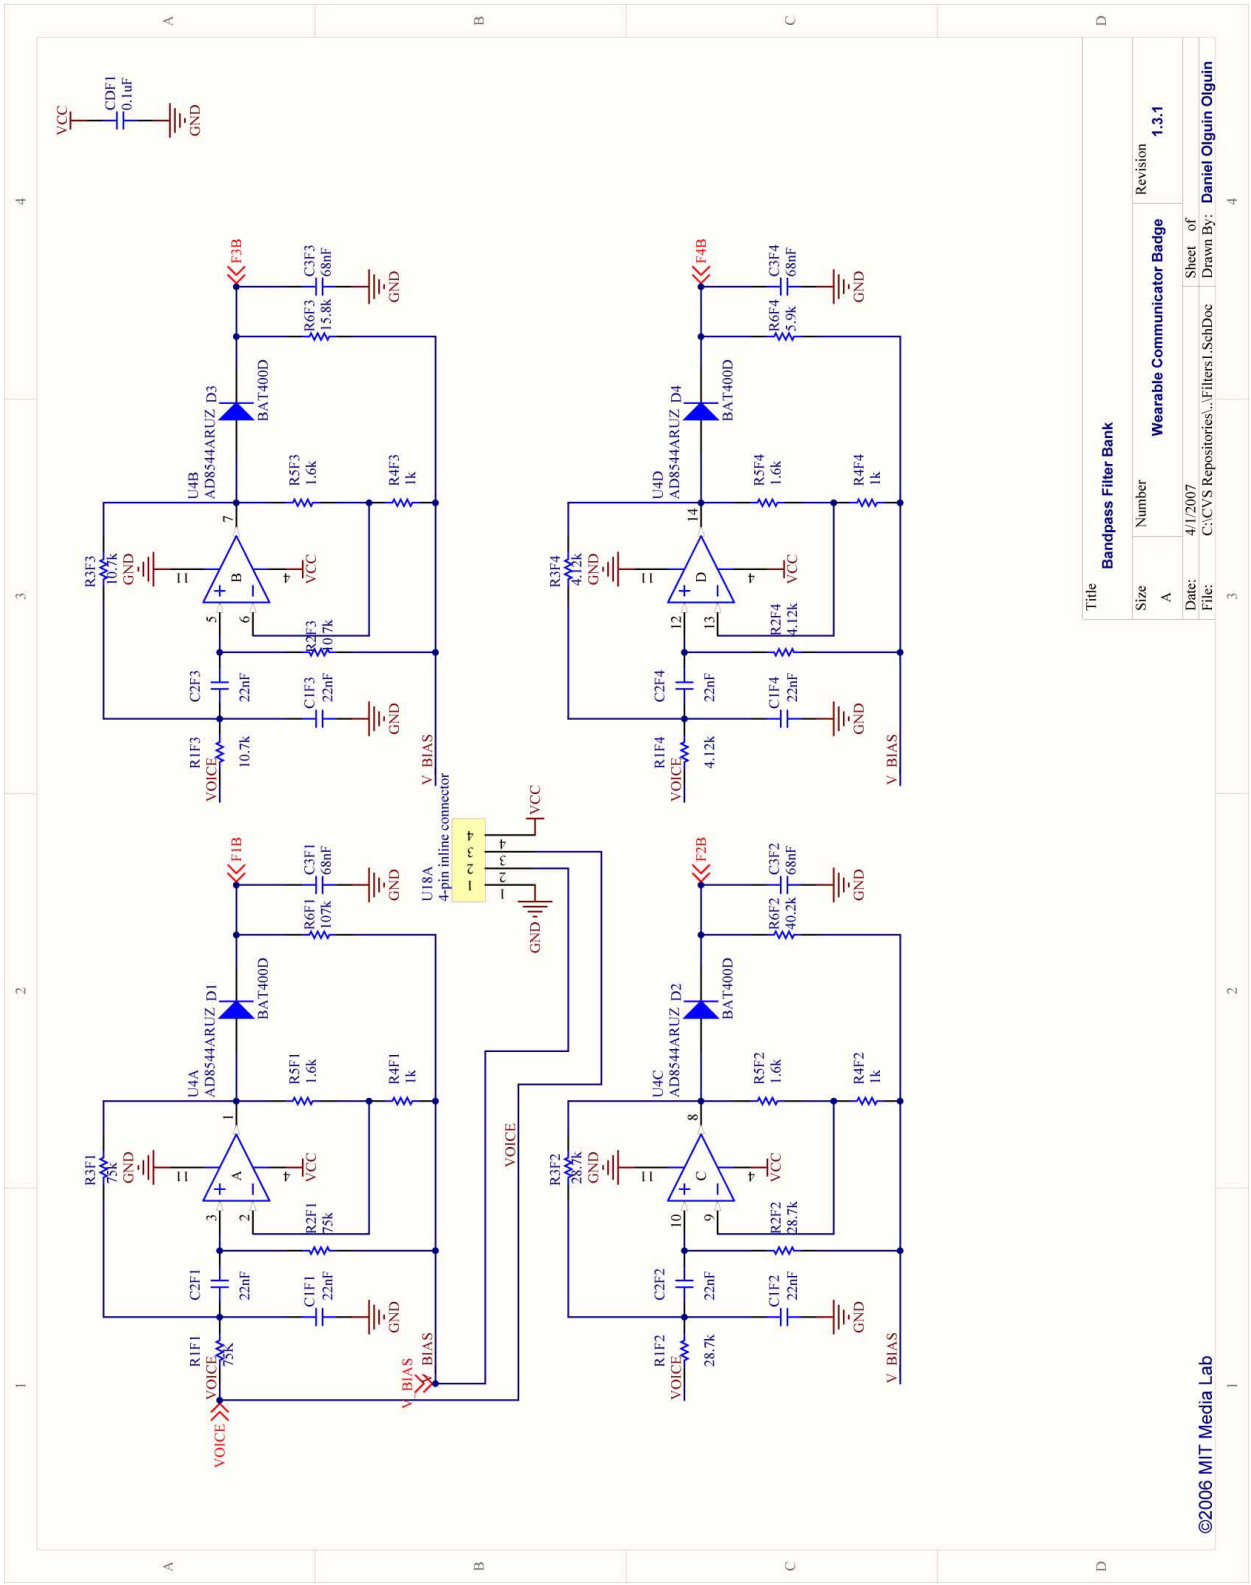

Figure C-5: Band-pass filter bank and envelope detectors.

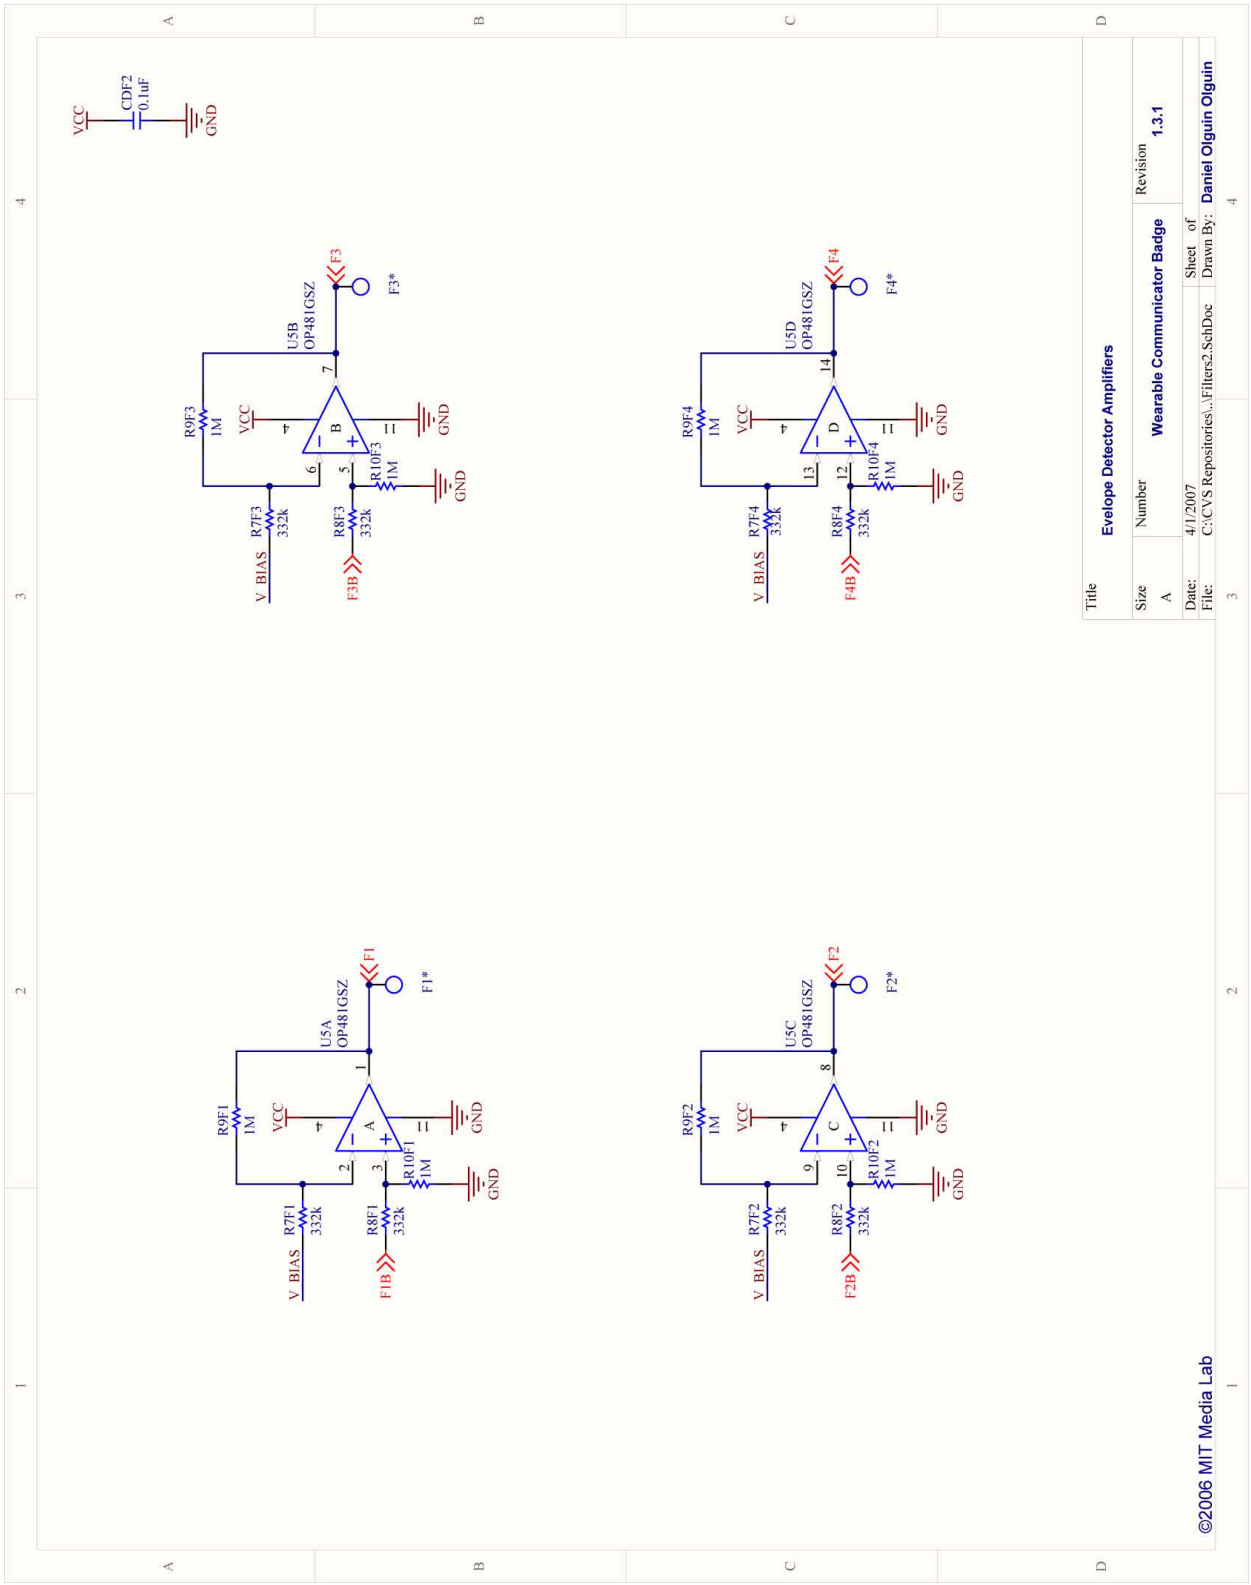

Figure C-6: Envelope detectors' amplifiers.

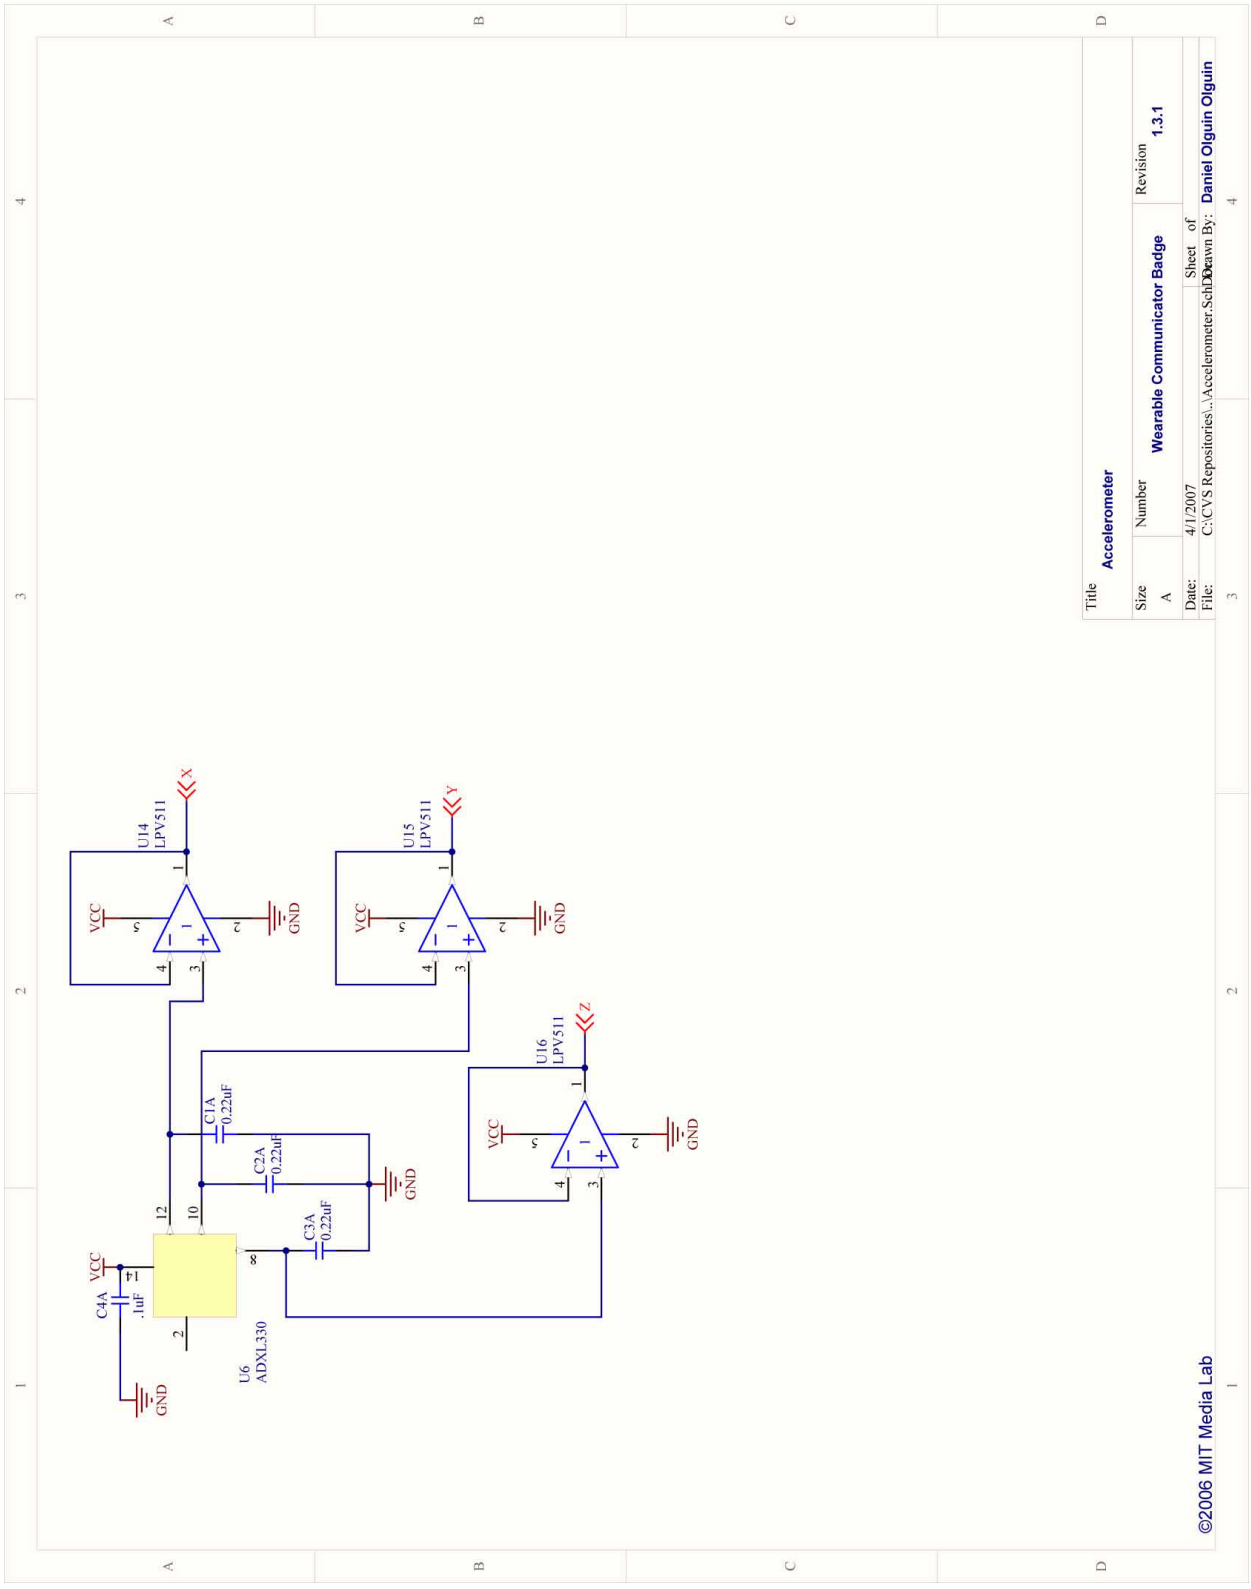

Figure C-7: Accelerometer.

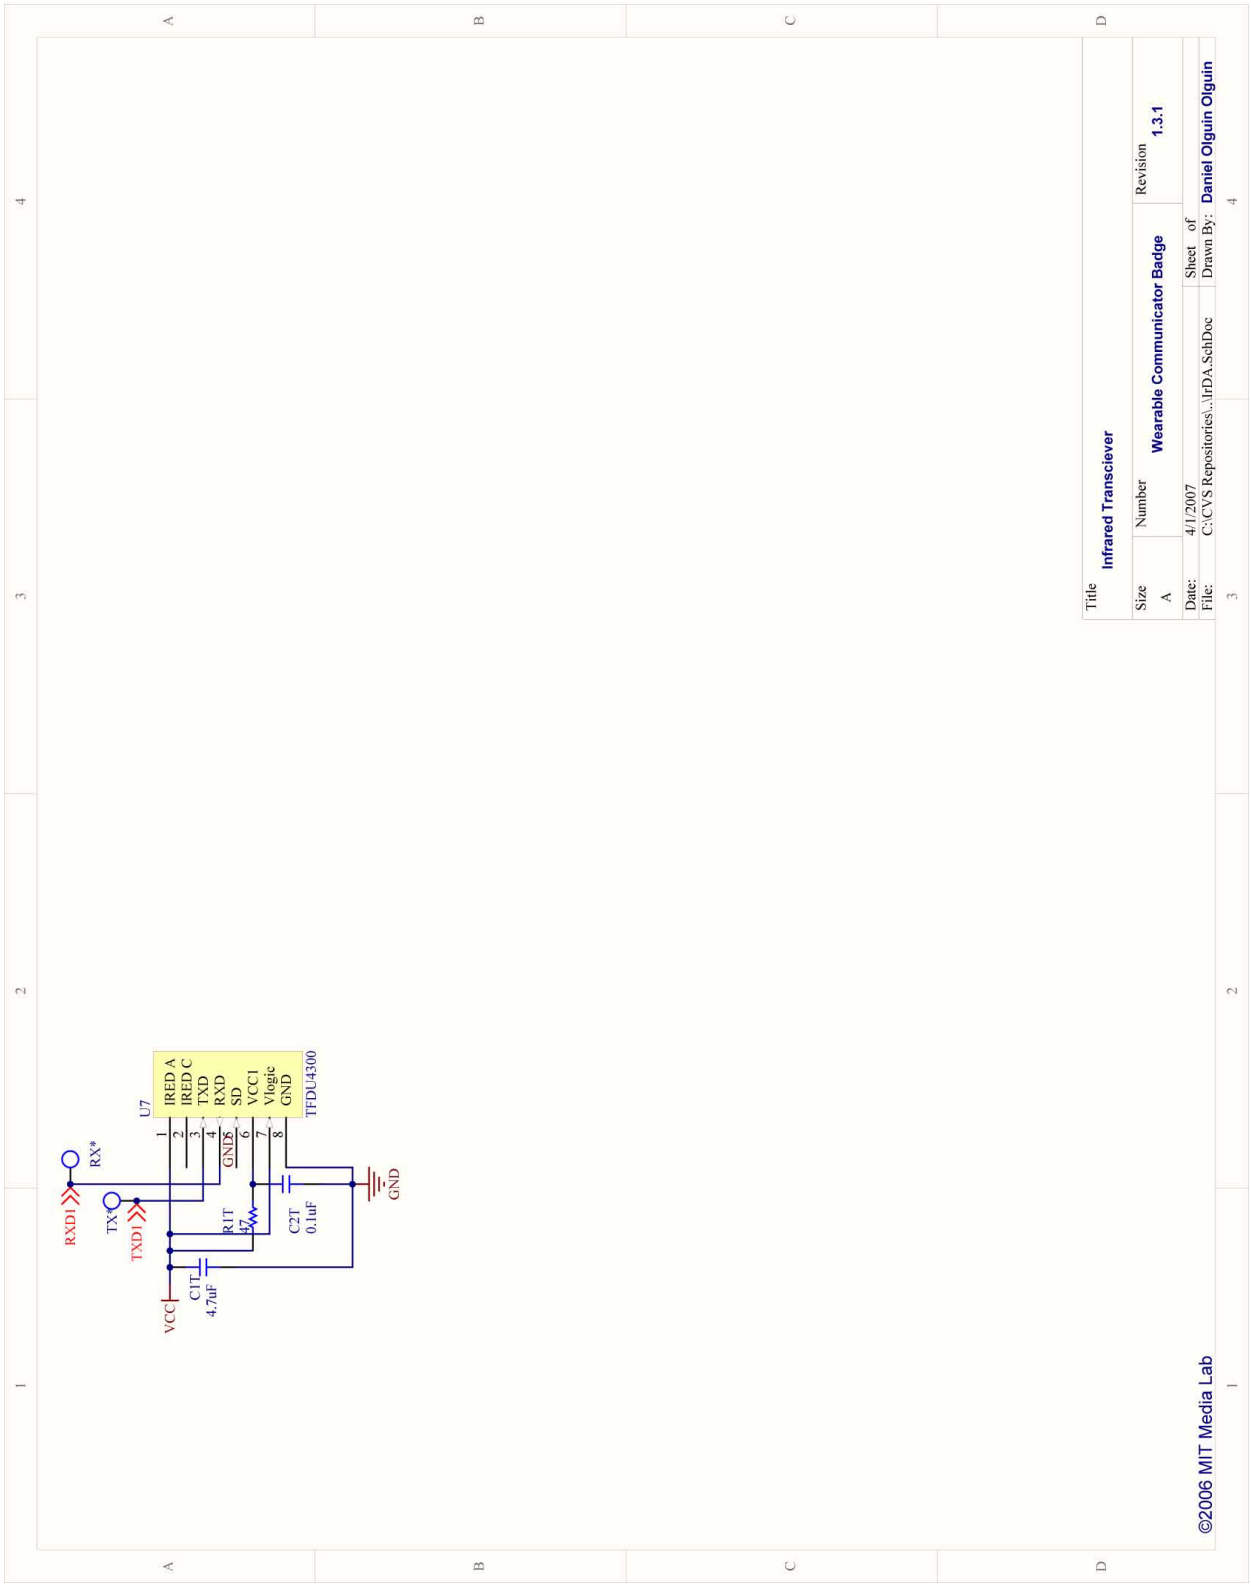

Figure C-8: IrDA transceiver.

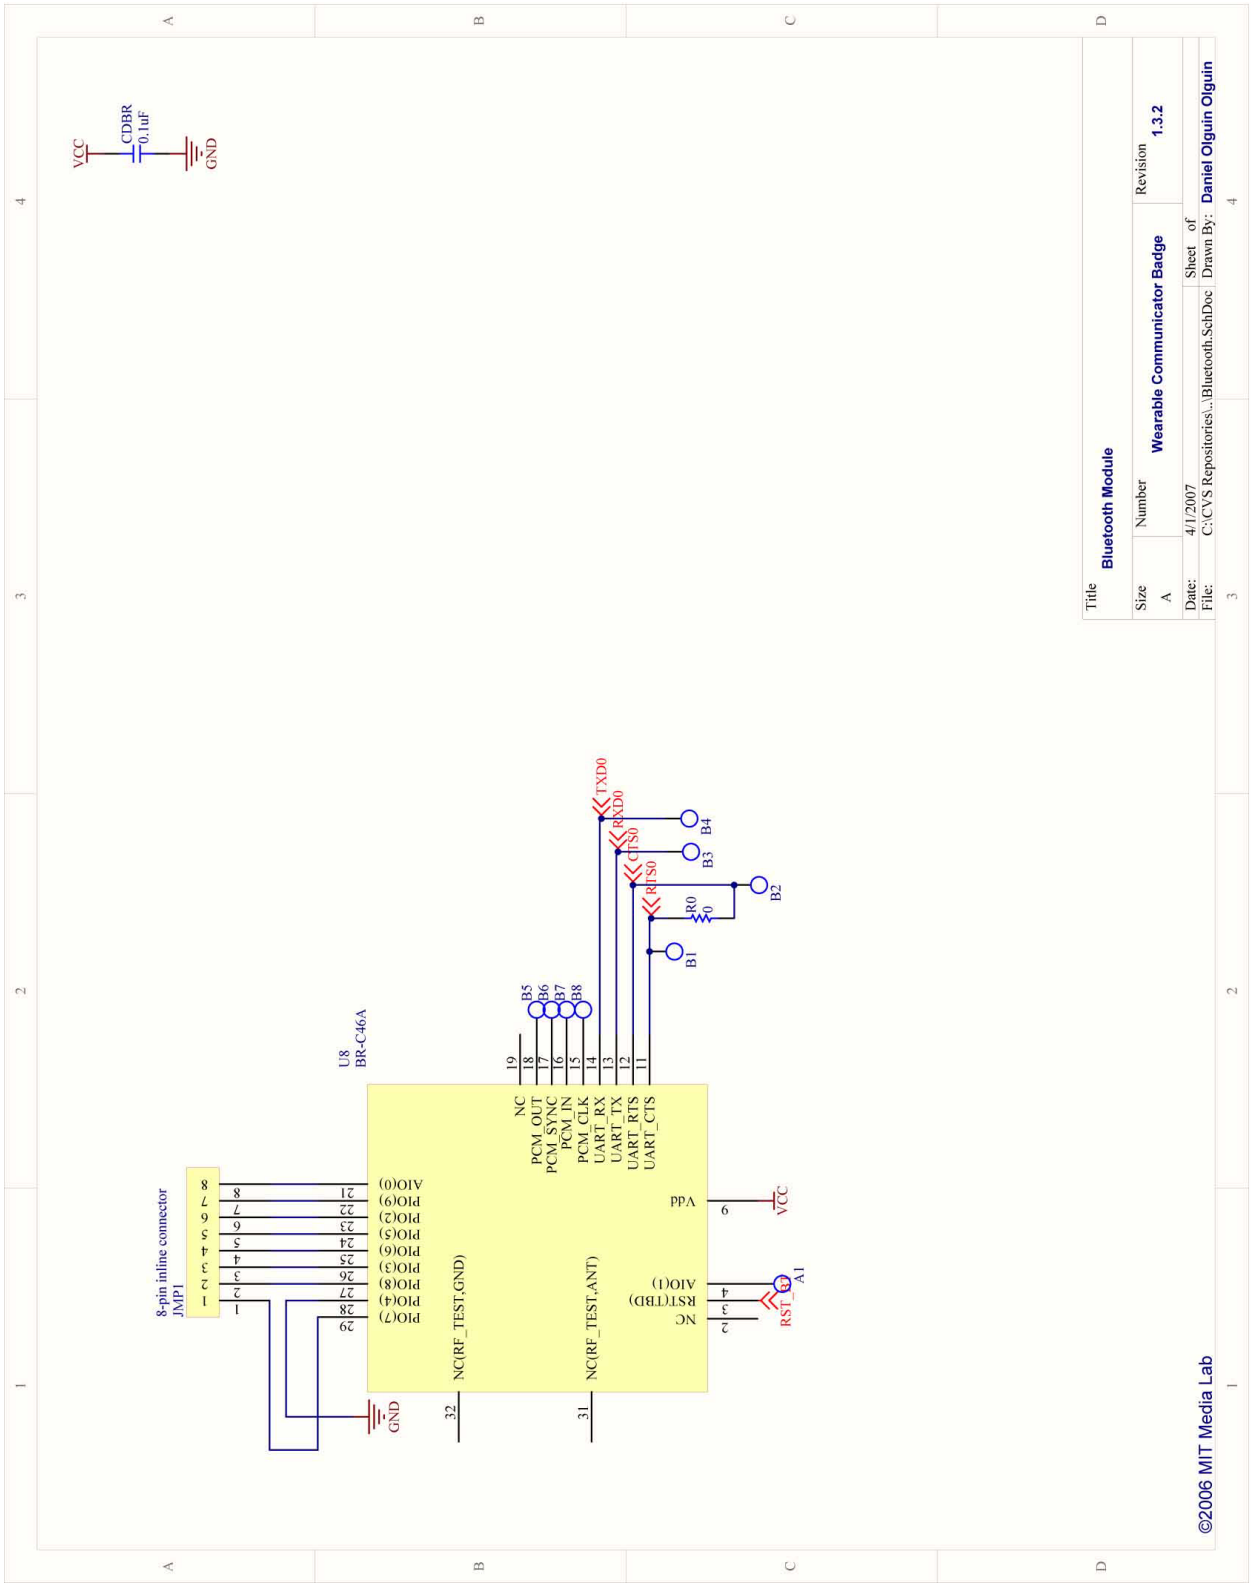

Figure C-9: Bluetooth module.

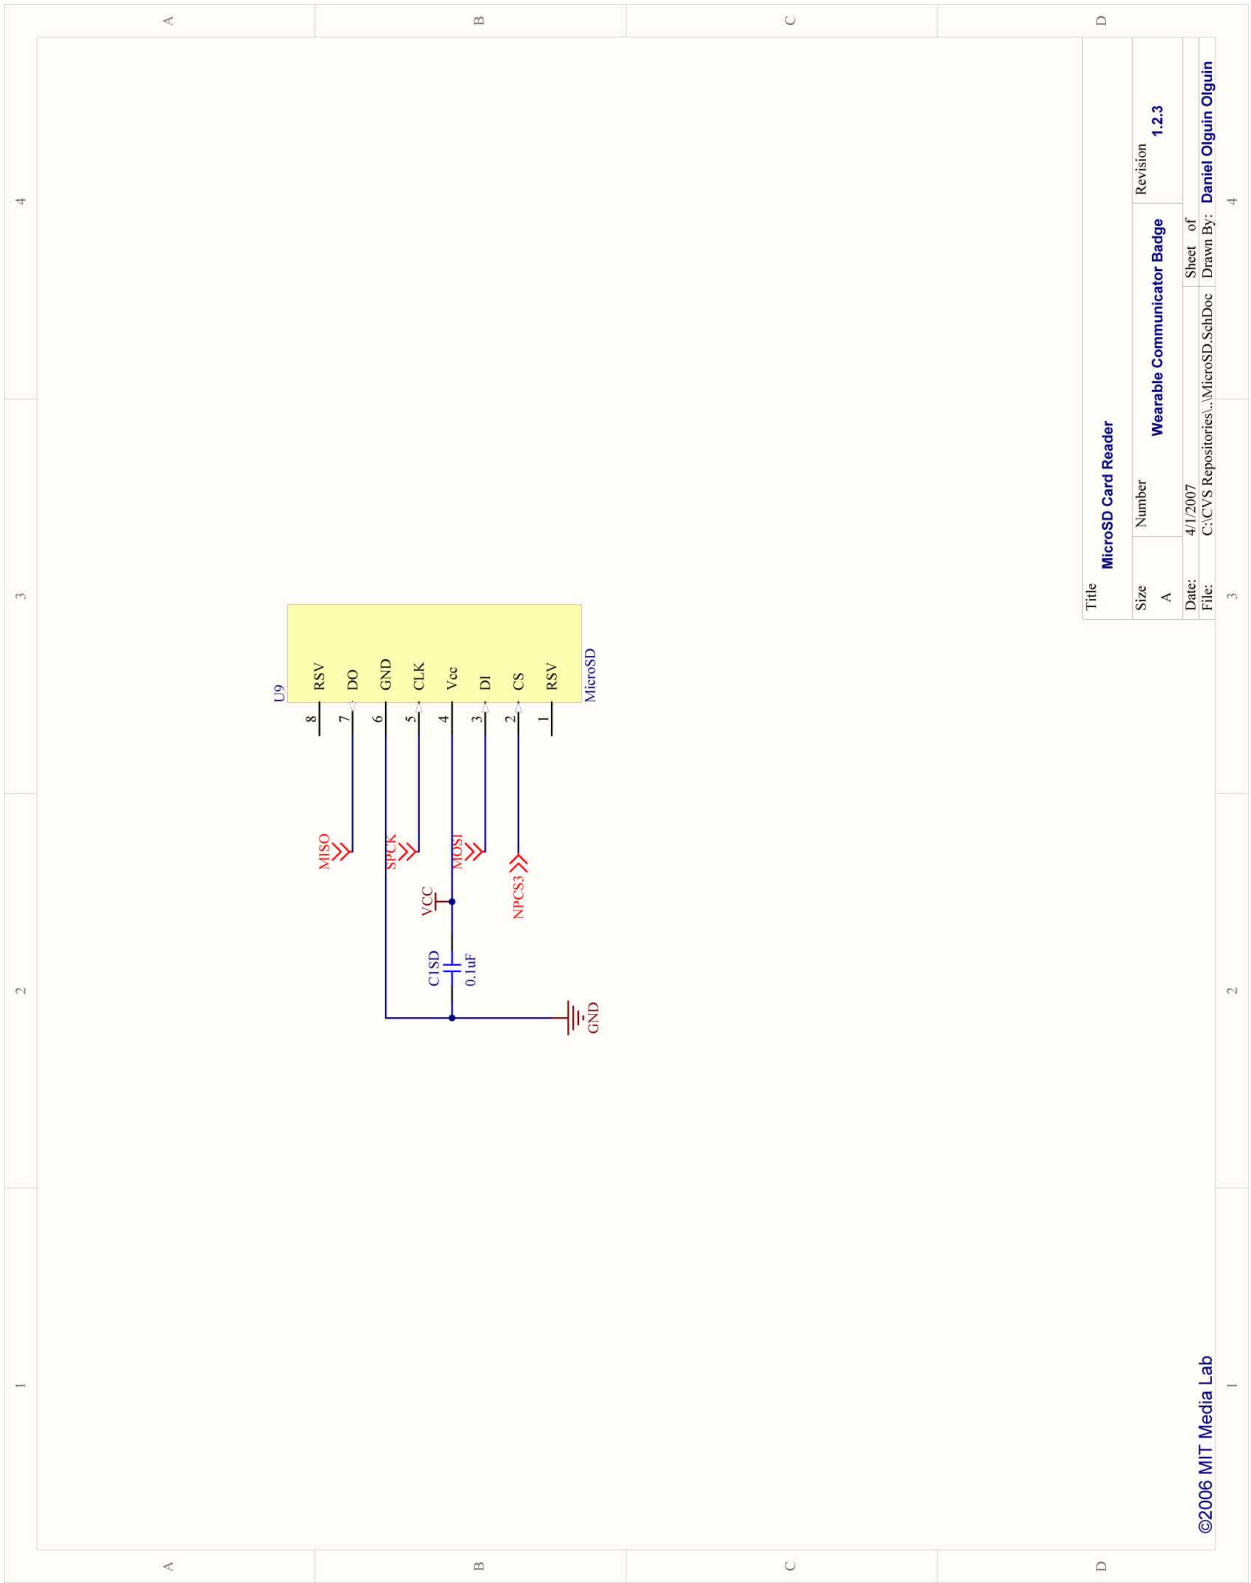

Figure C-10: microSD card socket.

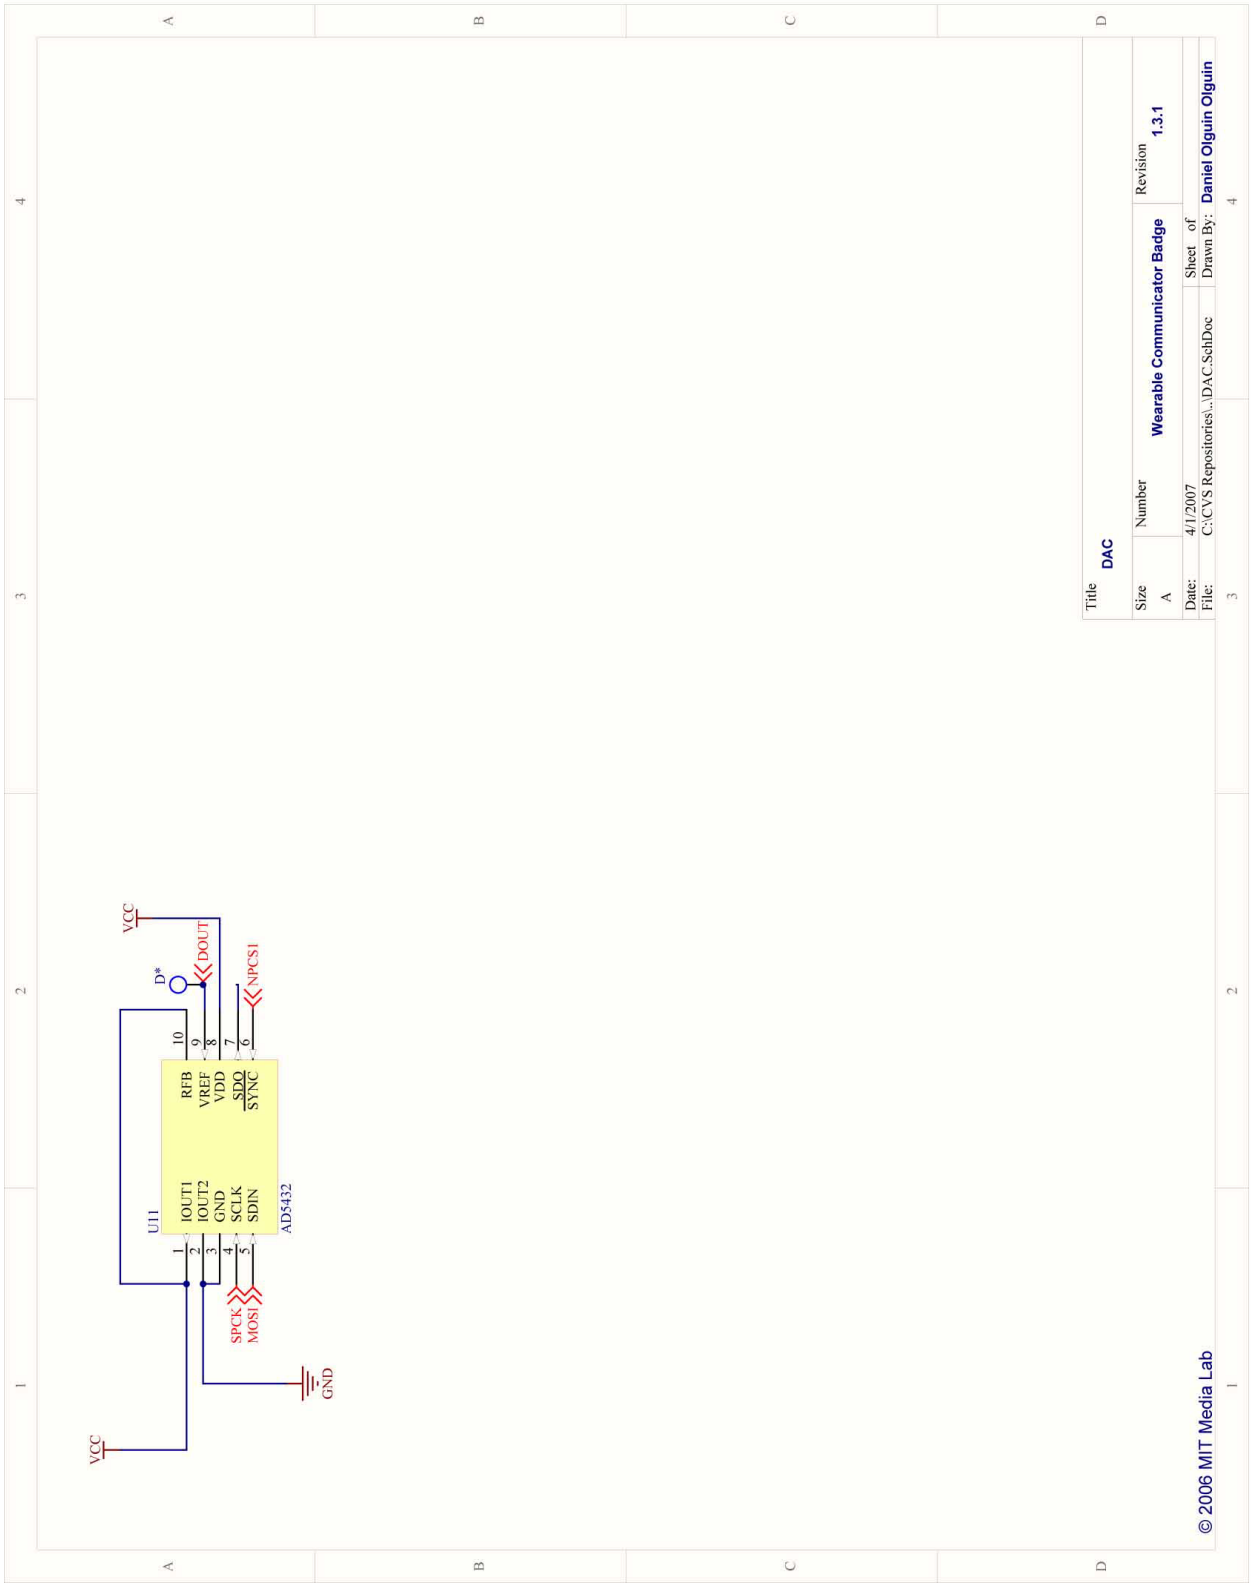

Figure C-11: Digital-to-analog converter.

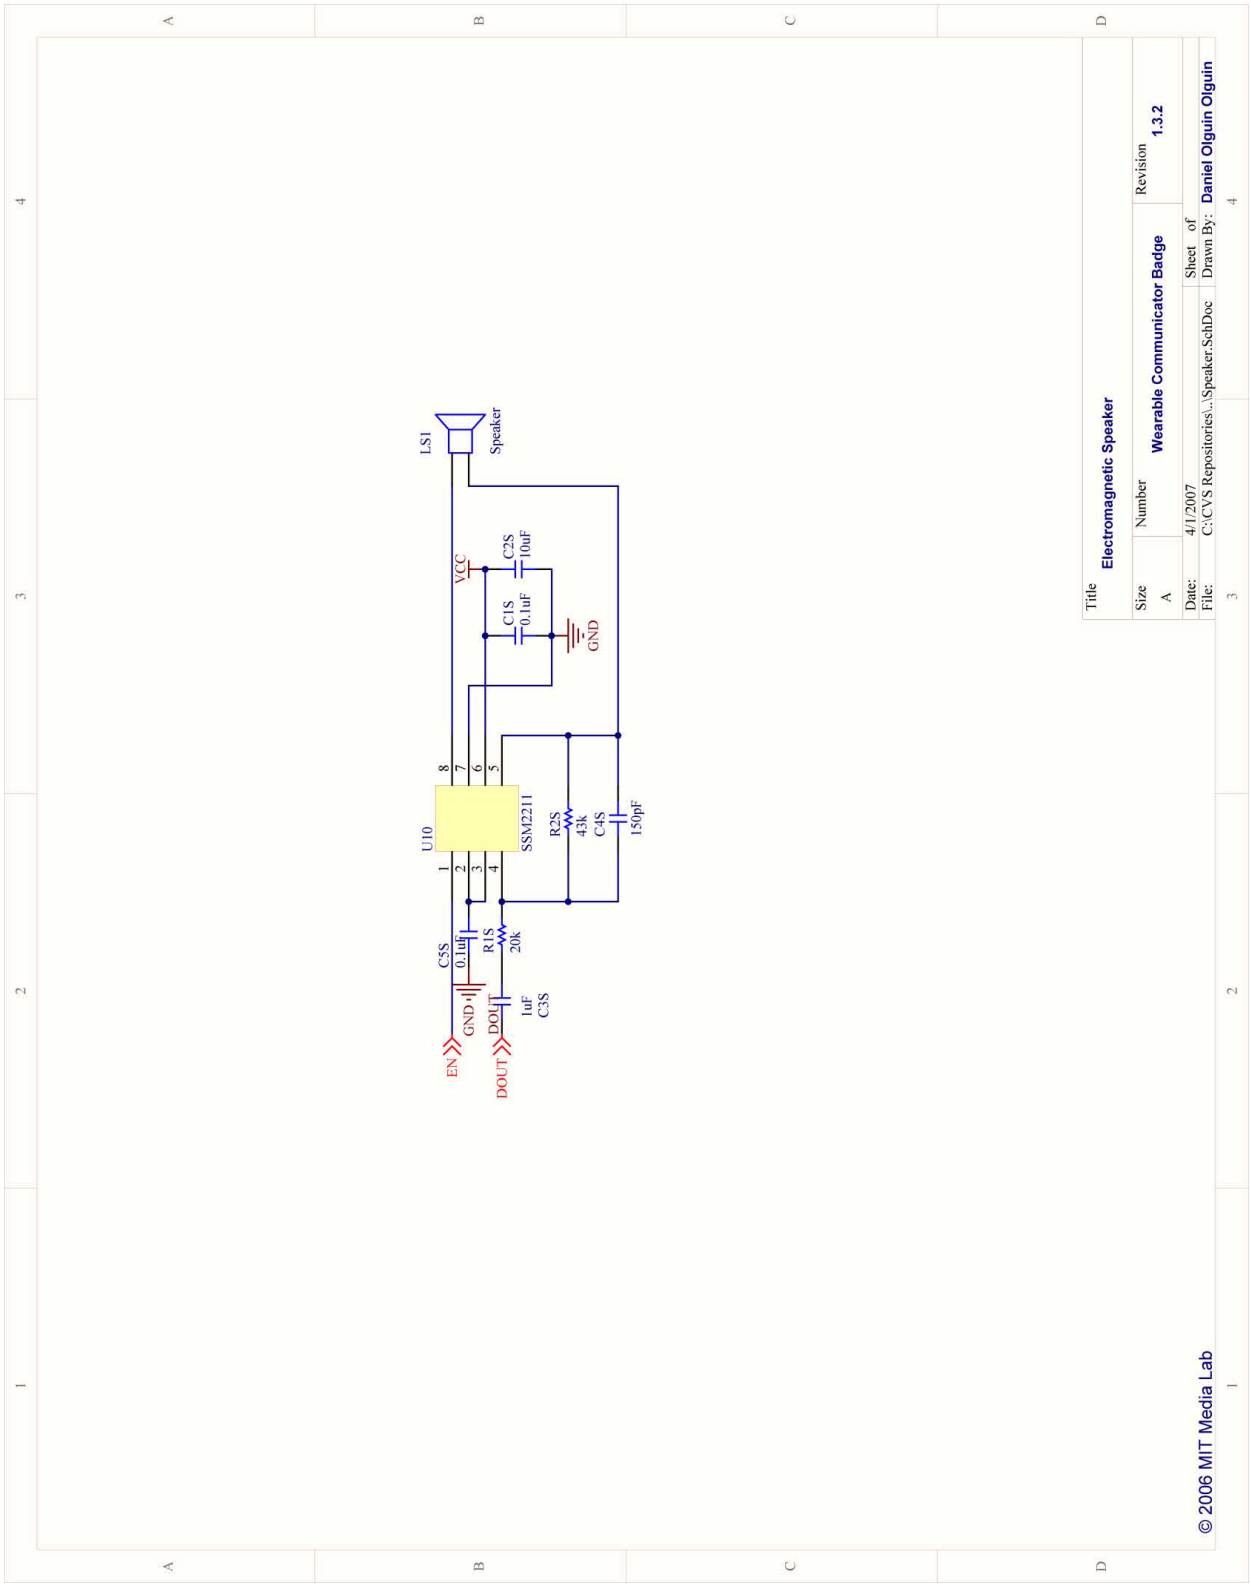

Figure C-12: Electromagnetic speaker with audio power amplifier.

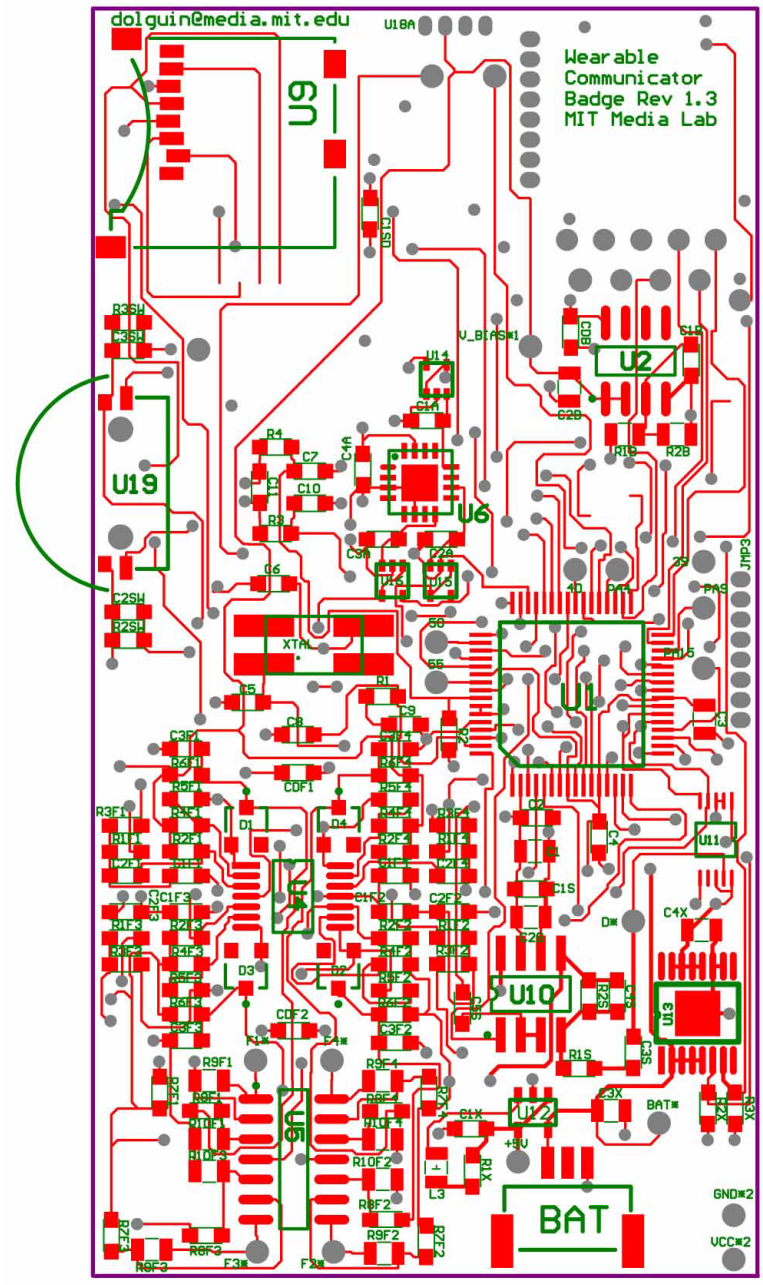

Figure C-13: Badge PCB top layer layout.

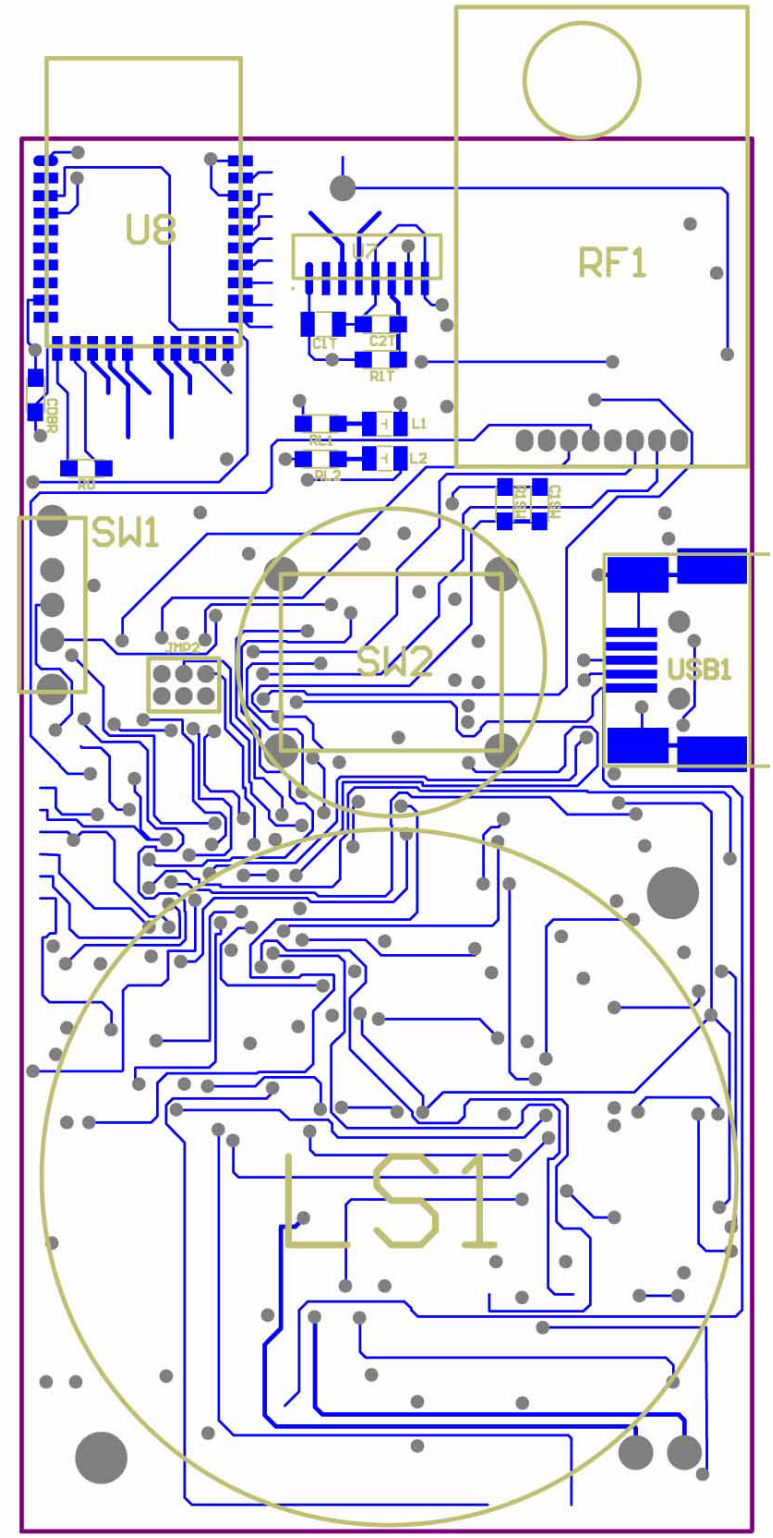

Figure C-14: Badge PCB bottom layer layout.

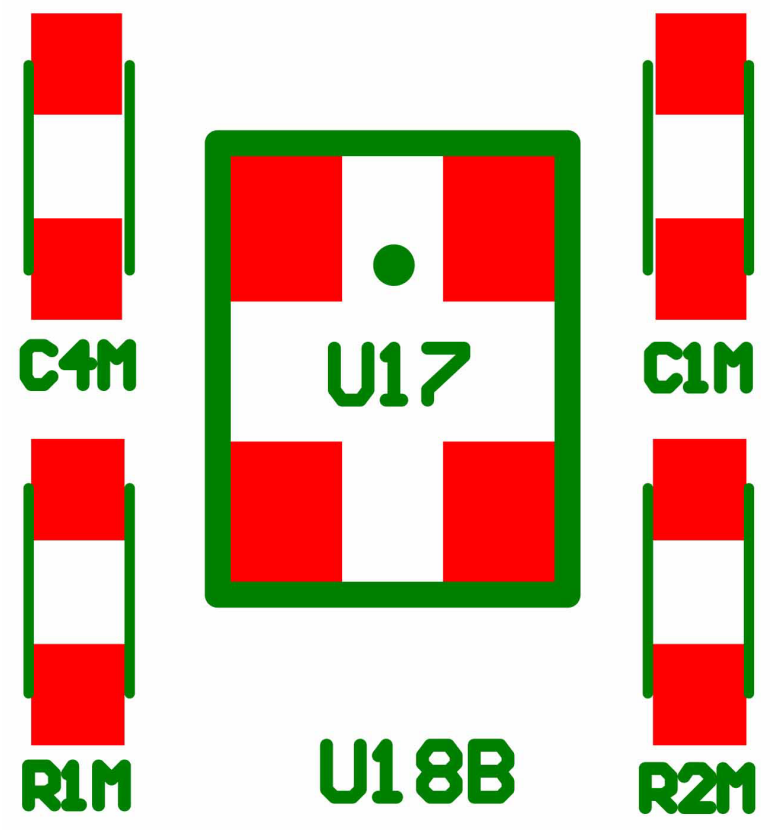

Figure C-15: Microphone PCB top layer layout

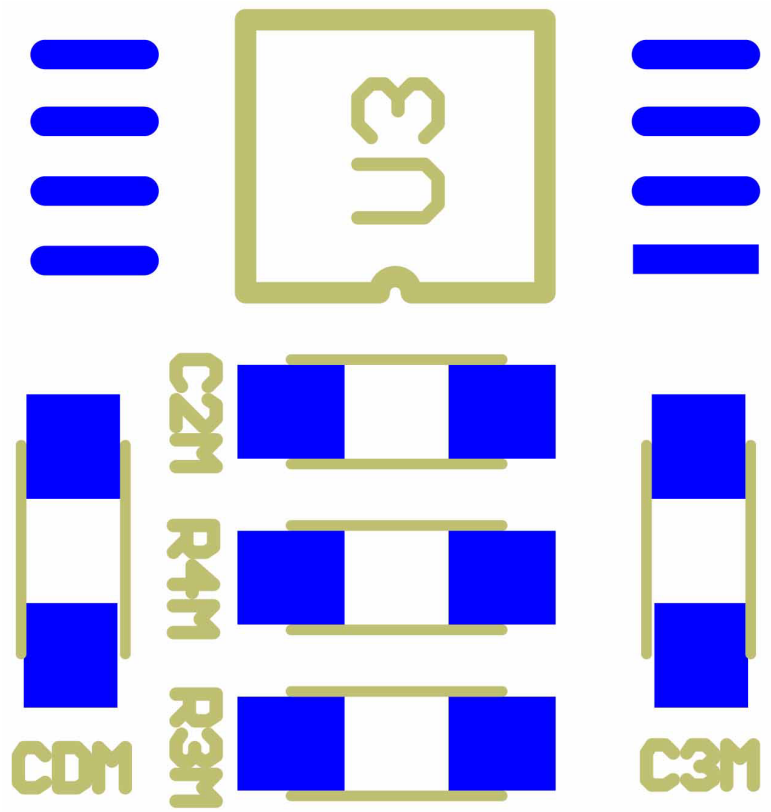

Figure C-16: Microphone PCB bottom layer layout.

Appendix D

Microcontroller's Peripherals Input/Output Configuration

Table D.1 contains the list of peripherals used with the AT91SAM7S256 microcontroller and the PIO controller configuration. Bold text indicates whether peripheral A or B is being used.

I/O Line	Peripheral A	Peripheral B	Comments
PA0	**PWM0**	TIOA0	Green LED
PA1	**PWM1**	TIOB0	Blue LED
PA2	**PWM2**	SCK0	Push button
PA3	TWD	**NPCS3**	MicroSD chip select
PA4	TWCK	TCLK0	
PA5	**RXD0**	NPCS3	Bluetooth TX
PA6	**TXD0**	PCK0	Bluetooth RX
PA7	**RTS0**	PWM3	Bluetooth CT
PA8	**CTS0**	ADTRG	Bluetooth RT
PA9	DRXD	NPCS1	
PA10	DTXD	NPCS2	
PA11	**NPCS0**	PWM0	CC2500 chip select
PA12	**MISO**	PWM1	CC2500 MOSI
PA13	**MOSI**	PWM2	CC2500 MISO
PA14	**SPCK**	PWM3	CC2500 SPCK
PA15	TF	TIOA1	
PA16	TK	TIOB1	
PA17	TD	PCK1	**AD0** Voice
PA18	RD	PCK2	**AD1** Accelerometer X
PA19	RK	FIQ	**AD2** Accelerometer Y
PA20	RF	IRQ0	**AD3** Accelerometer Z
PA21	**RXD1**	PCK1	IrDA RX
PA22	**TXD1**	NPCS3	IrDA TX
PA23	SCK1	PWM0	
PA24	RTS1	PWM1	
PA25	**CTS1**	PWM2	SEL1 (selector switch)
PA26	**DCD1**	TIOA2	SEL2 (selector switch)
PA27	**DTR1**	TIOB2	CC2500 IO2
PA28	**DSR1**	TCLK1	RST (charger)
PA29	**RI1**	TCLK2	EN (speaker)
PA30	**IRQ1**	NPCS2	CC2500 IO1
PA31	**NPCS1**	PCK2	DAC chip select

Table D.1: Microcontroller's peripherals configuration.

Appendix E

Bill of Materials

Sociometric Badge — Bill of Materials

Item	Quantity per badge	Value	Designator	Manufacturer	Manufacturer Part#	Description	Vendor	Vendor Part #	Unit Price in quantities of 100	100 Badges Cost = Unit Price * Quantity per Badge
1	2	10pF	C5, C6	PANASONIC	ECJ-1VC1H100D	Capacitor Ceramic 50V NP0 5%	Digikey	PCC100CVCT-ND	0.02110	0.0422
2	2	15pF	C10, C7	PANASONIC	ECJ-1VC1H150J	Capacitor Ceramic 50V NP0 5%	Digikey	PCC150ACVCT-ND	0.02110	0.0422
3	1	33pF	C11	PANASONIC	ECJ-1VC1H330J	Capacitor Ceramic 50V NP0 5%	Digikey	PCC330ACVCT-ND	0.02110	0.0211
4	1	150pF	C4S	PANASONIC	ECJ-1VC1H151J	Capacitor Ceramic 50V NP0 5%	Digikey	PCC151ACVCT-ND	0.02110	0.0211
5	2	1nF	C4, C8	PANASONIC	ECJ-1VC1H102J	Capacitor Ceramic 50V NP0 5%	Digikey	PCC2151CT-ND	0.04340	0.0868
6	1	10nF	C9	PANASONIC	ECJ-1VB1H103K	Capacitor 50V X7R 10%	Digikey	PCC1784CT-ND	0.01830	0.0183
7	8	22nF	**C1F1, C2F1, C1F2, C2F2, C1F3, C2F3, C1F4, C2F4**	AVX Corporation	06035C223JAT2A	Capacitor Ceramic 50V X7R 5%	Digikey	478-3722-1-ND	0.22750	1.82
8	4	68nF	C3F1, C3F2, C3F3, C3F4	PANASONIC	ECJ-1VB1C683K	Capacitor Ceramic 16V X7R 10%	Digikey	PCC1780CT-ND	0.02360	0.0944
9	13	0.1uF	C2, C1SW, C2SW, C3SW, C1S, C5S, C2T, C1SD, C4A, CDB, CDF1, CDF2, CDBR	PANASONIC	ECJ-1VB1H104K	Capacitor Ceramic 50V X7R 10%	Digikey	PCC2398CT-ND	0.05870	0.7631
10	3	0.22uF	C1A, C2A, C3A	PANASONIC	ECJ-1VB1A224K	Capacitor Ceramic 10V X5R 10%	Digikey	PCC1748CT-ND	0.03250	0.0975
11	0	0.47uF	C3M, C4M (Both caps are in Mic Board)	PANASONIC	ECJ-1VB1A474K	Capacitor Ceramic 10V X5R 10%	Digikey	PCC2275CT-ND	0.03250	0
12	3	1uF	C1B, C3S, C1X	PANASONIC	ECJ-1VB1C105K	Capacitor Ceramic 16V X5R 10%	Digikey	PCC2224CT-ND	0.06200	0.186
13	4	4.7uF	C1, C3, C1T, C3X	PANASONIC	ECJ-2FB1C475K	Capacitor Ceramic 16V X5R 10%	Digikey	PCC2323CT-ND	0.16050	0.642
14	3	10uF	C2S, C2B, C4X	PANASONIC	ECJ-2FB1A106K	Capacitor Ceramic 10V X5R 10%	Digikey	PCC2403CT-ND	0.52450	1.5735

Figure E-1: Bill of materials (part 1)

15	2	27ohms	R3, R4	Susumu Co Ltd	RR0816Q-270-D	Resistor 1/16W 0.5%	Digikey	RR08Q27DCT-ND	0.07540	0.1508
16	1	47ohms	R1T	Susumu Co Ltd	RR0816Q-47-D	Resistor 1/16W 0.5%	Digikey	RR08Q47DCT-ND	0.07540	0.0754
17	2	100ohms	RL1, RL2	Susumu Co Ltd	RR0816P-101-D	Resistor 1/16W 0.5%	Digikey	RR08P100DCT-ND	0.07540	0.1508
18	1	200ohms	R1X	Susumu Co Ltd	RR0816P-201-D	Resistor 1/16W 0.5%	Digikey	RR08P200DCT-ND	0.07540	0.0754
19	4	1k	**R4F1, R4F2, R4F3, R4F4**	Susumu Co Ltd	RR0816P-102-B-T5	Resistor 1/16W 0.1%	Digikey	RR08P1.0KBCT-ND	0.31080	1.2432
20	2	1.5k	R1, R2	Susumu Co Ltd	RR0816P-152-D	Resistor 1/16W 0.5%	Digikey	RR08P1.5KDCT-ND	0.07540	0.1508
21	4	1.6k	**R5F1, R5F3, R5F4, R5F2**	Susumu Co Ltd	RR0816P-162-B-T5	Resistor 1/16W 0.1%	Digikey	RR08P1.6KBCT-ND	0.31080	1.2432
22	3	4.12k	**R1F4, R2F4, R3F4**	Susumu Co Ltd	RR0816P-4121-B-I5-60H	Resistor 1/16W 0.1%	Digikey	RR08P4.12KBCT-ND	0.31080	0.9324
23	1	5.9k	R6F4	Susumu Co Ltd	RR0816P-5901-D-75H	1/16W 0.5%	Digikey	RR08P5.9KDCT-ND	0.07540	0.0754
24	3	10.7k	**R1F3, R2F3, R3F3**	Susumu Co Ltd	RR0816P-1072-T5-04C	Resistor 1/16W 0.1%	Digikey	RR08P10.7KBCT-ND	0.31080	0.9324
25	1	15.8k	R6F3	Susumu Co Ltd	RR0816P-1582-D-20C	Resistor 1/16W 0.5%	Digikey	RR08P15.8KDCT-ND	0.07540	0.0754
26	1	20k	R1S	Susumu Co Ltd	RR0816P-203-D	Resistor 1/16W 0.5%	Digikey	RR08P20.0KDCT-ND	0.07540	0.0754
27	3	28.7k	**R1F2, R2F2, R3F2**	Susumu Co Ltd	RR1220P-2872-B-M-T5	Resistor 1/10W 0.1%	Digikey	RR12P28.7KBCT-ND	0.31080	0.9324
28	1	40.2k	R6F2	Susumu Co Ltd	RR0816P-4022-D-59C	1/16W 0.5%	Digikey	RR08P40.2KDCT-ND	0.07540	0.0754
29	1	43.0k	R2S	Susumu Co Ltd	RR0816P-433-D	1/16W 0.5%	Digikey	RR08P43.0KDCT-ND	0.07540	0.0754
30	4	75k	**R1F1, R2F1, R3F1, R4M**	Susumu Co Ltd	RR0816P-753-B-T5	Resistor 1/16W 0.1%	Digikey	RR08P75.0KBCT-ND	0.31080	1.2432
31	5	100K	R1SW, R2SW, R3SW, R2X, R3X	Susumu Co Ltd	RR0816P-104-D	Resistor 1/16W 0.5%	Digikey	RR08P100KDCT-ND	0.07540	0.377
32	1	107k	R6F1	Susumu Co Ltd	RR0816P-1073-D-04D	Resistor 1/16W 0.5%	Digikey	RR08P107KDCT-ND	0.07540	0.0754
33	8	332k	R7F1, R8F1, R7F2, R8F2, R7F3, R8F3, R7F4, R8F4	Susumu Co Ltd	RR0816P-3323-D-51D	Resistor 1/16W 0.5%	Digikey	RR08P332KDCT-ND	0.07540	0.6032
34	2	1M	**R1B, R2B**	Susumu Co Ltd	RR1220P-105-B-T5	Resistor 1/10W 0.1%	Digikey	RR12P1.0MBCT-ND	0.31080	0.6216
35	8	1M	R9F1, R9F2, R9F3, R9F4, R10F1, R10F2, R10F3, R10F4	Susumu Co Ltd	RR1220P-105-D	Resistor 1/10W 0.5%	Digikey	RR12P1.0MDCT-ND	0.07540	0.6032
36	4	500mA	D1, D2, D3, D4	Diodes Inc	BAT4000-7-F	Shottky Diode	Digikey	BAT4000-FDICT-ND	0.43200	1.728
37	1	18.432 MHz	XTAL	CTS	406C35B18M43200	Crystal	Digikey	CTX668CT-ND	1.87500	1.875

Figure E-2: Bill of materials (part 2)

38	1		L1	Avago Technologies	HSMM-C170	Green LED 120 mCandela 3.3V	Digikey	516-1435-1-ND	0.43250	0.4325
39	1		L2	Avago Technologies	HSMN-C170	Blue LED 120 mCandela 3.3V	Digikey	516-1437-1-ND	0.31140	0.3114
40	1		L3	Lite-ON	LTST-C171CKT	Red LED	Digikey	160-1421-1-ND	0.09500	0.095
41	1		U1	ATMEL	AT91SAM7S64	AT91 ARM Thumb-Based Microcontroller	Digikey	AT91SAM7S64-AU001 ND	5.50000	5.5
42	1	3 contacts	BAT-A	Molex	51146-0300	Receptacle housing	Digikey	WM5401-ND	0.35860	0.3586
43	1	3 contacts	BAT	Molex	53780-0370	Right Angle SMT Header	Digikey	WM7801CT-ND	1.15280	1.1528
44	2		BAT-C	Molex	50641-8141	Crimp Terminal AWG 2830	Digikey	WM5508-ND	0.28880	0.5778
45	0.25	8-pin connector	RF1 (Daughterboard: CC2500 Radio)	Mill-Max	850-10-050-10-001000	Connector Strip Header 50 Pos	Digikey	ED8250-ND	5.71050	1.427625
46	0.17	6-pin in 2 lines connector	JMP2	Mill-Max	853-93-100-10-001000	Connector Strip Socket 2 x 50 Pos	Digikey	ED85100-ND	15.85250	2.65875
47	0.125	4-pin right angle connector	U18A	Mill-Max	850-10-050-20-001000	Connector Header Right Angle 50 Pos 0.05" pitch	Digikey	ED8650-ND	7.75130	0.9688125
48	1	Side Switch	SW1	E-Switch	EG1213	Sideswitch	Digikey	EG1908-ND	0.54540	0.5454
49	3	SC70-5	U14, U15, U16	National Semiconductor	LPV511MG	Voltage follower	Digikey	LPV511MGCT-ND	1.20700	3.621
50	1		U6	Analog Devices	ADXL330KCPZ	Small Low Power 3-Axis +-3g iMEMS Accelerometer	Analog Devices	ADXL330KCPZ	6.95000	6.95

Figure E-3: Bill of materials (part 3)

51	1		U5 (ABCD)	Analog Devices	OP481	Ultra Low Power Rail to Rail Output Operational Amplifiers	Analog Devices	OP481GSZ	3.39000	3.39
52	1		U2 (AB)	Analog Devices	OP281	Ultra Low Power Rail to Rail Output Operational Amplifiers	Analog Devices	OP281GSZ	2.59000	2.59
53	1		U4 (ABCD)	Analog Devices	AD8544	Quad Rail-to-Rail Input and Output, Single Supply Amplifier	Analog Devices	AD8544ARUZ	0.59000	0.59
54	1		U10	Analog Devices	SSM2211	Low Distortion 1.5 Watt Audio Power Amplifier	Analog Devices	SSM2211SZ	0.42000	0.42
55	1		U11	Analog Devices	AD5432	10-Bit High Bandwidth Multiplying DAC with Serial Interface	Analog Devices	AD5432YRMZ	2.81000	2.81
56	1		U7	Vishay	TFDU4300	IrDA Module	Mouser	782-TFDU4300	1.94000	1.94
57	1		U9	Molex	538-500873-0801	microSD card connector	Mouser	500873-0801	2.04000	2.04
58	1	8-ohm speaker	LS1	ICC/Intervox	S1505LA-R5W	Speaker, Round	Allied Electronics	623-0008	2.06000	2.06
59	1		SW2	NKK Switches	JF15CP-2H	Push Button (gray)	Allied Electronics	870-9355	1.65700	1.657
60	1		USB1	Tyco	1-440247-1	Mini-USB port	Allied Electronics	512-1829	0.72000	0.72
61	1		U8	BlueRadios	BR-C48A	Bluetooth	BlueRadios	BR-C48A	34.00000	34
62	1		U19	ITT	TPC1133GLFG	Tri-direction scan switch	Digikey	TPC1133GLFG	1.50150	1.5015
63	1		U12	MAXIM	MAX1555	USB Battery Charger	MAXIM	MAX1555EZK+	1.34000	1.34
64	1		U13	MAXIM	MAX1793	Voltage Regulator	MAXIM	MAX1793EUE33+	1.83	1.83
65	1	zero-ohm	R0	Yageo Corporation	RC0603JR-070RL	Resistor 1/16W 0.5%	Digikey	311-0.0GRCT-ND	0.0378	0.0378

Figure E-4: Bill of materials (part 4)

Extra costs										
	1							PCB	15.00000	15
	1							Assembly	45.00000	45
	1							Li-ON Battery	18.00000	18
	1							USB Cable	1.00000	1
	1							AC Wall Plug to 5V USB Adapter	3.95000	3.95
	1							microSD card	15.00000	15
	1							CC2500 Radio Board	25.00000	25
	1							Enclosure	15.00000	15
	1							Microphone board	12.00000	12
										Cost per Badge
										250.27

Figure E-5: Bill of materials (part 5)

Bibliography

[1] At91sam 32-bit arm-based microcontrollers devices. http://www.atmel.com/dyn/products/devices.asp?family_id=605#1586.

[2] Infrared data association. http://www.irda.org.

[3] ipod nano lanyard headphones. http://www.apple.com/ipodnano/accessories.html.

[4] Knowles acoustics. http://www.knowlesacoustics.com/html/sil_mic.html#advantages.

[5] Mvox technologies. http://www.mvox.com/.

[6] Nokia 5500 sport phone. http://europe.nokia.com/A4160003.

[7] Ntag interactive. http://www.ntag.com/.

[8] Plug sensor network. responsive environments group, mit media laboratory. http://www.media.mit.edu/resenv/plug/.

[9] Toolless plastic technologies, inc. http://www.toolless.com/index.html.

[10] Vocera communications system. http://www.vocera.com.

[11] E. Alpaydin. *Introduction to Machine Learning*. The MIT Press, Cambridge, MA, 2004.

[12] Analog Devices, Inc., Norwood, MA. *Ultralow Power, Rail-to-Rail Output Operational Amplifiers*, Mar 2003. OP281/OP481 Data Sheet Rev. B.

[13] Analog Devices, Inc., Norwood, MA. *Low Distortion 1.5 Watt Audio Power Amplifier*, Oct 2004. SSM221 Data Sheet Rev. C.

[14] Analog Devices, Inc., Norwood, MA. *Small, Low Power, 3-Axis ±3 g iMEMS Accelerometer*, Mar 2006. SSM221 Data Sheet Rev. 0.

[15] K. Ara, N. Kanehira, E. Megally, Y. Poltorak, G. Singh, R. Smith, D. Suzuki, M. Mortensen, M. Van Alstyne, and A. Pentland. Sensible organizations inspired by social sensor technologies. Technical report, MIT Media Lab and Sloan School of Management, May 2006. http://vismod.media.mit.edu//tech-reports/TR-602.pdf.

[16] R.J. Baken. *Clinical Measurement of Speech and Voice*. Taylor and Francis Ltd, London, UK, 1987.

[17] L. Bao and S. S. Intille. Activity recognition from user-annotated acceleration data. In *Proceedings of the 2nd International Conference on Pervasive Computing*, pages 1–17, Apr 2004.

[18] BlueRadios, Englewood, CO. *BR-C46AR class 2 Bluetooth module*, Mar 2006. BR-C46AR Data Sheet.

[19] M. L. Blum. Real time context recognition. Master's thesis, Swiss Federal Institute of Technology Zurich (ETH) & MIT Media Lab, Oct 2005.

[20] P. Bonato. Advances in wearable technology and applications in physical medicine and rehabilitation. *Journal of NeuroEngineering and Rehabilitation*, 2(2):1–4, Feb 2005.

[21] R. Borovoy, F. Martin, S. Vemuri, M. Resnick, Silverman B., and C. Hancock. Meme tags and community mirrors: Moving from conferences to collaboration. In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work*, pages 1–10, 1998.

[22] R. Borovoy, M. McDonald, F. Martin, and M. Resnick. Things that blink: Computationally augmented name tags. *IBM Systems Journal*, 35(3):488–495, 1996.

[23] J. L. Bowditch and A. F. Buono. *A Primer on Organizational Behavior*. John Wiley & Sons, Inc., Danvers, MA, 2005.

[24] O. Cakmakci, J. Coutaz, K. Van Laerhoven, and H. Gellersen. Context awareness in systems with limited resources. In *Proceedings of AIMS: Artificial Intelligence in Mobile System*, pages 1–8, July 2002.

[25] Chipcon Products from Texas Instruments, Cupertino, CA. *Single Chip Low Cost Low Power RF Transceiver*, Oct 2005. CC2500 Data Sheet Rev. 1.1.

[26] T. Choudhury. *Sensing and Modeling Human Networks*. PhD thesis, Massachusetts Institute of Technology, Media Laboratory, Feb 2004.

[27] T. Choudhury and P. A. The sociometer: A wearable device for understanding human networks. In *Computer Supported Cooperative Work - Workshop on Ad hoc Communications and Collaboration in Ubiquitous Computing Environments*, 2002.

[28] J. Curhan, A. Pentland, R. Caneel, N. Eagle, and M. Martin. Thin slices of negotiation. In *Proceedings of the Annual Meeting of the Academy of Management*, Aug 2005.

[29] R. W. DeVaul and S. Dunn. Real-time motion classification for wearable computing applications. Technical report, MIT Media Lab, Dec 2001.

[30] N. Eagle and A. Pentland. Eigenbehaviors: Identifying structure in routine. In *Ubicomp 2006: The 8th International Conference on Ubiquitous Computing*, 2006.

[31] J. Farringdon, A. J. Moore, N. Tilbury, J. Church, and P.D. Biemond. Wearable sensor badge & sensor jacket for context awareness. *Proceedings of the 3rd International Symposium on Wearable Computers*, pages 107–113, Oct 1999.

[32] S. Franco. *Design with Operational Amplifiers and Analog Integrated Circuits*. McGraw-Hill, New York, NY, 2002.

[33] Z. Ghahramani. An introduction to hidden markov models and bayesian networks. *International Journal of Pattern Recognition and Artificial Intelligence*, 15(1):9–42, Jan 2001.

[34] J. P. Gips. Social motion: Mobile networking through sensing human behavior. Master's thesis, Massachusetts Institute of Technology, Media Laboratory, Sept 2006.

[35] J. Greenberg and R. A. Baron. *Behavior in Organizations*. Prentice Hall, Upper Saddle River, NJ, eighth edition, 2003.

[36] B. Haskell. *Portable electronics product design and development: for cellular phones, PDAs, digital cameras, personal electronics, and more*. McGraw-Hill, New York, NY, first edition, 2004.

[37] International Components, Corp, Hauppauge, NY. *Micro-Miniature Low Profile Speakers*. S150SLA-R5W Data Sheet.

[38] D. M. Karantonis, M. R. Narayanan, M. Mathie, N. H. Lovell, and B. G. Celler. Implementation of a real-time human movement classifier using a triaxial accelerometer for ambulatory monitoring. *IEEE Transactions on Information Technology and Biomedicine*, 10(1):156–167, Jan 2006.

[39] N. Kern and B. Schiele. Context-aware notification for wearable computing. In *Proceedings of the 7th International Symposium on Wearable Computing*, pages 223–230, Oct 2003.

[40] T. Kim, A. Chang, and A. Pentland. Enhancing organizational communication using sociometric badges. In *Submitted to the 11th International Symposium on Wearable Computers*, Oct 2007.

[41] B. L. Kirkman, B. Rosen, P.E. Tesluk, and C. B. Gibson. The impact of team empowerment on virtual team performance: the moderating role of face-to-face interaction. *Academy of Management*, 47(2):175–192, 2004.

[42] Knowles Electronics, LCC., Itasca, IL. *Amplified Mini SiSonic Microphone Specification*, Nov 2005. SPM0103NE3 Data Sheet Rev. C.

[43] M. Laibowitz, J. Gips, R. Aylward, A. Pentland, and J. A. Paradiso. A sensor network for social dynamics. In *Proceedings of the 5th International Conference on Information Processing in Sensor Networks*, pages 483–491, Apr 2006.

[44] J. Lehtonen. *Designing Context-Aware Systems: Choosing Sensors.* Laboratory of Computational Engineering, HUT, P.O. Box 9203, FIN-021015, Finland, Feb 2004. T.121-900 User Interfaces and Usability Seminar Notes.

[45] J. Lester, T. Choudhury, N. Kern, G. Borriello, and B. Hannaford. A hybrid discriminative/generative approach for modeling human activities. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, pages 766–722, 2005.

[46] J. Lifton, M. Feldmeier, Y. Ono, C. Lewis, and J. A. Paradiso. A platform for ubiquitous sensor deployment in occupational and domestic environments. In *Proceedings of SPOTS '07*, 2007. In review.

[47] P. Lukowicz, J. A. Ward, H. Junker, M. Stager, G. Troster, A. Atrash, and T. Starner. Recognizing workshop activity using body worn microphones and accelerometers. In *Proceedings of the 2nd International Conference on Pervasive Computing*, pages 18–22, Apr 2004.

[48] A. Madan and A. Pentland. Vibefones: Socially aware mobile phones. In *Proceedings of the 10th International Symposium on Wearable Computers*, pages 109–112, Oct 2006.

[49] J. Mantyjarvi. *Sensor-based context recognition for mobile applications.* VTT Publications, Finland, 2003.

[50] M. J. Mathie, A. C. F. Coster, N. H. Lovell, and B. G. Celler. Accelerometry: providing an integrated, practical method for long-term ambulatory monitoring of human movement. *Physiological Measurement*, 25(4):R1–R20, 2004.

[51] U. Maurer, A. Smailagic, D. P. Siewiorek, and M. Deisher. Activity recognition and monitoring using multiple sensors on different body positions. In *International Workshop on Wearable and Implantable Body Sensor Networks*, Apr 2006.

[52] Maxim Integrated Products, Sunnyvale, CA. *Low-Dropout, Low IQ, 1A Linear Regulator*, Nov 2000. MAX1793 Data Sheet Rev. 1.

[53] Maxim Integrated Products, Sunnyvale, CA. *SOT23 Dual-Input USB/AC Adapter 1-Cell Li+ Battery Chargers*, Oct 2003. MAX1551/MAX1555 Data Sheet Rev. 0.

[54] Maxim Integrated Products, Sunnyvale, CA. *Charging Batteries Using USB Power*, May 2004. Application Note 3241.

[55] Maxim Integrated Products, Sunnyvale, CA. *Charging Batteries Using USB Power*, Aug 2005. Application Note 3607.

[56] R. Mayrhofer, H. Radi, and A. Ferscha. Recognizing and predicting context by learning from user behavior. In *Proceedings of the 1st International Conference On Advances in Mobile Multimedia (MoMM'03)*, volume 171, pages 25–35. Austrian Computer Society (OCG), September 2003.

[57] B. A. Miller and C. Bisdikian. *Bluetooth Revealed*. Prentice Hall, Upper Saddle River, NJ, 2001.

[58] D. Minnen, T. Starner, J. A. Ward, P. Lukowicz, and G. Troster. Recognizing and discovering human actions from on-body sensor data. In *Proceedings of the IEEE International Conference on Multimedia and Expo*, pages 1545–1548, Jul 2005.

[59] E. Munguia Tapia, N. Marmasse, S. S. Intille, and K. Larson. Mites: Wireless portable sensors for studying behavior. In *Proceedings of the 6th International Conference on Ubiquitous Computing*, Sep 2004.

[60] K. Murphy. Hidden markov model (hmm) toolbox for matlab. http://www.cs.ubc.ca/murphyk/Software/HMM/hmm.html.

[61] D. Olguin Olguin, J. A. Paradiso, and A. Pentland. Wearable communicator badge: Designing a new platform for revealing organizational dynamics. In *Proceedings of the 10th International Symposium on Wearable Computers (Student Colloquium)*, pages 4–6, Oct 2006.

[62] D. Olguin Olguin and A. Pentland. Human activity recognition: Accuracy across common locations for wearable sensors. In *Proceedings of the 10th International Symposium on Wearable Computers (Student Colloquium)*, pages 11–13, Oct 2006.

[63] D. Oster. Angewandte analyse sozialer netzwerke in unternehmen: Auswirkungen elektronischer kommunikation auf die performanz von teams. Diploma Thesis, University of Cologne Department of Information Systems and Information Management, Germany , 2007.

[64] S. Park, I. Locher, A. Savvides, A. Chen, R. Muntz, and S. Yuen. Design of a wearable sensor badge for smart kindergarten. *Proceedings of the 6th International Symposium on Wearable Computers*, pages 231–238, Oct 2002.

[65] A. Pentland. Automatic mapping and modeling of human networks. To appear in: *Physica A*, Dec 2006.

[66] A. Pentland. Social dynamics: Signals and behavior. In *Proceedings of the 3rd International Conference on Developmental Learning*, pages 263–267, Oct 2004.

[67] A. Pentland. Socially aware computation and communication. *IEEE Computer*, 38(3):33–40, Mar 2005.

[68] R. W. Picard. *Affective Computing*. MIT Press, Cambridge, Massachusetts, first edition, 1997.

[69] L. R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *IEEE Transactions on Signal Processing*, 77(2):257–286, Feb 1989.

[70] SanDisk Corporation, Sunnyvale, CA. *SanDisk SD Card Product Manual*, Nov 2004. Data Sheet Rev. 2.2.

[71] A. Sinan, E. Brynjolfsson, and Van Alstyne M. W. Information, technology and information worker productivity: Task level evidence. In *Proceedings of the International Conference on Information Systems*, Dec 2006.

[72] V. Stanford. Beam me up, doctor mccoy. *IEEE Pervasive Computing*, pages 13–18, Jul-Sep 2003.

[73] W. T. Stoltzman. Toward social signaling framework: Activity and emphasis in speech. Master's thesis, Massachusetts Institute of Technology, Media Laboratory, Sept 2006.

[74] M. Sung, C. Marci, and A. Pentland. Wearable feedback systems for rehabilitation. *Journal of NeuroEngineering and Rehabilitation*, 2:17–28, Jun 2005.

[75] K. Van Laerhoven and O. Cakmakci. What shall we teach our pants. In *Proceedings of the 4th International Symposium on Wearable Computers*, pages 77–83, Oct 2000.

[76] Vishay Semiconductors, Heilbronn, Germany. *Infrared Transceiver Module (SIR, 115.2 kbits/s) for IrDA Applications*, Dec 2005. TFDU4300 Data Sheet Rev. 1.5.

[77] B. N. Waber, D. Olguin Olguin, T. Kim, A. Mohan, K. Ara, and A. Pentland. Organizational engineering using sociometric badges. In *NetSci: International Workshop and Conference on Network Science (Contributed Talk)*, May 2007.

[78] M. Weiser. The computer for the 21st century. *IEEE Pervasive Computing*, pages 19–25, Jan-Mar 2002. First published in *Scientific American*, September 1991, pp. 94-104.