

# An Efficient, Full Stack Protocol for Body Area Mesh Networks

*Ramakrishnan Menon  
Jan M. Rabaey, Ed.  
Ali Moin, Ed.*

Electrical Engineering and Computer Sciences  
University of California at Berkeley

Technical Report No. UCB/EECS-2018-22

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2018/EECS-2018-22.html>

May 1, 2018



Copyright © 2018, by the author(s).  
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

---

**An Efficient, Full Stack Protocol for Body Area Mesh Networks**

Ramakrishnan S. Menon

---

**Research Project**

Submitted to the Department of Electrical Engineering and Computer Sciences,  
University of California at Berkeley, in partial satisfaction of the requirements for the  
degree of **Master of Science, Plan II.**

Approval for the Report and Comprehensive Examination:

**Committee:**



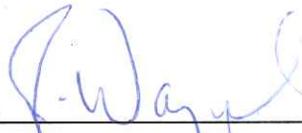
---

Professor Jan Rabaey  
Research Advisor

December 6, 2017

Date

\* \* \* \* \*



---

Professor John Wawrzynek  
Second Reader

12/8/2017

Date

## Abstract

An Efficient, Full Stack Protocol for Body Area Mesh Networks

by

Ramakrishnan S. Menon

Master of Science in Electrical Engineering and Computer Science

University of California, Berkeley

Professor Jan Rabaey, Chair

In anticipation of a new wave of continuous monitoring of the body and future human machine interfaces, research on Body Area Networks is becoming increasingly relevant. If on-body applications are to become ubiquitous, they must overcome challenges involving energy sparsity, security, and performance. Currently, IEEE 802.15.6 remains the only established standard designed with Body Area Networks in mind. The standard however caters towards data aggregation and does not fully address the range of applications that require reliable, low energy networking on the body. The Human Intranet takes a different approach and seeks to solve many of these challenges by providing a local, distributed network that can handle the processing and storage needs of sensors and actuators. In this report we explore an alternative protocol aligned with the vision of the Human Intranet. Specifically, we prototype a mesh network with off the shelf radios to evaluate the efficiency and reliability of a CSMA MAC layer and controlled flooding routing layer. We characterize the prototype network at the link level and benchmark the protocol at the network level for reliability and network lifetime. We find that CSMA based protocols can offer similar performance to TDMA based protocols under heavy traffic at higher power levels. A more realistic traffic pattern modeled after a potential Human Intranet use case reveals that the protocol can be reliable even at the lowest power levels in both static and dynamic environments.

## Acknowledgments

First I would like to thank Professor Rabaey for the opportunity to work on such varied and forward looking projects during my time at the BWRC. They have helped to spark my imagination and renew my excitement for the field. I would also like thank Ali Moin for his advice and help throughout this research project. He lent a helping hand in all aspects of the project, from the lab work to the report. Finally, I thank my parents for their guidance and support throughout my undergrad and graduate career here at Berkeley. I am grateful for all the effort they have put into providing me with an education second to none.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Background</b>	<b>4</b>
2.1	Body Area Networks . . . . .	4
2.2	Network Topology . . . . .	7
<b>3</b>	<b>Protocol Considerations</b>	<b>10</b>
3.1	MAC Layer . . . . .	10
3.2	Routing Layer . . . . .	13
<b>4</b>	<b>Prototype Implementation</b>	<b>14</b>
4.1	Hardware . . . . .	14
4.2	Firmware . . . . .	15
4.3	Debugging . . . . .	18
<b>5</b>	<b>Testing</b>	<b>21</b>
5.1	Power Baseline . . . . .	21
5.2	Link Level . . . . .	23
5.2.1	Channel Estimation . . . . .	24
5.2.2	Multihop . . . . .	26
5.3	Network Level . . . . .	27
5.3.1	Comparison to Network Simulations . . . . .	27
5.3.2	Realistic Human Intranet Traffic Pattern . . . . .	28

<b>6 Conclusion</b>	<b>31</b>
<b>Bibliography</b>	<b>33</b>

# List of Figures

1-1	Placement of Hubs and Nodes on the Body . . . . .	3
2-1	Illustration of Fully and Partial Mesh Topologies . . . . .	7
2-2	Illustration of Regular and Multihop Star Topologies . . . . .	8
3-1	Two nodes participating in a TDMA network . . . . .	11
3-2	Two nodes participating in a CSMA network . . . . .	12
4-1	Firmware Architecture . . . . .	16
4-2	State Machine of MAC Task . . . . .	17
4-3	State Machine of CSMA Routine . . . . .	17
4-4	State Machine of Flooding Module . . . . .	18
4-5	LCD Debug Screen via UART . . . . .	19
4-6	ROV View and Debug Window in Code Composer Studio . . . . .	19
5-1	Keithley 2612 connected to SensorTag and operated by Test Script over GPIB . . . . .	22
5-2	Power Consumption of CC2650 SensorTag . . . . .	22
5-3	Placement of Hubs in Human Intranet Mesh Network . . . . .	24

# List of Tables

5.1	Average RSSI of Packets Sent at 5 dBm . . . . .	24
5.2	Average RSSI of Packets Sent at 0 dBm . . . . .	25
5.3	Average RSSI of Packets Sent at -9 dBm . . . . .	25
5.4	Average RSSI of Packets Sent at -15 dBm . . . . .	25
5.5	Average RSSI of Packets Sent at -21 dBm . . . . .	25
5.6	Results of Multihop Test . . . . .	27
5.7	PDRs from Round Robin 1 kBps Network Test . . . . .	28
5.8	Parameters of Human Intranet Simulation [12, 9, 19] . . . . .	29
5.9	PDRs from Mock HI Network Test: Stationary . . . . .	29
5.10	PDRs from Mock HI Network Test: Active . . . . .	30

# Chapter 1

## Introduction

The Human Intranet describes a future for the convergence of wearables and the Internet of Things (IOT). It outlines a scalable, dynamic, and heterogeneous network of sensors and actuators located on the human body which may also interact with nearby external networks [22]. In contrast to the wearable and IOT implementations of present, the Human Intranet offers the promise of both a self-contained system and seamless integration with other networks. With these features the Human Intranet can act as an enabling framework for a whole host of new applications.

A canonical example of a currently untenable application is the portable and real-time control of a neuroprosthesis. In a typical setup, EEG acquisition systems or subdural implants record streams of neural data. The data stream undergoes a series of computationally intensive steps to translate brain waves into control commands for the prosthesis. The computation for data rich tasks like this often takes place off-body on a nearby PC, or on bulky hardware housed in a backpack [11, 16]. The Human Intranet remedies this situation by networking the existing hardware on the body and providing a platform for distributed computation and storage. The more efficient utilization of local hardware resources would result in more reasonable energy costs and reduced latency.

Another compelling application of the Human Intranet involves enhanced functionality during specific activities. Consider the case of endurance athletes training for grueling events like the Tour de France, a 21 day stage race that covers more than

2000 miles by bicycle. On its own, HI could provide invaluable continuous data to trainers and athletes about skin temperature, lactate levels, electrolyte loss, blood oxygenation, or even the onset of illness [10]. When temporarily augmented by the wireless network of sensors and actuators on a bike (GPS, cadence, power meters, electronic shifting), HI could further enhance training in real time. Imagine that the bike automatically changes gear to encourage higher cadence, or that the GPS could reroute a rider on a training ride to prevent over-exhaustion or fatigue based on vitals collected from a sweat sensor. This sort of integration between networks would enable synergistic human-machine interactions in a number of scenarios.

Underpinning one potential implementation of the Human Intranet is a hybrid network topology. The first topology is a star network that connects a series of 'nodes' to a proximal 'hub'. Since energy-sparsity is a major consideration for on-body applications, these nodes should be ultra-low power sensors that operate on harvested energy or inductive coupling and backscattering much like current RFID technologies. The hub would interrogate the sensor through intermittent polling rather than continuous measurement, essentially adopting a master-slave access scheme that is best served by a star topology. Recent developments show that in addition to the traditional array of sensors (ECG, temperature, accelerometer, gyroscope), bio-impedance, oximetry, and EMG sensors will soon be staples in on-body measurement and can operate in these energy starved environments [19].

The second topology, which serves as the glue for this HI implementation, is a series of hubs located around the body and configured as a mesh network. Hubs are mainly distinguished from nodes by their greater energy availability and the ability to store charge (by way of batteries). More concretely, the role of a hub may be satisfied by multipurpose SOCs, sensor-actuators like a neuroprosthesis, or even a less permanent fixture like a smartphone. Ultimately each of these hubs aggregates data from its local sensor network and relays the information wirelessly around the body as requested. In addition, the hubs can expose their nodes' and their own resources to the rest of the network to facilitate the distributed computation and storage that is key to the Human Intranet vision.

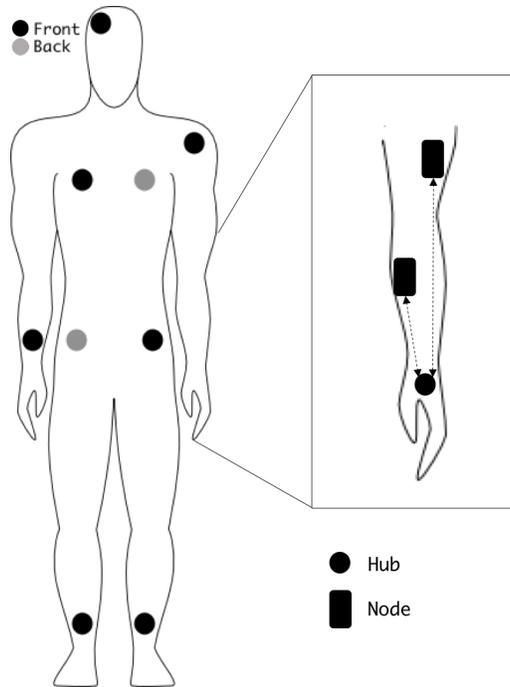


Figure 1-1: Placement of Hubs and Nodes on the Body

In this thesis, we focus on the design and implementation of a prototype mesh network and evaluate a potential protocol for low power Body Area Networks (BANs). The following chapter explains relevant details pertaining to BANs and mesh networks to understand prior work and the general constraints facing an HI implementation. We then discuss the specific design considerations for layers 2 and 3 of the custom protocol that we present. Next we outline the prototype architecture and steps involved in the implementation. The test setup, procedures, and results are discussed before the final chapter which concludes the report and notes directions for further research.

# Chapter 2

## Background

### 2.1 Body Area Networks

The Human Intranet is first and foremost a Body Area Network. Accordingly, it must be designed around a specific set of constraints not encountered in generic wireless sensor networks (WSNs). Perhaps the most burdensome constraint is the lack of available energy or 'energy sparsity' on the body. A simple solution would involve devices equipped with larger batteries. Real estate on the body, however, is at a premium and to gain mass adoption the device and power source must take on a small and unobtrusive form factor [9]. Additionally, these power sources must lengthen lifetime as frequent recharging of devices will also hamper adoption [19]. Barring the development of novel materials with higher charge density, energy harvesting is the most promising way to extend battery lifetime.

While there is diversity in energy harvesting sources (solar energy, temperature gradients, vibrations, and human propulsion are all viable) they are also terribly unpredictable due to the dynamic conditions surrounding a human during a typical day. Even assuming conditions based on a model of the average workday, the theoretical upper limit on power harvesting is mildly optimistic. An array of the most common energy harvesting devices, each occupying a square centimeter would produce a peak of 3 mWh around noon [9]. Of that, about 2.5 mWh is unsurprisingly harvested solar energy with the rest of the devices each contributing an order of magnitude less

power.

Another unique design constraint for BANs comes in the form of significant path loss. Path loss is the result of a number of effects as an electromagnetic wave propagates through space. The most basic effect is expansion of the signal as it travels a line-of-sight path through free space. In a less ideal environment, the atmosphere, buildings, and other objects cause a signal to reflect and refract along its journey. Multipath describes the effect where a signal encounters these obstacles and recombines either constructively or destructively with phase shifted and attenuated versions of itself. Generally these effects result in reduced signal strength at the receiver.

For BANs, the dramatic path loss is linked to absorption by the skin. A rough estimate for the incremental power reduction in decibels,  $L$ , by moving  $d$  meters from the source is given by the Log Distance Path Loss model [27].

$$L = 10n * \log_{10}(d)$$

Where  $n$  is the path loss exponent (PLE). The path loss exponent for free space is 2, but estimates for a channel in a Body Area Network are around 3.4 [26]. If the transmitting antenna is placed closer to the skin, the PLE can reach close to 4 [25]. Over a distance of a meter, a PLE of 3.4 equates to nearly a 40 dB loss in power at the receiver [4]. Moving body parts like arms and legs add further complication to the channel estimation. Ultimately, the body presents a poor and variable channel for wireless communication.

Though we addressed the most pressing design tradeoffs, a more comprehensive set of constraints must be considered for a BAN design to be widely adopted. A BAN typically has fewer nodes than a typical WSN, affecting the redundancy and quality of service of the network. The information being broadcast around the body must also be secured from eavesdroppers and the devices impervious to hacking. And if the emphasis on low power design was not underscored, limits on the Specific Absorption Rate, energy dissipated into the skin, restrict power usage even if the energy were available.

Despite the various standards proposed to tackle the efficient design of WSNs (Bluetooth, ZigBee, ANT), there is only one established standard for Body Area Networks. The IEEE 802.15.6 specification for wireless BANs introduced in 2012 provides guidelines for the hardware, software, and deployment of a BAN. It also addresses the secondary constraints we are ignoring for the purposes of this thesis. The standard is intentionally broad and allows for a number of different configurations but we will briefly discuss relevant characteristics of the default configuration.

The PHY or physical layer refers to the circuitry that facilitates the actual transmission and reception of data over the air. At this level of design, decisions are made about transmission frequencies, modulations, and power consumption. The standard specifies an ultrawideband (UWB) PHY, meaning the device can transmit in the 3.1 GHz to 10.6 GHz range albeit at restricted power levels. The use of this newly released high frequency range is most likely to avoid interference with WiFi or Bluetooth. The UWB PHY should also be capable of data rates between 487.5 kbps and 15 Mbps using on-off keying. Also of note is a mechanism for Clear Channel Assessment (CCA). CCA indicates whether the channel is occupied by other devices. This is usually achieved by measuring the amount of energy in a certain frequency band or by detecting packets using the same modulation scheme specified in the standard.

The Medium Access Control (MAC) layer coordinates access to the channel among the nodes to prevent them from talking over each other. The 802.15.6 standard coordinates nodes through beacon periods or what is essentially time synchronization. As the network starts up, a central node lets other devices join and ensures that they are synchronized to the same clock. Once all the nodes in the network are synchronized, they are assigned one of many frames (time slots) within a beacon period. Only the node assigned to that frame can transmit packets during that time while all others go to standby. The standard also describes a channel hopping scheme for added reliability.

The network itself takes on a star topology. At the center is the network coordinator or 'hub'. 'Nodes,' which lie at the periphery of the network, only talk directly to the hub. If a node needs to pass data to another node, the hub acts as a gateway

between the two. However, the primary purpose of a sensor network like this is data aggregation and a star topology makes intuitive sense. Data is collected from all the nodes to a central location and transmitted to the cloud or some other off-body computation resource through a single access point [1].

## 2.2 Network Topology

As mentioned earlier, a mesh network forms the backbone of the Human Intranet. The mesh topology carries with it unique implications on both the physical placement and logical function of network components. These implications make mesh networks well suited to a number of real world applications such as city-wide wifi, factory operations monitoring, and customer tracking in retail stores [5, 20]. We passed over the star topology used in the IEEE 15.6 standard in favor of a mesh to facilitate our overall goals for the Human Intranet; distributed computing, robust networking, and low energy consumption.

In a typical mesh network, node placement can be dynamic and relatively unstructured. A node may be in range of one or more neighbors and can route data to any of its neighbors. Logically a mesh network may be 'full' or 'partial' where the former is achieved only if each node maintains a direct link with every other node. For practical reasons, a partial mesh is implemented in a majority of applications.

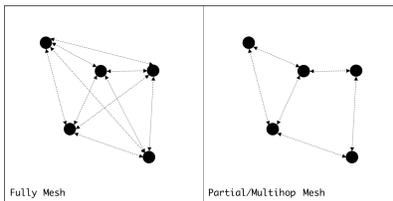


Figure 2-1: Illustration of Fully and Partial Mesh Topologies

A star network on the other hand adheres to a stricter hierarchical node placement where leaf nodes are placed radially around the central node. Logically only the central node acts as a router. Each leaf node communicates solely with the central node regardless of physical proximity to other nodes in the network. Typical use cases

include wired Ethernet networks, WiFi, and Bluetooth.

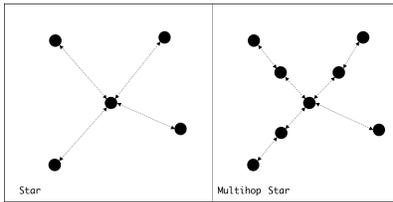


Figure 2-2: Illustration of Regular and Multihop Star Topologies

Unlike most WSNs where the primary function of the network is data aggregation, the Human Intranet is geared towards data processing and sharing. This translates into an emphasis on point to point connections rather than source to sink connections. As such the flat, non-hierarchical structure of a mesh topology where nodes have multiple neighbors suits the purpose of the Human Intranet far better than a star topology.

A mesh topology also outperforms a star topology in the HI context with regard to robustness. A robust network can continue to perform basic functions despite a few individual link or node failures. In practice, robustness is achieved through a combination of redundancy and resiliency. As the partial mesh in Figure 2-1 illustrates, there are redundant paths between each node. Most routing schemes can make use of the redundancy in a mesh network to bypass individual link or node failures. Furthermore, a decentralized mesh network protocol can also prevent any one node failure from having an outsized effect on overall network functionality. In a mesh topology this would involve nodes collaborating locally to quickly determine new routes. A star topology however does not benefit from the same robustness due to its two-tier hierarchy. Although a star network can sustain multiple leaf node failures, there is no redundancy for the central node since it mediates all communication.

A mesh topology is also better suited to implement multihop routing. Multihop refers to the use of one or more intermediate nodes as relays between two distant nodes rather than establishing a direct link between them. While the primary benefit of multihop in the HI context is energy savings, implementing multihop routing in a mesh network also enables design tradeoffs between transmission power, network

range, and throughput.

A multihop network has the potential to reduce power consumption among all the nodes in the system. Since the path loss to distance relationship is superlinear, shortening the distance between hops theoretically leads to overall savings in transmission power. Only the reception power of the device limits the number of hops until the technique no longer yields an advantage [23]. When applied to a star network, the same principle works although with certain complications. While multihop reduces overall network power, in a star topology the central hub becomes a power hot spot [4] and can adversely affect network lifetime. In fact a study [4] on multihop and single-hop star networks found that relaying could resolve the increased energy burden closer to the central node. In other words, the strict star topology was broken to resemble a mesh network.

A multihop network may also be optimized for maximum network coverage. In a typical urban mesh network, using longer hops at high power can improve SNR and channel capacity [8]. For BANs, a multihop route with shorter, low power hops can connect non line-of-sight nodes which may be placed on opposite sides of the body.

One disadvantage to increasing the number of hops on a route is the increase of latency and decrease in throughput. For each additional node in the route, the packet must be decoded and coded. The node may also need to process the packet to determine the next hop in the route. A multihop star network also introduces unnecessary latency through its tree-like hierarchy. Nodes that may be physically proximate must travel several hops back and forth through the hub to communicate. Certain links may also be shared by several multihop routes. This can lead to a decrease in the effective throughput for each route.

# Chapter 3

## Protocol Considerations

With the background above, we have the basis to make informed design decisions about the rest of the protocol. The goals as ever remain increasing network lifetime, maintaining robustness, and enabling the distributed nature of the Human Intranet. With the network topology set, the two remaining decisions involve the MAC layer and Routing layer. We know that the path loss in BANs makes transmissions costly and certain links unreliable. The multihop mesh topology also requires a routing strategy that can make use of all available routes.

### 3.1 MAC Layer

The MAC layer logically sits on top of the PHY and assumes reliable methods for sending and receiving raw bits. The next challenge involves establishing the logic that nodes within the network will use to arbitrate control of the medium. Design decisions made in the MAC layer can also have a significant impact on the energy consumption of the device and the latency. One concrete way to reduce energy consumption involves preventing collisions. Collisions occur when multiple nodes transmit simultaneously and result in the reception of garbled packets. Another typical energy reduction method involves minimizing the time spent listening in favor of waiting in standby. However because this protocol is intended for efficient wake-up style devices that can listen without penalty [15], we focus instead on minimizing collisions and

transmissions. There are two main approaches to the MAC layer in BANs; Time Division Multiple Access (TDMA) and Carrier Sense Multiple Access (CSMA).

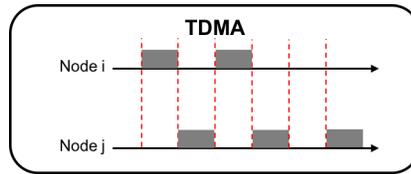


Figure 3-1: Two nodes participating in a TDMA network

In the most basic TDMA setup, each node is allocated a time slot in which it can transmit data. When a node is not transmitting it listens to the medium during specific time intervals to receive packets. Before the devices in the network can participate in this timely manner, they must follow a synchronization and scheduling process. The most basic TDMA network has a single device that acts as a coordinator. As the coordinator discovers devices, it must accommodate them with a time slot in each frame and then forward the updated schedule to the rest of the devices. The coordinator also periodically sends beacons for all nodes to stay synchronized [3]. WirelessHART, ZigBee, and SmartMesh are protocols that implement TDMA or more advanced time synchronization in mesh networks.

For a multi-hop mesh network, however, this simple TDMA setup is inefficient because there is a lack of spatial reuse. In other words, devices that are practically isolated from each other should not be barred from transmitting at the same time [6]. Additionally the coordinator will take longer to schedule and synchronize the network since it is not in direct contact with each device. To resolve this issue, the coordinator can be replaced by a distributed algorithm to find neighborhoods of nodes that are spatially disjoint and converge on a feasible schedule as information about neighborhoods propagates through the network [6]. In the HI this sort of scheme would impede the seamless integration of another sensor network and more importantly allocate a fixed bandwidth to each node leaving limited flexibility to handle irregular traffic patterns. A node that infrequently requires higher bandwidth must either request timeslots to accommodate its peak channel usage or request additional timeslots only when necessary. The first method reduces available timeslots for other

nodes. The second has the downside of added latency while the network converges on a new schedule for every request.

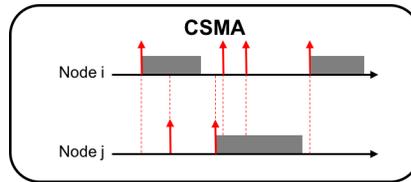


Figure 3-2: Two nodes participating in a CSMA network

CSMA/CA, the alternative, is often called a random access MAC. In this scenario, a device simply checks to see if the medium is busy before transmitting. A node may use Clear Channel Assessment if the PHY provides it or note the received signal strength indicator (RSSI) and timestamp of the last received packet. If the device determines that medium is busy from either of the two metrics, it waits for a backoff period before transmitting. The backoff period is usually designed to be twice the maximum time for a recognizable signal to traverse the distance between two nodes. If there is a collision, the device uses random exponential backoff to decide the duration of the next backoff.

As discussed earlier, many mesh sensor network protocols pass over CSMA in favor of TDMA. But in the HI context, CSMA has some unique merits. The weak and changing channels of a BAN would lead to convergence issues in a distributed TDMA scheduling algorithm. With CSMA, this dynamic network is handled on a case by case basis. A CSMA MAC also allows the flexibility for devices to send bursts of data without a change in any sort of schedule. Although TDMA outperforms CSMA/CA in heavy traffic situations, CSMA/CA does not lag far behind for networks with fewer than 10 nodes [28]. Additionally, a number of small but powerful extensions to the basic CSMA algorithm exist for further improvement of a working prototype [21, 7]. Simulations of CSMA and TDMA networks also show that while TDMA achieves higher packet delivery rates, a CSMA based network can achieve a longer lifetime [17].

## 3.2 Routing Layer

The partial mesh topology and dynamic nature of the body area environment requires a routing strategy that can improve the robustness of the network. As discussed earlier, that entails making use of the network's redundancy in a timely manner. Conventional routing strategies can be classified either as proactive or reactive and take different approaches to route discovery and maintenance.

Proactive routing involves a periodic evaluation of the network using some metric to decide which routes are ideal. Reactive route uses similar metrics to determine a route at the time of packet delivery. Proactive routing faces the same challenges as TDMA scheduling algorithms since the network is assumed to be dynamic and variable.

Reactive routing is more promising since the routes are determined as frequently as the network changes. Perhaps the simplest form of reactive routing is controlled flooding. In this approach, each hub simply rebroadcasts packets that they receive but are intended for other addresses. Each packet comes with a few fields to ensure a finite number of retransmissions. The first field is a set of flags to indicate the history of the packet. Nodes that retransmit a packet first mark them and will drop them if seen again. This promotes packet traffic to travel away from the source. The second field is a hop count which places a limit on the number of retransmissions. The limit is based on the hops needed to traverse the longest span of the network.

With these limits in place, the controlled flooding algorithm can operate with scarce energy resources. Further rules can help to create a hybrid of flooding and point to point routing methods. A cache of recently seen packets and fuzzy acknowledgements are both effective ways of limiting retransmissions while capitalizing on the natural overhearing that happens in wireless networks [24]. Finally simulations of a similar protocol setup (CSMA, Flooding) showed packet delivery rates well into the 90% range [17].

# Chapter 4

## Prototype Implementation

The prototype network we built for this project was intended to test and quantify the performance of one potential approach to the Human Intranet; specifically a multihop mesh network with a CSMA/CA MAC layer and controlled flooding for routing. The hardware was chosen for ease of use and flexibility while the software was implemented to be both portable and transparent.

### 4.1 Hardware

Since the focus of the project was evaluating the protocol, we looked for off the shelf hardware solutions that were flexible, low power, and easy to bring up. Eventually we chose the Texas Instruments CC2650STK SensorTag. The SensorTag board combines the CC2650 2.4 GHz band wireless micro-controller with a number of sensors and peripherals that were particularly useful for the prototype.

The SensorTag board has sensors that would be common in Body Area Networks; altimeter, light sensor, humidity sensor, microphone, thermal, motion. These sensors aided in data generation while we were testing the network under regular usage patterns. The board also integrates JTAG and UART interfaces, both of which were useful for debugging. In addition to the standard CR2032 battery terminals, the board also has easily accessible terminals which we used for measuring power consumption.

The CC2650 MCU itself offers a configurable proprietary protocol that we used exclusively in our implementation. The proprietary radio command API allows for custom preamble and payload lengths, choice of modulation, variable bitrate, variable transmission power, and optional whitening and CRC calculation. We opted to use GFSK modulation at a 2.4 GHz center frequency and 50 kHz frequency deviation for a 200 kbps bitrate and the standard BLE 1M PHY for the 1 Mbps bitrate. We also made use of multiple transmission power levels between 5 dBm and -21 dBm). The CC2650 also supports other common standards such as Bluetooth Low Energy and Zigbee. The Zigbee stack in particular would be useful when comparing our custom protocol with TDMA based protocols.

The SensorTag was also very well documented. The hardware datasheet clearly laid out details for power consumption, modulation characteristics, and system architecture. Most importantly the time spent on software bring up was minimized due to the abundance of example projects, documentation of TI-RTOS, and active online forum with support from actual TI engineers. Finally the simulations in [17] served as a point of comparison to the prototype implementation since both are based off of the same radio specifications.

## 4.2 Firmware

The protocol was implemented in C and made use of the Texas Instruments Real Time Operating System (TI-RTOS). Despite the proprietary nature of TI-RTOS, programs written on top of it can be ported to similar RTOSs like FreeRTOS. TI-RTOS also comes with comprehensive documentation and online support which proved invaluable during development. Like many RTOSs, TI-RTOS provides a Task based concurrency model along with a number of message passing and synchronization primitives.

Overall the firmware follows an Active Object design methodology. Each active object encapsulates a state machine and responds to events or messages from other objects. The firmware makes use of Tasks to divide the code into modules corresponding to the three logical layers of the protocol; the MAC layer, routing layer,

and application layer. Each Task responds to Events which may be generated internally or as the result of a message via the Mailbox primitive. Figure 4-1 illustrates the organization of the firmware and the interconnections between modules.

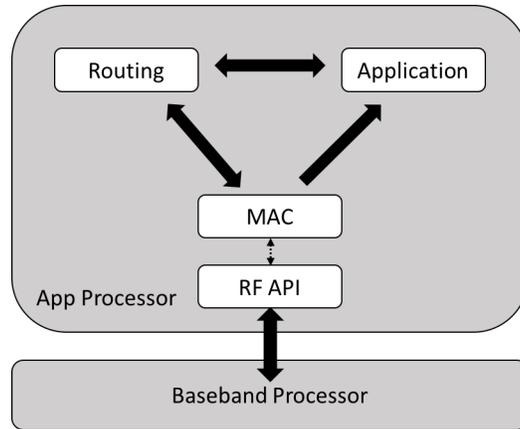


Figure 4-1: Firmware Architecture

All of the modules written for our protocol execute on the Cortex M3 application processor. The MAC task, which houses the CSMA logic and manages access to the radio, makes use of an API to communicate with the baseband processor. Since the radio control and protocol logic reside on separate processors, the radio can operate asynchronously from the other tasks. The MAC task does however use a semaphore to manage the radio resource and ensure transactions are serviced one at a time.

The state diagram in Figure 4-2 fully describes the functionality of the MAC layer. When idle, the task sets the radio to receive mode indefinitely. Since most of the radio's time is spent listening, we made use of the API's asynchronous receive command to allow the scheduler to preempt the MAC task. When the receive command triggers a callback there are a number of different paths. Malformed packets are ignored. Packets addressed to other nodes are delivered to the flooding task via a Mailbox and the MAC task returns to its previous state. Alternatively, packets addressed to the receiving node are passed directly to the application process. Immediately after the reception of a valid data packet, the node responds with an acknowledgement (ACK) packet. On receipt of an a ACK packet, the MAC task marks the end of a transaction by freeing the radio semaphore and indicating success to the higher layers.

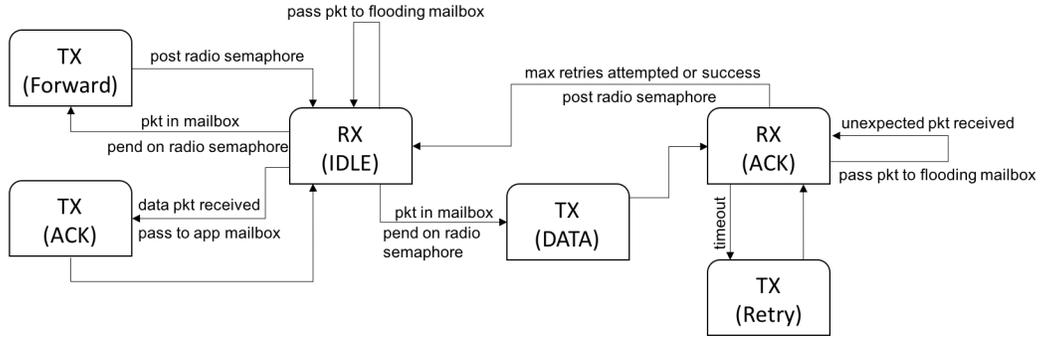


Figure 4-2: State Machine of MAC Task

When the MAC layer receives a packet from the flooding layer it is either starting a new transaction or acting as a relay for other nodes. In both cases, the packets are sent to the baseband processor for transmission and the radio semaphore is decremented. If the node was simply acting as a relay, the radio will return to receive mode and release the semaphore. If the node started a new transaction, it listens for an ACK for a finite duration before timing out. Timing out triggers retransmission of the data packet. If retransmission does not yield an acknowledgement, the application layer is notified of the failure and the radio goes back to receive mode.

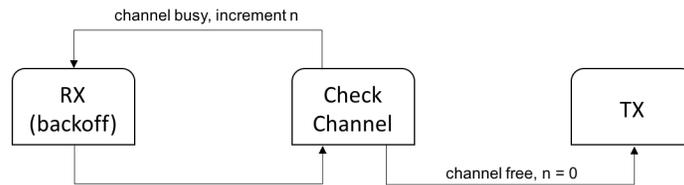


Figure 4-3: State Machine of CSMA Routine

Before all transmissions, excepting ACK packets in direct response to a valid data packet, the MAC layer runs through a CSMA routine. If the channel is determined to be busy through examination of the last RSSI and timestamp, the MAC layer sleeps for the duration of a backoff period. The MAC layer implements random exponential backoff, which means after  $n$  successive collisions the MAC layer sleeps some random amount between 0 and  $2^n - 1$  backoff periods before attempting to transmit. We used the True Random Number Generator provided by the RTOS to select the number of backoff periods and we limit  $n < 4$ . Parameters such as the RSSI threshold and

initial backoff period were informed by findings from reports on beaconless CSMA/CA implementations in IEEE 802.15.4 [13, 14].

The flooding layer interfaces with both the MAC layer and application layer. If a new data packet arrives in the transmit Mailbox, it first initializes the hop count and flooding flags before handing the packet to the MAC layer. If the flooding layer receives a packet from the MAC layer, it inspects the flooding fields to decide whether to forward the packet or not. If the packet should be forwarded, it is updated and passed to the MAC layer. Otherwise the packet is dropped.

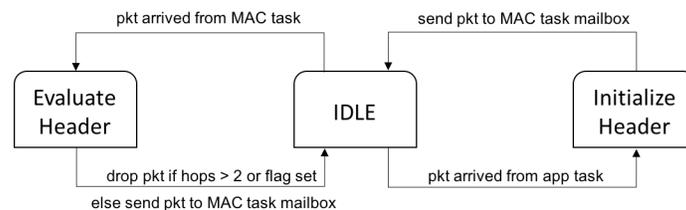


Figure 4-4: State Machine of Flooding Module

The application layer does not have any fixed functionality and was modified a number of times to accommodate different tests. The interfaces with the other tasks however did remain constant. A Mailbox lies between the application layer and flooding layer to facilitate the transfer of packets between the two tasks. The interface with the MAC layer is uni-directional as the application layer only receives packets. For the various tests, we used interrupts generated by timers or sensors to kick off packet transmissions.

### 4.3 Debugging

Debugging on an RTOS is always tricky and especially so when multiple targets are involved. We took a number steps to ensure the system worked as a whole. Each layer was implemented sequentially starting with the MAC layer. Using the JTAG connection to debug the target allowed us to step through each state of the state machine and ensure variables and flags were being set appropriately. In addition to setting watchpoints and breakpoints, we utilized the UART pins to attach an LCD

to each SensorTag. This allowed us to quickly identify obvious bugs by displaying relevant packet statistics such as the RSSI, sequence numbers, source and destination addresses, number of retransmissions, ACKs received, etc.



Figure 4-5: LCD Debug Screen via UART

Once all the tasks were written and basic functionality was verified, we started testing the network with 3 or more nodes. As we added more nodes into the network, memory usage and timing became more relevant to the debugging process. We switched to statically allocated memory to obtain more deterministic memory usage and to avoid stale pointers. Generally, run-time debugging proved informative but ill-suited to actually diagnosing problems in real time. We did use the Runtime Object Viewer in Code Composer Studio to inspect the state of semaphores, tasks, interrupts, and events. This usually allowed us to understand what might have caused deadlock or stack overflows. We also wrote to the system buffer to construct a trace of events for nodes that were not connected to a debugger.

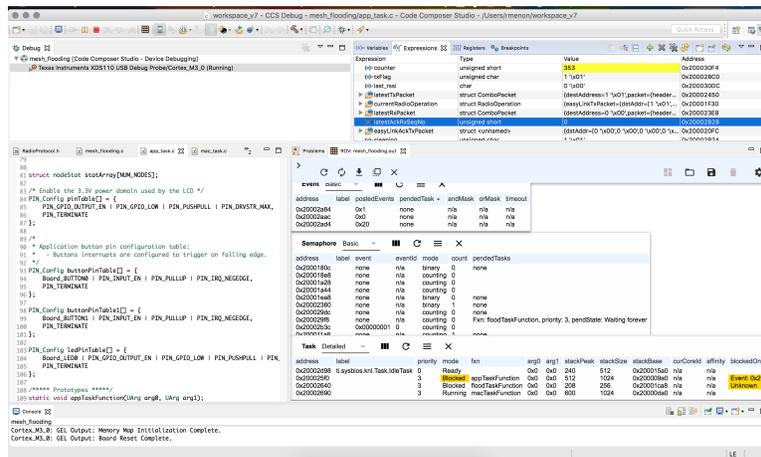


Figure 4-6: ROV View and Debug Window in Code Composer Studio

Both of these approaches however placed a burden on the target and noticeably slowed down execution at times. One alternative method we employed was to program another radio as a sniffer, so that we could obtain a trace of all packets sent in the network without affecting the timing of certain operations. Finally we also referenced the TI online forums and conversed with engineers to quickly understand which design decisions were causing odd behaviour. In fact, the forums allowed us to discover a bug in the RF API which would have been extremely difficult to diagnose without knowledge of the inner workings of the baseband processor.

# Chapter 5

## Testing

In order to evaluate the protocol and understand its performance in context, we focused on both link level and network level measurements. The following sections describe individual test setups, results, and potential caveats.

### 5.1 Power Baseline

We first set to measure the baseline power consumption of the CC2650 radio during transmission at different power levels as well as reception and standby. To record the power consumption, we used a Keithley 2612 Sourcemeter remotely operated via GPIB. A simple task written in the Keithley Test Script Builder sets the voltage source to 3.3 V and samples the current draw at 200  $\mu$ s intervals. Depending on the bitrate, transmissions lasted between 5 ms and 204  $\mu$ s. As a result we needed to adjust the sourcemeter's number of power line cycles (NPLC) setting to 0.001. The NPLC controls the AC noise integration time and directly affects how long the device waits before sampling the signal. Smaller time intervals equate to larger error bounds and the measurements below have a tolerance of 0.5%.

The firmware was modified in a number of ways to ensure that the overhead of context switching and computation would not add variability to the results. During the transmission tests, only the MAC task is active and a timer kicks off transmissions at 500  $\mu$ s intervals. CSMA routines and all receive commands were also disabled so



measure standby power, we put the MAC task to sleep for 2 seconds and forced the RTOS to remain in the idle task. The power consumption at standby was 11.3 mW on average. To measure RX power, we simply set the radio to receive indefinitely. After subtracting the standby power from the total power consumption, we found that the estimated 18.3 mW (29.6 mW - 11.3 mW) for the radio RX alone was slightly below the projected 19.5 mW (5.9 mA at 3.3 V) from the datasheet [18]. Similarly the measurements for TX at 5 dBm and 0 dBm marginally outperformed the datasheet.

One interesting takeaway from these results was that below 0 dBm, the transmission power consumption falls below reception power. This impacts our ability to measure the effectiveness of our protocol. Our main contention in choosing CSMA over TDMA involved eliminating the overhead of maintaining synchronization between nodes. While a TDMA network may have more transmissions, the time spent actively listening is drastically cut down. By contrast, the power consumption of the CSMA protocol running on the CC2650 will be dominated by reception even if we are able to use lower transmission powers. Thus, in anticipation of a wake up style radio that can actively listen without penalty, recording the transmission power level and number of transmissions would allow us to better evaluate our protocol in the context of a realistic HI implementation.

## 5.2 Link Level

After getting an idea of the power usage for each node across a number of parameters, we characterized the individual links in the network. In the following tests, we used consistent node placement to mimic the optimized network in [17]. Our topology is limited to the front plane of the body and uses five nodes as can be seen in Figure 5-3. Each node was placed in a rubber housing and attached to the body with a velcro strap. The SensorTag PCBs were approximately 1.5 mm above the surface of the skin and were oriented similarly as well. We should also note that the tests were conducted in a typical lab environment with many WiFi and Bluetooth devices around.

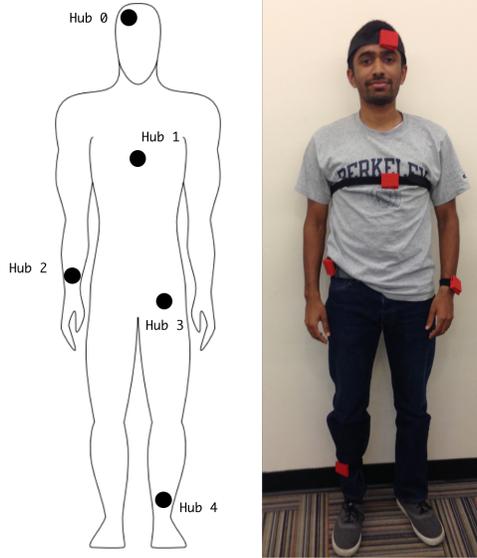


Figure 5-3: Placement of Hubs in Human Intranet Mesh Network

### 5.2.1 Channel Estimation

We measured the attenuation of each link in the mesh network by averaging the Received Signal Strength Indicator (RSSI) from every packet at the destination hub. Each entry in tables 5.1, 5.2, 5.3, 5.4, and 5.5 represents the average from 4 trials of 10,000 packets each. During the tests, we adopted a posture similar to what is shown in Figure 5-3; an upright stance with arms by the sides. We also took measurements in both directions in case of any significant asymmetry in the channel.

		Destination Hub				
		0	1	2	3	4
Source Hub	0		-60.6 dBm	-56 dBm	-65.6 dBm	-63.6 dBm
	1	-57.6 dBm		-63.6 dBm	-68.6 dBm	-59 dBm
	2	-61 dBm	-61.3 dBm		-74.6 dBm	-64.6 dBm
	3	-65 dBm	-64.3 dBm	-76.6 dBm		-66.3 dBm
	4	-63 dBm	-59.3 dBm	-69.6 dBm	-65.3 dBm	

Table 5.1: Average RSSI of Packets Sent at 5 dBm

Looking at the link quality across multiple transmission powers, it is interesting to note that the network maintains a fully connected mesh topology until the power drops to -15 dBm. At this power level and below most of the received packets fall under the typical receiver sensitivity rating of -81 dBm [18]. One metric that illustrates

		Destination Hub				
		0	1	2	3	4
Source Hub	0		-57.48 dBm	-66.45 dBm	-64.38 dBm	-65.77 dBm
	1	-62.93 dBm		-67.55 dBm	-72.86 dBm	-69.30 dBm
	2	-68.69 dBm	-68.43 dBm		-79.28 dBm	-70.97 dBm
	3	-70.40 dBm	-71.37 dBm	-82.06 dBm		-74.69 dBm
	4	-68.35 dBm	-63.19 dBm	-64.14 dBm	-76.33 dBm	

Table 5.2: Average RSSI of Packets Sent at 0 dBm

		Destination Hub				
		0	1	2	3	4
Source Hub	0		-72.44 dBm	-78.95 dBm	-74.27 dBm	-82.59 dBm
	1	-75.15 dBm		-78.94 dBm	-79.08 dBm	-80.03 dBm
	2	-77.32 dBm	-83.10 dBm		-86.67 dBm	-85.31 dBm
	3	-80.90 dBm	-81.99 dBm	-92.73 dBm		-79.80 dBm
	4	-76.76 dBm	-75.57 dBm	-74.54 dBm	-85.55 dBm	

Table 5.3: Average RSSI of Packets Sent at -9 dBm

		Destination Hub				
		0	1	2	3	4
Source Hub	0		-81.41 dBm	-85.56 dBm	-84.56 dBm	-85.57 dBm
	1	-80.75 dBm		-86.93 dBm	-86.01 dBm	-83.79 dBm
	2	-83.68 dBm	-90.18 dBm		-94.25 dBm	-90.94 dBm
	3	-87.22 dBm	-82.68 dBm	-96.30 dBm		-79.80 dBm
	4	-83.80 dBm	-81.64 dBm	-83.22 dBm	-91.25 dBm	

Table 5.4: Average RSSI of Packets Sent at -15 dBm

		Destination Hub				
		0	1	2	3	4
Source Hub	0		-86.08 dBm	-91.79 dBm	-88.97 dBm	-94.16 dBm
	1	-88.98 dBm		-90.93 dBm	-93.18 dBm	-92.96 dBm
	2	-89.66 dBm	-97.09 dBm		-101.30 dBm	-98.76 dBm
	3	-90.44 dBm	-96.24 dBm	-100.81 dBm		-86.98 dBm
	4	-90.18 dBm	-88.79 dBm	-87.82 dBm	-97.20 dBm	

Table 5.5: Average RSSI of Packets Sent at -21 dBm

the effect of transmissions below the receiver sensitivity is Packet Delivery Ratio (PDR). PDR is the ratio of packets received at the destination to packets sent from the source and it represents the probability that a packet will reach its destination. At -21 dBm, most links achieved PDRs between 70% and 80%, while at a TX power of -15 dBm they remained well over 95%.

On the individual link level, there were also unexpected findings. Despite their relative proximity, the link between the wrist and waist suffered from major path loss. Compared to any other link with a terminal on the waist, that link consistently experienced 10 to 15 dB more loss. Even the link from the head to ankle, which spans the largest distance across the body, had higher average RSSIs. It is clear that in certain configurations, distance between hubs matters less than maintaining an absolutely clear line of sight path. The most likely explanation for the extreme path loss between hubs 2 and 3 is that during testing the wrist remained slightly behind the front plane of the torso, causing the signal to travel around the circumference of the torso [2].

As the transmission power decreased, we also noticed asymmetry in the directionality of certain channels. The link from hub 4 to 3 suffered an additional loss of 6, 11, and 10 dB for TX powers of -9, -15, and -21 dBm respectively. Similar effects are visible in the link from hub 2 to hub 4, and from hub 2 to hub 1.

### 5.2.2 Multihop

Next we focused on power consumption of the nodes as function of the number of hops in the network and the packet delivery rate. We started by placing a node on the head and another on the ankle to create the link that spanned the greatest distance. The originating node at the ankle transmitted a fixed number of packets at 50 ms intervals while the node at the head listened and responded with acknowledgement packets. During the transmission, both nodes were battery operated and stationary. The test was repeated twice more, with a third node placed by the chest and a fourth node at the waist. As the number of nodes increased, the transmission power for the network was decreased to the lowest feasible power level which could achieve a PDR greater than 90%.

As is clear from table 5.6, the use of multihop routes allows for a decrease in transmission power and overall power consumption per node. In the 4 node path, each node saves approximately 10 mW per transmission compared to the 2 node path. While the total consumption of the path is lower for the 2 node path (65 mW vs 100.1 mw),

# Nodes	Tx Power	Packets Sent	Packets Received	PDR
2	0 dBm	1000	998	99.8%
3	-9 dBm	1000	957	95.7%
4	-18 dBm	1000	989	98.9%

Table 5.6: Results of Multihop Test

we get the added benefit of extended network lifetime and more computation.

## 5.3 Network Level

Our final set of tests evaluates the network as a whole. Nodes were placed at the locations specified in figure 5-3 and the network was tested under a variety of traffic patterns and placed in both static and dynamic environments. Each node recorded a few key statistics that would help understand the network’s performance in each scenario. As with the link level tests, nodes will record the number of packets sent and received from each address in order to calculate the Packet Delivery Ratios. In the network context, the PDR can be correlated with the throughput along each route. Another important metric is the number of retransmissions at the MAC layer. A high number of retransmissions coupled with a high PDR can indicate a lower effective throughput. The final statistic recorded at each node was the number of duplicate packets received. We used this to evaluate the efficiency of the controlled flooding routing layer.

### 5.3.1 Comparison to Network Simulations

The optimized networks presented in [17] serve as a relevant comparison point to the protocol presented in this thesis. Moin et al. used the Castalia network simulator along with a mixed integer linear program to evaluate a number of possible network configurations. Some of the highest performing networks used 4 to 5 nodes placed at similar positions as in Figure 5-3. All simulations were based off of the same CC2650 radio used in our prototype and operated at 1 Mbps and one of three distinct power levels (0 dBm, -10 dBm, and -20 dBm). The mesh network simulations followed a

particular traffic pattern where each node sent a 100 byte packet every 100 ms in a round robin fashion. We performed the same tests at varying power levels for comparison. For each power level, the CSMA threshold was modified to reflect the findings from the link level testing and retransmissions at the MAC layer were turned off to match the simulations. All the results below are the average of 3 runs lasting 5 minutes each.

Tx Power	Network	Hub 0	Hub 1	Hub 2	Hub 3	Hub 4
0 dBm	91.09	99.79	99.66	99.38	89.09	67.09
-9 dBm	81.16	72.35	98.62	99.15	98.28	37.31
-21 dBm	46.80	34.74	32.05	33.77	61.68	71.75

Table 5.7: PDRs from Round Robin 1 kbps Network Test

The results from the prototype network match the trends indicated in [17]. At 0 dBm the network achieved a PDR greater than 90%. After lowering the transmission power to -9 dBm, the network maintained a similar level of performance on some runs but averaged a PDR of 81.16%. Just as the simulations found, at -21 dBm the PDR drops off significantly and levels out around 40%.

Along with lower PDRs, lower transmission power also led to greater variability. At 0 dBm all three runs yielded PDRs within 2% of the mean. That figure dropped to 13% and 28% for -9 dBm and -21 dBm respectively. The increased variability also applies to PDRs for individual nodes. For all 6 runs above -21 dBm, hub 4 consistently had the lowest PDRs. The network was more or less unpredictable at the lowest power setting. Hubs that achieved PDRs of 80% in one run dropped to 40% in others.

### 5.3.2 Realistic Human Intranet Traffic Pattern

The last test was a simulation of a realistic use case. In this setup we assigned a unique datarate to each node. The scenario involves an ECOG sensor relaying control signals to a neuroprosthesis. The neuroprosthesis responds with temperature and accelerometer data as part of a feedback loop. A smartwatch sends EMG readings and blood oxygen readings to the smartphone which acts as a general sink for sensor

data. The chest node would naturally aggregate ECG and glucose level data to display on the watch face. Table 5.8 shows the assignment of nodes to sensor and datarate profiles that were used in the test.

Node #	Sensor	Datarate	Signal Frequency	Dest Node
0	ECOG	10 kbps	100 Hz	4
1	ECG, glucose	10 kbps	50 Hz	2
2	EMG	1 kbps	100 Hz	3
3	N/A	N/A	N/A	N/A
4	temperature, accelerometer	10 kbps	10 kHz	0, 2

Table 5.8: Parameters of Human Intranet Simulation [12, 9, 19]

In contrast to the previous tests, this use case is more realistic due to the datarates and limited destinations for each packet. Most sensors in our test (Glucose, Temperature, EMG) report data much slower than the 10 packet per second speeds in the simulation. However a few sensors such as the accelerometer and ECOG generate more data more quickly. In an HI implementation, each hub would also only address a subset of the other hubs in the mesh. Unlike the round robin, this traffic pattern can put higher stress on certain routes and can potentially monopolize individual links.

Tx Power	Network	Hub 0	Hub 1	Hub 2	Hub 4
0 dBm	99.89	99.93	99.87	99.87	99.88
-9 dBm	96.65	97.44	98.73	95.53	94.93
-21 dBm	92.13	88.03	94.91	95.96	89.65

Table 5.9: PDRs from Mock HI Network Test: Stationary

Overall the network performed quite well despite hubs 0 and 4 which dispatched packets 2 to 4 times faster than the other hubs. In fact due to the overall decrease in network traffic, the performance did not drop off as transmission power decreased. All network PDRs remained above 90% as did most hub PDRs.

We also tested the network in an 'active' scenario where we ran on a treadmill for the duration of the recordings. This was intended to model the varying path loss that can be present in a BAN. Despite the constant movement of hubs 2 and 4,

Tx Power	Network	Hub 0	Hub 1	Hub 2	Hub 4
0 dBm	99.71	99.87	99.75	99.44	99.77
-9 dBm	90.28	85.63	82.37	95.93	97.18
-21 dBm	87.57	79.3	91.38	95.94	83.67

Table 5.10: PDRs from Mock HI Network Test: Active

the effects on PDR were minimal and didn't necessarily correlate with transmission power. In fact the results from this test indicate that the protocol can be reliable even at -21 dBm.

# Chapter 6

## Conclusion

The Human Intranet is a rapidly nearing future. While the variety of sensors may not exist today, the appetite for continuous monitoring of our bodies is evident and the infrastructure for the technology is being built. We have shown that a Body Area Network intended for local computation and storage rather than data aggregation can benefit from our simple yet effective protocol. With CSMA and controlled flooding, our mesh network prototype achieved Packet Delivery Rates up to 91% under heavy traffic. Additionally we found that under normal loads, the network can utilize the multihop mesh topology to lower transmission power to -21 dBm while still being reliable. We also matched the trends found in network simulation, giving us a basis to compare our prototype to TDMA star networks.

Future directions for this protocol are numerous. First the MAC layer could be improved by employing a spread spectrum approach based on node location. The MAC could also employ adaptive CSMA thresholds and transmission power levels based on the history of success for individual links. For the routing layer we would test more sophisticated reactive protocols like AODV and RPL as well as location or resource based addressing.

With the existing protocol, we could also integrate real sensors in place of data generation. Another option would be to port the code to custom wake-up style hardware to observe the energy benefits of extremely low power consumption when the node is idle and receiving. Finally, we could perform a direct comparison to

TDMA style protocol like 802.15.4 on the same hardware to evaluate the distinct advantages each method has.

# Bibliography

- [1] IEEE Standard for Local and metropolitan area networks - Part 15.6: Wireless Body Area Networks. *IEEE Std 802.15.6-2012*, pages 1–271, Feb 2012.
- [2] Muhammad Mahtab Alam, Elyes Ben Hamida, Dhafer Ben Arbia, Mickael Maman, Francesco Mani, Benoit Denis, and Raffaele DErrico. Realistic simulation for body area and body-to-body networks. *Sensors*, 16(4):561, 2016.
- [3] SIG Bluetooth. Specification of the bluetooth system-version 4.0. *Bluetooth SIG*, 2010.
- [4] Bart Braem, Benoit Latre, Ingrid Moerman, Chris Blondia, Elisabeth Reusens, Wout Joseph, Luc Martens, and Piet Demeester. The need for cooperation and relaying in short-range high path loss sensor networks. In *Sensor Technologies and Applications, 2007. SensorComm 2007. International Conference on*, pages 566–571. IEEE, 2007.
- [5] Joseph Camp, Joshua Robinson, Christopher Steger, and Edward Knightly. Measurement driven deployment of a two-tier urban mesh access network. In *Proceedings of the 4th international conference on Mobile systems, applications and services*, pages 96–109. ACM, 2006.
- [6] Petar Djukic. *Scheduling algorithms for TDMA wireless multihop networks*. University of Toronto, 2008.
- [7] Amre El-Hoiydi and Jean-Dominique Decotignie. Wisemac: An ultra low power mac protocol for multi-hop wireless sensor networks. In *ALGOSENSORS*, volume 4, pages 18–31. Springer, 2004.
- [8] Martin Haenggi. Twelve reasons not to route over many short hops. In *Vehicular Technology Conference, 2004. VTC2004-Fall. 2004 IEEE 60th*, volume 5, pages 3130–3134. IEEE, 2004.
- [9] Mark A Hanson, Harry C Powell Jr, Adam T Barth, Kyle Ringgenberg, Benton H Calhoun, James H Aylor, and John Lach. Body area sensor networks: Challenges and opportunities. *Computer*, 42(1), 2009.
- [10] J Heikenfeld. Sweat sensors will change how wearables track your health. *IEEE Spectrum*. October, 22, 2014.

- [11] Andreas Inmann and Morten Haugland. Implementation of natural sensory feedback in a portable control system for a hand grasp neuroprosthesis. *Medical engineering & physics*, 26(6):449–458, 2004.
- [12] Yasser Khan, Aminy E Ostfeld, Claire M Lochner, Adrien Pierre, and Ana C Arias. Monitoring of vital signs with flexible and wearable medical devices. *Advanced Materials*, 28(22):4373–4395, 2016.
- [13] Tae Ok Kim, Hongjoong Kim, Junsoo Lee, Jin Soo Park, and Bong Dae Choi. Performance analysis of ieee 802.15.4 with non-beacon-enabled csma/ca in non-saturated condition. In *EUC*, volume 6, pages 884–893. Springer, 2006.
- [14] Benoît Latré, Pieter De Mil, Ingrid Moerman, N Van Dierdonck, Bart Dhoedt, and Piet Demeester. Throughput and delay analysis of unslotted ieee 802.15.4. *Journal of networks*, 1(1):20–28, 2006.
- [15] R Liu, J Naghsh Nilchi, Y Lin, TL Naing, and CT-C Nguyen. Zero quiescent power vlf mechanical communication receiver. In *Solid-State Sensors, Actuators and Microsystems (TRANSDUCERS), 2015 Transducers-2015 18th International Conference on*, pages 129–132. IEEE, 2015.
- [16] Pranav Mistry and P Maes. Sixth sense: integrating information with the real world. <http://www.pranavmistry.com/projects/sixthsense>, 20(1):0, 2010.
- [17] Ali Moin, Pierluigi Nuzzo, Alberto L Sangiovanni-Vincentelli, and Jan M Rabaey. Optimized design of a human intranet network. In *Proceedings of the 54th Annual Design Automation Conference 2017*, page 30. ACM, 2017.
- [18] CC2650 Simple Link Multi-Standard. 2.4 ghz ultra-low power wireless mcu. *Texas Instruments datasheet*.
- [19] Maulin Patel and Jianfeng Wang. Applications, challenges, and prospective in emerging body area networking technologies. *IEEE Wireless communications*, 17(1), 2010.
- [20] Stig Petersen and Simon Carlsen. Wirelesshart versus isa100. 11a: The format war hits the factory floor. *IEEE Industrial Electronics Magazine*, 5(4):23–34, 2011.
- [21] Joseph Polastre, Jason Hill, and David Culler. Versatile low power media access for wireless sensor networks. In *Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 95–107. ACM, 2004.
- [22] Jan M Rabaey. The human intranet: where swarms and humans meet. In *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition*, pages 637–640. EDA Consortium, 2015.

- [23] Jan M Rabaey, M Josie Ammer, Julio L Da Silva, Danny Patel, and Shad Roundy. Picoradio supports ad hoc ultra-low power wireless networking. *Computer*, 33(7):42–48, 2000.
- [24] A Rahman, W Olesinski, and P Gburzynski. Controlled flooding in wireless ad-hoc networks. In *Wireless Ad-Hoc Networks, 2004 International Workshop on*, pages 73–78. IEEE, 2004.
- [25] Laurens Roelens, Sara Van den Bulcke, Wout Joseph, Günter Vermeeren, and Luc Martens. Path loss model for wireless narrowband communication above flat phantom. *Electronics Letters*, 42(1):10–11, 2006.
- [26] Helder D Silva, Jose A Afonso, and Luis A Rocha. Body attenuation and path loss exponent estimation for rss-based positioning in wsn. *Wireless Personal Communications*, 94(3):835–857, 2017.
- [27] Sunil Srinivasa and Martin Haenggi. Path loss exponent estimation in large wireless networks. In *Information Theory and Applications Workshop, 2009*, pages 124–129. IEEE, 2009.
- [28] Sana Ullah, Daehan Kwak, Cheolhyo Lee, Hyungsoo Lee, and Kyung Sup Kwak. Numerical analysis of csma/ca for pattern-based wban system. In *Biomedical Engineering and Informatics, 2009. BMEI'09. 2nd International Conference on*, pages 1–3. IEEE, 2009.