Distance Estimation and Positioning Based on Bluetooth Low Energy Technology

JOHAN LARSSON

KTH ROYAL INSTITUTE OF TECHNOLOGY
INFORMATION AND COMMUNICATION TECHNOLOGY

Distance estimation and positioning based on Bluetooth low energy technology

Master of Science Thesis
by Johan Larsson
Kista 2015

Examiner
Mark Smith, KTH

Supervisor
Patrik Hillbo, Syntronic

Abstract

This thesis deals with location determining, also known as positioning, using Bluetooth Low energy radio. The goal is to implement a low power low cost indoor positioning system which utilize existing hardware. Two main methods are investigated and their viability assessed.

Previous efforts with indoor positioning systems concentrate on statistical fingerprinting methods, mainly using 802.11 (WLAN) as the platform. Some efforts have been made with purely signal strength based positioning, but indoor environments have shown to work unfavorably for these kinds of methods.

Contents

List of Figures **2**

List of Acronyms **5**

1 Introduction **6**
1.1 Introduction 6
1.2 Background 6
1.3 Motivation 7
1.4 Goals 8
1.5 Organization 8

2 Positioning and radio theory **9**
2.1 Bluetooth Low Energy 9
2.1.1 Radio interface 10
2.1.2 BLE and positioning 11
2.2 Positioning concepts 12
2.2.1 Proximity 12
2.2.2 Ranging 12
2.2.3 Angulation 13
2.2.4 Lateration 13
2.3 Fingerprinting 15
2.3.1 Beacon distance considerations 17
2.4 Multipath fading 17
2.5 Available positioning systems 18

3 Implementation **19**
3.1 Hardware platform 19
3.2 Software 19
3.2.1 Receiver 19
3.2.2 Beacon 20
3.3 Algorithm 20

4 Test setup **22**
4.1 Ranging 22
4.2 Fingerprinting 22
4.2.1 Fingerprint density calculation 23

5 Evaluation **25**
5.1 RSSI Distribution 25
5.2 Fingerprinting accuracy 26
5.2.1 Fingerprinting - Nearest neighbor 27

6 Conclusion **29**

7 Discussion **30**
7.1 Future works and use of BLE for positioning 30
7.1.1 Algorithm tuning 30
7.1.2 Beacon placement strategy and density 30
7.1.3 Determine useful fingerprint density 30
7.2 BLE Proximity based positioning 30

List of Figures

2.1 Major components of a Bluetooth 4.0 dual mode chipset 9
2.2 Channel configuration of BLE 10
2.3 Principle of triangulation 13
2.4 Trilateration based locationing system 14

4.1 Drawing of area used for measurements. Each fingerprint is marked with a red dot. The distance between the points are 1m 23

5.1 Distribution of RSSI readings from CSR1000 25
5.2 Distribution of RSSI readings from TI keyFOB 26
5.3 Distribution of RSSI readings from a Bluegiga beacon 26
5.4 Distribution of RSSI readings from a Bluegiga beacon, second series 27

GPS Global Positioning System
TOF Time Of Flight
TOA Time Of Arrival
TDOA Time Difference Of Arrival
BLE Bluetooth Low Energy
RSSI Received signal Strength Indication
UWB Ultra Wideband
GSM Global System for Mobile Communications
ISM Industrial, Scientific and Medical
DQPSK Differential Quadrature Phase Shift Keying
8DPSK 8-way Differential Phase Shift Keying
BR Basic Rate
EDR Enhanced Data Rate
IEEE Institute of Electrical and Electronics Engineers
GFSK Gaussian Frequency Shift Keying

Chapter 1

Introduction

1.1 Introduction

As cost of low power radio systems steadily decreases, their use in both cost and performance sensitive applications become more attractive. One such application is positioning. The problem of positioning has largely been solved for outdoor situations, with systems such as the Global Positioning System(GPS)[1]. Indoors however, the situation is another as GPS signals generally don't penetrate buildings. In this thesis, the possible use of Bluetooth Low Energy(BLE)[2] to position cargo and similar slow moving or semi static devices indoors will be studied. The thesis will be conducted at Syntronic Software Innovations AB(SSI) and based on work of implementing an indoor positioning system.

1.2 Background

Plenty of work has already been done regarding indoor positioning, most of it is however focused on IEEE 802.11[3], more commonly known as WiFi. While many successful results have been achieved, either they rely on a radio standard that is not low power(such as 802.11) or they rely on non commodity hardware as is the case of Motetrack[4] or Ultra Wideband(UWB)[5] based systems. For Bluetooth, much of the initial work showed little promise[6] [7], as many parts of the standard work unfavorably for use in positioning. The two largest problems are the accuracy of the clock and the implementation of the Received Signal Strength Indication(RSSI) value. The clock allows for a 1µs deviation according to the specification [2], this makes the expected accuracy very low and no article investigating the use of time based positioning was found at the time of writing. The RSSI problem seems to relate to both hardware implementation and specification. One study [6] found the correlation between distance and RSSI to be poor, while [7] found it usable but with limited accuracy.

Other studies focused on the concept of fingerprinting. The concept uses RSSI, but not to estimate distance. Rather, the whole facility of interest is covered with beacons that constantly advertise their presence. The device to be positioned then reads the signal strength of all received advertisements, and compares them to a pre-compiled database of locations, selecting the best match. Using fingerprinting techniques, accuracy of less than 3m can be achieved with 80% precision [8]. Even though some papers seem conclusive that fingerprinting is the best choice for Bluetooth positioning, it is important not to rule out other methods if performance can be made good enough. They also focus on Bluetooth, not BLE, which might have different properties making it more useful for other methods such as RSSI based ones. No studies using BLE were found.

1.3 Motivation

Several competing competing commercial indoor positioning systems exist on the market, but there is no clear standard. The vendors are usually not specific when describing their product and it is often not clear what kind of method is used or what accuracy is to be expected. Most of them however indicate they are utilizing 802.11-standards, making them inherently not low power. The common denominator for all these are that they focus on positioning people, either for navigational purposes or location specific information such as commercial offers at retail locations or tour information at museums.

Table 1.1 shows a small selection of positioning methods available for indoor use, with GPS as a reference. The numbers are rough and systems based on similar designs may vary in both power consumption and accuracy. The power figures quoted are when active and not directly comparable between each other. As some technologies allow for very short bursts of activity, the average power will be considerably less than a system with similar active state power but long setup times.

What can be found is that there is a gap for something with sub 10m accuracy that is truly low power. Two options look particularly promising, UWB and BLE based systems. Both should be able to fill the role for indoor positioning, albeit with different cost and accuracy expectations. For this thesis, the focus is solely on BLE.

BLE is expected to grow rapidly as it is deployed in mobile phones and wearable electronics. This promises many benefits from economy of scale, such as low cost and high availability of devices. It will probably not be able to compete with UWB for accuracy and resolution[5] but considering cost and availability BLE looks to get a more widespread use.

System	Target area	Active power usage	Accuracy
GPS and similar satellite based	Outdoor	50-200mW	3-15m[1]
Cellular proximity	Both	50mW	300-3500m
802.11 Fingerprinting	Indoor	50-300mW	3-10m[8]
802.11 Time differential lateration	Indoor	50-300mW	30cm
UWB Time differential lateration	Indoor	50-60mW	30cm
BLE proximity	Indoor	$<$10mW	10-50m
Bluetooth RSSI	Indoor	$<$40mW	5-10m
Bluetooth Fingerprinting	Indoor	$<$40mW	3-10m

Table 1.1: Small overview of positioning technologies using radio

1.4 Goals

The goal of this thesis project is to study the viability of using BLE for indoor positioning for static or slow moving objects. The central questions for the project are:

- What or which methods are suitable for implementation using BLE?
- What are the relative strengths and drawbacks of the methods?
- How can indoor positioning with BLE be further improved upon?

At the end of this project a functioning prototype positioning system is expected to have been implemented and evaluated. The system is expected to be able to reproduce some of the results found by others. The system should implement fingerprinting as it is the most complex signal strength based method commonly deployed. During development, measurements and tests mirroring other methods will be performed, covering the major positioning methods. The results will be analyzed according to it's accuracy, expected power usage for a large scale system and required beacon density for the desired accuracy. Problems faced during implementation and evaluation will also be discussed if they are relevant to the systems viability for deployment.

The system is expected to be able to position a device at a resolution of 5m and offer a battery life of at least a year for battery powered nodes, using low cost batteries.

1.5 Organization

- Chapter 2 introduces both common and project specific concepts used when discussing and implementing positioning
- Chapter 3 explains the implementation of the positioning system
- Chapter 4 describes the tests performed and how they are set up
- Chapter 5 presents the results from fingerprinting and ranging measurements. Also offers some discussion about their implications
- Chapter 6 summarizes the results and discuses future work and areas of interest
- Chapter 7 offers discussion about BLE based positioning and it's applications

Chapter 2

Positioning and radio theory

This chapter introduces the relevant parts of the Bluetooth 4.0 standard, as well as some radio concepts useful for understanding positioning with it's possibilities and limitations.

2.1 Bluetooth Low Energy

Introduced in the Bluetooth 4.0 standard, BLE is a low power radio standard from Bluetooth Special Interests Group. While maintaining Bluetooth as part of its name, it's not fully backwards compatible with Bluetooth 3 and earlier. With Bluetooth 4.0, the standard specifies two different radio interfaces, one that is compatible with the older standard, also known as Basic Data Rate(BR) or Enhanced Data Rate(EDR). The other radio interface is the one used for BLE. The specification allows for Bluetooth 4.0 devices to be dual mode, as well as single mode with support for either BR/EDR or BLE.

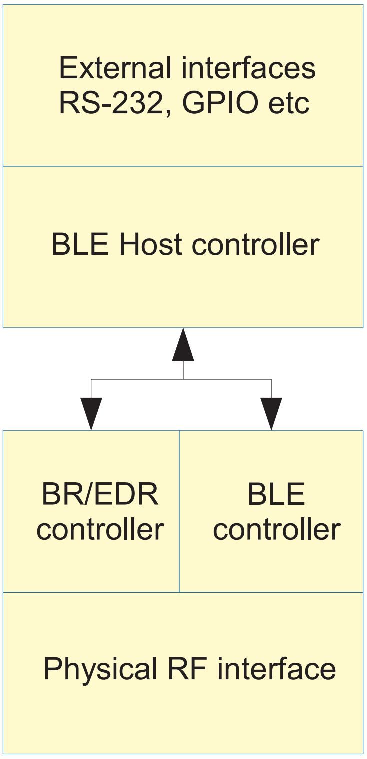

Figure 2.1: Major components of a Bluetooth 4.0 dual mode chipset

2.1.1 Radio interface

One of the largest deviations from the old[9] Bluetooth standard is the radio interface. BLE is still in the 2.4 GHz ISM band, but instead of being split up into 79 channels 1 MHz wide, BLE uses 40 channels each one 2MHz wide. To reduce scanning time the allowed broadcast channels are reduced to three, as illustrated in figure 2.1.

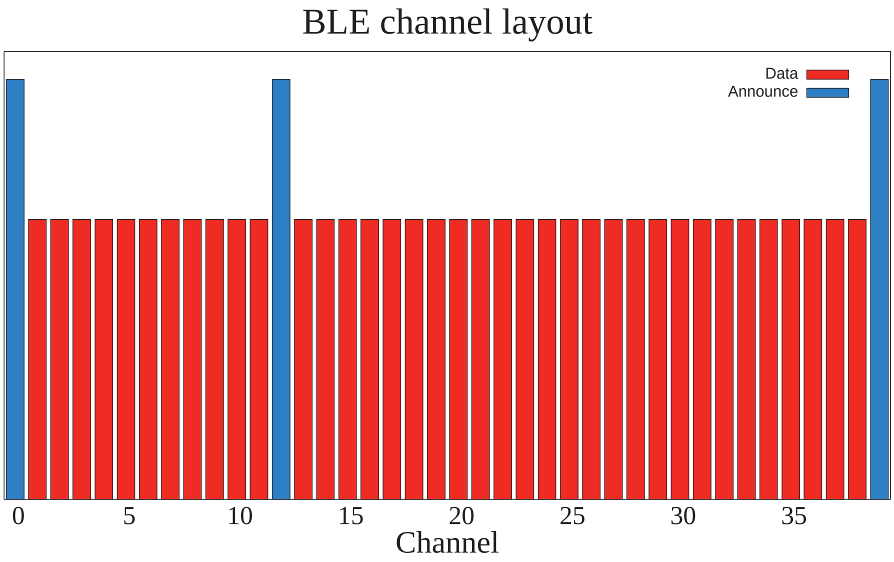

Figure 2.2: Channel configuration of BLE

The modulation scheme used is Gaussian Frequency Shift Keying(GFSK), the same as used in the low speed data mode of Bluetooth. However, the channel bandwidth differs, making interoperability impossible at the physical layer.

BLE is also capable of waking from sleep mode, sending a packet and then going back to sleep quicker than higher speed Bluetooth. Where BLE uses roughly 6ms, older standards would use more than 100 ms. Not only does this affect scanning time, it also serves as an optimization for battery life when sending small amounts of data.

2.1.2 BLE and positioning

The way that services are exposed on BLE devices is new and several profiles for services are included in the standard itself. This includes two service profiles relating to positioning, Find Me Profile(FMP) and Proximity Profile(PXP).

FMP is the "Find Me Profile", with the goal to make devices easier to locate if lost. It works by having two devices, one is the locator and the other is the target to be found. The target takes the role as server, meaning it listens for requests from the locator. When a request arrives, an alert of some kind is activated on the target device. The alert is in most cases a visual or auditory signal to make the device easier to find. It could also be used to trigger a broadcast or an attempt for the device to position itself and report back with the estimated position. This opens for the possibility to use it in ways not exactly as intended, but as part of a positioning system.

PXP or the Proximity Profile is the other location based service included in the standard. It defines the behavior of a device when it loses or establishes a connection with another device. There's also the possibility to take action on crossing a pre determined path loss/RSSI value. This service is more interesting for a positioning system, but it relies on the devices establishing connections. This means that both PXP and FMP suffer the same problem, the device to be found needs to constantly listen to be a functional part of the system. Both profiles might be useful in applications where battery life is not as important or regular recharging of the batteries can be expected.

There is also a commercial product/system called iBeacon from Apple computer, that aims to provide locationing service with compatible devices. This is not a standard profile, but many mobile phones and computers are expected to support it. This is a two way system, in some ways very similar to PXP. The beacons are expected to be able to respond and provide some data to a local application at the device being positioned. With this data delivery functionality and requirements of special software on the devices, this system is wider in scope than a low power positioning system. The target market has been specified as commercial services for improved shopping experience. As such, it is a positioning system albeit too big to consider for a low power purely positioning service. However, if it gains traction this will likely bring down the prices of BLE modules, indirectly benefiting other systems as well.

One consideration with practically all BLE enabled devices is the directionality of the antenna. Small devices utilize chip antennas and they exhibit non isotropic radiation patterns. Connecting an external antenna to such devices will likely not be an option in the majority of positioning scenarios due to size constraints. Part of the problems arising from antenna directionality could be compensated for by using relative signal strengths instead of absolute ones when using multiple signal strengths to determine position.

2.2 Positioning concepts

There are several important core concepts when discussing positioning and a brief introduction to them follows.

2.2.1 Proximity

The simplest form of positioning is to check whether or not the object to be positioned is present in radio coverage, with no regard to RSSI or time of arrival. The resolution offered is the same as the radio coverage, as it would only be able to differentiate between two states, present and not present. To increase accuracy it is possible to combine several radios at different positions and to utilize the overlap of their coverage areas to determine the most likely position of the object.

2.2.2 Ranging

Ranging is the process of determining the distance between two objects, without any regard to angle. It is also the base used for lateration(described in the next section). Several methods to accomplish this exists and the two most common in radio ranging are based on either received signal strength or time of flight. In free space, the relation between signal strength and distance can be expressed as

$$d = 20^{(RSSI+a)}$$

Where a is an offset based on the maximum received signal strength for the system. However, for indoor use, this expression is at best a poor approximation as indoor environments are filled with obstacles such as walls and furniture. the obstacles are the source of multipath effects that makes the relation between distance and RSSI complicated, making accuracy low[6]. Using time of flight based methods have proven more useful for indoor ranging[10][5][11] as multipath effects interfering with signal strength have less effect on measurements. The trade off is the cost of implementation, RSSI is a measurement that is integrated in virtually all modern radio controllers, but accurate time measurements require specialized hardware[10][5] as many popular radio controllers do not offer an accurate timestamping function[2], or even an accurate clock required in the hardware. BLE is one of the standards that lack accurate timestamping.

2.2.3 Angulation

Angulation is the process of determining where an object is using the angle of incoming signals. To do so, some level of directionality in the transmitter and/or receiver is needed, such as directional antennas or antenna arrays. For the simple case of one measurement, no distance information can be gained. With two measurements from objects with known coordinates the location of the third object can be found. This is referred to as triangulation. Illustrated in the figure below

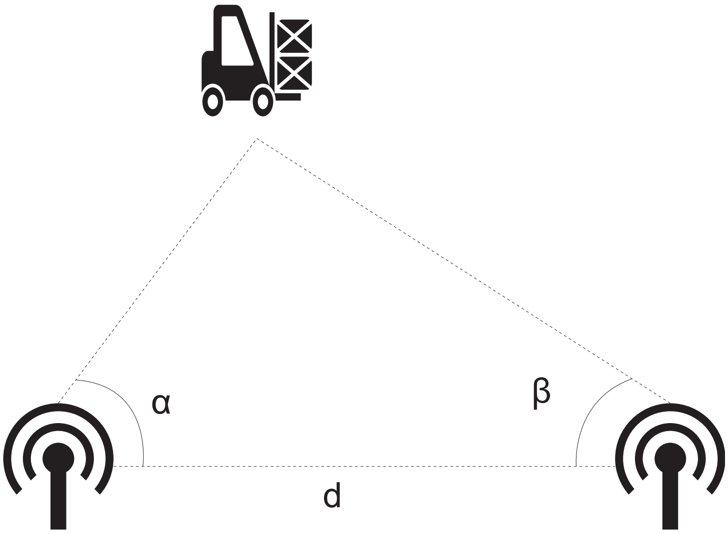

Figure 2.3: Principle of triangulation

2.2.4 Lateration

Lateration is the determining of one objects position using several range measurements. If the range measurements are done from several known positions, each measurement can be seen as a circle or sphere of possible positions at the measured range. Combining several such measurements and calculating the intersection of the measurements yields a position. The principle is illustrated below with three beacons, all at different distance from the object.

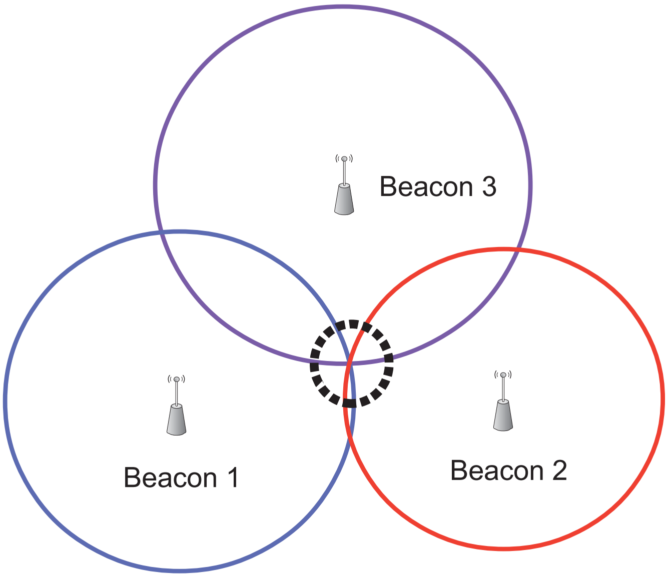

Figure 2.4: Trilateration based locationing system

2.3 Fingerprinting

Fingerprinting is a method based on using primarily RSSI data to identify the current location using pre-recorded data. The data consists of RSSI readings of different radio signals available in the area. The readings are done in the mobile device, the beacons don't need to be aware of any mobile device. In practice these signals consist of base station announcements from cellular networks or WiFi as well as broadcasts from beacons deployed for the purpose of enabling fingerprinting.

For indoor cases it is not practical to utilize external sources such as cellular networks. Even though it might seem attractive not needing to deploy any infrastructure for a positioning system, as the cellular networks are under the control of telecom operators there is no guarantee that the cell configuration would stay the same for any period of time making the recorded fingerprints obsolete at unpredictable intervals. Cellular networks also cover large areas, making the fluctuations in RSSI not as prominent as other shorter range radio systems. The smaller variations in values makes it harder to separate fingerprints from each other with high confidence.

WiFi already exists in many indoor locations and is a prime candidate for fingerprinting. It has indeed been extensively studied[3] and deployed, but the usefulness of WiFi-based systems have been limited by the size and cost of WiFi radios. WiFi fingerprinting has proven successful and is used for fingerprinting both as separate smaller commercial systems and as parts of larger systems in the case of Google Maps. There are also non commercial alternatives available such as RedPin.

If WiFi is not an option, an alternative that have had some success is to install special beacons specifically used for generating signals used for fingerprinting. The benefit of doing so is the added possibility to choose a radio technology, making it possible to tailor the system to a locations specific needs. These could be frequency interoperability restrictions, range, power consumption or many other parameters. This is where new low power radio standards become relevant, it is possible to deploy cheap and low power units to make up the fingerprint map. At the same time, they would offer the possibility to operate in a true low power mode, enabling them to run on small battery cells for years. This is of high importance for maintenance, as it could otherwise become a significant part of the systems total cost. Other benefits of the low power nature includes freedom of placement and lower installation cost, as no extra wiring is needed saving both material and labor costs.

2.3.1 Beacon distance considerations

If beacons are placed dense to achieve a higher data content in the fingerprints, they also start to provide the less unique information thus adding extra hardware but no extra benefit. Considering the extreme, where two beacons are placed on top of each other, the signal strength would be identical to any receiver within their coverage. Calculating the amount of redundant information would be possible if the radio signal propagation indoor could be easily modelled, however, this is not the case as has been showed several times before[7][6][3][12]. Therefore, this issue is not considered when designing the experiments to evaluate a fingerprinting system, but taken into account when evaluating it as a possible problem source.

2.4 Multipath fading

Multipath, or more accurately multipath propagation effects, make up one of the biggest problems for indoor radio communications. Multipath effects are caused by the same signal arriving from multiple directions at the same time. Because every path the signal propagates have different channel characteristics, it might interfere with another copy of itself at the receiver, both in a constructive and destructive way. Another name for destructive interference is fading. The different paths originate from the environments physical layout and materials, making the effects complicated to model as even small changes in dimensions or a change of material will have a large impact on the paths channel properties.

The effect is further amplified when the environment consists of objects that are not static, such as people, pallets and packets. Such dynamic environments cause several problems when modelling them. One is that a signal might have it's line of sight blocked by some object, attenuating the most direct path's signal. Other more subtle problems are the different characteristics of fading introduced for different kinds of obstacles. Some materials are prone to reflecting radio waves, while others simply attenuate it. As it is virtually impossible to identify a material and determine its properties in real time, multipath effects are better compensated for in a statistical sense.

To overcome this problem, simplifications are generally used and no effort to accurately model the radio environment is made. The effects should be most pronounced when using RSSI based systems, but they also affect TOF based ones, albeit not as greatly[11]

2.5 Available positioning systems

As can be seen in 1.1, even when looking at a small group of available positioning systems their accuracy varies by orders of magnitude. Positioning is also not limited to radio based systems, acoustic and optical ones are available as well

Chapter 3

Implementation

3.1 Hardware platform

Common for all tests and experiments performed are the hardware platforms used. A mix of hardware from two vendors was used, Texas Instruments(TI) and Cambridge Silicon Radio(CSR). A total of six beacons and one receiver were used. The TI beacons are based of the CC2541[13] chipset and the additional CSR beacons are based on CSR1000[14]. The TI beacons consisted of one CC2540-KEYFOB[15] reference module from TI themselves and three BLE112 [16] modules from BlueGiga. The TI module was pre-assembled and used as is, the BlueGiga ones however were slightly modified, adding larger battery and a breakout for the programming connector. The two CSR modules have part number CSR1010 and was used as is. The modifications done to the BLE112 modules adds significant amounts of metal in the form of batteries. This has a high probability of changing the antenna characteristics, both directionality and attenuation. As the exact characteristic of the antenna is not part of this study, the impact of any such effects are not taken into further consideration.

To receive the signals a TI CC2540 USB Evaluation Module Kit based on the CC2540 chipset was used. The receiver was connected to a personal computer via an USB extension cord to give flexibility in beacon placement, as well as moving it away from the laptop which otherwise might act as an RF shield.

3.2 Software

3.2.1 Receiver

The main part of the software resides in the the receiver controller. The receiver itself runs a very small C-program that passes raw BLE packets between the computer it is attached to and the BLE host controller. The controller program acts as a virtual serial port, which the program on the computer interfaces with. Packets are sent as raw streams of data and needs to be decoded in software before being useful.

To maximize code recycling a small library was written to handle the basic functions needed in all the experiments, such as initializing the host controller, initializing the radio and enabling/disabling of scanning. Limited support to properly encode and decode packets were added, but was quickly abandoned as it proved time consuming and there was no real benefit for the kind of positioning system implemented here. The library also contains classes to handle fingerprints, fingerprint maps and beacons.

The fingerprint map contains an array of fingerprint objects and the means to compare them for likeness. This class might be considered a singleton, as it only makes sense to have one map active at any one time.

The fingerprint object contains the RSSI readings of the broadcast signals from the beacons used while creating it. The beacon is stored with its name and its broadcast ID. The broadcast ID is simply a small message sent with each broadcast as to identify the beacon. Using a user selectable key like this rather than something unique like the device Universal Unique IDentifier (UUID) or serial number should make it easier to replace failing beacons in the field.

For the initialization part, a so called magic packet has to be sent to the BLE controller for it to wake up. This was not documented in the Bluetooth standards document, as it is left as a vendor specific procedure. It was not fully documented in the development manual for the CC2540 chipset either, instead it suggested the use of the pre-compiled BLE-stack. As the BLE-controller already was loaded with the serial pass through firmware, this was not an option. Therefore, a sample program from TI was loaded and the serial port monitored. The initialization procedure was captured to a file as a byte stream. Looking at the stream and comparing with the developers documentation the initialization packages were extracted and inserted as part of the library. As this was done without any regard to parameters sent, it is not known if the radio is operating in the same mode or with the same parameters as it would if it had been initialized by the BLE software stack provided for use in production code. As the communications worked, no further investigation of the initialization procedure was performed and no problems were found that could be traced back to it.

One central part that was common to all of the experiments was the decoding of announce packets. These packets were sent from the beacons and the library extracted the beacon name as well as RSSI value. These were the only values used in the majority of the experiments. As such, any impact of errors during initialization should impact all the experiments in an equal way.

3.2.2 Beacon

The beacons run a program that continuously broadcast their ID and then sleeps for 500 ms. The broadcast was intended to only be sent on the lowest of the three available channels. Because of the different hardware platforms used, the exact same program could not be used for all beacons, instead three different programs were written. As the complexity of the programs are low, this should not affect the overall system performance.

The idea behind the decision to use only one announce channel is to minimize the spread caused by slight variations in channel properties for different frequencies. This also lowers the fingerprint complexity, as each beacon otherwise must be considered three times, one for each frequency.

3.3 Algorithm

The first positioning algorithm implemented was a naive approach, simply selecting the fingerprinted location that had the least Manhattan-distance from the currently measured one. The Manhattan-distance is calculated according to:

$$d = \sum_{i=0}^{n} |m_n - f_n|$$

Where m_n is the currently measured RSSI value and f_n is the stored RSSI value from the time of fingerprint mapping. During mapping, 120 values were collected using 6 beacons, meaning an average of 20 values per beacon was used to calculate the average RSSI. To achieve lower response time when positioning, each point is measured collecting 24 values, or 4 per beacon, before reporting an estimated position and restarting data collection. Each measurement is then saved as a fingerprint, making it simple to retrieve it at a later moment and to try different algorithms or distance methods, such as using the Pythagorean formula to calculate distance instead of the Manhattan distance.

Chapter 4

Test setup

All measurements were carried out within an indoor office landscape environment with large open spaces and no individual offices. The beacons were placed around the perimeter of the area chosen for the experiment. The receiver was connected to a laptop and put on a rolling table cart, to ensure that the height would remain unchanged during all measurements. As the setup consisted of a laptop it means it was also battery powered enabling free movement and removing any error introduced by needing to also move a power adapter around.

The whole area to be mapped was measured in a 1m x 1m grid, the same distance the fingerprints were indexed with, with each intersection marked with a sticker for reference. These stickers were lined up with a marker placed on the rolling cart for each measurement. The stickers remained for the whole duration of the fingerprint measurements to make sure the same reference was used both for mapping and for positioning.

4.1 Ranging

The first group of tests are the ranging tests to evaluate the RSSI-distance relationship for the equipment and environment selected. They also form the foundation for the approximation of fingerprint density, described in 4.2.1. The setup had the receiver connected to both a laptop on the rolling table cart and to stationary computer with the transmitter put on the cart. In both cases, precautions were taken to ensure that the radio environment remained undisturbed during the experiment. That meant ensuring no other persons moving in the area of the experiment and completing whole runs of measurements to not be impacted by furniture rearrangement.

4.2 Fingerprinting

The fingerprinting tests were performed in a small part of an office building. A grid with a point distance of 1m was measured and marked in the area as illustrated in fig 4.1. All points were then mapped, or fingerprinted, in four major directions. Each fingerprint consists of 120 measurements, 20 per beacon. This value is then averaged and saved together with coordinate and direction for later use.

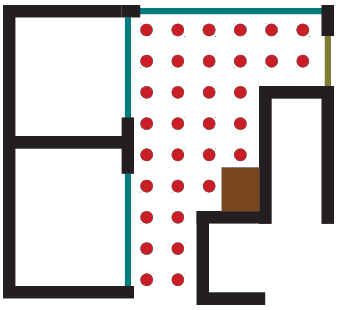

Figure 4.1: Drawing of area used for measurements. Each fingerprint is marked with a red dot. The distance between the points are 1m

4.2.1 Fingerprint density calculation

To determine the highest density of unique fingerprint nodes in a radio map, it is useful to estimate the amount of information contained in the RSSI readings that make up a fingerprint. Determining this would be of use when one is deploying or considering deploying a fingerprinting positioning system. It would give the minimum number of beacons needed, ensuring that hardware and installation costs are kept low, as well as keeping the size of the radio map small. The size of the map becomes especially important if the device running the fingerprinting algorithm has limited processing power and memory. Other benefits of having an estimate is to help determine if a fingerprinting solutions is feasible at all for some applications, given accuracy, reliability and cost constraints.

The limit in information obtainable from the RSSI readings stem from several sources, such as random noise during measurements, limited discoverable range and external RF interference. As the RSSI reading is not a central function in neither Bluetooth classic or BLE, none of the manufacturers specify the expected resolution of the reading or the noise associated with the readings. With no data on the expected error and resolution of this value from the vendor, an approximation of the resolution and error based on measurements needs to be performed. These measurements are taken from the ranging tests.

The standard deviation of the RSSI measurement is calculated from several measurements, using different transmitters at different distances. For each transmitter and distance several measurements are to be taken. From these measurements the average and approximate standard deviation should be calculated. If the behavior of the readings appear normal distributed, this model can be used to calculate a theoretical maximum density of fingerprints. This would be of use when installing such a system, as the hardware cost would be known and the accuracy could be estimated. The proposed method for calculating maximum density is iterative, starting with the lowest value:

$$q_0 - s_0 = RSSI_{\mathrm{min}}$$

Where $RSSI_{min}$ is the lowest RSSI reading possible. q_0 is the center of the first range that covers the lowest value of the available RSSI measurement range. The following ranges would be chosen as:

$$q_{\mathrm{n}+1} = q_{\mathrm{n}} + 2s_{\mathrm{n}} \ + 2s_{\mathrm{n}+1}$$

This means that the center of the next range would be chosen to not overlap with the previous range. This is repeated until the whole RSSI range is covered.

Chapter 5

Evaluation

5.1 RSSI Distribution

The first series of results were used to determine the distribution of the errors in RSSI readings. The theoretical work on RSSI resolution has been based on the assumption that the error was normal distributed, but, as can be seen and shown from the results that is not the case. Each set of measurements consisted of a total of 1000 readings. The beacon was placed at a fixed distance from the receiver and no moving objects were present within 20m during the measurements.

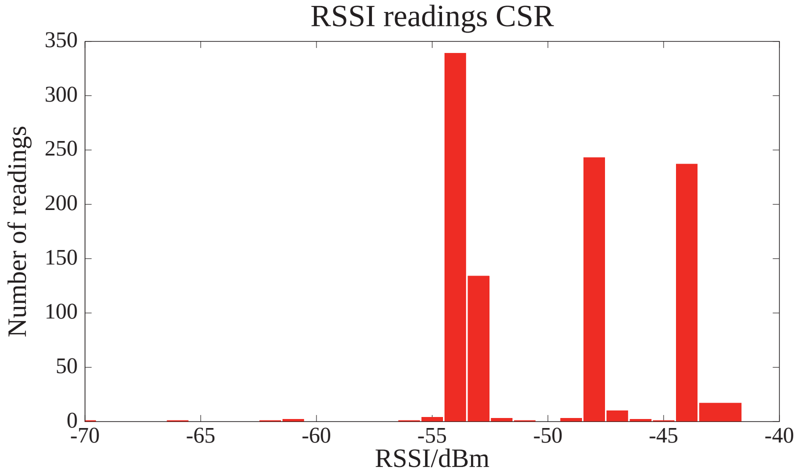

Figure 5.1: Distribution of RSSI readings from CSR1000

Visual inspection of the distribution indicates that it is not normal distributed. The three distinct readings in each measurement indicate that the distribution is not caused primarily by random noise. Even though the transmitters should only have transmitted at one frequency, they all transmitted at all three of the available broadcast channels. Since the channels differ in frequency, the channel properties, in this case attenuation, is clearly visible. This theory was further investigated and confirmed by looking at the individual data packets, all broadcasts channels were used. The varying channels do have implications for the fingerprinting system, as the readings are split between three different values where just one was expected.

The cause of the unwanted transmission was not further investigated and as there are three different programs controlling the beacons involved, it is not clear if this is a hardware limitation or a software bug. The cause might be different on the three platforms used.

The results also meant that the idea of dividing the RSSI range into smaller ranges to calculate the fingerprint density and information carried in the measurements was

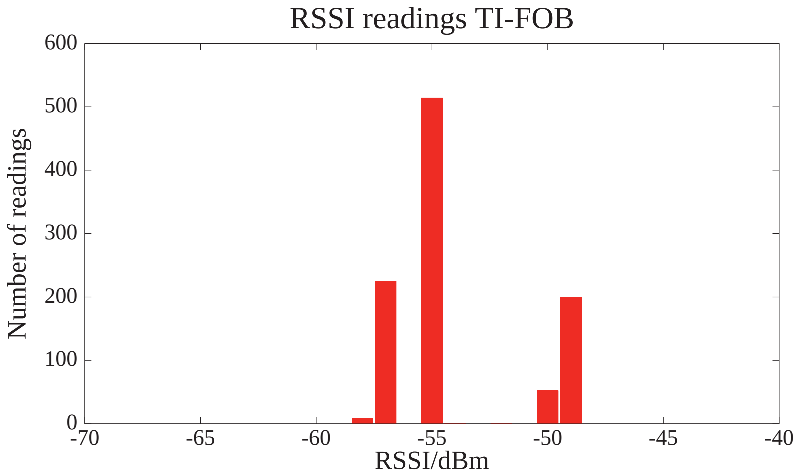

Figure 5.2: Distribution of RSSI readings from TI keyFOB

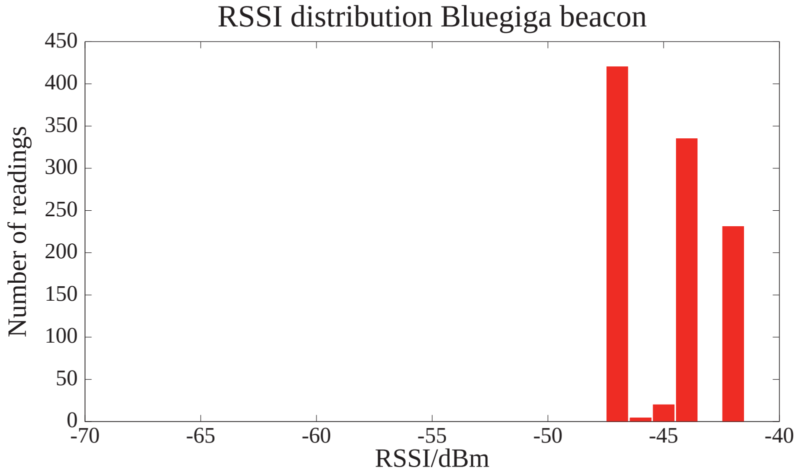

Figure 5.3: Distribution of RSSI readings from a Bluegiga beacon

no longer possible, as the distribution was not normal and varied very little between readings.

5.2 Fingerprinting accuracy

The fingerprinting data was collected in three series. Two of them was on a line, each measurement at a fingerprinted location. Those series differed in direction, but shared the locations. The third series measured a selection of n points, some of them in between fingerprint locations. The directions varied as well.

Here goes the fingerprinting results, new names for the subsections needed as well.

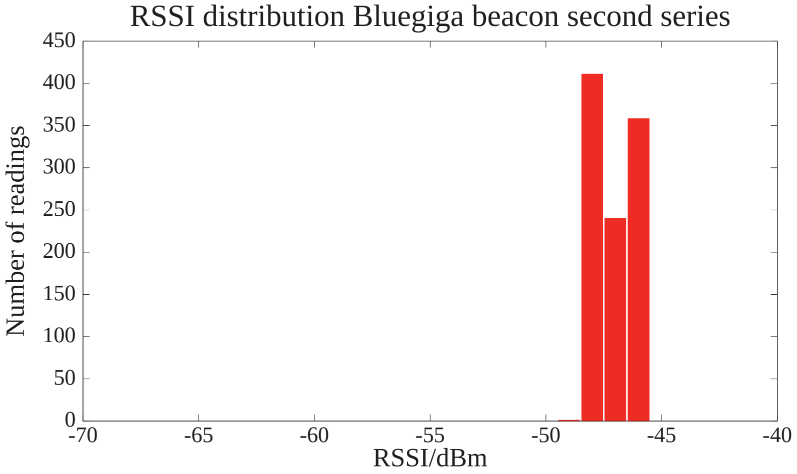

Figure 5.4: Distribution of RSSI readings from a Bluegiga beacon, second series

5.2.1 Fingerprinting - Nearest neighbor

In the most basic form, the system creates a fingerprint at the current position and then compares it to the ones stored in the fingerprint map. The one with the least manhattan distance, as described in x.x, is then chosen as the estimate. For a perfect system, the maximum error in position is the grid size of the fingerprint map divided by sqrt(2).

Three runs of tests were made. The first two runs consisted of the cart oriented in one direction and moved in a straight line while maintaining the same direction. Stops were made every 1m and 0.5m respectively. The last run was at four randomly chosen points in the grid. At each point, all four directions were tested. The results are summarized in the table below.

Test	n	Dir. hit	Pos. hit	Dir + Pos hit	Avg pos. error	Avg FP score
Line 1	13	31%	23%	15%	1.3m	14
Line 2	14	21%	36%	7%	1.4m	3
Points	16	57%	14%	14%	1.9m	9

Table 5.1: Results from nearest neighbor tests

Where n is the sample size of the test, "Dir. hit" is the percentage of correct directional estimates, "Pos. hit" the percentage of correct position estimate and "Dir + Pos hit" the combination of the two. Average positioning error expressed in meters follows, as well as the average fingerprint score. Lower score means a more likely match was found, as described in 3.3. The positional hit rate is counted as either a complete hit or a miss. A hit would be a fingerprint within 1/sqrt(2) m from the real position. The reason for using this tolerance is that the line 2 test was performed with 0.5m increments, putting half of the measurements between valid positions. For the directional accuracy, as only four possibilities exist, only exact matches count. Lastly there is the fingerprint score. The score is the calculated as the manhattan distance between the real position and the best matching fingerprint. Lower score indicates better accuracy.

The average position error is lower than expected, but as the sample size is low, the numbers should only be taken as indication as to possible performance. The conditions were made as close to ideal as possible, no persons or moving objects were permitted in the area close to the measurements. The cart was also fixed in it's layout and operated by the same person during all the tests as to not change the characteristics of the receiver. All this was done to make the experiments repeatable and have fewer unknowns, at the cost of similarity to a realistic scenario.

Not only do the accuracy results seem promising, the frequency dependence observed in the results presented in section 5.1 does not seem to impact the performance of the fingerprinting system too much. The results in 5.4 showed that a 10dB difference between the strongest and weakest channel was possible. With only 4 measurements during the test runs these differences could possibly make a big impact on the fingerprint score, which is the sole determining factor for estimated position.

Chapter 6

Conclusion

The results as presented in the previous chapter indicates that fingerprinting is viable for indoor positioning using BLE. Success was shown in a limited indoor area, although with some caveats.

The data suggests that BLE is suitable for simple positioning systems where accuracy is of low priority and low power consumption and low implementation cost is of higher importance. This is reflected in the commercial offerings that exist in the area, such as iBeacon. Such systems are expected to offer functionality similar to proximity based ones, albeit with limited ranging functionality.

With of the shelf hardware a fingerprinting system was constructed. The system as implemented worked and did provide the desired accuracy for a limited set of tests. The results are better than a purely ranging based solution and should with some improvements be able to perform similarly to WiFi based fingerprinting systems, with the potential of a lower power consumption.

Detailed power consumption could not be measured with the equipment available, but is estimated that with two AA cells rated at about 3Wh each the battery life for the beacons in a fingerprinting solution should measure in years.

Chapter 7

Discussion

7.1 Future works and use of BLE for positioning

There are several points for improvement in a BLE positioning system. For a fingerprinting system, the main points of improvement are

7.1.1 Algorithm tuning

Fine tune the algorithm for fingerprint comparison and classification or change it completely. During this project a simple method to compare fingerprints was used, while it has been shown[8] that more advanced methods yield better results.

7.1.2 Beacon placement strategy and density

Determine how dense beacons need to be placed for best performance. There is a possibility that too many beacons in a small area only adds to noisy measurements and degrades performance. It would also need to be clear before a commercial deployment to be able to calculate cost and power consumption of a full scale system. This needs to be done for different radio environments to get enough spread of the data, to determine worst case and best case scenarios.

7.1.3 Determine useful fingerprint density

Determine the practical density of fingerprint locations to be mapped. Too many fingerprints adds more data but might not necessarily increase accuracy but only increase implementation cost and the size of the fingerprint database.

7.2 BLE Proximity based positioning

There are however many downsides for fingerprinting at a larger scale. The extra accuracy gained might not be needed and simpler proximity or ranging-based systems could be used instead lower the cost of a complete system. The implementation cost of fingerprint is directly proportional to the area covered given a fixed density of fingerprints. This suggests that a pure fingerprinting system rarely would be the first choice for indoor positioning. However, utilizing some of the concepts from fingerprinting might help improve other solutions.

One such possibility is to utilize an overlapping quantized ranging model, where the RSSI from several beacons are utilized to determine the users location, similar to an overlapping proximity system, but using the fact that RSSI can be mapped to distance, albeit with low accuracy. Such a system would have practically no extra cost compared to a purely proximity based one, but still provide accuracy that could be good enough if the position only needs to be determined at room level or similar. This kind of system could also be used as a simple fingerprinting system later on, if the extra accuracy is needed. It could also be used as a hybrid system, where most of the locations are determined by the simple overlap, but some problematic locations are manually mapped and stored ahead of use as to improve accuracy.

Bibliography

[1] Department of Defense. *Global Positioning System standard positioning service performance standard*, fourth edition, 2008.

[2] Bluetooth Special Interests Group. *Bluetooth Core Specification Version 4.0*, first edition, 2010.

[3] Mikkel Baun Kjærgaard. Indoor positioning with radio location fingerprinting. *arXiv preprint arXiv:1004.4759*, 2010.

[4] Konrad Lorincz and Matt Welsh. Motetrack: A robust, decentralized approach to rf-based location tracking. In *Location-and Context-Awareness*, pages 63–82. Springer, 2005.

[5] Cemin Zhang, Michael Kuhn, Brandon Merkl, Mohamed Mahfouz, and Aly E Fathy. Development of an uwb indoor 3d positioning radar with millimeter accuracy. In *Microwave Symposium Digest, 2006. IEEE MTT-S International*, pages 106–109. IEEE, 2006.

[6] Josef Hallberg, Marcus Nilsson, and Kåre Synnes. Positioning with bluetooth. In *Telecommunications, 2003. ICT 2003. 10th International Conference on*, volume 2, pages 954–958. IEEE, 2003.

[7] Antti Kotanen, Marko Hannikainen, Helena Leppakoski, and Timo D Hamalainen. Experiments on local positioning with bluetooth. In *Information Technology: Coding and Computing [Computers and Communications], 2003. Proceedings. ITCC 2003. International Conference on*, pages 297–303. IEEE, 2003.

[8] Anja Bekkelien, Michel Deriaz, and Stéphane Marchand-Maillet. Bluetooth indoor positioning. Master's thesis, University of Geneva, 2012.

[9] Bluetooth Special Interests Group. *Bluetooth Core Soecification Version 3.0*, first edition, 2009.

[10] Lin Ji. Increasing accuracy of location determination: Exploiting phase change reconstruction and timing measurements. Master's thesis, Royal Institute of Technology, 2007.

[11] Ye et al. Experimental impulse radio ieee 802.15.4a uwb based wireless sensor localization technology. *ISSC 2011*, 2011.

[12] Alan Bensky. *Wireless positioning technologies and applications*. Artech HouseO, Norwood, MA, 2008.

[13] TI. *Texas Instruments CC2541 Data sheet*, revision d edition, 2012.

[14] Bluegiga. *BLE112 Data sheet*, fourth edition, 2014.

[15] TI. *Texas Instruments CC2540 Data sheet*, revision f edition, 2013.

[16] Bluegiga. *BLE112 Data sheet*, fourth edition, 2014.

TRITA-ICT-EX-2015:27