PhD-FSTC-2017-68
The Faculty of Sciences, Technology and Communication

DISSERTATION

Defence held on 13/10/2017 in Luxembourg

to obtain the degree of

DOCTEUR DE L'UNIVERSITÉ DU LUXEMBOURG

EN INFORMATIQUE

by

Walter BRONZI

Born on 4 December 1987 in Rome (Italy)

ENHANCING MOBILITY APPLICATIONS THROUGH BLUETOOTH COMMUNICATIONS

Dissertation defence committee

Dr. THOMAS ENGEL, Dissertation Supervisor
Professor, Université du Luxembourg

Dr. Raphaël FRANK
Research Scientist, Université du Luxembourg

Dr. Francesco VITI, Chairman
Associate Professor, Université du Luxembourg

Dr. Thomas SCHERER
Telindus S.A.

Dr. Jérôme HÄRRI, Vice Chairman
Associate Professor, Eurecom Communication Systems

January 14, 2018

Acknowledgements

This work was conducted at the SECAN-Lab of the University of Luxembourgs Interdisciplinary Centre for Security Reliability and Trust (SnT). The project has been realized in conjunction with Telindus Luxembourg through a Public-Private Partnership (PPP).

First and foremost, I would like to thank Prof.Dr. Thomas Engel, for having supported me throughout many years in his research group and for having allowed me to undertake an Internship at BMW in Munich. Second, I would like to thank Dr. Raphal Frank for his continuous and invaluable guidance for the past several years. Moreover, I would like to thank Dr. Thomas Scherer for his supervision and essential insights when they were most needed.

I would like to thank the other members of the VehicularLab especially Thierry, Sasan, Lara and Sebastien for making the work environment much more enjoyable. Also thanks to every past colleague, German, Maximilian and all the others, for the advices and support at one point or another.

A very special thank goes to my best friend and better half, Dervla Leonard, for being there through thick and thin when close and when far apart. I would thank all my family in Luxembourg, in Italy and in England for believing in me. Thanks to all my exceptional friends Ani, Yulia, Gabriela, Hamed, Maelle and most of all Martin and Mikkel for being who you are.

Without your support, family and friends, this milestone wouldn't have been possible.

Thank you.

Abstract

In the world of short and medium-range wireless technologies, Bluetooth has recently come to the forefront of innovation. Within the next five years its market presence, especially in its Low Energy variation, is expected to nearly double across all market segments. The technology is quickly and steadily gaining importance for a wide range of applications with a specific focus on Internet of Things (IoT) devices. The growing availability and variety of such devices constitute an untapped potential that we plan on exploiting. Our focus in this thesis is to understand Bluetooth's capabilities and explore its potential in mobile contexts. One specific field where this technology remains unexplored is Vehicular Ad Hoc Networks (VANETs). Because of the need to implement and moderate vehicular communications, the topic of Intelligent Transportation Systems (ITSs) is now trending more than ever.

In this thesis we propose two ways we can benefit from Bluetooth in a mobile environment. Firstly, we consider the technology as a *communication medium* to investigate how different mobilities affect the link performance between two devices. To do this, we define a set of communication experiments, in our case between two vehicles, to analyse how Bluetooth Low Energy (BLE) is affected by varying speed, distance and traffic conditions. We find that the maximum communication range between two devices can go beyond 100 m and that a robust connection, capable of handling sudden signal losses or interference, can be achieved up to a distance of 50 m. The experiments were conducted using a proof-of-concept mobile application for off-the-shelf smartphones that can be used to transmit data over multiple hops in various Vehicle-to-Everything (V2X) scenarios.

Secondly, we consider Bluetooth discovery capabilities as an *information medium* by using a connectionless approach to analyse different mobility frameworks. As there is an increasing need for vehicles and objects to become aware of their context, we implement Bluetooth as a sensing system to provide contextual information about its surroundings. Our challenge is to find out to what extent we can exploit the Bluetooth discovery and beaconing scheme for this purpose. We collect and analyse a dataset of Bluetooth Classic and BLE discoveries and evaluate their respective characteristics and ability to provide context-aware information from a vehicular perspective. By examining data recorded about encountered devices, such as quantity, quality of signal and device class information, we infer distinctive Bluetooth behaviours related to context and application. For this purpose, we propose a set a features to train a classification model to recognize different driving environments (i.e. road classes). Investigating the performance of our classifier, we were able to predict up to three classes (highway, city, extra-urban) by using only Bluetooth discovery data and no geographical information. This outcome gives promising results targeted at low energy and privacy-friendly applications and can open up a wide range of research directions.

In conclusion, in this thesis we present two ways of applying Bluetooth to mobile contexts for deploying novel human mobility applications.

Contents

Acronyms

A2DP Advanced Audio Distribution Profile.

ACL Asynchronous Connection-Less.

AES Advanced Encryption Standard.

AFH Adaptive Frequency Hopping.

AP Access Point.

API Application Programming Interface.

ATT Attribute Protocol.

AUC Area Under the Curve.

BC Bluetooth Classic.

BLE Bluetooth Low Energy.

BPSK Binary Phase-shift Keying.

CCM Cipher Block Chaining-Message Authentication Code.

CRC Cyclic Redundancy Check.

DSRC Dedicated Short Range Communication.

EBRS Event-Based Role Switching.

ECDH Elliptic Curve Diffie-Hellman.

GAP Generic Access Profile.

GATT Generic Attribute Protocol.

GFSK Gaussian Frequency-shift Keying.

GPS Global Positioning System.

HCI Host-Controller Interface.

HFP Hands-Free Profile.

HU Head-unit.

IaVC Intra-Vehicle Communication.

IFS Inter-Frame Space.

IoT Internet of Things.

IrVC Inter-Vehicle Communication.

ISM Industrial, Scientific and Medical.

ITS Intelligent Transportation System.

L2CAP Logical Link Control and Adaptation Protocol.

LE Low Energy.

LL Link Layer.

MAC Media Access Control.

MIC Message Integrity Check.

ML Machine Learning.

NESN Next Expected Sequence Number.

NTP Network Time Protocol.

OBU On-Board Unit.

ODB On-board Diagnostics.

OFDM Orthogonal Frequency-division Multiplexing.

OQPSK Offset Quadrature Phase-shift Keying.

OSM Open Street Map.

PAN Personal Area Networking Profile.

PH Personal Hotspot.

PHY Physical Layer.

RF Random Forest.

RFI Radio-Frequency Interference.

RL Reinforcement Learning.

RSSI Received Signal Strength Indication.

RSU Road Side Unit.

RTT Round-Trip Time.

SCO Synchronous Connection-Oriented.

SDK Software Development Kit.

SIG Special Interest Group.

SM Security Manager.

SN Sequence Number.

SVM Support Vector Machine.

V2I Vehicle-to-Infrastructure.

V2V Vehicle-to-Vehicle.

V2X Vehicle-to-Everything.

VANET Vehicular Ad Hoc Network.

VRU Vulnerable Road User.

WAVE Wireless Access in Vehicular Environments.

WLAN Wireless Local Area Network.

WMAN Wireless Metropolitan Area Network.

WPAN Wireless Personal Area Network.

WWAN Wireless Wide Area Network.

Chapter 1

Introduction

> The first step is to establish that something is possible - then probability will occur.
>
> *Elon Musk*

In the last decade we have observed a growing installed base of connected consumer electronics targeting all age categories. These devices have become a seemingly ubiquitous part of our lives. The most common piece of technology, for new and old generations, has undoubtedly been the smartphone. Together with tablets and smart wearables (smartwatches & smartbands), smartphones have been welcomed with open arms in households and workplaces all around the world.

These devices are increasingly equipped with a plethora of sensors and communication technologies that enables them to be used in many different scenarios. Their goal is to assist people in their day-to-day activities. Since the devices are carried by people, their applications are closely related to their mobility. It follows that, when designing and implementing wireless communication technologies, human mobility must be taken into consideration. Hence, specific applications will have different mobility profiles.

Within this world of wireless communication technologies, Bluetooth has been the protocol of choice for a variety of applications, such as cordless headsets and hands-free kits in cars. Although the technology we know as Bluetooth has been around since the 1990s, with the recent advent of its Low Energy (LE) version 4.0 and later 5.0, we were introduced to brand new and improved functionalities.

The Bluetooth Low Energy (BLE) technology was developed with a focus on devices that require a small amount of energy to run for long periods of time. As such, its main applications are in a broad range of Internet of Things (IoT) connected devices, ranging from smart wearables, such as heart-rate monitors and smart bands, to more advanced security systems, proximity sensors, payment machines and household appliances. Bluetooth beacons [1], a technology based on BLE, are a prime application that has been recently gaining in popularity [2]. Many stores and businesses are trying to find new ways to target customers' devices with promotions. Specifically, this application can be used to send offers and promo-codes directly to the customer. Low manufacturing costs and easily-deployable hardware are a strong point of this technology and make it attractive to many upcoming applications.

In this work we will view Bluetooth in a different light and analyse its potential in more scenarios

End Product Market Segment	**2017**	**2018**	**2019**	**2020**	**2021**	**Growth**
Cellular	2,501.8	2,646.8	2,775.9	2,894.6	3,003.5	20%
PCs/Peripherals	349.4	368.9	384.1	403.2	416.8	19.3%
Networking	18.6	29.0	39.8	51.1	62.7	237%
Smart Homes	20.8	39.5	71.3	120.2	179.1	**761%**
Automotive	74.2	86.4	92.3	99.0	103.4	39.3%
Wearables	146.4	174.5	202.4	232.1	265.0	81%
Beacons	26.2	83.1	217.5	372.0	565.3	**2057.6%**

Table 1.1: Bluetooth 5 enabled devices shipments forecast 2017 to 2021 (in millions) [3]

than just those it was originally developed for. By taking into account different mobility approaches we will study this technology from an active perspective, applied to Vehicle-to-Vehicle (V2V) communications, as well as from a passive perspective, by using its discovery capabilities to provide contextual information. In the following section we describe our research question and its relevance.

1.1 Research Question

The research question that we address in this thesis is the following:

How may we exploit Bluetooth in a mobile context?

Our goal is to find out how we can benefit from the Bluetooth technology in mobile environments. We define *mobile environment* as movements in our day-to-day activities that are closely reflected by the devices we carry with us (e.g. smartphones, smartwatches). Bluetooth has a considerable market presence and is the main communication method for a variety of IoT objects. Its ubiquity is expected to grow even further over the next five years [3], making it one of the most attractive technologies to explore in relation to human mobility. A recent study on Bluetooth's latest release, as seen in Table 1.1, shows how the expected number of Bluetooth-enabled devices shipped is supposed to dramatically improve for nearly all market segments in the coming five years.

BLE supports discovery of and interaction with IoT devices in a fast and seamless manner whilst offering a very advantageous range of 100 m and more, unlike previous Bluetooth versions. The combination of these features has allowed Bluetooth to redefine itself as a very dynamic technology that is implemented today in a variety of mobile contexts. For example, BLE is found in many of the wearable devices that we use every day to keep track of activities such as walking, running and cycling. Whatever the specific activity they were designed for, we carry these devices with us at all times, when driving, working or commuting on public transport. These devices are inherently exposed to a multitude of environments and are therefore subject to varying mobility profiles. As well as in objects that are naturally mobile, we find BLE also in many static items such as televisions and payment terminals.

This broad spectrum of implementations is the reason we want to test the limits of Bluetooth in different mobile environments.

Because BLE is still relatively new and is still evolving, little research has focused on its potential and on possible improvements. The objective of this work is to address this deficiency and see how Bluetooth could potentially fit into new application fields. Foreseeing three years ago the future proliferation of BLE-enabled devices, we decided to invest in research into finding and implementing solutions that would benefit human mobility studies. One concrete application is the dissemination of information in Vehicular Ad Hoc Networks (VANETs). As the deployment of VANETs is taking longer than expected [4], BLE could immediately be utilized and later embedded alongside other technologies to empower the new generations of V2V and Vehicle-to-Everything (V2X) communications. Since Bluetooth Classic (BC) is already present in car Head-units (HUs), it is safe to assume that BLE will soon make its way into production models. This is especially true as more and more car manufacturers provide interfaces to tightly integrate mobile devices within new vehicles (e.g. Apple CarPlay [5], Android Auto [6]). We believe the ubiquity of BLE will allow fast deployment of new Intelligent Transportation Systems (ITSs) in the near future.

Other applications can take advantage of the beaconing capabilities of new IoT devices to provide contextualized information, such as location, by collecting Bluetooth signals from nearby devices. This practice brings new challenges and new opportunities that are yet to be met. The high volume of information that will be available over the air will provide additional incentives for big-data studies with a specific focus on Bluetooth. Within the next few years we could see enough data becoming available to be exploitable in many research and commercial topics, as has happened with Wi-Fi in the past decade.

In this work we will explore the above applications and distinguish between using BLE *actively* for vehicular communications, by pairing and exchanging messages, and *passively* for characterizing different environments by simply observing other Bluetooth devices nearby. In the next section we provide details on the methodology and our contributions for these two approaches.

1.2 Methodology

Intuitively there are two ways we can benefit from Bluetooth. One is to leverage the technology in the role it was designed for, that is as a communication medium, by looking at its performance and robustness in different mobility scenarios. Secondly, given its application-centric nature, it is possible to take advantage of its discovery and beaconing scheme by using a more conservative connectionless approach. Given Bluetooth's ubiquity, we can utilize its radio scanning and beaconing procedures to learn about the surrounding context. Whilst using Bluetooth's communication paradigm as common ground, both approaches will take advantage of BLE's long-range beaconing capabilities that allow for communications up to 100 m.

In order to answer our original research question, we will explore these two approaches to ultimately better understand how Bluetooth (BC and BLE) can benefit us in a mobile environment.

In our methodology we first analyse the technology's usefulness as a communication medium. For this *active* approach we investigate the characteristics of BLE for exchanging information between devices. As a subset of this approach we will be analysing its potential with moving vehicles applied to vehicular networks, more specifically for Inter-Vehicle Communication (IrVC). Our goal is to develop and deploy BLE-compatible hardware and software for testing communications between devices and, in our specific scenario, between two vehicles. We initially analyse a single-hop message exchange, taking into consideration various parameters such as relative speed and distance between the two vehicles, as well as signal strength and round-trip time for the Bluetooth communication. We then validate the same scenario assuming multiple hops and the delays resulting from such a configuration. As a final step, we test the efficiency of Bluetooth's coexistence features in case of moderate and heavy Radio-Frequency Interference (RFI).

Secondly, we explore Bluetooth's *passive* capabilities as an information medium and sensing technology to enhance contextual awareness and provide information, such as the type of devices nearby as well as environment information. Because of the way BLE works, from an external observer's point of view, there are potentially a multitude of devices within communication range in various different environments. By passively sniffing advertisement packets during a data collection campaign, without the need for pairing, we collect information about the surroundings of 20 participants and use this knowledge to build a classifier able to output, through a Machine Learning (ML) process, approximate location information based only on Bluetooth discoveries. We evaluate the efficiency of our approach by comparing different scenarios and tuning our model to better fit our specific dataset.

1.3 Contributions

Through this work we discover that Bluetooth is an interesting means of communication in mobile contexts. Besides being widely available, it provides seamless pairing procedures as well as advantageous transmission ranges with extremely low beaconing delays. We also found many disadvantages to its use, such as relatively high pairing delays; these must be taken into account when considering possible applications, such as VANETs. By combining the strengths of Bluetooth as a communication medium with its role as an information beaconing tool we discovered synergetic advantages. One can learn, by exploiting Bluetooth discoveries, about the surrounding context. An example is identifying different environment properties that can characterize multiple road categories such as Highway and Urban. Moreover, there remains room for improvement, as we found the currently achievable accuracy remains limited and highly dependent on the market penetration of Bluetooth itself.

In more detail, we can summarize our two contributions as follows:

- **BLE platform for V2X**. In order to experiment with BLE's new communication paradigm, we engineered a two-way transmission scheme that utilizes two smartphones as hardware devices. Each phone implements the BLE roles *peripheral* and *central*. The platform allows for each device to automatically switch roles, depending on the situation. Whenever a certain event is triggered and data needs to be relayed to other vehicles, the transmitting device will change its mode to allow nearby phones to receive and act upon the new information. The application is capable of handling multi-hops, and can retransmit, if needed, the information to further devices.

For this experiment we established multiple static and driving scenarios to study how different behaviours would affect the communication channel. Specifically, we considered various distances and speeds between the devices and present our results concerning round trip delay times and power analysis. Further, we present a study on the coexistence capabilities of Bluetooth in the presence of considerable interference in the 2.4 GHz band.

Publications This contribution is presented in Chapter 3, and has been published in three different papers [7–9] and ultimately a Journal: *Bluetooth Low Energy Performance and Robustness Analysis for Inter-Vehicular Communications* [10].

- **Environment Classification Model**. We introduce a novel classification model which takes as inputs only Bluetooth discovery information, such as number and type of devices discovered, and outputs an environment, such as highway or city. In order to collect the required amount of data we developed another platform, **BlueScanner**, which included a mobile application based on our previous V2X contribution. Through this platform, we launched a data collection campaign that ran for two months and had 20 participants. The application recorded nearby BC and BLE beacons encountered whilst driving. The final dataset covered ≈13,500 km of roads in the country of Luxembourg, during which nearly 22,500 unique Bluetooth devices were encountered. After validating the viability of our data through a statistical analysis we applied ML techniques with the specific goal of building a model that would classify road categories by observing only information available over the air about nearby Bluetooth objects. By tuning the expected precision of the model we managed, to a certain extent, to distinguish fundamental changes in the kind of discoveries, depending on the environment.

  Publications This contribution is presented in Chapter 4, and has been published in two different papers: *Towards characterizing Bluetooth discovery in a vehicular context* [11] and *Characterizing Driving Environments Through Bluetooth Discovery* [12]. Additionally, future perspectives specifically applied to general user mobility were published in the following Journal: *Characterizing user mobility using mobile sensing systems* [13]

  The sensing platform used throughout the data collection campaign is freely available to the whole community under an MIT license, and is hosted on GitHub (`https://github.com/wbronzi/bluescanner`).

These contributions, linked to the chosen research question, are of great relevance for future applications and studies. Results from our BLE performance and interference study are of importance when considering the advancement of the vehicular industry. Moreover, our environment classification model is extremely relevant at a time when the growth of beaconing technologies calls for methods to quickly and intelligently analyse this data in order to identify various human behaviours (such as

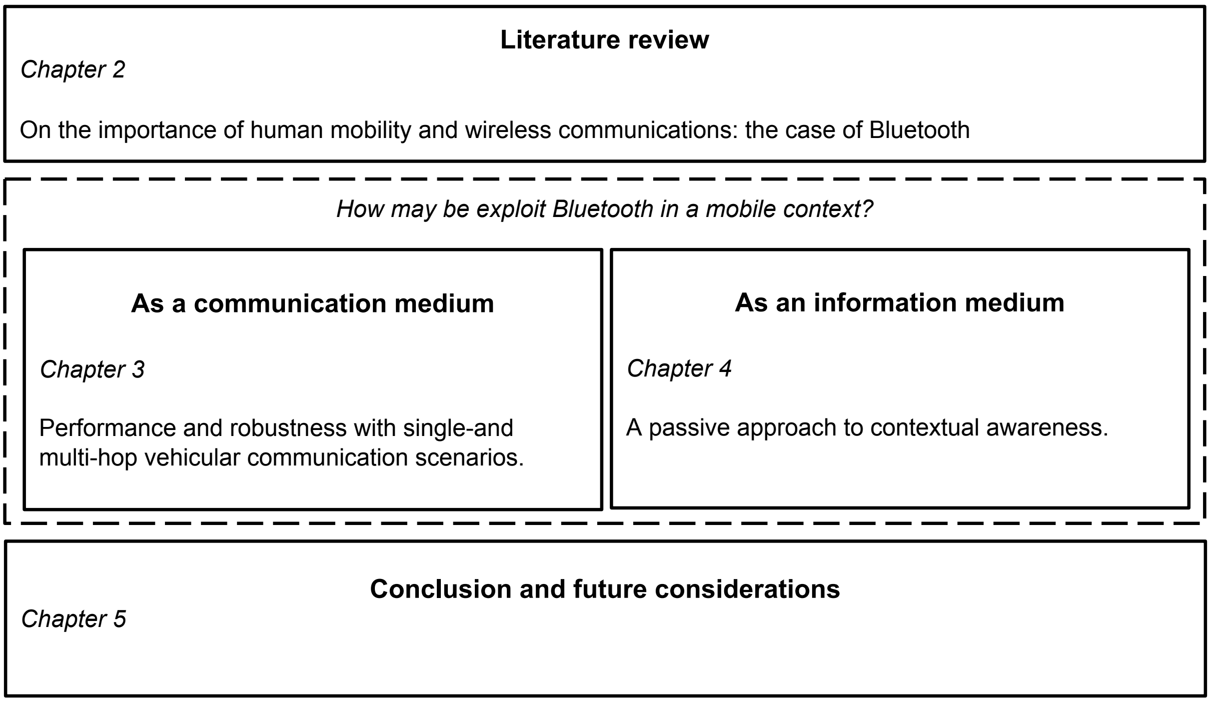

Figure 1.1: Thesis Overview

mobility). We forecast that the use of Bluetooth for sensing purposes will become much easier and more effective than other competing technologies.

1.4 Overview and Structure

The remainder of this thesis is organized as follows. In Chapter 2 we present a state-of-the-art study of Bluetooth technology, and of other relevant short-range wireless communication technologies. We also introduce different Bluetooth application fields and their corresponding mobility aspects.

We present a novel V2X platform based on BLE in Chapter 3. In Chapter 4, we study a dataset of Bluetooth discoveries and propose a classification model to detect different environments. We conclude with Chapter 5, which provides a final overview of the work, as well as giving future perspectives for the improvement and application of our research.

Furthermore, we provide in Figure 1.1 a structured overview of this thesis.

Chapter 2

Short-to Medium-Range Wireless Communication Technologies

Nowadays, we find a wide variety of wireless technologies in existence [14–16]. Some of these technologies are built for similar purposes and are therefore in competition with each other; but for the most part these standards are commercially implemented and applied in specific scenarios for different applications.

Everyday, we encounter many different wireless standards, two of the most common groups are Wireless Personal Area Networks (WPANs) [17] and Wireless Local Area Networks (WLANs) [18].

WPAN standards are developed with short-range communication in mind and can commonly be found in devices such as mobile phones, sensors and peripherals using technologies like Bluetooth [19] or ZigBee [20]. These protocols, Bluetooth in particular, are application-centric in nature. This means that what they advertise, before and during the pairing procedure, is in fact the application they provide. The advertisements can be interpreted as: *You can connect here to access the following service.*

WLAN, commercially known as Wi-Fi, are standards intended for wider area communications. These technologies are of a network-centric nature and, using protocols such as 802.11n, are deployed in systems that need to provide wireless access to other systems (*You can connect here to access the following network*). WLANs offer greater range and transfer speeds compared to WPANs.

There are other wireless standard groups, as shown in Figure 2.1, such as *Wireless Wide Area Network (WWAN)* and *Wireless Metropolitan Area Network (WMAN)* but these represent technologies with a very long range, such as cellular EDGE/LTE [21] and WiMAX [22], and will therefore not be part of our research.

In this thesis we will be focusing on short-to medium-range communication technologies. Specifically, we examine the original Bluetooth, to which we refer as Bluetooth Classic (BC), and Bluetooth Low Energy (BLE) and how they fare against other technologies such as ZigBee and Wi-Fi.

In this chapter we describe the Bluetooth protocol, with a specific focus on BLE, and compare specifications and usages between the different technologies listed above. Together with an overall state of the art, we outline how Bluetooth has been recently revived by its *Low Energy* branch and

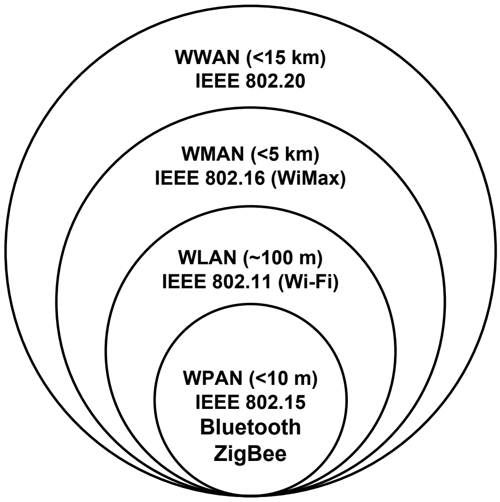

Figure 2.1: IEEE Wireless Standards Overview

is used today for a multitude of new applications. Many of these are concerned with the topic of human mobility. We finish this Chapter by considering how Bluetooth could be used specifically in a vehicular context.

2.1 Bluetooth Classic

BC is one of the most popular communication standards [23]. Originally invented by Ericsson in 1994 [24], this technology is a standard feature on a broad variety of devices and can be found in commonly-used tools such as keyboards, headphones, smartphones, car entertainment systems and many more. This versatile communication standard is currently managed by the Bluetooth Special Interest Group (SIG), which ensures that the protocol is maintained and properly updated.

Like WLAN protocols, BC operates in the 2.4 GHz Industrial, Scientific and Medical (ISM) frequency band. BC employs an Adaptive Frequency Hopping (AFH) [25] technique that allows devices to operate within the 2400-2483.5 MHz band by sending data packets over 79 BC channels spaced 1 MHz apart. As shown in Figure 2.2a, excluding guard bands [26], these channels go from 2402 MHz to 2480 MHz (2401 + n, where n, ranging from 1 to 79, represents the channel).

The AFH technique enables BC to coexist with many other users of the ISM band; these include not only Wi-Fi, but also, for example, appliances like microwave ovens that generate electromagnetic interference. A transmission uses a specific channel only for a short predefined amount of time before hopping to another. The hopping rate is 1,600 hops per second and the next channel is determined by a predefined pseudo-random hop sequence based on the Media Access Control (MAC) address of the master device. When interference is present, the lost data is retransmitted on a different channel. The protocol keeps an updated table to track which channels are to be avoided because of interference. Therefore, the communication will be somewhat robust in the case of encountering the same disruption again.

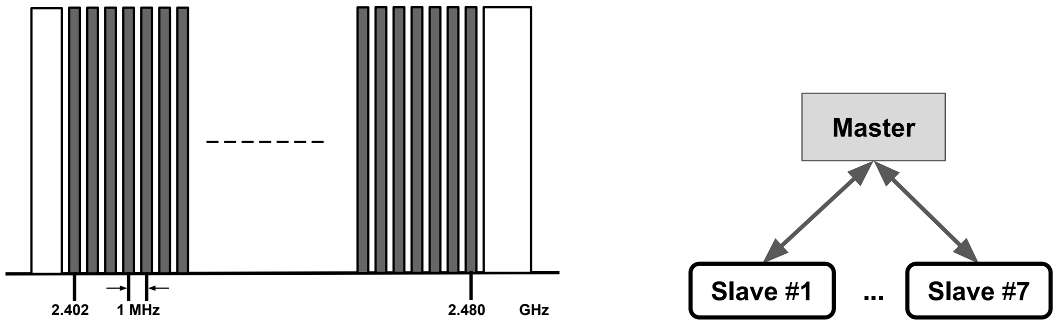

(a) Hopping Channels. (b) Architecture.

Figure 2.2: Bluetooth Characteristics

As represented in Figure 2.2b, BC is based on a master-slave architecture where one master can have up to seven simultaneous connections. This allows the creation of a network of BC devices, or *piconet*, where users can efficiently transmit data such as audio or images. BC implements a variety of standardized profiles (e.g. Advanced Audio Distribution Profile (A2DP), Hands-Free Profile (HFP), Personal Area Networking Profile (PAN), etc.) that allow manufacturers to build devices that use specific features for targeted applications. For example, the HFP is commonly used to interface car hands-free kits with mobile phones in the vehicle.

The vast majority of BC devices have a range of $\approx 10\,\mathrm{m}$ and are of type Class 2, as this provides the best trade-of between power consumption and range for day-to-day uses. The different power and range classes are:

- Class 1: Max. power 100mW (20 dBm), range 100m
- Class 2: Max. power 2.5mW (4 dBm), range 10m
- Class 3: Max. power 1mW (0 dBm), range 1m

2.2 Bluetooth Low Energy

BLE, initially introduced by Nokia in 2006 under the name Wibree, was standardized and incorporated into the Bluetooth Version 4.0 Core Specification in 2010 by the Bluetooth SIG under the Bluetooth 4.0 specification [27]. BLE drastically reduced the power required to transmit data, whilst keeping the same data rate. BLE identified several trade-offs between energy consumption and network performance that rely on the way the technology is fine-tuned [28]. Moreover, BLE completely eliminated the BC pairing procedure by introducing a seamless way to connect two devices without the need for user interaction.

BLE was developed as a single-hop communication technology with a multitude of different applications in mind, including healthcare, sport and fitness, consumer electronics, smart homes, security and proximity sensing.

Given the widespread availability of BC technology, it is fair to assume success for BLE based on

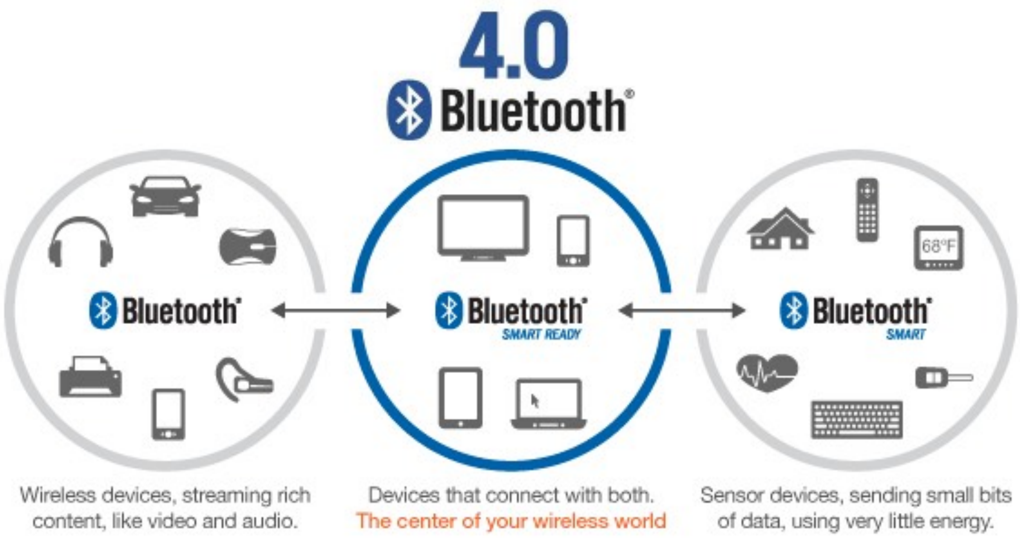

Figure 2.3: Bluetooth Dual Mode [31]

implementation similarities and existing penetration in different markets. The next release of BLE (version 5.0) promises further improvements, such as extended range (up to 4x) and faster speeds (up to 2x). This can only boost its potential, as we will see later in this work [29,30]

Some of the BLE features (version 4.0 at the time this research was done) are:

- 1 Mbps data rate (RF modulation symbol rate)
- 128bit Advanced Encryption Standard (AES) CCM Security
- Ultra low power consumption (around 1μW when sleeping and $< 20mW$ maximum consumption)
- Low Latency (6ms from non-connected state)
- Data throughput up to 128 kbit/s (Hardware specific)
- Range of over 100m (Unobstructed line of sight)

2.2.1 Protocol Stack

It is important to understand that BLE is not a simple version upgrade from BC; these technologies are incompatible with each other. BLE has different use cases and uses different chipsets to meet its hardware requirements. This means, that in some cases, BC and BLE are mutually exclusive.

In Figure 2.3 we see how BC-only devices and BLE-only devices interface themselves with *dual mode* devices that are capable of handling both Bluetooth flavours.

The BLE protocol stack, as we can see Figure 2.4, has multiple fundamental differences compared to BC.

The BLE stack is divided into two main parts; the Controller and the Host. In the Controller we can find the Physical Layer (PHY) and the Link Layer (LL); in the Host, the Logical Link Control and Adaptation Protocol (L2CAP), the Attribute Protocol (ATT), the Generic Attribute Protocol

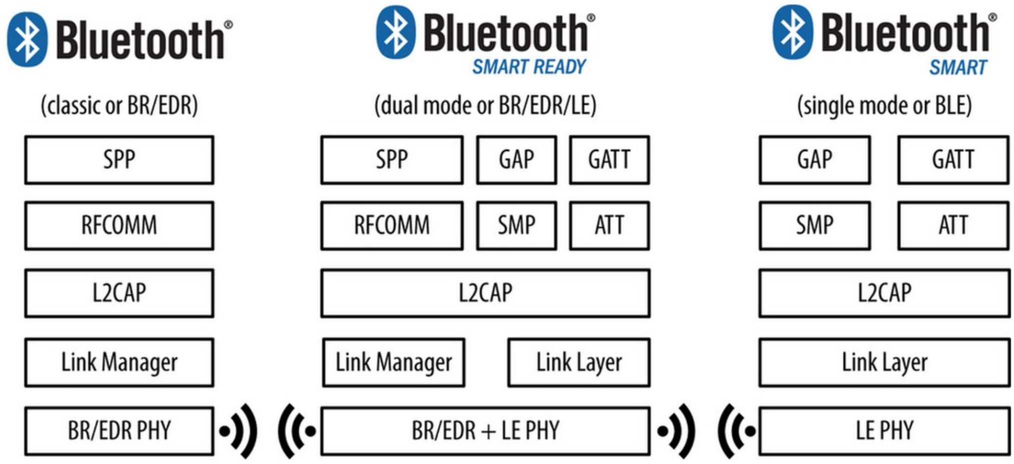

Figure 2.4: Bluetooth Stack [32]

(GATT), the Security Manager (SM) and the Generic Access Profile (GAP). The Host-Controller Interface (HCI) processes communications between the Controller and the Host.

The *Physical Layer* takes care of handling incoming and outgoing bits. BLE defines 40 channels, three dedicated to advertisement (37, 38, 39) and 37 data channels. Only one of these data channels is selected for transmitting data using the same AFH mechanism as BC to handle wireless propagation issues and interference.

The *Link Layer* is the medium used to establish bidirectional communication between two devices. Two roles are defined in this layer once a connection is established between two devices: master and slave. The master (or *central device*) is usually a device that scans for other devices (typically, a smartphone, tablet or laptop). The slaves (or *peripheral devices*) are devices advertising data (e.g. fitness tracker or smartwatch). Slave devices are in sleep mode by default, and wake up periodically in order to be discovered and to potentially pair to Master devices.

The Link Layer provides error and flow control. All data packets have a 24-bit Cyclic Redundancy Check (CRC) code to provide bit error detection, and a supervision timeout is in place to detect connection losses due to exceeding range or radio interference. A stop-and-wait packet flow acknowledges sequential packet numbers to ensure that no packet is lost.The *L2CAP Protocol* used in BLE is an optimised version of that of BC. The main goal of this layer is to encapsulate data channels for the upper layers that provide traffic management.

The *ATT Protocol* defines the communication between the master and slave. The master device is here referred to as *central* and the slave device as *peripheral*. The central device advertises data by maintaining a set of attributes. An attribute contains data managed by the GATT. The ATT protocol sets up communication between a server's attribute and a peripheral via a dedicated L2CAP channel.

The *GATT Profile* establishes a framework for discovering the data stored in the ATT protocol and

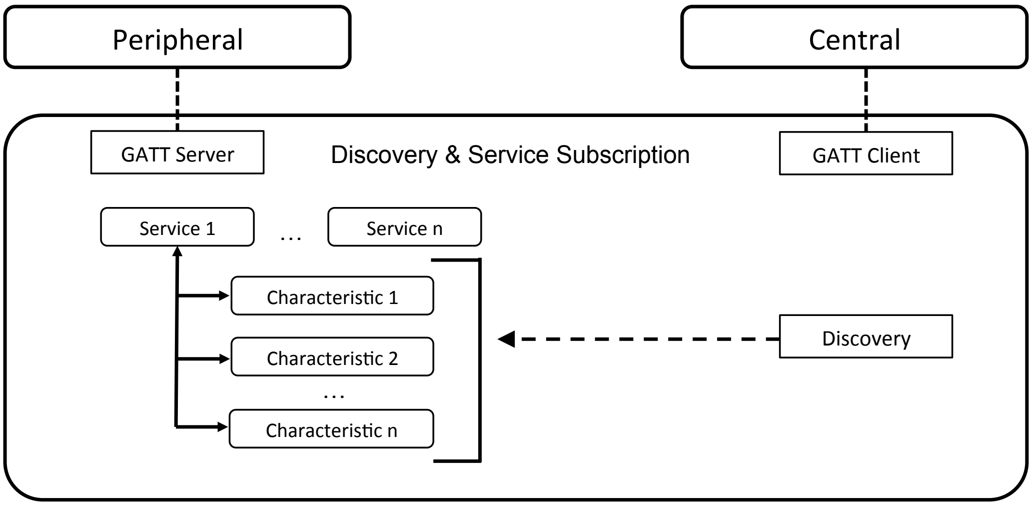

Figure 2.5: Generic Attribute Profile (GATT) Framework

exchanging characteristics between devices. Each peripheral has a set of attributes storing services and characteristics, where every characteristic contains a value and a number of properties. This configuration is shown in Figure 2.5.

As an example, if we have a heart rate monitor acting as the peripheral, its GATT server will contain a *Heart Rate Sensor* service with a *Heart Rate* characteristic. This characteristic will then include the reading of the sensor as a value and specify properties for the value, such as *Read Only*.

2.2.2 Energy Consumption

As BLE is mainly used for low power communications and applications, several studies have analysed its energy footprint, whilst also comparing it to other technologies [28,33,34]. Although we do not go into details about BLE power consumption in this work, we feel it is an important aspect of the technology and thus merits a brief discussion.

BLE energy consumption differs between roles (master and stave) as well as between the different states the devices are in (non-connected and connected). In its non-connected state, the slave device total consumption depends on the advertisement interval and number of advertising channels to use. For the master device, on the other hand, it is the scan duty cycle that impacts the total energy consumption. As an example from [35], when a slave device advertises with a 0.5s interval on only one channel, the power consumption is about 0.2mW; when advertising on all three channels this is nearly doubled to 0.36mW. It is possible to fine-tune both parameters depending on how fast we need the discovery to happen. A master devices with a 50% duty cycle consumes about 33mW during each scanning phase. The duty cycle is the probability that the master device finds a slave device during its advertisement interval. Each master's scanning phase covers the required number of advertising channels, depending on the specified duty cycle.

In a connected state the energy consumption is similar for master and slave, and depends on a series of parameters that are completely hardware-and application-specific. Every connection event in this state has specific power draws for every phase. The phases are pre-and post-processing, transmission, reception and Inter-Frame Space (IFS) delay. The total power draw in this state can quickly rise above 200mW per connection event [35].

Furthermore, analysing the energy footprint of BLE and comparing it to other technologies [35] it has been shown that its consumption is indeed very low, even considering the overhead introduced by additional layers such as IPv6 [36]. BC before the introduction of BLE, was put through a similar test against ZigBee [37,38] and Wi-Fi [39–41], from which it was shown that its classic *non-low-energy* approach already triumphed over its alternatives.

To see how these low energy characteristics can impact other parameters of BLE, we refer the interested reader to section 2.8.

2.2.3 Security and Coexistence

The BLE protocol was developed with two mutually-exclusive security modes built in: *LE Security Mode 1 & 2* [28]. These services provide encryption functionalities at the Link and ATT Layer respectively. With these two modes it is possible to transmit authenticated data over an encrypted (mode 1) or unencrypted (mode 2) Link Layer. With the first mode, the BLE link layer encrypts data by using a 128-bit AES block cipher together with the Cipher Block Chaining-Message Authentication Code (CCM) algorithm. A further Message Integrity Check (MIC) is then appended to the payload. The second mode, on the other hand, will just authenticate data over an unencrypted Link Layer by appending a signature to the payload at the ATT Layer. This signature is also computed with a 128-bit AES block cipher [27].

BLE, in its initial iteration (4.0), was known to have some security flaws. There were in fact multiple attacks aimed at the security aspects of the BLE protocol [42]. Some of these attacks show different techniques to eavesdrop on BLE communications, as well as demonstrating vulnerabilities against the key exchange protocol, effectively compromising the encryption. Indeed, several tools, such as the Ubertooth One [43], were developed to simplify sniffing procedures against BLE.

As this technology is constantly being worked on, with version 4.2 [44] and now 5.0 [45] the protocol has since vanquished many of these flaws. The pairing process has been enhanced with an Elliptic Curve Diffie-Hellman (ECDH) key agreement protocol [46,47]. The ECDH effectively neutralizes some eavesdropping attacks that are able to follow the hopping pattern if the initial pairing has been captured. In addition, a new privacy feature improves on BC by allowing a device to continuously change its MAC address, making it harder to track it.

Although security is an important topic to keep in mind, our research did not focus on this aspect. More important to our goals is to understand and validate the coexistence capabilities of BLE and how this technology deals with interference.

There has always been particular interest in the coexistence between Bluetooth and Wi-Fi [48–50]. As BC is still very predominant in the automotive field (Head-units (HUs)), the interest in this topic remains high. Some research has also introduced BLE into the equation [51, 52] but, since BLE hardware is able to reuse existing BC coexistence features such as passive interference avoidance schemes (e.g. AFH [53]), they assume that its behaviour is similar. This is partially true but, considering the applications of these technologies (BC for audio/file transfers and BLE for low-volume data), the average packet size is considerably smaller for BLE, allowing it to perform better than BC most of the time.

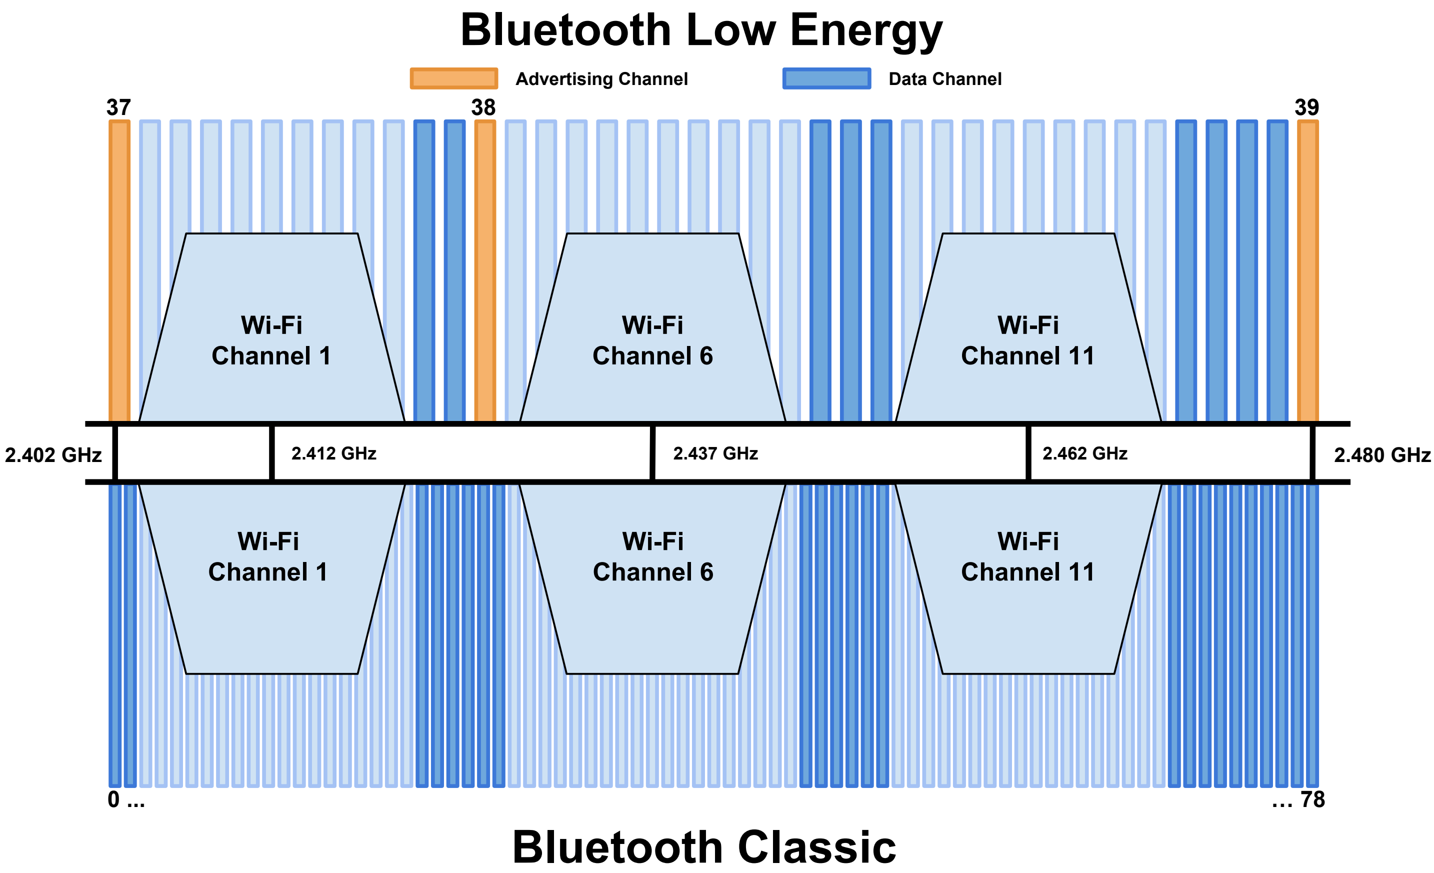

Figure 2.6: BC, BLE and Wi-Fi Spectrum Mapping

Moreover, BLE channels have a different spacing compared to BC's (2MHz for BLE and 1MHz for BC) and are of two types: data and advertising channels. Advertising channels are specifically chosen to be in the least congested zone of the 2.4 GHz band. Figure 2.6 shows how the different technologies occupy the spectrum.

Because of the shortcomings of existing research, we experimented first hand with the co-existence mechanisms of BC/BLE and present the results in Section 3.3.

2.2.4 Physical Limitations

When considering wireless communications applied to non-static scenarios, like in our case a vehicular context, there are certain physical limitations that should be taken into account. In particularly extreme applications that demand high reliability, when the relative speed between two communication points is very high, a Doppler shift [54] should be considered.

When travelling at higher speeds the Doppler effect can cause a shift of frequencies. This can inevitably lead to loss of packets and overall degraded transmissions as packets won't be in the expected frequency range. This issue is most researched for mobile networks were use cases are more prominent.

Studies have shown how technologies such as LTE can withstand, with reduced data rates, speeds of over 200 km/h [55].

As for the security aspect, this limitation was out of scope in our research and was therefore not further explored. Bluetooth was not designed with such scenarios in mind. It is probable that a high speed transmitter (over 200 km/h) on a certain 2 MHz channel (for BLE), will be observed on a different channel by a non-moving receiver. As in our experiments we do not reach, or even get close, to the aforementioned speeds, we can safely assume a negligible Doppler shift impact.

2.3 Vehicular Communications

Little research has focused on using BLE for Vehicular Ad Hoc Network (VANET) applications. In fact, the majority of works that refer to BLE only discuss Intra-Vehicle Communications (IaVCs) and not Inter-Vehicle Communications (IrVCs). Some propose modifying the protocol to make it more suitable for IrVCs by enhancing the performance to guarantee the worst-case latency required by automotive systems [56]. Others evaluate potential in-car applications [57] and the impact of in-vehicle interference on BLE caused by the simultaneous use of different ISM band wireless technologies [57]. All these studies also conclude that, amongst the alternatives, BLE is the most resilient technology when it comes to interference.

To the best of our knowledge we are the first to apply BLE in a extra-vehicular context.

There are many technologies currently in development for IaVCs. It was not our intention to replace these technologies, many of which have been years in development and have substantial corporations backing. Our goal was to take a technology with existing mass adoptions and see how it fared in this new environment, leading to the possibility of being implemented alongside mainstream vehicular communication technologies. We provide here a short list of these other technologies.

For an initial deployment of Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communications, two standards currently lead, promising to provide driver information and warnings to all equipped cars. Both these technologies, Dedicated Short Range Communication (DSRC) in the U.S. [58] and ITS-G5 in Europe [59], rely on the IEEE 802.11p standard [60].

Dedicated Short Range Communication (DSRC)

DSRCs is a suite of standards with the specific goal of providing a fast and reliable exchange of safety messages. DSRC relies on the IEEE 802.11p standard and thus operates in the 5 GHz frequency band, but within a specific channel range reserved for DSRC from 5.850 GHz to 5.925 GHz (5.9 GHz band). The frequency band is divided into seven channels and the protocol offers data rates between 6 and 27 Mbps up to 1 km range [61]. This technology was designed to provide minimum latency in the communication link whilst maintaining high data transfer rates.

The network is composed of two units; the Road Side Unit (RSU) and On-Board Unit (OBU). As the names suggest, one is mounted in a vehicle whilst the other is placed at a fixed spot in the road network. Vehicles use the Wireless Access in Vehicular Environments (WAVE) standards to communicate between each other and with the RSUs [62]. DSRC is set to be mandatory in new vehicles in the U.S. starting from 2020 [63].

ITS-G5

In parallel with DSRC in the U.S., Europe has developed its own standard; ITS-G5 [64]. Both approaches have many similarities. Like DSRC, ITS-G5 relies on the 802.11p [65] stack and operates in the 5.9 GHz frequency band. The difference between the two lies in the way the channels are divided within the spectrum. In ITS-G5, each channel set is assigned a letter from A to D (ITS-G5A,etc.), each with a different frequency span [66].

ITS-G5 utilizes an ad hoc multi-hop routing protocol that uses geographical coordinates for sending and forwarding messages. As with DSRC, ITS-G5 promises a range of up to 1km with default data rates of 6 Mbps. Vehicles implementing this protocol will be able to communicate with other vehicles as well as with infrastructure. ITS-G5 aims to improve the efficiency and safety of roads by relaying information gathered by vehicles and infrastructure regarding weather, accidents and other unexpected events. Unlike DSRC, an adoption roadmap for ITS-G5 in Europe has yet to be provided.

2.4 The Human Mobility Aspect

Given that the vast majority of wireless devices are carried by people, many of their applications are closely related to their mobility [67,68]. Each daily activity is very subjective and difficult to generalize for a large group of people. The balance between social and work activities varies depending on the lifestyle, age and geographical location of each observed case. Nowadays, it has become essential when deploying any kind of public service to have a general understanding of its application and uses. Modelling the mobility of users is therefore vital for completing this design process [69].

Human mobility is inherently very complex, making finding metrics to describe it challenging. It was found that the most effective way to find the essential properties of human mobility is to deduce them from real human mobility traces [70].

Thanks to the widespread availability of smartphones, there are many studies available today based on collected human mobility traces. These data collection campaigns pair GPS coordinates with sensor data in order to deduce mobility patterns by aggregating all information and associate it with specific activities.

Together with the previous metric, knowledge of encountered and used Wi-Fi networks as well as Bluetooth devices has recently gained traction to make the process of tracing human mobility even simpler [71].

The most prominent fields where we see Bluetooth being applied in this context are in the transportation and crowd-management sectors. In the transportation field, Bluetooth scanners are more and more used instead of loop counters to help cities understand and better manage the transport network by scanning Bluetooth devices to estimate travel times. The city of Brisbane is a prime example on how these technologies can be used to retrieve vehicle trajectories [72] and study traffic congestion [73]. Several other major cities in the United States have deployed similar technologies to improve their road traffic infrastructure [74].

In the field of crowd-management, Bluetooth is being actively used to observe pedestrian behaviours through Bluetooth monitoring [75]. The Transport for London (TFL) has concluded this year a pilot study to track the flow of people in and out of the underground network. Amongst the technologies used for tracking the crowds we also find Bluetooth [76]. Many companies are since offering

(a) Indoor Localization (b) Transport (c) Smart Payments (d) Automotive Industry

Figure 2.7: BLE Application Fields

products specifically targeting crowd-management by detecting nearby Bluetooth beacons [77].

BLE has, in our opinion, great potential in the quest to better understand human mobility. Compared to other similar technologies, such as ZigBee, its wide range of functions and applications are very attractive for a multitude of studies.

In the following section we discuss current and future Bluetooth BLE applications and how they relate to different types of human mobility, from static contexts to vehicular communications. We start with a rundown of applications, some which are still under development, which illustrate the potential of Bluetooth in the coming years. These applications cover diverse mobility profiles and show the broad reach of BLE in this field.

2.5 BLE Applications

When BLE was introduced in 2010, it targeted novel Internet of Thingss (IoTs) devices in the areas of healthcare, sport & fitness and smart homes. Since then much has changed, and many new applications have found their way onto the market.

When we think about BLE, we think about heartrate sensors or smartbands, but the reality is that there are many more BLE-enabled objects that are not commonly known. Our study in Chapter 4 reveals the many types of devices that are already in our homes and streets. Many more applications are currently in development and about to enter the mass market. To give an idea of the prospects for Bluetooth, here are some of the fields with the most interesting forthcoming applications.

Medicine. There is a growing number of BLE applications in the medical field [78]. Low-cost blood pressure monitoring systems [79] as well as blood glucose meters [80] propose a new approach to measuring health parameters at home by pairing the device with a smartphone.

Another application in development will provide, with the use of indoor localization, the precise location of a person in need of medical care during an emergency call. It is expected that the phone, once such a call is detected, will try to determine its precise position by observing nearby Bluetooth beacons and automatically communicate this information to the responder's team. With the help of BLE indoor localization, there is the potential for saving thousands of lives per year [81].

Indoor Localization (Figure 2.7a). By using beacons, retailers are able to interact with customers in new innovative ways. Stores can collect user data and push target ads and coupons to make for a more involved shopping experience.

Beaconing is said to have great potential to massively improve retail marketing [82].

Indoor navigation is another topic touched by this field and can improve personalized direction information in cultural centres (also enhancing tour guides), airports and other large transport hubs [83].

Transport (Figure 2.7b). BLE has great potential in the public transport field. As we will see in Chapter 4, and as others have pointed out [84], BLE is ideally suited for warning about arriving vehicles, counting and tracking passenger movements, and for handling ticketing and fare collection [85].

Innovative standalone apps have also been developed to facilitate some commuters' tasks, for example, discreetly requesting a seat when a pregnant woman is about to enter a bus [86].

Smart homes. The field of home automation systems has been revolutionized by BLE. There are many applications for smart homes intended to simplify tasks and for monitoring. *Smart locks*, which automatically unlock a door when a user approaches, are more and more common. [87]. Many *Sensors* for monitoring windows, fire and smoke detectors, motion detectors as well as lighting control, are now accessible on conventional smartphones though BLE.

Appliances such as fridges, microwaves and conventional ovens are also controllable, to make sure they are not left open and are turned off. Also in the kitchen, there are many other inventions, such as the smart fork [88], that helps you track eating habits and weight loss goals.

Smart payments (Figure 2.7c). This concept allows customers to pay directly in equipped stores by using a smartphone application [89, 90]. This concept can also be applied to other fields, such as public transport, for example, to pay for the ticket when entering a bus.

Tracking-related applications also show the potential of BLE. *Smart Tags* can be attached to keyrings or put into luggage to track their location [91]. Specific trackers have also been made for tracking a pets behaviour by attaching a BLE tag to the animal's collar [92].

Industry (Figure 2.7d). The automotive industry is one of the many that can benefit from the use of BLE. Besides the topic of this research, BLE has been implemented in many in-vehicles system, such as intelligent steering wheel and seats, connected mirrors and even keyless systems [93].

BLE has also made an appearance in agriculture. Fujitsu has developed a pedometer for cows to estimate the best time for insemination. It is said to improve fertility detection rates by up to 95% [94].

There are other fields and applications that have not been mentioned above. Among these are wireless charging, social network mapping, education and many more which would take too much space to describe here.

Next we present as a case study, an application that we developed, which allows users to share their cellular data plan by communicating with other BLE-enabled devices.

2.6 BLE Case Studies

To gain a fuller understanding of the underlying features behind BLE, we developed multiple applications to showcase its capabilities and potential. This process eventually helped us to realize our goal and focus on specific approaches.

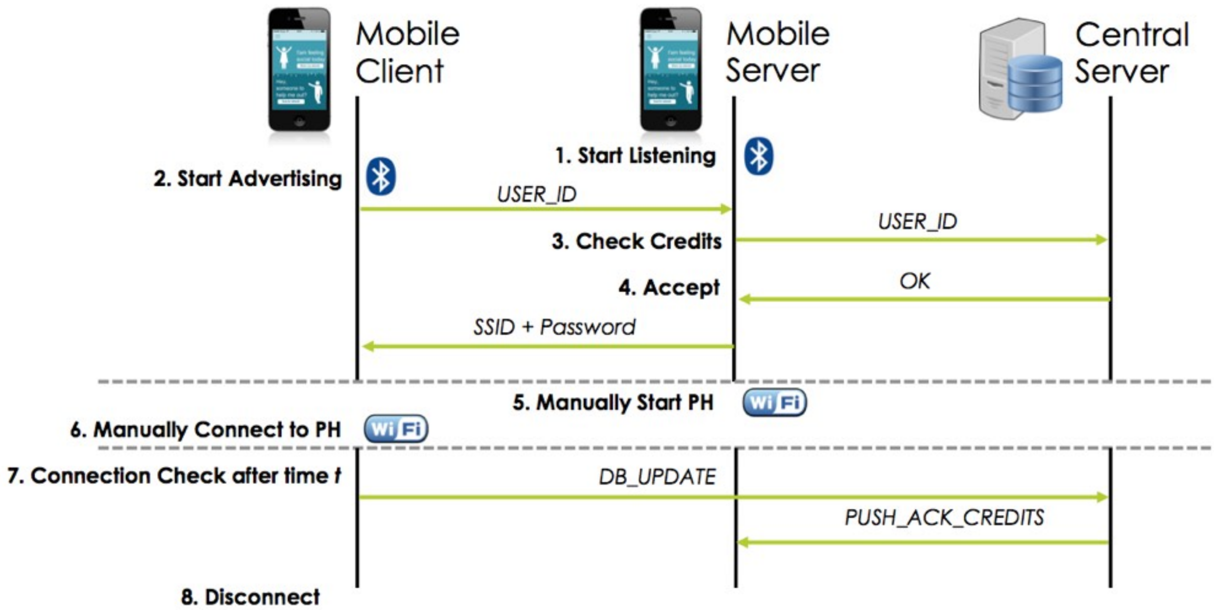

Figure 2.8: Credit-based Connectivity-sharing Platform using BLE

2.6.1 Connectivity Sharing

For this specific case study we developed a mobile application for allowing users to connect with nearby smartphones though an innovative handshaking procedure based on BLE. The goal of this application is to enable users to share their cellular data connection, using Personal Hotspots (PHs), with other travelling users that do not have access to the internet. Using the app, users are able to gain points (or credits) by sharing their cellular data plan and later spend them by using in return another user's connectivity.

The relevant aspect of this credit-based connectivity-sharing platform is the communication paradigm that allows members to connect with each other. We exploit BLE's pairing procedure to make seamless contact with all nearby devices that have the app installed. In Figure 2.8 we show how two devices, with the help of a central server for checking credits, assume different roles (client and server) depending if they are sharing or searching for connectivity. Using BLE's long-range beacons we automatically pair and exchange details, such as the PH name and password, without the need for user interaction.

Using this case study we experiment with the development of BLE enabled applications and establish fundamental requirements for the contributions outlined in this thesis.

2.6.2 Gate Opener

We deployed, as part of another study, a BLE receiver within a parking exit gate. The receiver was able to interact with our mobile application through our previously developed handshaking procedure. A user with enough credits was able to open the gate just by approaching the exit. This bare implementation was used a proof-of-concept to show how BLE could be embedded in objects that we

(a) Approaching Gate.

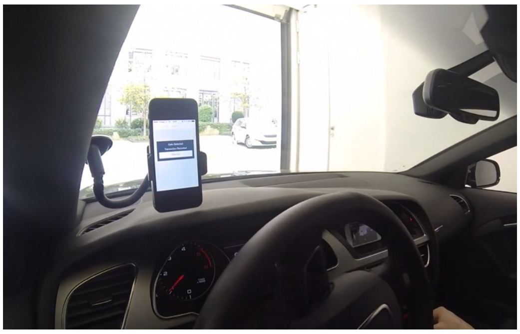

(b) Notification of Opening after Payment.

Figure 2.9: BLE Gate Opener

interact with whilst driving. A similar implementation can be used for a tool booth on the highway for example.

In Figure 2.9a we show a vehicle approaching the exit and which, without having to stop, is able to open the gate after receiving a confirmation of payment on the mobile phone (Figure 2.9b). The opening only occurs at a certain distance which is based on the signal strength from approaching devices. In order to have a seamless user experience, the appropriate signal threshold is chosen depending on the environment.

2.7 The Future of Bluetooth

Bluetooth 5 is still in its early adoption stage and thus has had little research effort invested in it. Aside from the standard specification [45], detailed studies are hard to find [95]. At the time of writing, there is only one production smartphone which supports this new technology [96] and only few development kits are available (e.g. Nordic Semiconductors nRF52840 [97]). Such kits can be used to benchmark the long range capabilities of Bluetooth 5 where, with a transmit power of just 8 dBm, a transmission range of more then 600 m was achieved [98]. It is expected that, by the end of 2017, many more Bluetooth 5 devices will appear on the mass market.

Since its introduction 20 years ago, Bluetooth has undergone major enhancements for each new version release. The most significant came in version 4.0, with the introduction of BLE. Bluetooth 5 further improves BLE's capabilities, showing a clear emphasis on this technology compared to Bluetooth Classic, for which nothing new was introduced in this new version.

In the specifications we also find hints at future audio streaming capabilities over BLE instead of BC, thanks to its higher transfer speeds.

Bluetooth 5 improves the physical layer by giving more coverage at the expense of data rate. The technology offers more range and higher data transfer speeds but not both at the same time.

The higher data rate (2x) is achieved by doubling the normal BLE data rate to 2M symbols per second, meaning that the same amount of data takes half the time to transmit as previously. In this

	BC	BLE 4.0-4.2	BLE 5	ZigBee	Wi-Fi	DSRC (802.11p)	ITS-G5 (ITS-G5)
Band (GHz)	2.4	2.4	2.4	2.4*	2.4/5	5.9	5.9
Ch. Bandwith (MHz)	1	2	2	2	20/40	10	10/20
Total Ch.	79	40	40	16	11	7	7
Modulation	GFSK	GFSK	GFSK	BPSK/OQPSK	OFDM	OFDM	OFDM
Max data rate (Mbps)	3	1	2	0.25	600	27	12
Max power (mW)	100	10	100	1	100	800	2000
Max Range (m)	100	100	400	40	100	1000	1000
Latency (ms)	100	6	6	30	5	<10	<10

*Also 784/868/915 MHz in China/Europe/US,Australia.

Table 2.1: Technologies Comparison

case, the radio is operating for a shorter time, meaning that energy consumption is further improved. The increased range (4x) is achieved by using the old BLE data rate, explaining why increased data rate and increased range are mutually exclusive. It uses a lower packet coding rate of 125 or 500 kbps and an increased transmit power of up to 20 dBm.

Furthermore, Bluetooth 5 allows advertising packets to be transmitted over both data and advertising channels.

2.8 Short-to-Medium Range Technologies Comparison Overview

To conclude this Chapter we present Table 2.1, which details all the technologies discussed. Using this table, we can compare key characteristics and use them as reference when considering different approaches concerning which technology to apply.

In our case, it is important to compare Bluetooth with other technologies. By doing this, we can better understand the shortcomings of each. For example DSRC/ITS-G5 infrastructure was built to be reliable for safety applications, whilst BLE, because of its widespread availability, can be used for general awareness broadcasts and low-priority driver information messages.

Each protocol listed was built upon an IEEE standard. Some share fundamental similarities and were forked to become something different to be used for specific purposes. We often find technologies such as Wi-Fi, Bluetooth and ZigBee being compared with regard to coexistence [19, 39] or energy consumption [35, 99]. These technologies have to coexist in most populated areas and as such are always under scrutiny for problems related to unique applications.

The contribution of this section is to put these technologies in perspective not only with Bluetooth's new versions(4 & 5) but also alongside VANETs. Since our purpose is to evaluate Bluetooth outside its original application area, it is vital to assess its performance in a vehicular environment.

2.9 Summary

In this Chapter we introduced Bluetooth, with particular focus on its low energy implementation. We characterized this technology by going into details on its spectrum, features and interference avoidance strategies. We contextualized its attributes with respect to other relevant wireless communication technologies such as VANETs. We also provided an overview of current and future applications with their different mobility aspects.

In the next two chapters we present our main contributions. We first investigate Bluetooth as an *active* technology for V2V communications. We benchmark BLE in specific vehicular environments and discuss its potential in such applications. Secondly, we focus on its *passive* use in conjunction with human mobility behaviour for context awareness applications. In this second study we also apply Machine Learning (ML) techniques.

Chapter 3

Bluetooth as a Communication Medium in a Mobile Context

A primary interest area for Intelligent Transportation Systems (ITSs) is that of vehicular communications [100]. Although technologies such as Dedicated Short Range Communication (DSRC) and ITS-G5 were standardized for this specific purpose, after more then a decade of research and development, their full adoption can only be expected at best within the next five to ten years.

As we presented at length in Chapter 2, Bluetooth has not been considered as a technology for mobile contexts until recently. While the protocol was originally designed for short-range one-hop communications, we show in this Chapter, how with Bluetooth Low Energy (BLE), it is possible to send short messages from one device to another up to a maximum distance of 100 m. Furthermore, a robust connection, capable of handling sudden signal loss or interference, can be achieved up to a distance of 50 m even for varying traffic and driving conditions.

By means of experiments, we first evaluate the characteristics of the wireless channel, then we define a set of driving scenarios to analyze how BLE is affected by varying speed, distance and traffic conditions.

We developed a proof-of-concept mobile application that uses off-the-shelf smartphones to show how BLE can be used to send data over multiple hops, significantly increasing the scope of the application.

To further analyze BLE in its current state for Inter-Vehicle Communications (IrVCs), we show how co-existence between Wi-Fi and BLE plays out in practice, and investigate how resilient BLE communications are to interference on the same frequency band from IEEE 802.11 devices operating at maximum capacity.

To do this, we built a static interference testbed composed of six Raspberry Pis, of which three were Access Points (APs) assigned to non-overlapping Wi-Fi channels 1,6 and 11. This drastically reduced the number of BLE data channels available to the protocol, while still leaving the advertising frequencies (BLE ch. 37, 38, 39) unperturbed. We observed, amongst other things, how the association time between the two mobile phones varies with the different scenarios.

During our experiments, we measured performance in terms of delivery ratio and round-trip time for multiple dynamic vehicular scenarios and static interference scenarios. This allowed us to identify and discuss several shortcomings that make the current version of BLE unsuitable some types of applications.

We conclude that BLE can, indeed, be used to exchange information between vehicles while driving. This makes BLE an interesting candidate for specific deployment but, given the current pairing delays, it cannot be considered as a complete replacement for on-board communication interfaces such as DSRC.

In the following sections we present a novel approach for utilizing BLE in a vehicular context. We propose an experimental set-up for implementing such technology alongside existing technologies and discuss our results. We also analyse BLE coexistence behaviour when heavy Wireless Local Area Network (WLAN) interference is present.

3.1 BLE for Vehicular Communications

BLE was not developed with vehicular applications in mind. Nevertheless, a wireless ad hoc network linking Bluetooth devices (also called piconet), between one or more vehicles can be engineered using off-the-shelf smartphones fixed to a car's dashboard. We can anticipate that, in the near future, where entertainment systems in cars provide a stronger bond between smartphone and car (e.g. Apple CarPlay [5], Android Auto [6]), the BLE interface from the smartphone or from the car itself can be used for fast and reliable communication with neighbouring vehicles.

In this chapter we aim to show that, at this moment, an exchange of information in a mesh of moving vehicles is reliably achievable up to a certain distance between each hop.

Different scenarios were put in place to show how connection and link performance is affected by the cars' distance and speed. Experiments on the road were conducted for single-hop scenarios and a proof-of-concept for multi-hop communications was developed and tested in our lab. In the following subsections we describe our single and multi-hop scenarios in more detail.

3.1.1 Single-Hop

For the single-hop scenarios we developed a mobile application (see Section 3.2.1) and used two identical smartphones (iPhone 5S) with the same OS version (iOS 7.1), which was the latest available hardware/software combination at the time the test was made. For all the experiments we fixed both phones inside the vehicle in a phone holder.

The application implements both BLE central and peripheral modes and handles role switching automatically. Each smartphone starts as a central device and will manually wake up its peripheral role via a button on the App's user interface. Note that, once a device switches to peripheral, the device still maintains a central role in a different queue, allowing it to discover other peripherals in the vicinity even while broadcasting.

We defined four different scenarios to thoroughly benchmark BLE for IrVCs:

- Static
- Constant distance & speed

- Constant speed
- City driving

The first set of tests consisted of measuring the channel performance with gradually increasing distance between the two *static* vehicles (from 5 to 100 m in steps of 5 m). The second scenario, *Constant Distance & Speed* consisted of two vehicles driving on the highway at constant speed (80 km/h) whilst increasing the distance by steps of 10 m from 20 to 80 m. The third scenario, *Constant Speed*, was conducted with one vehicle parked at the side of the road whilst the second vehicle was driving on the same road at constant speed. The final scenario, *City Driving*, consisted of a normal drive through the city of Luxembourg allowing changes in speed, distance and considering local traffic (e.g. other cars between sender and receiver).

3.1.2 Multi-Hop

The Multi-Hop scenario was conducted in our laboratory to give a *proof-of-concept* demonstrating its feasibility in a real-life driving scenario.

The application used required a few additions to its implementation, such as Event-Based Role Switching (EBRS), and the ability to trigger special messages and to handle them accordingly on the receiving device.

Event-Based Role Switching

As shown in Figure 3.1, for this application each device start in central mode and will activate its peripheral role automatically upon receiving a message to rebroadcast. In the scenario shown, the first vehicle on the left manually triggers a message that is received by all cars in range. The car in the middle rebroadcasts the messages to the third car (out of range of the first) and has both peripheral and central roles active simultaneously for a certain period of time. Once the middle smartphone receives the first message as a central device it activates its peripheral manager and rebroadcasts messages, whilst still receiving new ones from the first car. EBRS is a necessary feature for enabling vehicles to send and receive messages, automatically and at the same time, after being paired to each other.

3.2 Performance Evaluation

For each scenario described in Section 3.1, we analysed data taken from our mobile application logs regarding delivery ratio, round-trip time, distance between the vehicles and packet rate, as well as end-to-end delay for our multi-hop experiment.

The role of the mobile application was crucial for the performance evaluation since data connection and messages exchanged were directly logged on the devices used for all scenarios. These logs were then extracted after each test run and combined to examine the scenario results.

The application was developed for iOS devices since at the time it was possible to develop code for both Central and Peripheral modes. For other platforms such as Android, the option of making the phone act as a Peripheral device was not possible at the time, so these devices were therefore not

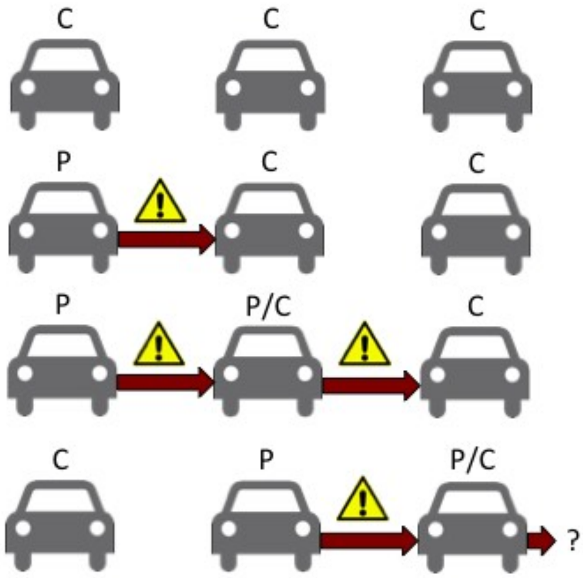

Figure 3.1: Event-Based Role Switching. (P=Peripheral Mode, C=Central Mode)

considered for this part of the research.

Although energy consumption is not a relevant factor in our study it is important to underline BLE's performance with respect to this parameter. As the detailed study by Siekkinen, Hiienkari and Nurminen [35] shows, the total energy consumption for BLE devices in non-connected states will vary depending on a variety of communication parameters. These parameters, such as connection interval, slave latency and connection supervision timeout, can be optimized for a vehicular scenario to reduce the power drain. As devices remain in a connected state only for a short period of time (typically a few milliseconds, depending on the application), we assume that the consumption in this state negligible. Since an analysis and optimization of the power usage was not required for our experiments, we did not investigate its impact on the devices available to us further.

3.2.1 Mobile Application

As mentioned, the mobile application was developed for iOS devices. The Core Bluetooth framework from the iOS Software Development Kit (SDK), at the time it was developed, offers not only the possibility to implement both central and peripheral BLE modes but also provides a simple and flexible platform to interact with BLE functionalities.

Testing of single-hop scenarios was conducted outdoors, thus requiring the iOS Location Manager to take advantage of the Global Positioning System (GPS) for tracking position and speed. For these tests, the application was deployed on two identical phones fixed each on the dashboard of one car, as can be seen in Figure 3.2a. One device was always kept in peripheral mode, having two characteristics—one readable and one writable—and constantly updating the readable characteristic

(a) Smartphone location

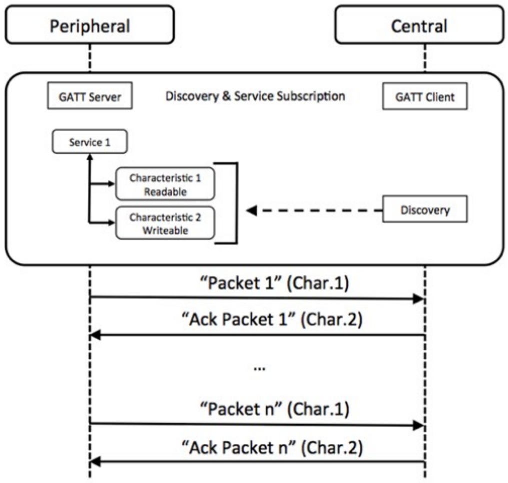

(b) Data exchange between devices

Figure 3.2: Smartphone position and configuration

every 100 ms (10 Hz) after successfully pairing with a Central device. The flowchart for this configuration can be seen in Figure 3.2b. We choose a beaconing frequency of 10 Hz because this meets the requirement of most ITS applications [101].

The readable characteristic contained the peripheral's GPS location, speed, and a sequential ID. When the central read this information it calculated the distance between the two devices using the latitude and longitude coordinates provided by the peripheral. This distance and the relative speed were then displayed on the screen of the receiving device, allowing better control of the experiments.

The central device also stored readings of the Received Signal Strength Indication (RSSI) corresponding to the peripheral's messages, as well as coordinates from both devices. After the peripheral had updated the readable characteristic with a message, the central device read and acknowledged the packet by writing its ID into the peripheral's writeable characteristic.

For multi-hop testing, our application behaved slightly differently. Compared to our single-hop scenarios, where we needed timestamps only on the peripheral device, we now needed to synchronize internal clocks on all devices regardless of their role. Using the single-hop scenario, we could calculate the round-trip time using the peripheral's characteristics by computing the difference between the timestamp of the sent message and the timestamp of the acknowledgment (when the central wrote on the second characteristic). For the multi-hop scenario, we choose not to have such a characteristic since we were interested in the end-to-end delay, which is more relevant for ITS applications.

We synchronized timestamps on all devices by programmatically updating the internal clocks of the iPhones using a local Network Time Protocol (NTP) server in our laboratory. Using this method, connection and disconnection events as well as sent and received messaged were logged with a negligible clock shift between devices.

For this proof-of-concept our three identical iPhones were programmed with different RSSI thresholds

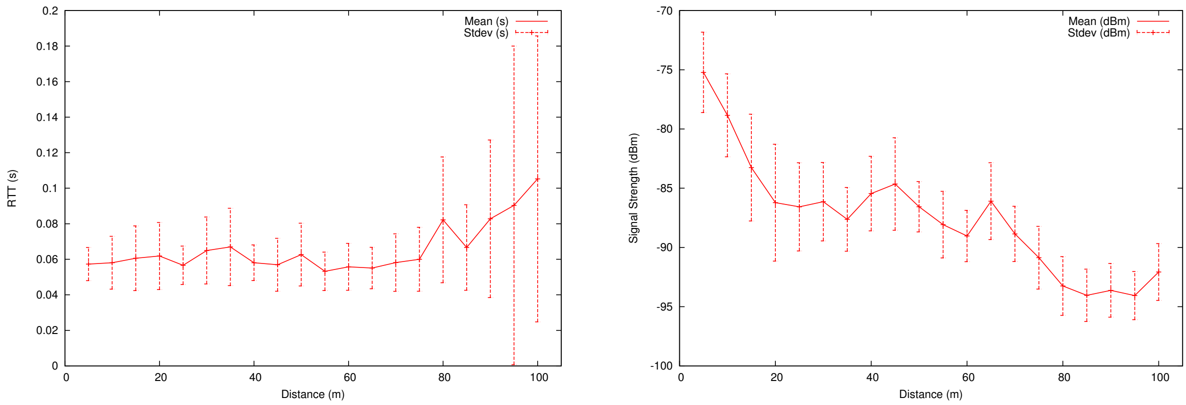

(a) Round-Trip Time

(b) Received Power

Figure 3.3: Results of the Static Experiment.

and connection settings to make sure that, at a distance of 0.5 m between each other, the first phone would not see the third one and vice versa. As we will see later, the second phone had to switch roles after a certain time in order to rebroadcast messages coming from the first device.

3.2.2 Single-Hop Scenarios

To evaluate the performance of our single-hop scenarios we used the logs extracted from our mobile application. Different information is relevant for each of our tests:

Static

This scenario was our first introduction to BLE capabilities [7]. For this scenario, we used two cars on a large outdoor car park, and gradually increased the relative distance between them from 5 m to 100 m in steps of 5 m. At every measurement step, we initiated a communication, as described in Figure 3.2b. The application specified one service with two characteristics. For this preliminary test we used the packet ID as the first read-only characteristic. Upon receiving the beacons, the client device acknowledged reception by writing the received message ID into the second characteristic. This allowed us to measure the Round-Trip Time (RTT) of the data packets (after the connection set-up). The duration of each test run was set to one minute, resulting in around 600 beacons sent per execution.

The results are shown in Figure 3.3. The mean RTT and its standard deviation (Y-axis) for an increasing relative distance (X-axis) are shown in Figure 3.3a. One can see that, up to 80 m, the connection remains stable, with an RTT of around 60 ms. Beyond that point the RTT starts to increase, but the connection remains active. The higher standard deviation above 80 m is therefore due to the increased unreliability of the communication channel. Note that during those tests no packet loss was experienced, up to a distance of 95 m.

Figure 3.3b shows the mean signal strength and its standard deviation (Y-axis) for the same experiment. As expected, the signal strength decreased continuously with increasing distance but

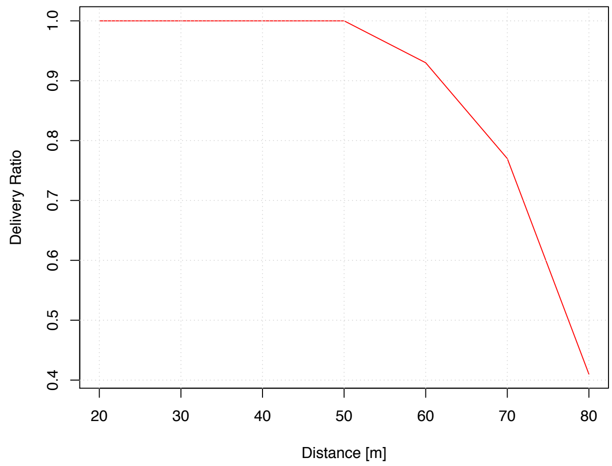

Figure 3.4: Delivery ratio at constant speed for an increasing distance.

remained strong enough to allow the successful decoding of packets at a distance of 100m. The standard deviation shows no severe fluctuation for all the tested distances. We found the RSSI reading to be very reactive with distances up to 20 m but somewhat unstable for longer distances.

Constant Distance & Speed

For this scenario we had both vehicles driving at constant speed on a stretch of highway. After some trials at different speeds (up to 120 km/h), we found that 80 km/h was not only the safest speed according to the highway code (typical stopping distance comparison: [102]) but also the most appropriate, considering the traffic conditions the day of the tests. Furthermore, in contrast to distance, the speed at which the vehicles were moving did not impact the quality of the communication link, provided that the relative speed was maintained at as close to 0 km/h as possible.

We established some rules for this particular scenario in order to minimise human errors:

- Both vehicles' speed was fixed at 80 km/h via cruise control and fine tuned with the help of GPS speeds and relative distance.
- Distance between the cars was constant for the entire duration of each distance-step run.
- If another vehicle came between the two cars, disrupting line-of-sight, results were discarded and the run repeated.

The central device, always in the hindmost car, was in charge of maintaining a constant distance, since its mobile application had feedback on the relative speed and space between the two cars by using the latitude and longitude transmitted by the car in front. The distance between the cars was gradually increased from 20 to 80 m by steps of 10 m. For every step, we initiated communication for at least 180 s. The results of this scenario can be seen in Figure 3.4.

We can see that, up to 50 m, the two devices were able to achieve 100% delivery rate. From this

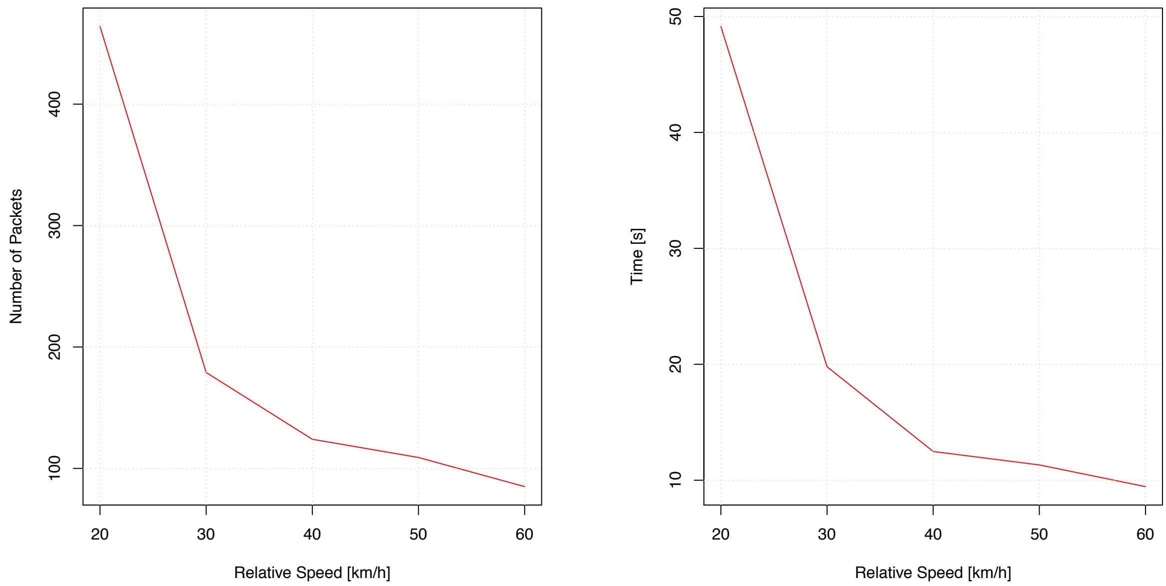

(a) Packets Received

(b) Connection Time

Figure 3.5: Packet delivery and connection time for an increasing relative speed.

point on, the pairing and communication between the two devices should be considered unreliable, as the delivery rate dropped drastically to around 40% at the distance of 80 m. This was due to the way BLE Link Layer handles information regarding connection supervision and which data channels are to be used in the event of a reconnection delay due to a change in the pattern of interference.

Constant Speed

For our second scenario, we took a stretch of road about 300 m long, with one vehicle parked on the side at the half-way point, whilst the second vehicle passed by at constant speed, starting from 20 km/h and repeating with steps of 10 km/h up to 60 km/h.

Figure 3.5 shows how the pairing and message exchange performed at different speeds. When the car was driving by at 20 km/h, the pairing started at 110 m with an overall connected time of nearly 50 s before going out of range. This gave enough time for the peripheral device to transmit 464 packets.

With increasing speed, the number of packets acknowledged and the connected time reduced significantly. At 30 km/h, the connection lasted for 19.8 s for 179 packets, continuously reducing until the 60 km/h test, where the connected time was 9.5 s and 85 packets were transmitted.

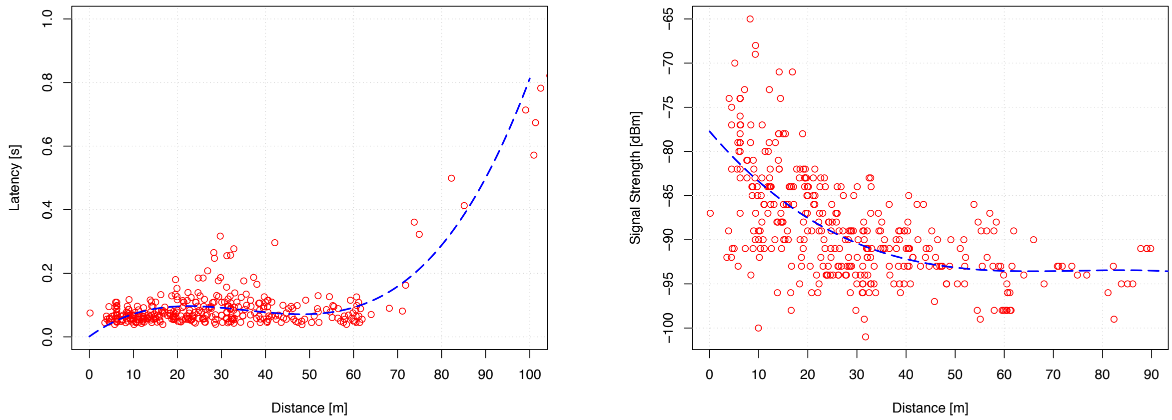

(a) Round-Trip Time

(b) Signal Strength

Figure 3.6: City Scenario Results

City Driving

This final single-hop scenario did not present driving constraints, and was intended to test BLE performance in an urban scenario. For this test, both vehicles drove on a busy road in Luxembourg city, continuously varying distance and speed and sometimes having other traffic interposed.

Figure 3.6a shows the RTT distribution and trend line (polynomial regression of degree 3) for different distances (X-axis).

Depending on different environmental radio interference impacting the signal propagation, variations in latency were present at any given speed. This graph shows a stable exchange up to 50-60 m with some messages also acknowledged in the range of 100 m. This graph shows a similar behaviour to our static experiment observed in Figure 3.3a.This experiment, having a larger number of environmental variables due to its nature, shows a clear cut of transmissions after the 60 m mark.

Figure 3.6b shows the RSSI with trend line (also degree 3) for different distances (X-axis). This reading varies greatly from device to device. Even with identical phones, it proved to be a challenge at times to have consistent values. Nevertheless, taking both interference and signal propagation into account, the overall fit clearly shows a degrading signal strength up to 50 m, which confirms the results from our RTT calculations. Moreover, similarly to the static experiment (Figure 3.3b) we notice how from 20-30 m the signal strength reaches a level after which it's nearly indistinguishable regardless of the distance.

Figure 3.7 shows the packet rate (10 Hz) and distance between the vehicles over time (Y-axis) for the duration of the scenario (X-axis). Here, we can see how the BLE protocol behaved in a vehicular environment. Almost no messages were lost, up to 60 m.

With the distance increasing beyond the 70 m mark, at around 300 s, we notice that the signal quickly degraded to the point where it is completely lost, and the packet rate fell to zero. Communication automatically resumed after the devices are again in range and re-pair.

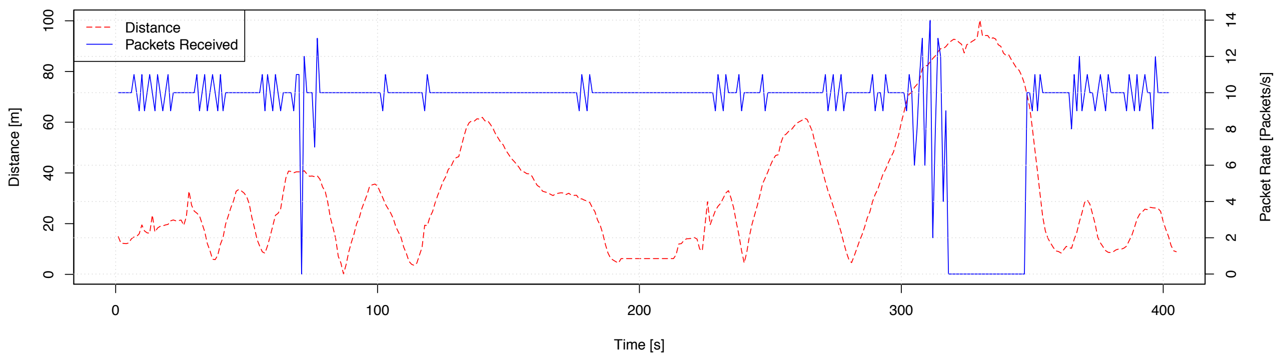

Figure 3.7: Distance between Peripheral and Central vs. Packet Rate.

In this graph we can also note that the packet rate fluctuated at times (e.g. between 100 and 200 s), showing how a message was retransmitted after a loss. This behaviour is better understood by looking at how the BLE Link Layer handles retransmissions [28]. Connections at the Link Layer are based on a stop-and-wait flow control mechanism regulated by cumulative acknowledgements, providing the protocol with error recovery. Each data channel packet header contains a Sequence Number (SN) and a Next Expected Sequence Number (NESN). If a packet is received by a device, the NESN of its next packet is increased, serving as negative acknowledgement of the current packet. This behaviour allows a BLE link to provide robust communication when obstacles such as other cars, obstruct the line of sight between the two vehicles.

3.2.3 Multi-Hop Proof-of-Concept

To evaluate our multi-hop scenario, we named our three smartphones device A, B and C, as shown in Figure 3.8.

The purpose of this test was to measure how much time an event-triggered message from device A would need to reach device C.

We initiated the communication by triggering a message simulating, as an example, a hard braking event. After switching to peripheral mode, device A searched nearby central devices. After pairing, in our case, to device B, device A started updating a readable characteristic with a sequential ID at the same frequency as the single-hop experiments (10 Hz). Upon reception of the first message, device B switched itself automatically to peripheral mode and started looking for new neighbours. Whilst it searched for other central devices, device B still kept receiving the remainder of the messages from device A. Once device C paired with device B, it started to forward the 100 messages initially sent by device A.

The performance of this scenario is shown in Figure 3.8. At t=0 ms, we started advertising with device A in peripheral mode. Device B immediately discovered it (t=39 ms), performed pairing (t=95 ms) and received the first message (t=1370 ms).

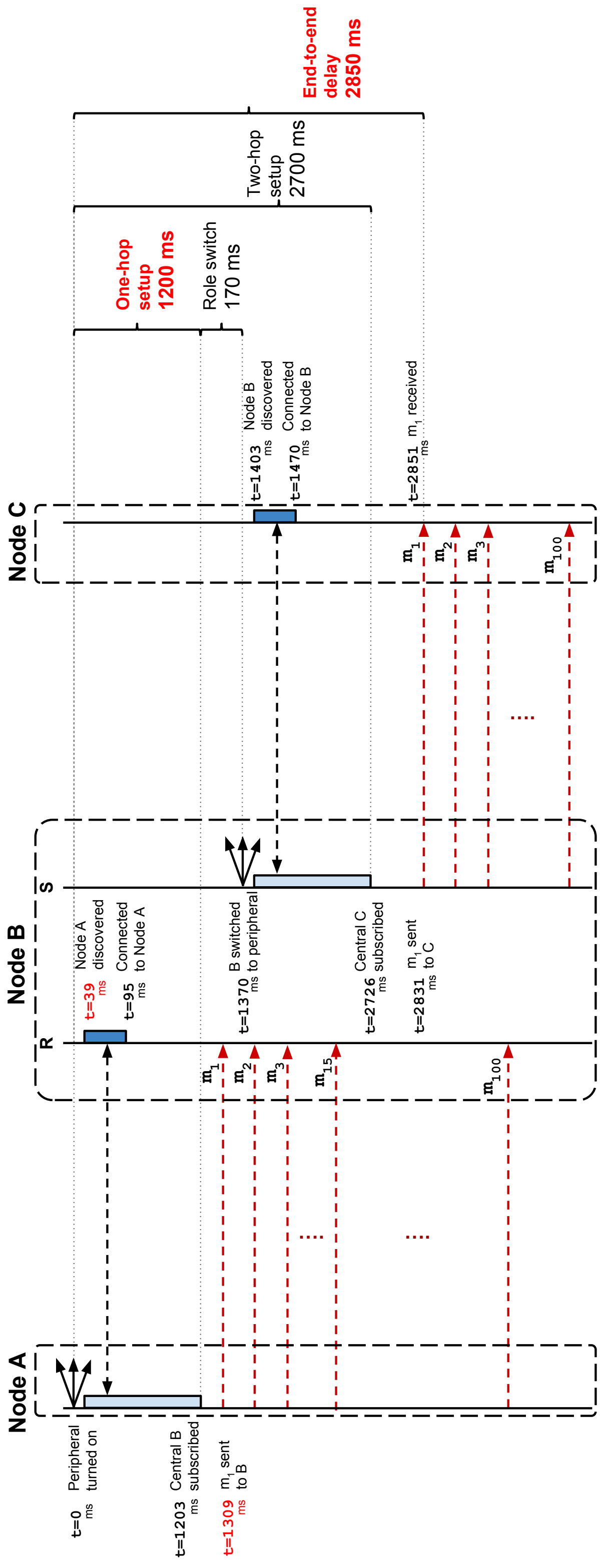

Figure 3.8: BLE Multi-Hop Message Exchange Graph

This delay between pairing and reception of the first message is determined by how the central device discovers and subscribes to the peripheral's characteristics according to the vendor-specific implementation of BLE within the Core Bluetooth framework. The pairing process takes around 1.2 s and must to be performed twice: once from device A to B, and again from B to C. This means that from the instant we triggered the event on the first device it took roughly 1.37 s for device B to acknowledge this message and turn its Peripheral manager on.

The end-to-end delay for a two-hop exchange, from the moment we triggered the event to the reception of the first message by device C, was around 2.85 s, which is in-line with our previous one-hop timing.

Nodes that are in peripheral mode switched back to central only after the 100 messages had been sent to all paired devices.

All three BLE devices were able to handle multiple communication attempts by keeping a local table of nearby neighbours to avoid redundant reconnections. If two simultaneous attempts at connection from two different peripheral devices were received at the same time by a central device then it had a probabilistic chance of connecting to either, based on their duty cycle (probability that the master finds the slave during one of its advertisements [35]). Once a communication link has been established multiple connections could coexist thanks to the Adaptive Frequency Hopping (AFH) scheme and the way the protocol handles interference (See also in Section 3.3).

3.2.4 Discussion

Communication via BLE can also be achieved without pairing [28] by advertising data through the three advertising channels. Unfortunately there is currently no way of doing this on a smartphone. The closest technology match is Apple's iBeacon [103], which uses advertising channels to transfer information but does not provide the tools to insert a data payload into the messages. Tracking the IDs of sent and received messages is vital for our application, so that it can know which message to rebroadcast to a given destination, making iBeacon inapplicable.

These results show that BLE, using bidirectional data communication that requires pairing, is not an optimal solution for delay sensitive safety applications such as signalling sudden emergency braking or a danger on the road ahead. The end-to-end delay is too long, mainly due to the pairing, to apply this technology at high speeds on a highway. Each pairing hop requires, with the current hardware, roughly 1.2∼1.4 s, making its use impractical. The pairing process could be optimized by reworking Apple's implementation and studying connection specific timeouts in the characteristics discovery phase.

What BLE provides is a relatively fast, low-power and very reliable solution for IrVC in dense traffic situations or when the pairing delay is not important (e.g. non-safety applications).

3.3 Robustness Analysis

As previously demonstrated, the current version of BLE might mainly be applied to low relative speed vehicular scenarios (e.g. city intersections, urban environments). We took the study of its behaviour in a congested RF spectrum as a challenge by reproducing an interference testbed in our laboratory. Previous work [104] shows how, in highly populated areas, one can easily detect more than 200 APs occupying mainly the non-overlapping channels 1,6 and 11. It is thus interesting to further study how a similar environment could affect BLE communications. As shown in Fig. 3.9, the interference testbed was comprised of six Raspberry Pis, three in AP mode assigned to channels 1, 6 and 11. This drastically reduces the number of BLE data channels available to the protocol, while still leaving the advertising frequencies (channels 37, 38, 39) unperturbed. Amongst other things, we were able to observe how the association time between the two mobile phones varies with the different scenarios.

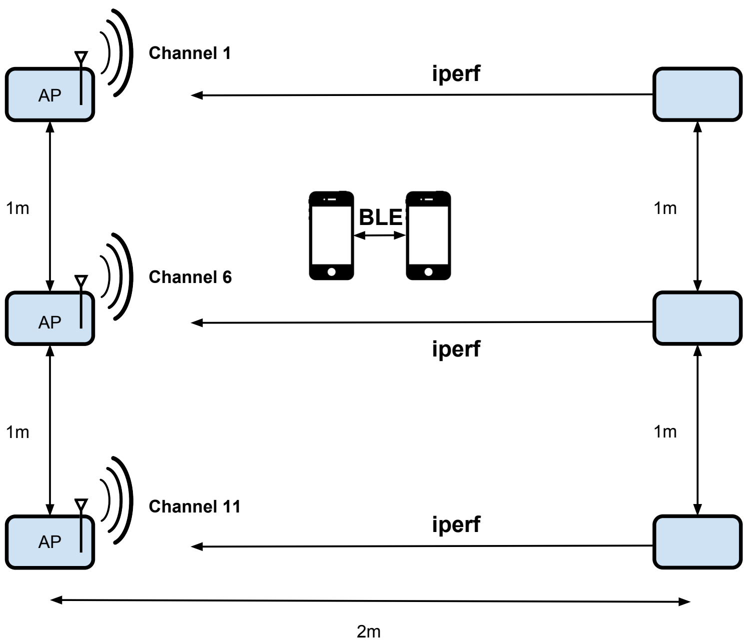

Figure 3.9: Raspberry PIs Interferences Testbed

For the purpose of benchmarking the behaviour of BLE in an environment with a sufficient amount of interference we built a relatively compact (2 m x 2 m) testbed.

On the APs we installed *hostapd* and *dnsmasq* to configure an IEEE 802.11g radio and assign different channels (1, 6, 11). To generate a sufficient amount of traffic to perturb communications, we utilized *iperf* from the clients to the APs to create a maximum-throughput stream (≈ 35 Mbps).

For the BLE communication, the mobile devices were running modified software from our previous experiment, in order to log packet exchange during a time frame of one minute. The mobile phones where physically placed in the middle of the testbed to keep the scenario relevant to our purpose and maximize interference.

We used our best efforts to optimize the test environment to ensure no other human or environmental interference would spoil the results.

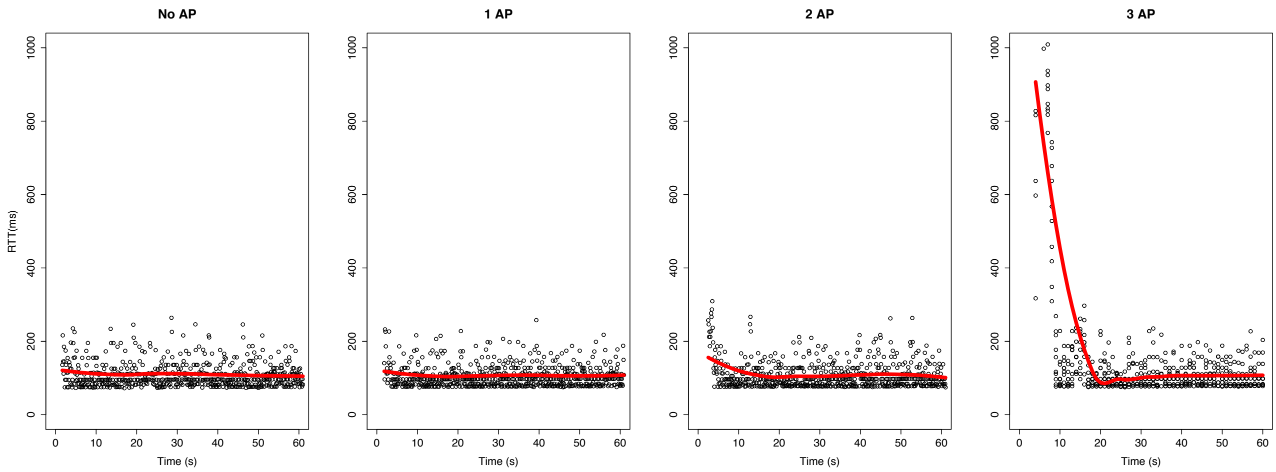

Figure 3.10: BLE Round Trip Time

3.3.1 BLE Interference Results

To maximize the throughput of the BLE communication, we chose a constant packet size of 158 Bytes (three for the header and 155 for the payload) at a rate of 10 Hz, allowing us to achieve optimum reliability and stability, given our hardware (≈ 12.6 Kbps). This packet size was the maximum that was achievable with our combination of hardware and software (iPhone 6 - iOS 8.1).

We defined four different interference scenarios:

- No Access Point
- One Access Point with iperf traffic (Channel 1)
- Two Access Points with iperf traffic (Channels 1, 6)
- Three Access Points with iperf traffic (Channels 1, 6, 11)

Keeping in mind the channel mapping between BLE and Wi-Fi shown in Figure **??** in the previous Chapter, we notice that, excluding the three advertising channels, the number of free BLE channels without interference is reduced to a minimum of nine in the case of three APs. In Table 3.1, we show the results of a one minute BLE communication for each scenario. The experimented time frame includes the association time required by the two devices before exchanging packets.

We notice that, even though the BLE advertising channels are theoretically not congested, the initial association time is longer in scenarios with higher interference. With no active AP, we recorded an average association time of 1.5 s, which reflects our previous driving scenarios results, where there was minimal interference. This behaviour can be interpreted as adjacent channel interference [105], as well as environmental reflections affecting the advertising channels.

We also notice that, even though the packet delivery rate was always 100%, the number of sent and received packets decreases because of the longer association time.

By further investigating the RTT behaviour in Fig. 3.10, with the help of a local polynomial regression fitting function, we observe that the set-up was apparently unaffected by interference coming from only one access point, but by adding a second, and particularly a third source of frequency conflict, the first batch of packets was initially delayed and then retransmitted later. In these circumstances we can assume an occupied data channel was picked for communication after association, explaining the initial higher delay. The chances of landing on a free channel are lower for our three AP scenario (only nine of 37 channels are free). As soon as the packet exchange starts, for every BLE connection interval (≈ 30 ms for iPhones) a new, better, channel was selected by using the AFH technique. The hopping sequence repeated until all missing packets had been retransmitted and an optimal link was found. For this reason, Table 3.1 also includes mean and standard deviation, excluding packets received before a stable data channel is discovered (first ≈ 2.5 s for zero and one AP scenarios, ≈ 5 s for two APs and ≈ 10 s for three APs).

From this data we can conclude that BLE, once a stable configuration is reached, is very resilient to interference.

3.3.2 Bluetooth Classic and Interference

In a similar way, we wanted to compare how Bluetooth Classic (BC) performs in the presence of heavy interference. Since we know that BC is mainly used for audio (phone calls and music) streaming [50], we tested both Synchronous Connection-Oriented (SCO) and Asynchronous Connection-Less (ACL) link modes by establishing different communication channels between two Bluetooth devices. For this test we chose a different approach. Since it is fairly easy to limit the channel-set utilized by BC, we reduced the operational frequency range of the protocol to achieve approximately the same ratio as BLE between channels with and without interference (nine out of 37). To achieve this, we reduced the number of BC operating channels to 24 whilst still leaving six free of interference (from BC channels 23 to 47). The goal was to generate interference only on one Wi-Fi channel (Ch. 6), while leaving some free channels for BC to attempt to communicate.

We ran four scenarios with identical time frames to show delivery ratios of both SCO and ACL communications with and without interference. In Figure 3.11 we can see a representation of Bluetooth packets grouped by channel. We notice immediately that the AFH scheme tries its best to avoid channels where Wi-Fi is present by using the only the six free channels on either side. Furthermore, as we can see in more detail in Table 3.2, the delivery ratio drops to 58.47% and 43.6% for SCO and

Scenario	**Association Time (s)**	**Sent/Received Packets (Delivery ratio)**	**StDev RTT (s)**	**StDev RTT (s)***	**Mean RTT (s)**	**Mean RTT (s)***
0 AP	1,56998	594 (100%)	0.03445	**0.03412**	0.10959	**0.10911**
1 AP	1,71192	592 (100%)	0.03539	**0.03079**	0.10894	**0.10613**
2 APs	2,39517	586 (100%)	0.03924	**0.03368**	0.11014	**0.10663**
3 APs	3,98211	570 (100%)	0.24149	**0.03849**	0.18194	**0.11121**

*Excludes packets until optimal data channel is found.

Table 3.1: BLE Interferences Results

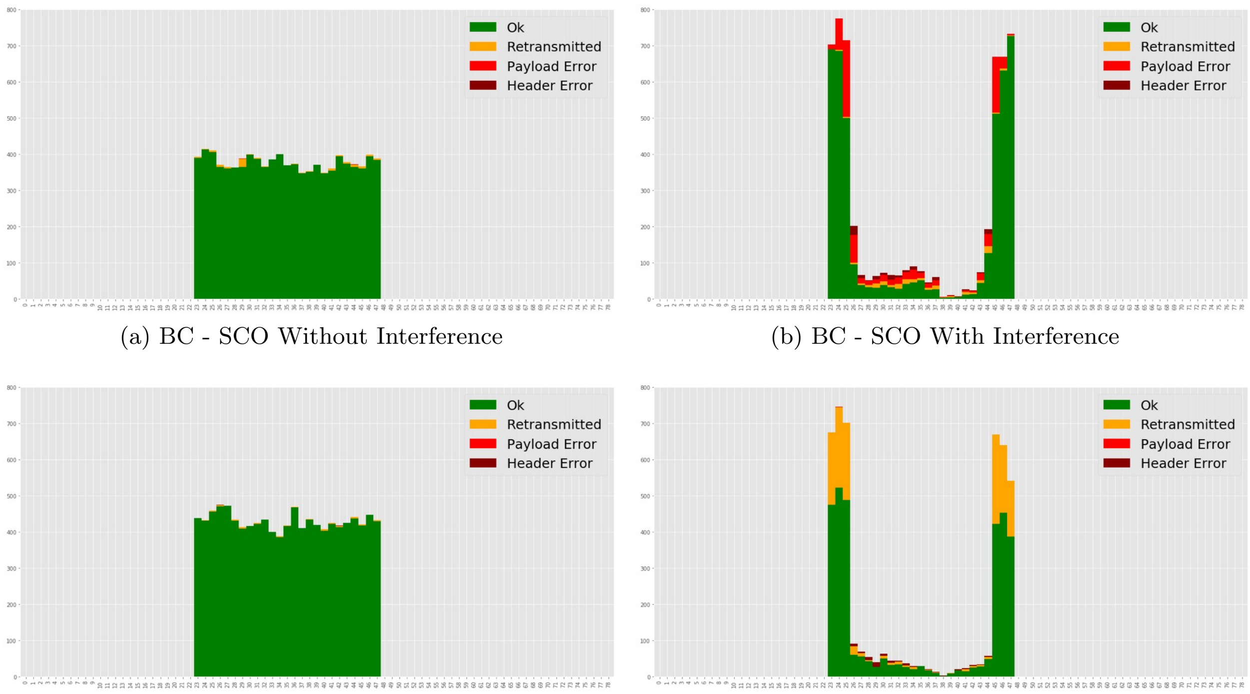

(a) BC - SCO Without Interference

(b) BC - SCO With Interference

(c) BC - ACL Without Interference

(d) BC - ACL Without Interference

Figure 3.11: BC Interference Scenarios (X-axis represents channels from 0 to 78)

Scenario	Sent/Received Packets (Delivery ratio)	Ok Packets	Retransmitted Packets	Payload or Header Error Packets
Without Interference SCO	9479 (100%)	3988	87	4
With Interference SCO	5542 (58.47%)	4441	159	942
Without Interference ACL	10765 (100%)	10715	47	3
With Interference ACL	4693 (43.6%)	3302	1313	78

Table 3.2: BC Interferences Results

ACL modes respectively, showing that the protocol cannot keep up a consistent communication link.

Although both BC and BLE use the same AFH scheme for coexisting with other signals, BC performance is lower, due to the bandwidth requirements needed by the protocol to enable audio streaming.

3.4 Conclusion and Final Remarks

In this chapter, we have studied the potential of Bluetooth Low Energy (BLE) for Inter-Vehicle Communication (IrVC). We developed a mobile application specifically intended to work for hardware already on the market, and which implements multi-hop communications between moving vehicles.

We measured performance in terms of delivery ratio and round-trip time for multiple single-hop scenarios and end-to-end delay for a multi-hop proof-of-concept application. We also analysed the impact on interference on BLE by examining its coexistence capabilities with other IEEE 802.11 devices on a crowded 2.4 GHz radio spectrum. We have shown that BLE is more resilient than other technologies due to its properties, and has more or less the same range. Its only weak spot (for low data applications) lies in the time required for two devices to associate.

BLE provides a fast, low-power and very reliable solution for IrVC in dense traffic situations, or as a back-up to other technologies (DSRC/ITS-G5). Possible applications include all scenarios where delay-sensitive data delivery is not a requirement (e.g. intersection, city traffic management).

As we previously discussed, BLE is a fast-evolving technology. At the time this research was carried out, the capabilities of BLE were constrained by the available hardware implementations. The technology has since evolved, and a new case can be made for using BLE more actively as a Vehicle-to-Vehicle (V2V) communication technology and with more sensitive IrVC applications.

With this prospect, further investigation is needed regarding how the new Bluetooth version (5.0) behaves in respect to association delays between devices. Moreover, since the amount of payload data transferable within an advertising packet has also been extended by supposedly 800% [30], broadcasting information, completely eliminating pairing delays, can again be considered as a faster way of sending information around. However, we advise caution with such techniques since, by only relying on the three advertising channels, susceptibility to outside interference or jamming could easily become problematic.

In the next Chapter, we present a means of exploiting Bluetooth's discovery capabilities as an information medium, in contrast to this Chapter's communication-based approach. In more detail, we focus on using Bluetooth for passively capturing nearby beacons by exploiting its application-centric nature. After describing a Bluetooth-focused wardriving [1] experiment, we propose a novel approach to classifying environment using Bluetooth in a mobile context.

[1] Bluetooth wardriving is the act of searching for Bluetooth radios by a person in a moving vehicle, using a laptop or a smartphone.

Chapter 4

Bluetooth as an Information Medium in a Mobile Context

Following our previous contribution, where we explored Bluetooth Low Energy (BLE) as an *active* communication technology in mobile contexts, this chapter describes a different, connectionless approach, using Bluetooth as a medium for information rather than purely for communication.

Bluetooth's beaconing nature is application-centric compared to other technologies, such as Wi-Fi, that are more network-centric. With BLE beacons, we are advertising an application (some data) whilst for Wi-Fi we advertise access to a network. BLE overall provides more information, allowing automatic registration to services over the air, in contrast to alternative technologies, which simply provide information for network access.

Our goal in this chapter is to outline how we can exploit nearby Bluetooth radios for human mobility-related applications (*What can we find out about the world by listening to it?*). As an example, it is becoming increasingly important for objects or vehicles to be aware of their environment. Having knowledge of nearby Bluetooth devices (e.g. inside and outside vehicles) can provide the listening device or vehicle with enough data to learn about its surroundings. This allows such an application to provide a hybrid location system for Internet of Things (IoT) devices with limited resources.

As a first step we analyse multiple datasets of Bluetooth Classic (BC) and BLE discoveries collected during a series of data collection campaigns.

By changing the mobility of the observer device we sequentially explored scenarios with increasing mobility, from walking to hybrid (walking & bus), and finally driving. For all the data collection campaigns we examined a variety of data about encountered devices, such as Global Positioning System (GPS) location, quantity and quality of signal, and specific device class information. In a statistical analysis we show the distinctive behaviours between BC and BLE relative to context. We ultimately focus on a further analysis of our most complete dataset, which includes Bluetooth-based traces collected by 20 participants.

In a vehicular scenario, where mobility is high, changes in environment can occur more frequently and can therefore be applied to providing context-aware information to the user. For this purpose, we propose a set a features to train a classifier for the recognition of different driving environments (i.e. road classes).

Such a classification should be capable of outputting a certain driving context when raw Bluetooth discovery data is provided as input. The purpose of this work is to benchmark the precision with which we can estimate *in what type of environment* we are by looking at *which devices* are around us. By the means of Machine Learning (ML) algorithms ,we score different configurations using diverse sampling methods. Using Bluetooth instead of Wi-Fi or GPS provides a privacy-friendly and relatively low-energy way of approaching the many problems that affect mobility worldwide. Specific applications could see this technology used for profiling driver behaviour (travel time estimation), passenger flow estimation (multimodal transportation), and for conducting general human mobility studies. In this thesis we focus mainly on the case of road class estimation.

Comparing the performance of our classifier with different sampling parameters to fine-tune our feature selection, we are able to predict with reasonable confidence general environments such as highway, city and extra-urban by using only discovery data and no geographical information.

In the following sections we explain the thinking behind our transition between implementing Bluetooth in an active scenario to its application as a passive sensing technology. We further describe the sequence of all experiments that led to us applying ML to a relatively new technology.

4.1 From Active Usage to Passive Scanning

In Chapter 3 we benchmarked BLE and its direct use in vehicular communications. Although it is clear that range and performance at different speeds are definitely strong points of BLE, the initial association delay between two new devices is something that heavily restricts its uses.

Using BLE for safety applications, such as relaying messages about emergency braking or obstacles on the road, is therefore not recommended.

We conclude from these results that BLE could play a more important role if its usage is expanded to more than Vehicle-to-Vehicle (V2V) communications. BLE is in fact a perfect candidate for low-to-medium speed Vehicle-to-Everything (V2X) communications and can be exploited for data exchanges with devices outside and in the vicinity of the scanning device. Scenarios such as intersections, roads with bike lanes and pedestrian populated roads are perfect candidates for this new communication paradigm. Knowledge about the presence of nearby Vulnerable Road Users (VRUs) can be extremely valuable in many decision-making situations.

As a first step we took the challenge to see what devices were actually out there; what quantity and type of Bluetooth signals we can discover doing daily tasks such as walking, driving or taking the bus. For this purpose we defined different scenarios with distinct mobilities with the goal of testing the characteristics of typical environments such as urban, countryside or highway.

We analysed at first the behaviour of BLE in different walking scenarios, such as in a shopping centre and in the city. We then analysed a trace from a hybrid (walking + public transport) trip and finally collected data from a driving scenario. Although we had promising results from these tests we decided to go one step further and expand BLE discoveries with BC discoveries as well. Doing so we further broadened the quantity of potential nearby devices and effectively improved the original driving scenario.

Our first step in this approach was to launch a data collection campaign within Luxembourg. By the means of an Android mobile application, participants were asked to use the App while driving in order to discover Bluetooth (Classic and Low Energy) devices.

4.2 Sensing Systems and Context Awareness

As we saw in Chapter 2, the use of Bluetooth is wide-ranging. In this Chapter we will focus on its use in sensing systems. Such systems generally allow the collection of large amounts of data, where the most conventional parameters are a sampling and recording frequency that vary according to the metrics collected and the sensors used. This type of system is primarily passive (or opportunistic) and allows the collection, analysis, processing and storage of a set of data. The numerous architectures that have been proposed for this type of system generally vary according to the technologies considered, the desired level of privacy and security, and more importantly the needs of the application (e.g. [106,107]).

With smartphones, the easiest and most flexible equipment to use in research, there are three main categories of metrics [108]: 1) Motion Metrics - usually transcribed by accelerometers or GPS; 2) Physiological Metrics – less prevalent but becoming increasingly popular, as with heart rate sensors that can go as far as inferring the stress level of a user; and 3) Contextual Metrics - making it possible to evaluate various characteristics of a users physical environment. Note, however, that contextual metrics require microphones and cameras to continuously process sound and images. Therefore, just like GPS, they expend an inordinate amount of battery life. They must also be accurately positioned on the user.

Recent studies have shown that the use of network traces makes it possible to understand different contextual and environmental characteristics. For example, Wi-Fi traces can be used to identify places [109] or to understand different mobility patterns [110]. Bluetooth traces, at the heart of this thesis, are also promising as they allow not only to measure interactions between devices, but also to retrieve details (varying depending on the version) of the nature and properties of devices.

Using BC for context-awareness was experimented before the introduction of BLE, for detecting daily activities [111] as well as for building context specific privacy profiles for mobile communications [112].

The aim of this part of our research is to outline ways in which a technology as simple as Bluetooth can be used to identify complex environmental characteristics in a vehicular context.

Our claim is that the quantity and the quality of traces obtained by Bluetooth are strongly dependent on the mobility (or speed) of the user. If this assumption is verified, it is certainly possible to identify these locations or certain mobility characteristics based solely on Bluetooth traces, which do not consume energy and which have the great advantage of preserving privacy.

The way Bluetooth works makes it so that from a vehicle perspective we can passively scan all nearby devices. Once discoverability is enabled, nearby objects can be tracked sometimes even if they are not being actively used.

Currently, there are different studies that target vehicular communication with VRUs [113] with a particular focus on pedestrians and bicycles [114] but, to the best of our knowledge, none using Bluetooth to its fullest extent with BLE.

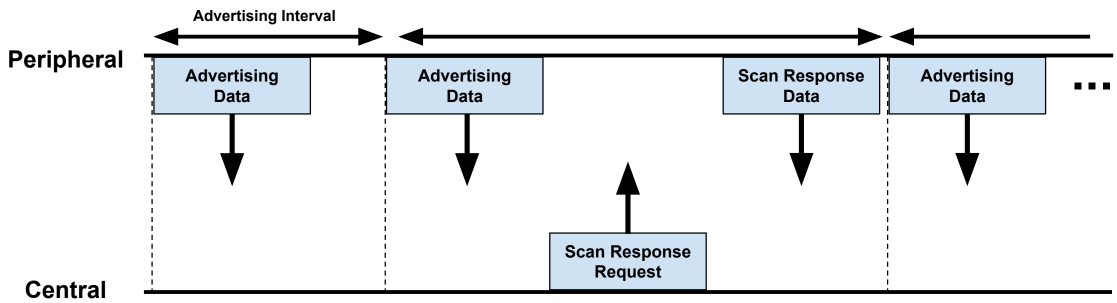

Figure 4.1: BLE Advertising Process

4.3 Low and Hybrid Mobility Field Experiments

For the beginning of our experiments we decided to sniff BLE advertisements whilst on foot to get an idea about the class of devices that are most commonly found. Advertisement packets include information such as MAC address of the device, received signal strength, details about connectivity, manufacturer data and others (more specifics in Chapter 2). Secondary details such as device name and services, as shown in Figure 4.1, are only received by answering to these advertisements with a scan request. These requests do not need pairing and are sometimes lost when the discovered device goes out of range or disappears before receiving the scan request. This is the reason why in the following experiments we cannot always resolve the names of all discovered devices.

The hardware used for these experiments was comprised of a Raspberry Pi 2 with Bluetooth and GPS dongles powered by a 2 amp output, 10000 mAh, power bank. The system would automatically start sniffing BLE when started and record traces locally. Through the BLE advertisement packets we can extrapolate unique information about device names and manufacturer data.

The first walking survey was inside a shopping centre. To best cover the area and maximize the amount of discoveries, we followed a snake walking pattern. As expected, most of the discoveries were in the electronics department where the newest consumer gadgets are commonly found. During this test we found 55 unique devices to which 47 answered with a scan response providing us with the device names.

Here are few examples:

- UE65JU7000 (Television)
- s97 (Traking)
- Flex (Fitness band)
- Surge (Fitness watch)
- BKB50 (Keyboard)
- dica.lu/2AUCHANXXXX (Digicash - payments)

We notice from this list how BLE beacons can be found not only in wearables, but also in computer

(a) Bus trace path (b) BLE traces

Figure 4.2: Bus trace

peripherals and televisions. Moreover, the most interesting item in this list is the last one. This device (Digicash) is in fact a payment terminal used at the checkouts and was discovered multiple times most likely in the proximity of checkouts. This first analysis already displays to which extent BLE radios are integrated in off-the-shelf devices.

For the following experiment we decided to combine public transport and walking in the city centre of Luxembourg. This trace, shown in Figure 4.2, is comprised of a trip on the bus and a walk around the city during lunchtime. Time and visited streets were specifically selected to pass the maximum number of shops and restaurants. On top of information about the discovered devices, this time we also had GPS data recording the location of each discovery.

By observing the BLE traces in both these scenarios we can reconstruct with good accuracy the predefined path. An interesting observation during the public transport trace (in our case a bus), is one device with a very strong signal strength (green trail) which is present from the start but suddenly disappears. This event would suggest that this device, a smartband, was owned by a nearby passenger who happened to exit the bus at the station where it disappeared from the trace. Such information could be crucial in studies around public transport such as passenger flow and behaviour estimation.

The trace around the city, shown in Figure 4.3, revealed the following statistics: 160 unique devices discovered to which only 40 replied with the name details. Nevertheless, manufacturer information is revealed by the advertisement packet directly and can therefore be used to partially identify the device.

Here are the most interesting discoveries:

- ALCATEL ONETOUCH POP C7 (Smartphone)
- dica.lu/2PHAXXXXX (Digicash - payments)

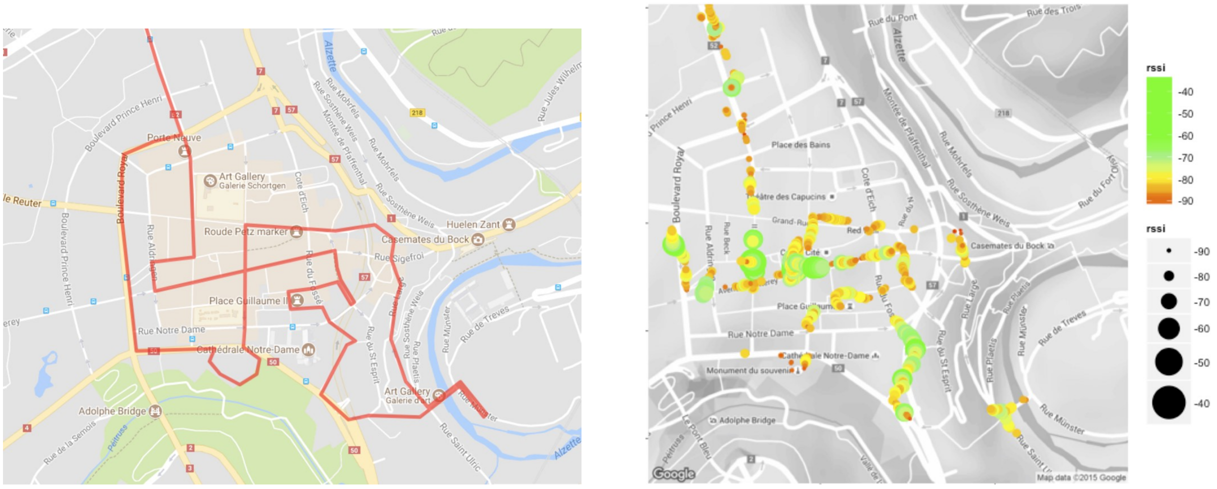

(a) Walking trace path

(b) BLE traces

Figure 4.3: City walk trace

- Chipolo-DXXXXX (Tracking)
- UE48J5510 (Television)
- IDD-212B XXXXX (ODB Car dongle)
- Kontakt (iBeacon)
- Orbit (Fitness band)
- GALAXY Gear (3E88) (Smartwatch)
- One (Headphones)

The majority of devices without name come from the same types as the ones with names. These are mainly smarphones and smartbands. As for our previous shopping centre scenario, we discover similar categories of devices such as televisions, headphones, trackers, payment terminals, as well as general purpose beacons. One of the most interesting discoveries for this trace is the On-board Diagnostics (ODB) car dongle. This device plugs into a vehicle to serve logging purposes for the on-board computer. We can safely assume that this kind of device represents an operating vehicle since the beacon wouldn't have been active otherwise.

This first part of our field experiments was aimed at validating the *quality and type* of BLE discoveries. Knowing the current reach of BLE is essential for answering our initial research question. Results are very promising as a wide variety of Bluetooth devices were discovered in different environments.

For the following section we will be exploring new scenarios different mobility profiles. We will initially use the same sniffing hardware for a high mobility driving scenario and following, we discuss and implement improvements to such solution by the means of a data collection campaign.

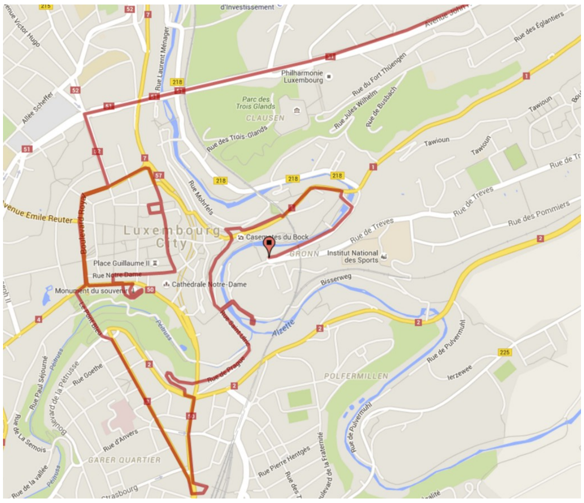

(a) Driving Trace

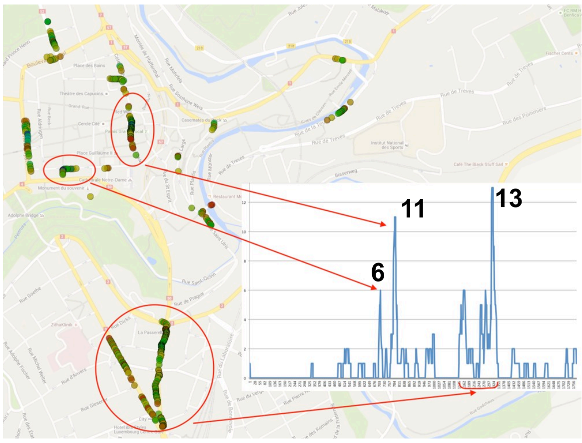

(b) BLE Simultaneos Discoveries

Figure 4.4: City walk trace

4.4 High Mobility Field Experiment

Having established that BLE beacons can provide relevant information about each discovered device, we decided to follow up on our previous low mobility experiments with a vehicular approach. This time we equipped a car with our hardware and drove through populated areas around the city of Luxembourg. Figure 4.4a shows the GPS trace from this trip.

During this experiment we focused not only on the type of discoveries but also the quantity. Figure 4.4b shows how in some areas, which correspond to highly populated streets with shops, we discovered up to 13 devices simultaneously within communication range. The device types are very similar to our previous finds with wearables as most popular but also payment terminals and television being very present in certain streets. Interestingly enough one beacon (BLE Edge 1000) was generated by a bike specific GPS device, indicating once more, that with this methodology it is possible to discover a multitude of accessories applied in different environments.

We notice from this experiment how roads in between our Bluetooth hot-spot areas do not have any discoveries. The assumption is that these roads are mainly populated by cars and are therefore nearly void of BLE traces.

Cars in general, to this date, rely on BC for audio transmission between smartphones and the hands-free kit in the vehicle. For this reason we decided to include BC specific discoveries for our future experiments.

As we discussed in Chapter 2, BLE's discovery is fundamentally different from BC's and therefore our BLE set-up is not capable to handle BC discovery in its current format.

To achieve our goal of capturing both BC and BLE traffic, as well as establishing a system that could be deployed in scale for further testing, we decided to build such a system in the form of an

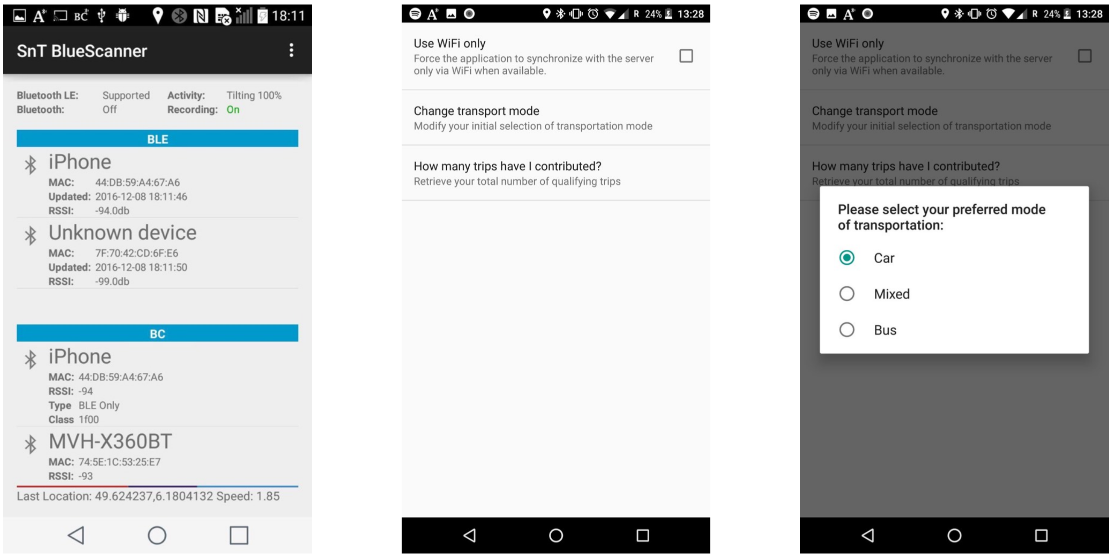

(a) Main screen (b) Settings page (c) Setting page (2)

Figure 4.5: The BlueScanner App

Android smartphone application.

The mobile app, named *BlueScanner* (Figure 4.5), is available for free under a MIT License on GitHub [1].

In order to a bigger number of traces we launched a data collection campaign advertised internally within our research group as well as to friends and family. Overall, we deployed the application on 20 smartphones from different brands belonging to the participants of the study.

20 users was a statistically sufficient amount of people, in our case, for conducting our empirical and exploratory analysis. A higher number of participants would have resulted into higher running costs outside of the research budget. Our goal was not to profile participants but to look at the network, the type and amount of data we can extract from it and observe unique features correlated to the kind of discovered Bluetooth devices.

For a period of two months, we collected data from daily drives. The length of the data collection campaign, although not being relevant in the scope of the process, was chosen to allow enough time for the participants to collect a sufficient amount of data without having them rush or work towards a specific goal or result.

4.4.1 Sensing System

We decided to implement our sensing system on devices which are widely available today and come with a multitude of sensors. For this purpose, we developed a mobile application which was ideal for this study. We wanted to use a device that users already own and that could be programmed with a

[1] `https://github.com/wbronzi/bluescanner`

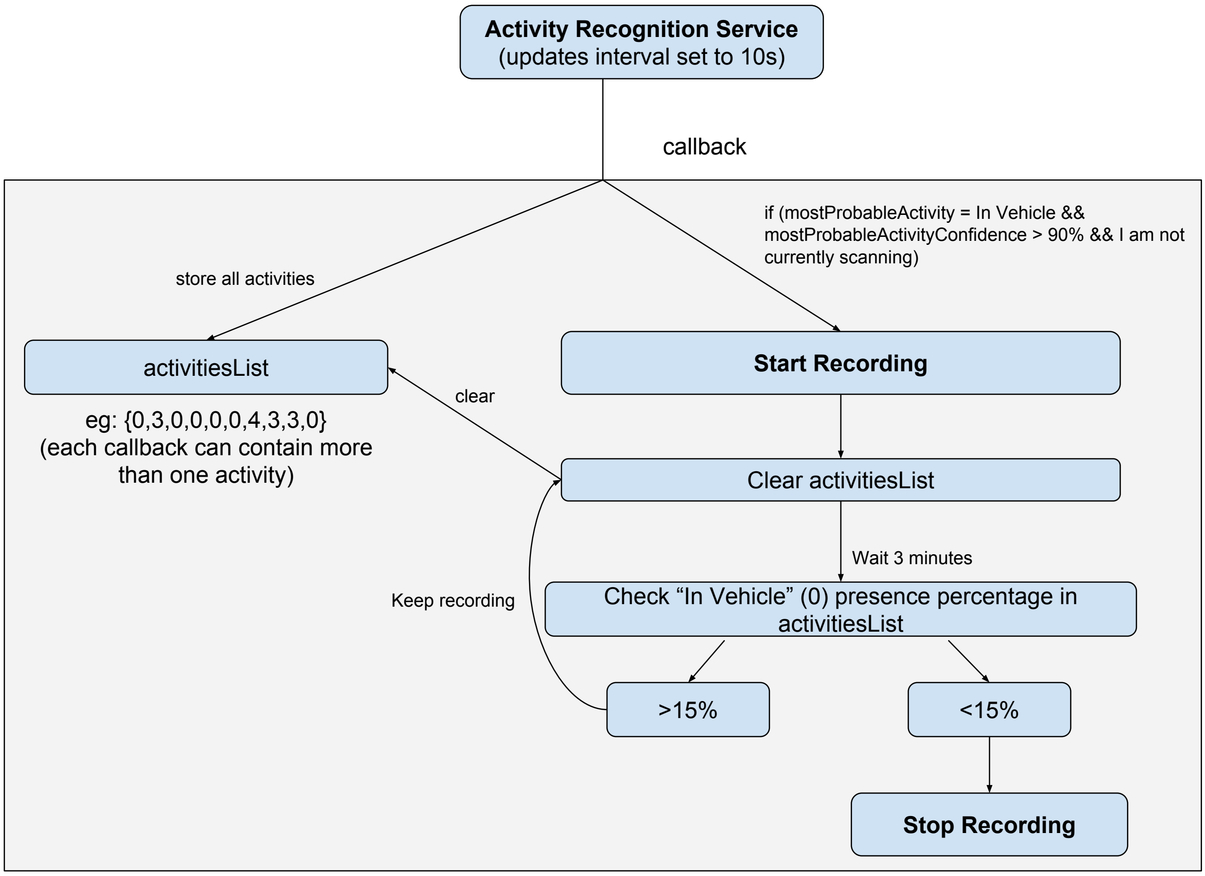

Figure 4.6: Session handling procedure with activity recognition

certain flexibility and control over the Bluetooth hardware. Our implementation is obviously not on a par with more expensive sniffers such as the Ellysis Explorer [2], but provides enough precision to prove our claim and to enable concrete and easily deployable applications.

Figure 4.5 shows screenshots of the application developed for the occasion. In particular Figure 4.5a portrays the main screen presented to the user during the sniffing process and finally Figure 4.5b & 4.5c show the settings page. Users could choose in which way to upload traces, either via Wi-Fi or by using cellular data, retrieve their total number of uploaded traces, as well as specifying their preferred mode of transportation (car, bus or mixed).

The application has been developed to run passively and scan the environment during driving. In order to minimize polluting the sessions and to keep the user experience fast and seamless, we developed a logic able to recognize the driving activity from our participants and automatically start

[2] `http://www.ellisys.com/products/bex400/`

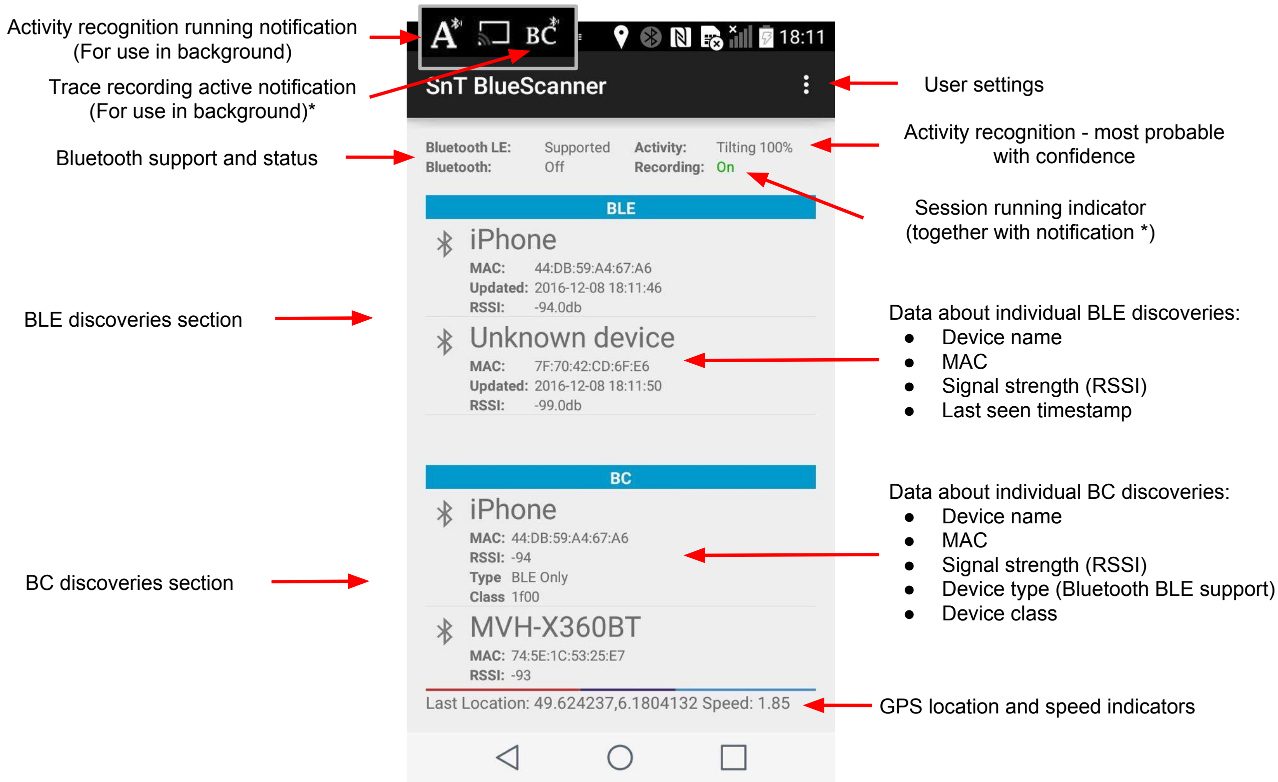

Figure 4.7: Main screen details

and stop the scanning procedure. This policy is portrayed in Figure 4.6.

Since we wanted to keep user interaction to a minimum, so as to avoid unnecessary distraction at the wheel, our application uses the Android activity recognition Application Programming Interface (API) [3] to automatically start discovering for devices once the in-vehicle activity is detected with a certain precision.

The device, upon opening the app, starts observing all activities from the API (new updates every 10 s). Once a confidence level of over 90% on the *In Vehicle* activity is reached, the application starts looking and recording Bluetooth traces. At this moment we also start keeping track of the user activities within set intervals of three minutes. If we notice that the *In Vehicle* readings coming from the activity recognition API represent below 15% of the overall activities during this set interval, we can safely assume the user is no longer driving. These values were chosen carefully after many tests in traffic scenarios, where the *In Vehicle* activity confidence drops drastically. With such configuration the recording would withstand different congestion levels (slow speed & start/stop traffic) in order to continue recording traces.

Once a driving activity is detected and a recording session is started the application will display to the user information about nearby discovered devices. Figure 4.7 shows in details what is displayed

[3] https://developers.google.com/android/reference/com/google/android/gms/location/DetectedActivity

Type	**Example**	**Details**
MAC	AA:BB:CC:DD:EE:FF	Used for cross identification (unique only for BC)
RSSI	-50 dBm	Received Signal Strength
Device Name	Audi_MMI_XXX	Assigned device name if defined
Class Identifier (BC only)	200408	Type of device (e.g. hands-free kit)
Manufacturer Information (BLE only)	4C 00	Manufacturer (and sometimes model) of the discovered device (e.g. 4C corresponds to the Apple).

Table 4.1: Discovery Data

to the user on the screen.

Because of how the discovery process works between the two technologies; BC devices are discovered during a 12 sec discovery cycle (Android platform specific[4]) which is then looped for the length of the session. For BLE instead we keep the device in discovery mode which in return will log every received advertisement at any time. Furthermore, throughout the whole session, we log GPS information such as coordinates, altitude and bearing.

Furthermore, Table 4.1 shows the recorded information for every discovery. The class identifier provides data about the type of BC device (e.g. class 200408 represents a hands-free device, which in our case is a vehicle). On the other hand, BLE manufacturer information provides details about the brand and model of the discovered device (e.g. 4C correspond to the Apple company identifier[5]).

Paired with the front-end application we also built a back-end solution for storing and browsing traces from all study participants. The upload of traces from the application to our server was also automated, along with the handling of the sessions. The application, unless specified otherwise by the user in the settings (Figure 4.5b), tries to constantly notify our back-end about running sessions and discovered devices. This allows us to have a live feedback of running sessions as well as what devices have been discovered and where. Every trace is also logged locally and uploaded later in case there is an issue with the user's connection. Moreover, to secure the connection between front-and back-end, the upload of traces was encrypted with an SSL certificate before being transmitted.

Once more, the thinking behind these automation design choices was to reduce user interaction to a minimum, to allow customization of settings for different uses and to standardize a way of starting and stopping scanning sessions in the same way for everyone.

Traces are stored in our database for fast and reactive tracking during the collection campaign. Throughout two months we pushed 14 updates to our application tweaking various aspects of the

[4] `https://dev.android.com/reference/android/bluetooth/BluetoothAdapter.html`
[5] `https://www.bluetooth.com/specifications/assigned-numbers/company-identifiers`

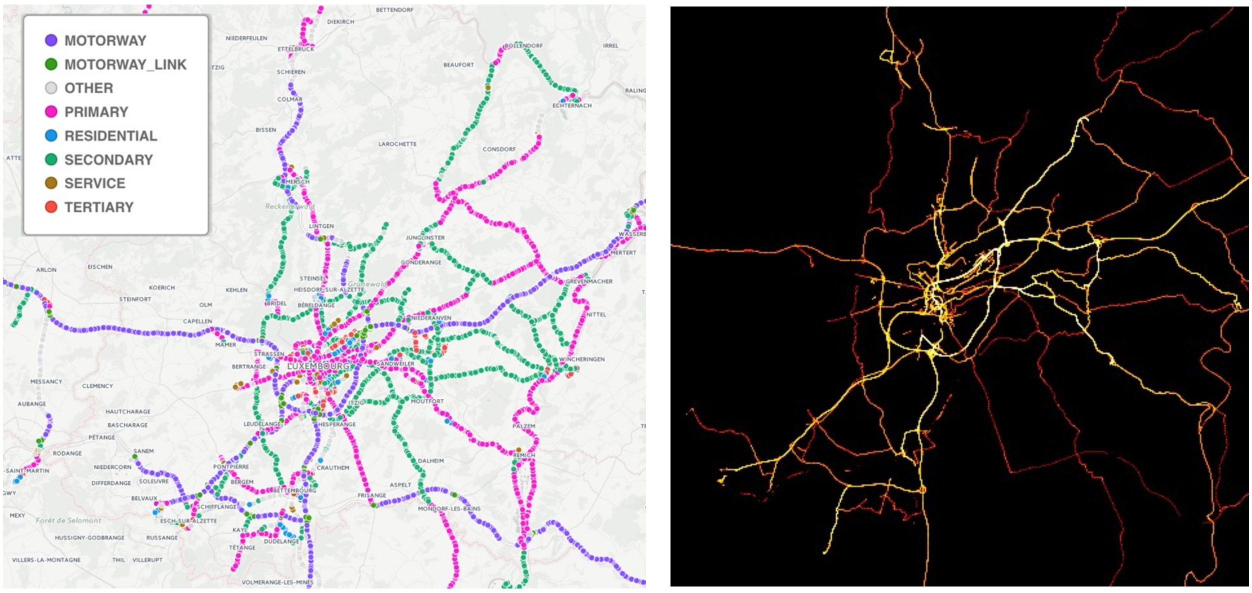

(a) Locations mapping (b) Locations density

Figure 4.8: Dataset locations

collection of traces and adapting to user's feedback.

The back-end was developed with a fully-fledged interface that allowed us to quickly visualize data concerning single traces as well as overall information about the contributions of each participant.

4.4.2 Data Collection Campaign

Because of the nature of this study, sensitive information about the participants had to be recorder. Information such as locations and details about the users' device was securely stored on our back-end

Number of participants	20
Number of recorded sessions	794
Location entries	284,975
Total coverage	13,486 km
BC entries	64,099 (16,201 Unique)
BLE entries	184,497 (6,241 Unique)
Average session duration	24.93 min

Table 4.2: Dataset Summary

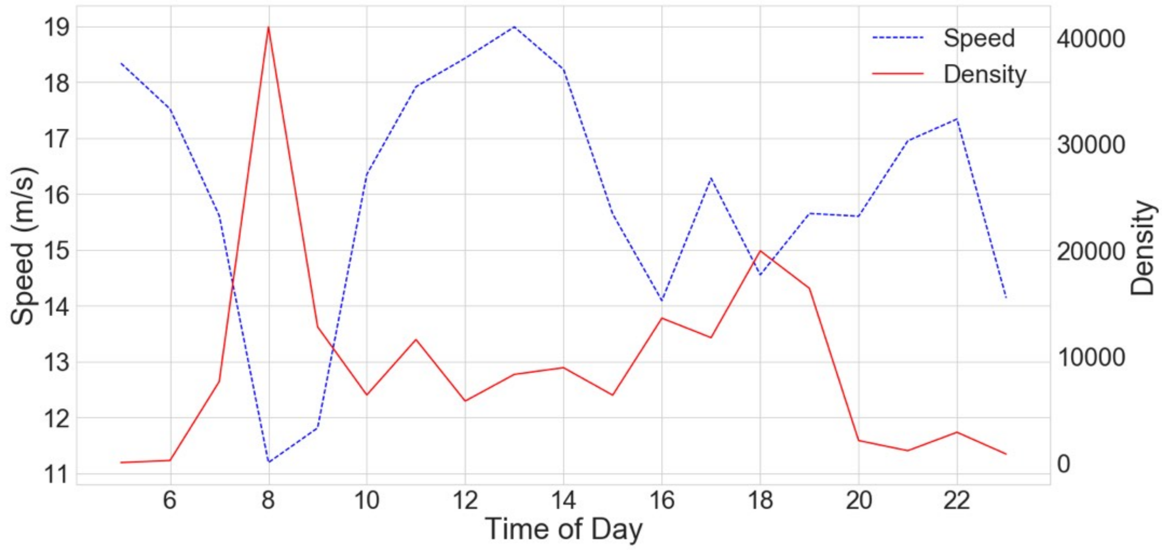

Figure 4.9: Hourly breakdown of the average speed with the locations density

to be later utilized. All of our users were prompted to read a terms and condition form and accept the privacy policy of the study.

Moreover, we only instructed the participants to keep the phone on a holder on the windshield of their vehicle. Typical driving routes include to and from work plus extra activities that require using the car.

Although the application gave the option for different mode of transportation, such as bus or mixed (Figure 4.5c), we found that all the users took the car making this option superfluous. If the study were to be extended and relaunched with more participants this option could definitely shed some light on differences between discoveries whilst driving the car compared to taking public transport.

4.4.3 Dataset Statistics

Table 4.2 provides an overview of the final statistics after our two-month campaign. The total number of entries reflects the differences between discovery behaviours of the two Bluetooth technologies. For more information on how the two protocol stacks handle discovery, interested readers can refer to Chapter 2 and the work of Gomez at al. [28].

For an initial validation of the dataset, we plotted the recorded latitudes and longitudes to Open Street Map (OSM) road classes directly using the *osm2po* tool [115, 116].

This conversion takes into account that sometimes latitudes and longitudes are not precise or consistent enough with the actual road structure. For example, if we are on a motorway and there is a perpendicular road that bridges over the motorway, a simple latitude and longitude to road class query at that particular spot might tell us that we are on the perpendicular road instead of the motorway. For this reason we take into consideration here GPS bearing as well as knowledge about immediately adjacent latitudes and longitudes. Doing so we can overrule with high precision

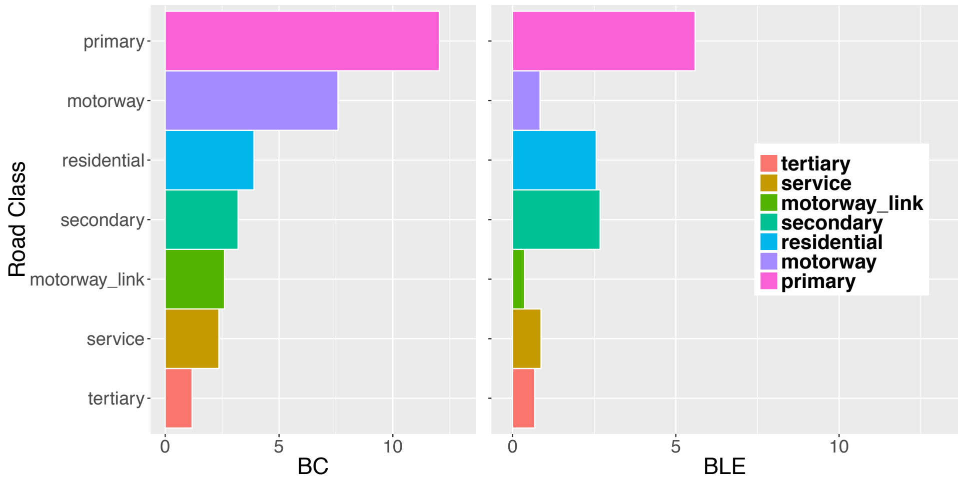

Figure 4.10: Average number of unique devices discovered per session (X-axis represents count)

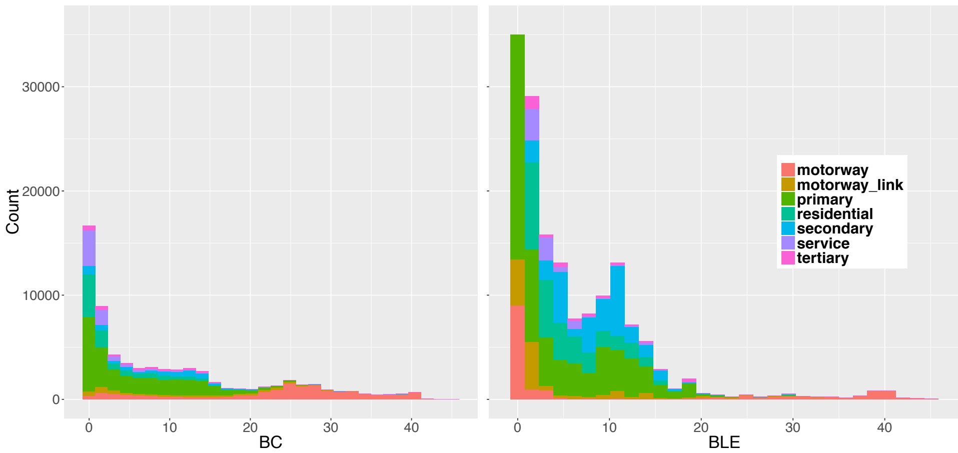

Figure 4.11: Speed (m/s) distribution per road class (X-axis represents speed in m/s)

any conflicting coordinates and correctly classify them with the appropriate OSM road class using *osm2po*.

As we can see from Figure 4.8, the density and spatial distribution of trips provides a large coverage of diverse road types within the country of Luxembourg. An hourly breakdown of location entries and average speed can be seen in Figure 4.9. As expected, these numbers reflect the morning and evening rush hour trend on the Luxembourgish roads, during which the average speed goes down and the count of encountered devices goes up [117].

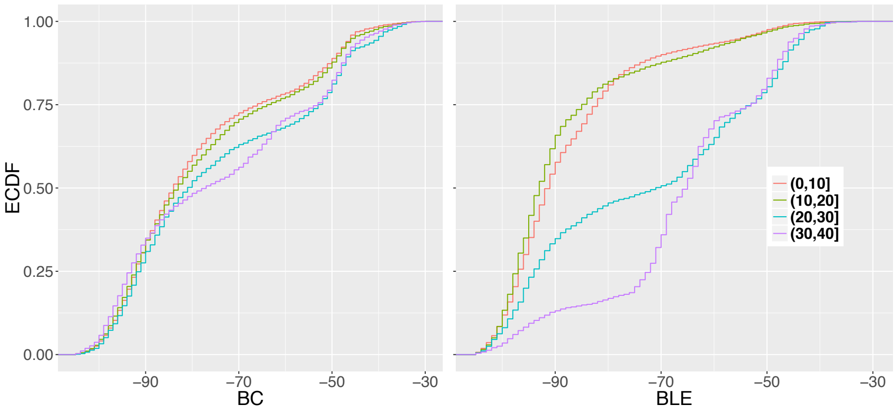

Figure 4.12: RSSI (dB) distribution comparison between speed-bins (m/s)

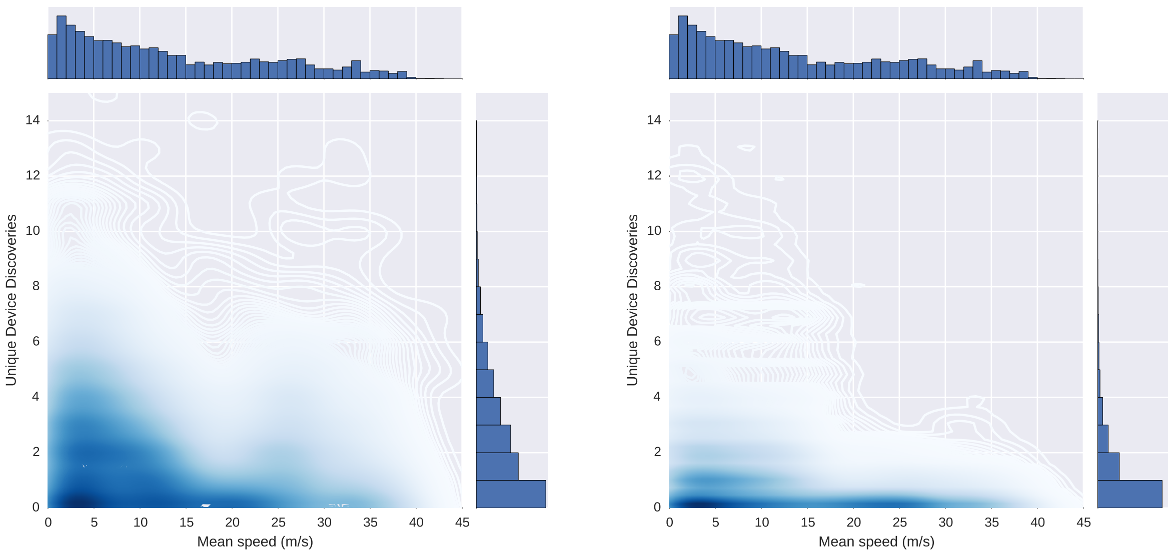

(a) BC

(b) BLE

Figure 4.13: Unique device discovery count vs. mean travel speed within 60 second time windows

4.5 Statistical Analysis

It is important to characterize the different behaviours of both BC and BLE in every respective road class. In Figure 4.10 we notice how, looking at the average number of unique devices discovered per session, BC (left) can be found mostly in primary and motorway roads whilst BLE in primary, residential and secondary roads. This already gives us an insight of where and for what purpose these protocols are used. On one side we have BC radios that represent mostly cars (hands-free and music streaming) and can thus be found with more ease on highways and primary roads. On the other side

BLE (right) which is mostly implemented in objects that pedestrians are expected to carry and that therefore are more likely to be encountered on roads with slower speed limits such as residential, and secondary roads.

The speed distribution per road class (Figure 4.11) further strengthen this assumption. BC (left) shows how motorway is not only better represented compared to BLE (right) but also relevant from speeds over $20\,m/s$ (usually primary, highway roads speeds). The high number of motorway entries observed at low speed in BLE represents dense traffic situations where more vehicles are found within communication range, i.e. within $100\,m$ for BLE compared to $10\,m$ for BC.

Moreover, we see how primary, secondary and residential roads represent the bulk of the discoveries for BLE proving how this technology is better applied to V2X scenarios where a car needs to gather knowledge of objects and pedestrians around it (mostly in cities).

The definition of roads as primary roads is sometimes ambiguous. At times this road type, at least within Luxembourg, is represented by both extra-urban roads that connect villages within the country and roads usually very crowded with pedestrians inside the city. We took this into considered when building the scenarios for our classification study.

We can conclude that primary, secondary and residential roads represent the bulk of the discoveries for BLE further proving how this technology is better applied to medium to low speed V2X scenarios where a vehicle needs to gather knowledge of objects and pedestrians around it.

In Figure 4.12, we look at the distribution of RSSI values with respect to speeds. We split the data into bins of $10\,m/s$, and show the ECDFs of BC and BLE. While the BC distribution seems homogeneous with respect to speed, the BLE distributions appear to be speed-dependent. Note that the higher speed-bins in BLE correspond to increased RSSI values. More specifically, the median value for the high speed bins is located at about $-68\,dB$, while the lower speed bins have a median value below $-90\,dB$. We believe this behaviour indicates that at higher speeds discovered devices are more sparse but RSSI is higher on average, likely due to more close-by encounters and BLE being more sensitive to speed (only high RSSI signals are received).

We also evaluated the average number of devices discovered in the course of 60 second time windows in both technologies. Fig. 4.13 displays average speed (over the measured 60 seconds) on the x-axis, and the number of discovered devices on the y-axis. The kernel density estimate shows which combinations of speed and number of device discoveries are more frequent. We can see that BC yields more discoveries along the entire velocity spectrum, while BLE has a clearly distinguishable *fold* in the discovery probability at approximately $20\,m/s$. This matches the BLE ECDFs in Fig. 4.12, where the higher speed bins ($> 20\,m/s$) showed higher median RSSI. We believe this indicates that in this speed range, discovered devices are more sparse but RSSI is higher on average, likely due to more close-by encounters.

Lastly, we analysed the types of devices encountered within a specific road class. Figure 4.14 & 4.15 show these statistics for BC and BLE. In the case of BC, although a good number of devices cannot have their class decoded (Unknown device), we notice that hands-free devices are the second most found type. This class is found uniquely inside cars which as a matter of fact are mainly found in

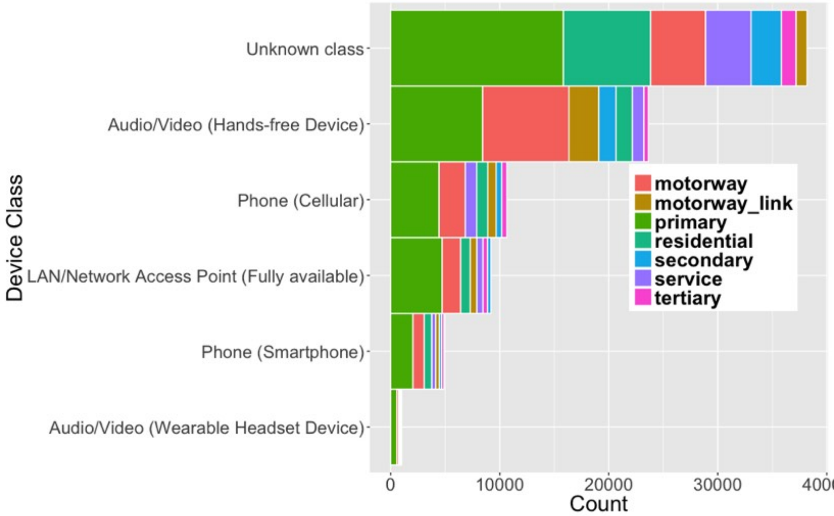

Figure 4.14: Encountered BC device classes per road class

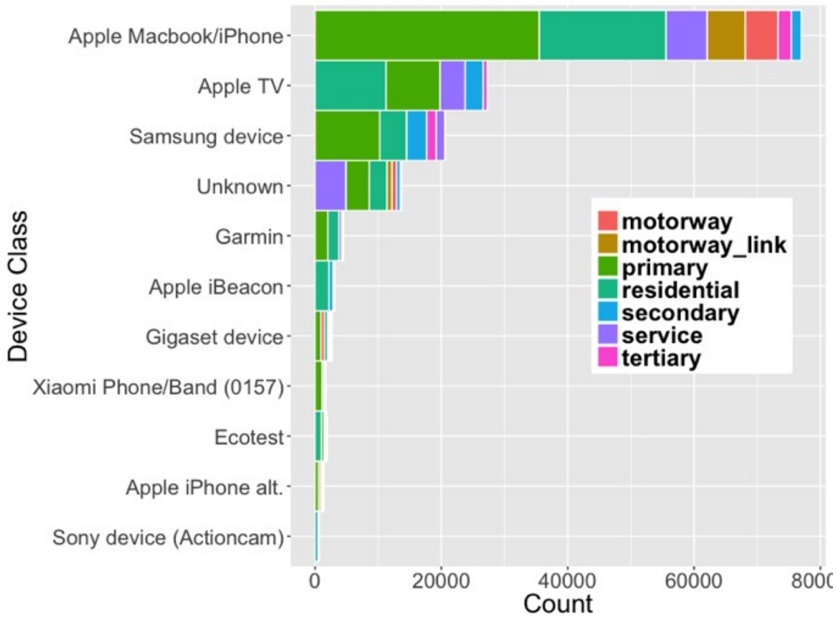

Figure 4.15: Encountered BLE device manufacturers & models per road class

highway and primary roads as the graph correctly portrays. A same analogy can be done for BLE's discoveries where the majority of beacons are found in primary and residential roads, mostly in devices that we carry around or we keep at home; smartphones, smartbands, televisions, etc.

In conclusion, from this analysis, we can derive that the dataset is consistent and large enough to be exploitable. The presented metrics such as the class of devices and the density of discoveries are complete enough to open the way for classification techniques towards estimating the environment (road class) through the BC and BLE discovery behaviour.

4.6 Classification

Our statistical analysis of the collected data points to different behaviours from BC and BLE dependent on the context in which they are discovered. We can distinguish from our data that BC is identified mostly in moving vehicles whilst BLE in pedestrians and fixed road side objects. With the help of additional information such as the device class (BC only) and manufacturer information (BLE only), we can further expand our knowledge of the nature of devices discovered nearby (e.g. smartphones, hands-free from vehicles, smartwatches, TV, bicycles).

Using this information we dwell in the field of ML to build a classification model able to take as input Bluetooth data alone and classify a generic environment (Highway, City e.g.).

4.6.1 Machine Learning for Contextualizing Discoveries

In the process of identifying environmental characteristics in a vehicular context we will make use of ML. The interest in ML in general is driven by the incredible amount of data collected around the world for nearly every application we can think of. The cost of storage is at an all time low enabling companies to store information on customers and products without worrying too much about costs [118].

The concept of ML is to make computers program themselves [119]. In traditional programming we develop a program and run it with some data to produce a certain output. In ML we run directly the data with the expected output and we obtain as a result the program. This program can then be used as input of the traditional programming effectively cutting the developer effort in the process.

The standard procedure to execute a ML task can be divided in 5 stages:

- **Data Collection**. The first step is to collect data to be used in the training process. The more data we have, the better our features will be. Having consistent and relevant features will further improve our final evaluation.
- **Data Analysis**. This step is mostly analytical. It is meant to determine the quantity and quality of data to make sure there are no inconsistencies or outliers. The goal is to determine that the dataset is coherent before the training process starts.
- **Model Training**. During training we can pick a specific algorithm and format the data to create a model. The data is divided into training, used for developing the model, and testing data, used as a reference to score the model.
- **Model Evaluation**. The test data is used here to determine the precision of the model prediction. The score will depend on the algorithm used and on the quality of the training data. Often we use *accuracy* as a metric for the evaluation but, as we will see during our research, this is often not precise enough especially when dealing with multiple classes. Other metrics include *F1 score* and Area Under the Curve (AUC) which we will both use in our work.
- **Model Performance Enhancement**. Basing ourselves form the scoring metrics from the previous step we can tweak our feature set, modify specific algorithm variables or change the algorithm all together.

4.6.2 Machine Leaning Types & Algorithms used

There are three main types of ML:

- **Supervised Learning**. Simply put, in this kind of ML, the desired output is included in the training data. Clear instructions are provided to the *predictive model* on what is expected in every case. Some examples of algorithms used are: Decision Trees, Support Vector Machine (SVM), Regression, Nave Bayes etc.
- **Unsupervised Learning**. As the name suggests, in this type of learning we do not include the desired output in the training data. The model is therefore *descriptive* and no single feature is more important the other. Examples of algorithms: K-mean clustering, hierarchical clustering etc.
- **Reinforcement Learning (RL)**. RL is a continuous learning process that aims at making specific decisions to solve business problems. It continuously retrains itself by interacting with its environment. The main difference with supervised learning is that RL algorithms learn from past experience and the continuous trial and error process rather than having desired outputs as examples.

In our research we used supervised learning algorithms and in particular Random Forest (RF) and SVM. RF is an algorithm [120] which runs by creating decision trees at training time. The more trees we create, the more confidence the accuracy results inspire. By using multiple decision trees we avoid an overfitting problem that we would otherwise have with a single big decision tree. Random forest achieves this by creating subtrees from random features subsets to be later recombined together.

There are no hyper parameters to tune in RF except the number of trees which its usually the more the better. There are many other advantages to using RF. Amongst them, the algorithm can handle multiple classes out of the box, as well as being efficient even with random missing data. These characteristics will be needed later on for our research.

Even though in the end we went with RF because of its simplicity and performance, we also tried SVM. SVM is a supervised learning algorithm and like RF [121], it attempts to distinguish two classes by using a hyperplane with the largest margin. In SVM, compared to RF, we have many parameters to tune; Kernel, slack variable, regularization penalties etc. Similarly to RF we also have multi-class classification in SVMs by using a *one-vs-all* approach.

4.6.3 Classifier Description

The raw discovery entries taken as inputs from our dataset (Figure 4.16) are preprocessed for our classification in order to provide features for different training timeslots. More precisely, the features are (see Table 4.1 for feature details):

- Total number of devices
- Unique number of devices
- Device Class or Manufacturer Information
- Received Signal Strength Indication (RSSI)

For each feature, we consider both the BC and BLE metrics. All different configurations of technologies

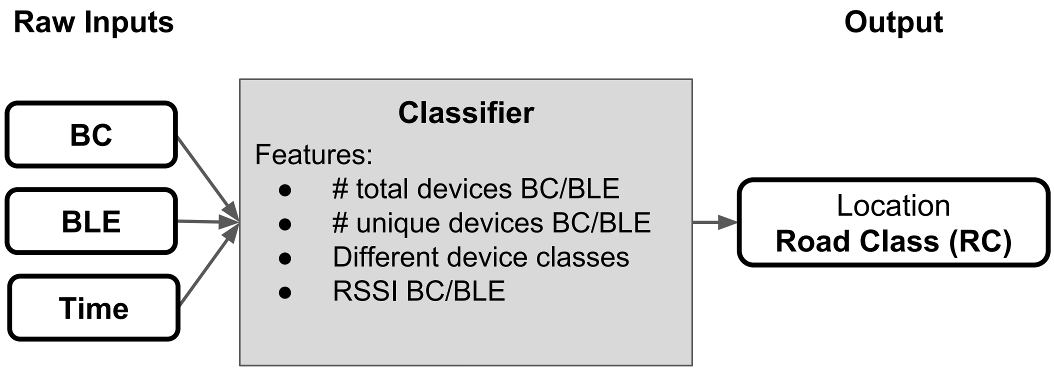

Figure 4.16: Classifier description

and feature sets were evaluated but ultimately the combination of both protocols proved to perform better. The RSSI, although proven somewhat inconsistent, has statistical significance when comparing both BC and BLE values as we saw in Figure 4.12.

To correctly homogenize the data, we ordered all entries temporally and only considered rush hours (from 8:00 to 10:59 and from 16:00 to 19:59). Moreover, we divided our training and test sets respectively 80% and 20%. This selection correctly matches the typical traffic demand during a work day in Luxembourg, as previously shown in Figure 4.9 and as observed by Codeca et al. [117].

4.6.4 Scoring metrics

As is the case in many classification approaches, we consider the accuracy as a starting point for scoring our classifier. This metric is produced by dividing the number of correct predictions made by the total number of predictions. Given the true positives (TP), true negatives (TN), false positives (FP) and false negatives (FN):

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

We quickly notice though, as seen in our statistical analysis, that we are dealing with an imbalance of classes in our dataset (On average highway and primary roads cover 60% of the total amount locations, cf. Section 4.5). Secondary and tertiary roads are greatly under-represented and though we might get a very high accuracy this is most likely due to a classifier always classifying the most represented classes. For this reason other performance measures must be considered .

The F_1 score is overall considered more useful than accuracy especially with an uneven class distribution. This metric takes better into account FP and FN by providing an harmonic mean of Precision and Recall.

(a) Scenario A - 2 Classes

(b) Scenario B - 3 Classes

(c) Scenario C - 4 Classes

Figure 4.17: Overview of the three scenarios

$$Precision = \frac{TP}{TP + FP}, Recall = \frac{TP}{TP + FN}$$

$$F_1 = 2 \times \frac{precision \times recall}{precision + recall}$$

The recall can be used to measure the integrity of a classifier. This score is obtained by dividing the number of true positives with the sum of true positives and false negatives. A low recall implies a high number of false negatives. The specificity or true negative rate, is represented by dividing true negatives with the sum of true negatives and false positives. Again, a low specificity implies a high number of false positives.

$$Specificity = \frac{TN}{TN + FP}$$

By combining these two metrics we obtain a 3rd metric; The AUC which is a performance measure between 0 and 1:

$$AUC = \frac{Recall + Specificity}{2}$$

We use in this case the AUC score for a single point on the Receiver Operating Characteristic (ROC), the curve plot of the recall against the specificity. A classifier with an AUC score > 0.5 performs better than random guessing [122]. By adopting this metric, we can give unbiased scores to our classifier by correctly representing minority classes.

4.6.5 Classifying Environments

We wanted the output of our classifier to provide a low precision context at first. Doing so provided us with a proof-of-concept of what we could achieve with our particular dataset and feature selection. For this purpose we propose three different scenarios built upon the data we collected.

These scenarios, as described in Figures 4.17 have different labels and can describe driving environments with different precisions.

As Table 4.3 shows, our simplest configuration (Scenario A) only has two labels indicating that only highways and urban environments are taken into consideration. For every scenario, we utilize a *osm2po* class-to-label mapping to determine the road type label (Figure 4.18). In scenario A, we can find the following OSM road classes: *highway* and *highway_link* represented by the label Highway whilst *residential* goes under Urban. Additionally, for the Urban label, we also consider every location within a certain radius (3.5 km for Luxembourg city) from the geographical centre of the city. This class is fine tuned for our scenario but can easily be adopted for other cities with similar road topology.

For scenario B, we have a similar configuration with an added class representing everything that is not highway or residential roads (mainly primary, secondary and tertiary roads). Finally, for our last configuration (Scenario C), we further divide the Extra-Urban group into two sub-groups: Extra-Urban (EU)-Village and EU-Road. A location is labelled as EU-Village when at least two *residential* roads are found within 350 m. After careful testing, we found this tuning to be the best at representing the road network structure.

For our classification we used two of the most popular and well-known set of models and multi-class learning algorithms; SVM and RF [123]. In our case case RF provides consistently better accuracies

Table 4.3: Labels

	Scenario A	**Scenario B**	**Scenario C**
Highway	✓	✓	✓
Urban	✓	✓	✓
Extra-Urban		✓	
EU-Village			✓
EU-Road			✓

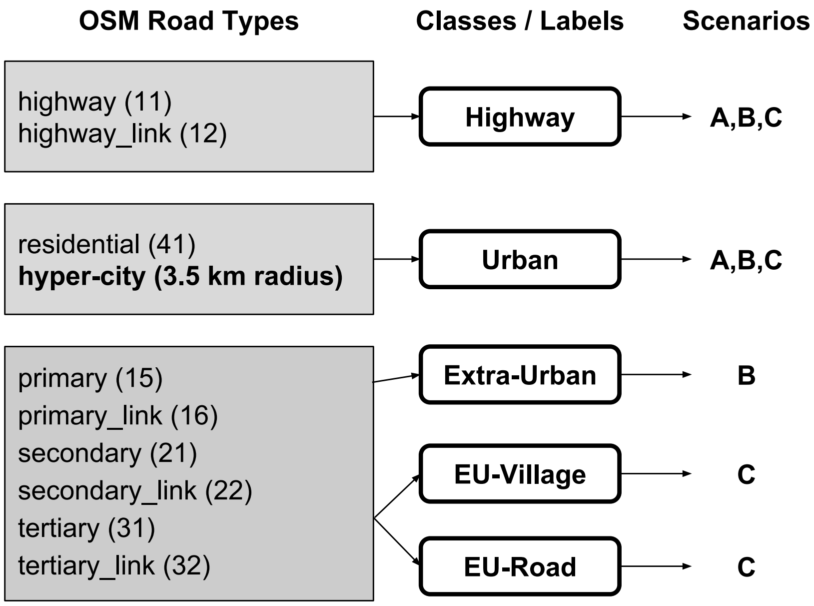

Figure 4.18: Open Street Map to Scenario Mapping

(1 to 5% higher) compared to SVM whilst also being better at handling an imbalance between classes. RF is a method used for classification and regression tasks that makes use of decision trees at training time [124]. We chose this particular algorithm because, as well as being fast and scalable, it can also handle multiple classes (up to four in our case) and behaves well with non-linear features. Using this algorithm allowed us to avoid hyperparameter tuning and instead to focus on resampling our dataset with different timeframes and sliding windows to observe the behaviour of the classifier. Moreover, for all our tests, we used 100 trees for each forest. A higher number of trees does not improve the accuracy.

Having exported and ordered our dataset temporally as described in Section 4.6.3, we then resampled it for every scenario using five different window lenghts (60s, 150s, 300s, 450s, 600s) and 5 different window overlaps (12.5%, 25%, 50%, 75%, 100%), giving us a total of 75 possible configurations (3 scenarios x 5 timeframes x 5 sliding windows). In Figure 4.19 we visualize how the sliding window length and overlap are used to sample our data. As the average of a session is $\approx$ 25 minutes, it did not make sense to try and sample windows bigger than 10 minutes.

In Figure 4.20 we show the entirety of the results for all timeframes and sliding windows configurations. For every single combination, we provide the results from an average of 50 runs of the RF algorithm.

For the final evaluation we decided to consider as scoring metrics accuracy and AUC score. The F_1 score results are directly correlated to the AUC score results and are therefore redundant for this purpose. We notice, as expected, that our Scenario A, with two road classes, is our best performing

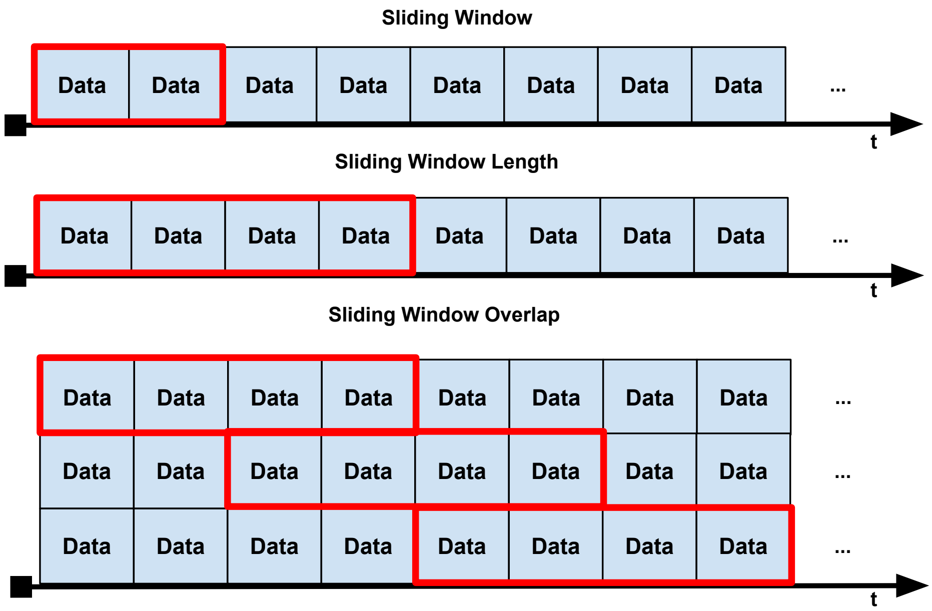

Figure 4.19: Sliding Window Length and Overlap

scenario. With a maximum accuracy of $\approx 88\%$ and ≈ 0.85 AUC score, we can say that our feature set can be successfully used to predict the road category. The more labels, hence the more road classes, we try to predict the more we see the accuracy of classifier decline. With three labels (Scenario B) we reach $\approx 74\%$ and ≈ 0.78 AUC score.

With the last scenario, when trying to predict four road classes, we observe that we cannot fully complete our result matrix. This means that in the empty cases one or more of the classes was rare enough that, regardless of the overall high accuracy, it could not be classified. The higher the time window is, the lower is the chance of a rare class being correctly represented. Although we could use, in specific cases, our classifier to predict up to four road categories, we found that with our dataset, the configuration with three classes achieves the best trade-off between number of potential predicted classes and overall accuracy of prediction.

4.7 Conclusion and Perspectives

We live in a world where everything is increasingly interconnected, at a time when both industry and researchers are trying to offer personalized services to users. Although the availability of opportunistic systems is extremely promising, the available data sources are currently not used to their full potential. The downside of standard sensors and communication interfaces that collect user data, such as GPS and Wi-Fi, is that they often have a high energy consumption and are more susceptible to compromising user privacy. In this chapter, we have taken on the challenge of exploiting the natural growth in the connectivity of our everyday devices using Bluetooth traces, which have a customizable

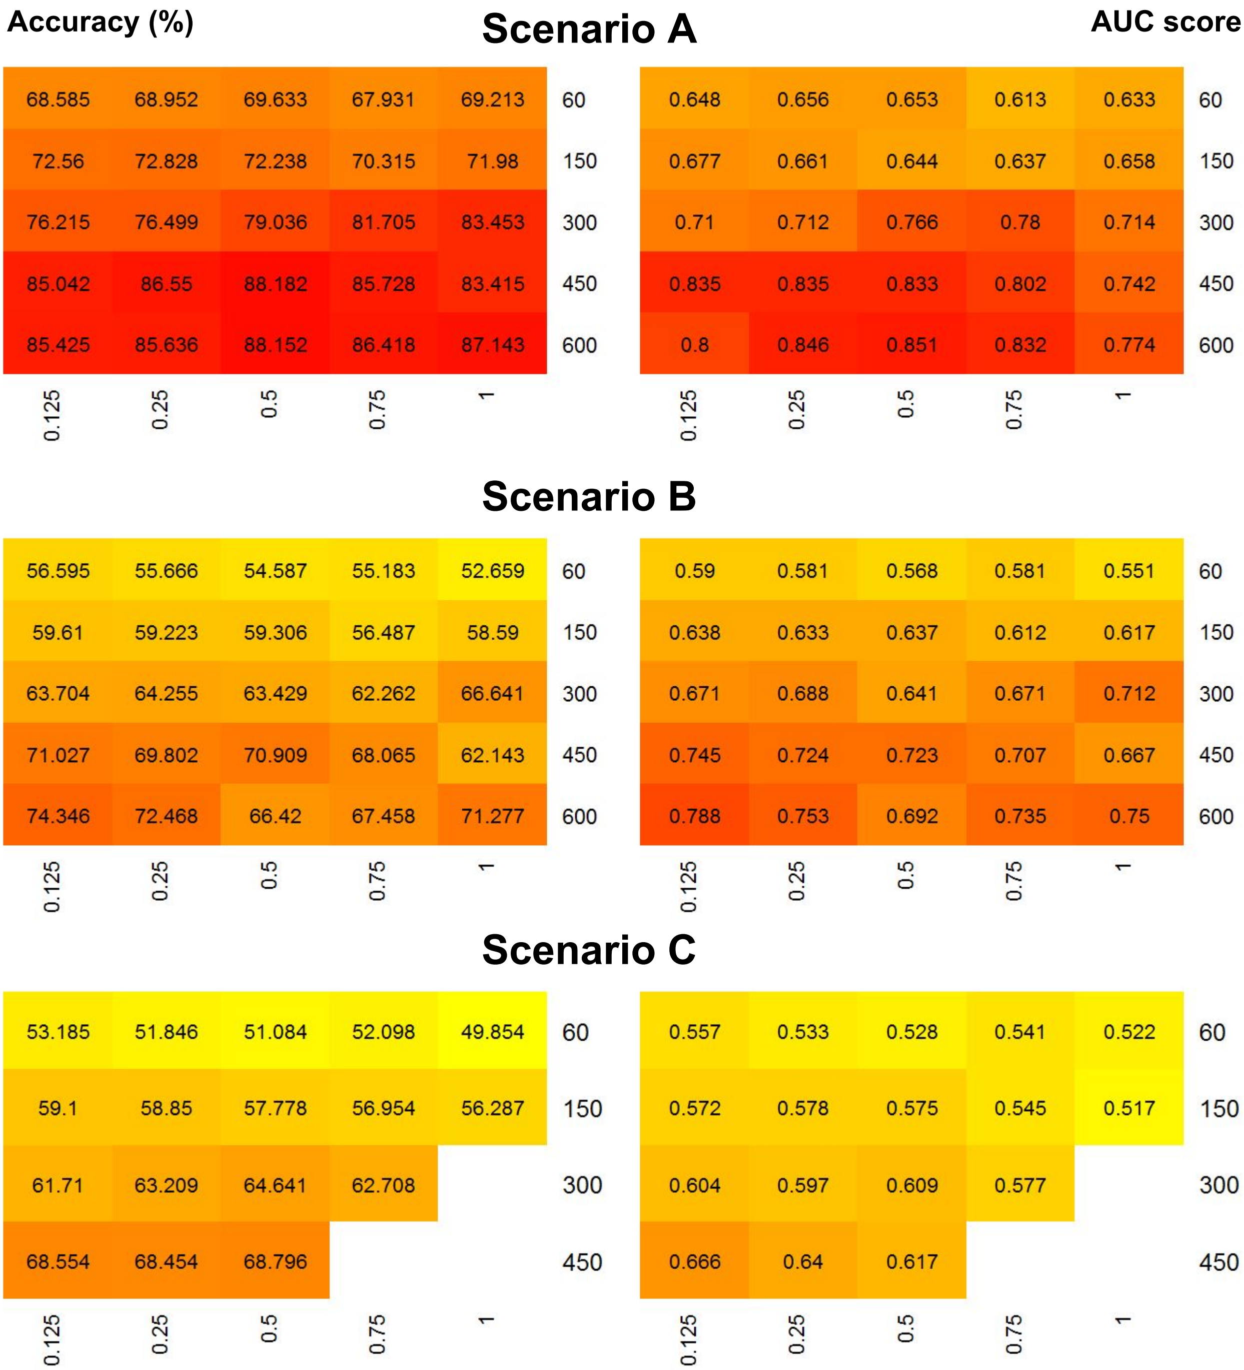

Figure 4.20: Classification results
(X-axes represent sliding windows and Y-axes represent timeframes)

sampling frequency and are one of the least expensive data sources in terms of energy consumption. Using a sensing system made freely available to the community online, we characterized Bluetooth discovery in a proof-of-concept vehicular scenario, with the aim of identifying different environmental characteristics. Such scenario allowed us to quickly cover a geographical area in Luxembourg, to demonstrate that the quantity and the quality of traces obtained by Bluetooth are strongly dependent on the mobility of the user and on the places he or she frequents. In particular, we implemented a mechanism that identifies different types of environments, such as highway, urban and extra-urban roads, by applying different sampling methods capable of obtaining promising accuracies.

This result is a definite step forward, considering that we are dealing with pure Bluetooth discovery data that is linked to human mobility and thus stochastic by nature.

This mechanism can open the way to a wide variety of applications and can easily be extended beyond vehicular systems (e.g. to pedestrians). More and more applications are using passive, opportunistic sensing systems to automatically characterize different aspects of their users. Most of the time, these sensing systems aim to provide personalized services for the user, for instance through a recommendation system. In addition, they aggregate a wide range of metrics related to their users' habits and environment.

The potential of using the mechanism proposed in this Chapter is threefold: 1) it is energy-efficient (compared to technologies like GPS); 2) it requires a low sampling frequency and minimal disk space and, above all; 3) it is privacy-friendly. An interesting example of this is the estimation of passenger flow in public transportation systems. If a user is moving while surrounded by other users, on a bus for example, it is possible to transcribe this behaviour by observing Bluetooth devices that remain in the vicinity of the user for a certain period of time, and by comparing this value with the speed estimated by the other devices (i.e. those only encountered a few times). Another application example is driver behaviour profiling by observing travel patterns. By detecting changes of environment (e.g. highway to urban and back to highway again), it is possible to identify recurring patterns like the journey from home to work and vice versa. Combining such knowledge with traffic data could allow the system to provide recommendations on route alternatives when needed. For instance, proposing a new route, at the discretion of the user, when a certain driving pattern is recognized within a known congested area.

The application developed for this article can be implemented as part of existing services that detect user contextual information and can easily be extended to cover more than just driving.

Future perspectives for this methodology includes proposing new methods of identifying other user characteristics, including current traffic state and type of vehicle, as well as finding new classifiers for different mobility scenarios such as cycling and walking. Studying relationships that may exist between the discovered devices would be of particular interest. Assuming that each device is a node, and that every time two nodes are scanned together, a link is created between them, we can apply graph theory to explore new approaches. This theory is not new, and has proven its effectiveness in describing a vast amount of data, especially in transportation-oriented research (e.g. [125]).

A further continuation of this work could look at identifying other characteristics of the discovered devices. Assigning a *mobility index* to a known kind of device would allow us to utilize the information in a knowledge database when selecting features for our classifier. Such an index could be compiled, based on repeated encounters with a specific device, by looking at locations and establishing a centre of gravity for all discoveries. This would allow us to establish if a certain class of devices is more likely to be fixed (such as a television) or mobile (smartphone, smartband).

Finally, exploring the possibilities of Bluetooth 5.0, a new standard that will provide devices with a far wider communication range, has the potential to significantly increase the precision of the classification process.

4.8 Contribution and Final Remarks

The goal of this chapter has been to outline the application-centric nature of Bluetooth (in both its classic and low energy modes) and to propose novel applications related to human mobility that exploit such technologies.

We presented an original data collection approach to exploit the beaconing aspect of Bluetooth. By first using a battery-powered embedded system and later a smartphone application, we collected data from multiple scenarios with various mobilities. This data collection contributed to validating the relevance of this research and stimulated further interest in the potential of how this data could be exploited.

With the help of a broader data collection campaign, we gathered nearly 800 traces from 20 participants with the objective of using Bluetooth discoveries to classify environments without the help of GPS.

Our approach determines the users' environment, with an increasing number of classes (from two up to four road types) by using ML algorithms on a predefined feature set containing only Bluetooth information.

The results show that using RF we can classify, with accuracies above 85%, two classes and, with around 75% accuracy, three. We tested with a maximum of four classes and found this to be the limit with our dataset.

In the future, we expect more and more devices to implement Bluetooth version 4.0 and higher. This will therefore make it more possible to add more classes and obtain higher accuracy with a more precise classifier.

Finally, we proposed applications scenarios that could see our research applied to real life solutions immediately.

Chapter 5

Conclusions

In this thesis we have presented how to use Bluetooth for human mobility applications; specifically, its use as a communication medium between devices in various dynamic scenarios and as a tool for contextualizing the environment by exploiting Bluetooth's application-centric nature. The goal we achieved in this research has been to outline specific benefits of utilizing the Bluetooth technology in mobile environments.

In Chapter 2 we characterized aspects of the Bluetooth technology, with a particular focus on Bluetooth Low Energy (BLE). We also introduced current and future application fields together with their corresponding mobility aspects.

In Chapter 3 we presented our first contribution, which makes use of BLE's active communication capabilities to build a platform to enable devices, in our case embedded in vehicles, to communicate with each other. In Chapter 4 we introduced our second contribution, a novel classification model which relies on passively-recorded Bluetooth data about nearby devices to provide contextual feedback on location.

This chapter summarizes our contributions and identifies future improvements. This work provides a first milestone in the research on BLE for communications in mobile contexts and environment sensing. Recent interest and improvements reaffirm our strong faith in the potential of this new Bluetooth version.

5.1 Summary of Contributions

The research question posed at the beginning of this thesis was:

> *How may we exploit Bluetooth in a mobile context?*

We proposed two ways we can take advantage of the ubiquity of Bluetooth in mobile environments. Firstly, by taking advantage of the technology as a communication tool, and secondly as an information provider. The following sections detail our contributions.

5.1.1 BLE as a V2X Communication Medium

Our initial studies focused on using BLE as an *active* communication medium for vehicular networks. In this context we developed a new platform to leverage BLE's transmission scheme by implementing a viable and easily-deployable application for off-the-shelf smartphones. Each device has the capability of being part of a wireless network where information regarding safety and traffic management can be transmitted to all other participating devices within range. By switching between BLE predefined master and slave roles we can quickly, and, most importantly without user interaction, pair two devices and transmit data through *undirected advertisements*. In the process of validating our platform we experimented with multiple scenarios to thoroughly study the behaviour and performance of BLE. Specifically, we conducted multiple experiments, both with constant and with varying distances and speeds. We measured performance in terms of delivery ratio and round-trip time for single-hop scenarios, and end-to-end delay for a multi-hop proof-of-concept trial. The maximum achieved communication range between devices went beyond 100 m and a robust connection was achieved up to a distance of 50 m, even in varying traffic and driving conditions.

After having analysed measurements taken in a variety of settings, we identified city scenarios with low to medium speed as the optimal environment to allow BLE to excel. This outcome suggested the need for a further study of the behaviour of Bluetooth and its coexistence mechanisms when subject to heavy interference like that encountered in the 2.4 GHz radio frequency spectrum in cities.

This study has been published in three papers [7–9] and ultimately a Journal: *Bluetooth Low Energy Performance and Robustness Analysis for Inter-Vehicular Communications* [10].

After having found the best performing environments for Bluetooth, we decided to broaden our focus and evaluate the technology as an information medium in specific mobile contexts. For this purpose we focused our next contribution on evaluating the suitability of Bluetooth's advertisement capabilities for recognizing characteristics of the environment by utilizing undirected broadcasts.

5.1.2 Bluetooth Classification Model for Contextualizing Environment

Bluetooth is found in devices that are inherently mainly present in populated areas. Since our previous work brought us closer to such environments, we decided to explore Bluetooth in a different light by building on our work to investigate the potential of BLE's discovery process. The first step toward this goal was to launch a data collection campaign, for which we developed a smartphone application capable of collecting both Bluetooth Classic (BC) and BLE discoveries. The campaign lasted two months, during which we observed 800 driving traces from 20 participants, who covered 13,500 km of roads in Luxembourg and neighbouring countries. The average driving session, usually commuting to/from work, lasted for 25 minutes and the total number of Bluetooth discoveries amounted to 22,500. Next, we analysed the complete dataset to validate its integrity by checking statistics such as number of discoveries, different speed distributions, and the types of devices discovered when travelling different road categories. As expected, we found BC to be more prevalent on highways and primary roads, where we find the majority of vehicles, and BLE in primary, secondary and residential roads, where the population is mostly concentrated. Having coherent and thus exploitable data was a requirement for proceeding further.

Finally, through Machine Learning (ML), we train a novel classification model to take the Bluetooth discovery information as input and produce a determination of the type of environment as output. In particular, with our prediction model, we were able to predict with reasonable confidence up to three environment classes (highway, city, extra-urban) by using only discovery data and no geographical information.

This contribution has been published in two papers: *Towards characterizing Bluetooth discovery in a vehicular context* [11] and *Characterizing Driving Environments Through Bluetooth Discovery* [12]. In addition, future perspectives specifically applied to general user mobility were published in the following Journal article: *Characterizing user mobility using mobile sensing systems* [13]

5.2 Final Remarks

Bluetooth is a technology that has been around for over 20 years; when we think about hands-free kits and cordless headsets we immediately think of Bluetooth. However, very few people are aware of its recently introduced Low Energy (LE) mode, which enhances the technology to a whole new level, bringing many new potential applications for IoT scenarios.

In this thesis, we have proposed two different applications of Bluetooth that were not within its original scope. We realized, in the early stages of our research, the potential of using such technology in human mobility related applications and, for this reason, we decided to experiment with different approaches to address current needs in vehicular and sensor networks.

We have provided the research community with an initial experimental study on the capabilities of BLE in environments where the technology had not been previously evaluated. Bluetooth is constantly evolving and our research can and should be updated when new specifications are released. Adapting to new features also means improving the BLE communication software as new chips (hardware) becomes available to mass-market manufacturers. The software developed for this thesis, which has now been made open-source, must therefore be optimized before being incorporated into larger frameworks.

5.3 Future Perspectives

The topics covered in this work are still extremely novel and our contributions only represent the first step towards employing Bluetooth in the fields studied. Wireless technologies evolve quickly and Bluetooth is no exception.

As an immediate short-term perspective, we can enhance the complexity of our classification model leading to the introduction of a new model to provide congestion information. Initially, our model can be revised to marginally improve the overall accuracy of the existing environment classifier by further fine-tuning our features and classification algorithm. A second and more effective improvement, would be to introduce a parallel model that, instead of giving information on the environment, would provide an indication of the traffic situation, for example *free-flow* or *congested*. Such a model could be paired with that described in this thesis to provide more precise and complete feedback regarding the mobile context by combining both road and traffic information (highway congested, city free-flow, etc.).

There is great medium-to long-term potential in this field, as few car manufacturers and industry giants have started adopting BLE specifically in vehicular applications [126]. Our study can be extended to include new mobility profiles for different classes of discovered devices. Every discovered device can be categorized with a different mobility probability, helping a faster and more accurate context analysis.

Another long-term development, since this research was conducted with Bluetooth version 4.0, would be to repeat some of the experiments with Bluetooth version 5, which was only recently released. The new release brings quadruple range, double speed and increased broadcasting capacity, all features that could be very significant for our current and future work.

Both of our contributions are impacted in some way by this new standard. For our V2X approach, the most interesting aspects are the extended range and increased broadcast capacity. Having the capabilities of transferring more information than before and achieving a reliable link up to 400m will definitely allow BLE to be revisited for a wider range of applications on the road, including safety and platooning. Without the need for pairing, we can completely eliminate the association delays that we recorded in our experiments.

Likewise, our driving data collection experiment could benefit from the improved broadcasting mechanism. With Bluetooth 5 hardware, we would be able not only to capture devices further away, but also to record more discoveries overall. Having greater range and better discovery capabilities could therefore help us build a bigger dataset to improve the definition of our classifier. The availability of Bluetooth 5, added to the increasing adoption of Bluetooth beacons, will almost guarantee an exponential growth of research interest in this technology. There will soon be a huge quantity of data available from wireless sensor networks, especially Bluetooth, that has the potential of being a game-changer in new human mobility research fields.

Further down the line, we can predict BLE being embedded within a fusion of sensors for autonomous driving. Bluetooth could help to gather knowledge about the surroundings with a particular focus on other vehicles and Vulnerable Road Users (VRU) such as bikes and pedestrians.

List of Tables

List of Figures

Acronyms

A2DP Advanced Audio Distribution Profile.

ACL Asynchronous Connection-Less.

AES Advanced Encryption Standard.

AFH Adaptive Frequency Hopping.

AP Access Point.

API Application Programming Interface.

ATT Attribute Protocol.

AUC Area Under the Curve.

BC Bluetooth Classic.

BLE Bluetooth Low Energy.

BPSK Binary Phase-shift Keying.

CCM Cipher Block Chaining-Message Authentication Code.

CRC Cyclic Redundancy Check.

DSRC Dedicated Short Range Communication.

EBRS Event-Based Role Switching.

ECDH Elliptic Curve Diffie-Hellman.

GAP Generic Access Profile.

GATT Generic Attribute Protocol.

GFSK Gaussian Frequency-shift Keying.

GPS Global Positioning System.

HCI Host-Controller Interface.

HFP Hands-Free Profile.

HU Head-unit.

IaVC Intra-Vehicle Communication.

IFS Inter-Frame Space.

IoT Internet of Things.

IrVC Inter-Vehicle Communication.

ISM Industrial, Scientific and Medical.

ITS Intelligent Transportation System.

L2CAP Logical Link Control and Adaptation Protocol.

LE Low Energy.

LL Link Layer.

MAC Media Access Control.

MIC Message Integrity Check.

ML Machine Learning.

NESN Next Expected Sequence Number.

NTP Network Time Protocol.

OBU On-Board Unit.

ODB On-board Diagnostics.

OFDM Orthogonal Frequency-division Multiplexing.

OQPSK Offset Quadrature Phase-shift Keying.

OSM Open Street Map.

PAN Personal Area Networking Profile.

PH Personal Hotspot.

PHY Physical Layer.

RF Random Forest.

RFI Radio-Frequency Interference.

RL Reinforcement Learning.

RSSI Received Signal Strength Indication.

RSU Road Side Unit.

RTT Round-Trip Time.

SCO Synchronous Connection-Oriented.

SDK Software Development Kit.

SIG Special Interest Group.

SM Security Manager.

SN Sequence Number.

SVM Support Vector Machine.

V2I Vehicle-to-Infrastructure.

V2V Vehicle-to-Vehicle.

V2X Vehicle-to-Everything.

VANET Vehicular Ad Hoc Network.

VRU Vulnerable Road User.

WAVE Wireless Access in Vehicular Environments.

WLAN Wireless Local Area Network.

WMAN Wireless Metropolitan Area Network.

WPAN Wireless Personal Area Network.

WWAN Wireless Wide Area Network.

Publications

W. Bronzi, R. Frank, G. Castignani, and T. Engel, "**Bluetooth Low Energy performance and robustness analysis for Inter-Vehicular Communications**," *Ad Hoc Networks*, vol. 37, pp. 76–86, 2016.

——, "**Bluetooth Low Energy for Inter-Vehicular Communications**," in *Vehicular Networking Conference (VNC), 2014 IEEE*. IEEE, 2014, pp. 215–221.

W. Bronzi, F. Sébastien, R. Frank, and T. Engel, "**Characterizing Driving Environments Through Bluetooth Discovery**," in *The 8th International Conference on ICT Convergence (ICTC 2017)*, Oct 2017.

W. Bronzi, T. Derrmann, G. Castignani, and T. Engel, "**Towards characterizing Bluetooth discovery in a vehicular context**," in *Vehicular Networking Conference (VNC), 2016 IEEE*. IEEE, 2016, pp. 1–4.

W. Bronzi, "**Bluetooth Low Energy Robustness Analysis for V2V Communications**," *Ulmer Informatik-Berichte*, vol. 1, no. 6, p. 9, 2015.

R. Frank, W. Bronzi, G. Castignani, and T. Engel, "**Bluetooth Low Energy: An alternative technology for VANET applications**," in *Wireless On-demand Network Systems and Services (WONS), 2014 11th Annual Conference on*. IEEE, 2014, pp. 104–107.

S. Faye, W. Bronzi, I. Tahirou, and T. Engel, "**Characterizing User Mobility Using Mobile Sensing Systems**," *International Journal of Distributed Sensor Networks*, vol. 13, no. 8, p. 1550147717726310, 2017.

S. Jafarnejad, L. Codeca, W. Bronzi, R. Frank, and T. Engel, "**A Car Hacking Experiment: When Connectivity meets Vulnerability**," 2015.

Bibliography

[1] S. Kajioka, T. Mori, T. Uchiya, I. Takumi, and H. Matsuo, "Experiment of indoor position presumption based on RSSI of Bluetooth LE beacon," in *Consumer Electronics (GCCE), 2014 IEEE 3rd Global Conference on*. IEEE, 2014, pp. 337–339.

[2] Kontact.io, "The complete Beacon industry report: Current Technology and Recent History," https://kontakt.io/blog/beacons-2017-forecast-7-experts-share-their-thoughts/, Jan 2017.

[3] Bluetooth Special Interest Group, "Bluetooth 5 - Rethinking the future document." https://www.bluetooth.com/~/media/files/marketing/bluetooth5-rethinking-the-future.ashx?la=en, Jan 2017.

[4] O. Altintas, F. Dressler, H. Hartenstein, and O. K. Tonguz, "Inter-Vehicular Communication - Quo Vadis (Dagstuhl Seminar 13392)," *Dagstuhl Reports*, vol. 3, no. 9, pp. 190–213, 2013.

[5] Apple Inc., "Apple CarPlay," URL: https://www.apple.com/ios/carplay/.

[6] Google Inc., "Android Auto," https://www.android.com/auto/.

[7] R. Frank, W. Bronzi, G. Castignani, and T. Engel, "Bluetooth Low Energy: An alternative technology for VANET applications," in *Wireless On-demand Network Systems and Services (WONS), 2014 11th Annual Conference on*. IEEE, 2014, pp. 104–107.

[8] W. Bronzi, R. Frank, G. Castignani, and T. Engel, "Bluetooth Low Energy for Inter-Vehicular Communications," in *Vehicular Networking Conference (VNC), 2014 IEEE*. IEEE, 2014, pp. 215–221.

[9] W. Bronzi, "Bluetooth Low Energy Robustness Analysis for V2V Communications," *Ulmer Informatik-Berichte*, vol. 1, no. 6, p. 9, 2015.

[10] W. Bronzi, R. Frank, G. Castignani, and T. Engel, "Bluetooth Low Energy performance and robustness analysis for Inter-Vehicular Communications," *Ad Hoc Networks*, vol. 37, pp. 76–86, 2016.

[11] W. Bronzi, T. Derrmann, G. Castignani, and T. Engel, "Towards characterizing Bluetooth discovery in a vehicular context," in *Vehicular Networking Conference (VNC), 2016 IEEE*. IEEE, 2016, pp. 1–4.

[12] W. Bronzi, T. Derrmann, R. Frank, and T. Engel, "Characterizing driving environments through bluetooth discovery," in *The 8th International Conference on ICT Convergence (ICTC 2017)*, Oct 2017.

[13] S. Faye, W. Bronzi, I. Tahirou, and T. Engel, "Characterizing user mobility using mobile sensing systems," *International Journal of Distributed Sensor Networks*, vol. 13, no. 8, p. 1550147717726310, 2017.

[14] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless sensor networks: a survey," *Computer networks*, vol. 38, no. 4, pp. 393–422, 2002.

[15] V. Garg, *Wireless Communications & Networking*. Morgan Kaufmann, 2010.

[16] J. Yick, B. Mukherjee, and D. Ghosal, "Wireless sensor network survey," *Computer networks*, vol. 52, no. 12, pp. 2292–2330, 2008.

[17] C. F. García-Hernández, P. H. Ibarguengoytia-Gonzalez, J. García-Hernández, and J. A. Pérez-Díaz, "Wireless sensor networks and applications: a survey," *International Journal of Computer Science and Network Security*, vol. 7, no. 3, pp. 264–273, 2007.

[18] Lifewire, "WLAN Standards 802.11a, 802.11b/g/n and 802.11ac," https://www.lifewire.com/wireless-standards-802-11a-802-11b-g-n-and-802-11ac-816553, May 2017.

[19] E. Ferro and F. Potorti, "Bluetooth and Wi-Fi wireless protocols: a survey and a comparison," *IEEE Wireless Communications*, vol. 12, no. 1, pp. 12–26, 2005.

[20] P. Kinney *et al.*, "Zigbee technology: Wireless control that simply works," in *Communications design conference*, vol. 2, 2003, pp. 1–7.

[21] G. Fettweis and E. Zimmermann, "Ict energy consumption-trends and challenges," in *Proceedings of the 11th international symposium on wireless personal multimedia communications*, vol. 2, no. 4. (Lapland, 2008, p. 6.

[22] J. G. Andrews, A. Ghosh, and R. Muhamed, *Fundamentals of WiMAX: understanding broadband wireless networking*. Pearson Education, 2007.

[23] "Bleutooth History," https://www.bluetooth.com/about-us/our-history, Sept 2017.

[24] J. Haartsen, "Bluetooth - the universal radio interface for ad hoc, wireless connectivity," *Ericsson review*, vol. 3, no. 1, pp. 110–117, 1998.

[25] N. Golmie, O. Rebala, and N. Chevrollier, "Bluetooth adaptive frequency hopping and scheduling," in *Military Communications Conference, 2003. MILCOM'03. 2003 IEEE*, vol. 2. IEEE, 2003, pp. 1138–1142.

[26] Bluetooth Special Interest Group, "Specification of the Bluetooth System Core - Frequency Bands, Guard Bands and Channel Arrangement," http://grouper.ieee.org/groups/802/15/Bluetooth/core_10_b.pdf, Nov 1999.

[27] ——, "Bluetooth Specification Version 4.0," URL: https://www.bluetooth.org/docman/handlers/downloaddoc.ashx?doc_id=229737.

[28] C. Gomez, J. Oller, and J. Paradells, "Overview and evaluation of bluetooth low energy: An emerging low-power wireless technology," *Sensors*, vol. 12, no. 9, pp. 11 734–11 753, 2012. [Online]. Available: http://www.mdpi.com/1424-8220/12/9/11734

[29] J. Leonard - Get Connected Blog, Nordic Semiconductors, "Bluetooth 5 - What to expect," http://blog.nordicsemi.com/getconnected/bluetooth-5-what-to-expect, April 2017.

[30] Bluetooth Special Interest Groupf, "Bluetooth®5 - Quadruples Range, Doubles Speed, Increases Data Broadcasting Capacity," https://www.bluetooth.com/news/pressreleases/2016/06/16/-bluetooth5-quadruples-rangedoubles-speedincreases-data-broadcasting-capacity-by-800, June 2016.

[31] R. Gondane - EDGEFX Official Blog, "Introduction to Bluetooth Technology, its Working and its Applications," http://www.edgefxkits.com/blog/bluetooth-technology-and-its-working/.

[32] "Getting Started with Bluetooth Low Energy," https://www.safaribooksonline.com/library/view/getting-started-with/9781491900550/ch01.html, Sept 2017.

[33] S. Kamath and J. Lindh, "Measuring Bluetooth Low Energy power consumption," *Texas instruments application note AN092, Dallas*, 2010.

[34] A. Dementyev, S. Hodges, S. Taylor, and J. Smith, "Power consumption analysis of Bluetooth Low Energy, ZigBee and ANT sensor nodes in a cyclic sleep scenario," in *Wireless Symposium (IWS), 2013 IEEE International*. IEEE, 2013, pp. 1–4.

[35] M. Siekkinen, M. Hiienkari, J. Nurminen, and J. Nieminen, "How low energy is bluetooth low energy? comparative measurements with zigbee/802.15.4," in *Wireless Communications and Networking Conference Workshops (WCNCW), 2012 IEEE*, April 2012, pp. 232–237.

[36] J. Nieminen, T. Savolainen, M. Isomaki, B. Patil, Z. Shelby, and C. Gomez, "IPv6 over Bluetooth®Low Energy," Tech. Rep., 2015.

[37] M. Blom, M. Ekstrom, J. Castano, and M. Lindén, "Bluetooth energy characteristics in wireless sensor networks," in *Wireless Pervasive Computing, 2008. ISWPC 2008. 3rd International Symposium on*. IEEE, 2008, pp. 198–202.

[38] R. de Francisco, L. Huang, and G. Dolmans, "Coexistence of ZigBee wireless sensor networks and Bluetooth inside a vehicle," in *Personal, Indoor and Mobile Radio Communications, 2009 IEEE 20th International Symposium on*. IEEE, 2009, pp. 2700–2704.

[39] J.-S. Lee, Y.-W. Su, and C.-C. Shen, "A comparative study of wireless protocols: Bluetooth, UWB, ZigBee, and Wi-Fi," in *Industrial Electronics Society, 2007. IECON 2007. 33rd Annual Conference of the IEEE*. IEEE, 2007, pp. 46–51.

[40] S. Gollakota, F. Adib, D. Katabi, and S. Seshan, "Clearing the RF smog: making 802.11n robust to cross-technology interference," in *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4. ACM, 2011, pp. 170–181.

[41] R. Friedman, A. Kogan, and Y. Krivolapov, "On power and throughput tradeoffs of Wi-Fi and Vluetooth in smartphones," *IEEE Transactions on Mobile Computing*, vol. 12, no. 7, pp. 1363–1376, 2013.

[42] M. Ryan, "Bluetooth: With low energy comes low security," in *7th USENIX Workshop on Offensive Technologies*. Berkeley, CA: USENIX, 2013. [Online]. Available: https://www.usenix.org/conference/woot13/workshop-program/presentation/Ryan

[43] Great Scott Gadgets, "Ubertooth," https://github.com/greatscottgadgets/ubertooth/, April 2017.

[44] Bluetooth Special Interest Group, "Bluetooth Specification Version 4.2," URL: https://www.bluetooth.org/DocMan/handlers/DownloadDoc.ashx?doc_id=286439.

[45] ——, "Bluetooth Specification Version 5.0," URL: https://www.bluetooth.org/DocMan/handlers/DownloadDoc.ashx?doc_id=421043.

[46] Texas Instruments, "How Bluetooth 4.2 can help enable product security," https://e2e.ti.com/blogs_/b/connecting_wirelessly/archive/2016/08/11/how-bluetooth-4-2-can-help-enable-product-security, Aug 2016.

[47] Ravikiran HV, "Security considerations for Bluetooth smart devices," https://www.design-reuse.com/articles/39779/security-considerations-for-bluetooth-smart-devices.html.

[48] S. White - ECN, "Wi-Fi and Bluetooth coexistence," URL: http://www.ecnmag.com/articles/2012/03/wi-fi-and-bluetooth-coexistence, Feb 2012.

[49] Jeff Paynter, "Can Bluetooth and 802.11b/g/n Wi-Fi Devices Coexist?" URL: https://web.archive.org/web/20160323053451/http://intelligenthospitaltoday.com/can-bluetooth-and-802-11bgn-wi-fi-devices-coexist/, June 2013.

[50] A. Mourad, F. Heigl, and P. A. Hoeher, "Performance Evaluation of Concurrent IEEE 802.11 Systems in the Automotive Domain," in *IEEE LCN 2016*. IEEE, 2016.

[51] P. Smith - Convergent Promotions LLC, "Comparing Low-Power Wireless Technologies," URL: http://www.digikey.com/en/articles/techzone/2011/aug/comparing-low-power-wireless-technologies, 2011.

[52] R. Lall - Instrumentation Newsletter, National Instruments, "Ca se bouscule dans le spectre 2,4 GHz!" URL: http://www.ni.com/newsletter/51685/fr/.

[53] N. Golmie, O. Rebala, and N. Chevrollier, "Bluetooth adaptive frequency hopping and scheduling," in *Proceedings of the 2003 IEEE Conference on Military Communications - Volume II*, ser. MILCOM'03. Washington, DC, USA: IEEE Computer Society, 2003, pp. 1138–1142. [Online]. Available: http://dl.acm.org/citation.cfm?id=1950654.1950739

[54] A. M. Sayeed and B. Aazhang, "Joint multipath-doppler diversity in mobile wireless communications," *IEEE Transactions on Communications*, vol. 47, no. 1, pp. 123–132, 1999.

[55] J. Rodriguez-Pineiro, P. Suárez-Casal, M. Lerch, S. Caban, J. A. Garcia-Naya, L. Castedo, and M. Rupp, "Lte downlink performance in high speed trains," in *Vehicular Technology Conference (VTC Spring), 2015 IEEE 81st*. IEEE, 2015, pp. 1–5.

[56] A. Kandhalu, A. Xhafa, and S. Hosur, "Towards bounded-latency bluetooth low energy for in-vehicle network cable replacement," in *Connected Vehicles and Expo (ICCVE), 2013 International Conference on*, Dec 2013, pp. 635–640.

[57] J.-R. Lin, T. Talty, and O. K. Tonguz, "On the potential of Bluetooth Low Energy technology for vehicular applications," *IEEE Communications Magazine*, vol. 53, no. 1, pp. 267–275, 2015.

[58] J. B. Kenney, "Dedicated short-range communications (DSRC) standards in the United States," *Proceedings of the IEEE*, vol. 99, no. 7, pp. 1162–1182, 2011.

[59] European Commission, "Cooperative, connected and automated mobility (C-ITS)," https://ec.europa.eu/transport/themes/its/c-its_en, June 2017.

[60] D. Jiang and L. Delgrossi, "IEEE 802.11p: Towards an international standard for wireless access in vehicular environments," in *Vehicular Technology Conference, 2008. VTC Spring 2008. IEEE*. IEEE, 2008, pp. 2036–2040.

[61] Y. L. Morgan, "Managing DSRC and WAVE standards operations in a V2V scenario," *International Journal of Vehicular Technology*, vol. 2010, 2010.

[62] K.-Y. Ho, P.-C. Kang, C.-H. Hsu, and C.-H. Lin, "Implementation of WAVE/DSRC devices for vehicular communications," in *Computer Communication Control and Automation (3CA), 2010 International Symposium on*, vol. 2. IEEE, 2010, pp. 522–525.

[63] A. Festag, "Standards for vehicular communication from IEEE 802.11p to 5G," *e & i Elektrotechnik und Informationstechnik*, vol. 132, no. 7, pp. 409–416, 2015.

[64] ——, "Cooperative intelligent transport systems standards in Europe," *IEEE communications magazine*, vol. 52, no. 12, pp. 166–172, 2014.

[65] E. G. Strom, "On medium access and physical layer standards for cooperative intelligent transport systems in Europe," *Proceedings of the IEEE*, vol. 99, no. 7, pp. 1183–1188, 2011.

[66] ETSI, "ETSI EN 302 663 - Access layer specification for Intelligent Transport Systems operating in the 5 GHz frequency band," http://www.etsi.org/deliver/etsi_en/302600_302699/302663/01.02.00_20/en_302663v010200a.pdf, Nov 2012.

[67] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of human mobility on opportunistic forwarding algorithms," *IEEE Transactions on Mobile Computing*, vol. 6, no. 6, 2007.

[68] M. Zhao, L. Mason, and W. Wang, "Empirical study on human mobility for mobile wireless networks," in *Military Communications Conference, 2008. MILCOM 2008. IEEE.* IEEE, 2008, pp. 1–7.

[69] C. Song, T. Koren, P. Wang, and A.-L. Barabási, "Modelling the scaling properties of human mobility," *Nature Physics 6, pp. 818-823*, 2010.

[70] I. Rhee, M. Shin, S. Hong, K. Lee, S. J. Kim, and S. Chong, "On the Levy-walk nature of human mobility," *IEEE/ACM transactions on networking (TON)*, vol. 19, no. 3, pp. 630–643, 2011.

[71] P. Sapiezynski, A. Stopczynski, R. Gatej, and S. Lehmann, "Tracking human mobility using Wi-Fi signals," *PloS one*, vol. 10, no. 7, p. e0130824, 2015.

[72] G. Michau, A. Nantes, A. Bhaskar, E. Chung, P. Abry, and P. Borgnat, "Bluetooth data in an urban context: Retrieving vehicle trajectories," *IEEE Transactions on Intelligent Transportation Systems*, 2017.

[73] T. Tsubota, A. Bhaskar, E. Chung, and R. Billot, "Arterial traffic congestion analysis using bluetooth duration data," 2011.

[74] Department for Transportation - Washington State, "Error Modeling and Analysis for Travel Time Data Obtained from Bluetooth MAC Address Matching," URL: http://www.wsdot.wa.gov/research/reports/fullreports/782.1.pdf.

[75] Y. Yoshimura, A. Amini, S. Sobolevsky, J. Blat, and C. Ratti, "Analysis of pedestrian behaviors through non-invasive bluetooth monitoring," *Applied geography*, vol. 81, pp. 43–51, 2017.

[76] Transport for London, "Review of the TfL WiFi pilot," URL: http://content.tfl.gov.uk/review-tfl-wifi-pilot.pdf.

[77] Libelium, "Smartphone, cellular and hands-free mobile phone detection," URL: http://www.libelium.com/products/meshlium/smartphone-detection/.

[78] A. H. Omre and S. Keeping, "Bluetooth Low Energy: wireless connectivity for medical monitoring," *Journal of diabetes science and technology*, vol. 4, no. 2, pp. 457–463, 2010.

[79] Z.-M. Lin, C.-H. Chang, N.-K. Chou, and Y.-H. Lin, "Bluetooth Low Energy (BLE) based blood pressure monitoring system," in *Intelligent Green Building and Smart Grid (IGBSG), 2014 International Conference on*, April 2014, pp. 1–4.

[80] M. Ali, L. Albasha, and H. Al-Nashash, "A Bluetooth Low Energy implantable glucose monitoring system," in *Microwave Conference (EuMC), 2011 41st European.* IEEE, 2011, pp. 1265–1268.

[81] E. Azih - Bluetooth SIG, "John Oliver Questions Outdated 911 Call Centers Bluetooth Is The Answer," https://blog.bluetooth.com/john-oliver-questions-woefully-outdated-9-1-1-call-centers-bluetooth-has-the-answer, June 2016.

[82] P. Lewis - Think With Google, "How Beacons Technology Can Reshape Retail Marketing," https://www.thinkwithgoogle.com/marketing-resources/retail-marketing-beacon-technology/, Aug 2016.

[83] R. Faragher and R. Harle, "An analysis of the accuracy of Bluetooth Low Energy for indoor positioning applications," in *Proceedings of the 27th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+14)*, 2014, pp. 201–210.

[84] Y. Hadas and B. B. Moshe, "Bluetooth Low-Energy based System for Automatic Public-Transport passengers' Movement data collection," *Second International Workshop Automated Data Collection Systems*, 2016. [Online]. Available: http://www.northeastern.edu/transitdata2016/wp-content/uploads/2016/08/Hadas_Yuval.pdf

[85] W. Narzt, S. Mayerhofer, O. Weichselbaum, S. Haselböck, and N. Höfler, "Be-in/be-out with Bluetooth Low Energy: Implicit ticketing for public transportation systems," in *Intelligent Transportation Systems (ITSC), 2015 IEEE 18th International Conference on.* IEEE, 2015, pp. 1551–1556.

[86] Z. Kleinman - BBC News (Technology), "New seat alert app for pregnant commuters," http://www.bbc.com/news/technology-38690811, Jan 2017.

[87] T. Andersson, "Bluetooth Low Energy and Smartphones for Proximity-Based Automatic Door Locks," p. 10, 2014.

[88] HAPI.com, "HAPIfork: Eat slowly, lose weight, feel great!" https://www.hapi.com/product/hapifork, Aug 2017.

[89] Let's Talk Payments, "7 Exciting Use Cases Highlighting BLE based Payments," https://letstalkpayments.com/7-exciting-use-cases-highlighting-ble-based-payments/, Jan 2015.

[90] "Digicash Mobile Payments NFC + BLE," http://www.digicash.lu/digicash-beacon, Aug 2017.

[91] Tile Inc., "Tile tracker tag," https://www.thetileapp.com/, Aug 2017.

[92] PitPatPet Ltd, "pitpat - Dog Activity Monitor," http://www.pitpatpet.com/, Aug 2017.

[93] Nordic Semiconductor ASA, "Nordic Semiconductor launches nRF51824 Bluetooth low energy SoC for latest connected car intelligent automotive applications and wireless in-car charging," http://bit.ly/2h5yNuw, June 2016.

[94] D. Gilbert - International Business Times, "Connected Cattle: How wearables and the cloud help farmers get their cows pregnant," http://www.ibtimes.co.uk/connected-cattle-how-wearables-cloud-help-farmers-get-their-cows-pregnant-1499220, Apr 2015.

[95] M. Tailor - Wireless Connectivity Blog, Laird Technologies Inc., "Whats New with Bluetooth 5.0?" http://www.summitdata.com/blog/whats-new-with-bluetooth-5-0/, Feb 2017.

[96] S. O'Kane - The Verge, "The Samsung Galaxy S8 is the first phone with Bluetooth 5.0," https://www.theverge.com/2017/3/29/15112646/samsung-galaxy-s8-bluetooth-5-headphones, Mar 2017.

[97] Nordic Semiconductor ASA, "nRF52840 Advanced multi-protocol SoC," https://www.nordicsemi.com/eng/Products/nRF52840, Feb 2017.

[98] O. Bauck - Nordic Semiconductor ASA, "nRF52840 Long range demo," https://devzone.nordicsemi.com/blogs/1076/nrf52840-long-range-demo/, Feb 2017.

[99] Y. Xiao, P. Savolainen, A. Karppanen, M. Siekkinen, and A. Ylä-Jääski, "Practical power modeling of data transmission over 802.11g for wireless applications," in *Proceedings of the 1st International Conference on Energy-efficient Computing and Networking*. ACM, 2010, pp. 75–84.

[100] P. Papadimitratos, A. De La Fortelle, K. Evenssen, R. Brignolo, and S. Cosenza, "Vehicular communication systems: Enabling technologies, applications, and future outlook on intelligent transportation," *IEEE Communications Magazine*, vol. 47, no. 11, 2009.

[101] J. Ploeg, B. Scheepers, E. van Nunen, N. van de Wouw, and H. Nijmeijer, "Design and experimental evaluation of cooperative adaptive cruise control," in *Intelligent Transportation Systems (ITSC), 2011 14th International IEEE Conference on*, Oct 2011, pp. 260–265.

[102] Department for Transport Driver and Vehicle Standards Agency, United Kingdom, "The Highway Code: typical stopping distances," URL: https://assets.publishing.service.gov.uk/media/559afb11ed915d1595000017/the-highway-code-typical-stopping-distances.pdf.

[103] Apple Inc., "Apple iBeacon," URL: https://developer.apple.com/ibeacon/.

[104] E. Giordano, R. Frank, G. Pau, and M. Gerla, "Corner: a realistic urban propagation model for vanet," in *Wireless On-demand Network Systems and Services (WONS), 2010 Seventh International Conference on*. IEEE, 2010, pp. 57–60.

[105] E. G. Villegas, E. Lopez-Aguilera, R. Vidal, and J. Paradells, "Effect of adjacent-channel interference in IEEE 802.11 WLANs," in *Cognitive Radio Oriented Wireless Networks and Communications, 2007. CrownCom 2007. 2nd International Conference on*, Aug 2007, pp. 118–125.

[106] S. Faye, N. Louveton, G. Gheorghe, and T. Engel, "A two-level approach to characterizing human activities from wearable sensor data," *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA)*, Sep. 2016.

[107] N. D. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, and A. T. Campbell, "A survey of mobile phone sensing," *Communications Magazine, IEEE*, vol. 48, no. 9, pp. 140–150, 2010.

[108] S. Faye and T. Engel, "Understanding user daily mobility using mobile and wearable sensing systems," in *International Conference on Information and Communications Technology Convergence 2016 (ICTC 2016)*, Oct. 2016.

[109] H. Liu, J. Yang, S. Sidhom, Y. Wang, Y. Chen, and F. Ye, "Accurate Wi-Fi based localization for smartphones using peer assistance," *IEEE Transactions on Mobile Computing*, vol. 13, no. 10, pp. 2199–2214, 2014.

[110] T. Ellersiek, G. Andrienko, N. Andrienko, D. Hecker, H. Stange, and M. Mueller, "Using Bluetooth to track mobility patterns: depicting its potential based on various case studies," in *Proceedings of the Fifth ACM SIGSPATIAL International Workshop on Indoor Spatial Awareness*. ACM, 2013, pp. 1–7.

[111] Y. Chen, Z. Chen, J. Liu, D. H. Hu, and Q. Yang, "Surrounding context and episode awareness using dynamic Bluetooth data," in *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, 2012, pp. 629–630.

[112] P. Aditya, V. Erdélyi, M. Lentz, E. Shi, B. Bhattacharjee, and P. Druschel, "Encore: private, context-based communication for mobile social apps," in *Proceedings of the 12th annual international conference on mobile systems, applications, and services*. ACM, 2014, pp. 135–148.

[113] J. J. Anaya, P. Merdrignac, O. Shagdar, F. Nashashibi, and J. E. Naranjo, "Vehicle to pedestrian communications for protection of vulnerable road users," in *2014 IEEE Intelligent Vehicles Symposium Proceedings*. IEEE, 2014, pp. 1037–1042.

[114] General Motors Inc., "GM Developing Wireless Pedestrian Detection Technology," URL: http://media.gm.com/media/us/en/gm/news.detail.html/content/Pages/news/us/en/2012/Jul/0726_pedestrian.html, June 2012.

[115] C. Moeller, "osm2po - OpenStreetMap converter and routing engine for Java," URL: http://osm2po.de/.

[116] M. Haklay and P. Weber, "OpenStreetMap: User-generated street maps," *IEEE Pervasive Computing*, vol. 7, no. 4, pp. 12–18, 2008.

[117] L. Codeca, R. Frank, S. Faye, and T. Engel, "Luxembourg SUMO Traffic (LuST) Scenario: Traffic Demand Evaluation," *IEEE Intelligent Transportation Systems Magazine*, vol. 9, no. 2, pp. 52–63, 2017.

[118] M. A. Hall, "Correlation-based feature selection for machine learning," *PhD Thesis Contribution, University of Waikato Hamilton*, 1999.

[119] T. M. Mitchell, *The discipline of machine learning*. Carnegie Mellon University, School of Computer Science, Machine Learning Department, 2006, vol. 3.

[120] A. Liaw and M. Wiener, "Classification and regression by randomforest," *R news*, vol. 2, no. 3, pp. 18–22, 2002.

[121] B. Scholkopf and A. J. Smola, *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2001.

[122] P. O. Glauner, A. Boechat, L. Dolberg, R. State, F. Bettinger, Y. Rangoni, and D. Duarte, "Large-scale detection of non-technical losses in imbalanced data sets," *arXiv preprint arXiv:1602.08350*, 2016.

[123] C. M. Bishop, *Pattern recognition and machine learning.* Springer, 2006.

[124] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.

[125] S. Faye and C. Chaudet, "Characterizing the topology of an urban wireless sensor network for road traffic management," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 7, pp. 5720–5725, Jul. 2016.

[126] A. Balakrishnan - CNBC, "Apple has an idea for car sensors that could drastically reduce crashes," URL: https://www.cnbc.com/2017/08/17/apple-patent-shows-connected-car-sensors-v2v.html, Aug 2017.