Translating BLE Received Signals across different Smartphones

Internship Report
Geomatics for the Built Environment

Dimitris Xenakis

TU Delft

October 2018

Cover:
Illustration of different Smartphones "reacting" to the same Bluetooth signal, in a different way.

Translating BLE Received Signals across different Smartphones

Internship Report
Geomatics for the Built Environment

by

Dimitris Xenakis

Date:	October 2018	
Supervisors:	Edward Verbree	TU Delft
	Alex de Mulder	Crownstone
	Bart van Vliet	Crownstone
	Anne van Rossum	Crownstone

This report forms the written record of an internship carried out in the period 16/7/2018 - 16/9/2018. The internship is an elective of the MSc Geomatics for the Built Environment, Delft University of Technology, awarded with 5ECTS.

Abstract

This report constitutes the final product of my internship, where research was conducted on the differences of the Bluetooth Low Energy signal reception between 2 different phones *(p1, p2)* and eventually, the possibility of developing a translation function that could be used to predict the signal strength reception of *p1*, by considering the signal strength reception of *p2*. Such model would be particularly useful to applications related to Indoor Localization/Positioning, as these are often based on the BLE signal strength.

For the development of this model, the influence of several parameters was assessed, such as: a) the distance between a phone and a beacon, b) their orientations and c) the number of concurrently broadcasting beacons, and all were found to be significant. Furthermore, it was discovered that as long as there is no movement in the system, the BLE signal reception at a specific channel has low variations and so, even a few samples can be representative for each channel.

The evaluation of the translation functions was quite promising. Ultimately, by taking advantage of a specific Android's behavior during the training phase, it became possible to identify the channels of incoming BLE signals. This information was then used to significantly enhance the performance of the translations under specific circumstances (i.e. the channels can be identified during the operational phase too).

List of Figures

List of Tables

Contents

1

Introduction

1.1. Company's operation description

Crownstone is a new (founded in 2016) start-up company, the head office of which is located in Rotterdam. It is powered by Almende Investments (`http://almende-investments.com/`), an independent research and investment company that conducts since 2000, research and development activities in several ICT domains, ranging from healthcare to manufacturing systems. At the same time, Crownstone is also supported by DoBots, a company specializing in robotics-orientated software that offers solutions for Smart-Buildings and Smart-Robots (`https://dobots.nl/`).

Crownstone offers its expertise in the field of IoT and Indoor Localization via their specially designed modules (the crownstones), that can monitor and control the power supply of devices connected to them, based on user's proximity and/or automated or on-demand calls. These modules operate normally as nodes in a network and their coordinated functionality allows for numerous Indoor Localization applications. For example, one could easily ensure that a specific device will not be left turned on or on stand-by, when not being present.

The company has a long-established tradition of getting involved in research (especially thanks to DoBots and Almende), with many interesting topics related to the undertaken one, having emerged. Some recent examples are:

- **Indoor Localization and Fingerprinting across Multiple Smartphones** [1]
- **Indoor localization using BLE - Using Bluetooth Low Energy for room-level localization** [2]
- **Human SLAM - Simultaneous Localisation and Configuration (SLAC) of indoor Wireless Sensor Networks and theirs users** [3]

1.2. Project description

Crownstone's Indoor Localization service is fundamentally based on the Bluetooth Low Energy (BLE) Technology. More specifically, when a set of crownstone modules, broadcasting BLE signals, is deployed in space (forming a mesh), a radio distribution occurs within that space. Crownstone binds each location to the corresponding signal signature at that location (during a fingerprinting training phase), creating a reference database which enables the localization through a reverse lookup - from signal signature to location. A smartphone is typically the device used to scan all these signals and process them, delivering ultimately an estimation for current location.

Unfortunately, the above summarization only reveals half the story. The other half lies in the complexity of receiving a consistent signal signature at a given position (also evident during the experimental phase from Haagmans (2017) [4]), consequently affecting the accuracy of the positioning estimation. There is a vast number of random parameters affecting the signal propagation and its final form by the time it reaches the phone's Bluetooth antenna. To begin with, BLE operates in the 2.4 GHz ISM band, where the signal is respectively subject to various reflective and absorbing obstructions. These are caused by different types of barriers, the interference potential of which may be less (e.g. wood), or higher (e.g. metal), with the human body (being made up of mostly water) lying somewhere in the middle. For example, one can expect significant signal fluctuations in a building crowded with people moving around. Moreover, the same 2.4 GHz frequency is also being used by many other radio devices intended for industrial, scientific, or medical (ISM) requirements. From Wi-Fi routers and microwave ovens, to fluorescent lighting in offices, all can potentially become sources of interference.

However, even if all environmental factors could be eliminated from the equation, different Bluetooth-enabled phones (having different hardware/firmware) would still “react” to an incoming BLE signal, in a different way. Undoubtedly, trying to match a new-read signal signature at a given position with one inside the reference database, would require that between the training and operational moment, all parameters above are as similar as possible. Otherwise, an analogous uncertainty would be introduced to the positioning estimation. For that reason, Indoor Positioning services often require that, at least, the same device is used during both the training and the operational phase; and this is also the case for Crownstone. Therefore, the leading object of this research internship is to approach a solution for that problem. Namely, the necessity of using the same devices among these two phases, to be able to achieve accurate positioning estimations. This suggests the development of a model that could support “translations” between signal readings from different devices.

Ideally, the following basic scenario should be feasible: A smartphone could be used to initially train the reference database, while scanning under specific known conditions (e.g. phone model, device placement, etc). At a later moment (operational phase) and under different known conditions, another smartphone should be able to utilize the initially trained database, to produce an accurate-enough positioning estimation.

To assess the degree to which the above scenario could be accomplished, several research questions need to be addressed. These include:

- **How should the research methodology be defined? Aspects to consider are the data collection, their analysis and the selection of a proper statistical model.**
- **Which parameters should optimally (having good “value-complexity” ratio) be considered and how?**
- **Is it possible to develop a model that could support "translations" between signal readings?**
- **Which measure of performance should be used to evaluate these signal translations?**

2

Research methodology

As mentioned, the general objective is to see whether there are phones of different BLE reception trends, such that this difference could be modelled and generally used to translate a received signal from one phone to the other. For that reason, trying to accurately capture this difference is crucial for the performance of the translations. There are factors, however, such as Transmitter-Receiver distance and their orientations, the unpredictableness of which might introduce ambiguity while searching for these differences. For example, would it be sufficient to only place the phones and beacons at a single orientation while gathering the data? Or will this produce biased conclusions? Therefore, it makes sense to start determining the proper methodology and thus, answering the 1st research question, by researching how these could affect the RSSI. Since in a real case scenario, the beacon-phone distance and especially their orientations are unknown, it would eventually be practical if a translation could be achieved without being actually aware of them.

2.1. Assessing the influence of phone's orientation

The first phase is to check whether (and how) the phone's orientations could affect the differences in RSSI between 2 phones. For example, for a specific distance, if the difference in the RSSI distributions between 2 phones did not remain the same while changing their orientations in the same way, it means that the phone's orientation (as a factor) plays a significant role. Contrariwise, if the differences remained the same, then this would be an indication that the orientations do not affect the RSSI distributions as much as other factors.

To investigate this, 2 mount-stands were built. One to place a specific beacon (B1) and another one to place each phone during the RSSI gathering. These mounts were then placed in a fixed (2m) distance in the middle of an open field (Figure 2.1) with closest (out-of-system) reflecting surface being further than 40 meters. It should be noticed that B1 was broadcasting every 100ms at a specific Tx Power, which has been measured that can effectively send a signal to a maximum distance of 50 meters. During the whole process, the beacon remained fixed and the only changing parameter was the orientation of the phone that was being mounted. In total, 4 different phones were sequentially used for fast (of low latency) data gathering and, on each case, all other Radio adapters were disabled. The orientation placements that were tested are shown in Figure 2.2 and on each case, the phone was remotely (the user was leaving first) gathering RSS values for 90 seconds.

During the data gathering and analysis, several interesting facts became evident. To begin with, the following was observed. Normally, when Android devices are listening for BLE signals, they are switching (cycling) between the 3 BLE Advertisement channels (37-38-39 of Figure 2.3) very fast (≈2 rotations per second). In this specific case, however, whenever the signal was found in the beginning of the scanning to be stable enough (i.e. there was no significant movement in the system) and also the WiFi adapter was disabled, then the device was listening on a specific channel for 5 seconds and then it changed to the next one for another 5s and so on (i.e. 1 rotation every 15 seconds). For that reason and to avoid being biased towards a specific channel, the data that were further used in the analysis were clipped based on a scan duration that is a multiple of 15. As such, the first 60 seconds were considered on every different orientation scan.

Figure 2.1: Sampling Field

Figure 2.2: Orientations being tested (ZXY Angles in degrees)

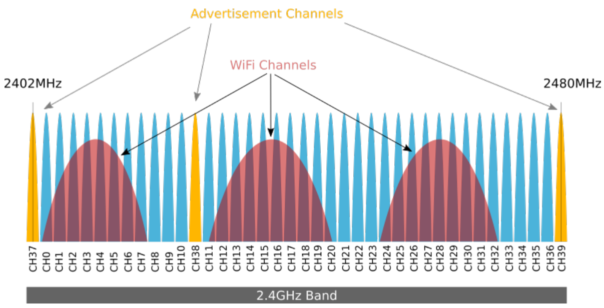

Figure 2.3: The 3 BLE Advertisement Channels in 2.4GHz Band

Furthermore, it was noticed that almost on every orientation placement, the RSSI at each distinct channel was softly variating around a specific dBm level. Due to this fact, a kernel density estimation (a mixture of several Gaussian distributions instead of a single one) was used to estimate more accurately the probability density function (PDF) of the RSSI. This can be seen in the following group of graphs (Figure 2.4 - Figure 2.8) that show a sequence of 45-degrees rotation on the Y axis graphs. For example, in the 1st case (phone's rotation 90-0-0), the 3 black arrows show (for the Sony Xperia Z2) the 3 main dBm levels for the 3 corresponding channels. These graphs (top right plots) also illustrate clearly the 15 second periodic channel shift.

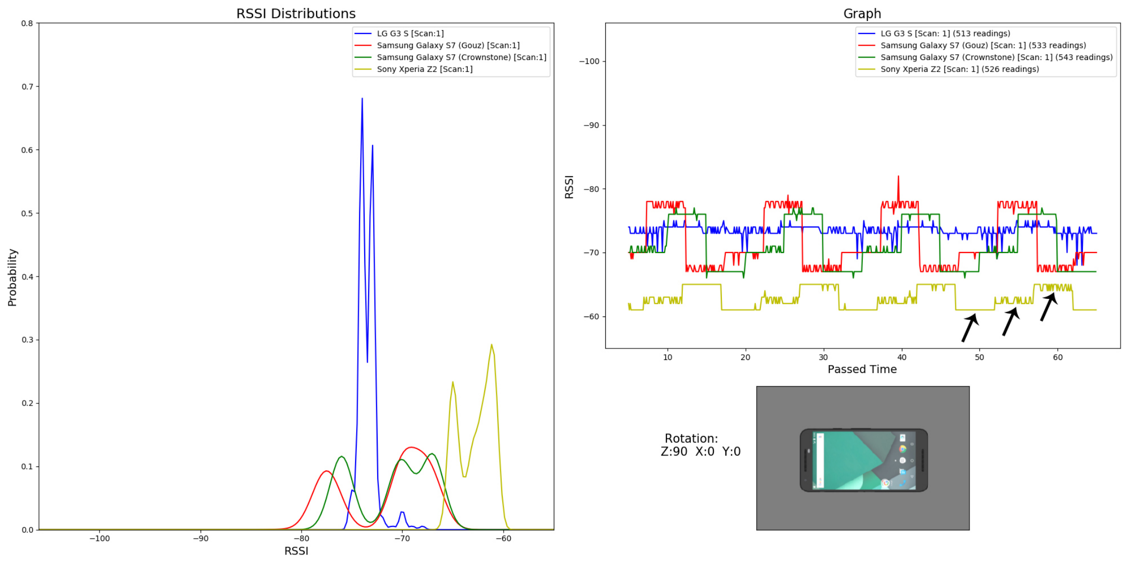

Figure 2.4: RSSI analysis (Phone's Orientation ZXY: 90-0-0)

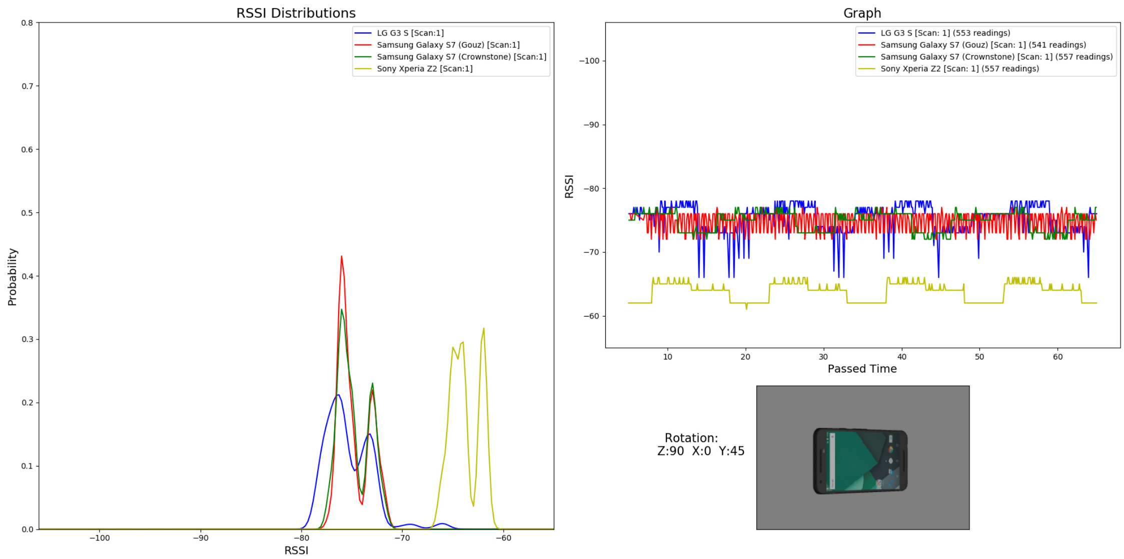

Figure 2.5: RSSI analysis (Phone's Orientation ZXY: 90-0-45)

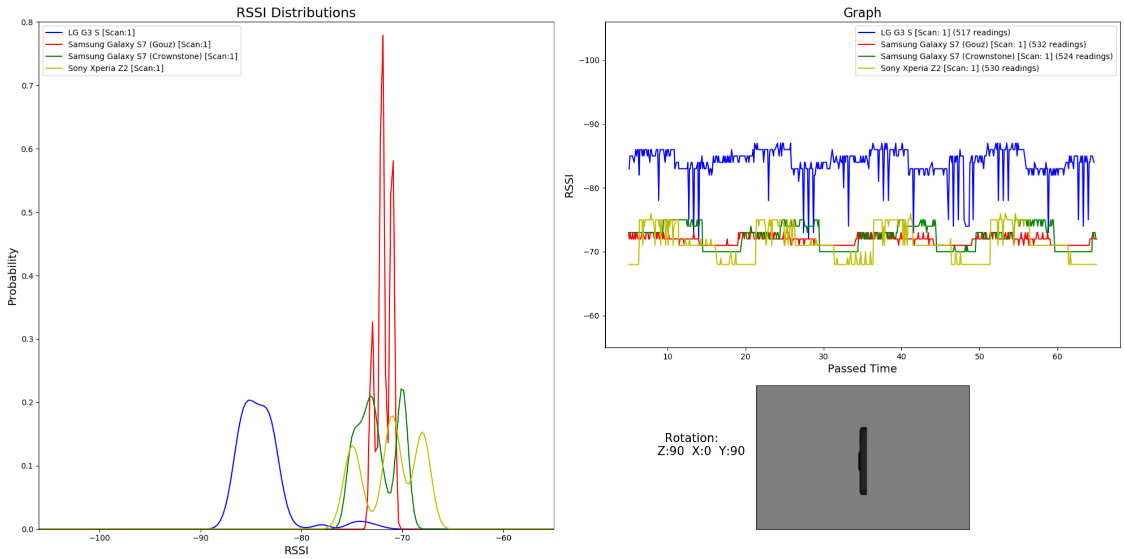

Figure 2.6: RSSI analysis (Phone's Orientation ZXY: 90-0-90)

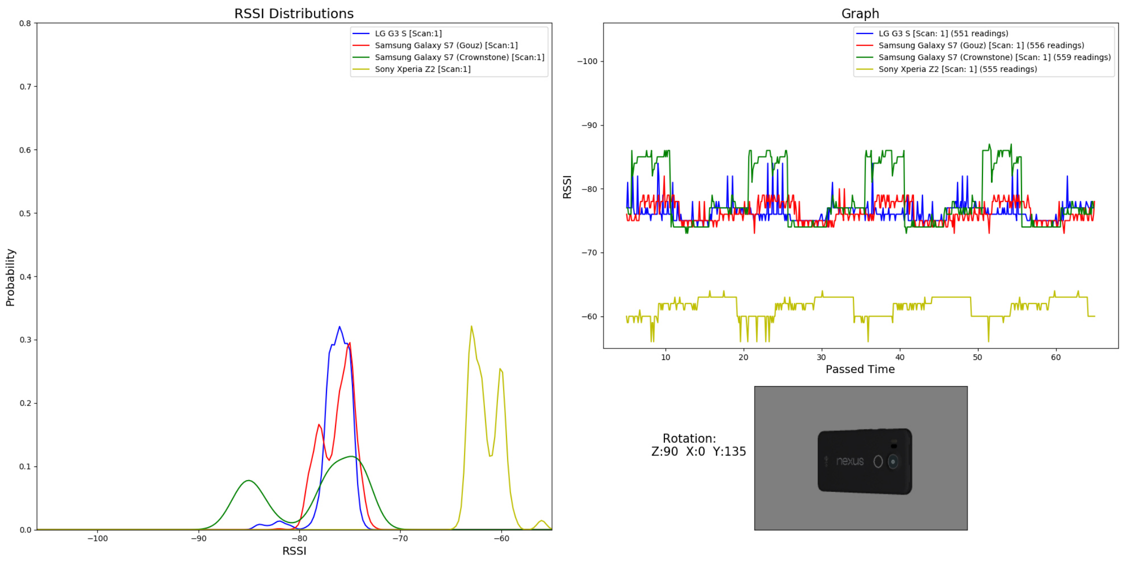

Figure 2.7: RSSI analysis (Phone's Orientation ZXY: 90-0-135)

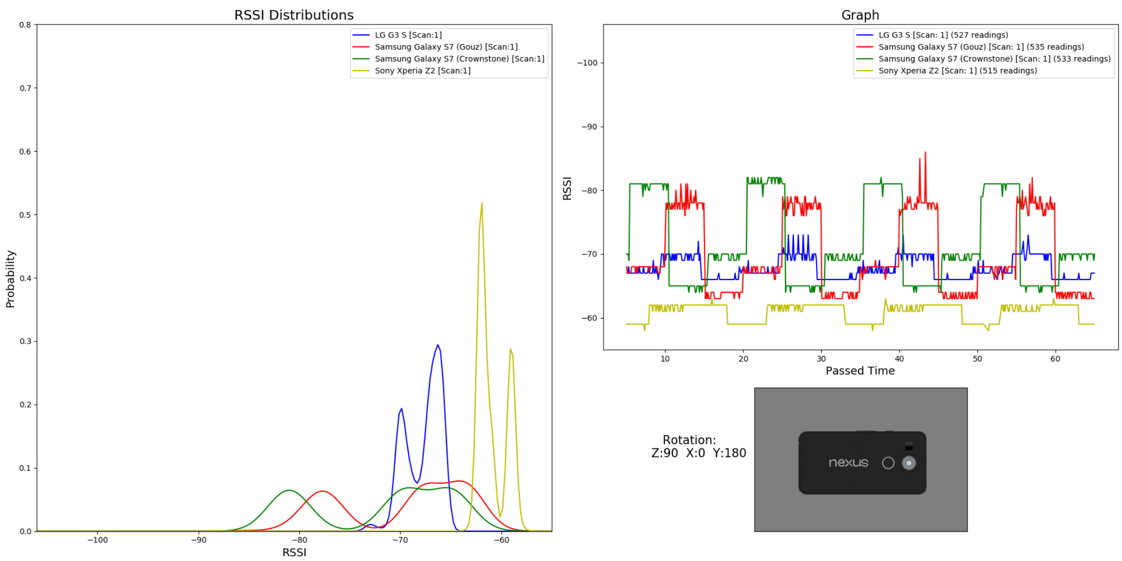

Figure 2.8: RSSI analysis (Phone's Orientation ZXY: 90-0-180)

Using a Channel-Reference technique that will be described in section 2.4, it was found that (for each phone) the relationships between the 3 main dBm levels were changing as soon as the orientation was changing (i.e. at some orientations, a specific channel was delivering the strongest signal and at other orientations the same channel was delivering the weakest). That, along with the fact that Androids do not reveal the channel of a single RSSI, might be a good reasoning why the RSSI records should better be considered combined and possibly as a single Gaussian distribution (to minimize this way, the computational complexity).

Proceeding with the examination of the differences between the RSSI distributions, across different phone's-orientations, the above graphs show that, indeed, between 2 phones, the orientation plays a significant role. For example, although the Sony Xperia seems to have often the strongest signal, in the 3rd case (90-0-90), its reception was similar to the Galaxy S7 devices. Conclusively, by examining also the rest of the orientation cases of Figure 2.2, it became clear that the RSSI distributions were so sensitive (and "random") to phone's orientation changes, that checking only a single phone's-orientation placement has to be avoided (in order not to become biased) and instead, consider many orientations during a single scan (leading inevitably to even wider RSSI distributions).

However, which exactly should these orientations be, since most of them (i.e. their RSSI distributions) differ significantly from the rest ones. Without doubt, in a real case scenario, the orientation of a phone (in respect to the beacon's position) would be very difficult to predict. As such, gathering RSS data from as many phone orientations as possible would be the way to go, while at the same time, this should be done uniformly (unbiased towards specific orientations). This, led to purchasing the mechanical (to avoid any electro-magnetic interference coming from the motor) 360° in 1 hour rotating mount of Figure 2.9, which offers one complete linear-rotation on a single axis. Ideally, maybe, a 3-axis gimbal mount (similar to the one showed in Figure 2.10) should be used during a 360°x360°x360° Scan, however such equipment is not easily available.

Figure 2.9: 360° mechanically rotating mount for the phones

Figure 2.10: An even better (probably) concept for a mount

2.2. Assessing the influence of beacon's orientation and uniqueness

The analysis above considered the phone's orientation in respect to the beacon. However, the possibility that the beacon's orientation in respect to the phone's position could also affect the difference (among phones) in (their) RSSI distributions, was evident. Therefore, the same experiment as above was repeated (namely the: Duration: 60sec, Distance: fixed, Tx Power: Fixed, Broadcast Interval: Fixed, Beacon's Orientation: fixed, Phone's Orientation: alternating), but at a reversed form (i.e. Duration: 60sec, Distance: fixed, Tx Power: Fixed, Broadcast Interval: Fixed, Beacon's Orientation: alternating, Phone's Orientation: fixed). Examining this factor was important, because if results would show that the beacon's orientation does affect the RSSI distribution differences, then the same approach as in the case of the phone's orientation should be used for the beacons too (i.e. sampling from different beacon's orientations too). The orientations that were put to the test are the ones being shown in Figure 2.11. Additionally, it is worth mentioning that in each case, the readings across all channels were combined into a single Gaussian distribution to make comparisons simpler.

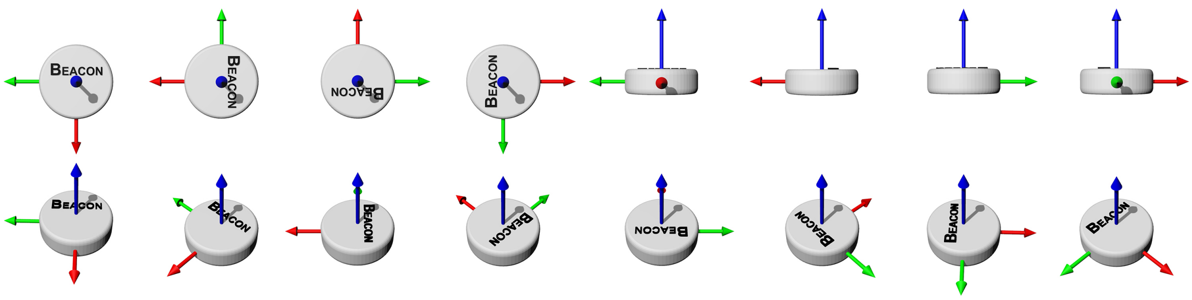

Figure 2.11: Beacon orientations that were tested

The aforementioned experiment was repeated 2 times, for 3 different beacons of same model (i.e. 2x3 = 6 times). This alternation was used to reject (or not) the null hypothesis that different beacons of same exact model, orientation and hardware/software configurations, do not lead to different RSSI distributions on a specific Phone. This was also important because, eventually, multiple concurrent beacons were intended to be used, to simultaneously log multiple emitted signals (and thus, fasten the whole process). If, however, this experiment rejected the null hypothesis, it would mean that different beacons (of same type) broadcast in a different way, which suggests that the use of multiple beacons should be avoided.

The results of this twofold experiment gave early answers to both research questions above. First of all, similarly to the phone's orientation case, the beacon's orientation has also great impact on the differences between the RSSI of two different phones. For example, just by rotating for a few degrees the beacon, a phone that had stronger signal (compared to another phone), could now easily have weaker signal. This becomes clear by examining the relationship shift between the Sony Xperia (brown distribution) and LG G3 (blue distribution) in Figure 2.12, after rotating the beacon for 90 degrees.

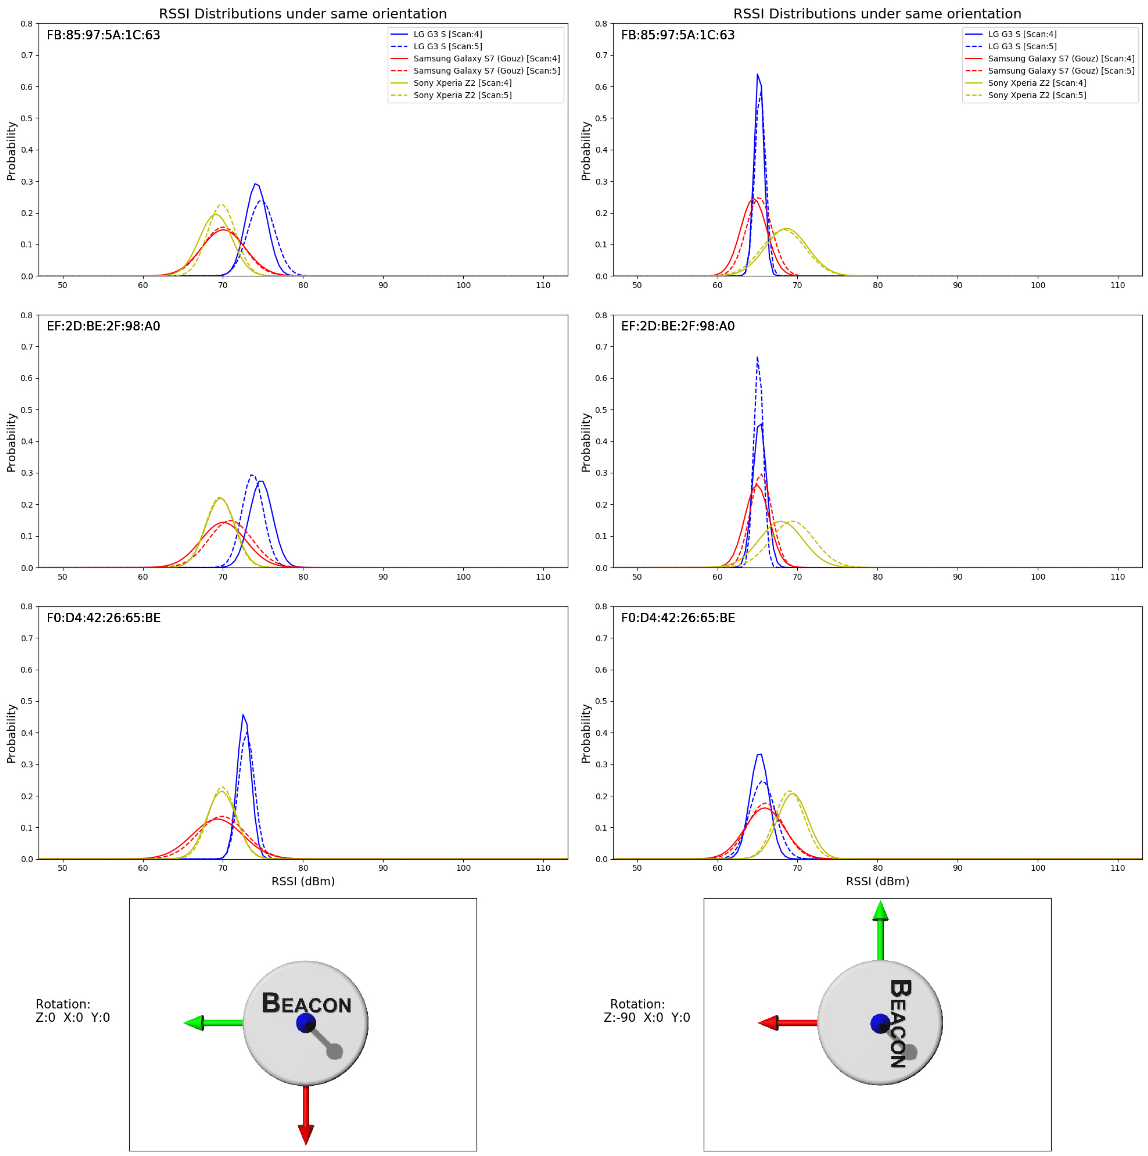

Figure 2.12: Comparing the signal of three different Beacons, while having 2 different orientations

Similarly to the 1st experiment (influence of phone's orientation), these new results also suggest that recording only the emitted signal from a single beacon orientation, is not enough. Ideally, another 3-axis gimbal mount mechanism should be used for the beacon too. This, however, leads to a great number of rotational combinations between the phone and the beacon and thus, much time needed for the data gathering. For example, if each angle step (i.e. the rotational accuracy) was 10 degrees, then the total rotational combinations to check would be: 36°x36°x36° phone rotations x 36°x36°x36° beacon rotations, equalling to more than 2 billions rotations (although, by considering only the front hemisphere of the beacon, as the back one is always attached to a surface and never visible, would lower that number). For that reason, having considered the corresponding "value-complexity" ratio (subject of the 2nd research question), the only sensible option was to use at the same time, many beacons placed at various random orientations.

The same figure also reveals another finding, regarding the influence of beacon's uniqueness. It is clear that all three beacons (having distinct MAC addresses) had both times (orientation cases) very similar broadcast footprints, which means that each beacon of the same model behaves the same. Finally, as already mentioned, the experiment was repeated twice (straight line with ID Scan:4 and dashed line with ID Scan:5) to ensure that the results were consistent. A consistency that is again evident in the same figure.

2.3. Examining the effects of a non-uniform (orientation-wise) sampling

As it has been shown, the various phone's/beacon's orientations can easily influence the BLE RSSI differences between 2 phones. However, to get a deeper understanding on the possible effects of a non-uniform (orientation-wise) sampling, when trying to develop an accurate RSSI translation model between 2 phones, the following experiment was performed: The 360°/hour mechanical rotator shown in Figure 2.9 was placed on a vertical mount. All around it (forming a circle), 20 beacons were placed on their own vertical mounts, at an incremental distance starting from 0.5m to 10m (i.e. incremental step of 0.5m). All beacons had the same exact orientation in respect to the phone mount, for their broadcasting to be comparable. Two phones were then sequentially placed on the centered mount, each, at 3 different orientations (Horizontal, Vertical, Flat). In total, this setup produced 120 sets of 1-hour scan readings (2 phones x 3 orientations x 20 beacons). Finally, a Gaussian distribution was fitted to each set and the combined results are presented in Figure 2.13.

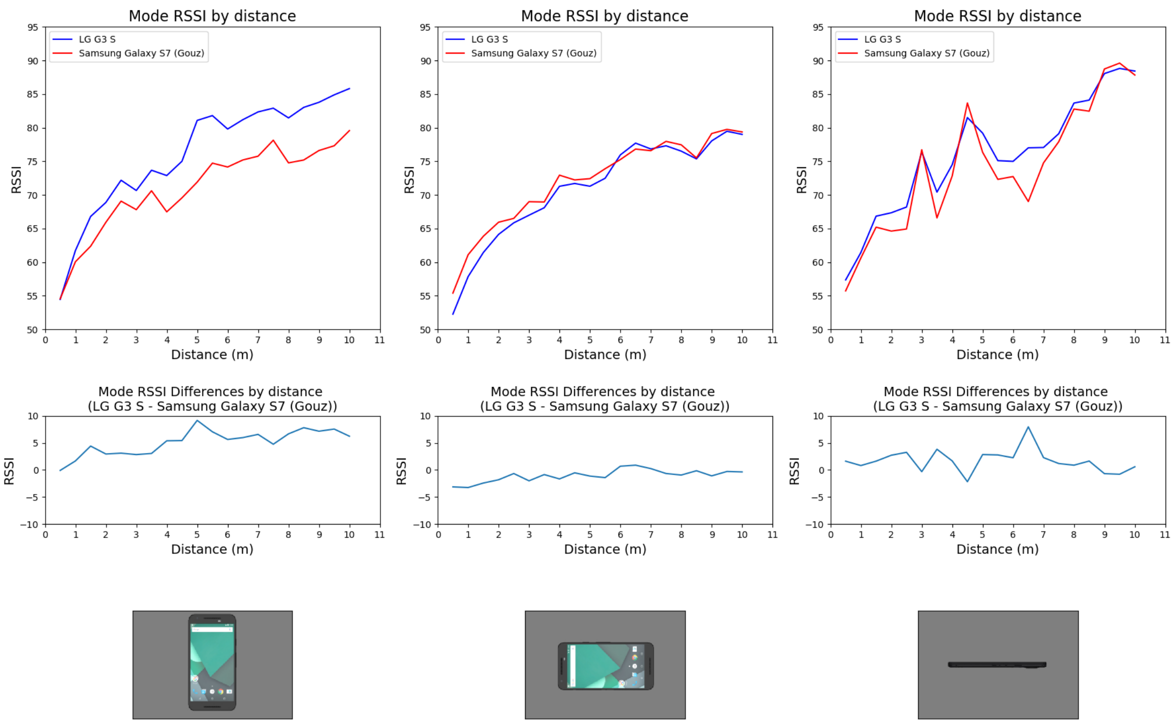

Figure 2.13: Comparing the RSSI difference (based on distance), between 2 phones

The above plots show the RSSI differences of two phones (LG G3 - Samsung Galaxy S7) based on their distances to the BLE source. As already noticed, in practice, it is not possible to know this parameter during the operation mode, which adds a level obscurity to the readings. However, having considered now the distance parameter too, provides us with an extra level of information (the trend based on distance) and thus, makes it easier for us to intuitively understand the effects of a non-uniform (orientation-wise) sampling.

More specifically, assuming that the mount was slightly turning every 1 second, then within a 360°/hour rotation, approximately 3600 (60 seconds x 60 minutes) different angles were sampled, which indeed sounds a lot and thus, possibly sufficient. However, this motion still corresponds to only 1 degree of freedom (Yaw), without considering the other 2 (Pitch & Roll). As mentioned, this scanning process was executed 3 times (Horizontal, Vertical and Flat placement of the phone) and these 3 cases were compared in Figure 2.13. One might expect that in all 3 cases the trends would be similar (consistent), however, this is not the case. At the Horizontal and Vertical placements, the differences remain more or less the same as distance increases, but at the Vertical placement, the differences seem to increase. Hence, if the Vertical placement hadn't been checked at all, this (valuable to the translation model) information would be missing, producing less accurate results. Equivalently, it is yet unknown what other "valuable" information we are missing, having only checked 3 placements. Moreover, this problem also illustrates the great difficulty in developing and utilizing an accurate Signal Strength to Distance model, such as the well-known Log-Distance Path Loss model [5].

The importance of considering not only the two orientation parameters (beacon's & phone's), but also the parameter of distance (as we saw from the left graph, it contains valuable information on the RSSI differences between phones), is evident. In section 2.2, it was concluded that many beacons should be used for parallel logging. However, since the parameter of distance has also been introduced, these should ideally be placed at various distances from the recording phones. Moreover, due to the logarithmic change of the signal strength as distance increases (which can also be seen in Figure 2.13), their placement should also favor the short distances. That, in order to have balanced RSSI samples (i.e. about the same number of weak and strong signals) available, for the development of the translation model. These conclusions form an educated response to the 2nd research question.

2.4. Estimating the ideal scanning duration

In the group of graphs Figure 2.4 - Figure 2.8 it was shown that during a scan, the readings at a specific channel did not have significant variations. This, suggested that collecting data for too long would not add much information to the corresponding distribution and so, the ideal scanning duration had to be determined. This, would ultimately decrease the total time needed for the future scanning processes.

However, such task also required knowing the exact channel at which each signal was detected, which is something that Android's API do not provide yet. Instead, 3 BLE development kits from NORDIC (Figure 2.14) were utilized as Channel-References, each programmed to broadcast on a specific only channel. Then, by exploiting the discovered fact that Androids were continuously scanning for 5 seconds at each channel, it became possible to identify which channel that was, because within the same time-span, only a single Channel-Reference had also been detected.

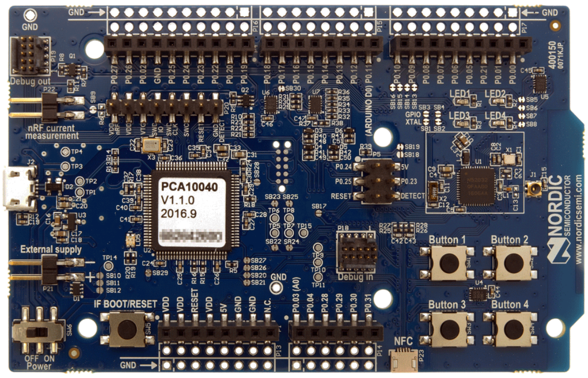

Figure 2.14: nRF52 Development Kit from NORDIC SEMICONDUCTOR

To eventually calculate the ideal scanning duration, the following procedure was implemented. In 2 different indoor environments (IE1, IE2), 20 beacons were (in each IE case) placed at random positions and at random orientations. Additionally, the 3 BLE Channel-References were also placed in the middle of each area. Then, 4 random positions were chosen (again in each IE case) to sequentially place 4 Android devices (i.e. each phone device was placed once on each one of the 4 positions). All phones that were placed at a specific position, shared the same exact orientation, which was different at each position. In the end, 1920 scans were collected (i.e. 2 Areas x 4 Positions x 4 Phones x 20 Beacons x 3 Channels), each containing 25 minutes of RSSI records (corresponding in total to 800 hours of RSSI records). It should also be noted that during each scanning, there was no movement within the environment.

The collected data were processed to develop a sampling minimization function, which could be advised for finding an efficient and effective scanning duration. More specifically, the following data processing was applied: At first, the 1920 scans (having 25-minutes duration each) were checked in terms of their sample size (RSSI records) and this information was plotted along with the underlying Gaussian distributions (some results can be seen in Figure 2.15). It was noticed that no phone had detected the 3 channels of a single beacon, the same amount of times (the 3 numbers at each top right legend), while additionally, no channel was observed to be in general the dominant one. At some degree, this variation might be expected because of several reasons. First of all, a specific channel could possibly be received as "bad", as the phone's threshold-line between detecting/not detecting a weak signal. Also, the transmitting/receiving intervals of the 2 devices are neither perfect nor synced (on the phone side, they are even depended on the available CPU sources). However, in a few cases this phenomenon was more excessive (red underlines showing the number of samples of the least detected channel), or even extreme (green circle).

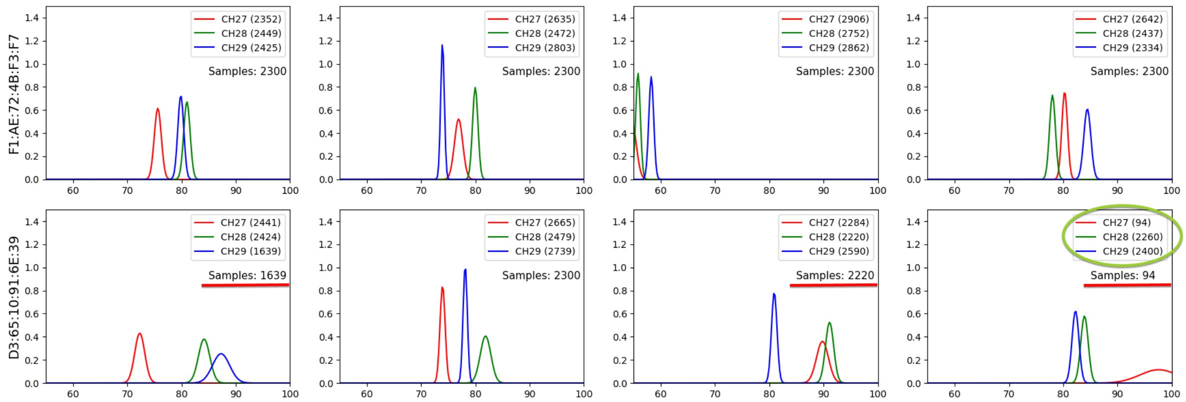

Figure 2.15: LG G3 RSSI records at channel level, from 2 (out of 20) different beacons.
The 4 plot columns correspond to the 4 different positions used.

Next, it was calculated that in ≈95% of the cases, there were at least 2300 samples available. Observing that (on average) it was taking 165ms for the next RSSI to be received during a low latency scanning, it can be deduced that the 2300 samples correspond to ≈19 minutes of recording (2300 samples x 165ms x 3 channels). Then, the mean (μ_c) of the underlying Gaussian distribution of each one of the cases that had at least 2300 samples (1804 out of 1920 scans), was considered as a reference point for the following Monte Carlo method:

- Create an empty **vector-list** *(a list of eventually 1804 vectors, each one of which containing 50 means)*
- For each **scan-case** set *(1804 in total)*
  - Create an empty **mean-list**
  - For each **sample-size** in the sequence: **500 to 10 with step -10** *(i.e. 500, 490, 480, .., 10)*
    - Create an empty **diff-list**
    - Repeat 10000 times
      - Randomly select **sample-size** number of samples from the **scan-case** set
      - Calculate the **new-mean** of the above subset
      - Calculate the absolute difference between the **new-mean** and the corresponding μ_c
      - Add the above difference to the **diff-list**
    - Calculate the **mean** of the **diff-list** *(at this point, its size will be 10000)*
    - Add the **mean** to the **mean-list**
  - Add the **mean-list** *(at this point, its size will be 50: the number of steps)* as a vector to the **vector-list**
- Average all vectors within the **vector-list** *(a final vector containing 50 mean averages will be produced)*

The above algorithm produced the function presented in Figure 2.16. This function shows the relationship between the size of a sub-sample set (of a single channel) and the absolute RSSI difference (i.e. accuracy) between the mean of this sub-sample set and the mean of the entire sample set ("ground truth"). The figure makes it clear that for as long as there is no movement in the system, the signal reception is very stable. More specifically, the average of even just 10 samples has less than 0.25dBm difference with the average of more than 2300 samples.

Nevertheless, the ideal accuracy/samples ratio, seems to be at around 125 samples, where the curvature change decreases. This sample size corresponds to about 20 seconds of scanning on a specific channel. However, as phones do not scan only on 1 channel, this duration needs to be multiplied by 3 (i.e. 1 minute).

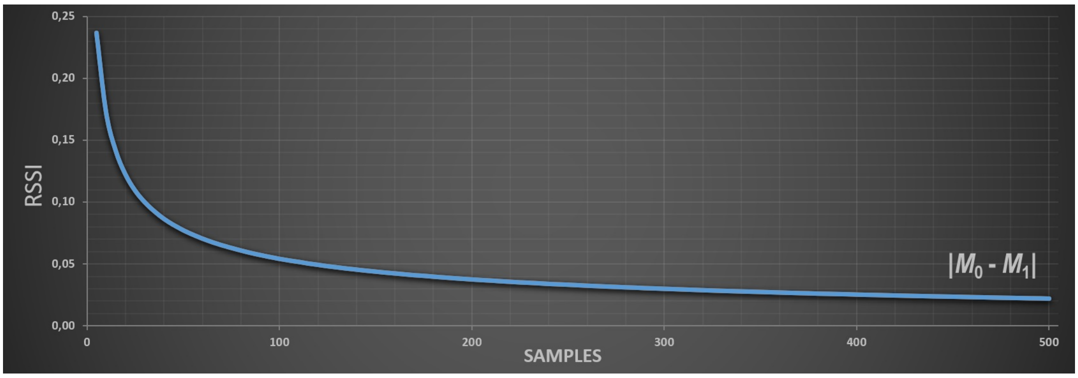

Figure 2.16: Minimization function for the samples

2.5. Assessing the impact of multiple beacon transmissions

In the previous section, it was concluded that 1 minute of scanning offers a good "accuracy/samples" performance, which corresponded to 125 samples per channel. However, as many beacons were about to be simultaneously utilized, it was meaningful to test whether an increased number of beacons could possibly influence the number of RSSI records on a phone. To research that, the following experiment was repeated twice (to confirm consistency) and for 3 phone devices: A device (along with the 3 Channel-References next to it) was placed in a room to do 2 scans, lasting 1 minute each. During the 1st scan, 60 beacons were placed around it, while during the 2nd scan, only 3. The results showed that, indeed, by increasing the number of beacons, they get detected less times. For example, Figure 2.17 shows the results for LG G3, where the left group (1st scan) contains considerably less RSSI detections (samples), when comparing to the right group (2nd scan). A phenomenon that could possibly be depended by the computational power of the device, or even the saturation of the antenna. Accordingly, this was taken into consideration for all future scans.

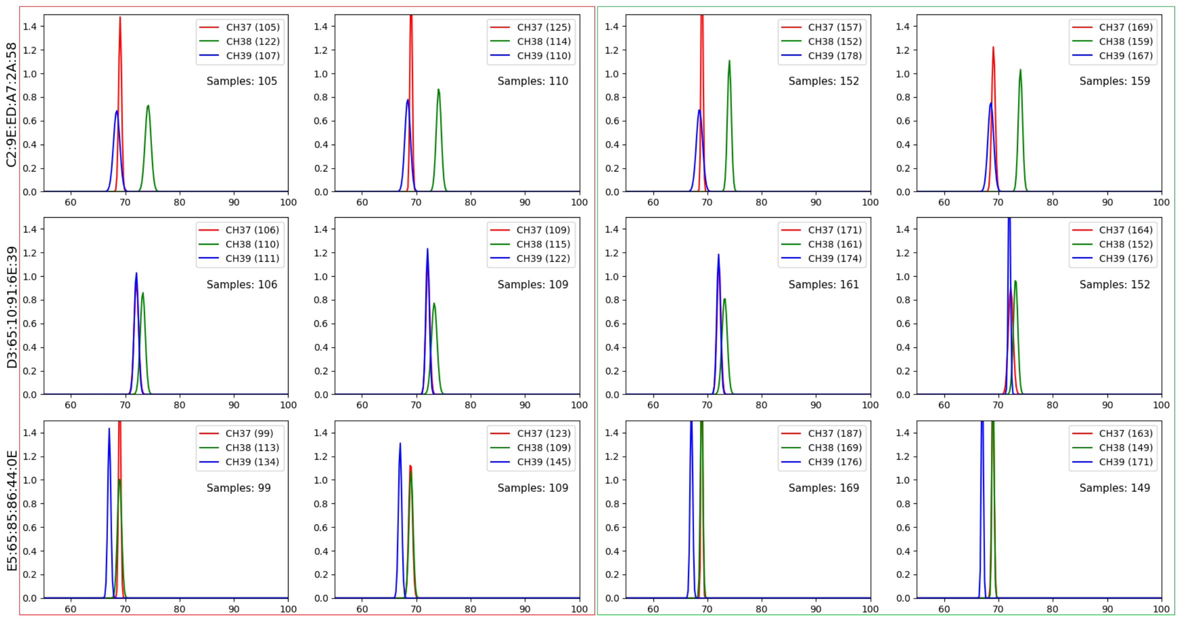

Figure 2.17: Impact of multiple beacon transmissions to the detection response.
(Total transmitting beacons at Left Group: 60 beacons vs Right Group: 3)

3

Resulting translation model and its evaluation

Having defined the data collection methodology, the next step was to finally gather the data needed for the development of the translation function between the phones. Namely, the Sony Xperia (which was found to be the strongest in the the group of graphs Figure 2.4 - Figure 2.8), the LG G3 and the Samsung S7. With that said, in an open area, 60 beacons were placed, forming a circle, at random orientations and positions. In the center, along with the 3 Reference-Channels, the 3 phones were sequentially placed to scan at 32 uniformly fixed orientations, for 100 seconds each. This was repeated 3 times in total and each time, the 60 beacons were randomly repositioned (to increase the orientations variety). A procedure that practically took more than 8 hours to finish.

To convert the data into a translation function, the first step was to merge per BLE channel *(3 Ch)*, per beacon *(60 Bcn)* and per random re-positioning *(3 Pos) (i.e. 540 «Ch-Bcn-Pos» total combinations)*, the RSSI records from all 32 discrete phone orientations, into a single one. Their fusion offers at some degree an averaged (and less biased) RSSI, which is a) closer to the average of all possible (continuous) orientations and b) away from outliers. After that, each «Ch-Bcn-Pos» case from a phone was compared with the same «Ch-Bcn-Pos» case from another phone. This revealed indeed some underlying RSSI differences between phones, based on which, the development of a translation function could become possible.

More specifically, it was noticed that the biggest difference existed between the Sony Xperia Z2 and the LG G3 S and so, this pair became the subject of the translation model. Their differences are presented in the following 2 groups of plots ***(GP1, GP2)*** (Figure 3.1 & Figure 3.2), where ***GP1*** shows the differences from phone A to phone B, and ***GP2*** shows the differences from phone B to phone A. In each group of plots, the first two scatter plots ***(pA, pB)*** show the same RSSI-pairs distribution. Their difference is that the second scatter plot ***(pB)*** includes the information of the specific channel. Last, the third scatter plot ***(pC)*** shows the differences from phone A to phone B. As expected, between ***GP1*** and ***GP2***, the first 2 scatter plots are basically reversed. However, this is not the case of ***pC*** too, due to the calculation of their differences.

After plotting the RSSI pairs, the first thing that became evident was the great importance of the channel (i.e. the frequency) when examining the difference in the signal receptions of the two phones. For example, in ***pC*** of ***GP1***, it is clear that the signals coming on channel 39 are generally received better by LG G3 S (when compared to the Sony Xperia Z2). Contrariwise, the Sony generally receives better the other 2 channels.

The last part of the first research question had to do with the selection of a proper statistical model, based on which, the translations would be performed. For that, different functions were used to check how well they describe the RSSI relationships and the best performance was proved to be a simple linear regression (least-squares of vertical offsets [6]). This performance check was done in terms of complexity and coefficient of determination *(R^2)*, which is the proportion of the variance in the response variable that is explained by its (linear) relationship with the explanatory variable [7]. The produced functions are presented in Figure 3.1 & Figure 3.2. In both cases, the ***pB*** plots include 3 functions (1 per channel), whereas in the ***pA*** plots, the channel is considered unknown, and so, a (single) less descriptive translation function has been produced.

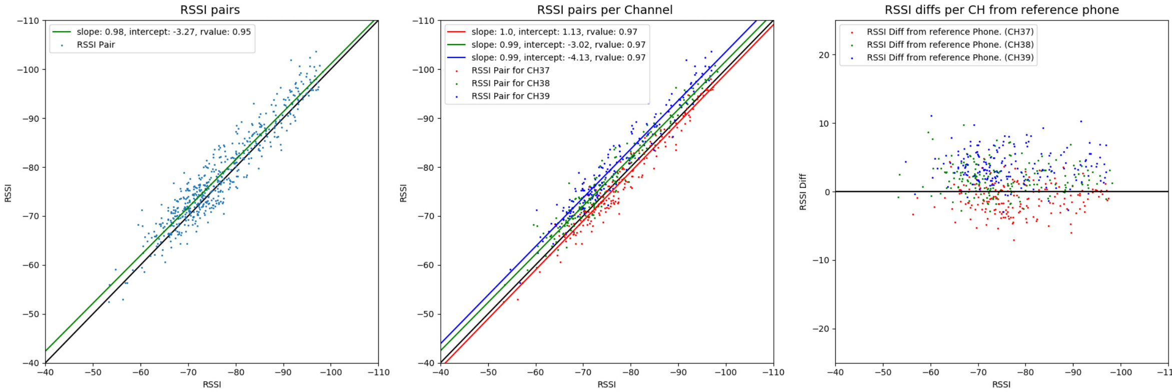

Figure 3.1: RSSI Difference from Sony Xperia Z2 to LG G3 S [GP1: Training Phase]

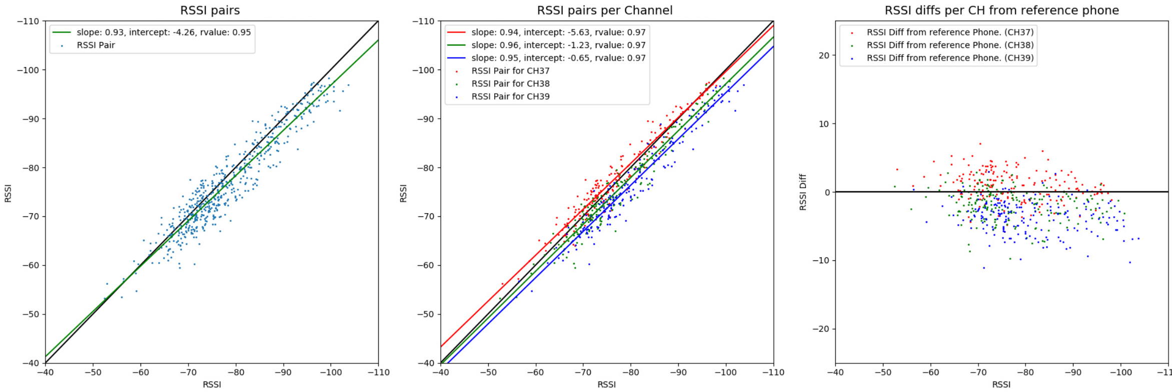

Figure 3.2: RSSI Difference from LG G3 S to Sony Xperia Z2 [GP2: Training Phase]

As already mentioned, between ***GP1*** and ***GP2***, the data of ***pA*** and ***pB*** are reversed. However, their corresponding functions are not also reversed (all of them have positive slopes). For that to happen, instead of a «least-squares of vertical offsets» approach, a «least-squares of perpendicular offsets» approach should have been followed (also known as Deming regression [8]). However, this model was outperformed during the evaluation phase by the currently chosen one, and so, it was rejected.

So far, only the pair having the most significant RSSI difference was discussed. However, it would be meaningful to also provide an example of phones with small RSSI differences and such example is the one showed in Figure 3.3, between the LG G3 S and the Samsung Galaxy S7. It is clear that the calculated functions (even at channel level) are highly coinciding with the y=x slope, which means that their receptions are very similar.

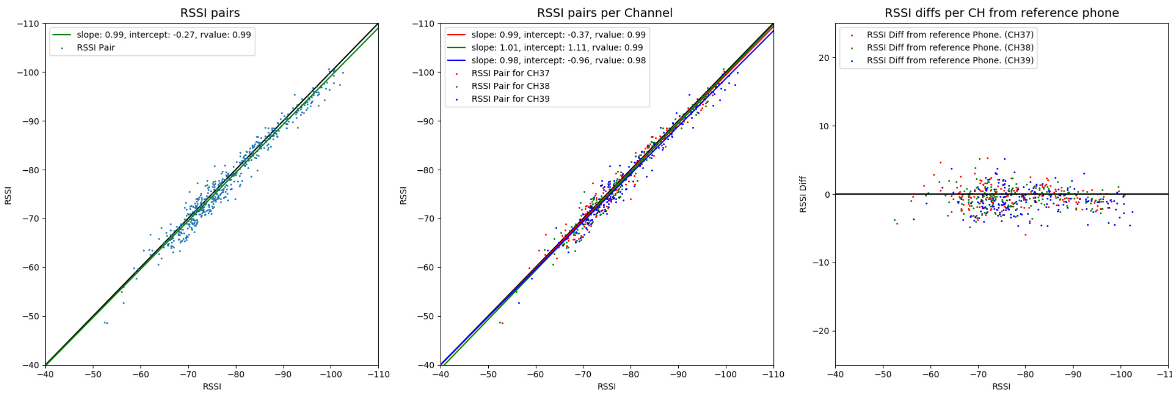

Figure 3.3: RSSI Difference from LG G3 S to Samsung Galaxy S7 [GP3: Training Phase]

The next step was to see whether a similar RSSI difference between the Sony Xperia Z2 and the LG G3 S (to the one found in Figure 3.1), could also be found in other case scenarios. If so, then, by utilizing the produced translation function and the readings from only 1 phone, it would in theory be possible to predict, with some accuracy, the RSSI readings of the other phone. Ultimately, if this accuracy was higher than the accuracy we would have if no translation was done, then the 3rd research question would have an affirmative response.

To test that, two new RSSI samplings were performed in a big indoor environment using the two selected phones. In both cases, 60 beacons were randomly scattered throughout the area (at random orientations), with the 3 Channel-References being additionally placed in the center. In the 1st case, each phone was consecutively used at 3 specific positions to record RSSI data under 6 different orientations. This smaller amount of orientations (compared to the 32 ones that were considered for the model training) was used to check whether these differences of interest could even be found during a less analytic sampling. Going one step further with this check, during the second sampling, an even more abstracted (and possibly error-prone) scenario was considered. Instead of placing the phones at fixed positions, a user, holding each device on one hand, walked for 1 minute a specific path. This process was repeated 10 times and each time, the 2 phones were swapping hands. In the 1st case, the 6 different orientations were combined in the same way as in the processing of the training data. However, in the 2nd case, there was no merging to be done, as the user introduced a serious distance and orientation averaging by walking around during a single path. On the evaluation data above, the same linear regression analysis (as the one used to train the translation functions) was applied and the results are presented below (Figure 3.4 & Figure 3.5).

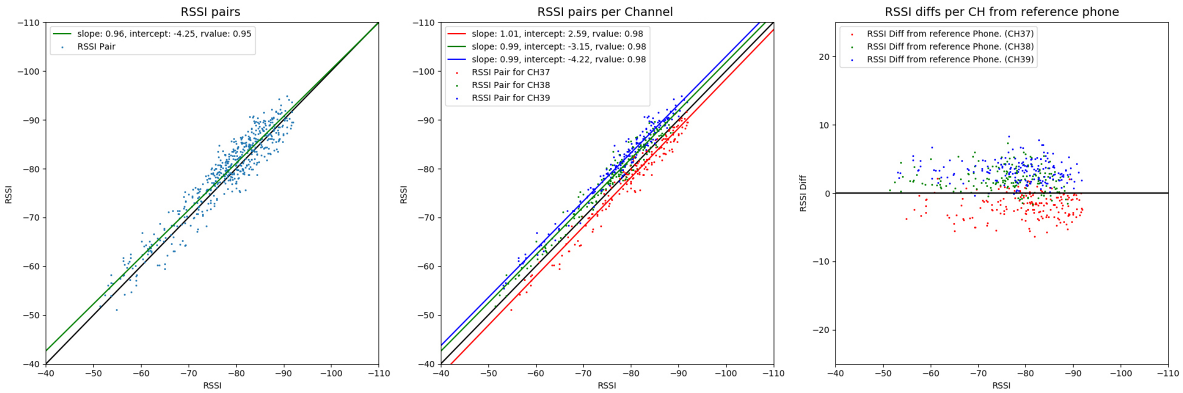

Figure 3.4: RSSI Difference from Sony Xperia Z2 to LG G3 S [GP4: 1st Evaluation]

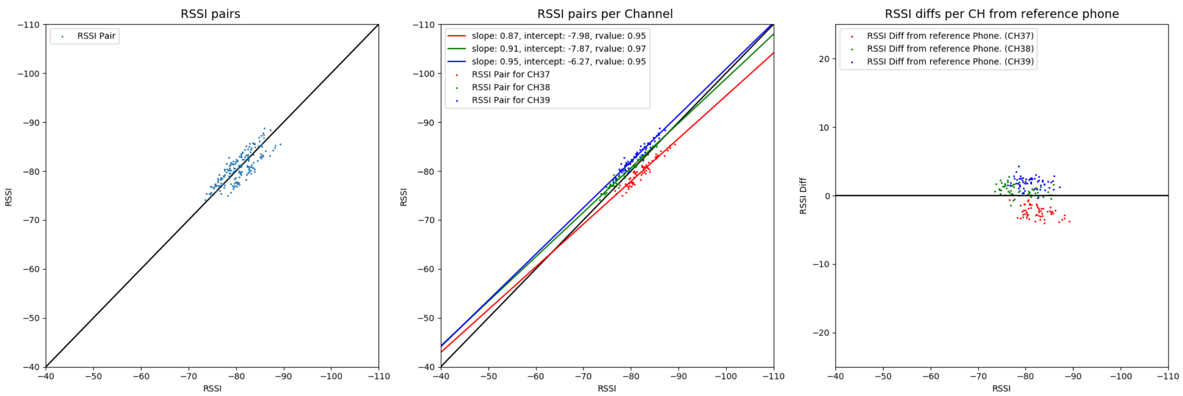

Figure 3.5: RSSI Difference from Sony Xperia Z2 to LG G3 S [GP5: 2nd Evaluation]

The translation functions shown in these plots are very similar to the ones shown in ***GP1*** and ***GP2***, which is quite promising for the trained model. Nevertheless, the ***pA*** plot of ***GP5***, where the 3 channels are not considered separately (but rather as a single set of measurements), does not include a function. That, because a minimum coefficient of determination of 0.95 (which has been generally set as a threshold) was not achieved.

The objective of the last research question was the selection of a proper measure to quantify the performance of the produced model. Regarding that, 3 measures of performance became the leading candidates. Namely, the *Mean squared error (MSE)*, the *Mean Absolute Error (MAE)* and the *Root Mean Squared error (RMSE)* (Table 3.1), each offering some advantages over the other. For example, the *RMSE* has generally the benefit of penalizing more the larger errors (residuals). However, in this case, such weighting wouldn't introduce better interpretation to the evaluation. Instead, a measure that could weight (penalize) more the errors at strong signals (having high dBm), would probably be the ideal case. That, due to the logarithmic nature of the signal propagation and thus, the importance of not "missing" the strong signals (especially if Trilateration is used for the localization algorithm). Such specialized measure, however, was not found (and probably hasn't even been developed yet) and thus the *MAE* was eventually selected, to take advantage of its easy and straightforward interpretation [9].

Table 3.1: Measures of performance

Mean absolute error	$\text{MAE}=\frac{1}{n}\sum_{t=1}^{n}\lvert e_t\rvert$
Mean squared error	$\text{MSE}=\frac{1}{n}\sum_{t=1}^{n}e_t^2$
Root mean squared error	$\text{RMSE}=\sqrt{\frac{1}{n}\sum_{t=1}^{n}e_t^2}$

The following table presents for each evaluation test *(Eval_A, Eval_B)*, the Mean Absolute Errors of the (estimated) RSSI values of the LG G3 S, which were predicted by using the (real) RSSI values of the Sony Xperia Z2 and the translation models that were presented in (***pA*** and ***pB*** of ***GP1***). The model of ***pA*** is evaluated within the "All Channels Combined" column, while the 3 models (1 for each channel) of ***pB*** are evaluated within the remaining (green) columns. After translating the signals, the resulting RSSI accuracy ("With Translation" entries) is compared with the accuracy that we would have if no translation was done and so, the RSSI values of the Sony Xperia Z2 were used ("Without Translation" entries).

		All Channels Combined	**Ch 37**	**Ch 38**	**Ch 39**	
Eval_A	*Without Translation*	2.52	2.05	2.25	3.26	
	With Translation	2.2	1.54	1.34	1.3	***Channel Average: 1.4***
Eval_B	*Without Translation*	1.73	2.38	0.9	1.92	
	With Translation	1.82	1.46	1.19	1.77	***Channel Average: 1.47***

Table 3.2: Evaluation results for predicting the signal of the LG G3 S (Mean absolute errors of RSSI)

The above results show that in the first evaluation case, the translation model always predicted RSSI values that were more accurate, when comparing to not having translated the signal at all. More specifically, without knowing the channels, the model improved the signal by 0.32dBm. However, as soon as the channels were known, the average RSSI improvement became 1.12dBm (2.52dBm - 1.4dBm). In the second evaluation case, where the distance to each beacon and the orientations were significantly mixed (and averaged), the model only improved the signal in the case where the channels were known (by 0.26dBm). In the other case, the accuracy of signal was decreased by 0.09dBm (1.73dBm - 1.82dBm).

4

Conclusions & Recommendations

In the beginning of this project, several research questions were formulated. Their response was considered fundamental for eventually reaching the final goal. Namely, to assess the degree to which, a model could be developed that would support correcting BLE RSSI translations between 2 phones, leading to more accurate comparisons between them. Each research question has been already extensively discussed and several conclusions have been reached.

First of all, due to different hardware & firmware specifications, every phone receives the 3 BLE channel signals, differently. This difference, however, is also far from stable. It varies unpredictably among all possible orientation combinations, between the phone and the beacon. In addition to that, another parameter that significantly affects the BLE RSSI differences across 2 phones, is the distance between the device and the beacon. As such, to acquire a representative estimation, many diversified orientations and distances have to be considered. A process that is definitely time-demanding. However, it was also discovered that for as long as there was no movement within the system, the received signal on each distinct channel was quite stable. Therefore, just a few samples per channel were sufficient each time (which helped accelerating the whole data gathering process), to receive an accurate signal representation. It is also worth mentioning, that although Androids do not state the BLE channel on which a signal has been received, yet, by taking advantage of a specific Android's scanning behavior, it became possible to identify the channels of incoming BLE signals. Another observation that has been made, is that, by increasing the number of transmitting beacons (and consequently, the number of BLE signals reaching the phones), each distinct beacon was detected less frequently. A phenomenon having probably to do with the available computational sources.

All the aforementioned research led eventually to the development of the aimed translation functions. Several prediction models were tested to see how well they describe the RSSI relationships, but the best performance (in terms of complexity and coefficient of determination) was proved to be a linear regression that minimized the vertical offsets. During the data processing, the LG G3 and the Sony Xperia Z2 were observed having the biggest RSSI difference. For that reason, the evaluation of these functions was done by comparing the RSSI accuracy between a) the predicted BLE signal of the LG G3 (based on the signal from the Sony Xperia Z2) and b) the signal from the Sony Xperia Z2. This comparison was used because during Indoor Positioning applications, the phones are considered to have the same BLE reception properties. Something that has been disproved as the evaluation results showed that, almost every time, the translation model introduced some improvement to the accuracy of the signal and especially when the channels were known.

So far, no BLE signal translations have been practically implemented in Indoor Positioning applications. However, since it is indeed possible to develop a model that could support beneficial translations between BLE signals from different phones, it might be a good idea for companies offering such services, to do it. Especially, as soon as Google identifies that by stating also the BLE channel, it would have direct application to Indoor Positioning.

5

Internship Reflection

To be able to undertake this project, someone should already be familiar with several aspects related to the topic. To begin with, a strong background in programming was necessary, as more than 3000 lines of code had to be written in Python and Java, not only for the data analysis, but also, for the logger application running on the phone. Furthermore, some adequate knowledge in statistics was required too, in order to be able to process the signals. Finally, some prior understanding of how BLE localization works, was mandatory, in order to be able to properly direct the whole research. Indeed, all these prerequisites had already been sufficiently acquired from the Geomatics MSc, before starting the internship.

During this research project, lots of different literature materials were considered and most of them, on statistical approaches. Topics related to properties of various data distributions, distribution comparisons, statistical measures of performance, prediction models, etc, were among the most reviewed ones. Several references have been made to them and these can be found within the bibliography.

At this point, the great importance of doing this project in Crownstone should be highlighted. Not only because, individually, it would be nearly impossible for me to have access to all the crucial equipment that was used, but also, for the effective assistance and advisements from all supervisors, when required. Additionally, I even had the opportunity to be very close and gain valuable insight into other interesting AI projects. For all these, and generally for the great environment that I found there, I couldn't be more grateful.

I finish this internship, having enriched my knowledge in several fields, which are quite important to me. First of all, I got to deepen my knowledge and understanding in many statistical aspects, such as the ones mentioned above. However, the most important knowledge I gained, is that indeed, it is possible to achieve BLE signal corrections based on trained models. Something that companies offering Indoor Positioning services, should invest in further researching. Especially, if Smartphones in the future start providing information regarding the identity of the BLE channel.

Before closing, it is worth mentioning that, more MSc Geomatics students are expected to become in the future interested in a topic close to this one, or Indoor Localization in general. If so, then in case the student is also familiar with Java, it would be useful to him to follow the «Smart Phone Sensing» course (IN4254) taught by Z. Zamalloa at the EEMCS faculty. Having followed this course, I acknowledge that I did find it quite useful during this internship. Therefore, I believe that it should (by default) be listed in our elective courses.

Bibliography

[1] M. Nima, "Indoor Localization and Fingerprinting across Multiple Smartphones," National Institute of Advanced Technologies of Brittany, Rotterdam, Tech. Rep., 2016. [Online]. Available: https://www.ensta-bretagne.fr/jaulin/rapport2016_mehdi.pdf.

[2] T. v. Dijk, "Indoor localization using BLE - Using Bluetooth Low Energy for room-level localization," Delft University of Technology, Rotterdam, Tech. Rep., 2016. [Online]. Available: https://crownstone.rocks/attachments/internships/tomvandijk.pdf.

[3] W. Bulten, "Human SLAM - Simultaneous Localisation and Configuration (SLAC) of indoor Wireless Sensor Networks and theirs users," PhD thesis, Radboud University, 2015. [Online]. Available: http://theses.ubn.ru.nl/handle/123456789/264.

[4] G. Haagmans, "A statistical analysis on the system performance of a Bluetooth Low Energy indoor positioning system in a 3D environment," PhD thesis, Delft University of Technology, Delft, 2017. [Online]. Available: https://repository.tudelft.nl/islandora/object/uuid%3Af7646f3b-9a99-4207-97a6-20ef83c6a571.

[5] T. Val and A. Thaljaoui, "BLE Localization using RSSI Measurements and iRingLA (regular paper) BLE Localization using RSSI Measurements and iRingLA," *Industrial Technology (ICIT), 2015 IEEE International Conference on*, 2016. DOI: 10.1109/ICIT.2015.7125418.

[6] E. W. Weisstein, "Least Squares Fitting," Tech. Rep. [Online]. Available: http://mathworld.wolfram.com/LeastSquaresFitting.html.

[7] O. Renaud and M. P. Victoria-Feser, "A robust coefficient of determination for regression," *Journal of Statistical Planning and Inference*, vol. 140, no. 7, pp. 1852–1862, 2010. DOI: 10.1016/j.jspi.2010.01.008.

[8] T. Therneau and M. Clinic, *Deming, Theil-Sen, and Passing-Bablock Regression*, 2018.

[9] C. J. Willmott and K. Matsuura, "Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average model performance," *Climate Research*, vol. 30, pp. 79–82, 2005. DOI: 10.3354/cr030079.