MSc Human Media Interaction
Master Thesis

Developing a smart tile to improve train platform safety

Max Pijnappel

Supervisor: Juliet Haarman

Chair: Dennis Reidsma

External Supervisor: Mark Vlutters

April, 2024

Department of Human Media Interaction
Faculty of EEMCS,
University of Twente

UNIVERSITY OF TWENTE.

Contents

Abstract

Train platforms are becoming more crowded, raising safety concerns by ProRail. KITT Engineering has developed a smart tile that aims to spread people across the train platform using nudges. This tile has a screen and LEDs to give those nudges to the passengers and an embedded Doppler-radar to detect crowded situations. The tile functions fully autonomously using a battery and solar energy for power and LoRaWAN for communication. The purpose of this research is to investigate whether the power and radar are sufficient for the tile to perform its intended function. Furthermore, we will take a look at the noticeability of the tile, to see if a tile is the correct medium to convey the nudges.

The radar was tested at the Den Bosch train station to investigate how well the data could be used to detect activity by comparing it to data from people trackers placed by ProRail. In this comparison, the radar was shown to be capable of reliably detecting up to a distance of 1.25 metres.

In addition, the available power budget of the tile was calculated, showing that it is capable of operating for an entire year without intervention. Finally, observations of how many people looked at the tile when a nudge was given show that 33% of the people who walked past the tile visibly reacted to the tile. The conclusion of this research is that the tile is capable of its intended purpose. The next for the tile will be researching the nudges, seeing which nudges work and how effective they are.

Keywords: autonomous, tile, train, platform, low-power, solar, nudging, Doppler-radar, detection

Chapter 1

Introduction

During rush hour, the train platform becomes a crowded place. When a train arrives, all those passengers want to get on the train as quickly as possible. At the same time, there are people on the train who want to get off the train and get onto the platform. This means that many people will be in the same space and push each other around, which can lead to dangerous situations near the rails of the train [1]. To avoid these dangerous situations and improve the safety of the train platform, ProRail has started the SBIR project [2], to challenge companies to find a solution to improve the safety of the train platform.

KITT Engineering tried to solve this problem using smart tiles containing a Doppler-radar, LEDs, and an LCD screen. The screen and LEDs are in the tile to create nudges for passengers on the train platform. The Doppler-radar was placed in the tile to detect movement to determine when an area is crowded. Together, they should detect a crowded situation and nudge the passengers to spread more across the train platform.

To achieve these goals, we identify two areas of investigation. First, what nudges are most effective in impacting passenger behaviour? We investigated this question in a previous work [3]. Second, are the technical capabilities of the smart tile sufficient to perform these nudges over long periods of time? In this study, we investigate the second question.

The primary research question will be "To what extent can the smart tile be used to improve train platform safety?". We divide the research question into three subquestions:

- RQ1: How reliable is the Doppler radar in detecting activity?
- RQ2: How does the power system of the smart tile compare to its power usage?
- RQ3: Do people notice the smart tile on the train platform?

For the first research question, we will start with taking a close look at the data. This is done by seeing how the data reacts during experiments in a lab setting. The experiments will focus on distance, walking speed, and number of people. After having a better idea of how the radar data behaves, the tiles will be placed on Den Bosch. In Den Bosch, the tiles will collect data for two weeks. This data will be compared with people tracker data from ProRail, which will also be gathered from the same place at the same time. The data from ProRail will be considered as the ground truth, as this is the closest to knowing the actual situation. Since this data is verified by ProRail for their research on people flow. The radar data will be compared to the ProRail to see how reliable the radar is in detecting people.

The second research question consists of two parts. First, how much power is available to the system year-round. This will be calculated based on how much power the solar panels generate throughout the year and the capacity of the battery. The second part is the power consumption side of the tile. To calculate this, the power consumption of all components and states will be measured. At the end these two parts will be compared to know if the available power is sufficient for the tile to operate.

The final research question will focus on the tile as a communication medium. To test this, the tile will be placed in the actual location, being a train platform in Den Bosch. With passengers as they would be in the actual situation, like for example in a hurry to catch the train or on their phone. This will tell how well passengers will notice the tile while nudging. During this experiment, the researcher will display something on the tile while observing how many people show a visible reaction to the tile.

Since these three research questions are so fastly different, the second and third research will have its own section, and the first will have two sections. Each of these sections will have their own methods, results, and discussion. In the end there will be a general discussion and conclusion that focusses on the primary research question and the state at which the tile is.

Chapter 2

Previous work

Before looking at the technical aspect of the tile, I looked [3] at how best to change the behaviour of passengers and what the limitations of signaling on a train station are.

2.1 Limitations

ProRail has a rule book [4] to communicate with passengers on the train station. This includes specific colors to convey a message with an associated purpose. And which icons are being used or where signs have to be placed and how often. All signage on the train station has to follow these rules, to make communication as clear and consistent as possible in dutch train stations. This is why it is good practice to try and have the tiles adhere to the same rules.

In addition to these rules of ProRail, it is important to note that the tiles can also be seen by the train operators, since the tiles are located near the border of the platform. This means that the tiles cannot interfere with the job of the train operators. Since they receive signals from lights located near the train border. And a signal from a tile should not be mistaken for a signal from the lights from there signals.

When looking at the limitations it seems the blue and orange colors caused the least amount of conflict. The arrow is the only icon that can be used out of the rule book. This is due to the limited resolution of the LCD screen placed in the tile. Furthermore, the color blue could even help signal train operators where the train needs to stop. Since, the color blue is used to indicate to the train operator where they need to stop on the platform.

2.2 Behaviour change

We looked at the best method to make people perform the desired behavior. There are many methods to try and influence people's behavior. In the setting where the tile is going to be used, the desired behavior is not something that can be obliged. So something like a sign with rules that tells them not to stand near each other is too intrusive. And the rule is also something that is not always needed, only when is gets busy and dangerous situations can happen. The solution needs to be more towards guiding people away from certain locations.

There are a few methods to try and steer people to perform a desired behavior (nudging [5], coaxing [6], enticing, insisting and requiring [7]). The most non-intrusive way is called nudging [8] and seemed to be the best method. Some research has been done on the effect of nudging. One of those investigations placed green feet on the ground with a path to the nearest trash can, in an attempt to reduce street litter[9]. This decreased the amount of litter on the street. Another example is the fly sticker in the toilet to make the male bathroom less messy. People are not told to aim while peeing, but people aim for the fly, which has the same effect. But in this way, a person does not feel forced to perform the behavior, they do it of their own will.

On a train platform, there are many situations where a nudge could be used to improve a situation. Improvements not only to the safety of the platform but also to its efficiency or usability. But the place and the desired behavior can differ depending on the situation. This is why all tiles are capable of showing multiple signals. Not only are there multiple situations, the location of the situation when a nudge has to be given can also change over time. This is why a dynamic solution like the tile can work so well. The tiles can be placed all over the platform and they all capable of signaling for multiple situations. But they can also blend fairly well in the background, so they can be made noticeable to the passenger when needed.

2.3 Scenarios

I looked at some example scenarios of nudges and what people think they would mean. This was done with a survey showing scenarios with an image and some text to support it. There were two versions of the survey, one in which only an icon was shown on the tile, and the other in which there was also a supporting color. This was to test whether the colors were needed to help understand the intent of the nudge. From this study, it was clear that colors are not needed to understand the nudge. But this does not tell anything about the LEDs being needed to attract the attention of the passenger.

From the survey, we found a few scenarios with associated nudges that seemed to work. The most relevant was designed to improve platform safety by better distributing people. This was done by using arrows pointing to less crowded tiles. This scenario is the most related to solving the problem of reducing crowded situations and having the people spread more across the train platform. The intended goal of the survey was to understand what most people thought the signals were supposed to mean. At the moment it is unclear if the tile is capable of performing these nudges. This is what will be researched during this research.

Chapter 3

Hardware overview

The focus of this research is on some components of the smart tile. These components are the radar, power system, LCD screen, and LEDs. Other components are also in the tile, but they play a supporting role or are not used at all during this research.

The power system consists of a battery, solar panels, and circuits to manage those. Extra components are needed to ensure that the voltage of the solar panels is changed to the correct voltage for the battery and other components in the system. The power system plays an important role in the tile power budget, but the hardware of which cannot be controlled in the context of this research. The LCD and LEDs are for giving the nudges to the passengers, and the Doppler-radar is used to detect crowded situations. The other parts are sensors that can measure the temperature and humidity, GPS for time and location, and a LoRaWAN chip for the communication. The temperature sensor, humidity sensor, and GPS are not used during this research.

For this research there also was a datalogger developed to store the data generated by the radar. This is not part of the tile itself and will only be used during this research.

3.1 Radar - input

The tile has two types of radar, a Doppler radar [10] to detect moving objects and a continuous wave radar [11] to detect static objects. The Doppler radar (SMR-313) [10], detects moving objects by looking at a frequency shift. The principle behind the radar is that it sends a wave at a certain frequency. In this radar it is 24GHz and measures the returning wave. When the wave bounces back from a moving object, the transmitted frequency will shift. This is the Doppler effect; this effect changes depending on the speed of the object. The effect happens because the object compresses the waves, changing the wavelength and frequency. The radar module converts this frequency shift to an analogue value that the tile microcontroller can read and convert to a digital value. Because the type of wave is electromagnetic, it travels at the speed of light, which means that the radar can be measured many times per second. In the tile, this is done 1400 times every second. This Doppler radar is the radar that will be investigated during this research.

The second radar is called a frequency-modulated continuous wave radar (SMR-314) [11] and is used to detect static objects. It also transmits an electromagnetic wave around 24GHz, but it increases frequency over time in every measurement. Since it increases the frequency over time, there is a pattern in the signal, and this pattern can be measured by the radar. It can then also know how long it took for this signal to return and thus the distance. Currently, while the tile is being tested, this radar has not yet been implemented. And it cannot be used for testing.

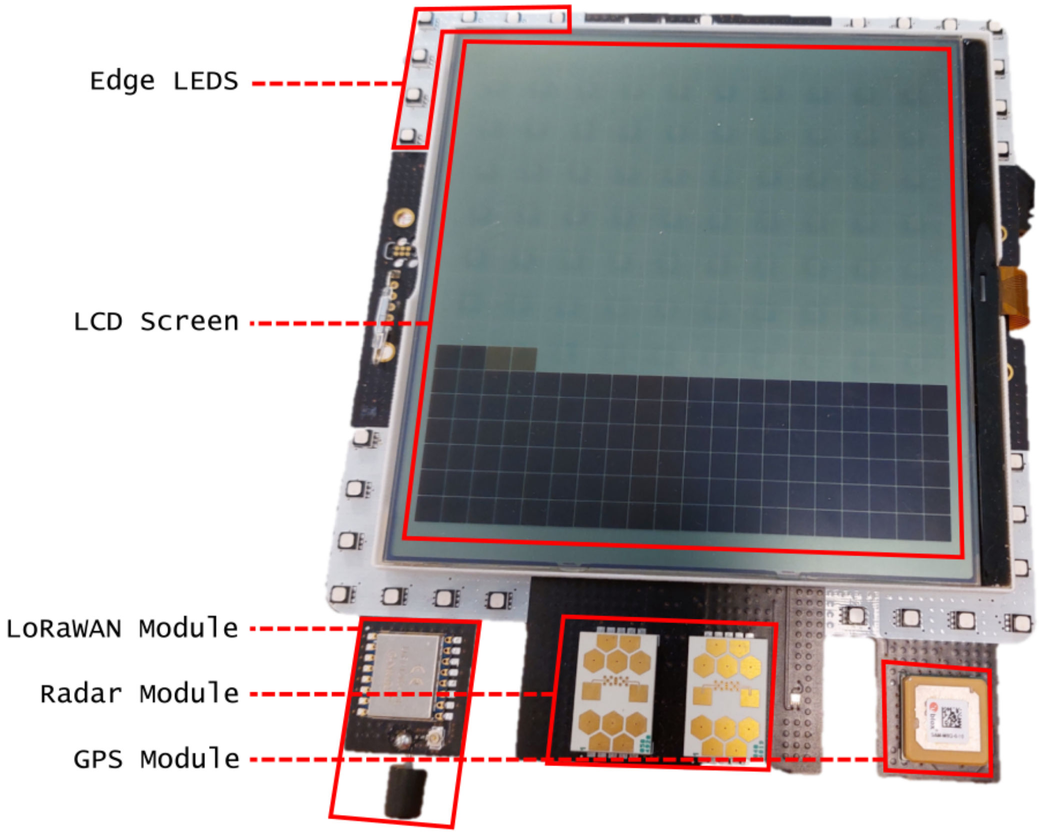

FIGURE 3.1: The hardware that is inside the smart tile

3.2 Nudges - output

The smart tile has an LCD screen, LED matrix, edge LED's that can present the nudge. They can be used to show icons and display colours that can be used to signal a certain behaviour. The LCD screen has 20 by 20 pixels which can be off (see-through) or on (black). Right behind the LCD screen is the LED matrix. This matrix has 10 by 10 LED's, resulting in one LED for every 4 pixels of the LCD screen. All LED's can display red, green, blue, or a combination of these three colors. There also is brightness control for the LED'S independently for each colour. Even with the LED matrix having a lower resolution than the LCD, it can still show basic icons. However, it can also be used as a backlight for the LCD screen by turning on either all LED's or only the parts where the LCD screen is see-through.

In addition to the LED matrix behind the LCD, there are LED's on the corners of the tile. These LED's can flash very brightly for a short time, for basic signals or to attract attention. They can also show red, green, blue, or any combination.

3.3 Communication

Communication with tiles is supported by LoRaWAN using the SX2612 chip[12]. This chip handles all communication with the gateway that is used to connect the tile to the Internet. The tile interacts with this module to exchange commands and sensor data.

The LoRa protocol [13] supports three classes A, B, and C. The tile can be in class A and class B mode. In class A, the tile sleeps for a set amount of time and wakes up to send a status message. After sending a message, it will listen for a short time to see if the gateway has any messages for it. After this it can go back to sleep mode. This mode is the most energy efficient, but also makes communicating slow. The communication speed will be determined by the duration of sleep. This can be altered remotely with a LoRa message.

The other class that the tile supports is class B; In this mode, the tile will not enter sleep mode. In this mode, every couple of minutes, the tile will receive a time beacon to be in sync with the gateway. This increases the power consumption of the tile, but decreases the delay when communicating with the tile. In the time beacon, there are also time slots for the next couple of minutes. During these time slots, the tile has to listen for possible new message for a few seconds.

Class C is not used in the tile, but only in the gateway, as it continues to listen to messages all the time. This consumes a lot of power and is not made for battery-operated devices. For example, the gateway that connects the tiles to the Internet. This enables communication to and from tiles from anywhere in the world. The software that manages distance communication for the smart tiles is called Chirpstack [14]. This gives insight into the states of the various devices connected to a website. It also enables communication over various protocols, so it can be automated with the use of programming. This is how communication among tiles can be orchestrated, so the correct nudge is shown in every tile. A application can listen to the status of the tiles, and can use those displays to show nudges on them. Since it can track how crowded each area is, and nudge people from the crowded to the non-crowded one.

3.4 Power components

In order to be independent from continuous external power, two main components are needed: a component to generate power and a component to store power. For power generation, three 1W solar panels are placed on top of the tile. The panels are protected with a see-through cover since they are embedded in the floor. The panels include the necessary components to regulate battery charging and to ensure that all voltages are correct. Although the panels have a combined peak output of 3W, the actual peak power is assumed to be 2W because of the inefficient angle to the sun and the fact that they will get dirty from people walking over them. This will be used in further calculations. Power storage requires a battery with sufficient capacity to buffer in case the solar panels are underperforming. We added two batteries in the tile with a combined capacity of 13600mAh. They provide this capacity in all weather conditions.

3.5 Other Measurements

In addition to the major parts that are the focus of this research, there are some extra sensors in the tile. These being an environmental sensor, voltage sensor and light level sensor. When the sensor gives an update over LoRa, it will also give the latest values of these sensors. First the environmental sense, this sensor has the ability to measure the temperature and air humidity in the tile. This is mostly to keep track of the tile, it can detect overheating or water damage. But the temperature sensor could also tell something about the ground temperature when placed in the shadow. This can be helpful in the winter months for detecting possible icy situations on the train platform.

The voltage sensor measures the voltage of the battery to keep track of the capacity. It is used to predict how long the tile is still able to run. But it is also used turn the tile into a sleep state when almost empty. It will only turn on when the battery is above a set voltage again. This is to decrease wear on the battery but also ensures the tile is not enter a on/empty loop while charging after being drained completely.

The last sensor present is the light level sensor, it will give a value in lux. It gives insight in how much sun the solar panels receive and thus charge. It could also be used to dynamically set the brightness of the tile. Because if there is less sun the LEDs can be less bright, which safes battery consumption.

3.6 Power states

The tile has a few states that it can enter. The most efficient state is the sleep state, in this state is turns of all peripherals except for the internal clock. This clock is used to wake up from the sleep state after x time. It will wake up from the state every few minutes to give an update over LoRa. This state is used to conserve power and extend the battery life of the tile. The next state is the idle state; in this state, most peripherals are turned on or can be turned on quickly. Things like the radar, LCD and LED's are turned off in this state by default. The other states are idle state with a combination of the radar, LCD LoRa or LED's turned on.

3.7 Data logger

For the purposes of data collection in the field a data logger will be embedded in the tile. This is mostly due to the limited bandwidth LoRa offers and can not be used to send the raw radar data. This data logger has to use as little power as possible to reduce the impact on the battery. To do this, the tile will store about a minute of data and then write it to the data logger. When the data logger is not used, it will be turned off to reduce power consumption. When the data logger is used, the data will be sent as fast as possible. This data will be checked to see of it was received correctly, and if all the data were sent over correctly, they will be written to the SD card. The data logger and tile will communicate with each other to make sure that everything is sent correctly. If all is done, the tile will turn of the data logger again until it needs to store data again.

To make this process as energy efficient as possible, a balance between power consumption and speed was found. Lower power consumption means that the data logger has to be on for a longer time. Vice versa, having a higher power consumption means that the data logger does not have to be on for as long. Since the important part is the average power consumption over the entire time, data is logged, not only when writing the data.

Chapter 4

Radar - Detection

A big part of nudging is knowing when to nudge; this could be done manually, but this is not a long-term solution. This is why the tile has radar modules inside. But at the moment there is no way of using the radars to detect how many people there are in the area around the tile. So we need to investigate the reliability of the radar to detect activity. To try and understand what is possible and how the radar data can be used, some controlled tests will be carried out. After these tests, a test will be performed on a train platform in Den Bosch to see how reliable the Doppler data is for different ranges, compared with a dataset provided by ProRail.

4.1 Controlled environment

Some controlled tests will be performed in a laboratory setting. First, we will look at the raw radar data and how it can be best interpreted. This makes it possible to analyse later, since this is not possible in its raw form. Now we can take a look at how to radar reacts to some situations. For this, some tests will be done where we will vary in walking distance from the radar, walking speed, number of people, and some walking patterns. In the end we should have gained insight in how the radar behaves, and what can be gotten from the data. But also how it can best be used to detect crowded situations.

4.1.1 Methods

In this section, we will discuss how we tested all the parameters that could influence the tile. The goal is to understand how the Doppler radar responds to the various parameters, and how to process and interpret the radar data.

Participants

The researcher participated in the experiments when only one person was required. If extra participants were required for the tests, the participants would be acquired from KITT Engineering, the company that wants to develop the tile. No personal data was registered or recorded during the experiments and the radar also does not collect data that could be connected to a participant.

Materials

To perform the experiments, a special smart tile, a micro controller and a laptop was needed. The main part was the special smart tile, which has the capability to send the raw data generated by the radar over a USB cable. This can be read by the laptop and store it so it can be analysed later. The programm that stores the radar data will also add timestamps with every sample. The microcontroller was needed for the two buttons placed on the ground. These buttons had to be clicked at the start and of the run. When the microcontroller detects them being clicked it sends a signal to the laptop. The laptop will store the button presses and the time. These button presses will be used to filter the relevant data. Between the two presses is the radar data where a person walked past it. This is the radar data which needs to be analysed. .

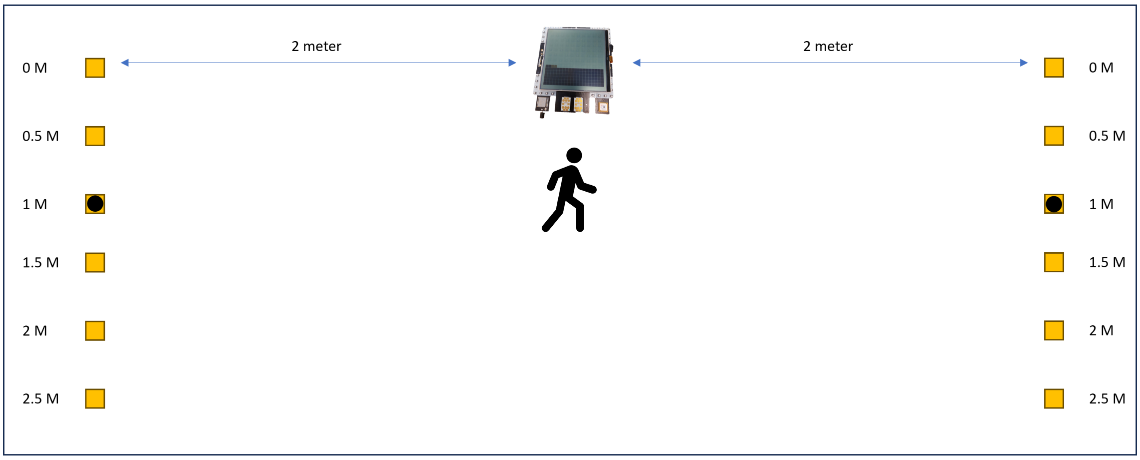

FIGURE 4.1: The experiment setup, the tile is at the top. The yellow squares indicate the various distances for the tests. The black dots are the buttons.

Procedure

Each variable discussed in the outcome measure was tested separately. Keeping all other factors as consistent as possible. Each test consisted of noting down the test and all the variables. The name of the test was changed in the programme that stored the data and the recording was started. At the beginning of each pass along the tile, a button was pressed. At the end of a pass the other buttons was pressed. During each test, the participant will take multiple passes in both directions. Per test, 8 passes will be made, 4 passes per direction. After the test, the data will be processed to see if the data was recorded correctly. The setup of the tests can be seen in figure 4.1.

Data analysis

The raw data that the Doppler radar generates is not in a form that can be used to gain insights in the data. Since the Doppler radar is a radar which works with frequency shifts. The raw data will be analysed using an FFT, this consists of 64 bins where each bin consists of 11Hz worth frequencies. Looking at all these bins independently would generate to much data. But to see how those bins influence the data, it has been plotted for a few tests. Figure 4.2 shows how the different frequencies behave during one of those tests. In this figure you can see changes in value when a person walks past the tile marked with blue bands. Only the lower frequencies are plotted; since the other higher frequencies show nothing, they can be seen in Appendix A. The y-axis shows the bin of the first 5 sums of the FFT, but to better compare the different frequencies, an offset was given depending on the frequency so they can be plotted together. Without the offset, its would be hard to differentiate between the different frequencies.

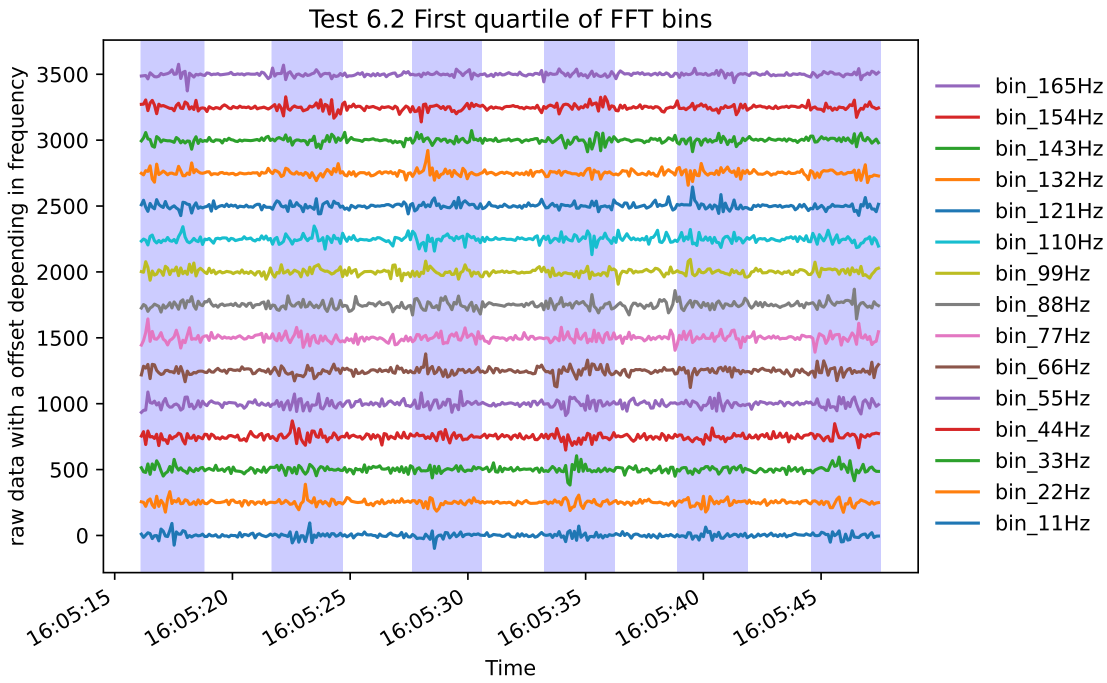

FIGURE 4.2: First quarter of FFT bins plotted overtime for normal walking speed and 1 meter distance. It shows that a specific frequency does not react stronger than others.

To better process the data, it is better to get one value for every FFT taken, a comparison between taking the highest bin value, power sum of the FFT, and sum of the first 5 bins is made. The same test was analysed as shown in Figure 4.2. The way the data look after they are processed with the three methods can be seen in Figure 4.3. In all three plots, there are clear peaks at the moment a person walks past the tile. Passes are indicated in light blue. The peaks seem most clear when summing the first 5 bins and less pronounced when taking the maximum bin value. When taking the power sum, the peaks are a bit longer, with less of a pause in between. To analyse the data, taking the sum of the first five bins seems to be the best. This will be used to analyse the rest of the experiments.

To gain insight into how the parameters behave on average, each experiment will be averaged. This is done by cutting the data into pieces and averaging the pieces. Each piece consists of one pass past the tile, and using interpolation all the passes will be made the same length. The result is the average curve per experiment or factor tested. The resulting curves will be plotted together per category, for example, range, to easily compare how the different ranges influence the data.

Outcome measures

The objective of this experiment is to see how passengers who walk on the train platform influence the radar. First we will test the range of the radar; the tested distances are 0 (over the tile), 0.5, 1, 1.5, 2 and 2.5 metres from the tile. These distances can be seen in the setup of the experiment in Figure 4.1.

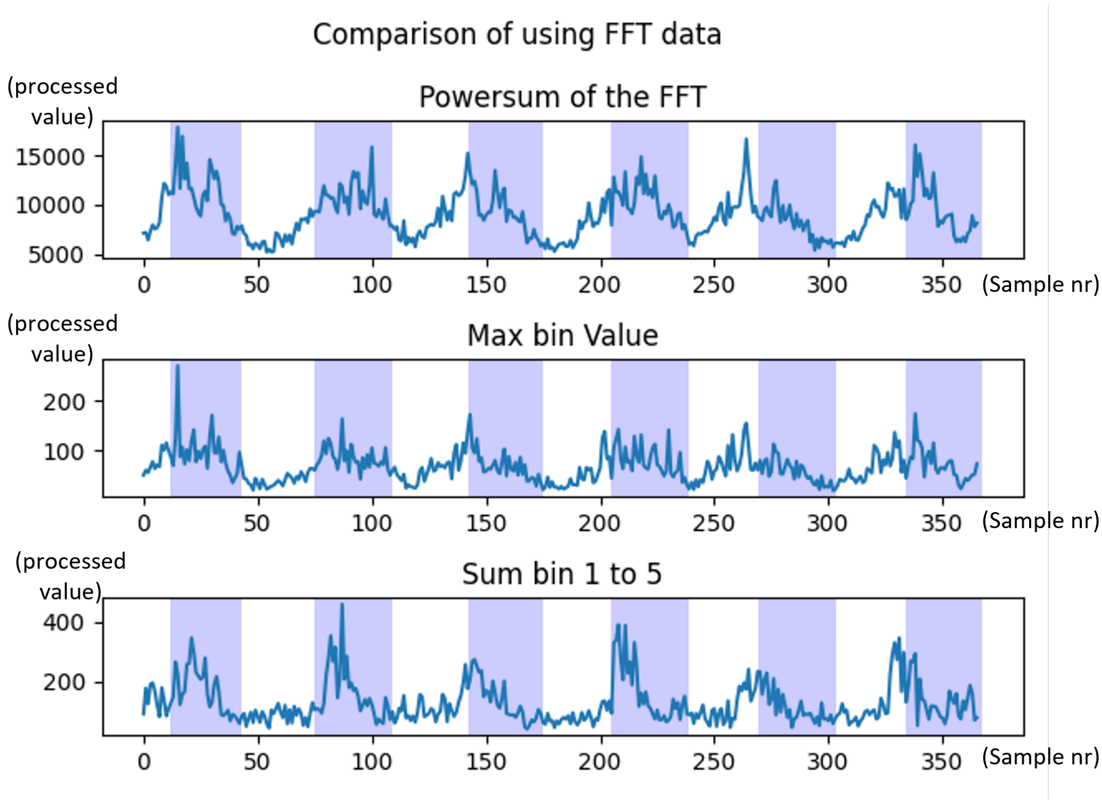

FIGURE 4.3: Comparison of processed data

After the radar ranges, the walking speed will be tested; we will test this with a distance of 1 metre from the tile and differentiated between very slow, below average, average, above average, and fast. No exact speed was taken since it was not possible to accurately measure or control this during the experiments. The tests were performed by one person to ensure that the various walking speeds were somewhat consistent.

These were the factors that are influenced by one person. However, a platform is usually quite busy, so the effect of the number of people will also be tested from 1 to 7. Each participant walked with an average walking speed and the centre of the group was 1 metre from the tile. People always walked as a cluster, since that best imitated how people tend to walk together.

Then the effect of walking patterns with two people are tested, since it can happen that two people not related to each other can pass the tile at the same time. The tested patterns are walking behind each other, besides each other, and crossing each other. The distances of the two participants will vary between 0, 1 and 2 metres from the tile.

In the end, each of the parameters listed above will be plotted in one graph. These will be used to visually inspect how each parameter affects the data. In the end, we will discuss how the parameters behave to see how the data could be used and analysed to detect activity on the train platform. As well as what kind of range could be expected from the Doppler radar in an ideal situation.

4.1.2 Results

The various average curves of the parameters have been plotted together and are shown here.

Effect of walking distance from the tile

The average curve of the tested ranges can be seen in Figure 4.4 A. It should be noted that the peak decreases as the participant walked further past the tile. The relation is not linear when you look at Figure 4.4.B. The difference in maximum value between the ranges decreases as the range increases.

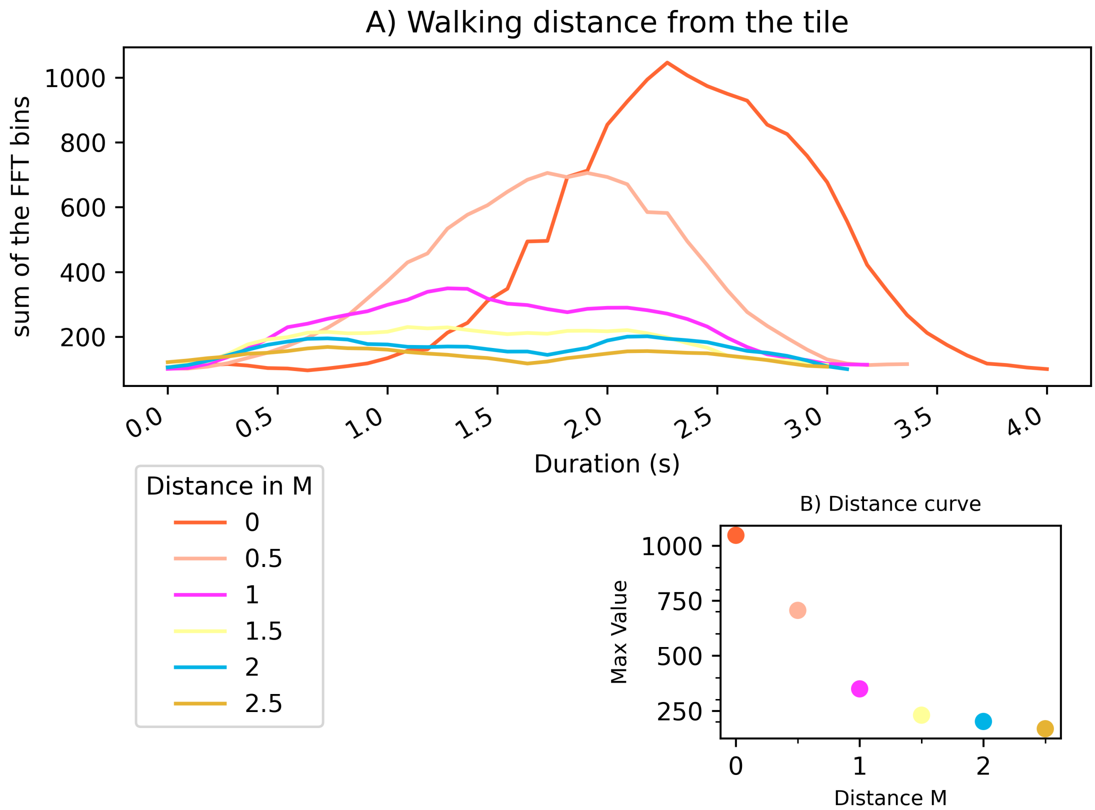

FIGURE 4.4: Relation FFT bin sum and distance from the tile

Effect of walking speed

When you look at the curves of the walking speeds tested shown in Figure 4.5, the only difference is the length of the peak. The faster the walking speed, the shorter the curve. Furthermore, there is a small dip in the peak at the lower walking speeds, creating two maxima in the curve.

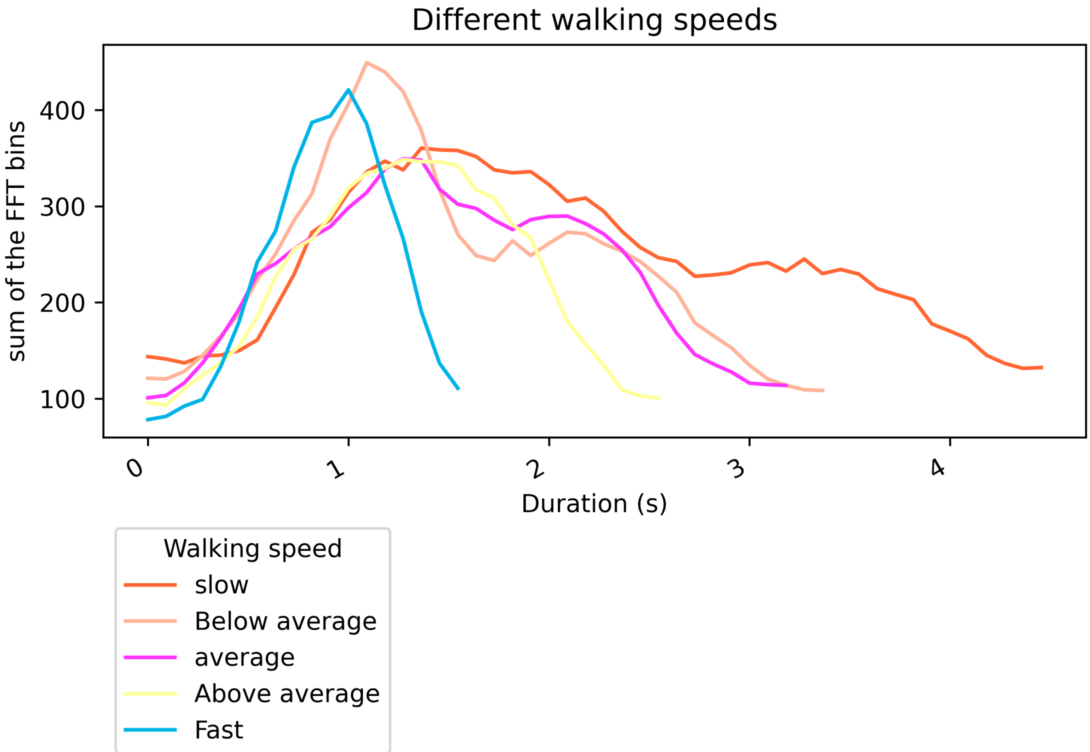

FIGURE 4.5: Relation between FFT sum and walking speed

Effect of number of people

The results of the tests, which vary in the number of people, can be seen in figure 4.6. The height of the peak increases when there are more participants in the test. The duration of the peak also increases as the number of people increases. Both these curve characteristics are influenced by the previous two parameters.

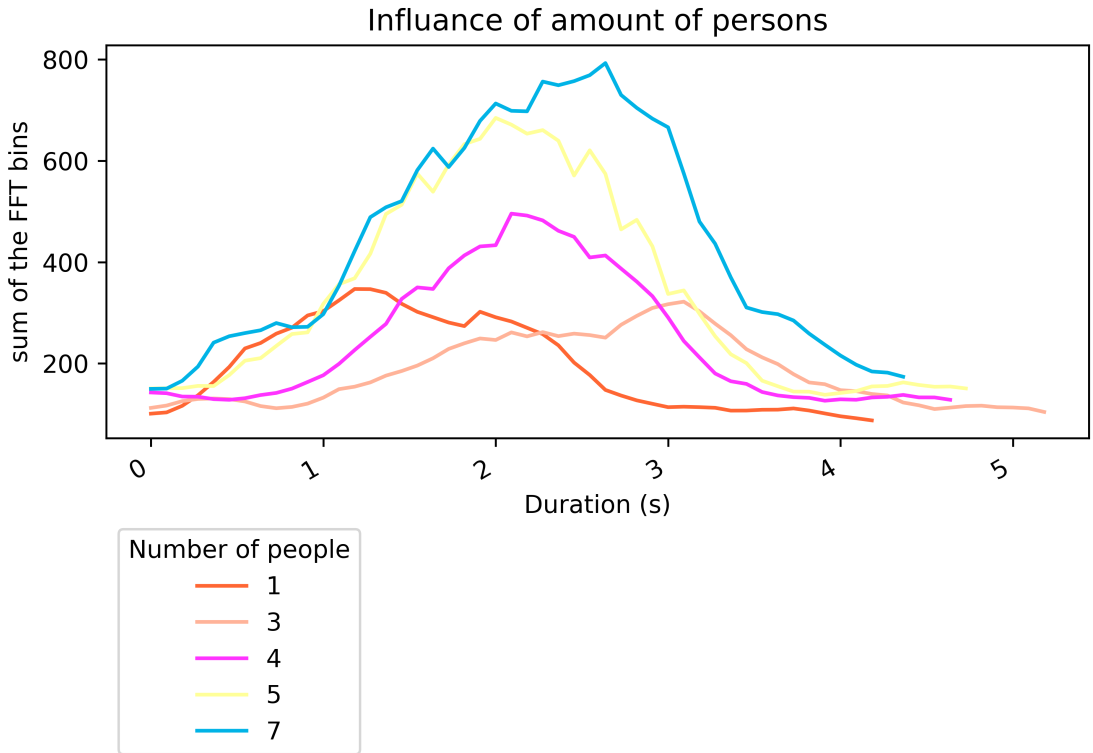

FIGURE 4.6: Relation FFT and number of people

Effect of Walking patterns - 2 persons

In figure 4.7 it can be seen that the height of the peak is correlated with the distance the person closest to the tile has. There is no difference in curve when walking in the same direction or walking in the opposite direction when looking at the same distances.

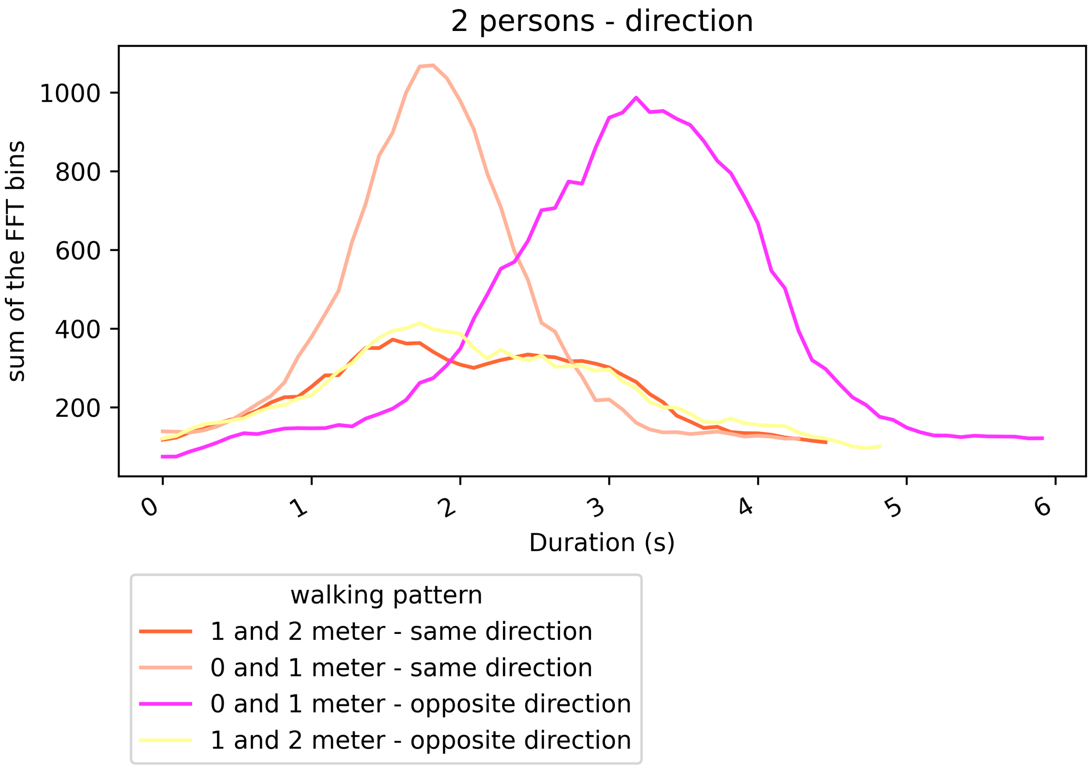

Figure 4.7: Relation FFT and walking direction

The second part of the walking pattern tests shows the effect of people that walk behind each other. The results are shown in figure 4.8. For the test in which one person walked one metre from the tile and the other one right over the tile, two distinct peaks can be seen. In the test where both participants walked one metre from the tile, a small bump can be seen. But with one person 1 metre from the tile and the other person 2 metres from the tile. It mostly looks like a very long curve with no distinct second peak.

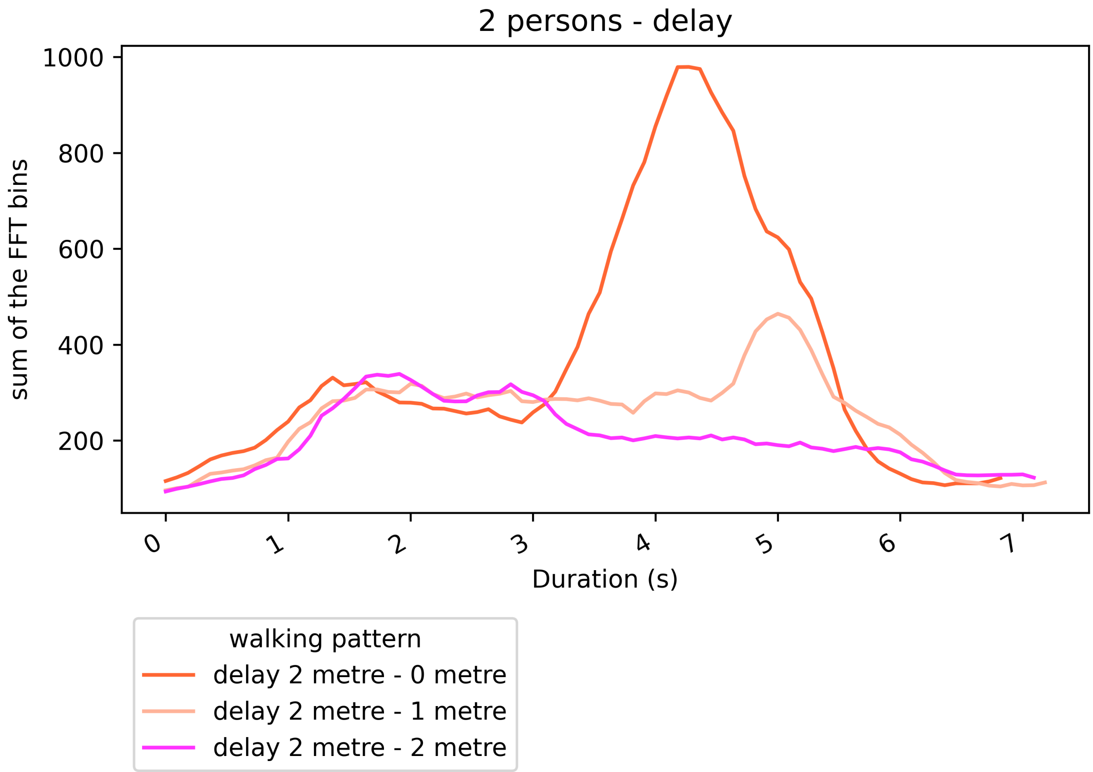

FIGURE 4.8: Relation FFT and 2 persons walking with a delay

4.1.3 Discussion

The aim of the experiments was to see how the radar behaves in controlled situations. Knowing how factors like distance and speed influence the measurements taken by the radar helps to see the limitations. The most important conclusion from the experiments is that the radar can detect passing objects, in this case persons. Thus it is at least possible to detect if something is within range. After a distance of 1.5 metres the peak generated by a passing person is decreasing. The effective range is somewhere between 1 and 1.5 metres in the lab situation. When analysing the data for the test on Den Bosch it might be best to start with a detection range of 1 metre since this is the largest range that still showed identifiable peaks.

The speed with which a person is passing the tile mostly seems to impact how narrow or wide the peak will be. This makes sense since the slower a person walks, the longer a person is within the detection range. The radar had no trouble detecting all walking speeds. That is, speed should not be a limiting factor when using the radar. It also did not seem to influence the height of the peak, meaning that speed and range should not conflict when analysing the data.

The goal is to know how busy the platform is, a good indicator would be to know how many persons walked past the tile at the same time. Detecting the difference between one and three people would better describe the situation. Looking at the peaks, it seems to get higher as more people walk past the tile at the same time. However, this could be caused by a larger group forcing a person closer to the tile. This is probably the case since the centre of the group is on the 1-metre line. Therefore, the outer people should be closer to the tile. The peaks also seem to get longer with more people. But since the walking pattern was a group, people also walked behind each other. This would make the group longer in the range of the radar. Both variables already conflict with the first two findings, meaning that it would be hard to differentiate between a group or one person walking close by and slowly. This should not affect the reliability with which the tile can detect activity, only the degree of activity, which will not be studied during this research.

Looking at the difference of two people walking next to each other or starting opposite of each other in the same direction. The curves match the curve of the range experiment and the curve has the same height as the range of the person walking closest to the tile. The person farther away seems to disappear when looking at the radar data. And since the radar is not aware of direction, the radar does not distinguish the person when they walk opposite each other at the same distance. The farthest person will still disappear behind the closer person.

When looking at walking patterns where two persons walk behind each other, both persons are visible in the data. Even when both people are 2 metres away from the tile, which is at the outer edge of the radar range. The first peak is always clearly separate from the second peak, which depending on the range will be higher. That is, it could be possible to detect people walking behind each other with the radar on the tile.

All these tests were done with people moving past the tile; it would have been interesting to see how the data behaves when a person is standing still beside the tile. This shows if they really disappear, or that small movements are still visible depending on the distance from the tile. It should be possible to detect small movements using a single channel continuous-wave Doppler radar, as observed by Withword [15]. This showed that small movements can be detected up to a distance of 2 metres. Having some experiments in which a person stood close to the tile and performed movements like scrolling in the phone. It would have shown if the tile is capable of detecting standing people that moved.

In addition to the continuous-wave Doppler radar, the tile also has a frequency-modulated continuous-wave radar. However, it was not implemented and could not be used for this research. The frequency-modulated continuous-wave radar can detect non-moving objects. This module would not help to detect multiple people, but should help to detect if people are standing still near the tile.

Doppler radars can have different frequencies; Mafredi [16] compared the accuracy of Doppler in detecting people in the forest for Doppler radars with 1GHz and 435MHz. Here, 435MHz proved to work better for detecting people in the forest. It could be that the 24GHz that the current radar module uses is not optimal for the situation. Since a higher frequency means that it could detect more precisely, this might not be needed to detect crowded situations. A different frequency might have a longer range or a lower power consumption. Or it could be better at detecting crowds than the frequency currently used. Perhaps in the next version of the smart tile, more in-depth analyses could be made of available radar modules to see which module is best suited for this use case.

Other radar systems use multiple frequencies to detect multiple object sizes. Wang [17] was able to see the respiration patterns of multiple people at the same time. Since it uses multiple receivers, the set-up knows where an object is in space, since the signal gets received from multiple angles. The multiple frequencies enable it to detect movement on different size scales. The tile does not have this capability, as it only has one receiver. The tile could better detect crowds if it can detect objects in the space surrounding it. This would enable it to better track 2 people as they walk on either side of the tile. Now they would be seen as one person. This would result in higher power usage, since it would need multiple radar modules, each of which uses 3mA. However, having, for example, four more directional radar modules should increase the range and also enable it to detect in four quadrants.

4.2 Den Bosch

The goal of this experiment is to see how well the radar performs in a real situation. The reference will be the data provided by ProRail and will be considered as the ground truth. The data will be analysed using time windows varying in duration. In addition to a variation in window duration, the threshold value for when it will be considered as activity will also vary. At the end a closer look at the range will be taken with the ideal window duration and threshold values.

4.2.1 Methods

Here we discuss what is needed to perform the experiment in Den Bosch. This includes the materials, what the goal is, and how the data will be analysed.

Participants

Participants in the train platform experiment were the passengers who walk on the train platform. There was no data collected that could be correlated with a person. There was nothing expected of the participants; they have to behave as they would always behave. For this experiment, ethical approval was needed since it involved experiments with people. To obtain ethical approval, an information poster was put up, this poster informed passengers that experiments are performed on that platform.

Materials

To perform the experiment, two things are needed, first, a special version of the tile. The tile should have a working radar module and a data logger to capture the radar data. It should also be aware of time, so timestamps can be added to each sample. This tile should also be on the train platform. The chosen location should be within the range of the ProRail people trackers so that the data can be compared to that data. In addition to the tile and its data, there should also be data on the actual situation during data collection. This will be provided by ProRail, who installed people trackers on the station. People are tracked with an ID that cannot be traced back to a person. The ID is only connected to an object as long as the person is within range of the trackers. ProRail data also needs to have timestamps so it can be compared to the radar data. In addition to these materials, the location of the tiles should be known in comparison to the ProRail data set.

Outcome measures

The goal of the experiment is to see which threshold value and window size can best be used to detect activity around the tile. This will be done by seeing which threshold and time window work most accurately without sacrificing the reactivity of the tile. Since a bigger time window means it can react slower to activity. If, for example, the time window is set to 10 seconds, the tile would need 10 seconds of radar data before making a decision. In the end, the optimal threshold value and time window will be decided. A closer look at possible ranges with the ideal values will also be taken. The decision about these values is explained in the data analysis.

Procedure

The tiles will be placed on the train platform and will be remotely told to start the radar and log the data on the SD card. They will also be periodically updated with the actual time. After two weeks, the SD card will be removed from the tile so that the data can be analysed.

There are three locations available where a tile can be placed. At all three locations, a tile was placed. Tile 10 is located near an elevator, Tile 11 is located near the stairs, and Tile 12 is located near the escalator. Tile 11 did not store any data on the SD card, so it did not provide any data that could be analysed. The tile near the elevator is not very busy but is a place where people can wait for the train. Tile 12 is placed where almost everyone who enters or leaves the platform will walk past. This means that there are not many people standing there and waiting there. Tile 10 will be called "tile near the elevator" and tile 12 is "tile near the entrance" from here on out.

Data processing

The raw data of the tile consists of 10 samples per second, with each sample containing 128 measurements. These samples will be converted to an FFT that consists of 64 bins of 11 Hz. To get one value that can be analysed, bins 1 to 5 will be summed together. This is what was also used for the experiments in the lab.

The data provided by ProRail consist of the x and y coordinates of the detected persons on the train platform, with a sample rate of 10 times per second. To be able to compare the data to the tile, the data needs to be transformed, so you have a number of people in an area around the tile. Since it also has 10 samples per second, just like the radar data, the sample rate does not need to change.

First, the location of the tile is found within the data provided by ProRail. This was done with the help of ProRail and the example code they provided. This example code plotted the coordinates of the detected people on a map of the train platform. On this map the tiles are visible, and thus the location of the tiles in the ProRail dataset could be found.

Second, it could be that the timestamp of the radar data and the ProRail data have an offset. This can happen because the timestamp for the tile has to be sent over LoRa. There can be a time between the gateway getting its current time and the tile receiving that same time. This is checked by plotting a short sample of the tile data and pro-rail data and looking for peaks that align with each other. The chosen sample can be seen in Figure 4.12, here the offset is -19 seconds for the tile near the entrance. This offset will be used for both tiles when comparing with the ProRail dataset. The tile near the elevator did show less correlation, so seeing the offset was more difficult and that is why it also has the -19 second offset.

Data analysis

Now the data is aligned it can be compared to each other. Deciding when the tile should detect something as activity will be based on its value being below or above a threshold. This threshold will be based on a baseline value, which is a safe maximum value of the noise. To get this value, a data sample of 15 minutes will be taken where according to the ProRail data no people were present. This sample will also be visually checked for anomalies. From this sample, the third-quartile value will be taken as the baseline on which the threshold will be based. This is done to ensure that lower peaks caused by, for example, a bird or plastic bad do not influence the lower threshold value.

The threshold will be a factor times this baseline value, varying from 1 to 10. To see how well the tile performs compared to the ProRail data, which will be considered as the ground truth. Each time window will be compared to see if one of them is an activity. The ProRail data will have activity if more than one person is in the chosen area size. The tile will be considered to contain activity if the value of that window is above the time window. From this comparison, a confusion matrix will be made, which represents the occurrences of true positive, true negative, false positive, and false negative. All confusion matrices will be collected and plotted as an ROC curve. In the ROC curve, the specificity or true positive rate will be compared to the false positive rate or fall-out. Specificity can be seen as the degree to which the tile correctly detected a person. A high value means that it rarely detects a person when there is none. This is in relation to the frequency with which the tile missed a person. The fallout is about how often the tile falsely detected a person and how often the tile correctly detected that there was no person.

The chosen range is 1000mm or 1m since this range still showed large peaks when doing the lab experiments. The threshold value will be chosen on the basis of the so-called elbow. This is the point where the fallout increases dramatically. Right before the fallout value becomes very high is typically the ideal value. To choose the ideal window size, the trend of the chosen threshold value will be checked. The ideal value will be where the trend flattens, and a longer time window does not give much better results to ensure the system can react as fast as possible.

The data will be fitted to a training data set and checked on a test data set. This is to make sure that the still performs well if the values are not chosen for that exact dataset. Furthermore, the chosen values will also be checked with multiple ranges, to see until what range the radar gives reliable results.

4.2.2 Results

Baseline value

From Figures Figure 4.9 and 4.10 it can be seen that the sample chosen to calculate the baseline value is a sample with nothing going on according to ProRail. When comparing regular data, as seen in Figure 4.11 and Figure 4.12 with the baseline, the value is much higher when there is a person present.

ProRail data was set to a range of 3 meters to make sure that nothing is in the area according to the ground truth. The samples are not in a window size; turning the data in windows will reduce the noise somewhat. However, the value will be automatically calculated based on the third quartile of the sample.

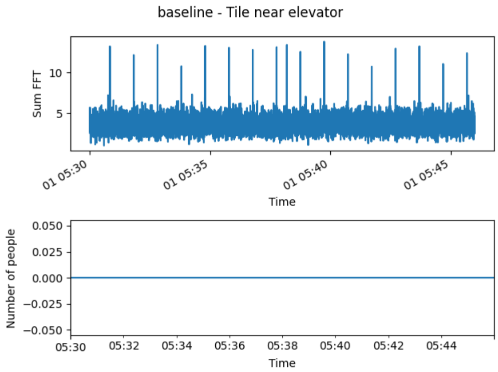

FIGURE 4.9: Baseline sample of the tile near the elevator. Below is the amount of people according to the ProRail data. values of the radar are low but noisy.

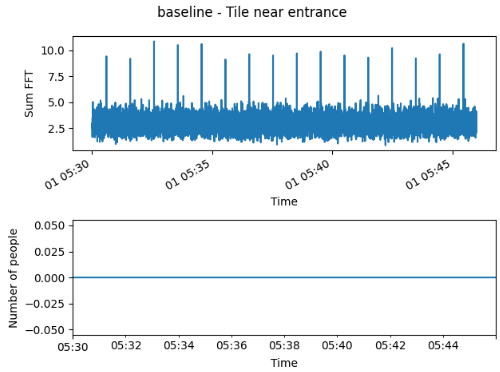

FIGURE 4.10: Baseline sample of the tile near the entrance. Below is the amount of people according to the ProRail data. values of the radar are low but noisy.

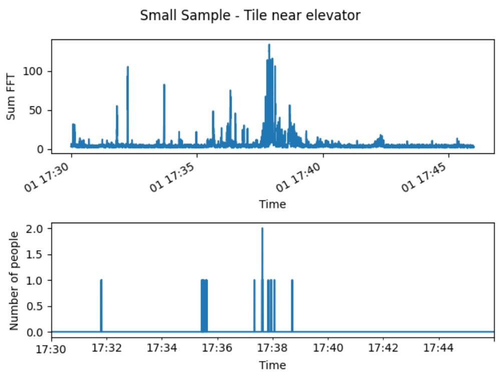

FIGURE 4.11: Small sample of the tile near the elevator. This is how the data looks when people are within detection range.

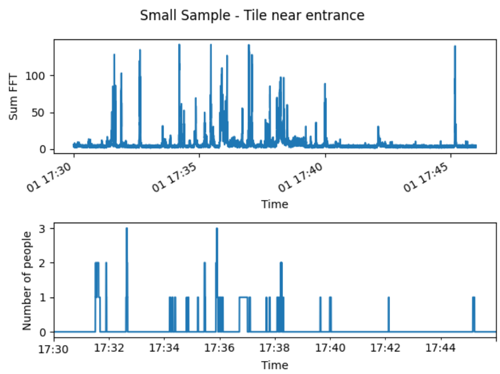

FIGURE 4.12: Small sample of the tile near the entrance. This is how the data looks when people are within detection range.

Threshold and time window

In Figure 4.13 and Figure 4.14 the ROC curves of the tile near the elevator can be seen, showing that it performs worse than the tile near the entrance, as seen in Figures 4.15 and 4.16. Looking at these ROC curves, the elbow is at threshold 2 for both tiles. A threshold value of 1 performs much worse than all the other threshold values in terms of fallout for both tiles. As for the time window duration, after 3 seconds the performance does not get much better. Both tiles have these same two values that appear to be ideal.

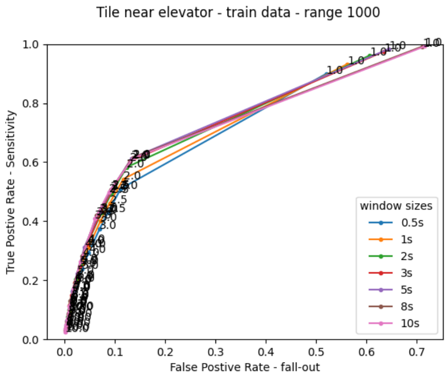

FIGURE 4.13: ROC curves tile 10 for window sizes and thresholds. The values in the plot represent the threshold factor.

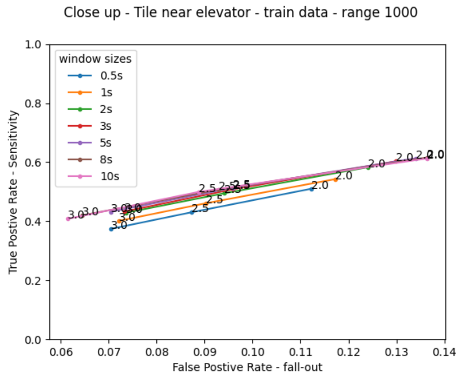

FIGURE 4.14: ROC curves tile 10 for window sizes and thresholds - close - up. The values in the plot represent the threshold factor.

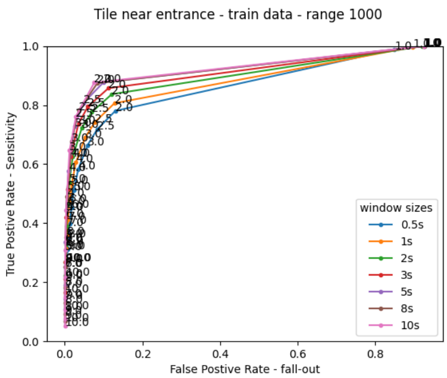

FIGURE 4.15: ROC curves tile 12 for window sizes and thresholds. The values in the plot represent the threshold factor.

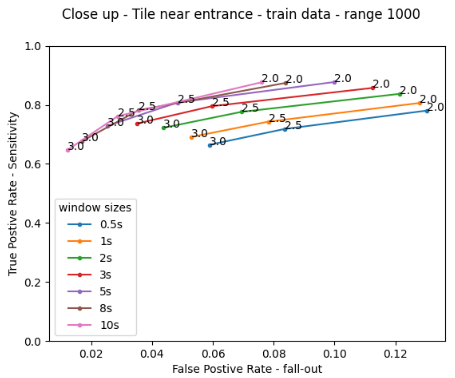

FIGURE 4.16: ROC curves tile 12 for window sizes and thresholds - close - up. The values in the plot represent the threshold factor.

Range and test data

Now that the values have been chosen, the performance can be tested for different ranges. A range of 500 means that the square of 1 by 1 metre is used as the detection range. See Figure 4.17 for clarification. The tile is in the centre of the area. This means that the distance from the tile to the edge of the area is 0.5M or 500mm. The range is given in millimetres. This is done on the training data set, as seen in Figure 4.18. But also in the test data set, as seen in Figure 4.19.

Using the elbow method, the range for the tile near the entrance seems to be around 1375mm. The tile near the elevator seems to have a range of 500mm. The curves of the two tiles also follow a different path. The tile near the elevator starts out curved and straits out. The tile near the entrance starts somewhat straight and becomes curved.

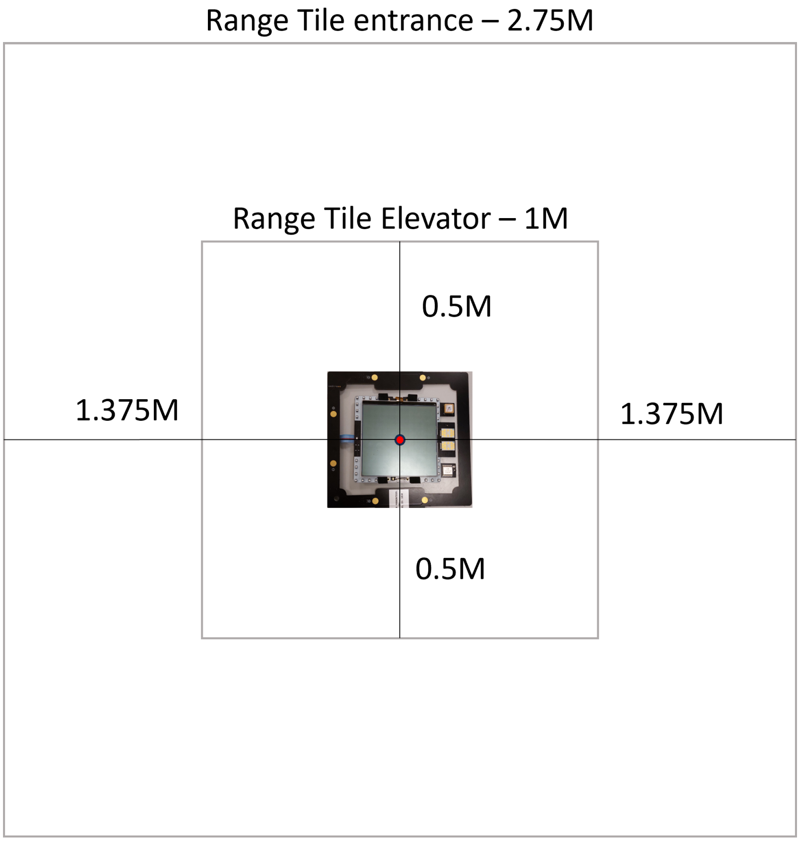

FIGURE 4.17: Illustration displaying the tiles range

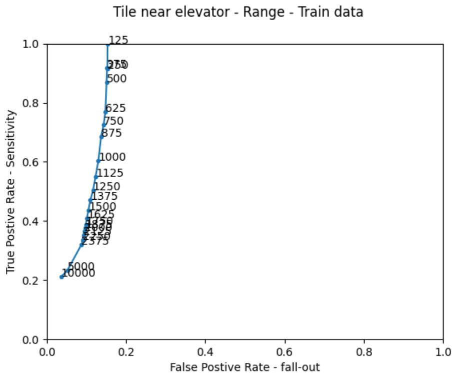

FIGURE 4.18: ROC curve tile 10 for different ranges on the train dataset

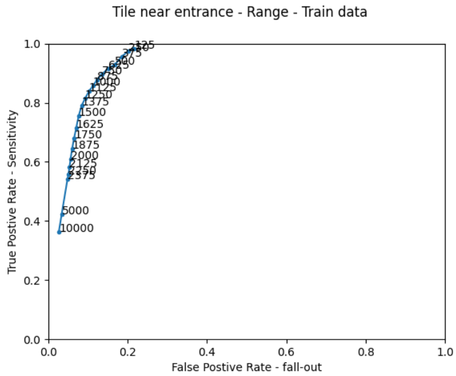

FIGURE 4.19: ROC curve tile 12 for different ranges on the train dataset

4.2.3 Discussion

The tests in Den Bosch had as their goal to see if the radar could be used in the real situation to detect activity. Looking at the results, the tile near the entrance should be able to. It can reliably tell if there is activity within a range of 1375mm every 3 seconds. The tile near the elevator did not perform as well and could only do it within a range of 500mm. The range of the tile near the entrance was the same as the range that showed good results in lab tests. This range is for both directions, meaning that the effective range is 2.5 and 1 meter. However, to monitor the entire train platform, a range of 2.5 meters is already small. A range of 1 meter would not suffice to have an accurate idea of the situation on the train platform. Unless the whole platform is full of these smart tiles.

There can be many reasons for the poor performance of the tile near the elevator. Due to the location, people tend to stand still near that tile. Since the continuous wave radar cannot detect objects that do not move, it also means that it will not see people standing still. It could be that the tile detects movement of people standing still that stand very close to the tile. And it can detect moving arms as activity. The tile near the elevator performs better for very small ranges with a higher specificity value and a lower fall-out value. The other tile is near the entrance, which means that people don't stand there waiting for the train. They only walk past the tile, which the Doppler radar can detect better since it is a large moving object. The elevator located near the poor-performing tile could also trigger the radar and be detected as activity, which could also interfere with the results.

With only two data points, it is hard to pinpoint if the tile near the elevator performs worse than normal or if the tile near the entrance performs better than normal. Both locations are also not the locations where the tile would be placed mainly. Most tiles would be placed near the platform border, but it was not possible to replace a tile there during the investigation. This is due to safety concerns when performing maintenance near the tracks. Having data from multiple tiles would have given a better insight into the performance of the radar. Or you can see if the data from multiple tiles could be combined to track how people move across the station. For future research, this could be beneficial and give a stronger conclusion to the performance of the tiles.

The ProRail data set contained the number of people that are near the tile every 10 seconds. Perhaps there is also a correlation between the value of the radar and the number of people near it. When quickly looking through the data, it did seem to correlate somewhat. For a future study, a closer look at this could be taken to see how reliable this is as a metric. If the relations are somewhat good, it could give better insight into how busy the train platform is.

When looking at other research on the detection of activity using a Doppler radar, many of them focus on using machine learning to detect multiple types of activity. Here, Chuma [18] used four neural networks to detect activity and reached high accuracy after some iterations. Alnaeb [19] was able to detect three types of activities after training a machine learning algorithm called SVM. Hanifi [20] was able detect states like lying down, heartbeat, and respiration rate using machine learning and the radar.

Since there is a lot of data since the tile collected data for 2 weeks, it could be possible to train a neural network to get more accurate results. It could be possible to detect how many people are near the tile. Or tell us something about the distance from the tile. It would be even better to have data from more than two tiles. Especially some tiles placed near the train platform, since this is the most interesting location for nudging. And this is also where people often stand still to wait for the train. The train itself can also influence the readings from the radar; if it does, it would need to be filtered out. But it could also be used to know where the train stops; knowing this over time could help in predicting where the train stops. This can be used to indicate where the train will stop and the doors will be, before the train has arrived at the platform.

Chapter 5

Power - Action

An important part of a system that must work autonomously is power management. The energy in the smart tile is generated using solar power, which at certain times of the year can be very limited. To bridge the times with no sun like during the night or winter, battery power is used, which is charged using the solar panels. First a calculation has to be made to see how much power is available on average. Second, the power consumption of all the components and states has to be measured. Finally, a calculation will be made to see how many nudges could be given during a normal day.

5.1 Methods

5.1.1 Materials

There are no materials needed to perform the power usage tests; calculating the available power budget will be done using a tool. To perform the power consumption tests, two things are needed; first, a smart tile that can be controlled to turn on or off certain components. Second, a method to measure the power consumption of the tile very accurately. This is needed since the tile is a low-power device that in certain states uses less than 1 mA. The Nordic Power Profiler Kit II is capable of doing this while also providing the power and the ability to store the measurements. During the tests, the power profiler voltage will be set to 3.7V, since this is the average voltage the battery will have. The power profiler can measure the current between 200nA and 1A; the resolution varies depending on how much current is drawn but is between 100nA and 1mA. All of this data will be stored and processed later to gain insight into the power usage.

During the LED tests for brightness and colour, the amount of lux will also be measured. This will be done using the mobilux USB - digital Luxmeter. This is to see how bright the perceived brightness is for the different colours and brightness levels. The lux sensor is placed directly on the LCD, and when the tile is placed on the train platform, there will also be a see-through lid placed on top. This will influence the actual brightness as seen by the passengers at the train station. However, this influence is the same for all tests. But it was not possible to measure the power consumption while the lid was placed on the tile.

5.1.2 Procedure

During each test, a specific component will be turned on and the current will be measured. The following components will be tested: (a) radar, (b) LCD screen, (c) communication, and (d) LED's. The radar will be tested with and without logging enabled, and communication will be tested in LoRa classes A and B. The LED's will be tested more thoroughly on brightness values, colours, and blink patterns. Finally, the idle state and the sleep state will be tested.

LED tests

The LED's have the most variables to test, as they can vary in colour, brightness, blinking pattern, and how many LED's are on. For all tests, an arrow will be used as an image since this is what will most likely be used to improve safety on the train platform. In all tests, the duration will be 24 seconds. During the brightness and blink pattern, the colour of the LED's will be green, since this is also the colour that will be used when testing the nudges. The brightness during the tests with colours and blinking patterns will be fixed to 6. This is the highest value to which the brightness can be set when giving commands over LoRa. In all tests, when the backlight is used, the corner LED's are also used.

5.1.3 Calculating available budget

The tile has to function autonomously using solar power and battery. A calculation has to be made to know how much power the tile can use on average per day. The goal is to ensure that the tile battery does not go below 40% for the entire year. This will increase the total battery lifetime. To calculate this value, a tool developed by the European Commission is used. The photovoltaic geographical information system [21] calculates how much power can be used on average based on the peak power of the solar panels, location, and battery capacity. You fill in a power consumption, and it will tell how many days the battery will be empty.

The tool will be used to test a few scenarios to see if the tile can run for the given period of time. A number of variables will be considered as constant; the first is that there will be two battery packs with each 6800mAh for a total of 13600mAh. The nominal voltage is 3.0V, meaning the capacity in Wh is 13.6*3.0 = 40.8 Wh. The second is the location; the tile is placed at the Den Bosch train station of which the coordinates are 51 ° 41'13.2"N 5 ° 17'34.8"E with an elevation of 11m.

The discharge cutoff point is set to 40%, as the battery will degrade less if it is not charged below this level. The slope at which the solar panel is placed will be 0 degrees, since the solar panel is placed flat in the tile. In an ideal situation, each of the three solar panels can capture 1W. Together, they should be able to generate 3W. However, the tiles are covered with a see-through lid that stops some light. And they are placed outside, and people walk over them. This means that dirt will slowly accumulate on top of the solar panel. Therefore, a total solar generation of 2W will be taken. There will be a variety in how much power the system can take on average and if it will be enough to always power the tile.For these calculations, the voltage of the system is set to 3V. The consumption per day will be (X mA * 3 * 24) / 1000. The steps of the average power consumption will be performed per 1 mA from 1 to 20.

5.1.4 How to verify best utilisation of the power budget

With how large the battery is and how much power the solar panels can generate, the estimated all-time available power budget is set to 15 mA. If this is the average power consumption all year round, the tiles should not be in a situation without power. The 15 mA will be used to calculate how many nudges can be done. In the end, a calculation will be done to see how many nudges can be done when considering all the measured components and their influence on power usage in a given daily scenario.

5.2 Results

For the tests with the LED's, the current measurement will be smoothed out with a moving average of 1000 samples. Since the readings of the LED's are very noisy and appear as a wide band. This smoothing will not be done for measurements with low power consumption. For these measurements, the short peaks are important and will be smoothed out too much if a moving average is used.

5.2.1 Idle states

The Tile can be in a few states, the most efficient being the sleep state. In this state, the tile turns off all its peripherals; this is done to reduce power usage as much as possible. The next state is the tile that gives a status update over LoRa; it will periodically go from the sleep state to an idle state and give a status update. It will also listen for a short period so that it can receive messages over LoRa. This LoRa mode is called class A and allows the tile to sleep between messages. The last state is the tile being in LoRa class B, during this state the tile cannot go to sleep. It will remain idle and listen periodically to messages; this mode can react faster to messages, but uses much more power. The amount of power these modes use over a 20-second window can be seen in figure 5.1. The amount of power these states use on average in a window of one minute can be seen in Table 5.1.

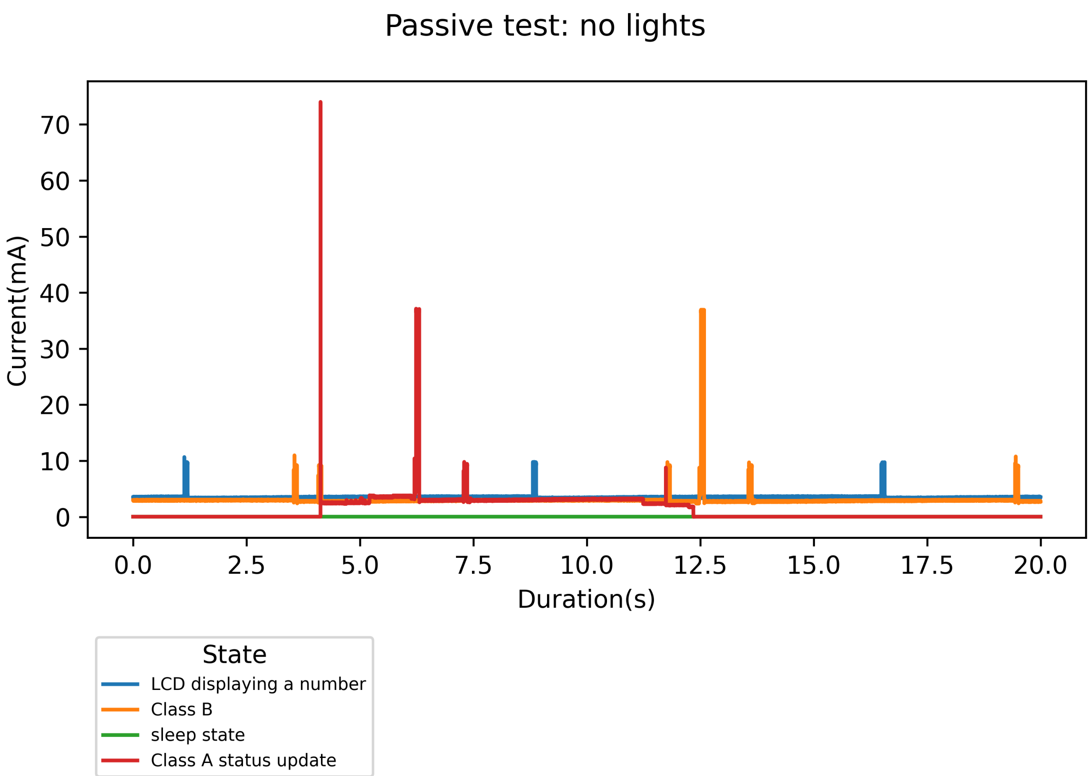

Figure 5.1: Power consumption of the various idle states or LCD turned on. The consumption is low, with some small peaks caused by LoRaWAN

When the tile is in an idle state, the extra effect of the LCD or Lora class B is negligible. The small difference is hard to see in the figure due to the short high peaks.

5.2.2 Radar

To detect when an area is busy, the radar must be turned on. The data is displayed in figure 5.2. For the purpose of this research, the tile has the option to log that radar data. As can be seen from a one-minute window in table 5.1 the difference in power usage is only 0.3mA. Compared to idle state, the radar uses an additional 3.2 mA of current, which is almost double. The tile also has to be in idle state,since the radar cannot be used in the sleep state.

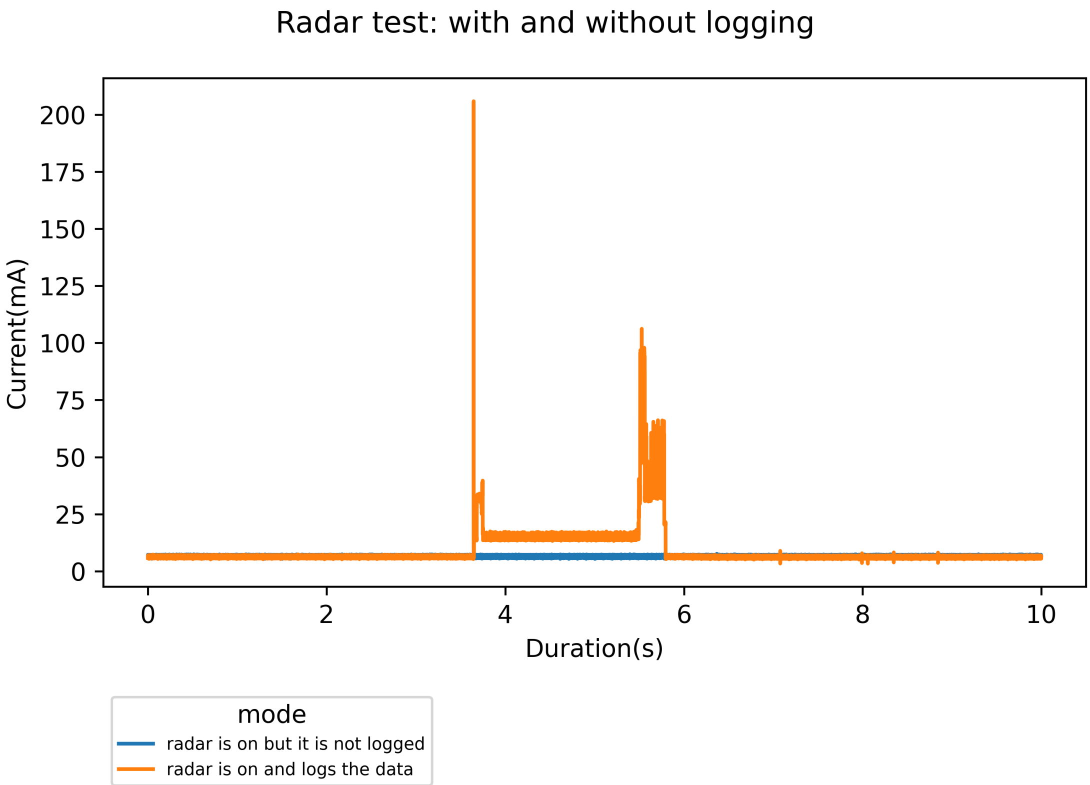

FIGURE 5.2: Power consumption of the tile with only the radar turned on. The elevated part in the orange graph shows the data being logged.

5.2.3 LCD

There are two tests done with the LCD. For the first one, the commando is used, which can also turn on the backlight. But the backlight is set to zero or off. This can be seen in figure 5.3, where the power consumption seems to be around 50 mA. The difference between the lowest backlight setting and the one without backlight is only a few mA. For the second test, a special command is used to display a binary number on the tile. This is shown in figure 5.1, the power consumption in this situation is around 3.5 mA. The difference between these two almost identical states is large, since the only difference is what is displayed on the LCD.

5.2.4 Brightness

The brightness is tested for all available values that can be given with commands, which is a number in the range of zero to six. Looking at the figure 5.3 A the tile starts to use a lot of power, even when the LED's are off. The difference between the lowest and off brightness levels is very small. Looking at figure 5.1, it should be around 3.5 mA. The LED's are also more efficient at brighter settings as the curve flattens as seen in Figure 5.3 B. This means that the power consumption changes almost nothing when the brightness does increase. The backlight also uses a lot of power compared to other components of the tile.

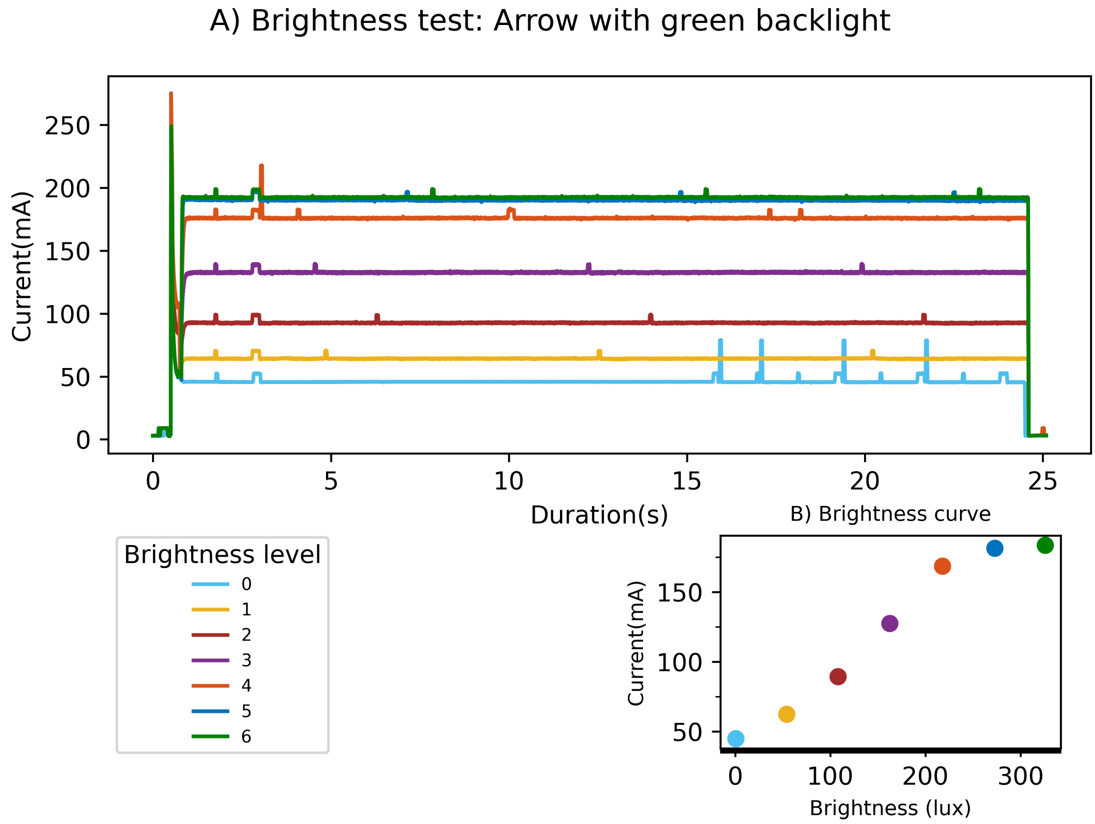

FIGURE 5.3: Brightness tests - colour green - icon arrow. It shows the influence of the brightness on the power consumption. It also shows the measured brightness.

5.2.5 Colour

Another variable that can change when using the backlight is the colour of the LED's and as can be seen in Figure 5.4 A, it does influence power usage. The tile has three colour LED's: red, green, and blue. All other colours are a combination of those. Blue is the most efficient colour, while white, which uses all three, uses the most power. Red seems to be the most inefficient when looking at the brightness curve in Figure 5.4 B. Both purple and red are above the curve when looking at brightness and power consumption.

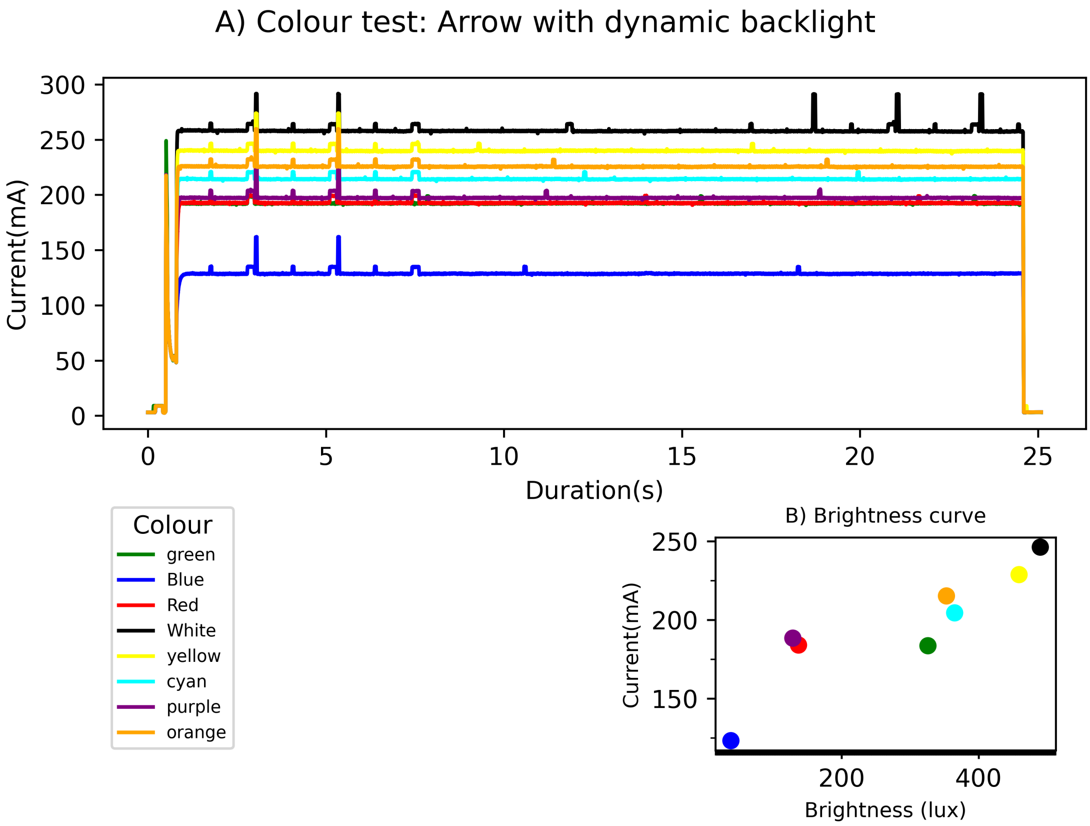

FIGURE 5.4: Color tests - brightness mode 6 - icon arrow. It shows the influence of the colour on the power consumption. It also shows how bright the colours are perceived.

5.2.6 Blinking

Looking at Figure 5.5, there are no anomalies. Power consumption reacts fast to the LED's being turned on or off, so there is no delay in turning it off. The lower the duty cycle, the lower the average power consumption, as can be seen in Table 5.1.

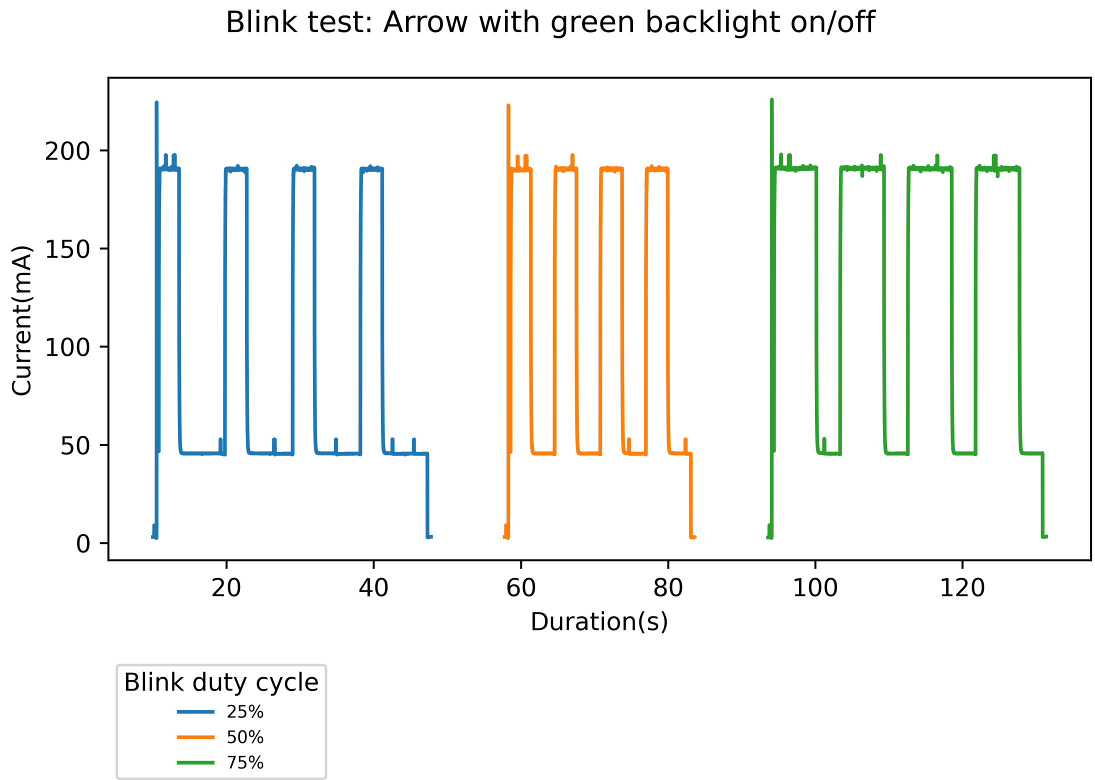

FIGURE 5.5: Blinking tests - brightness mode 6 - color green - icon arrow. It shows how the power consumption behaves during three different duty cycles.

5.2.7 Power budget

In Figure 5.6 it can be seen that the battery will have an energy shortage when the average power consumption is greater than 15 mA. At this point, there are days when the battery will be below the 40% discharge level. When using an average of 15 mA, the batteries still have a lot of days when they are full.

Taking a closer look at how the battery performs each month in Figure 5.7 when the average power usage is 15 mA. It can be seen that only January, February, November, and December produce to little power on average to sustain the 15mA. All the other months produce much more than the 15mA. A more detailed report can be seen in Appendix B

5.2.8 Power calculations

The tile can have many states without taking nudging into account. For the calculation, there will be three states in which the tile can be. The first and most efficient is the sleep state. In this state, the tile will be in deep sleep and uses almost no power. The power consumption in this state as seen in Figure 5.1, is 0.03. The next state is the active state, this is the most power-consuming state the tile can be in without using nudges. In this state, the radar will be on and the tile in class B, so it can react quickly when the platform becomes crowded. In this state, the power consumption will approximately be 6.3 mA. The last state is the idle state; the tile is active and will periodically wake up to check the state of the platform and give a status update in class A. From the test with the radar, a

Test	power consumption (mA)
Deep sleep	0.03
Radar no logging	6.218
Radar Logging	6.514
LoRa Class A - message every minute	0.477
LoRa Class B	3.033
Turn on LCD	3.496
Display arrow with command - brightness 0 - green	45.975
Display arrow with command - brightness 1 - green	62.401
Display arrow with command - brightness 2 - green	89.325
Display arrow with command - brightness 3 - green	127.467
Display arrow with command - brightness 4 - green	168.538
Display arrow with command - brightness 5 - green	182.167
Display arrow with command - brightness 6 - green	183.759
Display arrow with command - brightness 6 - blue	123.209
Display arrow with command - brightness 6 - red	183.297
Display arrow with command - brightness 6 - Yellow	228.718
Display arrow with command - brightness 6- cyan	204.411
Display arrow with command - brightness 6- magenta	188.350
Display arrow with command - brightness 6- orange	215.170
Display arrow with command - brightness 6 - white	246.182
Blink - Duty cycle 25%	89.588
Blink - Duty cycle 50%	109.992
Blink - Duty cycle 75%	135.778

Table 5.1: average power consumption over a minute for all tests

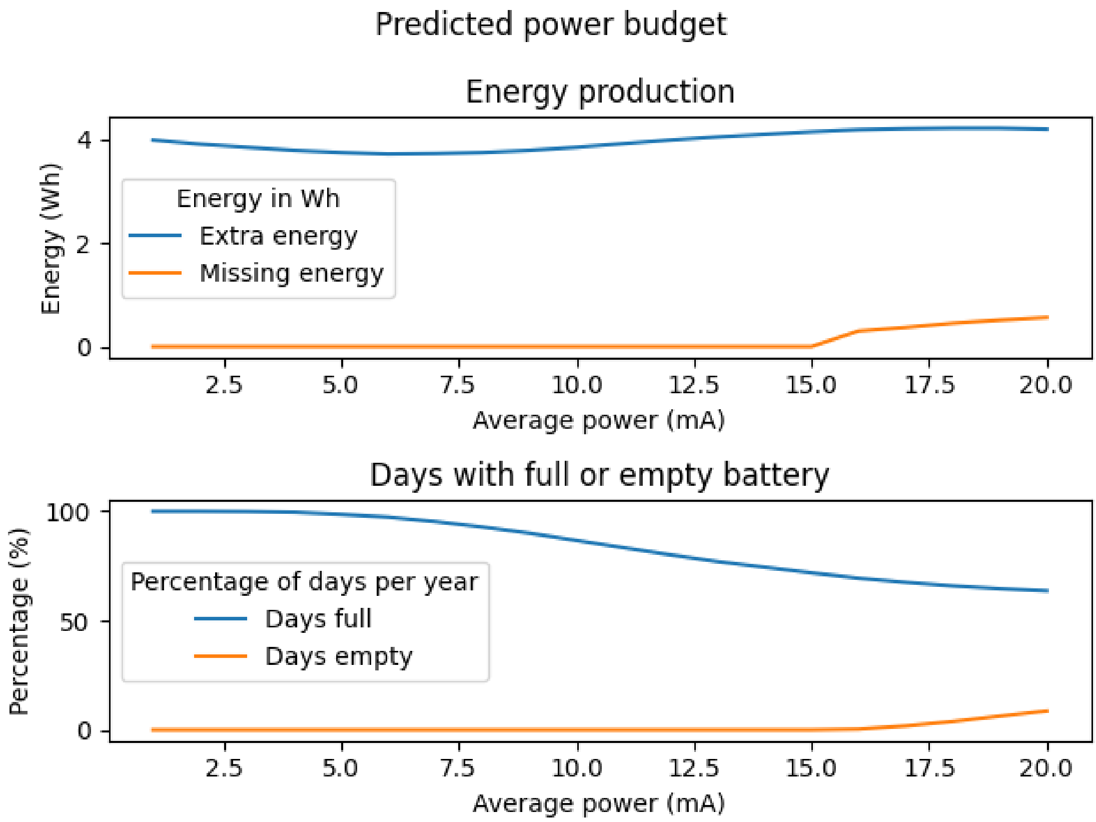

FIGURE 5.6: How the battery will perform for different average power consumption's

radar test needs 3 seconds worth of samples to be as accurate as possible. A possibility is that it would do this every minute, which means 5 seconds (10%) to turn on, perform the radar measurement, and send a status update. During these 5 seconds, it uses 6.3mA, the other 55 seconds (90%) it will sleep and use 0.03mA. This means that the idle state uses (6.3 * 0.1) + (0.03 * 0.9) = 0.66 mA on average over a minute.

Sleep state can be used in the period when no one is on the train platform, this is between 01:00 and 05:00 [22], as shown in figure 5.8. NS has a subscription called dalvoordeel (off-peak hours) [23] that is not active from 6:30 to 09:00 and from 16:00 to 18:30, those hours we will take as rush hour. These are typically the busiest hours of the day, and the active state will be used to react quickly to the situation. All other hours, it is required to be aware of the situation but not as fast, so the idle state can be used. This means that the energy consumption during the day will be ((4*0.03) + (5*6.3) + (15*0.66))/24 = 1.73mA. This leaves plenty of room for nudges.

To know how many nudges a tile could give on a day using the above calculation, the unit of mAh will be used. The power budget per day is 15 * 24 = 360 mAh, subtracting the average consumption of 1.73mA or 1.73*24 = 41.52 mAh. This leaves 360 - 41.52 = 318.48 Ah for nudges per day. If we take a green arrow blinking on full brightness with a duty cycle of 50% that takes on average 110.00 mA. The signal could be given for 318.48/110.00 = 2.90 hours or 10440 seconds. If a nudge consists of this signal blinking for 10 seconds, it could be given 1044 times every day.

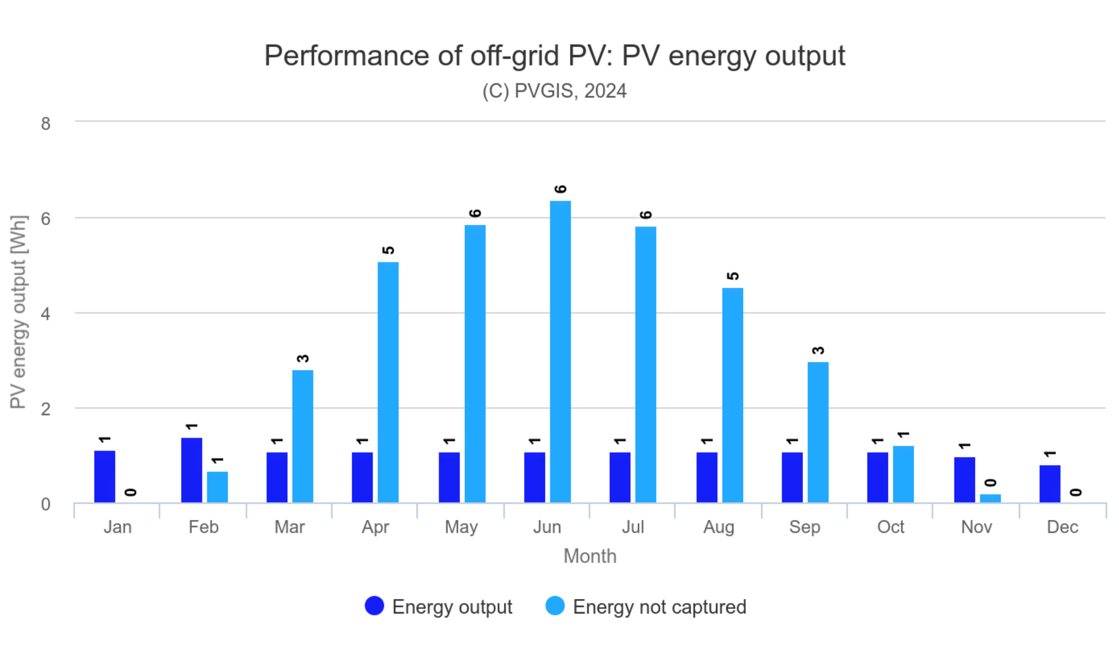

FIGURE 5.7: Power production per month

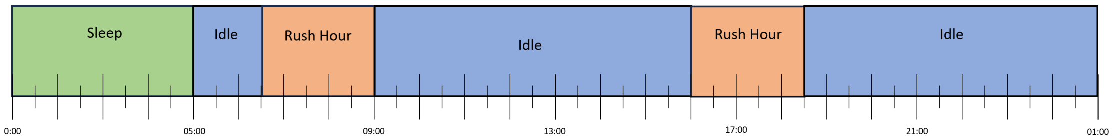

FIGURE 5.8: Proposed states during a typical day

5.2.9 Results long power test

The tile sends status updates during the day; this was also the case during the radar tests on Den Bosch. Part of the status is the battery voltage at that moment. How the battery was drained during the test can be seen in Figure 5.9. During this time, it was using the radar every day from 06:00 to 02:00 the next day. In addition to the radar, the tile also shows an animation in green with brightness 2 for 3 seconds every 3 minutes. If the tile is not collecting data, it will be in the sleep state.

During the day, the radar with logging enabled uses 6.51 mA for 20 hours every day. The animation uses around 89mA when looking at the power consumption of a green arrow with the same brightness. This 89mA is with the power of the tile itself, which is already on because the radar is on. The power consumption of LoRa class B can be subtracted from this value, so 89.69 - 3.03 = 86.66 mA. per 3 minutes, the tile would be 3 seconds, resulting in an average power consumption of (3/180) * 86.66 = 1.44 mA. The total power consumption per day would be ((6.51 + 1.44) * 20 + 0.03*4) = 159.12 mAh per day or 159.12/24 = 6.63 mA on average. This is below the allowed power budget. Looking at Figure 5.9, the battery from the tile near the elevator appears to slowly drain. However, this happens during the months when the tile is expected to lose energy, since it was done in December and January. During these 500 hours, it did not lose much, especially since the Li-ion discharge curve has a voltage plateau [24]. And according to the manufacturer [25], this voltage plateau for this battery is at 3.6V. which it did not reach.

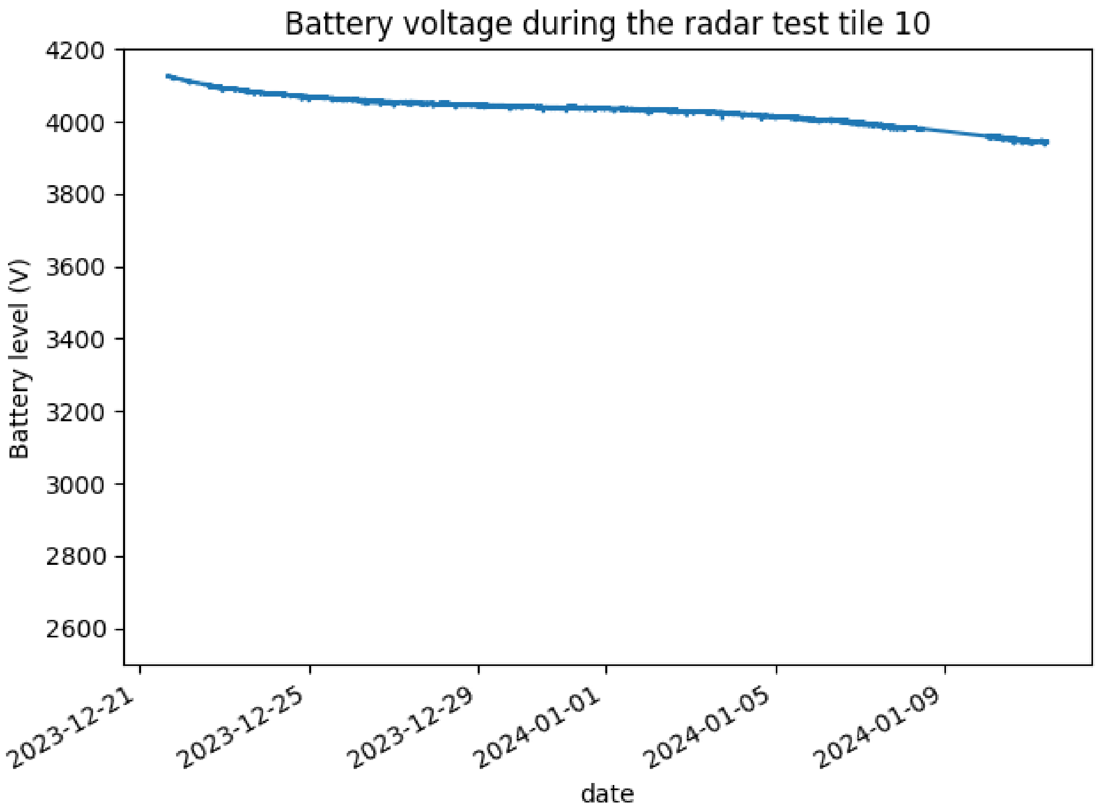

FIGURE 5.9: Battery voltage during the radar test on Den Bosch. It shows the battery voltage in the total range the battery can be.

5.3 Discussion

The goal was to see if the smart tile power system is sufficient for its power consumption when used for its intended purpose. When looking at the power consumption tool, the average power consumption that the tile can have to function throughout the year is 15 mA. The tile while performing its duties without taking nuding into account was only 1.73 mA. This should leave plenty of room for nudges. With an example nudge that takes 10 seconds, it could perform 1044 nudges.

This was with strange behaviour where the power consumption with the LEDs being of was higher than expected. Since there are two tests done with the LCD but without backlight, one test had an average power consumption of 3.5 mA, while the other had 45 mA. The state without backlight when blinking was closer to 45 mA, which means that it could be possible to lower the average power consumption while giving a nudge. It is not clear where this high power consumption originates from. Since the tile is capable of the lower power consumption when looking at the test, it seems to be a software-related problem and not a hardware problem.

Since the exact curve of the battery is unknown, it is unclear what the net power consumption was during the radar test. But the voltage drop seemed low, especially considering that radar and LEDs were used often, and it was during the worst months in terms of solar power generation. A longer test is needed to see if the tile is recharging starting in April. If that is the case, maybe some tests where the average power consumption is set higher each day and see how the battery holds up. At the moment we only know that solar panels generate less power than the tile used in that period.

If it turns out that the tiles generate less energy than calculated, there are energy-saving measures that can be taken. The first would be to reduce the brightness of the nudge. Because when looking at the power consumption of all components, LEDs use the most power. This is considering the baseline of the tile without nudging is sufficient. Using the brightness dynamically with the light level of the surrounding area. If it is less bright outside, the tile will generate less energy, but it is also easier to see the LEDs, so they can be turned down. This could greatly reduce power consumption during the winter months. Reducing the baseline power consumption also helps; this can be done by reducing the update frequency. The tile supports a command that changes how many minutes there are between each update. Another option would be to reduce the time windows of the radar; it performed best with a window of three seconds. However, it still performed fairly well with a shorter window. But this depends on the performance of the tile since there was already a big difference between the tiles currently used.

When looking at the calculations, it should be feasible for the tile to give enough nudges. However, it is hard to know how well the tiles will perform in reality. The real-world power production could be lower than the one used during the calculations. Only using the tiles for a long period will tell for sure how they perform.

Chapter 6

Noticeability study

The last aspect that needs to be researched is, whether people notice the tile. The previous study [3] researched which nudges work best. However, during this study, the focus was already on the tiles. It did not consider how well the tiles work as a medium to convey nudges. Since the study did not represent the actual situation, this is why this study will look at how well people see the tiles when placed in the actual location on a train platform on Den Bosch. The people will also be regular passengers going on about their day. That means people being on their phones or hurrying to catch the train. Doing it this way brings extra challenges, due to the unpredictable nature of a busy train platform. Trains are arriving, the weather can be bad and behaviour can differ depending on the location. But doing it here will give the most relevant result on whether a tile placed on the ground is suitable as a medium to convey information. The noticeability of the tile will be verified by counting the people who showed a visible reaction to the tile displaying something compared to people who did not show a reaction.

6.1 Methods

6.1.1 materials

The materials needed to perform the experiment are a method to display something on a tile and a piece of paper or a laptop to count people who walked past the tile when the nudge is shown.

6.1.2 participants

Participants in the experiment will be passengers who walk on the Den Bosch train platform. The partcipants are not made aware of the tiles. The only data collected will be if a person reacts to the tile or not. No descriptives or other information will be recorded that can be tracked to a person. ProRail has posters to inform people that there are experiments that are performed on a specific platform. No additional consent will be requested from participants.

6.1.3 Procedure

First, a pilot study will be conducted to see how the situation is and what to look for when designing the experiment. To get a general idea of how the system works. The tile showed a few festive images triggered by the researcher. This is to see how fast the system reacts and what kind of reaction people show if they notice the image on the tile.

For the actual experiment, a bright blinking image will be continuously displayed for 45 minutes. The blinking image can be seen in Figure 6.1. This should grab the most attention and also make sure as many people notice it as possible. From the pilot study, it was clear that not all commands are received. And sometimes it can take a while. By having it always show the signal, everyone has the change to see the signal.

During the experiment, the researcher will take note of when people showed a reaction that could indicate if they saw the tile or not, but only when it was clear. So, it was only marked as not a reaction when the person's face was visible the entire time. If this was not the case, it was not marked due to unclarity.

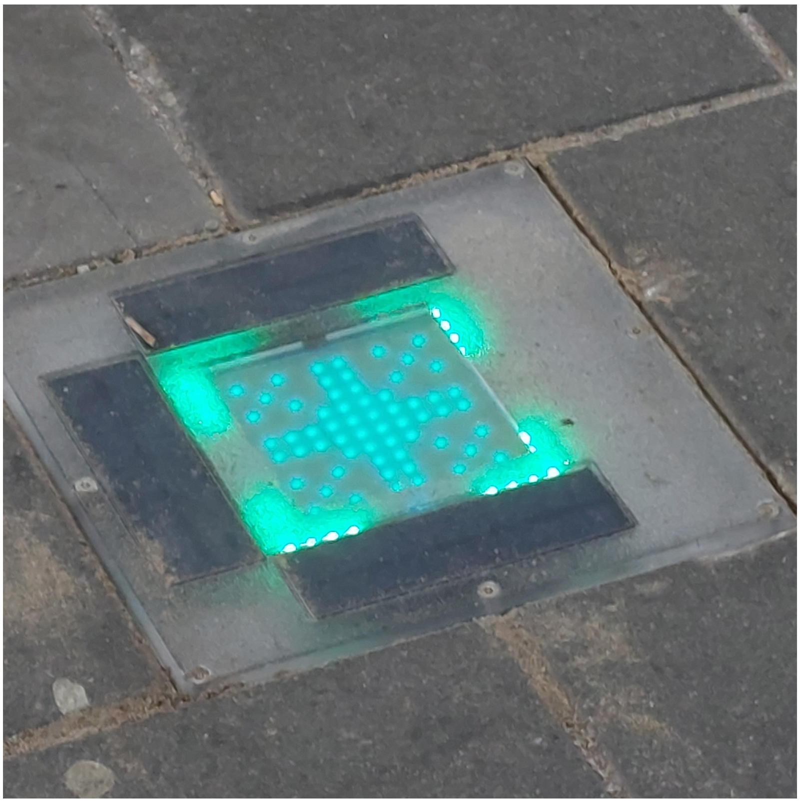

FIGURE 6.1: Nudge shown during observations

6.1.4 Outcome Measure

The goal of the experiment is to see if passengers will notice the tiles when something is displayed. In this case, the result will be a percentage of people who showed a visible reaction to the tile that displays an icon. Compared to people who did not show a visible reaction.

6.2 Results

A total of 166 people have been observed; of these 166 people 55 showed a visible reaction to the tile displaying something. This is 33.1% of the observed people. When looking at the tiles individually, the tile beneath the escalator had 37% of the people looking while near the elevator had 31%.

6.3 Discussion

The end goal of the tile will be to try to redirect passengers to a different place on the train platform. This can only happen if passengers notice the nudge given by the tile. If passengers do not, it will not matter how well the radar works or how many nudges can be given. The experiment carried out looked at a best-case scenario for the tile and was aimed at seeing if people showed any reaction to seeing the tile. On average 33% of the passengers observed showed a reaction to the tile. When comparing this with how many people notice advertisements at an airport, as researched by Wilson [26]. This percentage is higher since the average was 20% for that study. That percentage was even lower near the gate or in a busy area. The percentage of people who notice the tile could be even higher. Considering that there may be people who may have seen the tile but did not show an observable reaction, the score is good. The test was carried out during the winter months, which means that it was not very bright outside. The brightness was at the maximum. It could be that during the summer people do not notice the tile as well.

The location of the tiles was also not ideal, but it was the only place that was available during the experiment. One location was below the elevator and the main entrance to the platform. People will hurry past here and will not stand still. This gives less opportunities for people to stand still and notice the tile. The other location was near an elevator, more people stood still here, but it was not where the crowds would form. It is located further from the train platform border and is still a place where people enter the platform.

Many more tests could be performed to map the extent of the usability of the nudges. The first is the effect of colour; during this experiment, only the cyan colour was used. It could be that other colours are noticed better or worse. This can be due, for example, to the way certain colours are better detected by the eye [27]. How well wavelengths are observed by humans is called the eye sensitivity curve. Taking a closer look at this and finding the optimal colour could reduce power consumption or improve the way people notice the tile. Secondly, the brightness is also something that can help improve the performance of the tile. Lower brightness decreases battery consumption. Knowing what the lowest brightness is for a given environmental light can greatly increase the number of nudges.

Chapter 7

Discussion

At this moment, the system is not yet on a level where it can be deployed. However, looking at all the components that make up this system, it is certainly possible. And when looking for other low-power systems that can detect and nudge, not much could be found in literature. The concept of the smart tile is applicable not only on a train platform, but could also be used in other environments like cities. Having these tiles placed around the city could give insight into how people walk, redirect them to improve flow, and also provide environmental data (temperature, solar intensity) with great resolution. Currently low-power systems are mainly used only for data gathering, like monitoring the weather [28].

Looking at the results of the two radar tests, the Doppler radar is capable of reliable detecting activity within a range of 1.25M. However, this result was only for the tile near the entrance of the platform. The tile near the elevator performed worse and could only detect within a range of 0.5M with the same reliability. It should still be possible to detect whether the platform is crowded or not, but the range might be too limited. Because the results of the two tiles are so different, tests with more tiles should be done to get a better understanding of how the radar performs.

When considering a power budget of 15 mA, the tile in the current implementation is capable of performing more than 1000 nudges every day. The nudges can vary in how much power they use, but the chosen nudge was using the colour green, full brightness and a duty cycle of 50%. A different colour and brightness can change the power consumption of a nudge. But even if the nudge used twice as much power, the total amount of nudges would still be over 500. At the moment, I would consider the tile to be capable of being used to nudge people throughout the year. The restriction of 15 mA is mostly during the winter months; during the summer, the budget is much higher, and more nudges can be given.

The most important part is whether people even look at the tile. The first observations are positive, and the numbers are comparable to another study that focused on people looking at signage. To get a better understanding, the test could be repeated during the summer with more tiles located near the border. Near the border because that is where most tiles will be located, and during the summer to see if the LED brightness is enough when it is brighter outside. The current experiment was executed during winter, when it was darker outside and the tile was set to maximum brightness. The maximum brightness could be too low to attract attention during the summer months.

In a previous study [3], scenarios were developed that could be used to nudge people. The most promising scenarios are using arrows to steer people from crowded to non-crowded places and to indicate where the train stops. Steering people away from crowded places is something that can be looked at. The radar can be used to detect activity, the next step is to take that information to know where people have walked and see where people did not walk. This way you know between which tiles people should be standing or waiting. And this information can be used to distinguish between crowded and less crowded zones and direct people to spread.

The second scenario is also possible if the stopping location of the train is consistent and known. The radars can detect where the trains are stopping and with these data the stop position of the train can be predicted. And use that information to predict the stop location. Then use the tiles to indicate where this will be before the train arrives. Other scenarios that showed less potential but can still improve the efficiency and safety of the train platform, do not require the tile radar. They only need the control over the LCD and LED of the tile, and a battery that can run the tile the whole year. This is something the tile is capable of for sure.

The tile also has another radar that was not functional when doing the experiments. This radar is better able to detect non-moving objects. Considering that the problem is people waiting near each other, this radar could give a more accurate estimate of the crowdedness. The current analyses focused on detecting any activity or no activity. Over time, this could tell something about how crowded the platform is. However, for a single measurement, this does not tell if the situation is crowded or not only if a person walked near the tile. When doing further research on the radar, the correlation between the number of people and the value of the radar could be investigated. When looking at the data gathered from the lab tests, this correlation could be weak. But even a degree of correlation could tell more about the measurement. KITT can also make a second version of the tile, one that has multiple radars, so the tile can detect for multiple zones if there is activity or not.

The method of analysis was kept simple; it only looked at the average value and compared the time windows of the two data sets. It did not look at patterns or other statistics; there might be more to learn from the gathered data. It could be that a neural network can find other correlations and better predict the situation. According to the manufacturer, the chip used in the tile can take advantage of this [29].

So far, the combination of an autonomous system that can be used to nudge people has not been researched. There are examples of low-power autonomous systems with radar. Ditzel [30] tested an experimental wireless sensor network using a low-power radar. The goal is to see if it is feasible to monitor remote areas for animal monitoring or military purposes. The power target was in the same order as the tile, so it would run for a long time. But for this it was only gathering data, the system was not also used to display information.

The combination of a low-power battery-operated system intended for signalling has also been done. An example is the Intelligent Sign Control System from Torralba [31]. This is a system aimed at rural areas where there is not much power available. The system detects dangerous situations and can flash the LEDs to attract attention. This system is autonomous, just like the smart tile, but on a different scale. The tile package is much smaller and can be integrated more easily. The Torralba setup was on a large cart, not something that can disappear in the background when not used.

Chapter 8

Conclusion

KITT Engineering developed the smart tile to improve the safety of the train platform. The smart tile should do this by detecting a crowded situation and with the use of nudges direct the passengers to a quiet area. This is why the main research question will be 'To what extent can the smart tile be used to improve train platform safety?'. To reach this goal, three elements of the tile had to be researched: detecting activity with the tile; power consumption and availability of the tile; and the visibility of the tile.

First, 'How reliable is the Doppler radar in detecting activity?' Looking at the results of the Doppler radar, it can detect when one or more people walk past the tile. The range to reliable detect this differs per tile. The tile near the entrance could do this to around 1.25 metres in all directions. However, the tile near the elevator only up to about 0.5 metres. Going with the performance of the tile near the entrance shows that the radar should be capable of reliably detecting up to 1.25 metres. This was also the range found when performing tests in a lab setting. This means that the radar can be used to detect activity in its area and know if it is crowded or not. Next, 'How does the power system of the smart tile compare to its power usage?' The calculated available power budget was 15 mA on average per day. With this power consumption, the tile can function for the entire year. For a realistic power consumption without nudging, the tile should use around 1.73 mA. When looking at an example nudge and its power consumption, the nudge of 10 seconds could be done 1044 times per day. Meaning that the tile can give function and nudge during the entire year. Lastly, 'Do people notice the smart tile on the train platform?' On average 33% showed a visible reaction that they did see the tile performing the nudge. This is higher than the average number of people who notice advertisements at an airport. This means that the number of people is not particularly low.

The tile is in a state where looking at the effectiveness of nudging using the tile can be done. Since the three important elements needed to perform nudges are sufficient.

Bibliography

[1] A. Schneider, E. Krueger, B. Vollenwyder, J. Thurau, and A. Elfering, "Understanding the relations between crowd density, safety perception and risk-taking behavior on train station platforms: A case study from switzerland," *Transportation Research Interdisciplinary Perspectives*, vol. 10, p. 100390, 6 2021.

[2] prorail, "Perronrandbeveiliging | ProRail," https://www.prorail.nl/toekomst/innoveren/perronrandbeveiliging, accessed: November 28, 2023.

[3] M. Pijnappel, "Research topics: Improving the train station border safety using autonomous intelligent tiles." *Unpublished report*, June 2023.

[4] Prorail, "Handboek bewegwijzering | spoorbeeld," https://www.spoorbeeld.nl/beleid/station/handboek-bewegwijzering, accessed: 11-12-2023.

[5] P. G. Hansen and A. M. Jespersen, "Nudge and the Manipulation of Choice: A Framework for the Responsible Use of the Nudge Approach to Behaviour Change in Public Policy," *European Journal of Risk Regulation*, vol. 4, no. 1, pp. 3–28, 3 2013. [Online]. Available: https://www.cambridge.org/core/journals/european-journal-of-risk-regulation/article/nudge-and-the-manipulation-of-choice/D1ED64479FF868BD79FFE90E76A4AB54

[6] D. Postma, S. ten Brinke, R. van Delden, and D. Reidsma, "Coaxing: An Empirical Exploration of a Novel Way to Nudge Athletic Performance in Sports," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 13213 LNCS, pp. 174–189, 2022. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-030-98438-0_14

[7] R. van Delden, A. Moreno, R. Poppe, D. Reidsma, and D. Heylen, "Steering gameplay behavior in the interactive tag playground," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 8850, pp. 145–157, 1 2014. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-319-14112-1_13

[8] C. R. Sunstein, "Nudging: A Very Short Guide," *Journal of Consumer Policy*, vol. 37, no. 4, pp. 583–588, 12 2014. [Online]. Available: https://link.springer.com/article/10.1007/s10603-014-9273-1

[9] P. G. Hansen, "Green Nudge: Nudging Litter Into The Bin | iNudgeyou." [Online]. Available: https://inudgeyou.com/en/green-nudge-nudging-litter-into-the-bin/

[10] *Data Sheet SMR-313/333*, InnoSenT, 7 2020, version 2.2.

[11] *Data Sheet SMR-314/334*, InnoSenT, 2 2022, version 2.3.

[12] *Long Range, Low Power, sub-GHz RF Transceiver*, Semtech, 12 2021, revision 2.1.

[13] lora alliance, "Lora — lora documentation," url https://lora.readthedocs.io/en/latest/, accessed: 13-12-2023.

[14] Chripstack, "ChirpStack open-source LoRaWAN Network Server," https://www.chirpstack.io/, accessed: may 8, 2023.

[15] A. Whitworth, A. Droitcour, C. Song, O. Boric-Lubecke, and V. Lubecke, "Characterization technique for a doppler radar occupancy sensor," *Electronics 2023, Vol. 12, Page 4888*, vol. 12, p. 4888, 12 2023. [Online]. Available: https://www.mdpi.com/2079-9292/12/24/4888/htmhttps://www.mdpi.com/2079-9292/12/24/4888

[16] G. Manfredi, I. D. H. Sáenz, M. Menelle, S. Saillant, J. P. Ovarlez, and L. Thirion-Lefevre, "Measurements and analysis of the doppler signature of a human moving within the forest in uhf-band," *Remote Sensing 2021, Vol. 13, Page 423*, vol. 13, p. 423, 1 2021. [Online]. Available: https://www.mdpi.com/2072-4292/13/3/423/htmhttps://www.mdpi.com/2072-4292/13/3/423

[17] Y. Wang, Q. Liu, and A. E. Fathy, "Simultaneous localization and respiration detection of multiple people using low cost uwb biometric pulse doppler radar sensor," *IEEE MTT-S International Microwave Symposium Digest*, 2012.

[18] E. L. Chuma and Y. Iano, "Human movement recognition system using cw doppler radar sensor with fft and convolutional neural network," *2020 IEEE MTT-S Latin America Microwave Conference, LAMC 2020*, 2020.

[19] A. Alnaeb, R. Syamsul, A. R. Abdullah, A. A. Salah, A. Sali, N. E. A. Rashid, and I. Pasya, "Human activities detection and classification based on micro-doppler signatures near the baseline of forward scattering radar," *2018 International Conference on Radar, RADAR 2018*, 12 2018.

[20] K. Hanifi and M. E. Karsligil, "Elderly fall detection with vital signs monitoring using cw doppler radar," *IEEE Sensors Journal*, vol. 21, pp. 16 969–16 978, 8 2021.

[21] E. Commission, "Jrc photovoltaic geographical information system (pvgis) - european commission." [Online]. Available: https://re.jrc.ec.europa.eu/pvg_tools/en/tools.html#PVP

[22] treinposities.nl, "treinposities when are train on the platform," https://treinposities.nl/vertrekstaat/s-hertogenbosch/2024-02-05, accessed: 04-01-2024.

[23] NS, "NS when is rush hour," https://www.ns.nl/abonnementen/dal-voordeel, accessed: 11-12-2023.

[24] S. S. Zhang, K. Xu, and T. R. Jow, "Charge and discharge characteristics of a commercial licoo2-based 18650 li-ion battery," *Journal of Power Sources*, vol. 160, pp. 1403–1409, 10 2006.

[25] djbattery, "Dj18650-34h-3.7v3400mah cb certificates,3.7v-dj18650 and 26650 cells." [Online]. Available: https://djbattery.net/products/dj18650-34h-37v3400mah-lithium-ion-battery-82

[26] R. Wilson and B. Till, "Airport advertising effectiveness: An exploratory field study," *Journal of Advertising - J ADVERTISING*, vol. 37, pp. 59–72, 04 2008.

[27] D. Poelman, N. Avci, and P. F. Smet, "Measured luminance and visual appearance of multi-color persistent phosphors," *Optics Express*, vol. 17, p. 358, 1 2009. [Online]. Available: https://www.researchgate.net/publication/23764706_Measured_luminance_and_visual_appearance_of_multi-color_persistent_phosphors

[28] E. Murdyantoro, R. Setiawan, I. Rosyadi, A. W. Nugraha, H. Susilawati, and Y. Ramadhani, "Prototype weather station uses lora wireless connectivity infrastructure," *Journal of Physics: Conference Series*, vol. 1367, p. 012089, 11 2019.

[29] STMicroelectronics, "Optimize neural networks on stm32 with stm32cube.ai," https://www.st.com/resource/en/product_presentation/stm32cubeai_press_pres.pdf, accessed: 02-12-2024.

[30] M. Ditzel and F. H. Elferink, "Low-power radar for wireless sensor networks," *Proceedings of the 3rd European Radar Conference, EuRAD 2006*, pp. 139–141, 2006.

[31] A. Torralba, J. P. Garcia-Martin, J. M. Gonzalez-Romo, M. Garcia-Castellano, J. Peral-Lopez, and V. Perez-Mira, "An autonomous, intelligent sign control system using wireless communication and led signs for rural and suburban roads," *IEEE Intelligent Transportation Systems Magazine*, vol. 14, pp. 115–128, 2022.

Appendix A

Overview of FFT bins

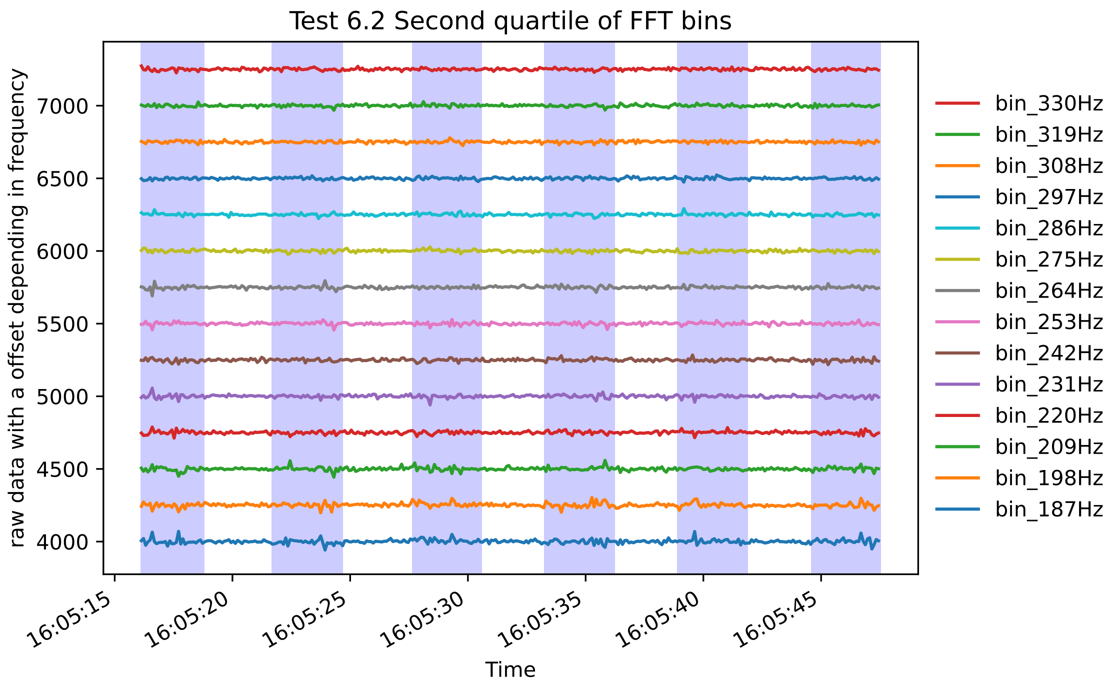

FIGURE A.1: Second quarter of FFT bins plotted overtime for normal walking speed and 1 meter distance

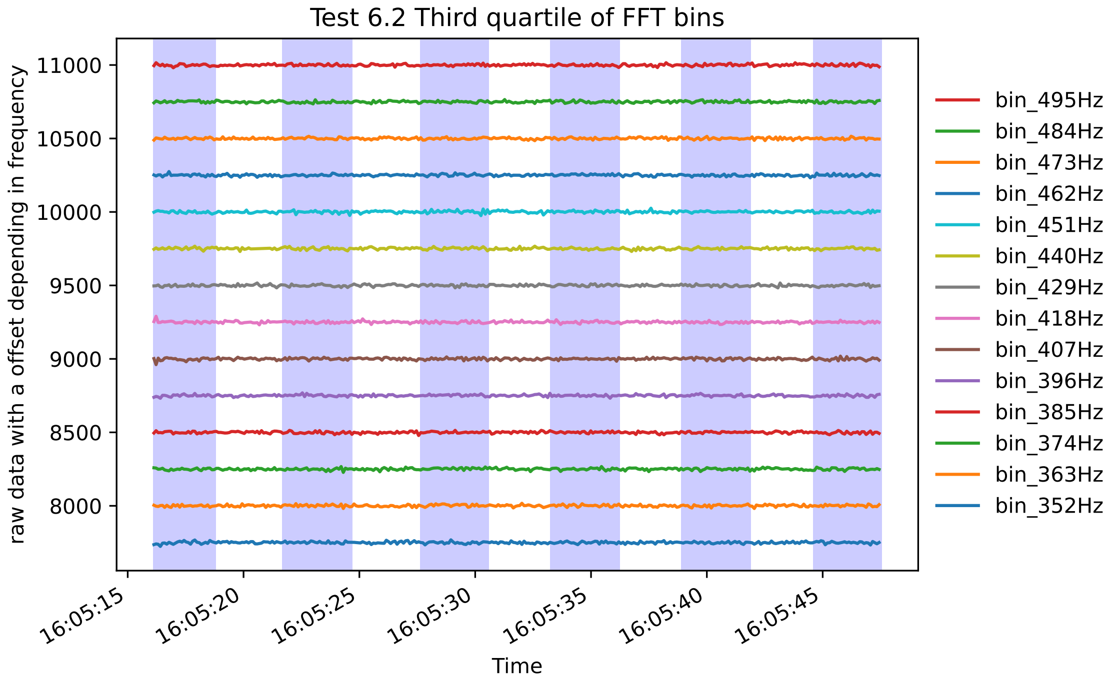

FIGURE A.2: Third quarter of FFT bins plotted overtime for normal walking speed and 1 meter distance

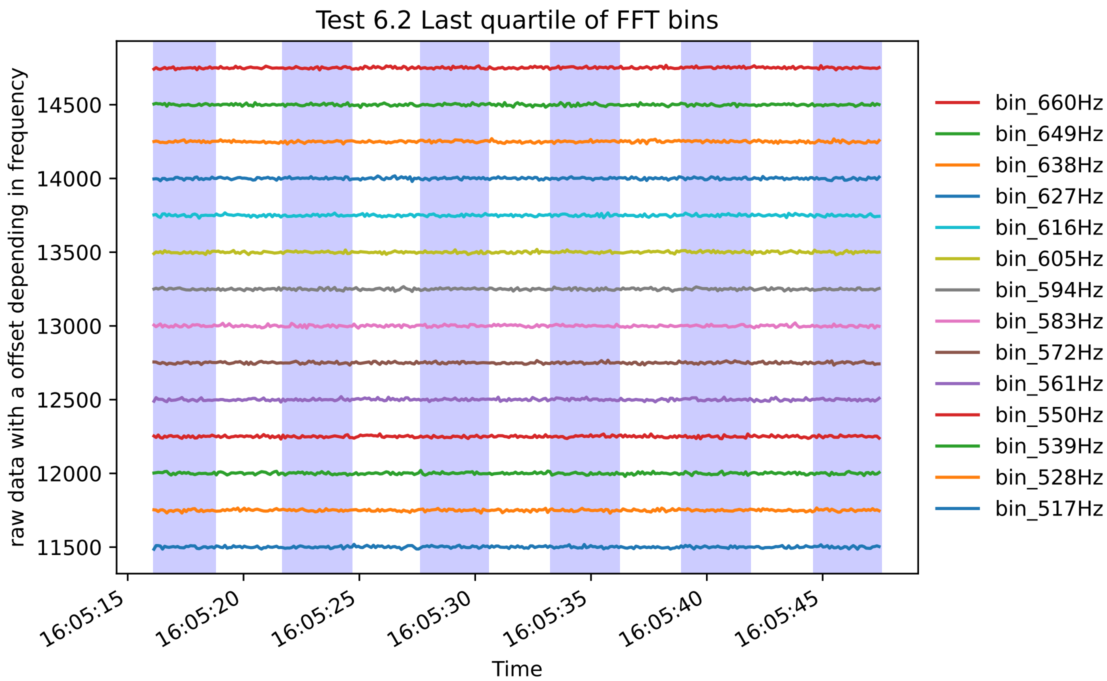

FIGURE A.3: Last quarter of FFT bins plotted overtime for normal walking speed and 1 meter distance

Appendix B

Detailed power prediction report

European Commission

Performance of off-grid PV system

PVGIS-5 estimates of solar electricity generation

Provided inputs

Latitude/Longitude:	51.687,5.293
Horizon:	Calculated
Database used:	PVGIS-SARAH2
PV installed:	2 Wp
Battery capacity:	40.8 Wh
Cutoff limit:	40 %
Consumption per day:	1.08 Wh
Slope angle:	0 °
Azimuth angle	0 °

Simulation outputs

Percentage days with full battery:	71.84 %
Percentage days with empty battery:	0 %
Average energy not captured:	4.15 Wh
Average energy missing:	0 Wh

Outline of horizon at chosen location:

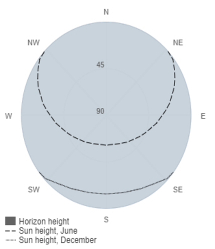

Power production estimate for off-grid PV:

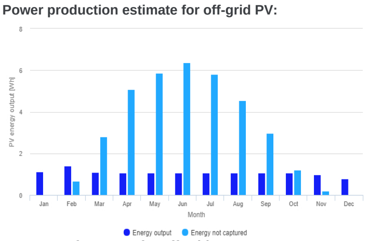

Battery performance for off-grid PV system:

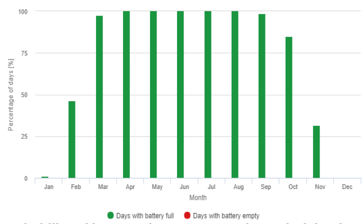

Probability of battery charge state at the end of the day:

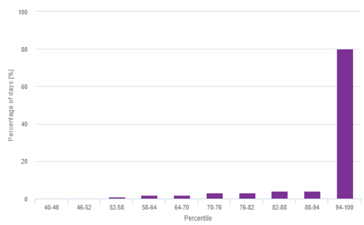

Monthly average performance

Month	E_d	E_l	f_f	f_e
January	1.1	0.0	1.0	0.0
February	1.4	0.7	46.4	0.0
March	1.1	2.8	97.6	0.0
April	1.1	5.1	100.0	0.0
May	1.1	5.9	100.0	0.0
June	1.1	6.4	100.0	0.0
July	1.1	5.8	100.0	0.0
August	1.1	4.5	100.0	0.0
September	1.1	3.0	98.8	0.0
October	1.1	1.2	84.9	0.0
November	1.0	0.2	31.5	0.0
December	0.8	0.0	0.0	0.0

E_d: Average energy production per day [Wh/day].
E_l: Average energy not captured per day [Wh/day].
f_f: Percentage of days when battery became full [%].
f_e: Percentage of days when battery became empty [%].

Cs	Cb
40-46	0.0
46-52	0.0
52-58	1.0
58-64	2.0
64-70	2.0
70-76	3.0
76-82	3.0
82-88	4.0
88-94	4.0
94-100	80.0

Cs: Charge state at the end of each day [%].
Cb: Percentage of days with this charge state [%].

Report generated on 2024/02/05