

Delft University of Technology  
Master of Science Thesis in Embedded Systems

# Dynamic Vital Sign estimation for Multiple Persons using mmWave technology

Dirk den Hoedt





# Dynamic Vital Sign estimation for Multiple Persons using mmWave technology

Master of Science Thesis in Embedded Systems

Embedded and Networked Systems Group  
Faculty of Electrical Engineering, Mathematics and Computer Science  
Delft University of Technology  
Mekelweg 4, 2628 CD Delft, The Netherlands

Dirk den Hoedt

06-10-2022

**Author**

Dirk den Hoedt ()

()

**Title**

Dynamic Vital Sign estimation for Multiple Persons using mmWave technology

**MSc Presentation Date**

14-10-2022

**Graduation Committee**

Marco Zuniga (chairman) Delft University of Technology

Francesco Fioranelli Delft University of Technology

Caitlin Ramsey MSc Delft University of Technology

Tom Goos Erasmus MC

## Abstract

When a patient is in a hospital, it is very important to monitor their vital signs. Doctors and nurses use this information to assess the condition of the patient. Most of the existing vital signs measurement devices need physical contact with the patient. This thesis focuses on a non-contact vital signs estimation method. Using a mmWave radar, one or more persons in the view of the sensor are being monitored. This monitoring consists of finding the chest region of a person, and monitoring this chest for vibrations. These vibrations are caused by breathing in and out, and the beating of the heart. Using signal processing, these vibrations can be converted to a heart rate and a respiration rate.

This thesis is about getting insight in the already available options regarding vital signs monitoring, programming the Texas Instruments IWR6843ISK mm-Wave radar module to estimate the vital signs of multiple persons and validating this project against trusted vital signs monitors.

The implemented solution which followed from this project is able to track multiple persons inside the radar view, and is able to measure the vital signs for up to four persons in real-time. The mean accuracy gained for one person heart rate estimation is 10.8%, the mean accuracy gained for one person respiration rate estimation is 7.6%. The mean observed accuracy for multiple person heart rate estimation is 13.4%, the multiple person respiration rate mean accuracy is 10.6%.



*“Never give up, for that is just the place and time that the tide will turn.”* –  
Harriet Beecher Stowe



# Preface

Because of my part-time job in the Erasmus MC, I have seen a lot of medical devices, most of them have one or more embedded systems inside of them. Because of my interest both for the embedded systems- and the medical world, I wanted to find a graduation project where both fields came into play. In the Erasmus MC I found Caitlin Ramsey, which did her PhD in the field of contactless vital signs monitoring of neonates using mmWave radar technology. Caitlin was looking for someone to research the ability to do the vital signs calculations on the chip hardware itself, where currently it was done using post-processing techniques. I was immediately interested in this topic, and after I found Marco Zuniga as my thesis supervisor from the TU Delft, I was ready to get started.

I want to thank all the persons that have helped me complete my graduation project. First of all, a special thanks to Caitlin for her great support during this project. I could always come to you with any questions I had, and you always answered them in great detail, both technical questions and more general questions. My thesis would not have succeeded without you. Also a big thanks to Marco. You were always able to keep an overview of the project, and helped me approach the project in an academic way. In moments of doubt, you were a great help to steer me in the right direction and to make the project as enjoyable for me as possible. I also want to thank all of the persons which helped me validate the project. Thank you for your time and patience while sitting very quietly in a chair for 5 minutes. I also want to thank all of the people who supported me along the way. Without your kind and encouraging words, I would have probably stranded somewhere halfway. In that regard, a special thanks to Geeske. You have always been there for me, by celebrating successes, and by cheering me up in difficult times. You always gave me the energy to push through and finish with a great result.

Dirk den Hoedt

Delft, The Netherlands  
6th October 2022



# Contents

<b>Preface</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Motivation . . . . .	1
1.2 Problem statement . . . . .	2
1.3 Objectives . . . . .	2
1.4 Outline . . . . .	3
<b>2 Background &amp; Related Work</b>	<b>5</b>
2.1 Background . . . . .	5
2.1.1 Vital sign detection methods . . . . .	5
2.1.2 mmWave technology . . . . .	7
2.1.3 Signal processing . . . . .	11
2.2 Related Work . . . . .	16
<b>3 People detection</b>	<b>17</b>
3.1 Overview . . . . .	17
3.2 People detection algorithm . . . . .	17
3.3 Existing solutions . . . . .	18
3.4 Implementation . . . . .	18
3.4.1 Input data format . . . . .	19
3.4.2 Noise removal . . . . .	19
3.4.3 Peak detection . . . . .	19
3.4.4 Result . . . . .	20
<b>4 Extracting Vital Signs</b>	<b>23</b>
4.1 Phase signal . . . . .	23
4.1.1 Radar parameters . . . . .	23
4.1.2 Phase unwrapping . . . . .	25
4.2 Vital signs estimation . . . . .	25
4.2.1 Data preparation . . . . .	26
4.2.2 Splitting the heart rate and respiration rate signal . . . . .	26
4.2.3 Circular buffers . . . . .	27
4.2.4 Peak counting . . . . .	29
4.2.5 Executing FFTs on the waveforms . . . . .	29

<b>5</b>	<b>Real-time implementation</b>	<b>33</b>
5.1	Overview . . . . .	33
5.2	How it works . . . . .	33
5.3	Building blocks . . . . .	34
5.3.1	Main sub-system . . . . .	34
5.3.2	DSP sub-system . . . . .	34
5.3.3	Inter-processor communication . . . . .	34
5.4	Implementation . . . . .	35
5.4.1	MSS program flow . . . . .	36
5.4.2	DSS program flow . . . . .	36
5.4.3	Timing . . . . .	36
<b>6</b>	<b>Validation</b>	<b>39</b>
6.1	Variables used . . . . .	39
6.1.1	Recorded variables . . . . .	39
6.1.2	External variables . . . . .	41
6.1.3	Person variables . . . . .	42
6.2	Testing setup . . . . .	42
6.3	Test subjects . . . . .	43
6.4	Results . . . . .	43
6.4.1	One person validation . . . . .	43
6.4.2	Two person validation . . . . .	46
6.4.3	Four person validation . . . . .	47
6.4.4	mmWave radar accuracy compared to test subject variables	48
6.5	Evaluation . . . . .	50
<b>7</b>	<b>Conclusions</b>	<b>57</b>
7.1	Future work . . . . .	58
<b>A</b>	<b>Tx beamforming</b>	<b>61</b>
A.1	Hypothesis . . . . .	61
A.2	Implementation . . . . .	62
A.3	Results . . . . .	62

# Chapter 1

## Introduction

### 1.1 Motivation

Vital signs monitoring belongs to one of the standard activities nurses perform to patients in a hospital. This has a very good reason, vital signs can give a good estimation of the state of a patient [12]. The most common methods to measure a patients vital signs require physical contact with the patient. A sensor has contact with the skin, and a wire runs from the sensor to another device which shows the heart rate. There are some downsides to this method of measuring. The sensor is always attached, so it could become a discomfort to the patient. Also, the mobility is limited because of the connected wires. Another problem could be that the sensor falls off. When this happens the device goes into an alarm state and produces a loud beep, which is annoying for the patient, especially at night. The nurse must also come and connect the sensor back to the patient.

Millimeter wave radar technology is already in use in the automotive industry [2]. This type of radar can help to detect other cars on the road. The radar module has antennae for sending the radar waves, and antennae for receiving the radar waves. A property that differentiates millimeter wave (mmWave) radar from other types of radar is the frequency of the transmitted signals. A typical mmWave radar can operate in the 30-300 GHz range. In this range, the length of one wave can be measured in millimeters, which explains the name of this radar type. This radar can detect objects on the road, even in difficult (weather) conditions.

MmWave radar can not only be used in the automotive industry, it could also be potentially used to estimate the vital signs of a person. The radar signals reflect on large bodies of water. 70% of the human body consists of water. This means that the mmWave radar signals can reflect on humans, a property that can be exploited to implement the vital signs estimation. When a person breaths, its chest moves up and down. The same goes for the beating of the heart. Every time the heart beats, the chest moves a tiny amount. The mmWave sensor can be setup in such a way that it can detect even those small vibrations of the chest. Using signal processing, those vibration signals can be transformed into a heart rate and a respiration rate.

Using this method solves various downsides compared to the conventional

way of measuring someones vital signs. The main advantage is that the sensor is non-invasive, it has no contact with the patient. This is especially beneficial for burn victims or babies that are born very premature. The patient doesn't even notice that he is being monitored. There are no cables connected to the patient which limit the movement, and no sensors which could fall off the body.

Depending on the antenna, this technology can have a wide field of view, opening up the possibility to measure multiple people using one radar module. For every person in the field of view of the sensor, a vital signs estimation can be performed. The benefit of this is that multiple persons can be monitored without investing in more hardware.

Using a FMCW radar for vital signs estimation also has its downsides. Radar data is very sensitive to noise, so it can only be used in a low noise environment. Because the sensor is contact-less, and the radar beams are invisible, the signal could be blocked by other persons or objects. Another downside that it is difficult to make distinctions between different persons in the radar data. This makes it hard for the sensor to figure out which vital signs belong to which persons.

The aim of this project is to have a mmWave sensor, which is programmed in such a way that it can detect persons, and measure the vital signs of one or more persons in the field of view of the sensor. All of these calculations will be done on the chip, and not on an external computer. Also, the accuracy of this sensor will be assessed.

## 1.2 Problem statement

The system must be able to measure the vital signs of multiple persons in reach of the sensor at the same time. There already are some projects which implement this function, but each has its own downsides. For some projects, the sensor is limited to only measuring one person at the same time. Other projects do succeed in measuring the vital signs of multiple persons at the same time, but only using a post-processing technique. This is not an optimal solution, because most of the time, the patients vital signs are needed right away.

The goal of this thesis is to implement a system which can measure the vital signs of multiple persons at once in real-time. Also, a thorough validation will be performed on the resulting system, to see if this solution can compete with the existing vital signs measuring solutions.

## 1.3 Objectives

For the implementation of this project, the Texas Instruments IWR6843ISK is used. This is a mmWave sensor which meets all of the technical requirements set for this project. Firmware for this module can be programmed using an IDE (Integrated Development Environment) provided by Texas Instruments.

From the initial problem statement, a set of technical requirements were established, which you can find below:

1. The sensor must be able to generate a heatmap from the radar data of the surrounding area.

2. The sensor must be able to detect one or multiple persons in this heatmap.
3. The sensor must be able to estimate the vital signs for up to four detected persons at the same time.
4. The sensor must compute these vital sign estimations in real-time.
5. The sensor must be able to track the measured person, such that the measuring continues even if the person moves to another location.
6. The outcome of the sensor must be displayed in a clear way using a GUI on a computer.
7. The project must be tested against a trusted measurement, to assess the accuracy of the project.
8. The project must be tested against a range of different persons, to see if the project performs the same on different kinds of persons.

## 1.4 Outline

In the introduction, which is Chapter 1, a brief introduction will be given about the project, along with the problem description and the objectives of this thesis.

Chapter 2 discusses the background information required to have a good understanding about the fundamental concepts this project is build on. Current vital signs measuring techniques will be discussed, but also some fundamental radar theory will be explained. It contains the explanation of the existing building blocks which are used to build this project, but also related work to this project, with the pros and cons of each work listed.

Chapter 3 focuses on the dynamic people detection algorithm. This algorithm detects persons in view of the sensor, and performs basic tracking to keep the persons in view and make sure that the vital signs estimation algorithms get the best possible data to work with.

Chapter 4 dives into the actual vital signs estimation functionality of this project. The algorithms used are explained and the program flow specific to the vital signs estimation will be visualized.

Chapter 5 discusses the real-time implementation of the system. The chip used in this project will be examined and a overview of the program flow of the whole project will be discussed and visualized.

Chapter 6 is the validation chapter. Now that a prototype has been built using the knowledge of the previous chapters, the prototype has to be validated against a trusted measurement to assess the accuracy of the prototype.

The last chapter, chapter 7, is the conclusion. There are also pointers provided for future work.

During this project, a lot of time has been spend on getting Tx beamforming to work. Sadly, while it should have worked well in theory, in practice it did not perform as expected. Appendix A goes into more detail about the idea behind Tx beamforming, and while the experiment didn't go as planned, a lot can be learned from this approach to the problem.



## Chapter 2

# Background & Related Work

In this chapter, the background and related work are discussed. In Section 2.1, fundamental concepts are explained which are used in this thesis. In Section 2.2 the building blocks are laid out, on which this project has been build.

### 2.1 Background

#### 2.1.1 Vital sign detection methods

In this section, a short summary is given about the most common vital signs monitoring systems, and the pros and cons of each method.

##### **Pulse oximeter**

The most well known method for measuring heart rate is using a pulse oximeter. In *Pulse oximetry: principles and limitations*[15], the principle of this way of measuring heart rate and pulse oximetry is explained. For this project, only the heart rate is used. This sensor is attached to the finger (the finger is most common, but this could be any body part) of a patient. In essence, the sensor consist of two parts: a light emitter and a light detector. The light emitter is located on one side of the finger and emits light, in most cases red LED light and infrared LED light. This light travels through the finger and is received by the light detector. The light detector measures the light intensity which ends up at the other side of the finger. A schematic view can be found in Figure 2.1.

This type of sensor works because of a property of hemoglobin. Hemoglobin is located inside of our red blood cells, and stores oxygen to transport it trough the body. A property of hemoglobin is that it absorbs more infrared light if there is a lot of oxygen encapsulated in the hemoglobin. It absorbs more red light if there is a low amount of oxygen encapsulated inside the hemoglobin. This makes it possible to measure the amount of oxygen inside of the body. But, this principle also works in the more general way. Every time the heart beats, blood is pumped through the blood vessels. Because of this pumping behavior, there are constant changes in pressure and volume inside the blood vessels. When

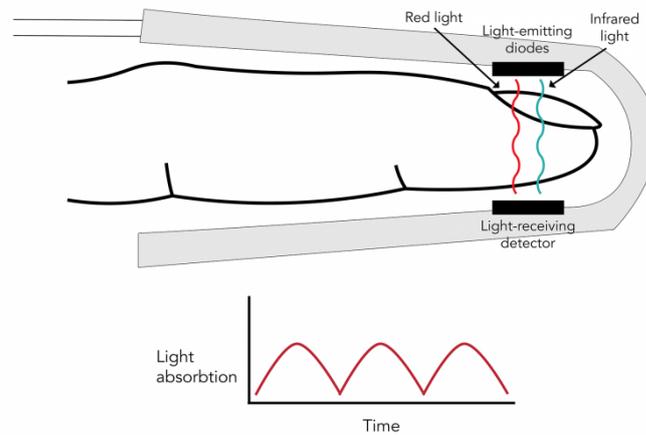


Figure 2.1: **Overview of the parts that make up a pulse oximeter. The different types of light travel through the finger from the emitter to the detector. From the amount of light arriving at the detector, the graph can be built. Image from [20].**

fresh blood is pumped through the blood vessel, there are more red blood cells present and more light, red and infrared, is absorbed. But that blood flows away and there is a period of less blood inside the vessels just before the next heartbeat. At that point in time, there is less light absorbed and more light detected by the detector. Using this principle, the heart rate can be detected. So, the absorption difference of the red and infrared light says something about the amount of oxygen in the blood, and the changes in absorption of both red and infrared light says something about the heart rate.

This method of measuring heart rate of a person is around for more than 40 years, so this technique is very reliable. The method used is very simple, just a LED and a light detector. This also means that privacy is not a concern, it is only made to detect heart rate and oxygen saturation, but it measures nothing else. There are however also some downsides on using an pulse oximeter. Because it uses a light detector, it could be influenced by the light in the room. Also, this method is less accurate for people with a darker skin color [7]. But the downside which this thesis focuses on is the patient contact. This method only works by having contact with the skin, and having a wire from the sensor on the patient to the device which is processing the data. This could be the cause for discomfort, the patients mobility is decreased because of the wire, and the sensor could fall off.

In more recent research [13], the pulse oximeter sensor can also be used to estimate the breathing rate of a patient. When breathing, small changes in pressure inside the thorax during a respiratory cycle change the amplitude of the blood pulsations that the pulse oximeter can detect.

### ECG monitoring

Electrocardiogram (ECG) monitoring is a more advanced way of monitoring someones heart rate [4]. The machine detects and monitors the electrical pulses

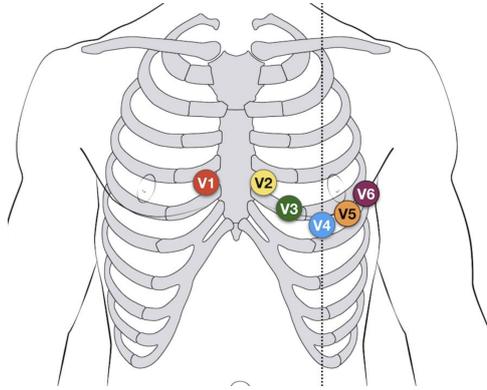


Figure 2.2: **Placements of the ECG pads on the chest region.** These pads can detect the electricity generated by the heart muscles. Image from [10].

that the heart muscles generate while contracting and relaxing. The detection happens by placing stickers on the chest in the heart region, depicted in Figure 2.2. These stickers detect the electrical signals generated by the heart muscles. Not only the heart rate can be monitored, but also the order of muscle contraction and the muscle power can be determined, which is beyond the scope of this project. ECG's can also introduce a small current and measure the change in impedance when a person is breathing. Because of the increase of air, the resistance of the tissue increases [6]. This is how the respiratory rate can be measured.

ECG monitoring is the golden standard in heart rate measurement. This is because of the extensive amounts of data that can be gathered. Doctors heavily rely on ECG's to monitor the heart and diagnose heart problems. There are however some cons that could be considered for this method. ECG's can be used for very preterm neonates, but it's difficult to apply the pads in the correct place. Also, the skin could be damaged by removing the pads. ECG's can also not be used for burn victims, because a clean connection to the skin is necessary for the ECG to work. Also, ECG's are relatively expensive compared to a mmWave radar. In a normal clinic there is only one ECG device which is shared with multiple patients.

### **Vital signs monitoring using radar**

Radar assisted vital signs monitoring is a relatively new research topic. A lot of research is still going on and real world implementations are still very scarce. In Section 2.2 an overview of the research related to this thesis project is given.

#### **2.1.2 mmWave technology**

This section goes into more detail about the type of radar which is being used in this project.

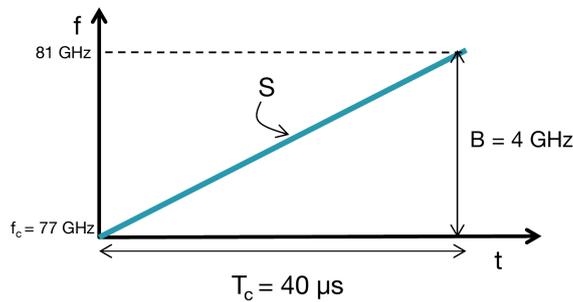


Figure 2.3: Example frequency/time graph of a chirp. The different variables of the chirp are labeled. Image from [5].

### FMCW millimeter wave radar

FMCW radar has been chosen over other types of radar because you can break range information into range bins. These range bins allow you to focus on one range of interest and ignore the noise of the environment around the person or object. The same technique can be applied in the azimuth direction. This functionality is perfect for this project, because the sensor can zoom in on the radar data from the chest region of the test subject, and ignore all other data.

This project uses millimeter wave radar which utilizes the FMCW modulation technique. Like the name implies, the frequency of these waves is in the range of 30-300 GHz, which means that the wavelength is in millimeters. This property makes these kind of waves a perfect fit for the detection of micro-movements.

The fundamental concept in radar systems is the transmission of an electromagnetic signal, which is reflected on a target and then received by the radar again. The basic concept of the FMCW modulation technique is that it sweeps frequency of the signal linearly over time. One frequency sweep is also called a chirp.

See Figure 2.5 for an example chirp waveform, with on the y-axis the current and on the x-axis time. Chirps can be defined by three parameters: frequency ( $f_c$ ), bandwidth ( $B$ ) and duration ( $T_c$ ). A chirp has also a slope ( $S$ ), which captures the rate of change of frequency. These variables can also be found in Figure 2.3.

These parameters and more can be programmed in such a way that the radar waves which are transmitted and received are tuned to the needs of the application. There are three building blocks involved, which can all be defined on their own and connected to each other. See for a visual representation Figure 2.4.

- **Frame:** the frame is the most high-level building block, and contains one or multiple chirps. The frame also defines how many times the chirps need to be repeated in each frame, how many frames of this type the chip should execute and the frame periodicity.
- **Chirp:** a frame can contain multiple chirps, but a chirp can only contain one profile. A chirp itself only contains a selector for which antennae to use, the rest of the properties of a chirp are defined in a profile. This makes it easier to use the same profile in multiple different chirps.

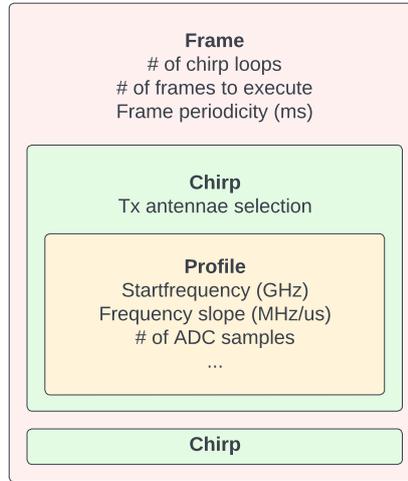


Figure 2.4: **The relations between a frame, chirp and profile, along with the most important properties of each element. See [19] for detailed information.**

- **Profile:** the profile contains parameters to define the chirp shape, such as the start frequency of the chirp, the chirp slope but also how many ADC samples need to be taken during the execution of the chirp.

Many parameters are available to modify the chirp for specific applications, for the multiple person vital signs application, only these variables are important: range and bin resolution. Range is the distance from the radar to the target, measured along the line of sight. For this application both the minimum and the maximum range is important. Bin resolution is the size of a bin, generally reported in hertz but can be converted to meters. A general bin resolution is a few millimeters for high frequency FMCW radars. Bin resolution should not be confused with resolution. A balance needs to be found between range and bin resolution. The signals from the sensor can reach up to 150 meters away, but the bin resolution will drop to around 4 centimeters. When the sensor is setup to have a range of 2 meters, the bin resolution can be as low as a couple of millimeters. The balance between maximum range and bin resolution is important for the multiple people tracking functionality. In Section 4.2 is explained how the sub-millimeter accuracy from the sensor is reached to measure vital signs.

### **IWR6843ISK Evaluation Module**

The type of radar used in this project is a frequency modulated continuous wave (FMCW) radar in the millimeter wave range, specifically operating in the 60-65 GHz range for this project. The radar module used in this project is the Texas Instruments IWR6843ISK, see Figure 2.6. It uses short-wavelength electromagnetic waves with a frequency between 60-81 GHz, which means that it is a free to use frequency band in the Netherlands [11]. The TI package that is

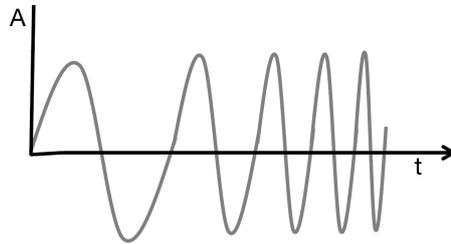


Figure 2.5: Chirp signal, with the current as a function over time. Image from [5].

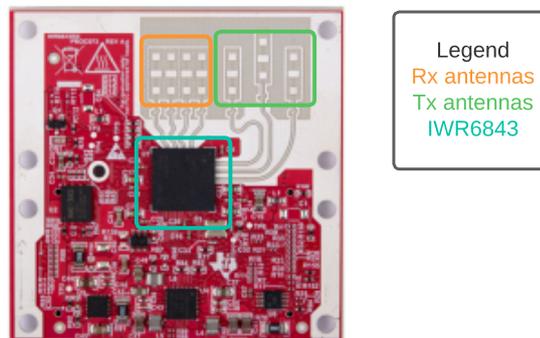


Figure 2.6: The Texas Instruments IWR6843ISK. This module is used throughout the project. See the legend for the most important functional components. Image from [17].

used in this project is an all-in-one solution. The most important part is the chip itself, which is the IWR6843. It also contains the Tx and Rx analog circuitry and antennae, see again Figure 2.6, to generate and capture the RF signals. The IWR6843 contains analog-to-digital converters (ADCs), micro-controllers (MCUs) and digital signal processors (DSPs) to make signal processing possible all on the chip. It also has some custom hardware accelerators build in to speed up standard radar calculations, such as the generation of fast Fourier transforms (FFTs).

A clear difference needs to be made between IWR6843 and IWR6843ISK. The IWR6843 is the chip from TI which is placed on the module. The IWR6843ISK is the module itself, with all of the supporting circuitry and antennae. In Figure 2.7, the inner components from the IWR6843 can be found. The left and middle section from the Figure are programmed by TI, and the behavior can only be modified by setting parameters. The rightmost column is where the programming is done for this project. The two processors are programmed to work together to create from the radar data as an input the vital signs data as output. More details about the embedded programming can be found in

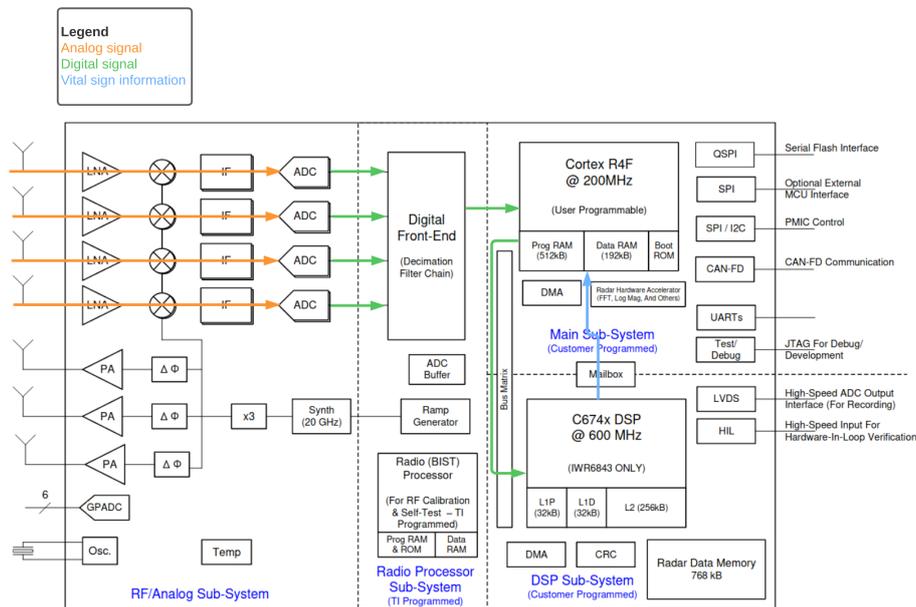


Figure 2.7: The inner components of the IWR6843. The path the signal takes through the chip is displayed. Image from [17].

Chapter 5. The bottom three antennae in the Figure are the transmitters. The signal is generated by the RF front-end and transmitted by two of the three transmitting antennae. Only the two outermost antennae are used because the middle antenna is for scanning in one additional dimension. This is why that antenna is slightly elevated, see Figure 2.6. This project only needs 2D radar images, this is why the middle transmitting antenna is not used.

The IWR6843ISK module contains the IWR6843 chip itself, but also the specific antenna configuration, voltage regulators, additional EEPROM chips and temperature sensors, since the chip can get quite hot during operation. See also Figure 2.8 for a functional diagram of the IWR6843ISK module. The IWR6843ISK has been chosen because it's an all in one embedded solution, allowing for real-time analysis and display.

### 2.1.3 Signal processing

This section goes into more detail about the signal processing steps required to extract usable data (for example a heatmap) from raw radar data. Three different types of FFTs will be explained. From the raw data, a range-FFT can be generated. From the range-FFT a doppler-FFT can be generated, and using the data from the doppler-FFT, the angle-FFT can be generated.

#### Range estimation

In Figure 2.9, a scheme of the inner workings of the FMCW chirp can be found. First, a chirp is generated using a ramp generator and then mixed with a syn-

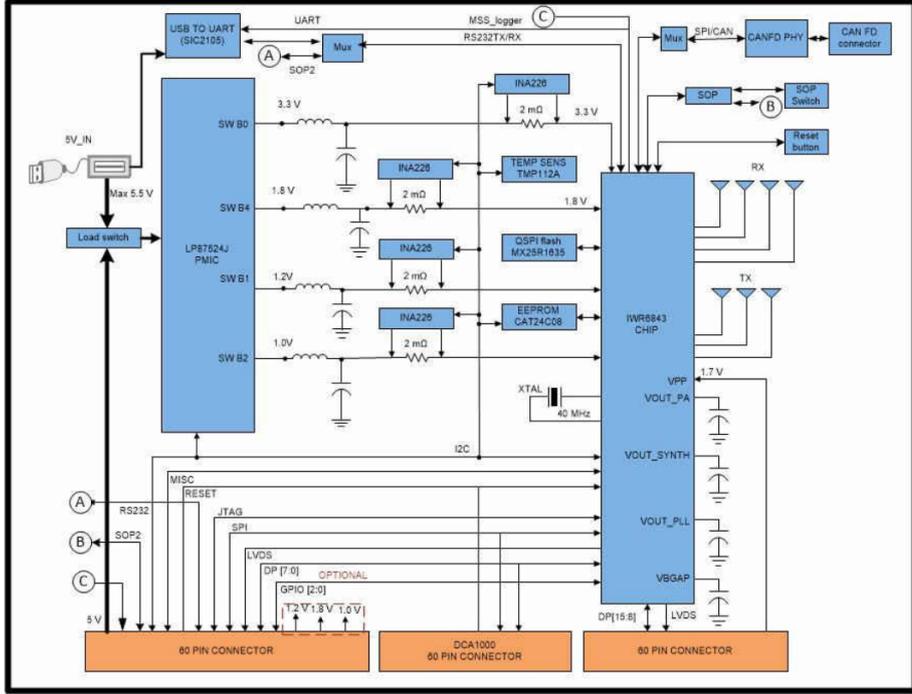


Figure 2.8: **Functional diagram of how the IWR6843 is fitted in the IWR6843ISK module. Image from [17].**

thesizer (1) to generate a 20 GHz signal, which is then multiplied to get a 60 GHz signal. This chirp is then transmitted over one or multiple Tx antennae (2). When reflected on an object, the signal returns and is captured using the Rx antennae (3). The transmitting and receiving of the signal introduces a time delay. The further the object is away from the sensor, the larger the time delay. The signal from the synth and the received signal from the Rx antenna are mixed (4). This mixer has two signals as input, and one signal as an output. The mixer outputs a signal with a frequency which represents the time delay in the frequency domain. This output signal is called the intermediate frequency (IF) signal. In Figure 2.10 the situation is displayed more visually. Using this information we can use Eq. 2.1.

$$\tau = \frac{2d}{c} \tag{2.1}$$

Where  $\tau$  is the time delay and  $c$  is the speed of light, to compute the distance to the object  $d$ .

This is an example where only one object is used. But in the real world, there are almost always more objects to be detected. All of these objects return another reflection, which means that one transmitted signal from the Tx antenna can result in multiple reflections reaching the Rx antenna, as depicted in Figure 2.11. In this instance, there are multiple IF tones at once, one for each reflected object. To differentiate between those, we make use of a fast Fourier transform. After processing this fast Fourier transform, it results in a

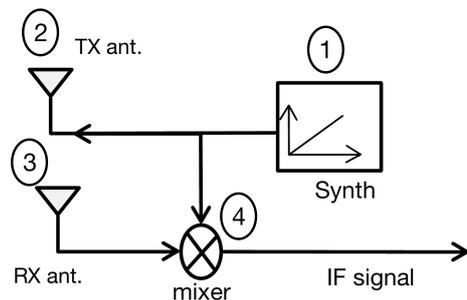


Figure 2.9: **Scheme of the inner workings of the FMCW chip. Image from [5].**

frequency spectrum in which each peak will point to one separate IF frequency. The IWR6843 has a hardware accelerator to speed up the fast Fourier transform generation.

### Doppler estimation

The range information, calculated using the range-FFT, is also known as the 1D-FFT. Most of the time, this is the first processing step for radar data. But, there is more information we can extract from the radar data. Sometimes it is useful to know the velocity of objects. The movement of the chest which is observed for this project is also a velocity. This velocity can be calculated by generating a doppler-FFT. The doppler effect is the change in frequency of a wave in relation to an observer who is moving relative to the wave source.

To calculate the velocity of objects, we need the data from the range-FFT. The radar transmits two chirps, spaced  $T_c$  apart from each other. Both of these chirps get reflected on the object and for both of them a range-FFT is calculated. Because the chirps are transmitted almost immediately after each other, the two FFTs have peaks in the exact same location. What matters now is the phase of the peaks. The output of a FFT is an array of complex numbers. The real part says something about the power (strength) of the signal, the imaginary part says something about the phase of the sinusoidal signal. The phase between the two chirps is changed slightly because even between those two chirps spaced  $T_c$  apart, the object moved. The velocity can be derived using Eq. 2.2, where  $\Delta\phi$  is the phase change.

$$v = \frac{\lambda\Delta\phi}{4\pi T_c} \quad (2.2)$$

Because the velocity is dependent on the phase value, the velocity could wrap around if the object is moving too fast. Therefore, the maximum speed the radar is able to record before the measurements start to be unreliable, can be calculated using the formula in Eq 2.3. Phase unwrapping can be used to deal with this issue.

$$v_{max} = \frac{\lambda}{4T_c} \quad (2.3)$$

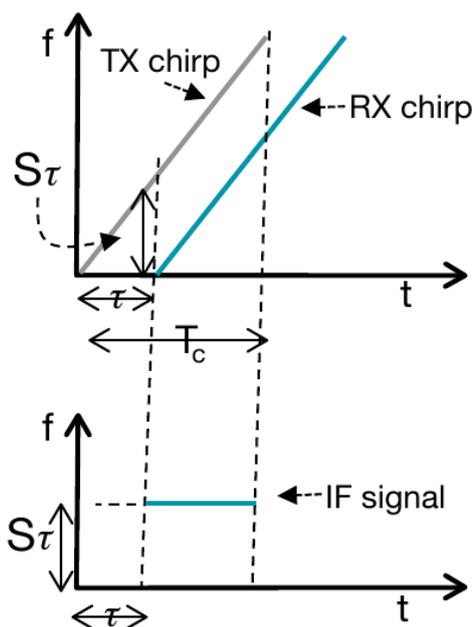


Figure 2.10: The calculation of the IF signal. In the top graph, the Tx and Rx chirp are plotted. In the bottom graph, the IF signal is plotted. Image from [5].

To determine the velocities of multiple objects, a FFT is used to differentiate between multiple objects. In the example above for only one object, two chirps are needed to calculate the velocity of the object. To distinguish between multiple objects, more chirps are needed. Each chirp is one input data element for the doppler-FFT, so the more chirps, the more accurate the result. All of those chirps together forming the data for the range-FFT and the doppler-FFT are called a frame.

### Angle estimation

Using the methods above, we can estimate the range and the velocity of multiple objects using range- and doppler-FFTs. But, there is still one more shortcoming of these calculations. They only work for one dimension. This means that if object 1 is 3 meters from the sensor and object 2 is 5 meters from the sensor, the objects can be differentiated. But, if two objects are at the same distance to the sensor but with a different angle, these two objects will be in the same range bin and cannot be differentiated. The solution to this problem is to try and calculate the angle of the object with respect to the sensor, also known as the Angle of Arrival (AoA). This way, the data collection from the sensor can be expanded from one dimension into two dimensions.

To calculate the Angle of Arrival, multiple Rx antennas are needed. In Figure 2.12 the functional diagram of how angle estimation works is displayed. A chirp is transmitted from a Tx antenna. The chirp is reflected on an object and received by the first Rx antenna. The distance between the object and the Rx

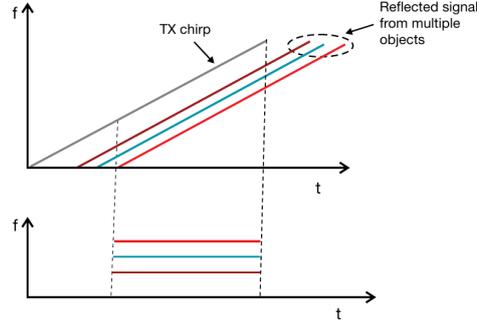


Figure 2.11: One Tx signal can result in multiple reflections on the Rx antennas from multiple objects in front of the sensor. Image from [5].

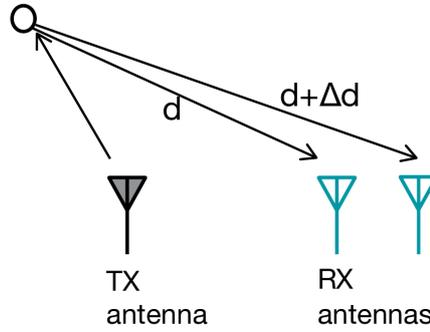


Figure 2.12: One Tx signal is received by multiple Rx antennas. Because the signal travels  $\Delta d$  longer to one Rx antenna compared to the other, the angle of arrival can be calculated. Image from [5].

antenna is  $d$ . When the signal arrives at the second Rx antenna, it has traveled  $d + \Delta d$ . Because of this  $\Delta d$ , the phase on the second Rx antenna is different from the first Rx antenna. A similar formula as the doppler estimation can be formed to tie these concepts together. This formula can be found in Eq. 2.4.

$$\Delta\phi = \frac{2\pi\Delta d}{\lambda} \quad (2.4)$$

Using some basic geometry and Eq. 2.4, we can deduce the formula for the angle of arrival, which can be found in Eq. 2.5.

$$\theta = \arcsin \frac{\lambda\Delta\phi}{2\pi l} \quad (2.5)$$

Using multiple Rx antennas, in the case of this project using the IWR6843ISK there are 4 antennas, we can calculate the AoA of different objects using a FFT. At this point, a 2D heatmap can be constructed using the calculations from the different estimations above. This heatmap is the starting point for this project.

## 2.2 Related Work

Using radar to estimate vital signs has already been researched before as shown by Changzhi Li et. al. [9] in 2009. This paper uses a transmitter to send a unmodulated signal to a patient and receive the modulated signal back. The signal gets demodulated and the heart rate and breathing rate can be extracted. In this attempt, the system is dependent on analog hardware to make the vital signs extraction work.

Zhicheng Yang et. al. [22] attempts something similar using two separate antennas, one for transmitting and one for receiving. The transmitting antenna does a first sweep through the room by physically moving the antenna to detect persons in the room. After that, the transmitter and receiver can zone in on the person and detect the vital signs of that person. The transmitting and receiving is done with large pieces of equipment.

In '*Remote monitoring of human vital signs using mm-wave fmcw radar.*' [3], the step to FMCW radars can be observed. The authors also use a chip from TI, but an earlier model (the IWR1440). They make use of a range-FFT to detect a person, and perform vital signs estimation on that person. A disadvantage of this approach is that the solution is only able to estimate the vital signs of one person. Another disadvantage is that the data gathered by the sensor is send to a computer for processing at a later time. This means that the vital signs estimation is not real time.

The paper which is closest to the research topic of this project, is from Adeel Ahmad et. al.[1]. The authors from this paper all work for TI, and made a proof of concept solution for multiple person vital signs estimation using a TI sensor. The sensor first looks for persons in the 2D space before the sensor. Then, it uses a post-processing beamforming technique to isolate the data for one person at a time to extract the vital signs. It does this for all of the persons within reach of the sensor. Again, there are some disadvantages with this method. Firstly, the beamforming technique they are using to isolate the different persons from each other is not shown in detail. The most important disadvantage is that the calculations are done in MATLAB (The Mathworks Inc., Massachusetts, USA) after the data gathering. First, the data is recorded using the DCA1000EVM. Then this data is transferred to a computer and after that the data is processed in post using MATLAB. In the medical world, it is important to have information about a patient right away, and not minutes to hours after they have been recorded. This also means that there is no feedback loop possible to tune the radar data on the fly. When using post-processing, the data is permanent. When using real-time data, the radar parameters can be tuned on the fly to improve the signal quality.

## Chapter 3

# People detection

The first step in dynamic multiple people vital signs estimation, is detecting and tracking persons in range of the sensor. In this chapter, the algorithm to detect persons is explained.

### 3.1 Overview

This chapter goes into more detail about detecting people using the radar data. But why is it needed to detect persons in the first place? Every time a measurement has been taken, this measurement consists of the whole field of view of the radar sensor. Apart from the reflections which bounce back from the people in range of the sensor, there are also other reflections, like reflections from static objects like walls, plants and other pieces of furniture in the room. The aim of this chapter is to isolate the persons in the radar data from all of the other data. Then the exact location from the persons is established inside of the radar data, this location can then be used in Chapter 4, where the actual vital signs estimation is explained.

### 3.2 People detection algorithm

For this people detection algorithm to work, a heatmap generated from the radar data is needed. This heatmap is a 2 dimensional array of bins. A bin is a 2D area in front of the sensor for which the signal strength is measured. The value in each bin denotes the amount of energy reflected in that location. If there is a high amount of energy that gets reflected back to the sensor, there is a high probability of an object present at that location.

The goal of this project is vital signs estimation on multiple people. Because this technology is in very early stages of development, it was opted to have a very low noise environment for the radar data gathering. Static noise will be filtered out, this will be explained in Section 3.4. So apart from the persons present in the view of the sensor, there will be no other sources of reflection present. This means that a fairly basic peak detection algorithm would suffice for this implementation. First, all of the peaks in the heatmap are found, then the peaks close to each other are grouped together. When these processing steps are done, each detected peak is the center location of a person present in the

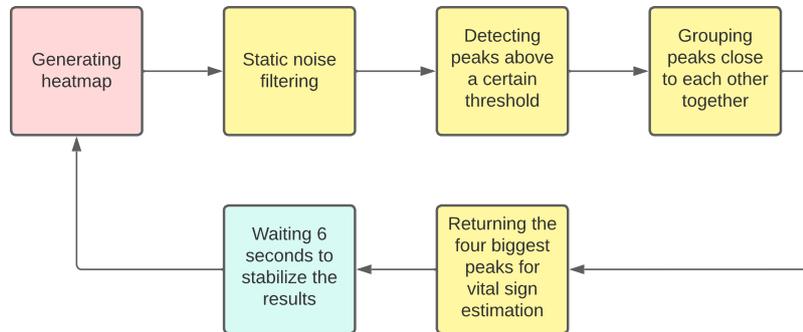


Figure 3.1: **Block diagram of the processing steps involved in the people detection algorithm.**

field of view of the radar sensor. See also Figure 3.1 for a block diagram of this person detection algorithm.

### 3.3 Existing solutions

There are existing algorithms to group together multiple points with high energy which are close to each other. These groups of points could be an object, a person or any other thing in the space which reflects the RF waves from the sensor. One example is a Constant False Alarm Rate (CFAR) algorithm. This algorithm is already implemented in the mmWave SDK of Texas Instruments [19]. This is a very difficult algorithm to set up however, it requires intricate knowledge about the inner workings of the algorithm and of all of the right data streams in the chip. Documentation on how to implement this algorithm was missing, only a working general version is provided. Because of these reasons, a detection algorithm was build from scratch. Because it was build from the ground up, it was ensured that all of the needs and output formats could be met.

### 3.4 Implementation

This section goes into more detail about the actual implementation of the algorithm on the chip. In each subsection, one step of the algorithm is explained. In Listing 3.1, the chirp parameters used in this project can be observed. The exact meaning of each parameter can be found in the mmWave SDK [19].

```

1 ...
2 profileCfg 0 60 250 10 40 0 0 98 1 64 2200 0 0 40
3 frameCfg 0 1 168 0 250 1 0
4 chirpCfg 0 0 0 0 0 0 0 1
5 chirpCfg 1 1 0 0 0 0 0 4
6 ...
  
```

Listing 3.1: **Portion of the parameters file which gets send to the IWR6843 to set it up.**

### 3.4.1 Input data format

The algorithm returns a list of coordinates with a maximum of 4 persons. It makes use of the build-in range-FFT, doppler-FFT and angle-FFT implementations, described in Section 2.1.3 to Section 2.1.3. Using these methods, a 2 dimensional heatmap can be generated. On the x-axis are all of the bins in the azimuth (angle) direction and on the y-axis are all of the bins in the range direction. So to recap the information found in Section 2.1, the data in the range direction comes from the range-FFT, and the information in the azimuth direction comes from the angle-FFT. Each bin contains information about the energy level in that bin. In other words, how much signal reflection has been observed in that bin location.

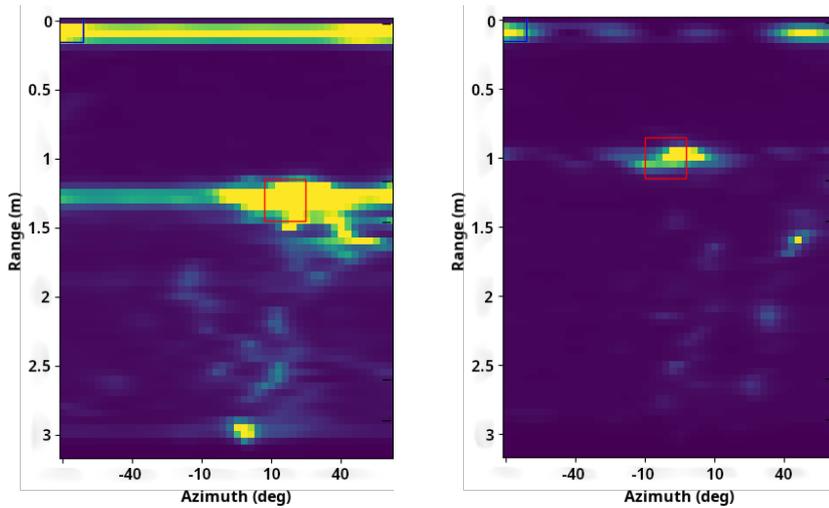
### 3.4.2 Noise removal

Before persons can be detected, the amount of background noise needs to be removed. Because with less noise, the persons in front of the sensor will stand out more and will be easier to detect. Also, the signal strength will be improved. For this project the testing is done in one location, so it can be assumed that the sensor will be in a static position and there will be no background change. For each new background, a new calibration round needs to be done. During this calibration, the sensor will scan the space in front of it. It is very important that no other persons or objects other than static ones are at that moment in view of the sensor. The sensor will perform 64 scans and take the average of those scans. This noise map will both be saved on the device and returned to the computer via UART. In this way, the computer can send the values along with all of the other parameters when the sensor is restarted, and the calibration round doesn't need to be run again. When the calibration data is in place, it can be used to remove the noise from each new frame that is coming in. This noise algorithm also corrects for whole rows being over saturated. Those rows can be seen clearly in Figure 3.2a. The difference between noise correction and no noise correction can be observed in Figure 3.2.

### 3.4.3 Peak detection

This algorithm works by detecting peaks. During the testing phase, it was determined what the minimal peak height is for a person in the frame. The maximum heat in one heatmap bin is 32000. The minimal peak generated by a person in a frame is 2000. These numbers have no unit, it is just a measure for signal strength were 0 is no signal and 32000 is the maximum observable signal. This is the first filter, only the peaks which are higher than the threshold are considered.

For this application, it is assumed that each person is sitting approximately half a meter to one meter apart, this makes it easier to detect individual persons. This translates in the grouping functionality of the person detection algorithm. For each new peak that is found, it is checked if there has been another peak found within half a meter of the new peak. If that's the case, the peak with the largest magnitude will be added to the list. This results in finding the biggest peak for each person, provided they sit half a meter apart. The biggest peak gives the information with the highest SNR value, such that the following



(a) Heatmap without any noise removal. (b) Heatmap with noise removal implemented.

Figure 3.2: These two heatmaps show the difference between noise removal and no noise removal.

algorithms will work in the most optimal form. The return value is an array with the coordinates of one or more bins, one bin per person.

This algorithm is executed once every 64 frames. Since the program processes around 10 frames per second, every 6.4 seconds there will be another person detection round performed. The algorithm is not run every frame because it takes relatively much time to compute, and the whole system is more stable in this way. In the future, it would be best to implement a sliding window and take the average of multiple frames to base the tracking on. For this project that could not be implemented, because there was not enough space on the chip and there was no time left for additional calculations. This could be solved by software optimizations.

### 3.4.4 Result

In Algorithm 1, an overview of the algorithm in pseudo code can be found. Figure 3.3 shows the heatmap result if two persons are sitting in front of the sensor. The red and yellow boxes are drawn exactly in the middle the peak that is detected. These peaks will be used internally for vital signs estimation, this is purely a visualization for the user. In Figure 3.4, four persons are detected at the same time. This is possible, but it becomes really crowded in this relatively confined space. The most optimal way for the users will be measuring one or two persons at the same time, which will also become apparent later in this thesis.

---

**Algorithm 1** Person finding algorithm

---

**Require:** *heatmap*

*maxPeaks*  $\leftarrow$  *list*()

**for**  $x = 0, 1, \dots, azimuthLength$  **do**

**for**  $y = 0, 1, \dots, rangeLength$  **do**

$bin \leftarrow heatmap(x, y)$

**if**  $bin > 2000$  **then**

**if**  $bin$  heat is larger than surrounding bins **then**

**if**  $bin$   $c$  in *heatmap* with distance  $< 0.5$  meter **then**

**if**  $bin > c$  **then**

            Add  $bin$  to *maxPeaks*

            Remove  $c$  from *maxPeaks*

**end if**

**else**

          Add  $bin$  to *maxPeaks*

**end if**

**end if**

**end for**

**end for**

---

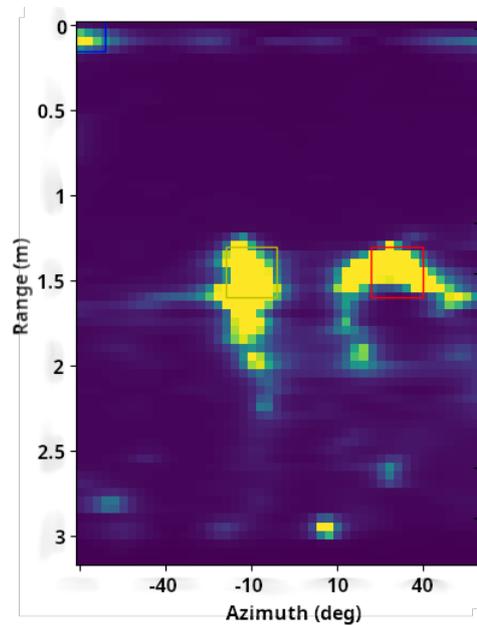


Figure 3.3: Two persons at the same time detected. Each detected person gets a box drawn around it with a color.

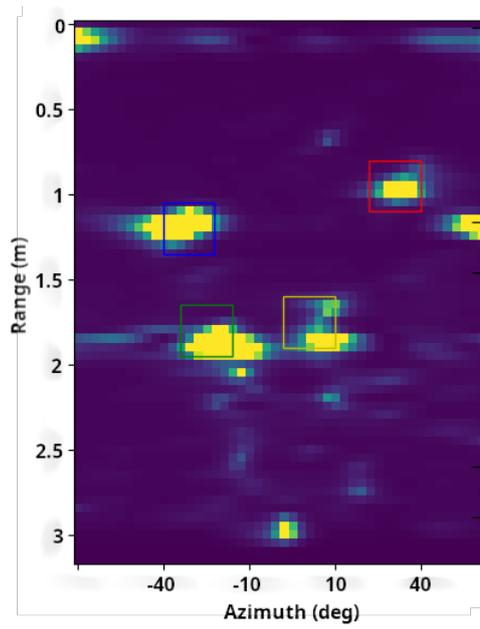


Figure 3.4: Four persons at the same time detected. Each detected person gets a box drawn around it with a color.

## Chapter 4

# Extracting Vital Signs

In this chapter, all of the signal processing techniques are explained to get from a selected bin in a heatmap to a heart rate and a respiration rate.

### 4.1 Phase signal

For accurate vital signs extraction and waveform analysis, the phase of the radar signal is needed. This section explains why the phase signal is needed and how it can be extracted.

#### 4.1.1 Radar parameters

Like mentioned in Section 2.1.2, every time the IWR6843 is restarted, a lot of parameters get send from the computer to the chip to properly setup different parts of the chip. An important part of these parameters are the chirp designs. These parameters among others determine the length of the chirp, the frequency range of the chirp and how many chirps are in one frame. The part of the parameters file which sets up the chirps and the frames can be found in Listing 3.1. To better understand how we can extract the vital signs from the radar data, we need to know what the resolution of the data is.

The exact meaning of these parameters can be found in the mmWave SDK documentation [19]. The ones important for this section are summarized in Table 4.1.

To determine the range resolution for the parameters used in this project, we make use of the formula provided by TI [16], as seen in Eq. 4.1.

$$d_{res} = \frac{c}{2B} \quad (4.1)$$

Where  $B$  is the bandwidth of the chirp in Hz, and  $c$  is the speed of light in m/s.  $B$  can be calculated by multiplying the slope of the chirp and the duration of the chirp, see Eq. 4.2.

$$B = S \times T_c = 98 \times 40 = 3920MHz \quad (4.2)$$

Using Eq. 4.1, it can be calculated that the resolution from the sensor is approximately 3.8 centimeters. The maximum range can also be calculated, by using Eq. 4.3.

Parameter	Value
Ramp End Time (us)	40
Frequency slope (MHz/us)	98
Start Frequency (GHz)	60
Number of ADC samples	64

Table 4.1: Some of the parameters from Listing 3.1 with their value.

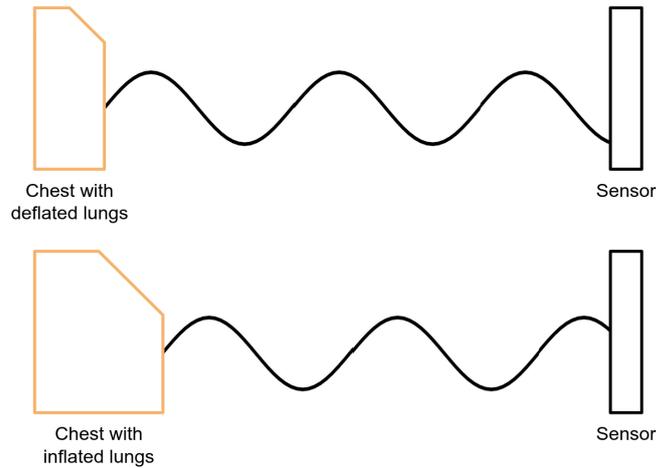


Figure 4.1: Exaggerated scheme on how the phase of the signal changes depending on the inflated or deflated chest of a person.

$$d_{max} = \frac{cA}{2B} \quad (4.3)$$

Where  $A$  is the number of ADC samples which are taken for each chirp. Using this formula, the maximum range that the sensor can reach using these parameters is 2.45 meters.

Bin resolution is too large for vital signs estimation. If a person is breathing, the chest only moves one millimeter or less, while the bin resolution is multiple millimeters. The heartbeats are even more fine grained, the movement of the chest for one heartbeat could even be fractions of millimeters. Therefore, another technique must be used to reach the resolution needed for vital signs estimation. When the radar data is processed into a FFT, the output is an array of complex numbers. This complex number not only says something about the magnitude of the input data, but also something about the phase of the input data. See also Figure 4.1. This phase is how the reflected waveform returns to the sensor. When nothing happens to the chest, this phase stays the same. But if the heart is in the middle of a beat, the chest is tens of millimeters closer to the sensor. These tens of millimeters result in a phase shift. This is data that can be used to estimate vital signs.

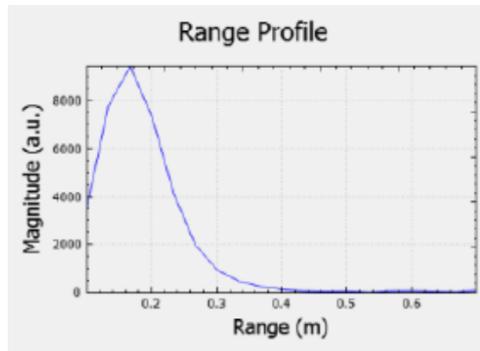


Figure 4.2: **Person detection using only the range-FFT. This screenshot has been taken from a TI demo.**

### 4.1.2 Phase unwrapping

Most existing solutions, like [9, 22, 3] make use of phase extraction like in Figure 4.3. From the radar data the range-FFT is calculated, on this FFT the peak is found. This peak denotes the location of the person, see also Figure 4.2. From that peak the phase is extracted and saved in a sliding window configuration. This same event repeats every 50 milliseconds. So, for each 50 milliseconds, the phase is extracted. All of these phase values form a waveform on their own, which can function as the input for the vital signs estimation.

In the case of this project, the phase must be extracted from the generated heatmap. However, this means that the radar data is put through two more FFTs other than the range-FFT, namely the doppler-FFT and the azimuth-FFT. There were some concerns that the phase information would be lost during the generation of those two other FFTs, but during testing, it became apparent that the vital signs waveform could also be extracted from the phase of one bin in the heatmap. This phase information is retained through the generation of the doppler-FFT and the angle-FFT. Because this is possible, information from a 2 dimensional space can be taken and processed to extract the vital signs. This is the second very important step that needs to be taken to arrive at the multiple person vital signs tracking goal. The first step is to track the persons in the radar view. This second step is to extract and unwrap the separate phase waveforms for each person. The next step is to take these waveforms and extract the vital signs data out of this waveform.

## 4.2 Vital signs estimation

Before starting with this section, I want to emphasize that I have copied these vital sign estimation implementations to a large extent from the TI Vital signs estimation demo [18] and the Matlab implementation from my mentor Caitlin Ramsey [14]. The main alterations I made was to make the code usable for multiple people at the same time. I chose for this approach because I am not very familiar with signal processing given my Embedded Systems background, so it is difficult to innovate in that field. Secondly, the main bulk of code is already implemented for this embedded platform. I think it is important to

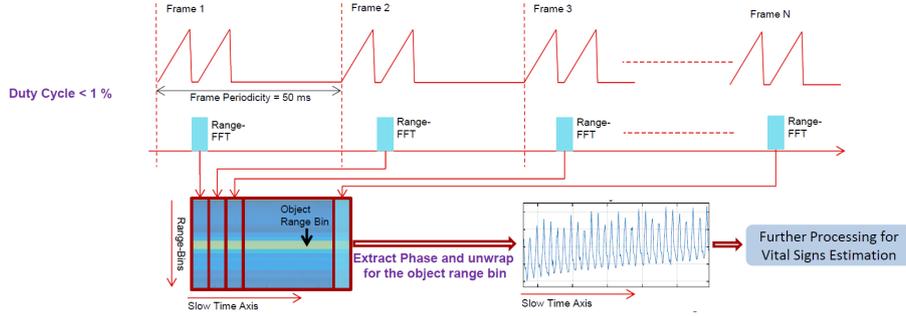


Figure 4.3: **Visual representation of how the phase gets unwrapped. Image has been taken from [18].**

explain the inner workings of this vital signs estimation implementation in this thesis, because it has a great significance to the project.

#### 4.2.1 Data preparation

The phase values coming from the heatmap are stored in circular buffers for further processing. But first, the data needs to be prepared. The algorithms work by using the phase changes. A phase change is computed by comparing the current phase by the previous phase value, like in Eq. 4.4.

$$P_{change} = P_{current} - P_{previous} \quad (4.4)$$

This results in an array with phase changes which can be used in processing. But first, the signal is cleaned up by an impulse noise filter. The data which is coming from the sensor is raw data. There can be unwanted spikes in the data, for example due to electromagnetic interference. This noise needs to be filtered out beforehand. This is done in a sliding window configuration, where the three most recent values are stored. If the ratio between the first and the second, and the ratio between the third and the second are above a certain threshold, the second value will become an interpolation from the first and third value. Figure 4.4 has a more clear visual representation. This filter makes sure that there are no sudden peaks in the data, but that the signal is still able to change shape. Now that the signal is cleaned up, the next step is to differentiate between the heart rate and the respiration rate.

#### 4.2.2 Splitting the heart rate and respiration rate signal

To do separate analysis on the respiration rate and the heart rate, the source waveform needs to be separated into two separate waveforms. The way this is done, is by using an IIR filter.

IIR filters are one of two common digital signal processing (DSP) filter types [21]. IIR stands for Infinite Impulse Response. The other common type is the Finite Impulse Response filter. These types of filters modify the frequency content of a signal. FIR filters and IIR filters are both part of the Linear Time Invariant (LTI) filter group. LTI filters can modify the phase and the amplitude of certain frequencies of a signal. These filters are most commonly used to remove

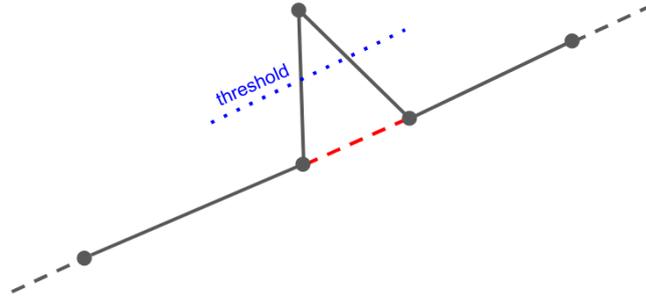


Figure 4.4: **Visual representation of how sudden spikes get filtered out of the phase waveform.**

undesired frequencies from an input signal. For example: the antenna of a car receives all of the radio stations at once, from 88 MHz to 108 MHz. To make it possible to only listen to a station at for example 89.9 MHz, a filter can be used to only allow the signal with a frequency of 89.9 MHz. The main difference between FIR and IIR filters, is that the IIR filters make use of a feedback mechanism. If the IIR filter is first zero and receives a 1 as an input, the filter output can (theoretically) never reach zero again. A FIR filter works with a certain amount of stages, and is able to reach zero after a certain amount of time. For this project an IIR filter is chosen because it is more efficient to implement for embedded devices, since it uses less memory and calculations compared to a similar FIR algorithm [8].

The IIR filter used in this project is an biquad cascade IIR algorithm. Biquad stands for bi-quadratic, which refers to the Z-domain, where the transfer function is a ratio of two quadratic functions:

$$H(z) = \frac{b_0 + b_1z^{-1} + b_2z^{-2}}{1 + a_1z^{-1} + a_2z^{-2}} \quad (4.5)$$

The filter used in this project can be written down as an equation:

$$y[n] = b_0x[n] + b_1x[n - 1] + b_2x[n - 2] - a_1y[n - 1] - a_2y[n - 2] \quad (4.6)$$

A more visual flow graph can be found in Figure 4.5. This algorithm is translated in C code. This filter is executed on the input signal two times, once for the respiration rate waveform and once for the heart rate waveform. The filter is therefore configured like a band-pass filter. It will remove frequencies which are not in the allowed frequency range. The frequency ranges that were used in this project can be found in Table 4.2.

So, there was one input waveform, which is now split into two different waveforms, one for heart rate and one for respiration rate. The next step is to concentrate on getting a beats per minute (BPM) reading out of this data.

### 4.2.3 Circular buffers

The data from these two separate waveforms are both stored in a circular buffer. This circular buffer acts like a sliding window. Based on the waveform in this

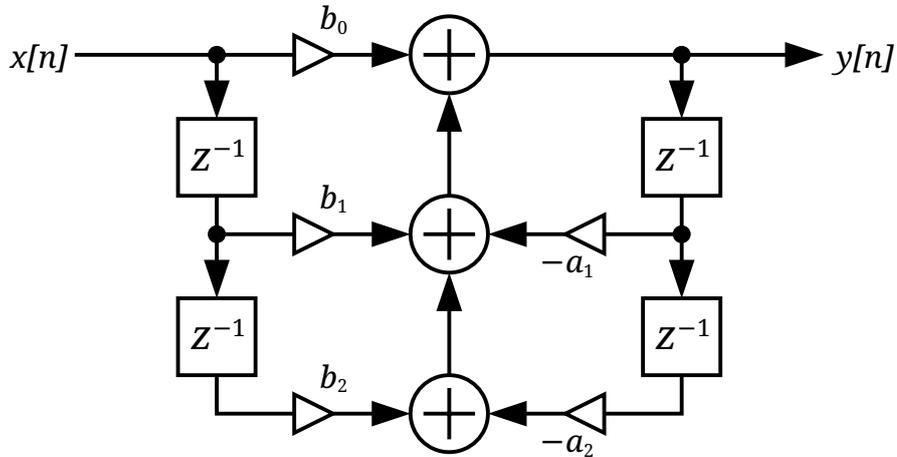


Figure 4.5: The flow graph for a biquad cascade IIR filter. A more visual representation from Eq. 4.6, how the input gets transformed to the output.

Parameter	Minimum frequency	Maximum frequency
Respiration rate	0.1 Hz	0.5 Hz
Heart rate	0.8 Hz	2.0 Hz

Table 4.2: Typical frequency ranges from the respiration rate and heart rate. These frequency ranges are used in this project to filter out the right waveforms to do analysis on.

buffer, the respiration rate and heart rate can be estimated. The length of these buffers are sent via the configuration file, as can be seen in Table 4.3. In this table, all of the vital sign parameters are listed, along with their value and a short explanation.

The buffer length is a balancing act. The more data is in the buffer, the more data can be input to the FFT and the more accurate the result will be. But, the longer the buffer is, the more data gets averaged together. This means that instantaneous changes in heart rate frequency will not be visible in the output data. This project makes use of a respiration waveform buffer of 256 data points. A measurement will be taken every 50ms, this is also the frame periodicity, as can be found in Listing 3.1. The total buffer length in seconds can also be calculated, as in Eq. 4.7.

$$256 * 50ms = 12800ms = 12.8seconds \quad (4.7)$$

Which means that a full buffer with respiration data will contain 12.8 seconds of data. Since the heart rate waveform has less amplitude because the chest doesn't move as much and the heart rate is faster than the respiration rate, the buffer is twice as long, resulting in a total buffer length of 25.6 seconds. This is done because we need more accuracy for the heart rate.

Because the heart rate is so fine grained, the heart rate waveform is added in

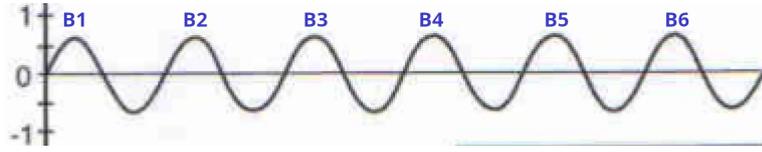


Figure 4.6: **This waveform could represent the respiratory waveform in the circular buffer. This waveform spans 12.8 seconds of data. 6 peaks can be detected, each peak is one breath.**

chunks to the buffer. If the chunk is too noisy, for example because the subject is moving a little bit, the chunk gets skipped. This mechanism makes sure that the waveform is as noise free as possible.

#### 4.2.4 Peak counting

The next step in the algorithm is the peak counting. Both waveforms are filtered, this means that for the respiration waveform there is only a peak when the chest is deflated and then inflated. For the heart rate waveform, there would be a peak for each beat of the heart. This peak counting is only a simple algorithm. For each three consecutive values, it checks if the middle value is bigger than the value on the left or on the right. If it is, a peak has been detected. An additional threshold filter is executed over the peaks to eliminate the invalid peaks.

Since it is known how many peaks there are in the waveform, and it is known how much time a full waveform buffer takes (see Section 4.2.3), the BPM could be calculated. For example, 28 peaks are detected. From Section 4.2.3 it is known that a full waveform buffer takes 25.6 seconds. With a simple calculation the estimated heart rate will be revealed, see Eq. 4.8.

$$\frac{28 * 60}{25.6} \approx 66BPM \quad (4.8)$$

For the respiration rate, Eq. 4.8 can be repeated in a similar way. Let's assume 6 peaks are detected in the respiration waveform, see also Figure 4.6. From Section 4.2.3 it is known that the full respiration rate buffer takes 12.8 seconds of data, so:

$$\frac{3 * 60}{12.8} \approx 13BPM(BreathsPerMinute) \quad (4.9)$$

The algorithm is completed.

#### 4.2.5 Executing FFTs on the waveforms

There is however another option that can be tried. The peak counting method from Section 4.2.4 is a bit rudimentary. Peaks are detected and filtered, but there could still be false positives or false negatives in this workflow. It would be best to execute a FFT on the waveform and extract the most present frequency from the FFT. Luckily, the chip is specialized in calculating FFTs. Each of the two waveforms is put through the FFT generation, and from each FFT output, the biggest peak is extracted. From this biggest peak, the BPM value can be

calculated in a similar way as in Eq. 4.8 and Eq. 4.9. Because this method is more advanced and less prone to errors, the outcome of the FFT method can in general be better trusted than the outcome of the peak counting method in Section 4.2.4. Both of these methods, peak counting and FFT generation, are applied at the data to end up with the vital signs. For the results of this project, the outcome from the FFT method is used.

Parameters	Values	Comments
Start range (meters)	0.3	The persons to be detected are expected to be within the start range and the end range from the sensor.
End range (meters)	0.9	
Respiration waveform size	256	Specifies the number of points within the waveforms. The more data points are in the circular buffer, the more data is captured and the more accurate the reading becomes. A downside of using a large waveform size is that it loses the ability to capture instantaneous changes.
heart rate waveform size	512	
Rx-antenna to process	4	Because there are 4 Rx antenna, this value could be max. 4. Because for this implementation the rangeAzimuth-FFT is used, only a value of 4 is allowed.
Alpha filter value for Respiration waveform energy computation	0.1	Alpha filter values for recursive averaging of the waveform energies based on the equation below where $x(n)$ is the current waveform value while $E(n)$ is the energy. $E(n) = \alpha x^2(n) + (1 - \alpha)E(n - 1)$
Alpha filter value for heart rate waveform energy computation	0.05	
Scale factor for the Respiration waveform	300000	Scaling factors to convert waveform values in floating points to 32-bit integers required for the FFT.
Scale factor for the heart rate waveform	300000	

Table 4.3: **Overview of all of the Vital Signs configuration parameters getting send to the chip. These parameters mostly control the sensitivity of the chip. For each parameter, the value used in this project is mentioned and a short explanation.**



## Chapter 5

# Real-time implementation

In this chapter, the functionality of the chip is explored, and the general program flow is discussed.

### 5.1 Overview

A big part of the innovation of this thesis is implementing all of the calculations on the chip itself. The chip comes equipped with a RF front-end, two processors and other hardware which must all work in unison to do all of the calculations in time, see also Figure 2.7 for the block diagram of the chip. Before the coding can start, the programmer needs to have a good knowledge of how all of the elements in the chip need to work together.

### 5.2 How it works

In a clinical environment, real-time devices have become indispensable. To give a patient the best care possible, vital signs need to be monitored. These vital signs need to be available right away, because if the patient's condition is deteriorating, immediate action is of the utmost importance. There are several examples of real-time clinical measurements. Monitoring of the heart with an ECG, heart rate and oxygen levels in the blood using a pulse-oxygen sensor. But also blood pressure, body temperature and respiration rate can be measured in real-time.

Because this real-time aspect is so important, it also is a big design goal for this project. As can be read in Section 2.2, there are some implementations which can measure the vital signs of multiple people using only one radar, but all of those methods are using a post-processing technique. This means that the data is first gathered using the radar and captured with a capture card. At a later stage, all of this data is analyzed and the vital signs are extracted. So there are vital signs measured, but they are available far too late from a clinical point of view. The implementation from this project is focusing on the immediate availability of the vital signs which are being recorded at that moment using the radar. In this chapter the designs, elements and techniques are highlighted to process these relatively large amounts of data using a low power chip.

## 5.3 Building blocks

The IWR6843 chip consists of two processors. One is a ARM Cortex R4F processor with a clock speed of 200 MHz. The other one is a Digital Signal Processor (DSP) specific C674x processor, made by TI. These two processors are separated into two sub-systems: the main sub-system and the DSP sub-system. See also Figure 2.7.

### 5.3.1 Main sub-system

The Cortex R4F processor is the generic all purpose processor in the system. This processor has a managing role in the system. The R4F is connected to all of the communication interfaces, like SPI, I2C, UART and CAN. It is also connected to the radio front-end of the chip. Via a special mailbox system it is also connected to the DSP.

The main sub-system (MSS) is responsible for all non digital signal processing tasks. It catches the chirp- and frame-interrupts from the radio front-end, and communicates them to the DSP. It also takes care of communicating with external devices. In the case of this project, UART is extensively used to communicate the measurement details with a computer. Another use case for the R4F are more general calculations on the radar data. The DSP takes all of the raw radar data in and processes it to a radar cube. In this radar cube are points in 3 dimensional space, which are detected by the radar. The R4F has also access to this data, and could perform additional algorithms on this data, like object detection, point grouping or even AI techniques.

### 5.3.2 DSP sub-system

The DSP sub-system (DSS) is responsible for the transformation between the raw ADC radar data and usable data for further processing, for example in a radar cube format. The DSS has also a hardware accelerator build-in, to be able to do specific radar signal processing steps as fast as possible. The Radar Hardware Accelerator (HWA) enables off-loading the burden of certain frequently used computations in FMCW radar signal processing from the main processor. FMCW radar signal processing involves the use of FFT and log magnitude computations to obtain a radar image across the angle, velocity, and range dimensions. This gives the flexibility of doing common calculations using the HWA, but still design your own algorithms.

The DSP has L1 and L2 caches to store and load calculations quickly. The L3 cache is shared with the MSS, and contains all of the radar data.

### 5.3.3 Inter-processor communication

Because the MSS and the DSS are separate processes but must be working in unison, there is a communication technique present in the chip called the *mailbox*. Using this mailbox, messages can be sent from one processor to the other, and vice versa. When a message arrives at a processor, an interrupt is triggered and the message can be processed. Another way to use this method is to use it for processor synchronization. Quite often, the two processors need to be in the same state to proceed to the next part of a calculation. When the

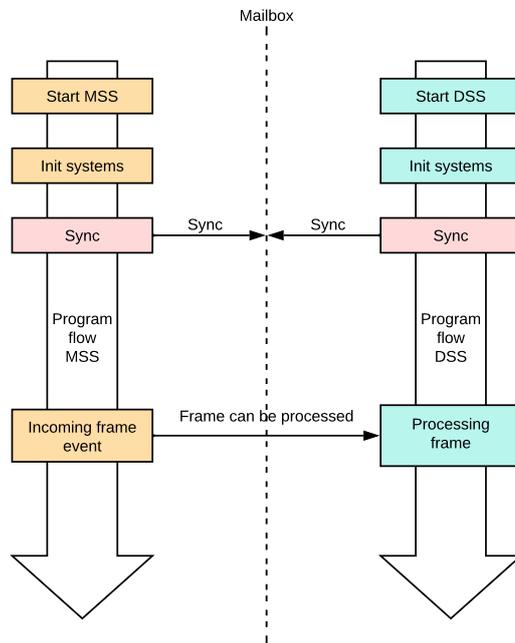


Figure 5.1: **Example of the mailbox principle. After startup and initialization, the two processors are synchronized such that they start at the same time. When a frame event ends up at the MSS, it can send a message using the mailbox to the DSS.**

two processors are synchronized using the mailbox, they both are at a known state in the program flow and can continue with the task at hand. See as an example Figure 5.1. The processors are starting separately, and are doing each their own initialization routine. After that routine is finished, they are waiting on each other (synchronized). Then they can start the program execution at exactly the same time. An example of sending a message from one processor to another is also given in Figure 5.1.

## 5.4 Implementation

This section goes into the program flows which make it possible to have real-time vital signs monitoring for multiple persons which are sitting in front of the sensor. A high-level overview can also be found in Figure 5.3.

Like mentioned in Section 5.3, the IWR6843 contains two processors, one generic ARM processor and one TI DSP. These processors are programmed separately, but the program flow is intertwined, both processors depend on each other to send and receive information.

### 5.4.1 MSS program flow

The program starts with an initialization phase, in this phase the processor gets setup to send and receive information from the right locations. The next step is to wait for the CLI input. The IWR6843 receives commands from the connected computer to provide the right parameters for the inner workings of the chip, see also Section 4.1.1. These parameters are processed, checked for errors and are getting send to the right locations. Some parameters are meant for the DSP, some are to program the radio front-end and some are for the MSS itself. After that step, the biggest task from this processor is done. There are two tasks that remain. The interrupts from the radio front-end which signal that a frame has been completed are send to the MSS. These messages are getting communicated along to the DSP, where the actual processing can start. All of the radar signal processing and the additional algorithm execution is done on the DSP. This program flow could be optimized a bit to divide the work between the two processors, but for this proof-of-concept design it was easier to do all of the calculations on one processor. When all of the data has been processed by the DSP, the heatmap and all of the vital signs information gets send back to the MSS. The MSS packages this information in a format which can be send over UART and sends this information to the connected computer, where it can be visualized using the GUI.

### 5.4.2 DSS program flow

The DSP also starts with an initialization phase, just like the MSS. But after that, it is very much a reactionary system. It gets impulses from external sources and reacts accordingly. Once the MSS sends the CLI parameters, it saves them into memory. This action can happen at any time during the runtime of the chip. The real important routine gets called when a signal is received from the MSS that the radar frame has been completed. This means that all of the information is present in memory to start another calculation round. First, the heatmap is generated using the HWA and the raw radar data. When this step is done, the people detection algorithm finds the persons in the room once every 64 frames. This happens only each 64 frames, to stabilize the person detection, persons in the room are in a static position anyway. Now that the positions of the persons in the radar view are known, the DSP proceeds to do a vital signs estimation on each of the persons in view of the sensor. After this has been done, the DSP sends this information back to the MSS.

### 5.4.3 Timing

For this project, timing is important. To process all of the radar data without creating a backlog, the system must process radar data at 10 frames per second. This means that each 0.1 second a new radar frame arrives and must be processed before the next frame arrives. The DSP is able to handle this load and always finishes with all the calculations before the next frame is ready. What could have posed a problem, is sending the data back to the computer using UART. The heatmap is 48 azimuth bins wide, 64 range bins long and each bin is 2 bytes. This means  $48 \times 64 \times 2 = 6144$  bytes of data which need to be send each 0.1 seconds. The vital signs data is an additional 80 bytes of data, including

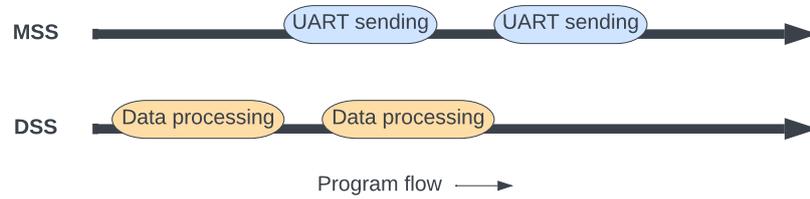


Figure 5.2: **Visual representation of how the two processors are working together to be as efficient as possible.**

the packaging of the data 6400 bytes of data need to be send. The UART port being used is running at 921600 baud, which means that 115200 bytes of data can be send each second. Doing the calculations,  $6400/115200 = 0.056$  seconds. With means that 56% of the time would be lost just sending the data. The solution to this problem is to use the MSS processor to send the data. This is possible because of the shared address space of the DSP and the MSS. So while the DSP starts doing the calculations for the next frame, the MSS is sending the data from the previous frame back to the computer. See also Figure 5.2.

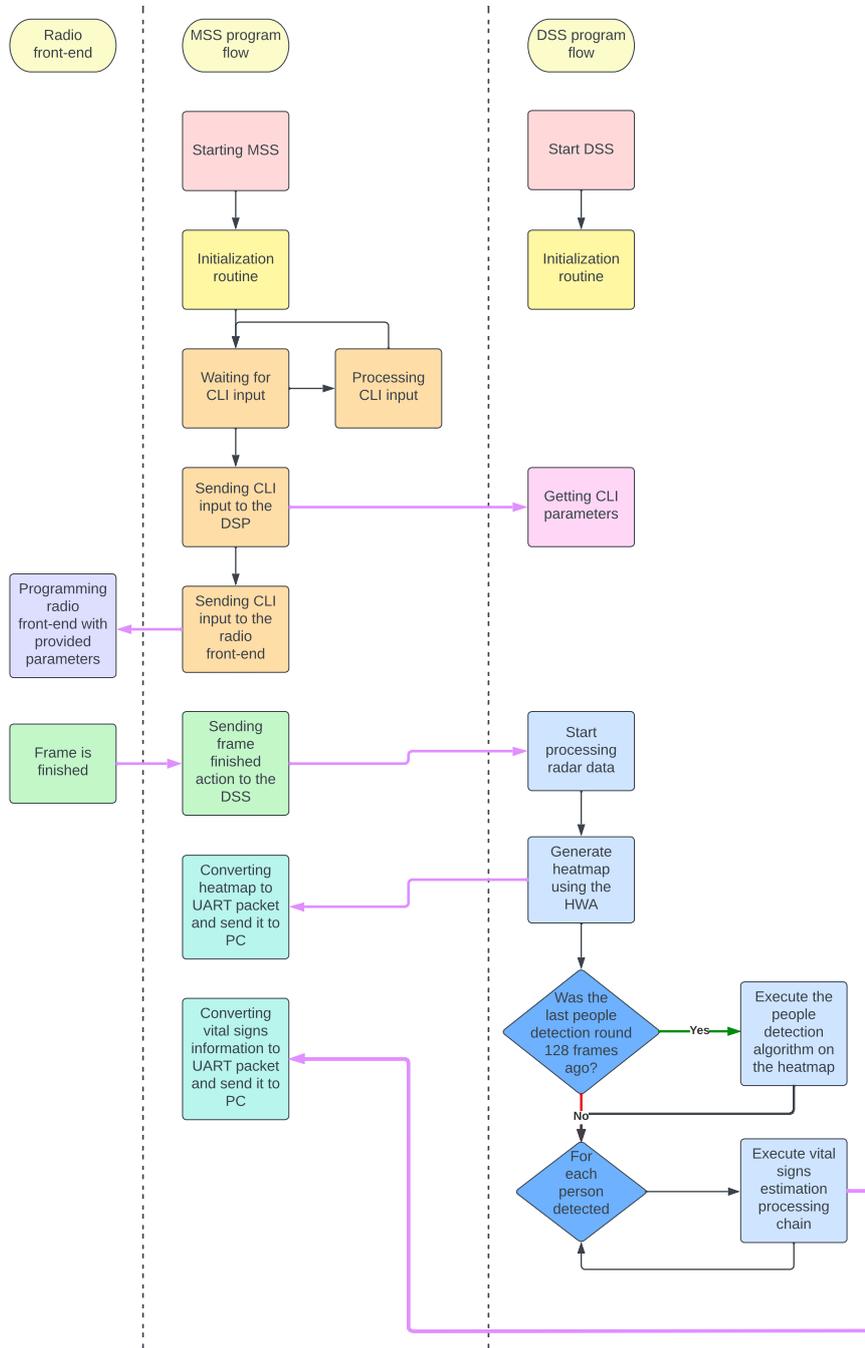


Figure 5.3: The program flow from the project in high-level overview. Steps with the same color belong to each other.

# Chapter 6

## Validation

Now that the prototype has been build, it needs to be evaluated to assess the accuracy of the implemented algorithms. This chapter goes into depth about the validation methods and the validation results.

### 6.1 Variables used

To validate this project correctly, the variables that are worked with must be established. With these variables in mind, some can be changed while others remain the same. In this way, it can be determined which variable has an impact on the accuracy of the sensor.

#### 6.1.1 Recorded variables

The recorded variables are the measurements that are being taken to measure the vital signs for a person. These measurements are coming from the sensor, but also from other devices to validate the measurements from the sensor. The sensor itself measures the heart rate and respiration rate from one or multiple persons. But not only this information gets send back to the computer. To make the visualization on the computer a bit better, different waveforms are also send to the computer. There is one unfiltered waveform, a heart rate waveform and a respiration rate waveform. The unfiltered waveform is the phase of the radar data for the bin in which the measured person is residing. The heart rate and respiration rate waveforms are the waveforms from which the heart rate and respiration rate can be determined. These waveforms are being generated after the filtering of the phase signal, see also Section 4.2. For a snapshot of the GUI on the computer, see Figure 6.1.

#### Heart rate validation

The sensor does output a heart rate, but there also needs to be a method to validate this heart rate measurement. This is done by also connecting the measured person to a pulse oximeter. The exact workings of this sensor are explained in Section 2.1.1. This sensor outputs two data streams: heart rate and blood oxygen levels. The only variable which is used for this project is the heart rate. The pulse oximeter used in this project is the *Nonin 9600*, see also

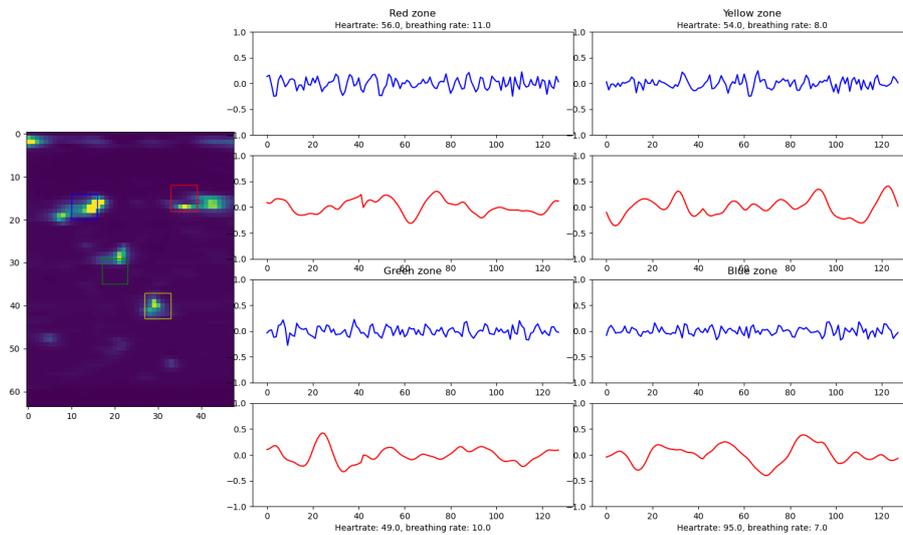


Figure 6.1: Snapshot of the GUI on the computer. On the left is the heatmap with the detected persons, on the right for each person the heart rate waveform in blue and the respiration waveform in red.



Figure 6.2: The Nonin 9600 pulse oximeter used as validation in this project.

Figure 6.2. This device has been chosen for various reasons. It is a medical grade device, which means that it is build following a high standard, and the measurements outputted by the sensor has very low tolerances. This device also has a serial port. Each second, the device outputs the heart rate and blood oxygen level to the serial bus. This data stream can be captured using a RS232 to USB converter, and saved to a file for later use. The only problem is that there are only two devices available. This is enough for a thorough evaluation of the vital signs estimation and also for measurements of one or two persons at the same time, but not for four persons. During the test with measuring four persons at the same time, two other pulse oximeters were used, the *iHealth Air Pulse Oximeter*, see Figure 6.3. These sensors don't have a connection available to send the data to a computer, so the sensors are read every minute, and the reading from the sensor is recorded.



Figure 6.3: **The iHealth Air Pulse Oximeter used as validation in this project.**

### **Respiratory rate validation**

While there are certainly devices which could measure respiration rate of a person, these devices were not present in the Erasmus MC, where the tests were performed. To still validate the respiration rate of a person, the following validation technique was developed. Since the person being measured is only sitting in front of the sensor and not doing any activities, the person is asked to count his own breaths. Every breath in and out counts as one breath, and each minute, this value is recorded and the counting starts again. This value counts as the validation for the respiration rate outputted by the IWR6843. This method has its downsides, people could forget to count, they could make a counting mistake, or they could start breathing non-naturally, since they are focusing on their breath. This is still the best way to measure the respiration rate during the project, but these downsides must be kept in mind when analyzing the validation results.

### **6.1.2 External variables**

Variables outside of the sensor could also have an effect on the quality of the measurements. These variables contain among others the position in front of the sensor, the amount of people in front of the sensor and movement in the room.

#### **Position with respect to the sensor**

The IWR6843 is set up to measure persons within 0.3 to 2.5 meter range of the sensor, in the whole azimuth range, see also Figure 6.4. The position of the person in the field of view of the sensor could play a role in the accuracy of the measurements.

#### **Amount of persons in front of the sensor**

The amount of persons that need to be measured is also a variable. It is possible that measurements of multiple persons close to each other will interfere.

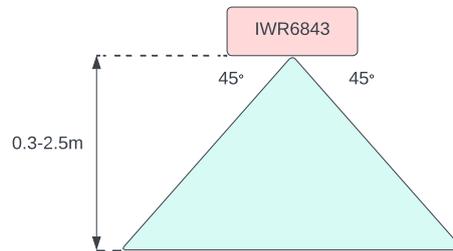


Figure 6.4: **Range and azimuth constraints of the sensor.**

### Background noise

The sensor is very sensitive to noise. For the best result, the person needs to sit very still, and there needs to be as less noise as possible in the field of view of the sensor, to get a better SNR ratio.

### Height of the sensor

The height of the sensor with respect to the chest region of the measured persons is very important. The sensor needs to be at the same height as the middle of the chest of the measured persons. Otherwise, the vibrations of the chest can't be picked up by the sensor and the measurements become unreliable.

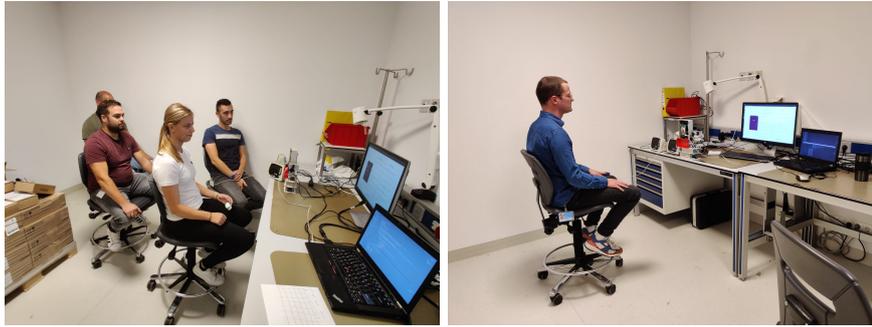
### 6.1.3 Person variables

The persons for which the vital signs are measured, all have different attributes. The most important ones are:

- **sex:** there exist physical differences between men and women, also in the chest region where this project is focusing on. This metric could determine if men and woman are easier or more difficult to measure.
- **age:** the age of a person could contribute to multiple factors. A child is small, has a small chest region and could have a faster beating heart. Older persons are more likely to have extra body fat to dampen the vibrations produced by the heart and lungs.
- **weight/BMI:** body mass could have a big impact on the measurements. Positively, persons with a bigger weight are more likely to have a bigger chest region, so are easier to measure. Negatively, persons with more body mass are more likely to have extra fat on top of the chest, which could dampen the vibrations caused by the heart and lungs.

## 6.2 Testing setup

The testing environment is kept as consistent as possible. All of the equipment is placed on a desk, and in front of the desk one or more chairs are placed. The IWR6843ISK is set on the desk, with the antennae pointing to the chairs. Next



(a) Picture taken during the 4 persons vital sign estimation test. (b) Picture taken during the 1 person vital sign estimation test.

Figure 6.5: Two different tests being executed.

to the IWR6843ISK the two Nonin 9600 pulse oximeters are placed, since the sensor leads are not that long. The IWR6843ISK and the two SpO<sub>2</sub> sensors are connected to the computer via serial buses. The computer is operated by the operator. See also Figure 6.6 and Figure 6.5.

All tests have a duration of 5 minutes, or 300 seconds. The amount of 300 seconds has been chosen for multiple reasons. It gives the sensor time to stabilize. Because multiple buffers are used inside of the sensor, it takes some time before the buffers are filled with valid measurement data. Because the test subjects have walked to the testing location, it also gives the heart rate of the test subjects some time to settle in the heart rate during rest, which is the most stable.

## 6.3 Test subjects

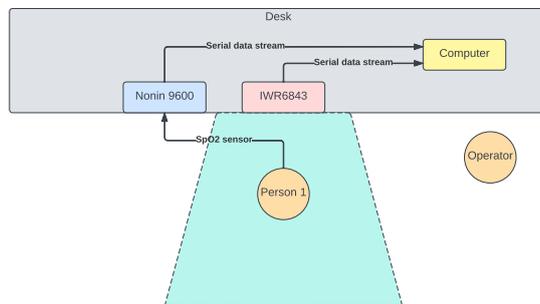
To properly validate the sensor, it has to be tested on persons. These test subjects were selected to reflect the validation plan. The test subjects are recorded for 5 minutes in a row, and along with the measurement data, additional metadata of the test subjects were recorded, as described in Section 6.1.3. This data includes age, gender, weight, length and additional comments. Using the weight and length, the BMI can be calculated. This data can be found in Table 6.1.

Because the data which is being recorded is highly sensitive medical data, the names of the persons are anonymized. All of the persons also have signed an informed consent. These consent forms are stored in the archives of the 4TU Centre of Research Data, along with the measurement data itself in an encrypted form.

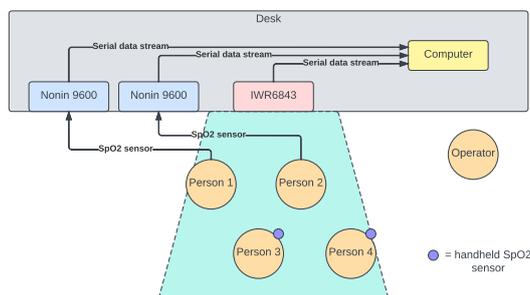
## 6.4 Results

### 6.4.1 One person validation

The validation of the vital signs of one person in front of the sensor is an important part of the validation process. Because the vital sign estimation of



(a) Test setup for vital sign estimation of 1 person.



(b) Test setup for the vital sign estimation of 4 persons.

Figure 6.6: Two test setups being used to validate the project.

one person is the same process as the vital sign estimation of multiple persons, the validation of one person was used to partially validate multiple person vital signs. To validate the general accuracy of the project, 12 persons have been tested, all individually. The test setup can be found in Section 6.2.

For each recorded heart rate and respiration rate from the IWR6843, it is compared against the control heart rate and respiration rate. For each test subject, the average accuracy percentage is calculated. The accuracy percentage is how much the measurement from the IWR6843 differs from the control measurement. The percentage is calculated using the formula in Eq. 6.1.

$$accuracy = \text{abs}(M_{IWR6843} - M_{control}) / M_{control} \quad (6.1)$$

These averaged percentages can be found in Table 6.2. The most accurate measurement and the least accurate measurement are drawn in this thesis. For the most accurate result, see Figure 6.9. For the least accurate result, see Figure 6.10.

In later sections, more details are provided of the relation between personal variables and the detected accuracy of the measurements. For now, only observations are done on individual measurements. The first observation is the difference in deviation between Table 6.2 and Table 6.3. This is because the percentage in Table 6.2 is calculated for each measurement. For 300 seconds of recording, 1200 measurements are gathered. These 1200 percentages get averaged into one percentage. For Table 6.3, the measurements from the control

Name	Age	Gender	Weight (kg)	Length (m)	BMI	Comments
P1	23	male	72	1.75	23	
P2	61	female	78	1.60	30	
P3	32	male	75	1.85	21	
P4	25	male	80	1.80	24	Wore big layers of clothing
P5	26	male	100	1.80	30	P5 was not detected very well by the sensor
P6	27	male	70	1.75	22	
P7	23	male	90	1.82	27	Was a bit restless during testing
P8	43	male	80	1.82	24	
P9	49	male	85	1.85	24	
P10	35	male	86	1.77	27	
P11	49	male	92	1.80	28	
P12	23	female	62	1.75	20	
P13	76	male	91	1.75	29	
P14	75	female	82	1.65	30	
P15	46	female	58	1.7	20	
P16	23	male	68	1.76	21	

Table 6.1: **All of the recorded metadata from the test subjects.**

and from the IWR6843 are first averaged, and then the deviation percentage is calculated. Because the measurements from the IWR6843 are quite noisy, the deviation percentages in Table 6.2 are higher than Table 6.3.

Another way to visualize the data is to generate a Bland-Altman plot. This plot is specifically designed to plot out the agreement of two different measurement devices. In the case of this project, that would be the radar measurement and the control measurement. In this graph, the mean difference and the standard deviation between the two measurement types are being given. The Bland-Altman plot for the heart rate comparison of all one person validation tests can be found in Figure 6.7. The Bland-Altman plot for the respiration rate can be found in Figure 6.8.

The average heart rate estimation accuracy of the sensor is 10.86%. This is mainly due to four persons where the measurement accuracy was worse compared to the other 7 measured persons. For three of the four persons a reason can be found for this lesser accuracy, see also the remarks in Table 6.1. *P4* had multiple layers of thick clothing on, which could distort the signal. *P5* was not detected by the people detection algorithm very well, so the quality of the data to do the vital signs estimation on was not very high. *P9* has lung problems, which express in a higher and more unstable respiration rate and heart rate. For the lower accuracy of *P13* there is no clear explanation, except an higher BMI, which will be discussed in Section 6.4.4. Taking out these four lower accuracy's, the average accuracy drops down to 3.30%.

The respiration accuracy averaged around all the persons tested is 7.61%. But, when looking at the breathing rate numbers in Table 6.3, the difference between the control- and the sensor measurement is either 0 or 1. Only at *P5*

### SpO2 versus Radar Heartrate Monitoring One Person Validation

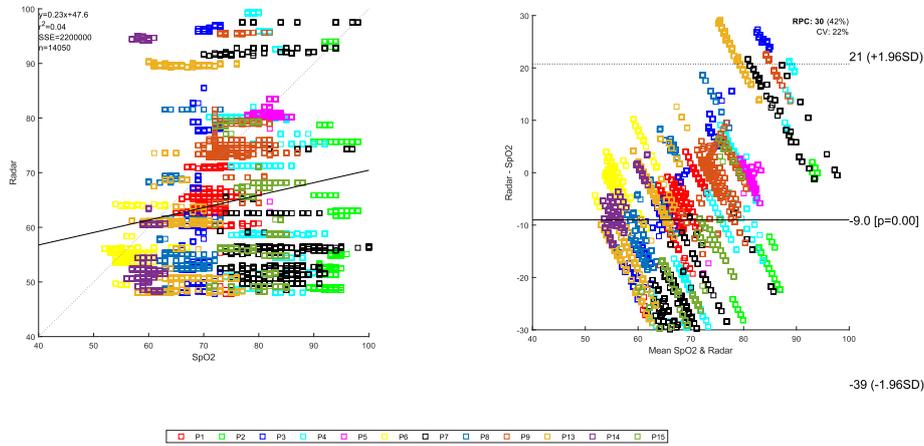


Figure 6.7: Bland-Altman plot of all one person heart rate validation results.

there exists a difference in measurements of 2, but this could again be due to the people detection algorithm not detecting the person properly. Since the control measurements are counted by the persons themselves, part of this accuracy number could be due to miscounting.

To conclude, the one person vital signs estimation works. Especially the respiration rate is exactly right or only one off. The heart rate estimation accuracy is around 10%. The reason for this lower accuracy could be that the heart rate is faster than the respiration rate and the magnitude of the vibrations is smaller. This accuracy could be improved by tweaking the algorithms and the algorithm parameters.

### 6.4.2 Two person validation

An important step for this project is moving on from measuring one test subject to two test subjects at the same time. The most important question to answer is if the measurement from one test subject can influence the measurement from another test subject. The test setup is the same as in Figure 6.6b, with the exception that only the first row is in use. The other deviation is that only one Novim 9600 was available to measure heart rate, so the other heart rate is measured using the iHealth device, see also Section 6.1.1.

The results from this validation can be found in Table 6.4. The first observation which immediately stands out is that the estimation for one test subject is almost perfect, while the estimation for the other test subject is far from perfect. In the next Section (6.4.3), four test subjects are measured at the same time, and there the accuracy is higher. No explanation can be found for this large difference in accuracy.

### Counting versus Radar Respiratory rate Monitoring One Person Validation

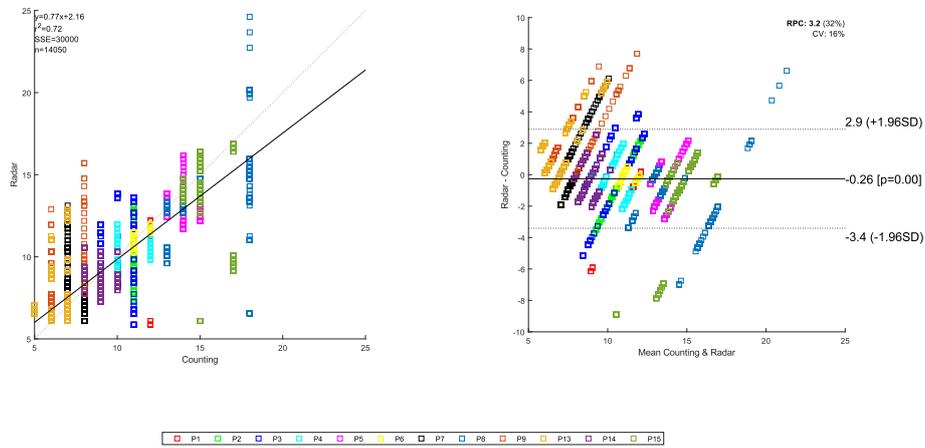


Figure 6.8: Bland-Altman plot of all one person respiratory rate validation results.

### 6.4.3 Four person validation

For this project, estimating the vital signs of 4 persons at the same time is the maximum. There are two reasons for this number. Firstly, because of the resource-heavy algorithms this is all the chip can handle at the moment because of timing constraints. Secondly, because the measurements take place in the range of 0.3 to 2.5 meters from the sensor, it becomes too crowded in front of the sensor with more than 4 persons.

To validate the accuracy of measuring four vital signs at the same time, four persons were put in front of the sensor, as can be observed in Figure 6.6b and Figure 6.5a. The rest of the validation procedure stays the same, the test is run for 5 minutes with the persons sitting statically. All of the chests of the four persons are at the same height of the sensor antennae.

The most uncertain part of this validation was if the person detection would work reliably. Not only one, but four persons must be detected and remain detected during the 5 minutes of testing time. Luckily, the person detection remained stable during the test, and apart from a few glitches the four persons remained detected during the test, see also Figure 6.11.

The graphs visualizing the measurement data for all four persons at the same time can be found in Figure 6.12 and Figure 6.13. The most outstanding observation which can be done, is that the estimations from the IWR6843 are really jumping up and down. The measurements from the control measurement are much more stable. A positive observation is that the control measurement is most of the time approximately the average of the noisy measurements from the IWR6843. This is also confirmed with the numerical results in Table 6.5 and Table 6.6. The numbers in this table are calculated in the same way as Section 6.4.1. In Table 6.5, each single result from the IWR6843 and the control measurement is compared to each other, which gives a large deviation because

Person	Heart rate accuracy	Respiratory rate accuracy	Average
P1	8.22%	6.56%	<b>7.39%</b>
P2	1.71%	9.53%	<b>5.62%</b>
P3	4.70%	2.01%	<b>3.35%</b>
P4	29.75%	11.11%	<b>20.43%</b>
P5	16.74%	8.51%	<b>12.63%</b>
P6	5.49%	15.74%	<b>10.62%</b>
P7	20.61%	23.93%	<b>22.27%</b>
P8	16.85%	13.49%	<b>15.17%</b>
P9	24.80%	7.89%	<b>16.35%</b>
P13	35.39%	13.12%	<b>24.26%</b>
P14	18.81%	16.64%	<b>17.72%</b>
<b>Average</b>	<b>17.44%</b>	<b>13.72%</b>	<b>15.58%</b>

Table 6.2: **One person vital signs estimation validation results. The percentage denotes the absolute difference between the IWR6843 heart- and respiration rate, and the control rates.**

of the scattered results from the radar sensor. In Table 6.6, all measurements are averaged in one number. In this case, the deviation is much smaller.

The two person validation and the four person validation are also summarized in Bland-Altman plots. From these plots the mean difference between the two measurement types and the standard deviation can be observed. The multiple person heart rate Bland-Altman plot can be found in Figure 6.14, respiratory rate plot can be found in Figure 6.15.

#### 6.4.4 mmWave radar accuracy compared to test subject variables

When performing the validation tests, some variables from the test subjects were captured, see also Section 6.1.3. An interesting observation would be to see if there is a correlation between those variables and the measurement accuracy. This would make sure that vital signs monitoring using mmWave radar would be a viable option for all types of persons. The accuracy is compared against the age, BMI and sex of a person. The age and BMI graphs can be observed in Figure 6.16.

##### Age

In Figure 6.16a the mean difference between the radar measurement and the control measurement is plotted against the age of the test subject. In this graph, no correlation can be observed. This means that the age of a person has nothing to do with the accuracy of the vital signs estimation.

##### Weight / BMI

In Figure 6.16b, the BMI of a test subject is plotted against the mean difference between the radar vital signs estimation and the control measurement. The BMI variable has been chosen, because persons with a higher BMI generally also

Person	Heart rate SpO2 sensor	Heart rate IWR6843	Heart rate deviation	Respiratory rate	Respiratory rate IWR6843	Respiratory rate deviation
P1	71	65	8.03%	10	9	6.70%
P2	81	81	0.54%	14	13	6.90%
P3	56	56	1.26%	11	11	0.54%
P4	83	62	25.83%	8	8	2.82%
P5	66	59	10.77%	15	13	8.29%
P6	75	76	1.36%	6	7	14.22%
P7	66	63	3.59%	6	7	19.19%
P8	60	57	3.80%	9	9	7.61%
P9	80	60	24.48%	15	15	4.61%
P13	94	61	35.26%	11	10	7.82%
P14	69	66	4.54%	10	10	5.00%
Average			10.86%			7.61%

Table 6.3: One person vital signs estimation validation results. The average heart rates and respiratory rates of all sensors from the whole reading, with the absolute deviation between the control measurements and the IWR6843 measurements.

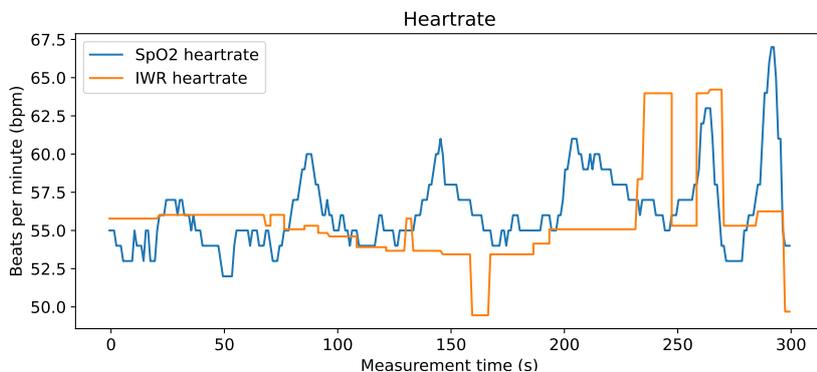
Person	Heart rate SpO2 sensor	Heart rate IWR6843	Heart rate Deviation	Respiratory rate	Respiratory rate IWR6843	Respiratory rate deviation
P15	83	57	31.20%	12	11	15.00%
P16	71	69	2.77%	10	10	8.33%
Average			16.98%			11.67%

Table 6.4: Two person vital signs estimation validation results. The average heart rates and respiratory rates of all sensors from the whole reading, with the absolute deviation between the control measurements and the IWR6843 measurements.

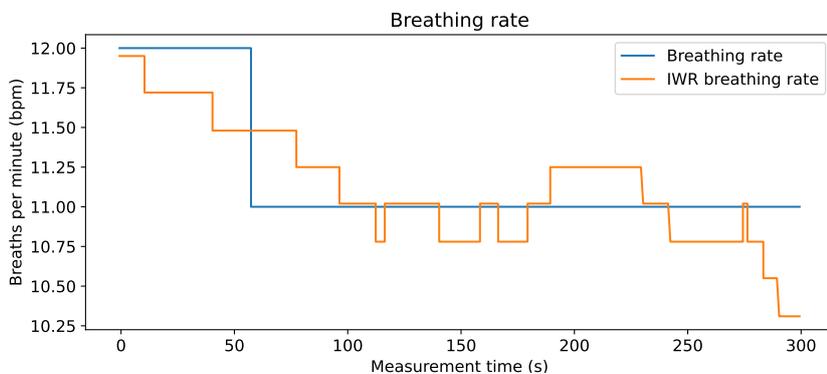
have more fat in their body. The hypothesis is that this fat could dampen the vibrations from which the vital signs are deduced. Again, no clear correlation can be found between this metric and the variable. This means that this type of measurement could be used for bigger persons, but also for more skinny persons.

## Sex

The mean difference between the radar estimation and the control measurement for female test subject's heart rate is 8.74%, for the respiratory rate it is 4.42%. The mean difference for the male test subject's heart rate is 11.76%, the respiratory rate difference is 9.04%. The difference between the measurements for the male test subjects is higher, but not by a large amount. This difference could also be due to measurement inaccuracies. Because the validation was performed on only 4 female test subjects compared to 12 male test subjects, more female



(a) The IWR6843 estimated heart rate and the SpO2 sensor heart rate from *P3* plotted out.



(b) The IWR6843 estimated respiratory rate and the counted respiratory rate from *P3* plotted out.

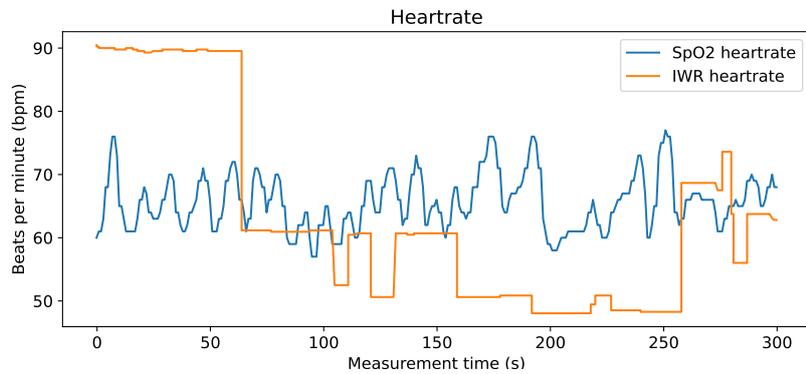
Figure 6.9: Validation data from *P3*. This is the most accurate result in the one person validation section.

validations need to be done to come up with a definite conclusion.

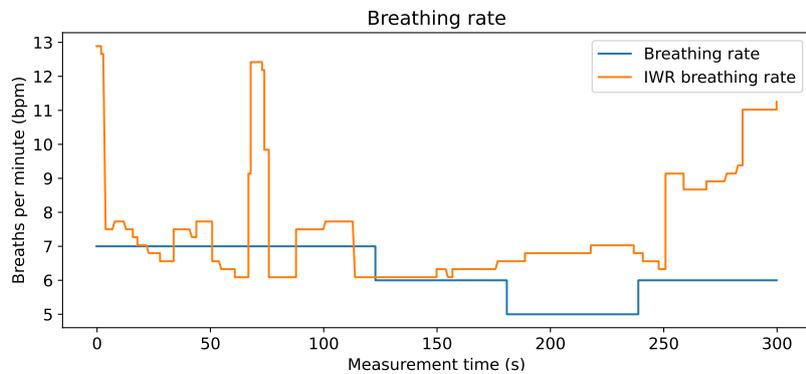
## 6.5 Evaluation

When observing the result from the one person-, two person- and four person validation, it could be noted that the accuracy reduces when there are more persons in the sensor view. This could have multiple reasons. Because there are more persons in the field of view of the radar, more noise is introduced. The test subjects all constantly move very slightly, this could introduce noise in the very sensitive phase information. Another option could be that the phase changes of the different persons interfere with each other. More knowledge about radar theory could probably substantiate that hypothesis.

Taking into account all of the measurements from the 16 test subjects in the different configurations, this gives a good representation of the accuracy of the prototype of this project. The algorithms could be tuned better to make the sensor measurements more stable and generally more accurate, but the main



(a) The IWR6843 estimated heart rate and the SpO2 sensor heart rate from *P7* plotted out.



(b) The IWR6843 estimated respiratory rate and the counted respiratory rate from *P7* plotted out.

Figure 6.10: Validation data from *P7*. This is the least accurate result in the one person validation section.

goal of this project is to see if the algorithms can be put on the sensor chip itself, and if it would hold up against the large stream of radar data which was coming in. This goal has been reached, because the sensor never missed a timing deadline during the validation phase.

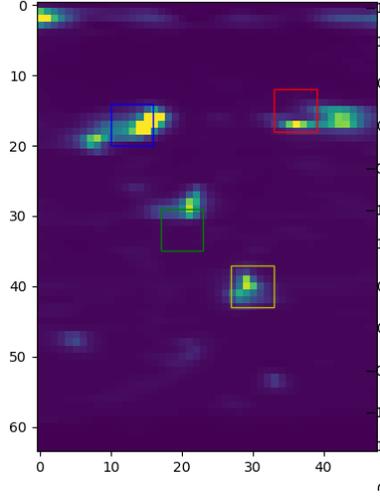


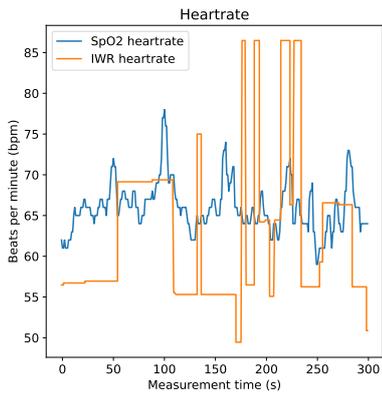
Figure 6.11: Four persons detected during the four person validation.

Person	Heart rate accuracy	Respiratory rate accuracy	Average
P3	12.15%	21.53%	16.84%
P10	21.19%	32.18%	26.69%
P11	24.34%	25.18%	24.76%
P12	18.06%	10.08%	14.07%
Average	18.94%	22.24%	20.59%

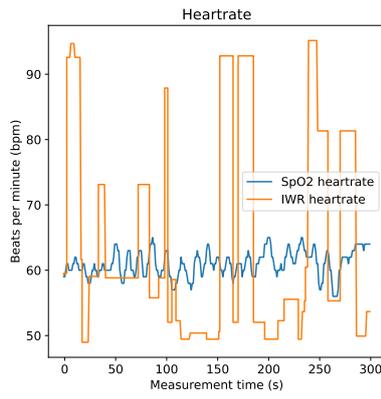
Table 6.5: Four person vital signs estimation validation results. The percentage denotes the absolute difference between the IWR6843 heart- and respiration rate, and the control rates.

Person	Heart rate SpO2 sensor	Heart rate IWR6843	Heart rate deviation	Respiratory rate	Respiratory rate IWR6843	Respiratory rate deviation
P3	61	64	5.64%	7	9	28.24%
P10	66	63	5.04%	10	9	6.02%
P11	69	58	16.09%	8	9	2.38%
P12	69	60	12.86%	11	11	2.14%
Average			9.91%			9.70%

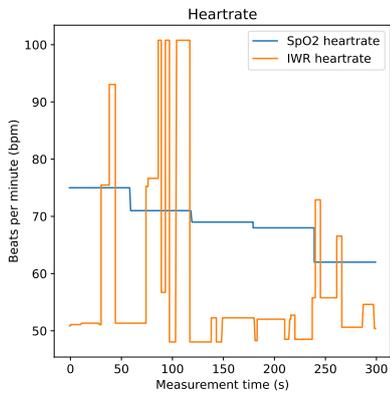
Table 6.6: Four person vital signs estimation validation results. The average heart rates and respiratory rates of all sensors from the whole reading, with the absolute deviation between the control measurements and the IWR6843 measurements.



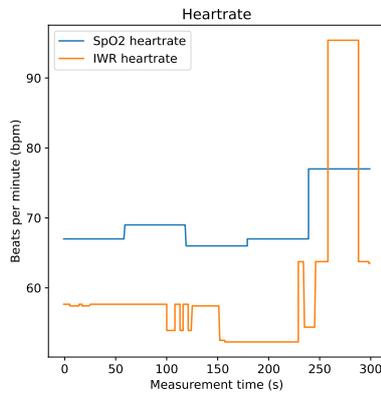
(a) heart rate of *P3*



(b) heart rate of *P10*

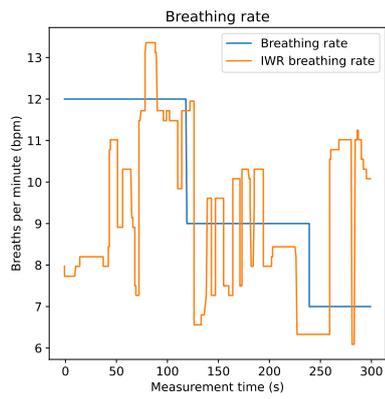


(c) heart rate of *P11*

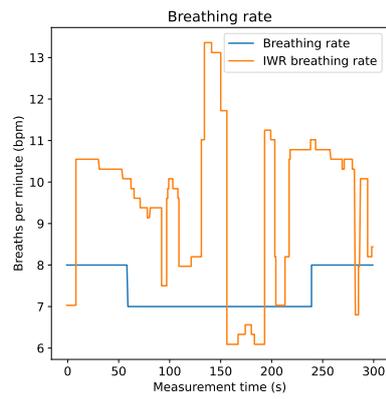


(d) heart rate of *P12*

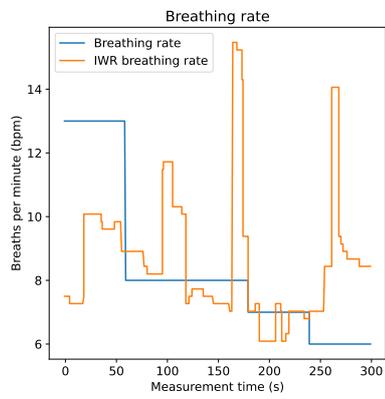
Figure 6.12: heart rate estimation from 4 persons at the same time.



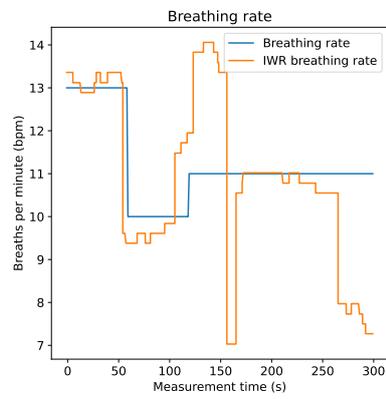
(a) Respiratory rate of *P3*



(b) Respiratory rate of *P10*



(c) Respiratory rate of *P11*



(d) Respiratory rate of *P12*

Figure 6.13: Respiratory rate estimation from 4 persons at the same time.

### SpO2 versus Radar Heartrate Monitoring Two Person Validation

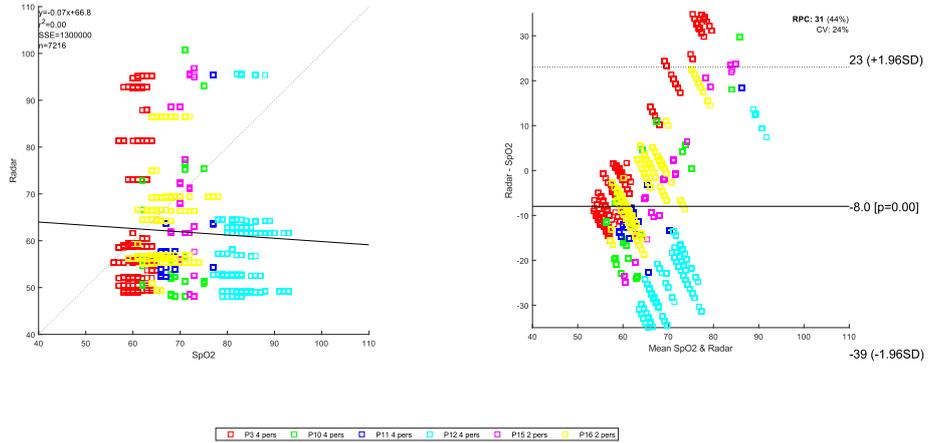


Figure 6.14: Bland-Altman plot of all multiple person heart rate validation results.

### Counting versus Radar Respiratory rate Monitoring One Person Validation

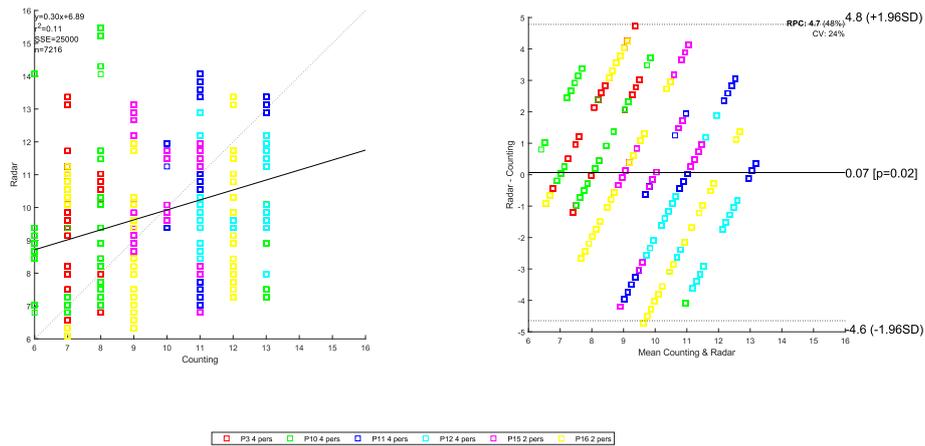
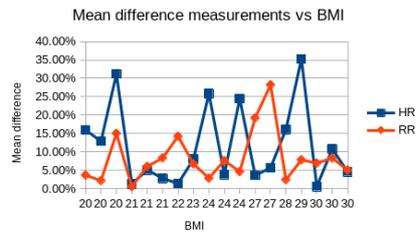
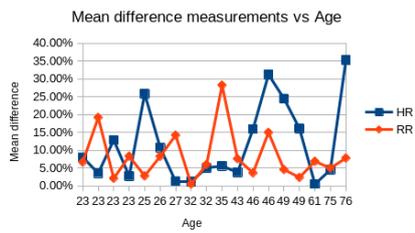


Figure 6.15: Bland-Altman plot of all multiple person respiratory rate validation results.



(a) Correlation between the mean difference of the measurements and the age of the test subject.

(b) Correlation between the mean difference of the measurements and the BMI of the test subject.

Figure 6.16: Correlation between the mean difference of the measurements and two test subject variables.

## Chapter 7

# Conclusions

The aim of this research was to measure multiple peoples heart rate and respiratory rate in real time with FMCW mmWave radar. First, the different existing methods are evaluated. Existing implementations make use of post-processing algorithms to evaluate the data. The existing real-time implementation of this task can only measure one person, because it only makes use of a 1D-FFT.

The implementation of this thesis is focusing on using the output from the range-azimuth FFT. This output can be further processed to form a 2 dimensional heatmap. Using the data from this heatmap, up to four persons can be detected dynamically. Because all of the complex data stays intact when forming the heatmap, the data needed for the vital signs estimation, the phases of the radar data, can be extracted from the heatmap data for each person detected. For each radar measurements these phases are collected to form a waveform. Two filters can be applied to this waveform, one to extract the heartrate and one to extract the respiration rate. Peak counting is applied to these two waveforms to finally end up with two numbers for each person in the radar frame: a heartrate and a respiration rate.

An important part of this thesis is the validation part. Because a new connection is used to extract the vital signs waveform from the radar data, this method needs to be properly tested. Tests have been done on individual test subjects, but also on two persons at the same time and four persons at the same time. There is also investigated if age, gender or the BMI of a person are connected to the accuracy of the measurement. The accuracy of the heartrate estimation using the sensor compared to the control measurement is around 10%. The accuracy of the respiratory rate estimation compared to the control measurement is around 8%. These accuracy's are not as high as we had hoped, but the main focus of this thesis is to show the vital signs of multiple persons can be estimated in real-time using only the embedded hardware. This goal has been reached, all of the calculations mandatory for the vital signs estimation of multiple persons are done on the chip itself. The relation between age, gender or BMI and the accuracy of the sensor cannot be proven using the validation data gathered during this thesis. This could denote that the vital signs estimation using the mmWave method can be used on a broad range of people.

In Section 1.3, the objectives for this project were given in eight points. Almost all of these objectives have been reached, there are only a few which are just partly completed. The chip is able to generate a heatmap from the radar

data and detect persons in the heatmap. For these persons, the vital signs are estimated in real-time. The resulting data gets send to the computer, where the data is visualized and saved for later analysis. When doing the validation of this project, the vital sign measurements were compared to medical grade device measurements, to assess the accuracy of the sensor. The tracking of multiple persons is only partly succeeded. When the test subjects are sitting statically, the sensor is able to adjust for small movements. However, if the test subjects are moving a lot, the heatmap becomes so noisy that person tracking is not possible anymore. For the validation part of this project, 16 test subjects have been tested, some even multiple times. This number of subjects gives a good impression of the overall performance of this implementation. One improvement could be the age distribution of the test subjects. No children have been tested for this project, these measurements could have given an interesting view on the performance on smaller chest surfaces.

## 7.1 Future work

The main improvement on the work done on this thesis are the vital signs estimation algorithms. Because the focus for this project was on the embedded programming and the real-time implementation, the implementation of the vital signs estimation algorithms could not really be improved upon because I lacked the required knowledge. The lower accuracy could be improved by tweaking the parameters of the algorithms, or implementing additional signal processing techniques.

The code for this project has been written as efficiently as possible. But due to the large code base, and the limited documentation of all the APIs included in the chip, it was very hard to have a grasp on the whole program flow. Developers with more experience on this chip could make the algorithm execution more efficient, by making more use of the specialized hardware in the chip, and the cooperation of the two separate processors. In this project, the DSP is used for all of the calculations, and the ARM processor is used for handling the UART data flow. To make better use of the processing power of the chip, the ARM processor could take over some tasks from the DSP. My estimate is that this could speed up the program by approximately 20%.

The reach of the sensor could also be improved. For now, the sensor can detect persons up to 2.5 meters away in front of the sensor. To measure for example the heart- and respiration rate of all persons in one room, the range of the sensor must be improved while keeping the accuracy required to estimate the vital signs. This of course also depends on the external requirement that this project is used for.

# Bibliography

- [1] Adeel Ahmad, June Chul Roh, Dan Wang, and Aish Dubey. Vital signs monitoring of multiple people using a fmcw millimeter-wave sensor. In *2018 IEEE Radar Conference (RadarConf18)*, pages 1450–1455. IEEE, 2018.
- [2] Ryota Aihara and Yasutaka Fujimoto. Free-space estimation for self-driving system using millimeter wave radar and convolutional neural network. In *2019 IEEE International Conference on Mechatronics (ICM)*, volume 1, pages 467–470, 2019.
- [3] Mostafa Alizadeh, George Shaker, João Carlos Martins De Almeida, Plinio Pelegrini Morita, and Safeddin Safavi-Naeini. Remote monitoring of human vital signs using mm-wave fmcw radar. *IEEE Access*, 7:54958–54968, 2019.
- [4] Selcan Kaplan Berkaya, Alper Kursat Uysal, Efnan Sora Gunal, Semih Ergin, Serkan Gunal, and M. Bilginer Gulmezoglu. A survey on ecg analysis. *Biomedical Signal Processing and Control*, 43:216–235, 2018.
- [5] Cesar Iovescu and Sandeep Rao. The fundamentals of millimeter wave radar sensors. [https://www.ti.com/lit/wp/spyy005a/spyy005a.pdf?ts=1657272691416&ref\\_url=https%253A%252F%252Fduckduckgo.com%252F](https://www.ti.com/lit/wp/spyy005a/spyy005a.pdf?ts=1657272691416&ref_url=https%253A%252F%252Fduckduckgo.com%252F), 2020. Last accessed: Oct. 1, 2022.
- [6] Peter H. Charlton, Drew A. Birrenkott, Timothy Bonnici, Marco A.F. Pimentel, Alistair E.W. Johnson, Jordi Alastruey, Lionel Tarassenko, Peter J. Watkinson, Richard Beale, and David A. Clifton. Breathing rate estimation from the electrocardiogram and photoplethysmogram: A review. *IEEE reviews in biomedical engineering*, 11:2–20, 2017.
- [7] John R Feiner, John W Severinghaus, and Philip E Bickler. Dark skin decreases the accuracy of pulse oximeters at low oxygen saturation: the effects of oximeter probe type and gender. *Anesthesia & Analgesia*, 105(6):S18–S23, 2007.
- [8] Iowegian International. Iir filter basics. <https://dspguru.com/dsp/faqs/iir/basics/>, 2020. Last accessed: Jul. 21, 2022.
- [9] Changzhi Li, Julie Cummings, Jeffrey Lam, Eric Graves, and Wenhsing Wu. Radar remote monitoring of vital signs. *IEEE Microwave magazine*, 10(1):47–56, 2009.

- [10] Mike Cadogan. Ecg lead positioning. <https://litfl.com/ecg-lead-positioning/>, 2022. Last accessed: Oct. 1, 2022.
- [11] Ministry of Economic affairs. Nationaal frequentieplan 2014. <https://wetten.overheid.nl/BWBR0035791/2022-03-31#Bijlage>, 2014. Last accessed: Jul. 08, 2022.
- [12] Wen Qi Mok, Wenru Wang, and Sok Ying Liaw. Vital signs monitoring to detect patient deterioration: An integrative literature review. *International journal of nursing practice*, 21:91–98, 2015.
- [13] Mohammad Abdul Motin, Chandan Kumar Karmakar, and Marimuthu Palaniswami. Selection of empirical mode decomposition techniques for extracting breathing rate from ppg. *IEEE Signal Processing Letters*, 26(4):592–596, 2019.
- [14] Caitlin Ramsey. Safety investigation of millimeter wave radar to monitor vital signs in the nicu. Master thesis, Delft University of Technology, Delft, The Netherlands, 2020.
- [15] James E. Sinex. Pulse oximetry: principles and limitations. *The American journal of emergency medicine*, 17(1):59–66, 1999.
- [16] Texas Instruments. Intro to mmwave sensing : Fmcw radars - module 1 : Range estimation. <https://training.ti.com/intro-mmwave-sensing-fmcw-radars-module-1-range-estimation?cu=1128486>, 2017. Last accessed: Jul. 14, 2022.
- [17] Texas Instruments. Iwr6843 intelligent mmwave sensor standard antenna plug-in module. <https://www.ti.com/tool/IWR6843ISK>, 2021. Last accessed: Oct. 1, 2022.
- [18] Texas Instruments. Vital signs lab. [https://dev.ti.com/tirex/explore/node?node=AKUOY-htBc6mwPY1fsOUvw\\_\\_VLyFKFf\\_\\_LATEST](https://dev.ti.com/tirex/explore/node?node=AKUOY-htBc6mwPY1fsOUvw__VLyFKFf__LATEST), 2021. Last accessed: Jul. 14, 2022.
- [19] Texas Instruments. mmwave sdk version 03.05.00.00. <https://www.ti.com/tool/MMWAVE-SDK>, 2022. Last accessed: Jul. 12, 2022.
- [20] Tunstall. What is a pulse oximeter? <https://blog.tunstallhealthcare.com.au/connected-health/what-is-a-pulse-oximeter/>, 2017. Last accessed: Oct. 1, 2022.
- [21] Vibration Research. Iir filtering. <https://vru.vibrationresearch.com/lesson/iir-filtering/>, 2022. Last accessed: Jul. 21, 2022.
- [22] Zhicheng Yang, Parth H. Pathak, Yunze Zeng, Xixi Liran, and Prasant Mohapatra. Monitoring vital signs using millimeter wave. In *Proceedings of the 17th ACM international symposium on mobile ad hoc networking and computing*, pages 211–220, 2016.

# Appendix A

## Tx beamforming

In this appendix, another function of the IWR6843 is explained: Tx beamforming. This could have been a very promising solution to the multiple person vital signs estimation problem, but sadly this solution did not work as expected. However, a lot of time has been spend researching, implementing and debugging this solution, so a short summary of the development will be given, along with some pointers to where this implementation might have gone wrong.

### A.1 Hypothesis

One function in the IWR6843 is Tx beamforming. This function is used to extend the range of the sensor beyond the limits during regular operation. This is achieved by making all three of the transmitting antennas work together. Each transmitting antenna has a build-in phase modulator, which can modulate the phase of the outgoing signal. Each transmitting antenna can modify its phase in such a way that the transmitting signals bundle together to form a beam. This beam has a much lower azimuth reach, but because it is directed at a specific direction it has a lot more range. This means that the range of the sensor could reach up to 150 meters away, compared to 50 meters in normal operation. Like mentioned, the azimuth reach is much lower compared to the normal operation. There is a way to achieve the normal azimuth reach, which is recording multiple radar frames at different angles with respect to the sensor, and stitching these frames together to one big radar frame.

This is the functionality that could have been exploited for this project, the range of the sensor doesn't have to be extended, but persons have to be separated from each other. This could be done by sending a Tx beam to the person. Because of the narrow azimuth reach, only data for this person would be captured.

So first, a normal radar image would be generated to detect the persons available in reach of the sensor. Then, the sensor would zone in on each person present, and record radar data to do the vital signs estimation on. After a certain amount of time, the sensor would switch to another person. See Figure A.1 for a visual representation.

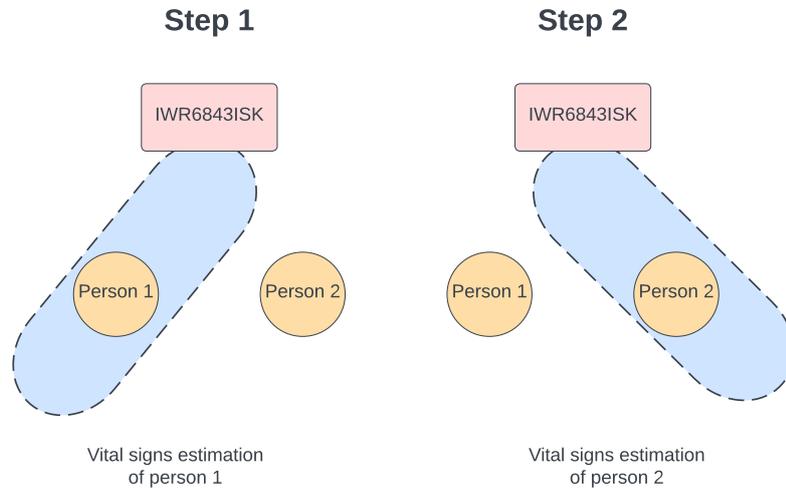


Figure A.1: **The hypothesis of how the Tx beamforming would work in theory.**

## A.2 Implementation

This hypothesis was implemented on the IWR6843, with the help of demo code provided by TI and the documentation available from the mmWave SDK [19]. However, the SDK documentation mentioned that this feature was implemented by the developers, but there was no testing done on this feature.

Also, when asking for help on the TI mmWave forum, they didn't have a lot of experience with this feature of the sensor. They had not tried this feature for this kind of application yet, so I wanted to give this idea a shot. I wrote a program for the IWR6843 which would form a Tx beam and steer it as much to the right as possible. The data which came back to the Rx antennae was visualized in a heatmap and send back to the computer.

## A.3 Results

Sadly, this implementation did not work as expected. While the sensor was setup to steer the Tx beam all the way to the right, persons were detected in the heatmap while standing in all angles with respect to the sensor. This meant that the idea of isolating persons by Tx beamforming would not work like we expected. Even if the distance between the sensor and the person was increased to up to 4 meters, the person was detected while standing all the way to the left of the sensor.

There are a few possible causes to this problem:

1. The beam formed by the sensor could be far wider close to the sensor than expected. The real beam would only form at for example a distance of 10 meters and beyond from the sensor, which would be too much for this project. Close to the sensor the beam would look like Figure A.2.

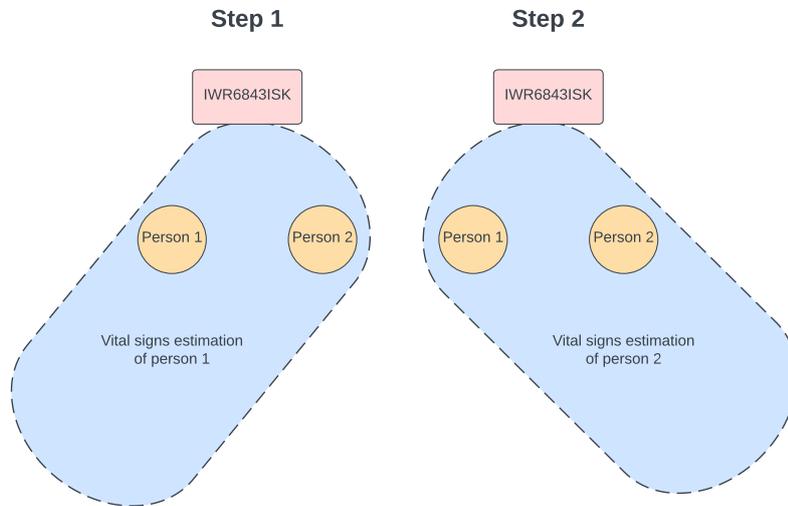


Figure A.2: **How the Tx beamforming performed during testing.**

2. Since the Tx beamforming functionality in the mmWave SDK has not been tested by the TI developers, there could be a malfunction in this technique causing it not to work properly.
3. Because of the limited documentation, I could have implemented this functionality in the wrong way. I did extensive research on this function, but because there was almost no documentation or examples available for this function, it was mostly a case of try and error.

In the end, it is still uncertain where the implementation of this idea went wrong. When after a lot of tries this idea still didn't work, I decided to change gears and focus on another method of getting the vital signs estimated for multiple persons at the same time. Luckily, this method did work, and you can read about all the details in the main part of this thesis.