Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Realtime Capacity Utilization

Group Thesis

Demian Jäger (jaegerde@ee.ethz.ch)
Lukas Sigrist (sigristl@ee.ethz.ch)

Distributed Computing Group
Computer Engineering and Networks Laboratory
ETH Zürich

Supervisors:
Samuel Welten
Prof. Dr. Roger Wattenhofer

June 28, 2013

Abstract

This thesis focus on counting people in a specified area by analyzing the local WLAN traffic. We use the fact that the majority of the Swiss are in possession of a smartphone.

We examined the traffic and found out that this counting method provides good results for recurring events but does not perform well in unique or rapidly changing situations.

Contents

Chapter 1

Introduction

Every student knows this situation: it's nearly twelve o'clock and you can't wait until the lecture has finished and you can go to lunch. But there will be thousands of other students heading to the same place. Wouldn't it be convenient to know the number of people in the canteens and where you get your food the fastest? This is only one of many situations where it is useful to know how many people are located in a place. Think of getting the last free seat in public transportation at rush hour. Or even more critically: where you have to direct all the people to when you have to evacuate a whole football stadium.

Counting people in a specific public area is the central part of our thesis. As given in the examples before, this could be the amount of people in a public building, in a park, in a train, in a bus or in many other possible places. The major problem of counting people in such a place is that the method should not restrict anyones freedom. In the best case, the counting isn't noticed by anyone.

There are many different approaches of counting or estimating the number of people in a specified area. To count the number of people taking the subway, for example, one could use turnstiles which give an exact number of passing passengers, but would be too restricting for only counting the number of people. An alternative to that would be light barriers which are accurate when everyone has to pass through a narrow enough point. However, if this method should be applied to a massively open area, the installation of many separate light barriers to count the passing people accurately would cost a lot of money. Additionally, there is still the issue of detecting whether someone is entering or leaving the area (this could be solved with two barriers behind each other, but this would double the costs again).

Another approach is the analysis of the video material from observation cameras. The people density could be estimated by image analysis and could be projected to the whole area in order to estimate the number of people. Because of the extrapolation of the measured people density, this method is less accurate than the both mentioned before, and requires a complex surveillance camera system and a lot of computational power to analyze the images in realtime. To

make the analysis more accurate one could take more and more cameras into account, which, in turn, would drive up the costs a lot.

A third method of measuring the number of people is not to count each human being itself, but only to count their mobile phones. The number of mobile phones indicates the approximate amount of people (as statistics show, 92% of the Swiss own a mobile phone[1]). To identify the number of mobile phones, the mobile phone communication could be recorded and analyzed in order to count the number of unique devices. However, analyzing the GSM traffic requires a complex infrastructure and a lot of computational power to keep track of the devices when they change their radio base station.

The approach we are following in this thesis: instead of analyzing the mobile phone communication, we only analyze the WLAN traffic of smartphones. As recent surveys show, two third of the population in Switzerland own a smartphone[2]. And a key feature of every smartphone is its WLAN capability. While observing the WLAN traffic, we recognized that smartphones are periodically sending so-called probe requests even if no WLAN network is reachable. We decided to analyze the WLAN traffic to estimate the number of smartphones and extrapolate it to get the number of people in the observed area. The advantage of this method is that it doesn't disturb people at all and works even in observing widely open areas. Compared to the analysis of the mobile phone communication, the hardware requirements are very low and by far not as much computational effort is needed. The hardware installation also doesn't require a lot of space and is easily installed in a predefined environment. However, as we can only recognize a certain part of all present people, the tradeoff we had to make is that it doesn't provide very accurate results.

This thesis is divided in three different parts: First of all, we have a closer look at what the WLAN standard says about our detected probe requests and whether there are other possibilities to detect attendant WLAN devices. Subsequently, we present the hardware and software environment we used to capture and analyze the WLAN traffic. At the end, we present results we investigated and other interesting information we gathered during our longterm WLAN traffic observation.

For capturing WLAN traffic and testing our analysis methods we installed a test system in our famous canteen called Gloriabar.

[1] Swiss Federal Statistical Office , `http://www.bfs.admin.ch/bfs/portal/de/index/themen/16/04/key/approche_globale.indicator.30103.301.html`, state of 2010

[2] bonus.ch, `http://www.bonus.ch/RDP-20130115_DE.pdf`, 15.01.2013

Chapter 2

WLAN Standard

In this section we describe the parts we used of the WLAN standard. We utilized the characteristics of this standard for our counting method. These are defined in the IEEE 802.11 standard.

2.1 Probe Requests

2.1.1 What is a Probe Request

A WLAN enabled device is permanently (even if already connected to a certain network) searching for so called basic service sets (e.g. your home WLAN router, or an access point of a company network). The device is hoping to find a better basic service set (BSS), for example one that corresponds to your home network, or one that corresponds to the network which the device is already connected to but has a better signal strength.

There are two different ways of finding a BSS: the active and the passive scanning mode.

Each BSS is sending so-called beacons in which it advertises its basic properties (e.g. 'Here you can connect to the ETH network'). In passive scanning mode the WLAN device is just listening. If it sees a beacon of a BSS which is of higher priority than the one which it is currently connected to, it will try to connect to that BSS.

In active scanning mode, the WLAN device is sending so-called probe requests in order to check if there are networks available. The BSS answers with a probe response. In the IEEE 802.11 standard, there is no time range specified how frequently this probe requests have to be sent[1]. In general, it isn't specified whether a device should be in active or passive scanning mode. The active scanning mode doesn't even have to be implemented. Both methods have their advantages and disadvantages.

[1] IEEE Standard 802.11-2012, page 108

2.1.2 Why we Capture Probe Requests

Even though it is not compulsory in the standard, most of the smartphones (especially devices on which runs Android or iOS) are at least partially in active scanning mode. Thus the amount of devices in active scanning mode is an indicator of the currently present smartphones.

In every probe request the WLAN device is sending its own MAC address to which the BSS replies with a probe response. Since the MAC address is unique, we can assume that as long as we capture probe requests from a specific origin this device must be near our capturing device.

2.2 Acknowledgment

2.2.1 When are ACKs sent

Acknowledgments (ACK) are sent by the receiver of a data packet to the sender of the packet in order to inform the sender that the packet was transmitted correctly.

2.2.2 Why we Capture Acknowledgments

Since in some places (especially at our testing place) most of WLAN enabled devices are connected to a network, this devices don't send probe requests permanently. For this devices, we check whenever a BSS acknowledges a packet. In this case it is not important who sent the packet (it is the MAC address of the BSS) but to whom it is addressed.

Chapter 3

Counting System Implementation

Our system is divided into two main parts. One part is our local capturing device which records WLAN traffic and sends the captured data to the server. The other part is a server which receives the data, analyze the data and shows the results on a webpage. An overview is given in Figure 3.1. In order to improve the results in a larger area, there can be used multiple capturing devices running simultaneously but transmitting data to the same server.

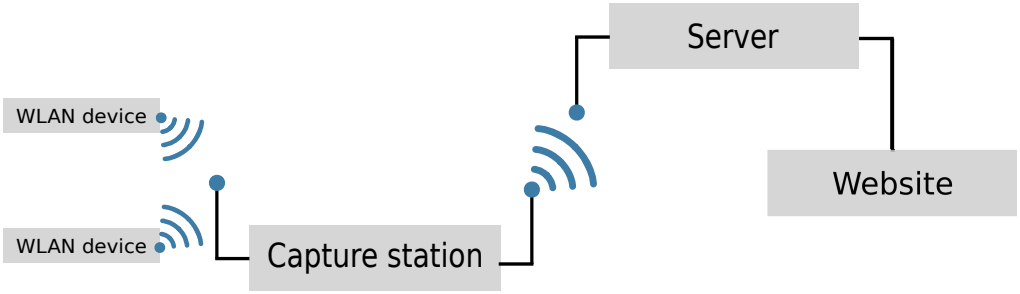

Figure 3.1: The counting system captures WLAN packets with its capture station, transmits the data to the server which presents the analyzed data on a website.

3.1 Capturing Device

The capturing device (here sometimes called station) we have chosen is a Raspberry Pi with two WLAN sticks attached to it over an USB connection. One of the WLAN sticks is used to capture the WLAN packets and one thread of the capturing program is writing this packets into a buffer. A second thread of the program simultaneously transmits the buffered packets over the second WLAN

stick to the server. A schematic overview of the capturing device is given in Figure 3.2.

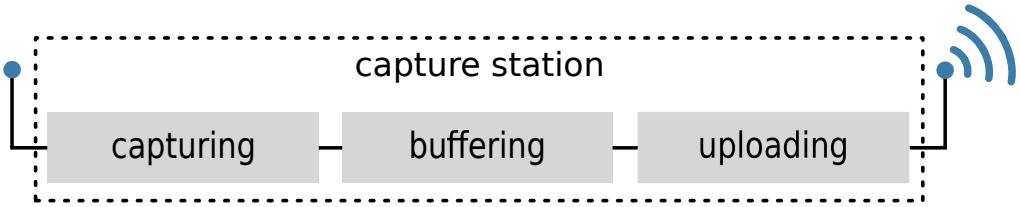

Figure 3.2: The capturing thread of the capturing station writes the data into the buffer. The buffer removes duplicated packets, while the uploading thread is transmitting data from the buffer to the server.

3.1.1 Capturing Thread

The capturing thread decides whether to keep a captured packet and forward it to the buffer or not. Since we decided to use only packets of the types probe request and ACK, this thread ignores all other packets. Additionally, this thread deletes all information of the packets which isn't relevant for our thesis. Only the MAC address, the capturing time, the signal strength and the packet type indicator are kept.

3.1.2 Buffer

In order to reduce the amount of data transmitted to the server, the captured data is written into the buffer.
Since probe requests are sent on multiple channels, they arrive several times at nearly the same time at the buffer. To prevent additional overhead, the buffer ignores new packets if the same address was already added within the last five seconds.

3.1.3 Uploading Thread

If the buffer is full or a specified time since the last upload has elapsed, this thread will empty the buffer and will transmit its data to the server by using a HTTP-post-request. If the server is not reachable for a certain time, the data will be written into a file which can later be transfered to the server by hand.

3.2 Server

The server has three main functions:

- store received data from the capturing device in a database
- analyze the captured data
- show statuses of the capturing stations and analyzed data on a webpage

3.2.1 Receiving Data from Stations

The server has to handle the uploads which are made from different stations. After receiving a packet from a station, it inserts the packets in the database with the following information: station which received the packet, packet type, capture time and the MAC address. In addition to this, there will be saved a flag which denotes whether the data should be used (e.g. we ignore MAC addresses which belong to routers or other devices which are permanently in our area and doesn't seem to belong to a person).

3.2.2 Analyze Data

This part of the server analyzes the data in the database and saves the results periodically. How the analysis works will be explained in Section 3.3.

3.2.3 Website

The website is able to display statistics over the past hours and the current estimated amount of people in the area. In a password protected area it is shown which of the stations currently is online and transmitting data. There is also an interface which allows to import earlier backed up data.

3.3 Analysis

The analysis of the captured data is done per time interval, which is moved by a constant shift. Figure 3.3 gives a short overview about the following explanation of the analysis. The shift between the interval start times is 30 seconds, as this is accurate enough for an analysis with a reasonable performance. The length of the analysis interval depends on different factors, which have to be defined previously to get good results. Before defining the analysis interval length, we have to go deeper into how we analyze the captured packets.

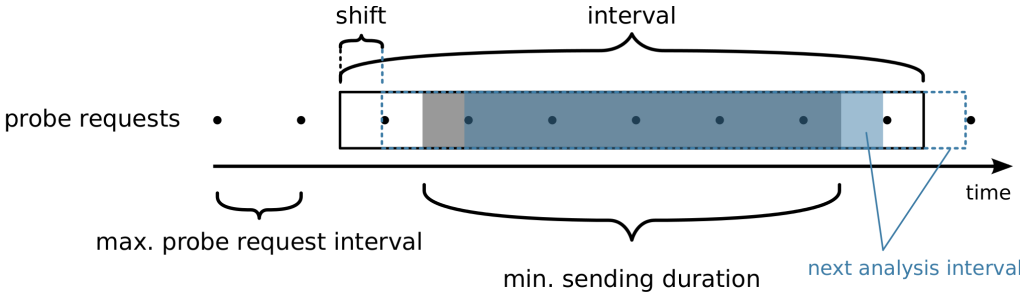

Figure 3.3: The analysis is done per time interval, which is shifted by a fixed amount of time. The interval itself has a length of the minimum sending duration plus twice the maximum probe request interval. A device is considered as present if it's transmitting packets at least as long as the minimum sending duration.

Our goal is to get the devices (identified by their MAC addresses), which attended our observed area for a specified amount of time. We are going to call this duration minimum sending duration. The minimum sending duration is the shortest time a device has to be in the area to be considered as present by our algorithm. With this minimum sending duration we can eliminate devices, which are passing through our area and are captured only a short time. It is obvious that the analysis interval has to be at least as long as our minimum sending duration. In addition, we have to consider that the probe requests are sent with a fixed period. Because the period start times are randomly distributed, we need to add twice the time of the maximum probe request period. This has to be done to ensure that we consider probe requests over at least the minimum sending duration, even in the worst case (last probe request is immediately before the interval start and the most recent isn't considered in this interval).

In Figure 3.4 you see the distribution of the average interval between two probe requests per device (below 10 seconds, we dropped repeated probe requests, above 3 minutes, we considered the device isn't available anymore). We've chosen a maximum probe request interval of 70 seconds, as it covers over 90% of all devices. This leads to our final interval length of minimum sending duration + $140s$.

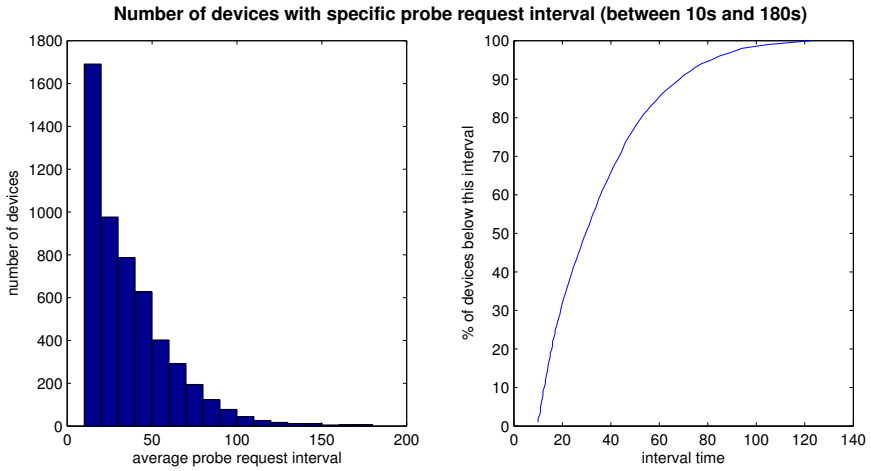

Figure 3.4: The distribution of probe request intervals shows that most of the WLAN devices which send probe requests do this at least every 70 seconds.

The analysis itself is done for each interval and all shift seconds. We analyze for each device occurring in the interval the duration of its availability and then split the devices into two groups: the counting devices which consists of all devices meeting the defined minimum sending duration restriction. All the other devices are identified as passing devices. In the end, we store the number of counting and passing devices. In addition we save the number of devices which are newly considered as counting devices and the devices which aren't considered anymore as counting devices, both compared to the analysis interval immediately before.

With the analysis results and a counting of people by hand, we can determine a constant factor to extrapolate our number of counting devices to an estimate of the number of attending people. This and the tuning of the interval parameters are discussed when presenting the results of our experiments.

Chapter 4

Experiments

4.1 Testbed

To test our system we installed two capturing stations in a canteen at ETH called Gloriabar. The advantage of this place is that it has only one entry. This made it relatively easy to count people entering and leaving by hand. The instability of the wireless network is one of the disadvantages. Therefore, we had to use backups what made it hard to provide realtime data. Another problem were the people who just checked the daily menu and then left again. This fact leads us to use a large minimum sending duration which causes a delay.

4.2 Raw Data

During two months of analyzing WLAN traffic, the stations submitted around three million packets of almost 25 thousand different MAC addresses excluding all the addresses which were only observed once. 200 devices are listed on a blacklist because they are either routers or other devices which are permanently present in the canteen or in the reach of the capturing devices.

Just looking at the raw data provided in Figure 4.2 there cannot be drawn any conclusions on the amount of people in a particular area. Also counting every device just once per interval doesn't show more than an increasing trend at 11:45 and there cannot be made any statement about absolute values (see Figure 4.3).

Therefore, we needed a more differentiated analysis method as presented in Section 3.3. Applying this algorithm there appear more useful information, as we will see in Section 4.3.

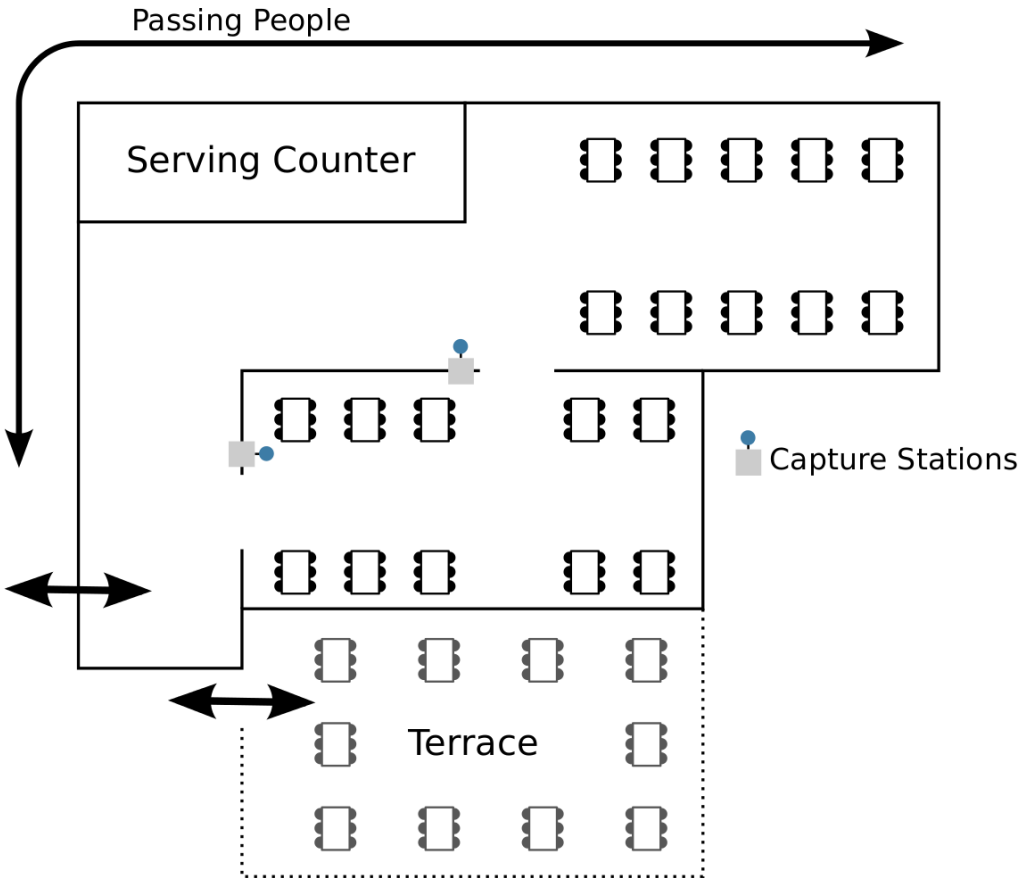

Figure 4.1: This schematic building plan gives an overview of our testbed. A big problem for the analysis were the people which were passing by on the depicted walkway.

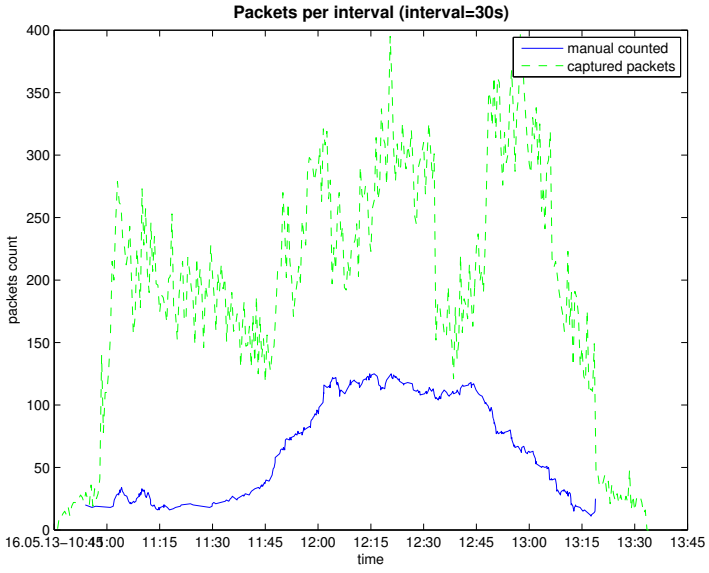

Figure 4.2: Just counting the amount of packets in a certain interval (here 30 seconds) doesn't show any helpful information.

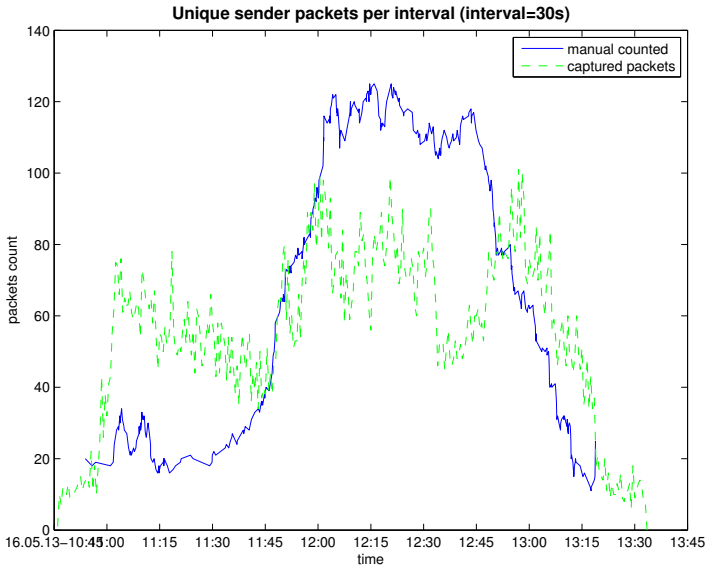

Figure 4.3: In contrast to Figure 4.2, every MAC address was just counted once per interval.

4.3 Results

4.3.1 Estimated Present Persons Versus by Hand Counted Persons

Our analysis method lead us to Figure 4.4 and Figure 4.5. Both figures show a comparison between the data we counted by hand and the data we estimated by observing the captured packets.

As mentioned above, people which were just passing by had a falsifying influence on our counting method. While Figure 4.4 shows the manual counted data versus an estimate with a minimum sending duration of 30 seconds, Figure 4.5 shows the same data with a minimum sending duration of 5 minutes. Note that in Figure 4.5 the amount of passing devices is significantly higher at beginning and end of lunch time, as a lot of people are passing by.

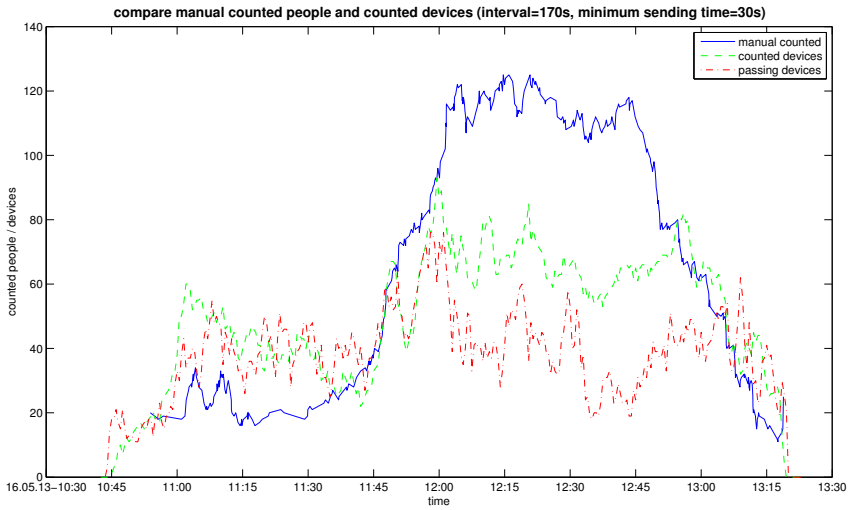

Figure 4.4: With a minimum sending duration of 30 seconds there are considered too many people counted which were just passing by on the way next to the canteen or just checking the menu.

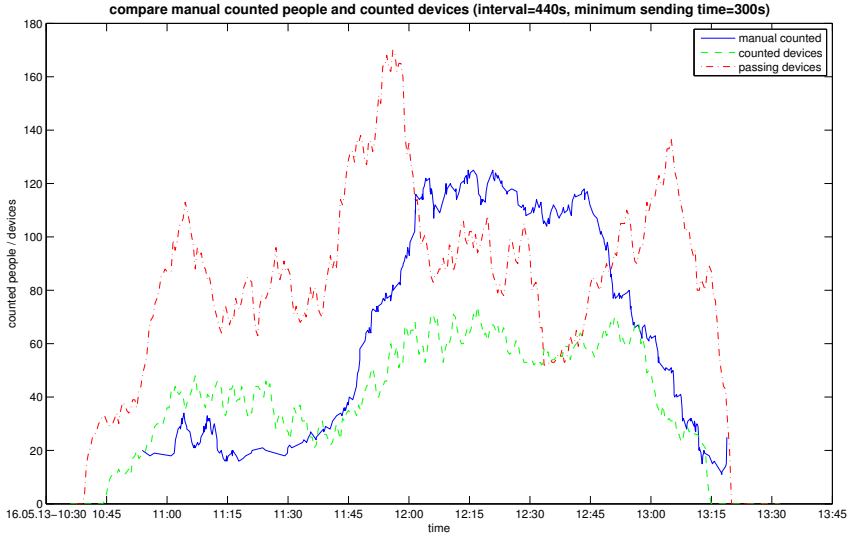

Figure 4.5: With a minimum sending duration of 5 minutes the amount of people considered as passing by is significantly higher than in Figure 4.4.

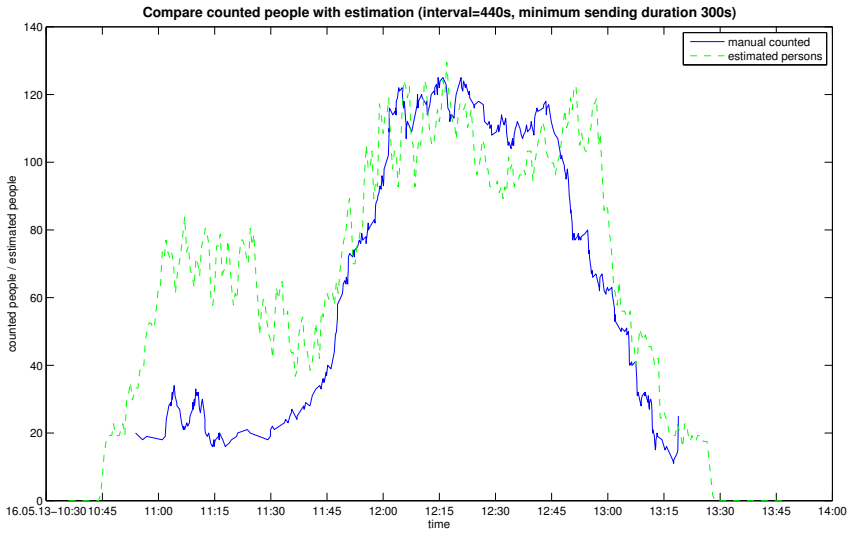

Figure 4.6: Assuming that one out of 1.75 persons carries a device with enabled WLAN, the estimate is fairly close to of the actual amount of people.

Figure 4.6 shows the same data as Figure 4.5 but with a scaling factor of 1.75. This means we suppose that one out of 1.75 persons carries a WLAN capable device with enabled WLAN. Once again, people staying near our observed area falsify our results. The increase of number of people at 11:15 could come from the fact that a library is located under the canteen and people were leaving.

4.3.2 Signal Strength

Even though we recorded the signal strengths of the packets, we didn't use them for our analysis as they were very volatile. There are several reasons for this highly varying values:

- different devices are sending with different signal strengths (laptops usually have more sending power than smartphones)
- if there is much other traffic and noise present, more sending power is required for reliable transmission of data
- the focus of the capturing antenna can have a large impact

Despite these facts, there is a remarkable correlation between the signal strength and the devices which are considered as passing. Figure 4.7 shows this correlation but also that the signal strength is varying over time. Due to the fact that the gap between the signal strength of the counted devices and the one of the passing

devices is small, it makes it hardly possible to decide whether we have to accept or refuse a device. Note that this information is not of importance to us since it wouldn't improve our results.

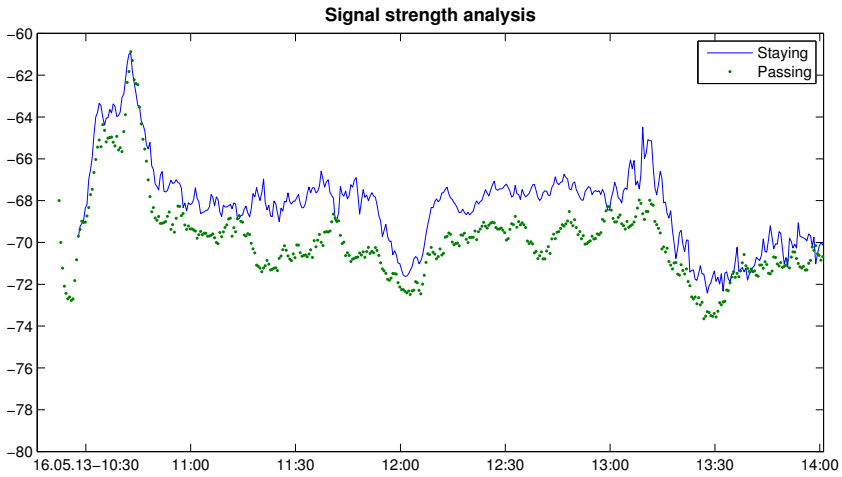

Figure 4.7: The signal strength of the devices which are considered as passing is nearly everywhere below the signal strength of devices which are counted.

4.3.3 Daily Average

As seen in Subsection 4.3.1 (e.g. Figure 4.6), the estimate can be very inaccurate. Figure 4.8 shows an average over the workdays. One can see that the lecture breaks are at xx:00 to xx:15 and many students are getting their coffee in the canteen.

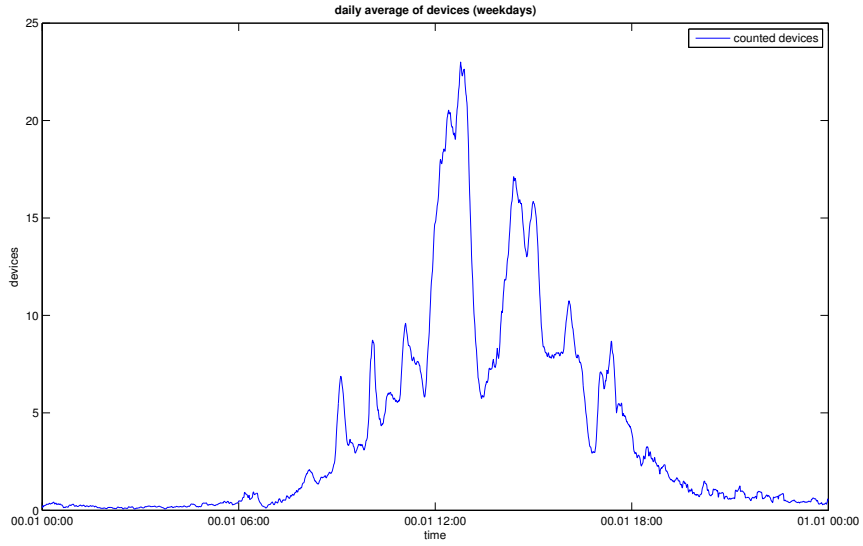

Figure 4.8: This average of counted devices on workdays shows the increase of present people during lecture breaks and during lunch time as well as the fact that the canteen is empty during the night.

4.4 Conclusion

As one can see in the daily average analysis, the captured data is reasonable. The analysis has shown that there are more people in the canteen during the lecture breaks and at lunch time. As expected, our canteen seems to be empty during night. Extracting the absolute number of people at one specific moment turns out to be a fault-prone task, especially at places where many people are just passing by. However, the analysis shows convincing trends and with a scaling factor the number of attending people can be estimated more or less accurately.

In conclusion, we state that this counting method provides good results for recurring events but doesn't perform that well in unique or rapidly changing situations.