

Spooft proof GPS timing

*A detection and mitigation system for GPS
time spoofing*

Aril Johannes Schultzen



Thesis submitted for the degree of
Master in Informatikk: programmering og nettverk
60 credits

Spooft proof GPS timing
Faculty of mathematics and natural sciences

UNIVERSITY OF OSLO

Autumn 2016

Spoof proof GPS timing

*A detection and mitigation system for GPS
time spoofing*

Aril Johannes Schultzen

© 2016 Aril Johannes Schultzen

Spoof proof GPS timing

<http://www.duo.uio.no/>

Printed: Representralen, University of Oslo

Foreword

Five years ago I started studying for my bachelors degree in computer science at Oslo and Akershus University college of applied sciences. I had no idea what I was doing and I can clearly recall a professor telling us that "you are here because you want to learn how to program". At that point, this was news to me. Five years later I am submitting my master thesis and all I want to do is thank the people who enabled me.

Thank you, Harald Hauglin for guiding me when I doubted myself and for being patient with me and my uninformed questions. Thank you, Sverre Holm for being my supervisor and for being honest when I needed it. Thank you, Tim Dunker for providing excellent assistance and support. Thank you, Bao Chi Tran Nguyen for teaching me how to program without letting me feel utterly useless. And thank you, Ingvild Totland Pettersen for being my loving partner and for dealing with all my frustration and immaturity.

Abstract

The goal of this project was to create a prototype of a smart atomic clock controller running on commodity hardware, capable of spoof proofing a GPS controlled atomic clock, i.e harden it against GPS jamming and spoofing attacks. The project aims to use a multi-layered approach to evaluate the integrity of GPS signals and to use knowledge about the atomic clock and its behavior when GPS disciplined, to detect and mitigate a spoofing attack. Many solutions have been proposed to solve the problem of GPS time spoofing [1], but as far as we know this is an issue rarely tackled using commodity hardware and simple techniques. We aim to demonstrate the efficiency of network enabling commodity GPS receivers and connecting them to a atomic clock controller by using a client/server architecture facilitating GPS and atomic clock data analysis. The preliminary GPS manipulation tests demonstrates the feasibility of our proposal.

Contents

1	Introduction	1
2	Background on GPS	3
2.1	Clocks	4
2.2	GPS signals and Time	4
2.3	Threat Models and countermeasures	6
2.3.1	Threats	6
2.3.2	Countermeasures	7
2.3.3	Summary	11
2.4	Benchmark clock spoofing attack	12
3	Our Proposal: Spoof proof atomic clock controller	13
3.1	Filtering and steering	13
4	Hardware	15
4.1	Atomic Clock	15
4.2	Atomic clock controller platform	15
4.3	GPS receiver	16
4.3.1	GPS receiver configuration	16
5	Atomic clock controller architecture	17
5.1	Client/Server model	17
5.2	Architecture description	18
5.2.1	Security	20
6	Atomic clock controller implementation	21
6.1	The Sensor server	21
6.2	Implementation description	21
6.2.1	Server core	23
6.2.2	Clock model implementation	26
6.2.3	Parser and Handler	28

6.2.4	GPS filter	32
6.2.5	Clock model filters	33
6.2.6	GPS data	34
6.2.7	Client connection	34
6.2.8	Shared memory & Semaphores	35
6.2.9	Client list memory segment	36
6.2.10	Client linked list	38
6.2.11	Client input parser	39
6.2.12	Ready check algorithm	40
6.3	Sensor Server Protocol	41
6.4	Atomic clock Communication	41
6.5	Data structures	42
6.5.1	server_data	43
6.5.2	server_synchro	43
6.5.3	command_code	43
6.5.4	nmea_container	44
6.5.5	list_head	45
6.5.6	transmission_s	45
6.5.7	client_table_entry	45
6.6	The Sensor Client	47
6.7	Interfacing	48
7	GPS manipulation tests	51
7.1	Logging filter and model data	51
7.2	Setup	52
7.3	Test one filter limits	53
7.4	Test one	54
7.4.1	Goal of test one	54
7.4.2	Description	54
7.4.3	Observations	56
7.4.4	Timing measurements	66
7.4.5	Test one results	68
7.5	Test two	68
7.5.1	Goal of test two	68
7.5.2	Change in setup	68
7.5.3	Test two filter limits	68
7.5.4	Description	69
7.5.5	Observations	70
7.5.6	Timing measurements	72
7.5.7	Test two results	72
7.6	Unplanned disturbance	74

8	Discussion	78
8.1	Test results	78
8.2	Shortcomings in current implementation	78
8.2.1	Resizing shared memory segments	78
8.2.2	Atomic clock management	79
8.2.3	External MJD calculation	79
8.2.4	External Atomic clock communication	80
8.3	Choice of programming language	80
8.4	Alternative approaches	81
8.4.1	Single computer, many GPS receivers	81
8.4.2	Store in database and analyze	81
9	Conclusion	82
	Appendices	84
A	Acknowledgments	85
A.1	Contributions	85
B	Clock model	86
B.1	Introduction	86
B.2	Input data for the clock model	86
B.3	Smoothing of sampled frequency steering data and estimates of the clock state	88
B.4	Phase jump filter – fast timing filter	89
B.5	Frequency correction filter	90
C	Data acquisition	92
C.1	CSAC Logger source code	92
C.2	GPS Logger source code	93
C.3	GPS Logger source code	95
D	Sensor server software	97
D.1	Client	97
D.2	Server	104
E	Scripts	179
E.1	Logger setup schematic	181
F	E-mails	182
F.1	Correspondence with Mr. Davis	182

G Figures	183
Complete Bibliography	186

Chapter 1

Introduction

Timing in the context of accurate and stable reference frequencies as used in timestamps, is both critical and valuable in the world of IT infrastructure as well as in the world of industrial control systems. The timing itself is often generated by clocks disciplined by GPS. This allows an inaccurate clock to obtain the long term accuracy and stability of a much more expensive and accurate clock. The only problem is that the GPS disciplining relies on an antenna with a clear view of the sky and uses a known civil GPS code structure. Since the code structure is known and antennas can not be hidden, a GPS receiver can be "spoofed" with a GPS-like signal generated by an attacker. GPS based timing may therefore be viewed as an unencrypted and physically unsecured port into industrial control systems.

An example of an application relying on GPS derived time is a phasor measurement unit (PMU). A PMU analyzes the waves on the electrical grid and uses a common time source for synchronization. This synchronization allows for real-time measurements between multiple points in the grid by multiple PMUs. The common time source (and why PMUs are relevant) is often obtained by using GPS. [2] The value of such a device is understood clearer by recognizing that the power grid is a complex, interconnected, interdependent network. In other words, errors and abnormalities in one part of the grid will have an effect on operation elsewhere in the grid, and in some cases lead to whole spread blackouts [3]. When the clocks in PMU applications are spoofed, they may end up causing damage in the network instead of protecting it [4].

GPS time is also used in telecommunication to synchronize base stations. The radio spectrum used by cellular phones is limited, making synchronization between base stations important in order to maintain efficient use of the spectrum. The ability to accurately time-stamp financial transactions is made possible all over the world using GPS and is crucial for traceability and

accountability [5].

Our goal in this report was to develop a spoof proof atomic clock controller prototype by combining multiple GPS receivers, a chip scale atomic clock and a client/server architecture. This allowed us to test some strategies for spoofing detection, knowing if the GPS signals are unsuitable or unsafe, and mitigation, how to maintain accurate timing when reliable GPS synchronization is unavailable.

Chapter 2

Background on GPS

The Global Positioning System (GPS) is a utility owned by the United States that provides its users with positioning, navigation and timing services. At the end of the sixties, the U.S Navy was developing the Polaris missile, a missile capable of being launched from a submarine. One of the requirements for launching the Polaris missile was exact knowledge of the submarines position. The problem led the Navy and The Applied Physics Laboratory at John Hopkins University to develop the Transit system, the earliest predecessor to the GPS system [6].

Today, roughly 40 years later, we are surrounded by GPS technology. In fields like emergency response, search and rescue, fleet management and even agriculture, it has become a vital tool of utmost importance to everyday operation. Satellite navigation can be found in most new cars and few phones are today sold without an internal GPS receiver. The European Space Agency estimated that there were 2 billion GPS enabled devices by 2012 [7]. What started out as a navigation tool for the U.S navy is now used by millions, if not billions of users both civilian and military all over the globe. A common misconception (that is often reinforced by Hollywood action movies) is that the GPS satellites track *you* by communicating with your GPS receiver. It actually works the other way around. You are, with your GPS receiver, tracking a set of satellites in order to establish your own position. At any given time, there are at least 24 GPS satellites each in its own orbit at about 20,372 nautical miles above your head [8]. In order for a GPS receiver to determine its position and obtain correct time, it will need four GPS satellites within line of sight ¹. The method used by your GPS receiver to determine its position is called *trilateration*. Trilateration is used in geometry as a process of determining the location (absolute or relative)

¹The line of sight requirement might seem unreasonable, but by the time the signal has reached earth, it has degraded to a minimum of -160 dBW [9]

of a point by measuring distance. It is often confused with triangulation which instead of distance, uses angles. Measuring the distance from the GPS satellites to a given position on earth is quite simple when using the equation:

$$Distance = Rate \times Time \tag{2.1}$$

The equation is simple to solve, first we need the rate. In this context the *rate* is how fast the signals travel. This is equal to the speed of light (299,792,458 m/s). The time the signal has used traveling from the satellite to earth can be obtained by analyzing the signal itself. A simple and slightly inaccurate description is that the signal contains a "time stamp" of when the signal was sent. By comparing this time stamp with the current time one can calculate the age of the signal and therefor how long it has spent traveling. This is explained in greater detail under (2.2) [10].

2.1 Clocks

What does a 10 USD wristwatch and a 100,000 USD atomic clock have in common? They do not stay accurate forever. This phenomena known as *frequency drift*, is when a clock no longer runs at the exact same speed as a reference clock and they drift apart. This property is a result of how they track time. In essence, all clocks work in the same way. They have a part that oscillates, a way to count the number of oscillations and a way to show the count. If we transfer this analogy to the typical "grandfather clock", the pendulum would be the oscillator, the counting mechanism the clockwork and the clock face and dials would be the display. In a typical wristwatch, the oscillator is a quartz crystal powered by a battery. The frequency of which the crystal oscillates is then divided down to a single Hertz by simple electronics. The purity of the crystal is among the decisive factors determining the accuracy of the clock. [11]. Although a completely different beast, the same principles apply to the atomic clock which uses the microwave radiation that electrons in atoms emit when they change energy levels. One of the most commonly used elements in atomic clocks, is *caesium-133*, an isotope of caesium.² [12].

2.2 GPS signals and Time

During the introduction of this chapter the properties of GPS as a tool for navigation was made apparent. This is however not the only use of GPS, it

²1 second equals 9,192, 631,770 cycles of the Cs-133 transition

is also used for timing. The GPS satellites transmits a *Coarse/Acquisition (C/A)* code and a restricted *Precision (P)* code. The C/A code is freely open for everyone and is transmitted at the L1 carrier frequency (1575.42 MHz) and the P code is transmitted at both L1 and L2 (1227.60 MHz) and is reserved for the military. The C/A code is a 1023 bit pseudo random code that is transmitted at 1.023 Mbit/s, which means it repeats itself every millisecond. Each satellite transmits a different pseudo random code, codes that does not correlate well with each other. This is important because it makes it possible to separate the satellites from each other. The way the receiver calculates its position was briefly mentioned earlier and is better explained here. The receiver calculates the distance from itself to the satellites by comparing the pseudo random code received from the satellite with an identical one that it generates itself. The receiver "slides" these codes over each other further and further in an attempt to match them up. The signals travel time is determined by how far the codes had to be slided before they matched. This is what is called *Code-phase GPS* and it is not without issues. Since the codes have a wide cycle width, almost a microsecond, there is still a significant uncertainty even though the codes match. At the speed of light a microsecond wrong translates to a roughly 300 meter error in the solved position. What many receivers do is that they start with the code-phase and move on to using measurements based on the carrier frequency. Since the frequency is much higher the uncertainty within the match decreases thus increasing the accuracy dramatically. This is what is known as *Carrier-phase GPS*. These signals can be seen as range signals, used to measure distance. This is however not the whole story. The GPS signals also include *navigation messages* like *ephemeris* as well as something called the *almanac*. The ephemeris data is information containing the orbit of every GPS satellite in the constellation. This is used by the receiver to calculate the satellites position. The almanac is a state report of the whole constellation. Alright, but what about time? We have already established that the key to GPS is measuring the travel time of a radio signal, but considering the consequences of a couple of microseconds of slack when dealing with light-speed, it is really putting some pressure on a GPS receiver's internal clock. As previously mentioned, all your receiver needs to do to find its position in a three dimensional space, are three GPS satellites. If the GPS receivers' internal clocks were perfect, the three satellite ranges would intersect at a single point, your position. But in the real world our clocks are everything but perfect. One could use atomic clocks in the receivers but that would make the receivers too expensive (even though chip scale atomic clocks (CSAC) are becoming increasingly affordable ³) for anyone to buy. The solution is to make a fourth measurement from a fourth satellite. This measurement will not intersect

with the first three when using an imperfect clock. The receiver can then try to find a correction factor it can subtract from its timing measurement in order to make the measurements intersect. By doing so, it also brings the receiver's clock back to sync with universal time. With the correct time it can also make correct and precise positioning. [13]

2.3 Threat Models and countermeasures

The threat models and countermeasures presented in this paper are based on the article *Reliable GPS-Based Timing for Power Systems: A Multi-Layered, Multi Receiver*[1]. The only exception is our own proposed countermeasure described in chapter3.

2.3.1 Threats

2.3.1.1 Jamming

By emitting a high-power signal at the frequencies used by GPS satellites, one can interfere with the signals received by the GPS receiver, effectively denying GPS receivers use of these signals. These signals are already weak considering their travel from space. Such an "attack", although effective, is pretty naive and easily recognized by the jammed party. If your equipment is operational and you do not have a signal, you are probably being jammed.

2.3.1.2 Signal-level Spoofing

Signal-level spoofing is when an attacker causes a receiver to lose lock on an authentic GPS signal by overpowering it with a false signal. This can be achieved by using a GPS simulator that matches the authentic signals phase, code delay and encoded data [14]. Knowing the signal that the victim is receiving is important in order to successfully spoof it. To anyone with access to the military-grade GPS signals this is less of an issue since military-grade signals are encrypted and harder to spoof. The civilian frequencies on the other hand are publicly known and readily predictable. Shepard, Humphreys and Fansler (2012)[3] describes in their paper *Evaluation of the Vulnerability of Phasor Measurement Units to GPS Spoofing Attacks* a way to successfully spoof a GPS signal used by a PMU. They describe how they "introduce" the counterfeit signal to the victim receiver by adjusting the power of the signal below the victim receiver's noise floor and then gradually raise it until it surpasses the authentic signal's strength. Once the victim's receiver locks on, the attacker has gained full control.

2.3.1.3 Data-level Spoofing

In data-level spoofing the contents (data) of the GPS signal is manipulated. GPS signals include ephemeris data used to solve the positions of each satellite in orbit and also the time and status of the satellite constellation. By altering this data the receiver solves incorrect velocity, location and most important in this context, clock offset [14].

2.3.1.4 Replay spoofing

Replay spoofing (or *meaconing*³) is a technique where GPS signals are intercepted and rebroadcasted. The rebroadcast can be delayed and used to confuse navigation or to cause delay in applications relying on GPS signals for time.

2.3.1.5 Malfunctions

Just like any tool or device a GPS receiver is prone to failure. This threat may not be posed by an external party, but is still a threat to normal operation. The ability to differentiate between an attack and a malfunction is important when deciding how to respond to such an event.

2.3.2 Countermeasures

2.3.2.1 Monitoring Signal Power

In any kind of attack, jamming or spoofing, a counterfeit signal must overpower the authentic signal in order for the receiver to lock onto it or in the case with jamming, denying access to the authentic signal. By monitoring the strength of the signal and detecting a spike or rise in signal power, a possible attack can be identified. This is a low-cost, low-complexity and independent (in contrast to for example using other receivers as a reference) countermeasure. It is however because of the unpredictable nature of signals, not considered to be a detection confident countermeasure and should therefore only be used along side other countermeasures.[1]

2.3.2.2 Checking solved position against known position

By checking the position solution against the known position of the receiver, both receiver errors and a replay spoofing attack can be detected. It does

³*Meacon* is portmanteau of *Masking Beacon*

however fall short when more sophisticated techniques like Data and Signal-level spoofing are used. These types of attacks when done properly (unless it is done with intention) will not alter the solved position. It is important to note that this is only relevant when only using *one* receiver. If the position solution from multiple receivers deployed in the same area are cross-checked, this countermeasure can still be considered effective. Consider the following scenarios when using three receivers:

- **None of the receivers are spoofed:** Each receiver's solved position matches their respective known position. They all solve the same time.
- **One or two receivers are spoofed:** The spoofed receiver(s) solve(s) different time compared to the receiver(s) not being spoofed.
- **All the receivers are spoofed:** As long as they are spoofed by the same spoofer, they will solve the same time but also the same position which again makes it possible to detect the attack.

A possible way for an attacker to avoid detection would be to use one spoofer per receiver. These spoofers would need to be synchronized and their signal power fine tuned to make sure that they only spoof their respective receiver. It is believed that such an attack would be too complex and costly to be considered practical [1].

2.3.2.3 Checking time solutions against receiver clock statistics

By comparing statistics created by monitoring the receiver's clock with the time solution, one can detect spoofing, as well as malfunctions. This is because the time solution is unlikely to be consistent with the statistics in event of an attack. Since this countermeasure relies on the receiver's clock which can be described as both unpredictable and stochastic, it should only be used along side other countermeasures [1].

2.3.2.4 Cross-checking navigation data among receivers

When under a data-level spoofing attack the navigation data is modified. By comparing one GPS receiver's navigation data with another, both data-level spoofing and malfunctions may be detected. This countermeasure can also prove useful during jamming attacks. The jammed receiver could use the data from other receivers in the event that it is unable to correctly decode navigation data, but still able to track satellites. This may enable the receiver to continue operation during an attack. [1]

2.3.2.5 Comparing navigation data and reverse-calculated satellite positions

The PMU's GPS receiver is never moved and its position is always known. By using their pseudorange measurements the satellites' positions can be reverse calculated by using trilateration. Since the reverse-calculated positions only match the positions calculated from the navigation data when both pseudorange and navigation data are correct, one can effectively detect replay spoofing and malfunctions. It is also worth noting that this countermeasure increases the difficulty of both signal and data-level spoofing, because it narrows down the possible valid (seemingly) spoofing signals. [1]

2.3.2.6 Cross-correlating P(Y) code

This countermeasure assumes two receivers with at least 1 km distance from each other that tracks a signal from a satellite visible to them both. It is also based on the assumption that the encrypted military P(Y) code cannot be forged by a spoofer. The receivers use the C/A code phase and timing relationship to the P(Y) code to obtain two samples from the same time frame of the received P(Y) code and then correlate the two samples. Even though the samples will be encrypted, noisy and perhaps distorted by narrow-band RF front-ends, a high correlation peak should be created when a cross-correlation is conducted as long as the receivers are not spoofed. A key conclusion of the research made by L. Heng (2013) as referenced by L. Heng *et alia* (2014)[1] was that the probability of detection errors using this method decreased exponentially with the length of the samples made from the P(Y) code and the number of receivers used as reference. This method has therefore proved itself effective against spoofing attacks, but ineffective against replay spoofing because the rebroadcast uses authentic GPS signals with correct P(Y) code. It is important to note that the implementation of this countermeasure relies on the GPS receiver's ability to output baseband samples and the samples' ability to be transferred over a data network. Because the sampling rate of the samples are fairly high, it is recommended that the spoofing detection is done periodically instead of continuously [1].

2.3.2.7 Position Aided (PIA) Tracking loops

Vector tracking is a receiver architecture that combine the tasks of signal tracking and position/velocity estimation into one algorithm. This is a contrast to the traditional way where the tracking methods track satellites independently as well as the position/velocity solution independently. Even

though this requires more computing power it increases immunity to interference and jamming. The vector tracking is aided by the fact the we know the PMU's GPS receiver's true location. The tracking robustness can be further improved by using a Kalman filter. Since a PMU and its GPS receiver remain stationary, the parameters of the tracking loops can be chosen to narrow the loop filter bandwidth which reduces noise and the effective radius of a potential jamming attack. Replay spoofing attacks will also fail since the PIA vector tracking depends on the knowledge of the GPS receiver's true position. In the event of such an attack the result would be that the vector tracking will fail to function [1].

2.3.2.8 Multi-receiver tracking loops

Building on the idea from *PIA Tracking loops* one can benefit from the networked nature of the GPS-timed PMU. In a multi-receiver vector tracking loop many receivers process information in collaboration. A key conclusion of the research made by A. Soloviev *et alia* as referenced by L. Heng *et alia* (2014)[1] showed that acquisition and tracking performance under low signal-to-noise ratio conditions were improved under multi-receiver signal accumulation. Multi-receiver phased arrays also improved the robustness against both jamming (2.3.1.1) and spoofing attacks (2.3.1.2,2.3.1.3) by "*Forming beams to satellites and steering nulls in the direction of attacking transmitters*" ([1], p.41). In addition to the increased robustness, it increases the ability to detect malfunction. A faulty receiver will usually not be consistent with other correctly functioning receivers. As with the countermeasure based on cross-correlating P(Y) code, this implementation also requires that the GPS receivers are able to output baseband samples. In this implementation the samples need to be transmitted continuously among the receivers which requires a capable data network such as a typical LAN. [1]

2.3.3 Summary

The table (2.1) shows the different threat models and the effect of the countermeasures discussed.

Table 2.1: The table shows the effectiveness of the covered countermeasures against threat models.

Counter Measures	Threat Models				
	JAM ⁴	SLS ⁵	DLS ⁶	RS ⁷	MF ⁸
Monitoring Signal Power (2.3.2.1)	N	X	X	X	N
Check pos. solution (2.3.2.2)	N	Y	Y	Y	Y
Check time solutions (2.3.2.3)	N	X	X	X	X
Checking nav. data (2.3.2.4)	X	N	Y	N	Y
Reverse calculated sat. pos. (2.3.2.5)	N	X	X	Y	Y
Cross-correlating P(Y) (2.3.2.6)	N	Y	Y	N	N
PIA TL (2.3.2.7)	Y	N	N	Y	N
Multi-receiver TL (2.3.2.8)	Y	X	X	X	X

Table 2.2: Legend for table (2.1)

Y	Effective	N	Ineffective	X	Auxiliary
---	-----------	---	-------------	---	-----------

⁴Jamming (2.3.1.1)

⁵Signal-level Spoofing (2.3.1.2)

⁶Data-level Spoofing (2.3.1.3)

⁷Replay Spoofing (2.3.1.4)

⁸Malfunctions (2.3.1.5)

2.4 Benchmark clock spoofing attack

In 2012 a team from The University of Texas at Austin published a paper [3] describing *The Civil GPS Spoofer*. It was a GPS spoofer and the first of its kind. The Civil GPS spoofer was able to precisely align both the C/A codes and navigation data in a counterfeit signal with those of an authentic GPS signal. The alignment capability allowed for sophisticated spoofing attacks that were hard to detect. The spoofer was implemented in software-defined radio with a digital signal processor. The user could control the fake GPS signal specifying both navigation and timing. By tracking and acquiring GPS signals and calculating navigation data, it could produce up to 14 spoofed L1 C/A signals by performing real-time prediction of the pseudo random C/A codes. Key features of the time spoofing attack:

- Seamless takeover: GPS spoofer tricks target into locking on to a replica of the authentic GPS signal, without loss of GPS lock or change in reported position.
- Once the target is locked on to spoofer signal, the attacker manipulates the apparent time of the target GPS clock.
- Low quality internal clocks enable aggressive manipulation of GPS clock timing, rapidly resulting in large apparent timing errors.

The sophistication of this spoofing attack is in part due to its technical complexity, but more importantly due to the demonstrated result: The attacked GPS controlled clock had no way to detect that it was being fed false GPS signals. The study was targeted to the use of timing to phasor measurement units used in the monitoring and control of power grid transmission lines. Similar aggressive manipulation of timing may potentially cause the malicious shutdown of transmission lines within a few minutes after the onset of the attack. The team tested the spoofer against a wide variety of applications using the civilian receivers, and they were always successful [3]. When designing spoof proof GPS systems the Civil GPS spoofer can be considered as a benchmark reference threat.

Chapter 3

Our Proposal: Spoof proof atomic clock controller

We propose to construct and use what we call a *Spoof proof atomic clock controller*. The spoof proof atomic clock controller consists of custom software running on commodity computer hardware, connected to at *least* two GPS receivers and an atomic clock. With the spoof proof atomic clock controller you can perform:

- Spoofing detection.
 - Having a stable clock that can be trusted makes it possible to verify the GPS timing solution.
 - Using two or more GPS receivers makes it possible to detect whether or not the signals originate from satellites in orbit or from a spoofer on the ground. See subsection 2.3.2.2 for more about this idea.
- Mitigation.
 - A stable clock will provide accurate timing for an extended time even in the absence of valid GPS correction.

3.1 Filtering and steering

The spoof proof atomic clock controller uses what we call *filters* for detection. The filters are algorithms used to detect a spoofing attack or just general abnormalities that might affect applications relying on GPS time. The filters can be divided into two classes or types. They are either GPS based or clock

model based. The GPS filters analyzes collected GPS data called NMEA data [15] and verifies their validity. This can be achieved by comparing the collected GPS data with known GPS data for the receivers. The clock model based filters are more sophisticated. By analyzing the behavior of the atomic clock used by the spoof proof atomic clock controller, a model can be built and used as a reference. This approach was actually suggested in subsection 2.3.2.3, but in that instance using the receiver's internal clock. This approach was considered auxiliary because of the receiver's clock's poor performance. This is in contrast to our proposal where we have a stable and accurate atomic clock. This makes it possible to evaluate the GPS signal used to discipline the atomic clock. The model can also be used for mitigation during an attack. Once an attack has been detected the atomic clock should no longer be disciplined by the signal. The clock should be usable for a while without disciplining, but with the clock model used for steering it will provide accurate timing for a longer time.

Chapter 4

Hardware

4.1 Atomic Clock

We decided to use the Symmetricom SA.45 as the atomic clock. This is an atomic clock measuring only 16cc with 1 pulse per second (PPS) output and 1 PPS input for disciplining. The SA.45's strength is its low power consumption (less than 120mW) and low price. The SA.45 also uses a built-in controller which can be communicated with over a RS-232 serial interface. The ability to communicate with the atomic clock, issue commands and collect data is paramount for the feasibility of our proposal. It is worth mentioning that any atomic clock, such as Cesium standard or even a Rubidium standard, could be used given that they have a means to communicate basic telemetry like phase difference and steer values and can configured by wire to change modes of disciplining. For more about clock performance, review the SA.45s' user guide [16] and data sheet [17].



Figure 4.1: Symmetricom SA.45s CSAC. Courtesy Symmetricom.

4.2 Atomic clock controller platform

We chose to cast the Raspberry Pi 3 Model B (RASPI3) in the role as the host running the atomic clock controller software. The RASPI3 is an interesting piece of equipment with an impressive list of specifications. It is a single board computer with a 1.2GHz 64-bit quad-core ARMv8 CPU, 1 GB of RAM, built-in 802.11n Wireless LAN and four USB ports [18]. As

with the Symmetricom SA.45, the RASPI3 is very affordable. The RASPI3 retailed at about 35 USD when this report was written. We also propose to use Raspbian [19], a Debian derived flavor of Linux optimized for the Raspberry Pi, as the operating system.

4.3 GPS receiver

We chose to use at least two GPS receivers. Both of the receivers should collect data and feed it to the atomic clock controller, but one of the receivers should also double as a 1 PPS disciplining source for the atomic clock. Considering the need for a stable 1 PPS source, we propose to use the u-blox NEO-M8T. This is a relatively affordable GPS receiver with a temperature compensated crystal oscillator (TCXO), 3 concurrent GPS reception and an external antenna [20]. In the current implementation of the atomic clock controller, only NMEA data is collected from the GPS receivers (see section 6.2.6 for more about NMEA data). However in the future it might be beneficial to collect and process raw data¹ from the receivers as well. Since most GPS receivers today follow the NMEA standard (to some extent) and raw data currently is not required, common and popular receivers like the u-blox NEO series should be more than sufficient for use in our implementation.

4.3.1 GPS receiver configuration

According to the u-blox NEO-M8T's manual [21], the device includes a feature called *Fixed Position*. This is a feature that must be enabled in order to put the device in *Time Mode*. This feature makes the device solve time with higher certainty even with fewer available satellites. The *Time mode* was not used, relied on or accounted for, in our solution.

¹The ublox M8T is capable of outputting RAW data. This is for example used by RTKLIB (<http://www.rtklib.com/>) instead of NMEA

Chapter 5

Atomic clock controller architecture

When discussing how to implement the spoof proof atomic clock controller's software the following features and requirements were identified:

- Speed and efficiency is important in order to detect attacks as early as possible.
- Extensive logging capabilities for forensics.
- Easy and fast access to all gathered data.
- Should support as many GPS receivers as possible.
- Easy interfacing. Should be possible to operate over Telnet/SSH.
- Easy to configure.
- Graphical user interface is not a requirement

5.1 Client/Server model

In order to connect many GPS receivers to the atomic clock controller, we decided to implement a client/server model. The atomic clock controller serves the same purpose as before, but doubles as a server. The data transmitted, will still be the same, the only difference is the interface and media. A network interface will be used instead of USB and the media can be whatever is available since TCP/IP is designed to be hardware independent. This opens up the possibility for using:

- Twisted pair
- Fiber optics
- WiFi
- 4G (cellular)

Because reaction time is a concern, high latency media is not recommended¹. The GPS receivers that we have chosen to use do not have a network interface, we instead use Raspberry PIs. A GPS receiver is connected to a Raspberry PI thus giving it a network interface. Considering how cheap single board computers have become, the price of a Raspberry PI can be justified even just to network enable a GPS receiver. The latest model of the Raspberry Pi has even got a built in wireless network interface. See 4.2 for more about the Raspberry Pi 3. The combination of the GPS receiver and the Raspberry PI are in this report abstractly seen as one device. We call this device a sensor. The sensor is responsible for reporting data to the Server and not much more. The sensor can be deployed using already existing infrastructure to communicate, thus making it easier to deploy than if each GPS receiver had to be cabled directly to the atomic clock controller. There is no denying that the implementation of a Server/Client model greatly increases the complexity of the atomic clock controller software, but it makes up for it by eliminating the need for long signal cables and amplifiers. We call this approach the *Sensor Server* model.

5.2 Architecture description

In order to connect the Sensors to the server, the atomic clock controller's software needs to be able to:

- Handle connections to clients.
- Update structures as clients' status changes (disconnects or gets kicked out).
- Receive data from clients

These are tasks that the atomic clock controller needs to be able to perform as well as controlling the clock and analyzing data. In order to implement the Server/Client model, the Server is implemented using the Linux Socket API.

¹IP over Avian Carriers as described in RFC 1149 is highly discouraged

The API is based on BSD sockets and is available in almost all Unix-like operating systems ([22], p.610). A socket is a handle that can be passed by a program to the network API in order to use the network connection. The plan was briefly to use Glib [23], a library that provides building blocks for libraries and applications written in C, instead of using low-level sockets. This idea was scrapped because I felt that it would be overkill for this implementation.

Having decided *how* the clients should connect to the server, the next challenge becomes how the server should communicate with the clients and vice versa. We decided to use what is called *blocking I/O* and *fork()*. This is a very common approach, but not the fastest ([24], p. 188). It is however quite easy to implement. Blocking I/O means that read operations on the socket block the main thread in the process until data has been received. This is not as dramatic as it might seem: If a Socket call cannot be completed immediately the process who issued the call will be put to sleep thus enabling the scheduler to schedule other processes for execution until conditions are right for the sleeping process ([24], p.435). Fork() is a system call used to create a new process. The new process is called a child and the process that called the fork() system call is called the parent. The child process is a duplicate of the parent process. The alternative to using fork() is to create a thread. The creation of threads are typically less expensive in terms of CPU cycles than the creation of processes, but presents their own challenges. Processes always have their own virtual address space as opposed to threads who share their address space with the other threads within the process. This makes programming with threads more complex. The result of a thread crashing may have a more severe impact on the other threads within the same process since they all shared the same data. The new process is used by the server to handle the newly created connection. This means that for every connected client, there is a process created. This of course means that this solution does not scale awfully well. This is however not a big problem for us since the time taken for a sensor to connect to the server really does not matter. Once all sensors in a setup have connected, they should not disconnect unless taken down for maintenance or replacement.

Using processes presented us with a new challenge, interprocess communication. A process is not aware of its neighbor. For all it knows, it is the only process running. Every time a client connects, a new process is born to take care of the communication with the new client. Since every process has its own virtual address space, the processes are isolated from each other. This posed a challenge because we wanted the atomic clock controller to collect data from sensors for processing. Inter process communication (IPC) is nothing new, and one way to accomplish it, is to use shared memory segments that are accessible for all the processes. The processes store the data

they have received from their respective clients in the shared memory segment. That way it is easy for one process handling a sensor to check up on data received from another sensor. This kind of functionality was seen as valuable, but is not utilized to the extent it was envisioned. If more filters are added, the architecture would probably make more sense as the current ones do not really use the shared memory to cross check data.

Using shared memory is by no means the only way to implement interprocess communication. Another approach to inter process communication that was considered, was using a **pipe**. The pipe can be seen as a unidirectional channel where data written to end of the pipe is buffered by the operating system until the data is read from the other end of the pipe. This would however mean that each process would not only have to listen to the client connected, but also to the pipes connecting them to the other processes. This would be the case with message passing and sockets as well. In contrast, the shared memory approach allowed for the server to act as a single unit even though many processes are at work. Allowing multiple processes to share the same memory segment is like asking for trouble. At some point the program will suffer from race conditions and unexpected behavior. The shared memory segments are therefor protected with semaphores.

The atomic clock controller does not have graphical user interface. Seeing how it was going to operate much like a service, the need for a graphical user interface was not prioritized. A user can interface with the system by logging on using telnet, and issue commands. In order to separate the intention of a connected client, be it to deliver GPS data or to send commands to system, roles were made. A client connected to the server may have two roles. It can choose to either be a sensor or a monitor. The sensor role is already explained. The Monitor role was added in order for a user of the system to connect to the server and check status or issue commands. For a client to assume the role of a Monitor, the client has to pick a negative integer as ID number. This way, the server does not expect you as a client to report any NMEA data the way it would with a sensor. The monitor role can also be used to interface with the server. See (6.7) for more.

5.2.1 Security

Implementing good security is not an easy task and should not be taken lightly. The goal when developing the Sensor server was to produce a rough prototype that could prove our concept. Time was therefor allocated towards producing features that would realize this goal instead of attempting to implement security mechanisms that would have been flawed anyway because of time constraints.

Chapter 6

Atomic clock controller implementation

6.1 The Sensor server

Figure 6.1 shows a simplified block diagram of the sensor server and its tasks. server tasks include:

- Handle connections to clients.
- Update structures as the clients' status' changes (disconnects or gets kicked out).
- Communication with the atomic clock and updating the atomic clock model.
- Sensor data analysis and filter updates.
- Raising alarms based on filter status.
- Controlling the atomic clock.

6.2 Implementation description

In the following section, the architecture and inner workings, data structures and key components of the Sensor server will be explained. The server core (6.2.1) consists of the source code in `sensor_server.c`. The Parser and Handler spans over `textttsession.c` and `textttactions.c`. The atomic clock model and the associated filters' source code can be found in `csac_filter.c`. The filter using GPS data is in `filters.c`. Figure 6.3 shows a simplified call graph for the Sensor server.

Sensor server overview

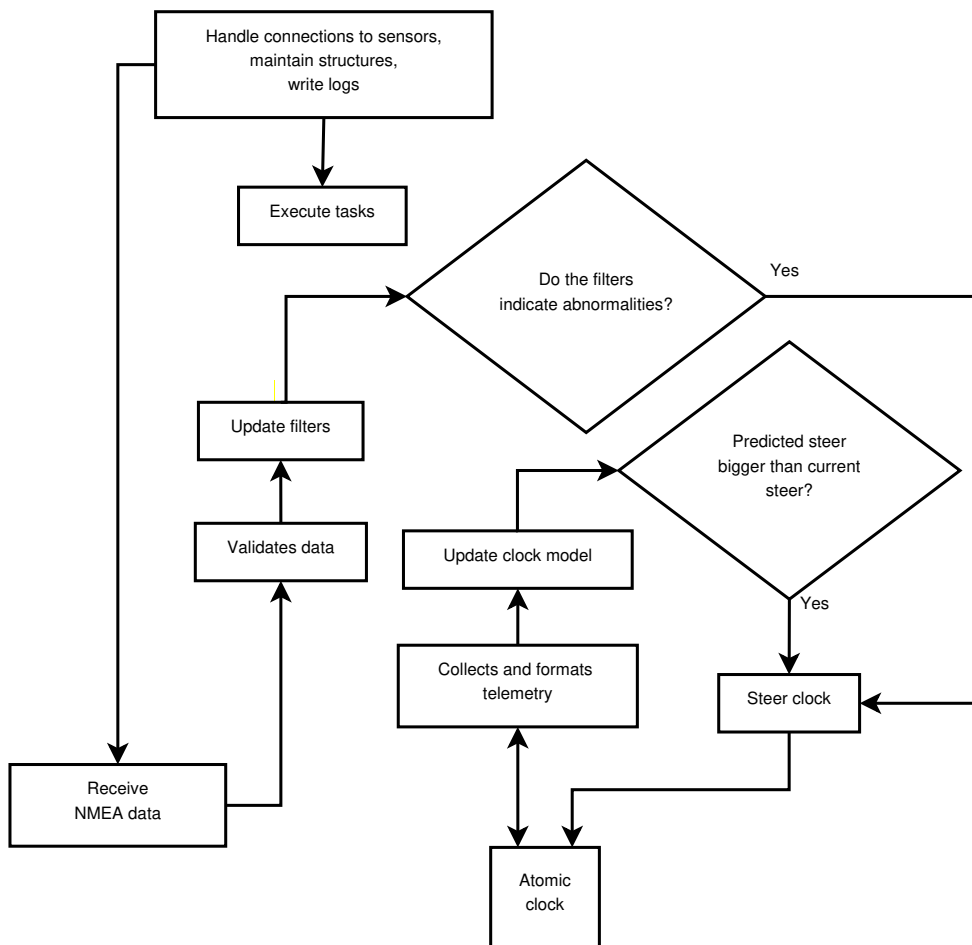


Figure 6.1: A block diagram of our proposed solution

6.2.1 Server core

The "Server core" is the main process and the parent of every process created during the life of the Sensor Server. It spawns the process maintaining the atomic clock model as well as new processes for every client that connects. Figure 6.2 is a block diagram of the Server core.

- The Server software takes one parameter to start, the port. If the parameters are missing or the parameters are illegal, a string containing usage information is printed and the program exits. Because the Server also is responsible for communication with the atomic clock, it needs the rights to access the serial port connected to the atomic clock ¹.
- The configuration is initialized and loaded. If the configuration fails to load the Server prints an error message and exits.
- The `client_list`, a shared memory segment containing a linked list containing all the clients, is initialized. See section 6.2.9 and 6.2.10 for more about the client list shared memory segment and linked list implementation.
- Shared memory segments and semaphores are initialized and allocated. If the allocation fails, the Server prints an error message and exits. See section 6.2.8 for more about semaphores and shared memory.
- The process responsible for maintaining the atomic clock model and filters is forked out. If the fork fails the server prints an error message and exits. See section 6.2.2 for more about the atomic clock model.
- Handler for SIGINT (interrupt) SIGTERM (terminate) and SIGCHLD (child process is interrupted or terminated) are registered.
 - When a process exits, a SIGCHLD signal is sent to the parent. When a client disconnects from the server the process that handled the client exits. Ideally, the client disconnects from the server by sending a protocol compliant "disconnect request". This way the server can handle the disconnect in a controlled manner and update its client list when the client disconnects. However, this is not always the case and if the client disconnects abruptly and without a warning, the server still needs to handle it. Therefore,

¹The easiest way to achieve this is to run the program as `sudo` or login as root. The best solution would probably be to add the user to a group with the right permissions. This is often the `dialout` group.

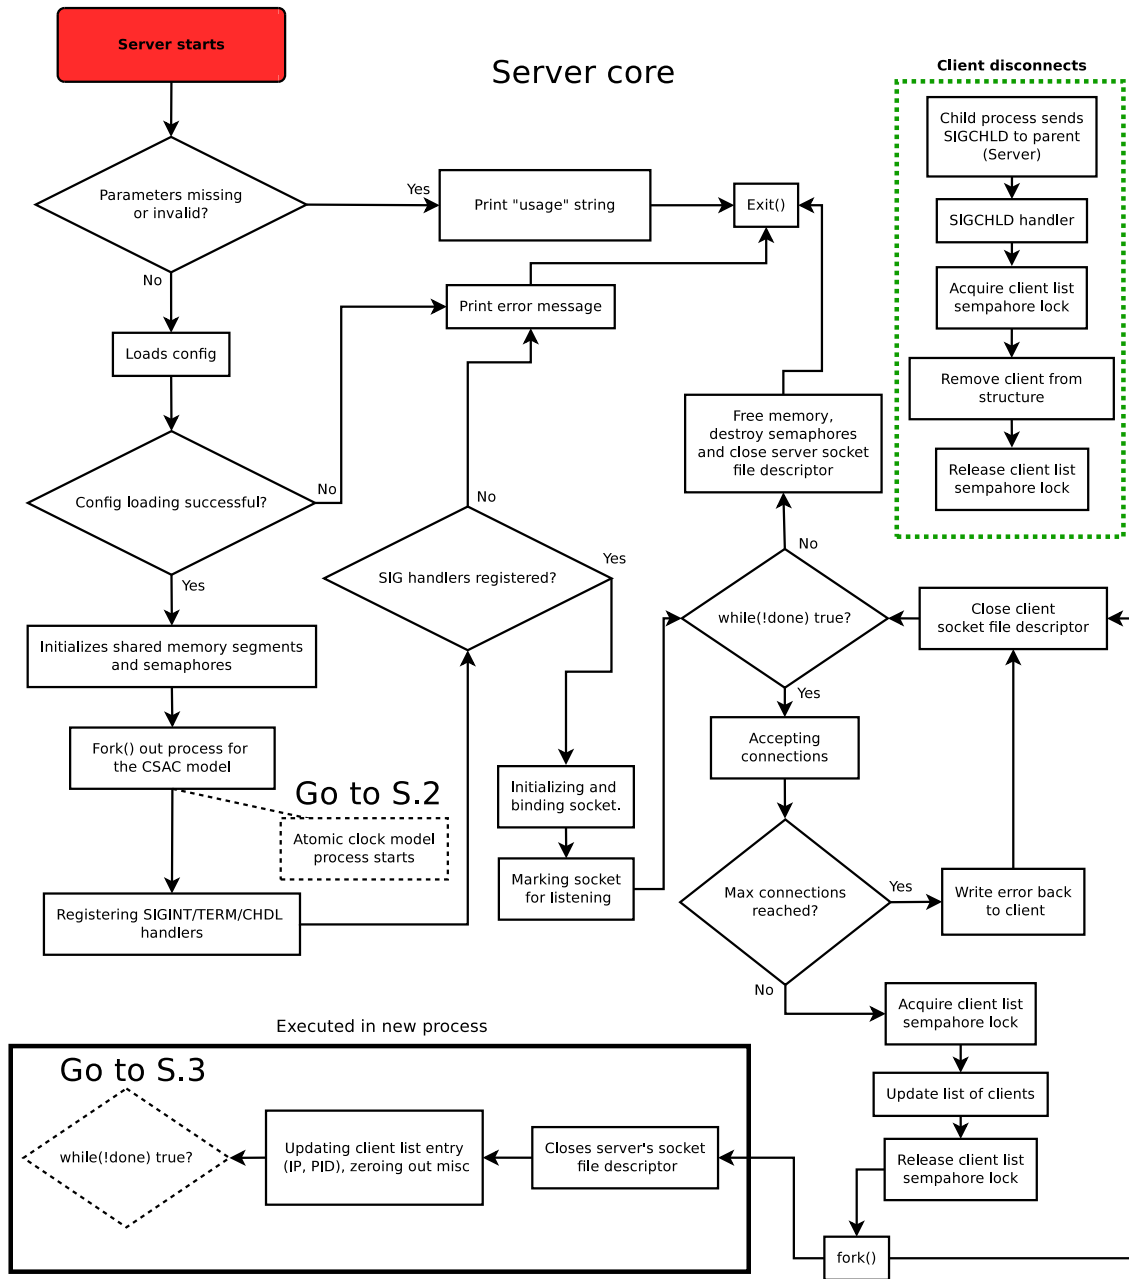


Figure 6.2: The block diagram shows an abstracted view of the Sensor Servers Core.

Simplified call graph

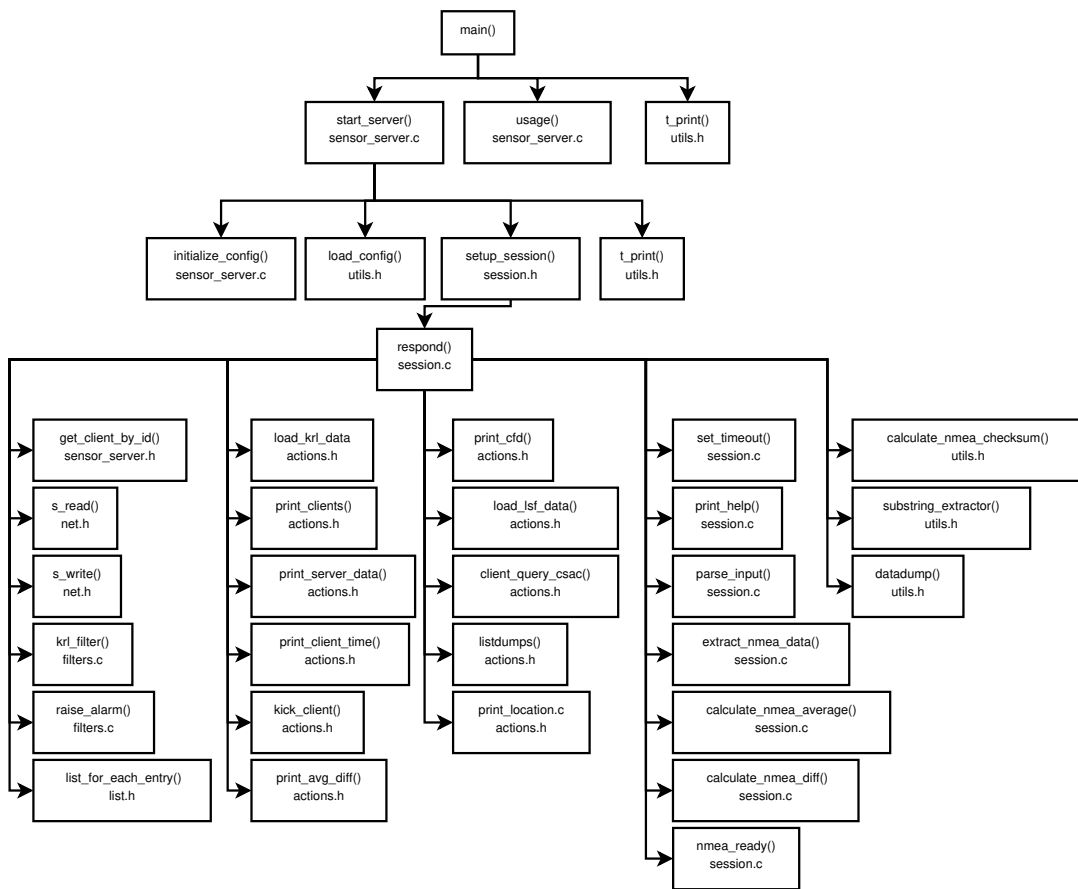


Figure 6.3: Simplified call graph for the Sensor Server.

when a SIGCHLD signal is received, the server iterates through its client list and finds the client whose PID matches the sender of the signal and removes it from the list. In figure 6.2.1 this routine is the part in the dotted box.

- The socket is initialized and marked for listening. See section 6.2.7 for more about sockets.
- A loop is entered. The loop breaks when `volatile sig_atomic_t done` equals 0. If the loop is broken, the semaphores are destroyed and all allocated memory is freed, the servers socket file descriptor is closed and the Server exits. See section 10 for more about variables used in synchronization mechanisms.
- When a client connects the server checks if the maximum number of clients has been reached. If this number has been reached an error message is written back to the client and the client's socket file descriptor is closed. If the maximum number of clients has not yet been reached, the client list is updated and a fork is performed. The parent process then resumes the loop.
- The process that just forked out for the new connection closes the server socket file descriptor it inherited from its parent.
- The new process then updates its client table entry and proceeds to the listening loop. This is covered in section 6.2.3. See section 6.5.7 for more about the client table entry data structure.

6.2.2 Clock model implementation

Conceptually, the atomic clock model and the filter that uses the model are two separate tasks. In practice however, they are intertwined as figure 6.4 suggests. See section 4.1 for more about the atomic clock and appendix B for the model itself.

- The configuration for the model is loaded. The configuration includes values defining the accepted limits for steering and phase difference. Optionally, it can also include state values from an "earlier" model ² that can be used to initialize the model further, thus reducing the warm-up time. If the loading of the configuration fails, an error message is

²The functionality to restore an earlier model is not without restrictions. It is typically meant for a situation where the Server needs to be taken down for a couple hours and a new warm-up period is considered undesirable. This is better explained in section B

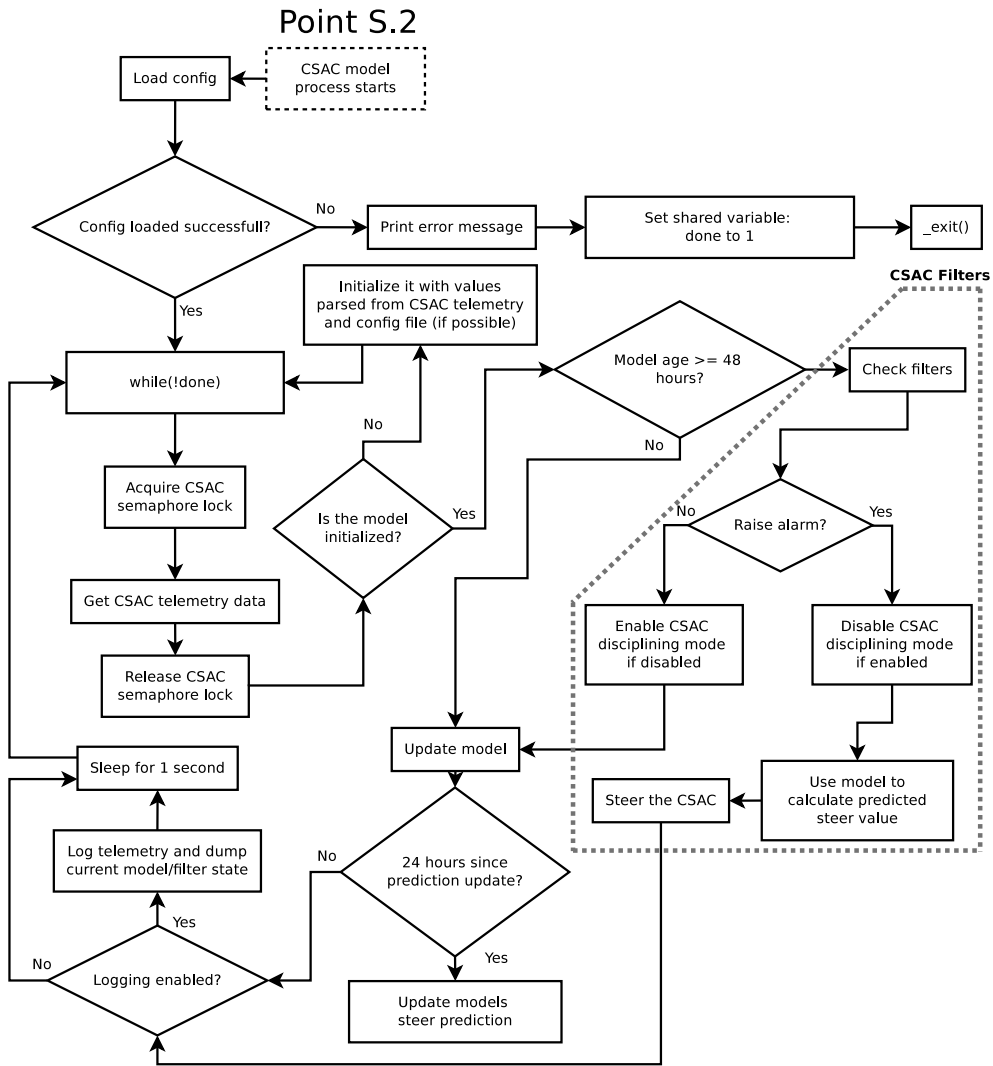


Figure 6.4: Block diagram showing the flow of execution in the atomic clock model

printed, the `volatile sig_atomic_t done` is set to 1 and the process exits. This also breaks the server core (6.2.1) loop. After all, the system needs the atomic clock to do the job it is intended for.

- A loop is entered. The loop breaks when `volatile sig_atomic_t done` equals 0.
- The atomic clock lock is acquired and once the telemetry is queried from the atomic clock it is released again.
- If the model is not initialized, telemetry data is used to initialize it. State values from earlier models will also be used if available in the earlier loaded configuration. See section 6.4 for more about communication with the atomic clock.
- A check is done to see if the model has been running for at least 48 hours. If it has, it is ready to be used as a reference for the filters. See section 6.2.5 for more about the atomic clock filters in particular.
- The telemetry received from the atomic clock is used to update the model.
- A check is done to see if 24 hours have passed since the prediction values were last calculated. The model keeps track of steer predictions for yesterday and today. By using these two points of data, an estimation of the predicted steer can be calculated.

6.2.3 Parser and Handler

The parser and handler are responsible for parsing data received from a client and making sure every request is protocol compliant. Once the goal of the request has been established, the request is executed. A function named `respond()` is called in the listening loop. Every time data is received from a client, the data is sent to the parser (`parse_input()`) to determine the validity of the request received. See subsection 6.2.6 for more information about the GPS data received from the Sensors. Figure 6.5 and 6.6 are block diagrams showing the execution flow for this part of the server. The protocol is explained in greater detail under section 6.3 and the parser is better explained under subsection 6.2.11.

- A loop is entered. The loop is broken when `volatile sig_atomic_t done` equals 0. This is the same variable that is used in the Server core 6.2.1 loop and for a good reason: If the server exits, its children should too.

Parser and handler

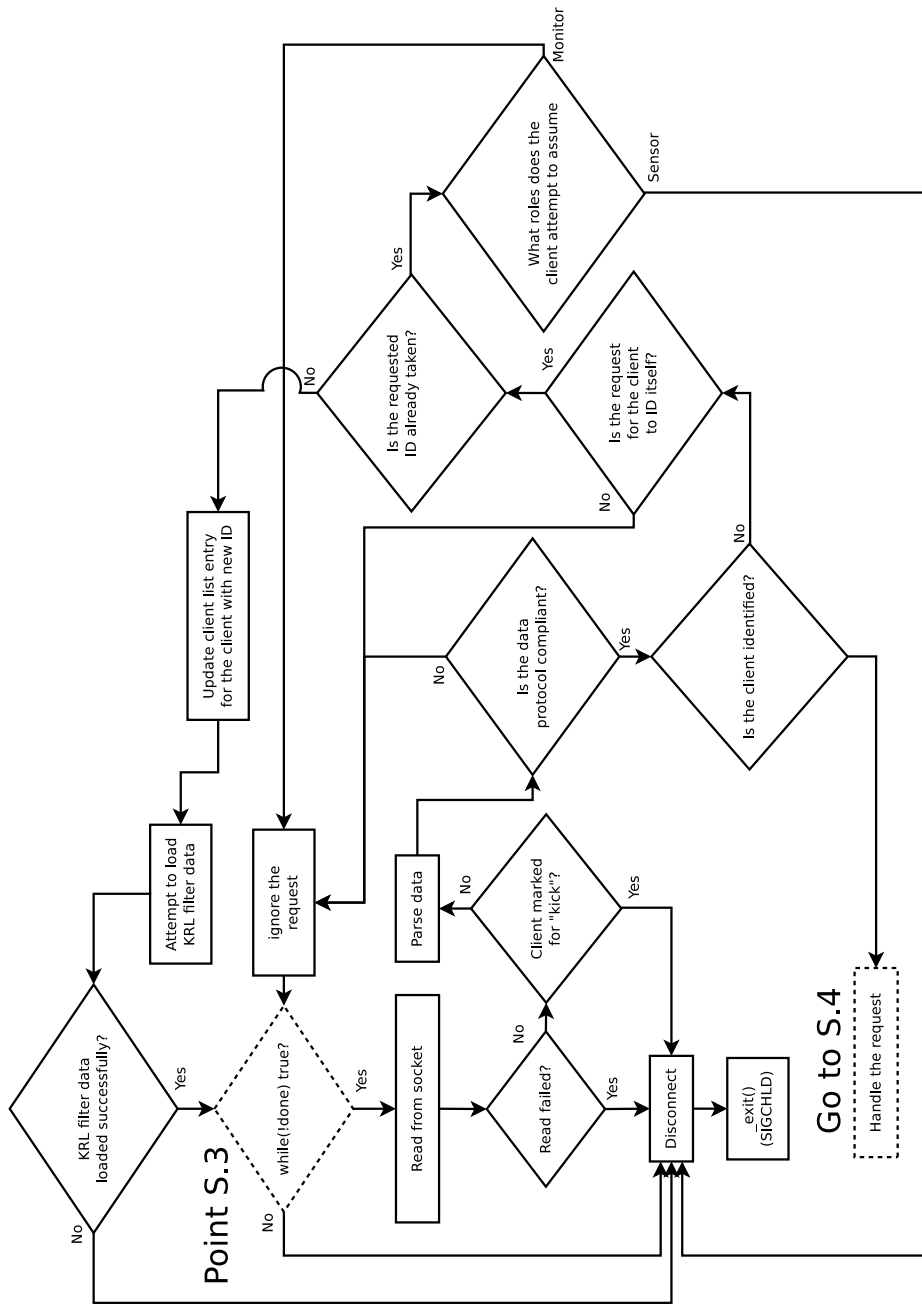


Figure 6.5: The block diagram shows an abstracted view of execution after a client has connected to the server and a `fork()` has been performed.

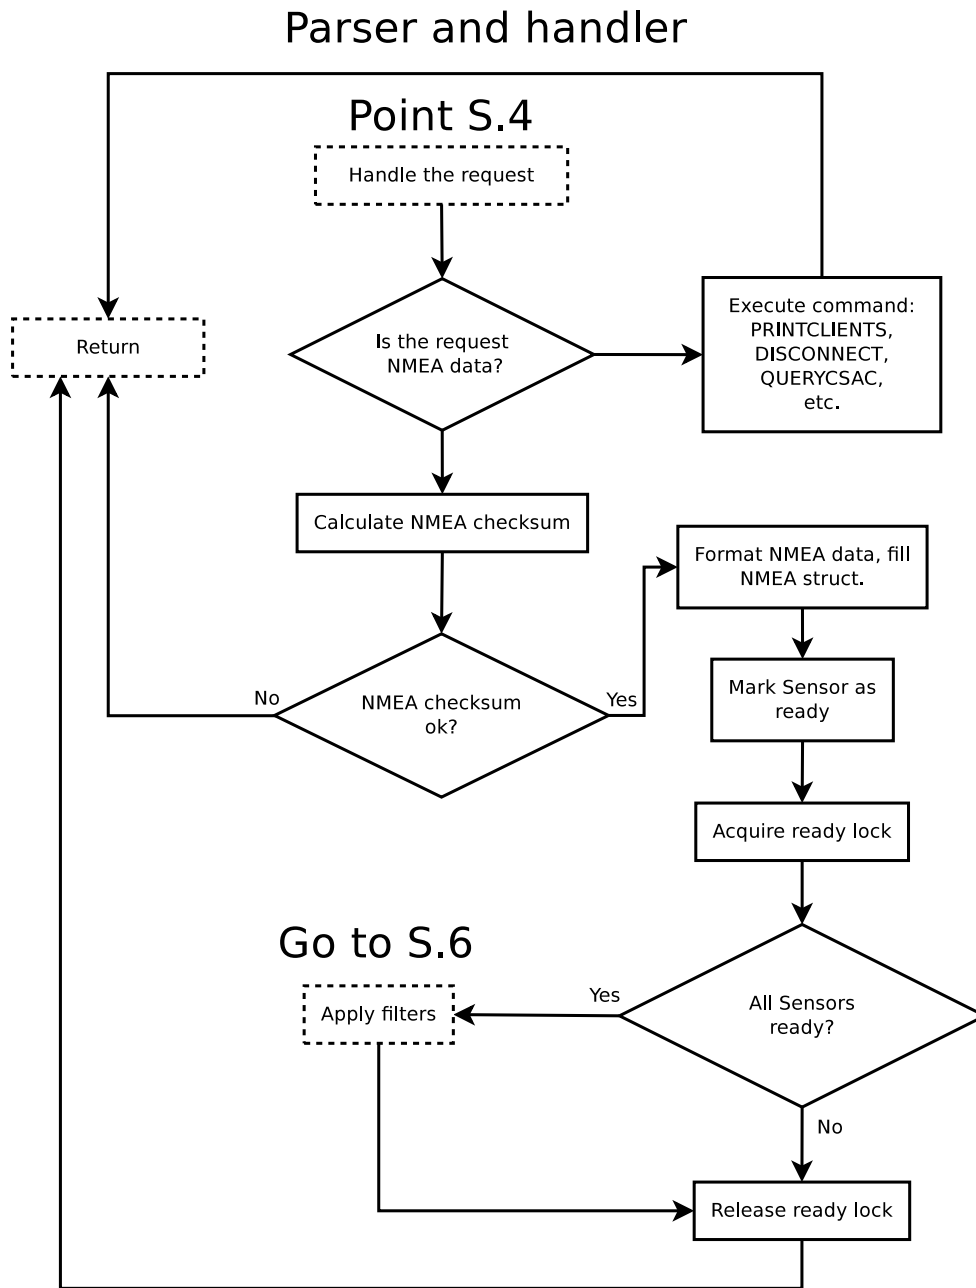


Figure 6.6: The block diagrams shows and abstracted view of the execution after data has been received from a client.

- If reading from the socket fails, the client gets disconnected. If it is successful, the client table entry for the connected client is checked to see if the client has been marked to be kicked. If so, the client is disconnected.
- The data received from the client is parsed and checked against the protocol. If the request received from the client does not match any of the supported requests or commands, the request is ignored and the loop is entered again. See section 6.3 for more information about the protocol and subsection 6.2.11 for more information about the parser.
- If the request received is valid, the clients ID is checked to determine whether or not it has identified itself. If the client has identified itself, the request gets carried out.
- If the request received is to handle NMEA data, the checksum for that NMEA data is calculated. If the checksum check fails, the data is discarded. If it succeeds, the data is copied into `nmea_container` struct (6.5.4) for easier handling.
- The client is marked as ready for processing and a ready check is commenced. The ready check is used to make sure that all the sensors have received NMEA data and are ready to have filters applied. Once the filters have been applied, the lock is released and normal execution resumes. The ready check is explained further under subsection 6.2.12.
- If the clients ID is 0, it is considered unidentified and any other requests but to identify itself, is ignored.
- If the client attempts to identify itself, the ID it attempts to use is checked against the rest of the connected clients. If the ID is already taken and the client attempts to assume the role of a Sensor, the client will get disconnected. If it attempts to assume the role of a Monitor, it will simply be notified that the chosen ID is taken and the request will be ignored. The reason why a Sensor would get disconnected and a monitor would not, is simple: When connecting the sensor to the server, it should not be any doubt whether or not it has been accepted. If something is wrong, it is better that a user of the system deals with the configuration at once rather than allowing the user to apply a faulty configuration that severely impacts the efficiency of the system. See section section 5.2 for more information about the Sensor and Monitor roles.

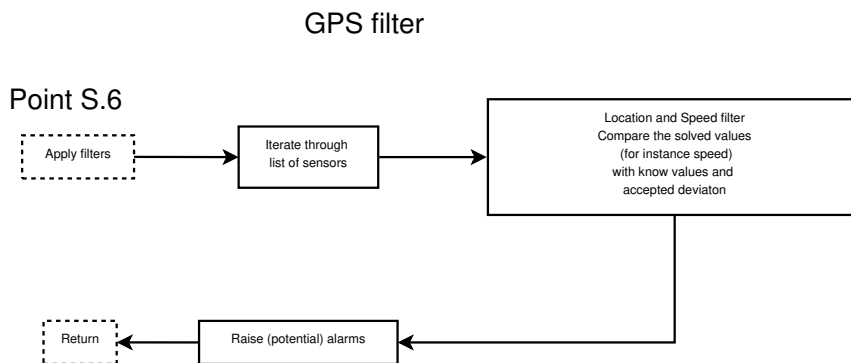


Figure 6.7: Block diagram show the order of execution in the GPS filter(s)

- If the attempted ID is available, the client table entry for the client is updated with the requested ID and the server will attempt to load data for the Location and speed filter data into the clients location and speed filter struct. The server will attempt to load this data from a file named `kr1_data_sensor<ID>`. If the load fails, the client will be disconnected. It is after all useless without its reference position data. For more information about location and speed filter, see section 6.2.4.1.

6.2.4 GPS filter

There is currently only one filter using GPS data in the Sensor Server implementation to this data. This is the *Location and speed filter*. Figure 6.2.4 is a block diagram showing the filter as a part of the Sensor Server.

6.2.4.1 Location and speed filter

By configuring the Sensor Server with the location of the different Sensors, the server may trigger an action if a Sensor reports an abnormal solved position. The filter is triggered when either latitude, longitude, altitude or speed is higher or lower than the reference value minus or plus a deviation. Listing (1) shows an edited sample of code taken from `filters.c`. The code in the sample is part of the algorithm used to check whether or not the latitude part of the GPS receivers solved position is within "safe" (not spoofed) range.

Listing 1: Listing shows sample of code taken from `filters.c` line 79. The code is part of the algorithm that compares the known position of a Sensor with the Sensor resolved position. The code has been edited for clarity.

```

1     if(latitude_current > latitude_reference + latitude_deviation) {
2         moved = 1;
3         lat_disturbed = HIGH;
4     } else if(latitude_current < latitude_ref - latitude_dev) {
5         moved = 1;
6         latitude_disturbed = LOW;
7     } else {
8         latitude_disturbed = SAFE;
9     }

```

When a Sensor connects to the Server and identifies itself, the Server will attempt to load a file containing the location of the Sensor as well as the accepted deviation values for the solved position. Listing 2 shows an example of such a file.

Listing 2: Location and speed filter configuration file

```

1     alt_ref: 123.8
2     lon_ref: 1102.1948
3     lat_ref: 5958.5448
4     speed_ref: 0
5     alt_dev: 10
6     lon_dev: 0.005
7     lat_dev: 0.005
8     speed_dev: 10

```

6.2.5 Clock model filters

Figure 6.4 shows both the atomic clock model and filter in one figure. As explained in section 6.2.2, the model of the atomic clock and the filters based upon, are closely related. See B for a description of the clock model.

- Once the model has been running for a configurable long time (48 hours was used during testing, see B for more about the model), it is ready to be used as a reference, filtering out abnormal data.
- The filters are checked. There are currently two filters implemented using the atomic clock model as reference:
 - *Phase jump filter*. This is a simple filter that compares the current phase reported by the atomic clock with a pre-configured limit. If the current phase is higher, the filter is triggered.
 - *Frequency correction filter*. This is a more sophisticated filter. It uses the atomic clock model to calculate a predicted steer value and compares it to the current steer value. If the current steer is too big, the filter is triggered.

- If any of the filters are triggered, the alarm is raised. The state of the alarm is expressed by an `int` that is either 0 or 1. If the alarm was not raised, the disciplining of the atomic clock is enabled if it was disabled, and the telemetry data is used to update the model.
- If the alarm is raised, the disciplining of the atomic clock is deactivated, thus protecting the atomic clock from being affected by the spoofing. The model is now used to calculate steer values that are applied to the atomic clock instead.

6.2.6 GPS data

The GPS receivers that we have chosen for our system and possibly all GPS receivers, follow the National Marine Electronics Association's 0183 standard [15]. NMEA data consists of sentences where the first word of the sentence, the type, defines how the sentence should be interpreted. The sentences used by the Sensor Server and Client, is the `GPRMC` for latitude, longitude and speed and `GNGGA` for altitude. See the u-blox manual [21] for more about the different sentences.

6.2.7 Client connection

Every time a client connects to the server, a process is forked for the new connection. Listing 3 is sample of code taken from the core of the Sensor Server, and can be explained like this:

- Line 4: The server waits for a connection. `accept()` is a blocking function. The code does not continue past this point before a client has connected.
- Line 12: A client has connected. The server creates a new process to handle the new connection. The new process is created using `fork()`.
- Line 13: Both the new process and its parent enters the "if statements" regarding its process identification (PID). The parent does not match the criteria of the "ifs" and returns to the top of the loop. The child on the other hand, matches the criteria for the "if sentence" at line 15.
- Line 16: The child process closes its parent's socket file descriptor and continues to setup the session at the next line.

Listing 3: Sample of code taken from `sensor_server.c` (D.2, line 494). The sample has been edited for clarity purposes.


```

1  listen(server_sockfd,SOMAXCONN);
2  int session_fd = 0;
3  while (!done) {
4      session_fd = accept(server_sockfd,0,0);
5      if (session_fd!=-1) {
6          if (errno==EINTR) continue;
7          t_print(ERROR_CONNECTION_ACCEPT,errno);
8      }
9      if(number_of_clients == max_clients) {
10         close(session_fd);
11     } else {
12         pid_t pid=fork();
13         if (pid!=-1) {
14             printf(ERROR_FAILED_FORK, errno);
15         } else if (pid==0) {
16             close(server_sockfd);
17             setup_session(session_fd, new_client);
18             close(session_fd);
19             _exit(0);
20         } else {
21             close(session_fd);
22         }
23     }
24 }

```

6.2.8 Shared memory & Semaphores

The architecture uses several shared memory segments. The pointers to the shared memory segments are declared as *extern* in `sensor_server.h`. The `extern` keyword means the the variable has an external linkage, making it visible from other files than the one in which it is defined. Listing 4 shows a code sample taken from `sensor_server.h` where the shared memory segments are declared.

Listing 4: Sample of code from `sensor_server.h`(D.2, line 44) where shared memory segments are declared.

```

1  extern struct client_table_entry *client_list;
2  extern struct server_data *s_data;
3  extern struct server_synchro *s_synch;
4  extern struct server_config *s_conf;
5  extern struct csac_filter_data *cfd;

```

These shared memory segments have different usage. The `client_list` points to a shared segment allocated for storing a list of all connected clients. `s_data` contains information about the server, `s_synch` contains semaphores

used to lock down shared resources, `s_conf` is the servers configuration and `cfid` is the data and state for the filters based on the clock model. Every process that forks out from the server is given access to these memory segment. One might make the point that this voids the idea of processes, and one might be correct (see 8). The shared memory is created using the GNU library's Memory Mapped I/O (MMAP). Although typically used to map files to a region of memory, MMAP can also be used to create an anonymous map which is not connected to file but rather for sharing data between tasks without using files. Listing 5 shows an example of MMAP being used to map an anonymous, shared map for the client list.

Listing 5: Listing shows the use of MMAP to create an anonymous map of memory to be used as a shared memory segment. Sample of code taken from `sensor_server.c`(appendix D.2, line 360)

```

1  client_list = mmap(NULL,
2                      (s_conf->max_clients * sizeof(struct client_table_entry)),
3                      PROT_READ | PROT_WRITE,
4                      MAP_SHARED | MAP_ANONYMOUS | MAP_NORESERVE,
5                      -1, 0);

```

6.2.9 Client list memory segment

Every time a client connects, it is given a piece of shared memory where it can store its data. This piece of shared memory is used as a node in a linked list. The linked list structure, stored in the shared memory segment is available to all the processes spawned by the server. This segment is static in size and is allocated once and is never changed during the whole life of the Server. The size of the segment is determined by the maximum number of allowed clients, a configurable value read from the configuration file every time the server is started. Read more about why the client list shared segment is static in size, under subsection 8.2.1.

6.2.9.1 Using the client list memory segment

As previously explained, the list of clients are stored in a shared memory segment using a linked list structure. In order to keep tabs on what part of the shared memory segment is in use, a simple map of the memory is used (6.8):

The map is an array containing pointers to the client list segment. These pointers are offset by the size of the `client_table_struct` and maps 1:1 to what can be thought of as individual *pieces* of memory in the client list




Figure 6.8: Figure is showing the relation between the shared memory segment containing the client list and the memory map




Figure 6.9: Figure is showing the relation between the shared memory segment containing the client list and the memory map after some of the pieces of shared memory has been given out.

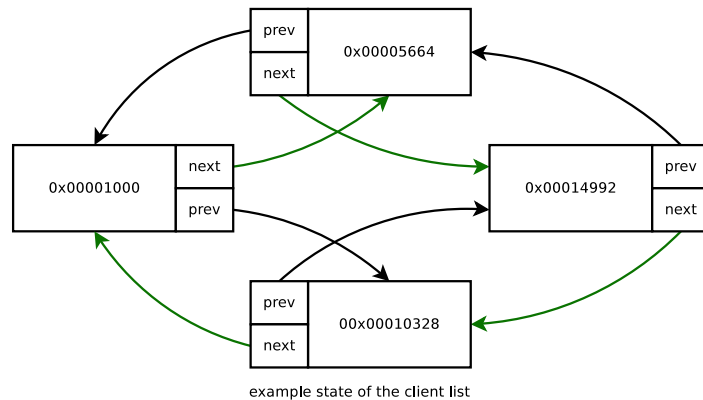
shared memory segment. A free piece of memory is in the map represented by a pointer that is not NULL. Figure 6.9 is an illustration showing the state of the map when four of the first pieces have been given out. Listing 6 shows the function that iterates through the map to find a free piece.

Listing 6: Sample of code showing the `get_mem_piece()` function. This function is used to find a free piece of memory in the client list shared memory segment. The code is taken from line 200 in `sensor_server.c` (appendix D.2)

```

1     static struct client_table_entry* get_mem_piece()
2     {
3         int i;
4         for(i = 1; i < s_conf->max_clients; i++){
5             if(client_list_map[i] != NULL){
6                 struct client_table_entry *tmp = client_list_map[i];
7                 client_list_map[i] = NULL;
8                 return tmp;
9             }
10            i++;
11        }
12        return NULL;
13    }

```



*all memory addresses are used as examples

Figure 6.10: Block diagram show an example of linked list state

6.2.10 Client linked list

Since the C standard does not provide data structures like linked lists, I had to choose between reinventing the wheel or finding some implementation to drop into the project. While studying another subject, I found a guide on how to use the linked list implementation from the linux kernel source code ([25]) in a user space program. Since the implementation was solid, well tested and had many useful functions, I decided to use it. The modified header file containing all the code, is GPL licensed. Figure (6.10) shows pieces of the shared memory segment linked together in a linked list structure.

6.2.10.1 Alternative to linked list

The initial reason for using a linked list to organize the connected clients, was to reduce the time spent iterating over a space of memory. With a linked list it would not matter where the different pieces of memory were located, a pointer to next piece would always be readily available anyway.

6.2.10.2 Semaphores

Having shared memory segments comes with a price. Whenever two or more processes are working on the same data set, they are prone to create race conditions, deadlocks and data corruption. Therefore, semaphores are used to lock the segments during read and write operations.

Listing 7: Function for removing disconnected clients from list of clients

```

1   void remove_client_by_id(int id)
2   {
3       struct client_table_entry* cli;
4       struct client_table_entry* temp_remove;
5
6       sem_wait(&(s_synch->client_list_sem));
7       list_for_each_entry_safe(cli, temp_remove, &client_list->list,
8                               list) {
9           if(cli->client_id == id) {
10              list_del(&cli->list);
11              s_data->number_of_clients--;
12          }
13      }
14      sem_post(&(s_synch->client_list_sem));
15  }

```

Listing 7 shows a typical example of a function locking down access to the shared memory segment containing the list of connected clients by using a semaphore. In the example (7) a client has been disconnected from the server and the list of connected clients is updated. The semaphore is necessary to make sure that another process is not attempting to read or write to the segment while the data is deleted. If another process had attempted to execute the `sem_wait()` on the semaphore, it would have been put in a queue. Depending on the operating system, it would most likely signal the scheduler to do a context switch since the resource was busy anyway and it therefor should relinquish control of the CPU. Once the semaphore is raised, it can be lowered again by another process. It is important to note that the semaphores are not a function of, or related to the memory segments by anything other than the name. The semaphores are just "flags" used to control access to a resource. There is no automatic raising or lowering of the associated semaphores by reading or writing to a specific shared memory segment. All functions in the sensor server use semaphores when dealing with shared memory segments in order to avoid deadlock and race conditions.

6.2.11 Client input parser

The parser used by the Sensor Server is a simple function. As explained in section 6.2.3, every time a client sends data to the Server, the data is sent through the parser in order to detect the purpose of the request. The parser uses the protocol and compares it with the input and looks for a match. Listing 8 shows a sample of code taken from the parser.

Listing 8: Part of the parser comparing the input data to the IDENTIFY command specified by the protocol. Sample code is taken from *session.c* (appendix D.2 line 164)

```
1 static int parse_input(struct client_table_entry *cte){
2     char *incoming = cte->transmission.iobuffer;
3     ...
4     else if(strstr((char*)incoming, PROTOCOL_IDENTIFY ) == (incoming)) {
5         int length = (strlen(incoming) - strlen(PROTOCOL_IDENTIFY) );
6         memcpy(cte->cm.parameter,
7             (incoming)+(strlen(PROTOCOL_IDENTIFY)*(sizeof(char))),length);
8         cte->cm.code = CODE_IDENTIFY;
9     }
```

The listing (8) shows how the parser attempts to find a match between the input buffer and the protocol defined IDENTIFY command. On line 6 it attempts to copy any potential parameter. The parser does not care if the parameter is missing, but it will attempt to extract it from the input buffer. If no matches are found, the function returns 0, letting the calling function (`respond()`) know that the input from the client was invalid or illegal. Comparing strings is hard work and a relatively CPU intensive task. This is why the parser sets a *command code* (6.5.3) as seen on the last line. The command code is an integer defined in the protocol and there is a command code associated with every command. The code is used in the listening loop by `respond` to determine what action to perform in case a valid request. Comparing integers are "cheaper" than comparing strings, and since the string comparison job is already done once, there is no need to do it again.

6.2.12 Ready check algorithm

Every time NMEA data is received and validated, the process that received the data will initiate a ready check. This is done by setting a flag in the client list entry called `ready` to 1, and calling the `nmea_ready()` function. The `nmea_ready()` function locks down the client linked list and iterates through it checking if the other processes are ready too. If all the other processes have received valid NMEA data from their respective client, the GPS filter is applied. If they are not ready, the lock is released and the process carries on. This means that the last process to receive NMEA data gets the job of applying the filters ³.

³Since NMEA data is generated every second by the GPS receivers, it might be more appropriate to say that the scheduler decides which process executes last. This of course also depends on the distance between the Sensor and the Client

6.3 Sensor Server Protocol

The Sensor Server protocol is quite simply a header file containing every keyword or command that the Sensor Server recognizes. The thought behind it was that a client application (or script) could import the protocol header file and easily use and recognize commands used between the Server and a Client. Some constraints are also defined in the protocol header file, like maximum parameter size, maximum and minimum command size. The *command codes* used by the parser are also defined in the protocol header file. See subsection 6.2.11 for more about the use of protocol.

6.4 Atomic clock Communication

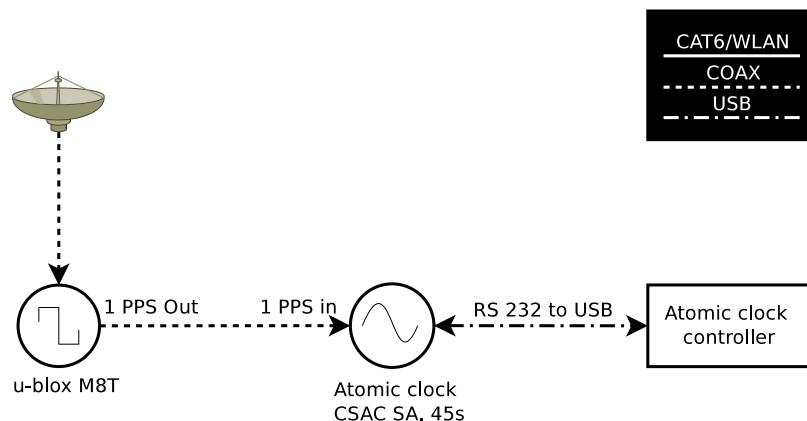


Figure 6.11: Block diagram showing the atomic clock connected to a PC

The SA.45 atomic clock includes a serial interface that enables communication with a PC by using a COM port. As mentioned earlier, our approach relies heavily on the ability to communicate with the atomic clock. Information may be queried by sending commands to the atomic clock. These commands are explained in table 6.1.

The Sensor Server communicates with the atomic clock by invoking a script called `query_csac.py`. This script can be found in the appendix (E). The script takes an argument which it sends to the atomic clock over serial, and prints the respond to the shell which the Sensor Server grabs. In this document the word *telemetry* gets mentioned quite often in context with the atomic clock. Telemetry can be obtained by querying the atomic clock. The telemetry is a string containing a plethora of information, but we are mainly interested in the following values:

Table 6.1: Commands for the SA.45 atomic clock

Shortcut	Description	Command
6	Return telemetry headers as comma-delimited string	!6[CRLF]
^	Return telemetry as comma-delimited string	!^[CRLF]
F	Adjust frequency	!F?[CRLF]
M	Set operating mode register bits	!M?[CRLF]
S	Sync atomic clock 1 PPS to external 1 PPS	!S[CRLF]
D	Set 1 PPS disciplining time constant	!D?[CRLF]
U	Set ultra-low power mode parameters	!U?[CRLF]
T	Set/report time-of-day	!T?[CRLF]

Source: [16]

- Phase, the difference between the atomic clock and the external signal signal at 1PPS in.
- DiscOK, the discipline status.
- Steer, the frequency adjustment.

Listing 6.1 show an example of telemetry received from the atomic clock. The fourth value from the right, is the `DiscOK`, the fifth is the `phase` and the seventh is the `steer`.

Listing 6.1: Example of a telemetry string received from the atomic clock

```

1 0,0x0000,1209CS00909,0x0010
   ,4381,0.86,1.573,17.62,0.996,28.26,-24,---,-1,1,1268126502,586969,1.0

```

The SA.45 uses a high-resolution phase meter to improve synchronization and to calibrate the frequency of the atomic clock. The phase meter measures the difference in time between the internal 1 PPS signal and the external reference, in our case a GPS receiver. This difference is the *phase* value. The atomic clock uses the phase value and steering algorithms to adjust the frequency of atomic clock's physics package thus simultaneously steering both the phase and frequency towards the external reference. This is called *disciplining* and is how the *steer* value is computed. The last value, *DiscOK* is simply the status of the 1 PPS disciplining routine [16].

6.5 Data structures

In the C programming language a "struct" is a complex data type that defines a list of variables to be placed under the structs given name in a block of

memory. This makes it possible for multiple variables to be accessed via a single pointer. Some crucial and often used structs is explained here.

6.5.1 server_data

Listing 9: Sample of code taken from `sensor_server_common.h` (appendix D.2, line 107) showing the `server_data` struct.

```
1     struct server_data {
2         int number_of_clients;
3         int number_of_sensors;
4         time_t started;
5         pid_t pid;
6         char version[4];
```

The `server_data` struct as shown in listing 9, contains information about the server. Some of the information like the PID, version and when the server was started, is just useful information about the server itself. The number of clients and sensors on the other hand are used to make sure that the server does not allow more connections than it can handle.

6.5.2 server_synchro

Listing 10: Sample of code taken from `sensor_server_common.h` (appendix D.2, line 116) showing the `server_synchro` struct

```
1     struct server_synchro {
2         sem_t ready_sem;
3         sem_t csac_sem;
4         sem_t client_list_sem;
5         volatile sig_atomic_t done;
6     };
```

The `server_synchro` as shown in listing 10 contains "flags" used by synchronization mechanisms (6.2.8) that the server and its children use to protect access to shared resources. The `csac_sem` is used to control serial access to the atomic clock, making sure that only one request is sent to atomic clock at a time. The `client_list_sem` is used by functions manipulating and reading the client list structure, and the `ready_sem` is used by ready check algorithm. See subsection 6.2.12 for more about this.

6.5.3 command_code

Listing 11: Sample of code taken from `sensor_server_common.h` (appendix D.2, line 34) showing the `command_code` struct

```

1     struct command_code {
2         int code;
3         char parameter[MAX_PARAMETER_SIZE];
4         int id_parameter;
5     };

```

The `command_code` struct as shown in listing 11, is used by the parser to more efficiently convey protocol compliant requests and related parameters. See section 6.2.11 for more about the parser.

6.5.4 nmea_container

Listing 12: Sample of code taken from `nmea.h` (appendix D.2, line 20) showing the `nmea_container` struct

```

1     struct nmea_container {
2         /* Raw data */
3         char raw_gga[SENTENCE_LENGTH];
4         char raw_rmc[SENTENCE_LENGTH];
5
6         /* Latitude */
7         double lat_current;
8         double lat_average;
9         double lat_avg_diff;
10        double lat_total;
11        int lat_disturbed;
12
13        /* Longitude */
14        double lon_current;
15        double lon_average;
16        double lon_avg_diff;
17        double lon_total;
18        int lon_disturbed;
19
20        /* Altitude */
21        double alt_current;
22        double alt_average;
23        double alt_avg_diff;
24        double alt_total;
25        int alt_disturbed;
26
27        /* Speed */
28        double speed_current;
29        double speed_average;
30        double speed_avg_diff;
31        double speed_total;
32        int speed_disturbed;
33

```

```

34     /* CHECKSUM */
35     int checksum_passed;
36
37     /* COUNTER FOR AVERAGE */
38     int n_samples;
39 };

```

The `nmea_container` struct as shown in listing 12 is used by the handler (see section 6.2.3) which dissects the NMEA strings and validates them before it fills the respective members of the structs. The purpose of it is just to make it easier for the other parts of the Server to use the NMEA data.

6.5.5 list_head

Listing 13: Sample of code taken from `list.h` (appendix D.2, line 70) showing the `list_head` struct

```

1     struct list_head {
2         struct list_head *next, *prev;
3     };

```

The `list_head` struct is shown in listing 13. The fields of the struct are pretty self explanatory. There is a pointer to the previous node and one to the next. One of the members of the `client_table_list` is a struct of type `list_head`. This is what makes it possible to traverse the list.

6.5.6 transmission_s

The `transmission_s` struct as show in listing 14 is used by the child process forked out by the Server in the server core to communicate with the client. The struct has two members, a file descriptor for the socket and a buffer to store incoming data.

Listing 14: Sample of code taken from `net.h`(appendix D.2) line 31

```

1     struct transmission_s {
2         int session_fd;
3         char iobuffer[IO_BUFFER_SIZE];
4     };

```

6.5.7 client_table_entry

The `client_table_entry` struct is what the name suggests, it is an entry in a list of clients. Every client connected to the server, no matter the purpose, has an entry in the client list. Listing 15 shows the complete struct.

Listing 15: Sample of code taken from `sensor_server_common.h`(appendix D.2, line 90) shows the `client_table_entry` struct

```
1     struct client_table_entry {
2         struct list_head list;
3         struct transmission_s transmission;
4         struct timeval timeout;
5         struct command_code cm;
6         struct nmea_container nmea;
7         struct filters fs;
8         pid_t pid;
9         time_t timestamp;
10        int client_id;
11        int client_type;
12        int ready;
13        int marked_for_kick;
14        char ip[INET_ADDRSTRLEN];
15    };
```

Beginning from the top:

- `list_head list` is used to traverse the linked client list. See 6.2.10 for more about the linked list implementation and 6.5.5 for more about the struct.
- `transmission_s transmission` is used by the process to network I/O with the connected client.
- `timeval timeout` is defined by the `sys/time.h` and is used in the Sensor Server implementation to hold a value defining how long a connected client can stay connected without sending any commands to the Server. When the timeout value is exceeded, the client is disconnected. The values are:
 - 5 seconds for a Sensor.
 - 1000 seconds for a Monitor.
 - 10 seconds for an unidentified client.

These values are defined in the `sensor_server.h` file.

- `command_code cm` is used by the parser to convey protocol compliant requests. See 6.5.3 for more.
- `nmea_container` is used to contain NMEA data in a more convenient way for the filters. See 6.5.4 for more about the struct, and 6.2.6 for more about NMEA.

- `filters fs` is a struct containing data for the GPS based filters. There is currently only one GPS based filter, the Location and speed filter. It is explained in subsection 6.2.4.
- `pid_t pid` is the process ID for the process forked out for the connected client.
- `time_t timestamp` is time stamp that is stamped every time the Sensor's filter data is processed.
- `client_id` and `client_type` is the client's ID and the client's type, either "SENSOR" or "MONITOR".
- `ready` is used to indicate that a Sensor has received NMEA data and is ready for filter processing.
- `marked_for_kick` is a flag. If it is 1, the client has been marked by a monitor indicating that it should be kicked. The next time the client sends data to the Server, the client will be disconnected.
- `char ip` is the connected client's IP address.

6.6 The Sensor Client

The sensor client software is a simple program written in C99 whose only task is to relay information read from the GPS receivers. Summed up shortly:

- The client software takes two parameters to start, the servers IP and port. If parameters are missing, the program exits.
 - Example: `./sensor_client -p 10000 -i 192.168.1.5`
- Initializes and loads configuration from configuration file. The configuration file includes path to the GPS receiver, the sensors ID number and a binary value for whether or not logging of NMEA should be done as well as path to the log file. If the loading of the configuration file fails, default values are used instead:
 - The ID number is chosen randomly but within legal limits.
 - Logging is disabled.
 - Maximum of server connection attempts are set to 10.

Client simplified call graph

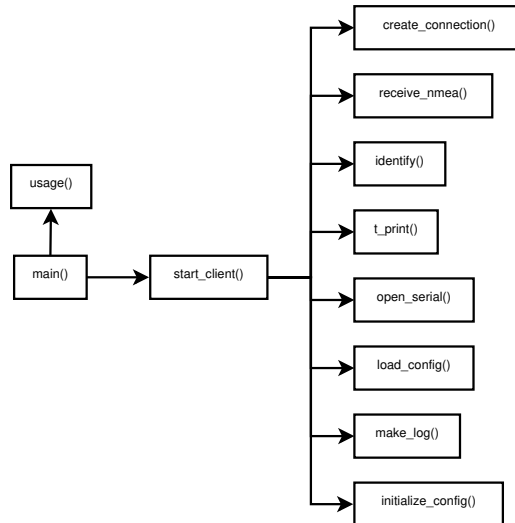


Figure 6.12: Simplified call graph for the Sensor Client.

- Path to GPS receiver is set to `/dev/ttyACM0`. This should be the path to the receiver unless another similar device is connected to the computer and given it is a Raspberry Pi running Raspbian.
- Establishes communication with GPS receiver, exits if it fails.
- Attempts to establish communication with the server, retries for a configurable amount of times at 1 second intervals.
- Identifies the client for the server according to protocol.
- Reads from the GPS receiver, scans for lines starting with either `$GNRMC` or `$GNGGA`. When both lines are found, the data is stored in a buffer.
- Sends the GPS data to the server according to protocol.
- Repeats.

Figure 6.12 shows a simplified call graph for the Client.

6.7 Interfacing

As mentioned earlier under section (5.2), a client can assume the role of both a monitor and a sensor. The sensor server does not differentiate between the

two roles other than when it routinely checks the status of its filters. This means that one could connect to the server using the sensor client software as a monitor by configuring the sensor client software to use a negative integer as ID. At this point, this kind of functionality is not very useful as there is currently no way to change the ID of a client unless the client explicitly issues the command to do so, but it opens for the possibility to interface with the server. One way to interface with the server is to connect to it using Telnet:

```
1 user@machine:/$ telnet 10.1.0.46 10001
2 Trying 10.1.0.46...
3 Connected to 10.1.0.46.
4 Escape character is '^]'.
5 ID -3
6 OK!
7
8 >
```

It is also possible to assume the role of a sensor by connecting to the server via telnet. This can be used to debug or troubleshoot the sensor server by manually feeding it NMEA data. The only requirement is that the communication is sensor server protocol compliant:

```
1 user@machine:/$ telnet 10.1.0.46 10001
2 Trying 10.1.0.46...
3 Connected to 10.1.0.46.
4 Escape character is '^]'.
5 ID 2
6 OK!
7
8 NMEA<GNRMC part of NMEA><GNGGA part of NMEA>
9 OK!
```

Another possibility is of course to write scripts that communicates with the server. An example script can be found in the appendix (E).

Table 6.2: sensor server available commands

Command	Short	Parameter	Description
HELP	?	NONE	Prints this table
IDENTIFY	ID	ID	Clients ID is set to PARAM
DISCONNECT	EXIT	NONE	Disconnect from the server
PRINTCLIENTS	PC	NONE	Prints an overview of connected clients
PRINTSERVER	PS	NONE	Prints server state and config
PRINTTIME		ID	Prints time solved from GNSS data received from sensor <ID>
PRINTAVGDIFF	PAD	NONE	Prints the difference between current solved position and the average reported for all sensors
PRINTLOC	PL	ID	Print solved position for sensor <ID>
LISTDATA	LSD	NONE	List all dump files stored by the server
DUMPDATA	DD	ID & FILE	Dumps state of sensor <ID>into a file named <FILE>
LOADDATA	LD	ID & FILE	Load state stored in file called <FILE>into sensor <ID>
QUERYCSAC	QC	COMMAND	Queries the CSAC with COMMAND.
LOADLSFDATA	LLSFD	ID	Load reference location data into Sensor <ID>
PRINTCFD	PFD	NONE	Prints CSAC filter data

Chapter 7

GPS manipulation tests

In this chapter we present the tests that were conducted. Ideally these tests would have been performed by using a GPS spoofer like the "Civil GPS spoofer" that was introduced in subsection 2.4. Unfortunately, or perhaps not, this kind of hardware is hard to come by and spoofing GPS signals is not legal. What we did instead was manipulating the time solution of the GPS receivers by moving the antennas, thus naively simulating a spoofing attack. By performing these tests we were able to not only test the sensor server architecture part of the atomic clock controller, but also the filters and the clock model. We also present data gathered during an unplanned GPS disturbance that occurred in between our planned tests.

7.1 Logging filter and model data

In order to create an accurate clock model of the atomic clock, it was necessary to log data from the atomic clock while it was running in a disciplined mode. In the disciplined mode the atomic clock corrects its frequency based on either a one pulse per second (PPS) signal or a 10 MHz signal as discussed earlier under section (6.4). A similar approach was used in order to collect GPS data. Data from the two u-blox NEO-M8T receivers was gathered over the same time period as the data gathered from the atomic clock. By gathering the data over the same period it was possible to detect any correlation between the time solved by the GPS receivers and any frequency adjustments done by the atomic clock. It also provided valuable data that could be used to tune the detection filters. The data gathering was done by simple Python scripts (appendix C.1 and C.2) running on a computer connected to the GPS receivers and the atomic clock. Figure (E.1) shows a block diagram of the setup.

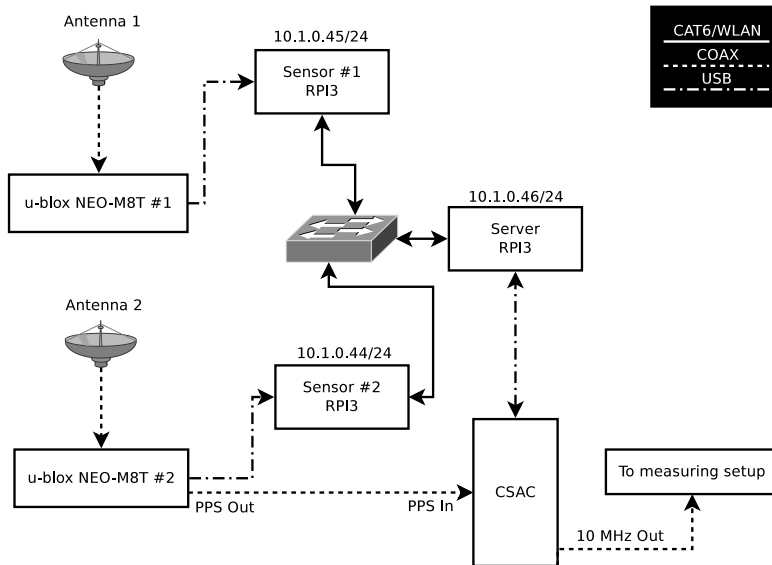


Figure 7.1: A block diagram showing the tested implementation.

7.2 Setup

Figure 7.1 shows how the server, sensors and atomic clock were physically set up. In order to assure good GPS satellite geometry, the antennas were placed at the roof of Justervesenet’s office at Kjeller. Antenna one was placed at a railing about one meter above the ground, antenna two was placed at ground level. Figure 7.2 shows the placement of the antennas. Antenna one was connected to GPS receiver one, which in turn was connected to sensor one. It was the same deal with antenna two, which was connected to GPS receiver two which in turn was connected to sensor two. The distance between the two antennas was about 35 meters. The antenna connected to sensor one was not placed as far away as the cable would have allowed. This is because it proved challenging to find a suitable place to securely place the antenna and at the same time use the full length of the cable. The sensors and the server were connected to a LAN through a Gigabit Ethernet switch. The server was configured to log telemetry data received from the atomic clock and the clients were configured to log all NMEA data received from their respective GPS receiver. GPS receiver two supplied the atomic clock with a one PPS signal. Because the atomic clock model needs live data over time to mature, the system was started the 7 October 2016 and the test was performed 10 October 2016. Figure 7.3 is a block diagram showing how the measuring setup was physically configured. The measuring was done using a Spectracom CNT-91 frequency counter using a 10 MHz reference

from Justervesenet’s atomic clocks. The CNT-91 was configured to measure a continuous gap-less frequency at one second time-intervals. Figure G.1 is a photograph of the measuring setup. Both sensors were configured to use Justervesenet’s internal NTP server and to use UTC instead of GMT+1.

7.3 Test one filter limits

Table 7.1 shows the filter thresholds used during test one.

Table 7.1: Filter thresholds used during test one.

	Sensor one	Sensor two
Altitude reference	123.8°	122.427°
Longitude reference	1102.1948°	1102.1934°
Latitude reference	5958.5448°	5958.5231°
Speed reference	0 knot	0 knot
Altitude deviation	10°	10°
Longitude deviation	0.005°	0.005°
Latitude deviation	0.005°	0.005°
Speed deviation	10 knot	10 knot




Figure 7.2: Block diagram showing the position of the antennas as reported by the receivers. Map courtesy of Kartverket [26]

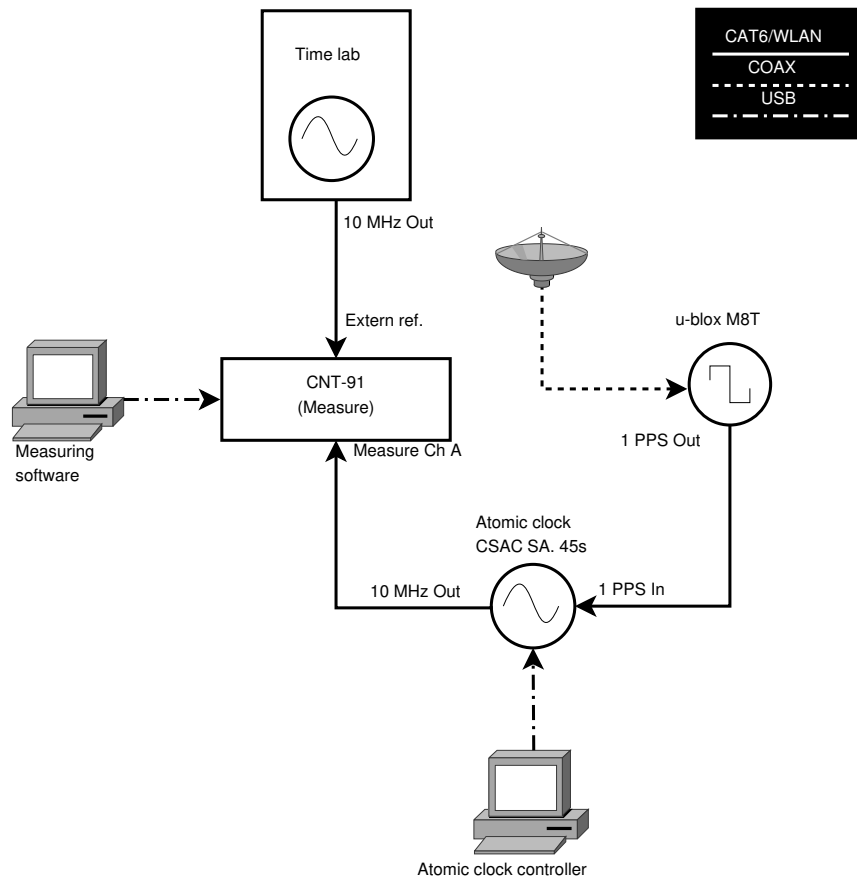


Figure 7.3: Block diagram showing the setup of the measurement equipment

7.4 Test one

7.4.1 Goal of test one

The goal with test one was to use the atomic clock controller to detect a simulated spoofing attack. By moving the antennas, the location and speed filter should be triggered as the solved longitude, altitude, latitude and speed change. See section 6.2.4.1 for more about the location and speed filter. The result is observable by analyzing the log files produced by the server and clients and the data collected by the measuring setup.

7.4.2 Description

The following is a step by step description of how the test was conducted. The time in the brackets was obtained from a wristwatch. The time is written

in the normal format (hours and minutes) as well as number of seconds after 10:48. This is because the data used to draw the graphs started at 10:48 but uses seconds on the x-axis, to ease comparison. It is important to note that neither the resolution nor the accuracy was of any notable concern when the time was noted down. The time was mainly noted to make it easier to find any correlation between the steps taken and patterns found in the log files.

- 10:58 - 600: Moved antenna 1 approximately 15 meters to the south.
- 11:03 - 900: Moved antenna 1 back to original location.
- 11:07 - 1140: Moved antenna 2 approximately 15 meters to the north.
- 11:12 - 1440: Moved antenna 2 back to original location.
- 11:14 - 1560: Waved antenna 1 around horizontally in a half circle motion at an increasing tempo.
- 11:18 - 1800: Waved antenna 2 around horizontally in a half circle motion at an increasing tempo.
- 11:20 - 1920: Covered antenna 1 with aluminium foil.
- 11:25 - 2220: Covered antenna 2 with aluminium foil.
- 11:28 - 2400: Removed foil from antenna 1.
- 11:33 - 2700: Removed foil from antenna 2.

Step one and two were designed to trigger the location and speed filter, especially the check of solved latitude, longitude and altitude but also the speed against known values. Step five and six were also designed to trigger the location and speed filter but more specifically the checking of solved speed. By waving the antenna around while standing still, the solved location should not exceed the configured limits except for speed. Step seven and eight were designed to reveal what would happen during a jamming attack as it was believed that covering the antennas with aluminium foil would block all signals out. A photography showing how the antenna was covered with aluminium foil can be seen in section G.2 in the appendix.

7.4.3 Observations

7.4.3.1 Sensor Server logs

By reviewing the log produced by the server, the following was observed:

- No false positives were reported. The filters were not triggered before the test started.
- The location and speed filter was triggered by sensor one at 10:59:19 and cleared at 11:04:35.

```
1 [10/10/16 - 10:59:17] [ ALARM ] Sensor 1 triggered LS filter!  
2 ...  
3 [10/10/16 - 11:04:35] [ ALARM ] Sensor 1 cleared LS filter!
```

- The location and speed filter was triggered again at 11:08:27, but this time by sensor one. The alarm was cleared at 11:13:43.

```
1 [10/10/16 - 11:08:27] [ ALARM ] Sensor 2 triggered LS filter!  
2 ...  
3 [10/10/16 - 11:13:43] [ ALARM ] Sensor 2 cleared LS filter!
```

- Once again, 11:22:03 the location and speed filter was triggered by sensor one and was not cleared until 11:29:21

```
1 [10/10/16 - 11:22:03] [ ALARM ] Sensor 1 triggered LS filter!  
2 ...  
3 [10/10/16 - 11:29:21] [ ALARM ] Sensor 1 cleared LS filter!
```

- Sensor two triggered the location and speed filter 11:27:31 and cleared at 11:34:16.

```
1 [10/10/16 - 11:27:31] [ ALARM ] Sensor 2 triggered LS filter!  
2 [10/10/16 - 11:34:16] [ ALARM ] Sensor 2 cleared LS filter!
```

- The last three seconds, sensor two triggered and cleared the location and speed filter.

```
1 [10/10/16 - 11:34:17] [ ALARM ] Sensor 2 triggered LS filter!  
2 [10/10/16 - 11:34:20] [ ALARM ] Sensor 2 cleared LS filter!
```

7.4.3.2 GPS data

The figures in this section are plots of the GPS data that was logged by the sensors during test one. Figure 7.4 shows the altitude in meters reported by sensor one. The only thing that stands out in the plot is when the antenna is covered in aluminium foil as shown at around the 2000 second mark. At

this point the receiver is no longer able to solve for a position thus generating valid but empty fields in the NMEA data. This creates a deep void in the plot. Because both receivers' antennas were covered in aluminium foil, the plots all share this trait. Figure 7.5 shows the altitude in meters as reported by sensor two. When compared with the altitude plot for sensor one 7.4, sensor one seems to correlate better. It remains a puzzle why sensor two reported to have been moved seven meters down which it clearly was not. When examining figures 7.6 and 7.7 showing the latitude for sensor one and two (in that order), one can clearly see that sensor one was moved to the south and sensor two was moved to the north. One might also notice that sensor two traveled further. The difference in travel is because sensor one's cable is shorter. This is explained in the description of test one. Figures 7.8 and 7.9 shows a plot for the solved longitude for sensor one and two. When examining figure 7.8, it seems as if the antenna were never moved. This is because the building the antenna was placed on is position in a north to south line and the receiver was moved along with the roof. The antenna connected to sensor two was not moved in straight line because the cable got caught in vegetation. The last two figures, figure 7.10 and figure 7.11 show the speed solved by sensor one and two. They both correlate well with when the antennas were moved. However, the antenna connected to sensor two seems to be slightly more sensitive than the antenna connected to sensor one.




Figure 7.4: The figure shows the solved altitude of the GPS receiver connected to sensor one, during test one




Figure 7.5: The figure shows the solved altitude of the GPS receiver connected to sensor two, during test one




Figure 7.6: The figure shows the solved latitude of the GPS receiver connected to sensor one, during test one




Figure 7.7: The figure shows the solved latitude of the GPS receiver connected to sensor two, during test one.




Figure 7.8: The figure shows the solved longitude of the GPS receiver connected to sensor one, during test one




Figure 7.9: The figure shows the solved longitude of the GPS receiver connected to sensor two, during test one




Figure 7.10: The figure shows the solved speed of the GPS receiver connected to sensor one, during test one




Figure 7.11: The figure shows the solved speed of the GPS receiver connected to sensor two, during test one

7.4.4 Timing measurements

Figure 7.12 shows the relative frequency offset (10^{12}) for the atomic clock. The thick dark plot in the middle is from telemetry data gathered from the atomic clock. Figure 7.13 shows the phase offset in nanoseconds for the atomic clock. The interesting thing to observe in both of these figures, is the jump in frequency and phase offset once the antenna was covered in aluminium foil. There is also a clear correlation between movement of the antenna and the relative frequency offset as seen in figure 7.12.

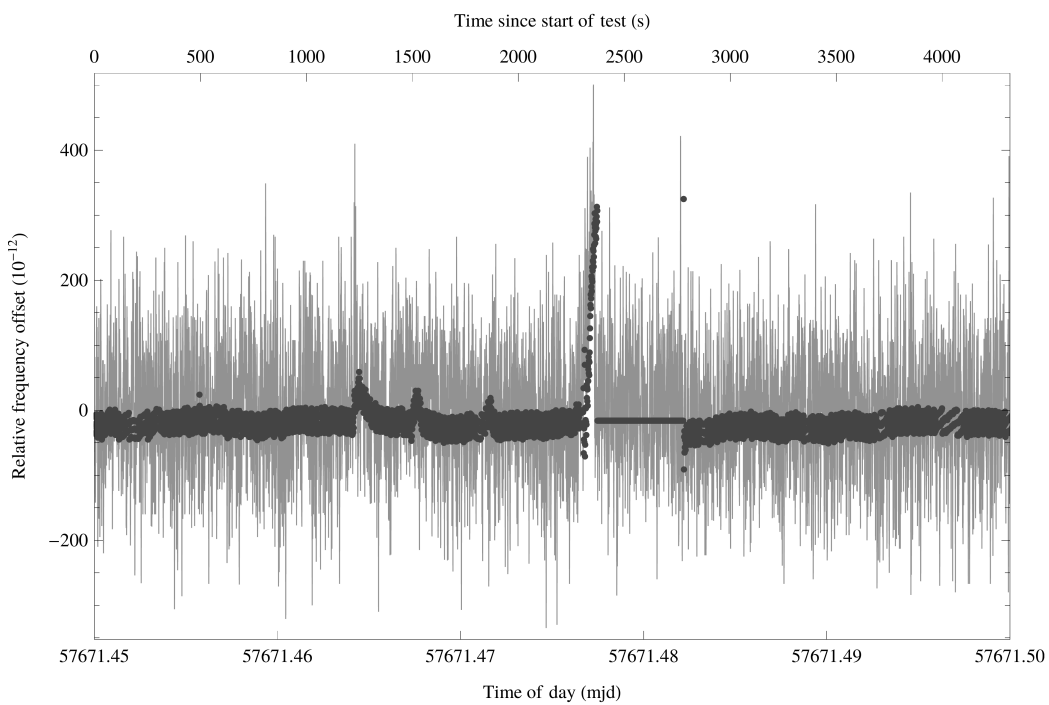


Figure 7.12: The figure shows the relative frequency offset for the atomic clock. The thick dark plot in the middle is from telemetry data gathered from the atomic clock. The lighter plot in the background is data from the CNT-91 frequency counter.




Figure 7.13: The figure shows the phase offset in nanoseconds for the atomic clock. The thick dark plot is from telemetry data gathered from the atomic clock. The lighter thin plot is data from the CNT-91 frequency counter.

7.4.5 Test one results

The observations from test one revealed some surprising results. When the antennas were moved horizontally the altitude solved by sensor two dramatically changed. The movement also resulted in a change in the timing solution. The GPS filter did, however, perform as expected. We discovered that the reason why the step five and six failed to trigger the filter, was that the limit for speed (see section 7.1 for more) was erroneously set too high.

7.5 Test two

7.5.1 Goal of test two

The goal of the second test was to test the clock model based filters. The test should make the atomic clock controller disable the atomic clock's disciplining mode and steer the atomic clock based on the clock model's predictions. We also wanted to test the location and speed filter again to make sure we had fixed the error in the filter configuration from the last test.

7.5.2 Change in setup

The setup in test one is quite similar to test two. There is however some changes. Sensor one is no longer connected to the atomic clock controller because the location and speed filter no longer was the main priority to test.

7.5.3 Test two filter limits

Table 7.2 shows the filter limits used in test two. Note that the speed deviation is set to 1 knot. Table 7.3 shows the limits for the clock model based filters.

Table 7.2: Filter thresholds used during test two

Config value	Sensor 2
Altitude reference	122.427
Longitude reference	1102.1934
Latitude reference	5958.5231
Speed reference	0
Altitude deviation	10
Longitude deviation	0.005
Latitude deviation	0.005
Speed deviation	1

Table 7.3: Clock model filter configuration

Phase limit	50
Steer limit	50
Time constant	10000
Warmup time	2
Prediction limit	200

7.5.4 Description

The description of test two is similar to test one (see section 7.4.2). One difference is that greater care was taken in obtaining an accurate time for the steps.

- 13:12:00 - 125: Started to move antenna two towards north.
- 13:12:45 - 170: Reached destination.
- 13:18:00 - 485: Started the move back to original location.
- 13:19:00 - 545: Reached destination.
- 13:23:00 - 785: Waved the antenna around at an increasing tempo in a half circle motion.
- 13:23:45 - 830: Stopped waving.
- 13:27:00 - 1070: Manually enabled the disciplining of the atomic clock.

7.5.5 Observations

7.5.5.1 Sensor Server logs

By reviewing the log produced by the sensor server the following was observed:

- No false positives, the filters were not triggered before the test started.
- The location and speed filter was triggered by at 13:12:10 and cleared at 13:19:54

```
1 [10/24/16 - 13:12:10] [ ALARM ] Sensor 2 triggered LS filter!  
2 ...  
3 [10/24/16 - 13:19:54] [ ALARM ] Sensor 2 cleared LS filter!
```

- The location and speed filter was triggered again at 13:23:12 for only a second. It was then triggered again two seconds later and was not cleared until 13:23:43.

```
1 [10/24/16 - 13:23:12] [ ALARM ] Sensor 2 triggered LS filter!  
2 ...  
3 [10/24/16 - 13:23:13] [ ALARM ] Sensor 2 cleared LS filter!  
4 ...  
5 [10/24/16 - 13:23:15] [ ALARM ] Sensor 2 triggered LS filter!  
6 ...  
7 [10/24/16 - 13:23:43] [ ALARM ] Sensor 2 cleared LS filter!
```

- More interesting was the fact that the *Frequency correction filter* was triggered.

```
1 [10/24/16 - 13:12:42] [ ALARM ] Steer > predicted!
```

7.5.5.2 GPS data

Figure 7.14 shows the plotted GPS data collected during test two. As with the GPS data observed during test one, there is a clear correlation between the plotted data and when the antenna was moved. The strange phenomena as described in test one (7.4.3.2) where the altitude was reported to be way too low, occurred in test two as well.




Figure 7.14: The figure shows (from the top) the latitude, longitude, altitude, speed and number of satellites as plotted from the GPS data collected during test two.

7.5.6 Timing measurements

Figure 7.15 shows measurements done during test two as two graphs with two time series each. The thin solid line in the top panel shows phase offset in nanoseconds. It reveals that the atomic clock was *not* steered based on the clock model once the disciplining was disabled. The panel at the bottom of figure 7.15 shows relative frequency offset as measured by the CNT-91 frequency counter. The thin darkly colored line is a plot of the data received as telemetry from the atomic clock. The plot shows that the disciplining of the atomic clock stopped after roughly 140 seconds.

7.5.7 Test two results

As mentioned under subsection 7.5.6, the disciplining of the atomic clock was successfully deactivated when the steer value reported by the atomic clock exceeded the steer value that was predicted by the model. The phase jump filter was not triggered because the phase offset never reached the configured limit. The frequency correction filter, on the other hand, was triggered because the steer value exceeded the configured limit (see table 7.3 for limits). It also shows that the atomic clock sadly was *not* steered using the clock model.

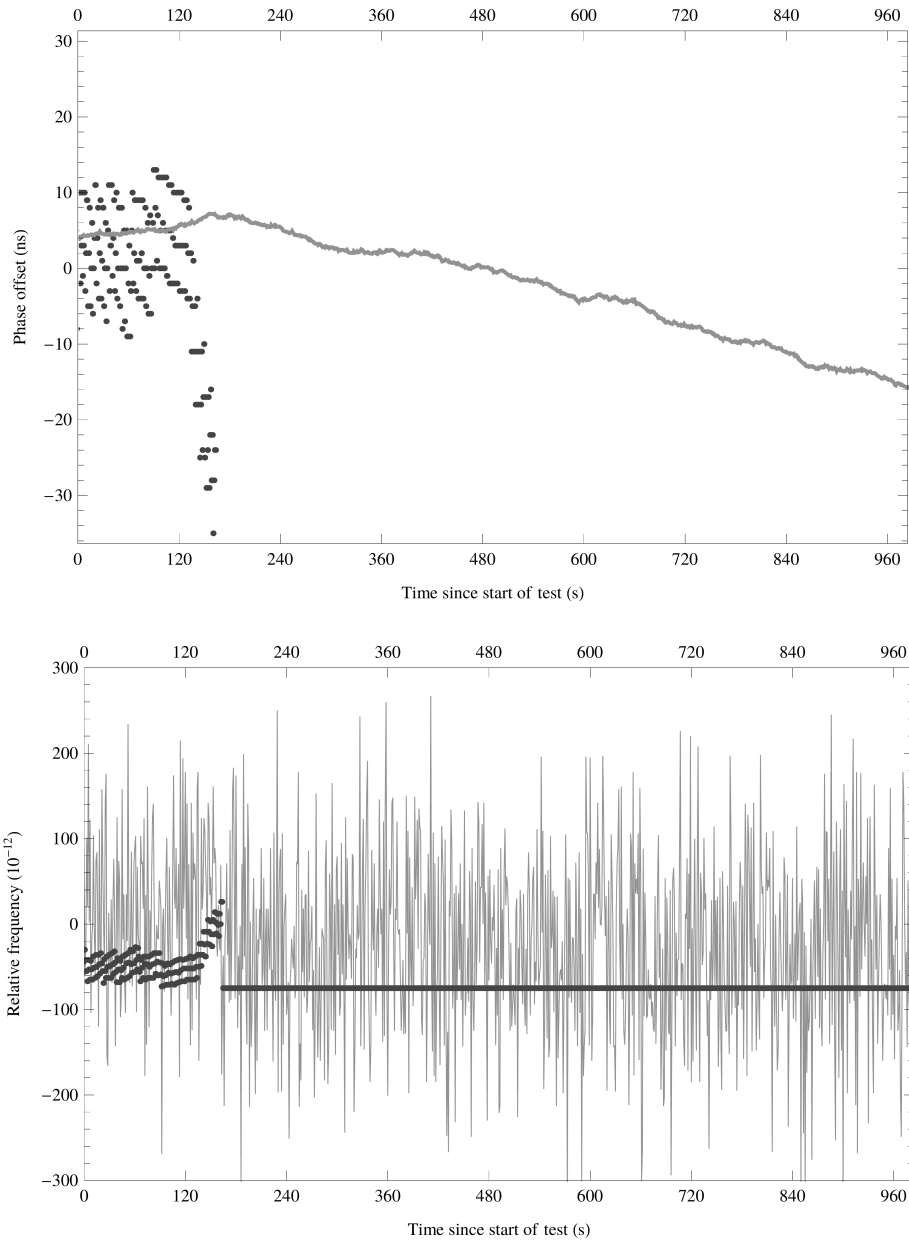


Figure 7.15: Timing measurements and clock telemetry data. The upper panel is phase offset in nanoseconds. The long, sloped light line is a plot of the measurements done by the CNT-91 frequency counter. The dotted dark plot, is also the phase offset but as reported by the atomic clock. The lighter plot in the background of the bottom panel is frequency steering as measured by the CNT-91 frequency counter. The dark plot is frequency offset as reported by the atomic clock.

7.6 Unplanned disturbance

The data presented in this section was gathered while the atomic clock controller was building the clock model for one of the planned tests. The data is interesting because it shows a disturbance from an unknown source. The data shown in the figures are from the 5 October 2016 to 6 October 2016. Figure 7.16 shows that about 10 minutes after midnight on 6 October 2016, the atomic clock entered "holdover mode", indicating that the 1 PPS signal from the GPS receiver was lost. Figure 7.17 shows the GPS receiver's solved altitude, longitude, latitude, speed and number of satellites. By examining the figure, it is obvious that the signal was lost minutes after the 57667 MJD (midnight) mark. The GPS receiver did not achieve consistent lock before approximately 57667.32 MJD (7:45 in the morning). In the meantime, the clock stability was impaired as figure 7.18 and 7.19 clearly show. Figure 7.19 is interesting because it shows how the atomic clock seems to have a delay in its steering algorithm. If the atomic clock controller applied the clock model filters while this data was collected, the phase jump filter would have triggered before the atomic clock would have been able to apply steering.

We have no idea what might be the origin of the disturbance. The following is therefore pure speculation: It is possible that a trucker spent the night parked by the road used a GPS jammer to hide his or her activities from an employer, and that the observed disturbance is a reflection of the jammer's signal. We did not find any traces of the disturbance when examining the logs from sensor one. Sensor one might have been shielded from a reflection because of its position. The time-frame for the disturbance also supports this theory since the disturbance started around midnight and ended at around 7:45, in other words "a good night sleep". This is of course only speculation as we do not have any data supporting it.




Figure 7.16: The figure shows the disciplined mode as reported by the atomic clock.

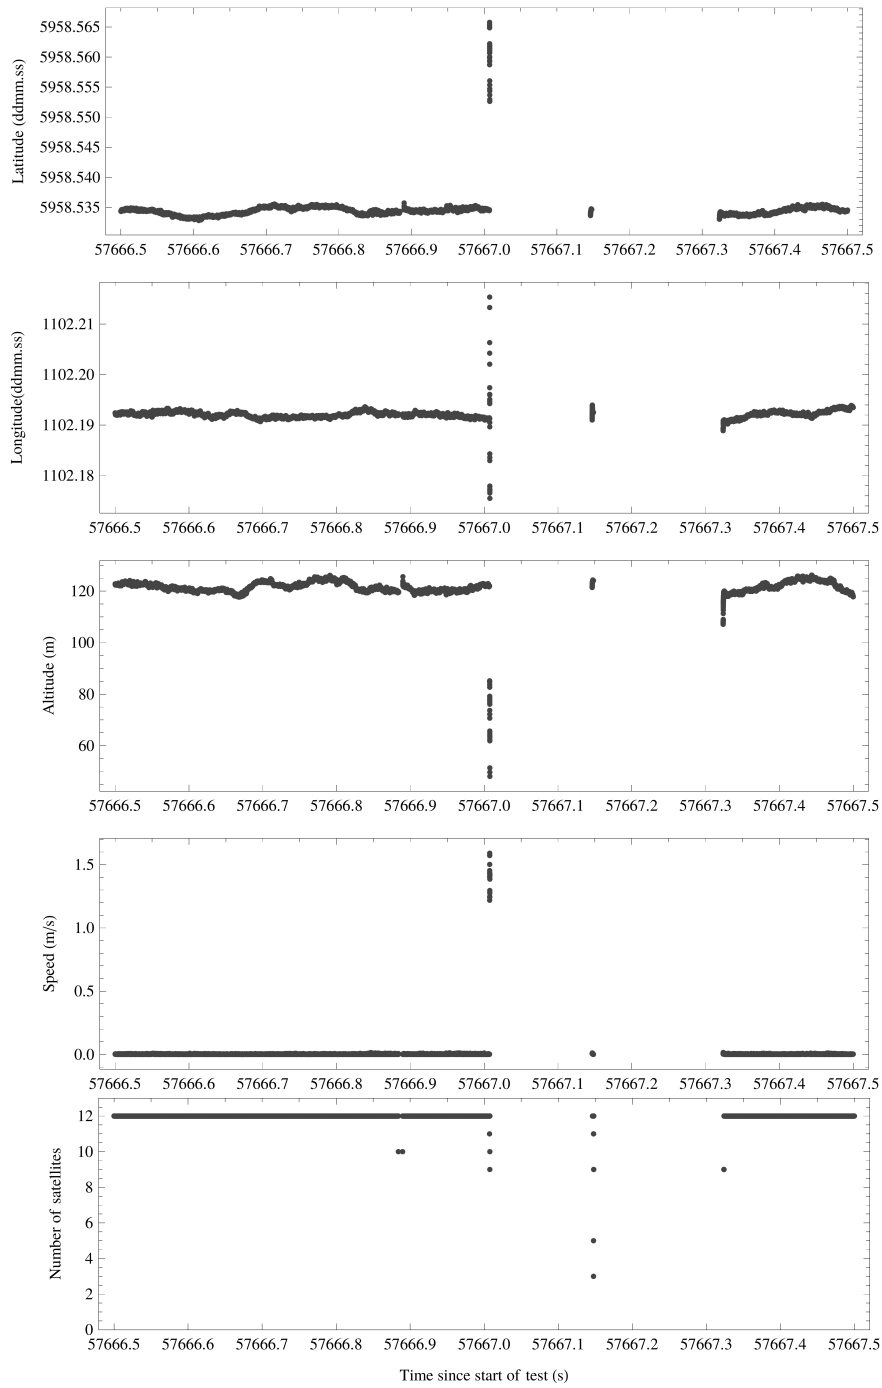


Figure 7.17: Figure shows solved position, speed and number of satellites as reported by the GPS receiver.

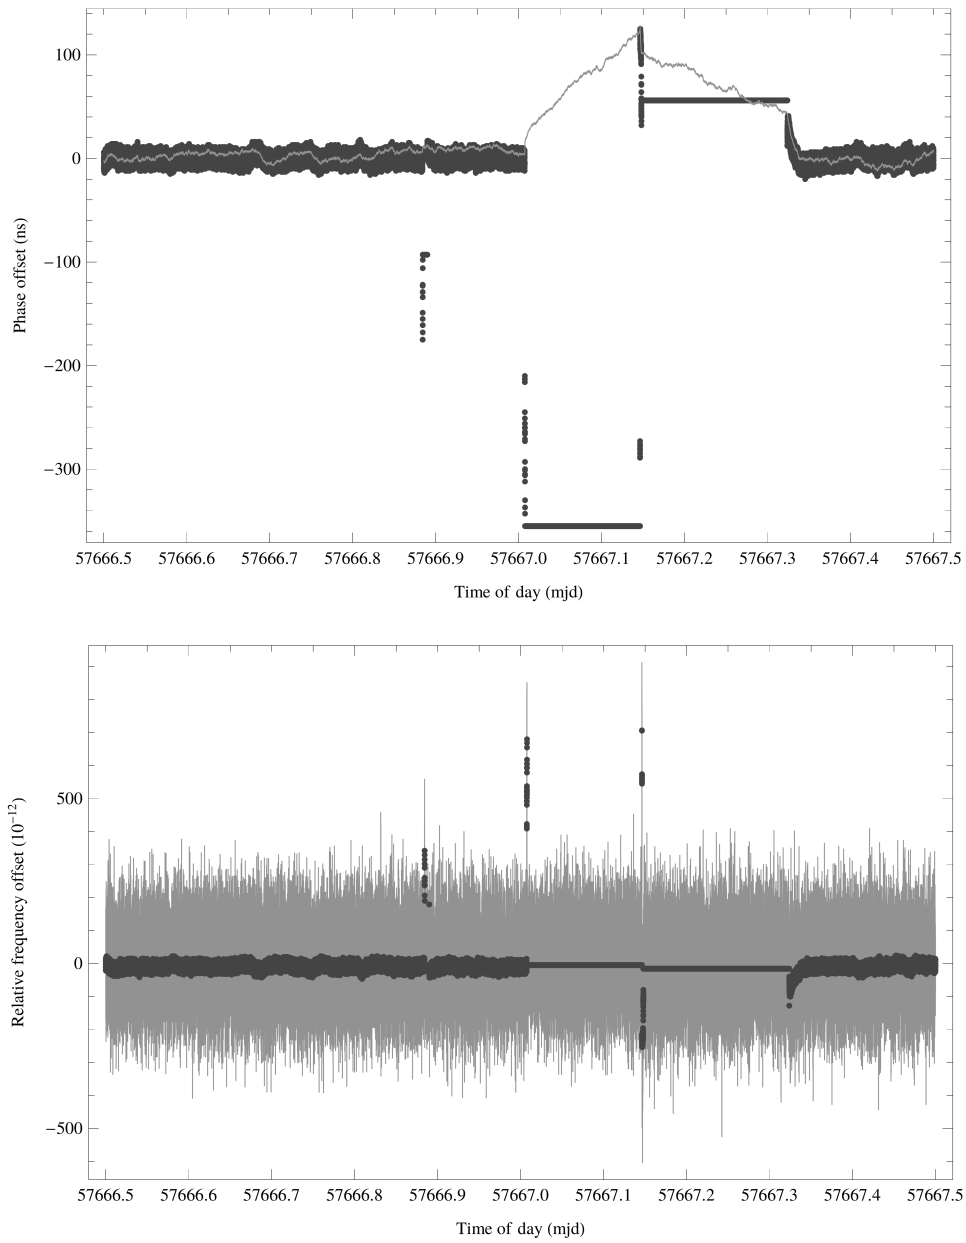


Figure 7.18: Top figure shows phase offset in nanoseconds. The thin lightly colored line is as measured by the CNT-91 frequency counter. The solid darker line is from the telemetry string from the atomic clock.

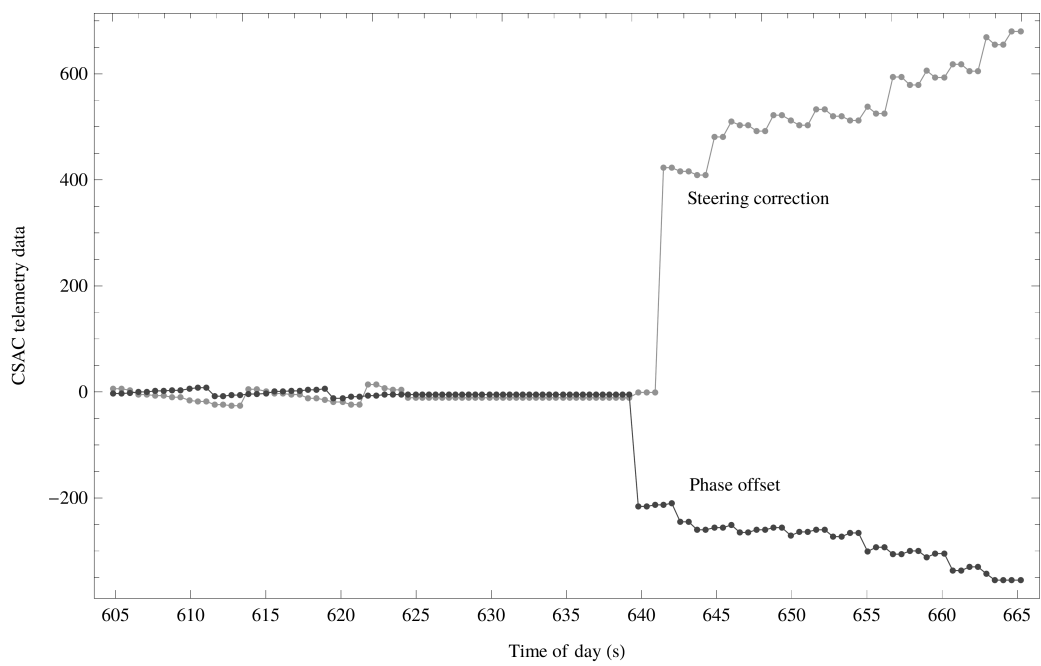


Figure 7.19: Figure shows the phase offset and steering correction as reported in the telemetry string from the atomic clock.

Chapter 8

Discussion

8.1 Test results

The GPS manipulation test as described in chapter 7 demonstrated the ability to quickly identify a GPS disturbance and also demonstrated the ability to protect the atomic clock by disabling the disciplining. The system did, however, fail to steer the atomic clock. When we attempted to reset and debug the system, we encountered a bug in the clock model as well. The bug affected the way the clock model initializes from the configuration file, which meant that the model would need at least 48 hours to rebuild. This is probably just a minor bug but because of the time constraint, we were forced to seize testing.

8.2 Shortcomings in current implementation

This section is used to discuss some of the shortcomings in the current version of the atomic clock controller and sensor server architecture. Some functions were never implemented and others were not finished in time.

8.2.1 Resizing shared memory segments

Ideally the shared memory segments containing the client list should be resizable and its size should depend on the number of connected clients. This proved difficult and I was not able to implement it. M. Kerrisk explains in his book "The Linux Programming Interface" [27] that most UNIX implementations do not support resizing of a memory map like the shared memory segments used in the sensor server implementation. There is however a non-portable and Linux-specific system call named `mremap()` that can be used

on Linux systems for this purpose. Unfortunately the address returned by `mremap()` might be different from the old address to the shared memory segment. This would mean that a pointer inside the shared segment might no longer be valid after a resize operation has been performed. A way to avoid this problem, caused by the remapping would be to use offsets instead of pointers when referring to addresses in the mapped region. While troubleshooting problems I had using `mremap()`, I stumbled upon a bug report in the the Kernel Bug Tracker [28] reported by someone with similar issues as I was having. This indicates that the trouble I had, might have been because of a bug in the Linux kernel. I have yet to confirm this but it did convince me to leave the implementation with its shortcomings rather than potentially wasting my time on something way out of my reach. The waste of memory would never be substantial anyway considering that the size of the `client_table_entry` struct is a modest 4664 bytes.

8.2.2 Atomic clock management

The sensor server should have used a separate process to handle the filters and communication with the atomic clock. This would free the processes handling the connections to clients from dealing with the filters, and make the filter abstraction more complete. This would also make the clock model cleaner. The atomic clock model is already logically separated from the filters associated with it, but because of the way the code is organized the use of the model implies the use of the filter. The atomic clock controller should still communicate with the atomic clock on its own like it does today, retrieving telemetry data, but the aforementioned process could keep track of the atomic clock's discipline status, steering and other functionality, thus creating a more generic way for the system to communicate with the atomic clock. The GPS based filter could greatly benefit from this approach because they today do not do anything but log occurrences where they were triggered.

8.2.3 External MJD calculation

Modified Julian Day (MJD) is a way to express both date and time as a single number. It's convenient when doing calculations with dates, for example: the difference between the MJD for one day and the next, is exactly 1¹. Modified Julian Date relies heavily on by the clock model (see B for more about the clock model). During testing a python script (see E for more about this

¹For example, the MJD for 24/10-2016 00:00 is 57685.0 and the MJD for 25/10-2016 00:00 is 57686.0. 12'o clock at the 25/10 would be 57685.5

script) was used to calculate MJD. This script was scheduled to be replaced by a module written in C, but was left in use because of time constraints. The script is called upon by using `popen()` which in turn calls `fork()` to run the script. The clock model is updated every second which means that this script is also invoked every second. This is of course counterintuitive considering that one of our goals for the software was to focus on efficiency.

8.2.4 External Atomic clock communication

For some reason, it proved to be a significant challenge to implement a solution in C to configure, read and write from the atomic clock. The best solution did not even provide a reliable means of communicating with the atomic clock even though communication with the GPS receivers which in theory should have been exactly the same, was no problem. Once time got tight, I made the decision to drop the development of the atomic clock serial communication module written in C, and decided to use `query_csac.py` by invoking with `popen()`.

8.3 Choice of programming language

The atomic clock controller software was originally planned to be written in Java since this was my most fluent programming language. Java is a great language. It is object oriented, it has a garbage collector and a lot of useful libraries. As development commenced it quickly became apparent that some parts of the code would be performance-critical and that portability really was not that important anyway. The platform was already chosen and I could not think of any reason for us to change it. I decided to look at other languages. Because performance was a concern Python was also quickly dismissed as an option. C++ would probably have been the best choice but never having written anything in C before made it sound more exciting and like a nice opportunity to learn something new. During the planning phase of the atomic clock controller development, raspbian-2015-05-07 was the latest build. It came with GCC 4.6.3 which only had experimental support for C11([29]). With C11 no longer considered an option, C99 was the obvious choice given its attractive features like:

- Variable-length arrays.
- Single line comments.
- `snprintf()` as standard [30].

8.4 Alternative approaches

When planning on how to execute our proposal, these were among the ideas that came up.

8.4.1 Single computer, many GPS receivers

A single computer is used to run the atomic clock controller software. The atomic clock controller does not include a server/client model, but the receivers used to collect data are all connected to the computer through whatever USB ports available or made available by the use of USB hubs. With this approach you are not dependent on a network, but it limits the number of GPS receivers you could connect as the USB specification limits the number possible endpoints to an absolute 127([31, pp. 3]) because of addressing. This does not mean that 127 devices can be connected because a single device might use more than one endpoint. It is also worth mentioning that a USB hub might "reserve" multiple endpoints. Depending on the GPS receivers and how they are made, this number might be reduced even further by the power usage of the connected devices. Depending on how far each GPS receiver is distanced from the atomic clock controller, a signal amplifier might be necessary to compensate for the signal attenuation. In some cases where a network is absent, this approach might be only option.

8.4.2 Store in database and analyze

With this approach, the idea of a GPS receiver and Raspberry PI as a single "sensor" unit is the same as with sensor server approach. The difference is that each sensor stores the collected data in a database. The atomic clock controller software monitors the clock directly as with the sensor server approach, but the data in the database is routinely queried and analyzed. The strength with this approach is that data is easily stored, shared and maintained by a single entity. The complexity of the client software would be the same as with the sensor server approach, but the atomic clock controller software could be implemented with less complexity as no client/server architecture or shared memory schemes would be necessary. During planning this approach seemed promising but was rejected because it was thought that it might not be time-sensitive enough. It was also some doubt concerning whether or not the ability to store data to a database actually was important. Once the different filters and algorithms was in place, it turned that the database functionality would have been nice but not of any real importance for the atomic clock controller to perform its tasks.

Chapter 9

Conclusion

The prototype developed and described in this report, demonstrates that a fully operational spoof proof atomic clock controller would resist a spoofing attack mounted with a sophisticated GPS spoofer like the Civil GPS spoofer [3]. We demonstrated the current implementation's ability to detect and in a limited sense mitigate, a simulated spoofing attack, using multi-layered defense mechanisms. The frequency steering filter made possible by the clock model would have detected steering attempts larger than 50×10^{-12} or 0.05 nanoseconds per seconds as shown in section 7.5. Even if a spoofing attack was done carefully and slowly enough not to trigger the clock model based filters, it would not have been able to spoof two different receivers without giving away the attack (see section 2.3.2.2). This would have required multiple spoofers spoofing individual GPS receivers. The spoofers would have to be meticulously tuned in order not to spoof neighboring receivers. This makes spoofing attempts very challenging to execute. Introducing a third receiver, which with sensor server architecture can be done with ease, would make it even harder.

During the GPS manipulation tests as described in chapter 7, the sensor server architecture proved its worth and worked perfectly. The system was stable and performed as designed over longer periods of time (over 4 days). The system logged data reliably and without fault, accepted connections and was responsive at all times tested. There were no memory leaks or segmentation faults and the system facilitated the detection of our GPS manipulation attempt with low latency. The system would probably react even faster if it received data from the atomic clock and GPS receivers at a higher frequency. The communication with the atomic clock also worked flawlessly, providing data for logging and the model as expected.

We have demonstrated the efficiency of creating a detection network using the sensor server architecture running on commodity hardware receiving

generic GPS data from commodity receivers. We have demonstrated that it is possible to detect GPS disturbances without using an atomic clock as reference but just by using GPS receivers. The detection network combined with an atomic clock and associated model may provide both detection and mitigation during a spoofing attack.

Appendices

Appendix A

Acknowledgments

A.1 Contributions

Harald Hauglin, Chief engineer at Justervesenet was responsible for the following:

- Describing the clock model algorithm as seen under chapter B.
- The concept of CSAC SMACC (Chip scale atomic clock smart miniature atomic clock controller) that this work is built on.
- Created the graphs used in section 7.4.3.2, 7.4.4, 7.5.5.2, 7.5.6 and 7.6.
- Calculating the preliminary filter reference values, both for clock model based filters and GPS based filters. These were based on analysis of data gathered as described under subsection 7.1. Optimal choice of filter parameters is considered to be outside of the scope of this work.

Mr. Matt Davis is the author of `jduutil.py`[32], a library written for converting dates and time to MJD and back. Permission to use the `jduutil.py` was granted by Mr. Davis. The correspondence with Mr. Davis is included in the appendix under F.1.

Appendix B

Clock model

Clock modelling and clock filters in the spoof proof clock controller V 2.0
20161027 HHA

B.1 Introduction

The goal of clock modelling is to provide an estimate of two key parameters of the clock "state", the frequency offset and the clock drift, i.e. the rate of change of the frequency offset. The clock model will be used for two purposes: (1) As a reference for the clock frequency correction filter, in order to determine whether the current clock correction, as calculated and applied by the CSAC disciplining algorithm, is consistent with normal behavior of the clock; (2) In the case that a valid external disciplining pulse is lost, the model will be used to calculate the frequency corrections to be applied to the CSAC by the "spoof" proof clock controller. The model described below is a simple way of modelling the clock state, motivated by being easy to implement, more than being optimal for the task.

B.2 Input data for the clock model

Input data for the clock model comes from the CSAC telemetry string, sampled nominally every second. The key data is contained in the field identified by "Steer". This is the CSAC disciplining algorithm's current estimate of the frequency correction (in relative units) that has to be applied to the CSAC microwave synthesizer to correct for the frequency error of the CSAC "physics package". Ref [16]. Sample telemetry string and explanation of the data fields below are from the CSAC User Guide[ref]. Note that the "Steer" value reported by the CSAC in disciplining mode is the frequency correction

that is applied to make the CSAC in sync with the external applied reference pulse. If the reference pulse is accurate – such as the 1 PPS pulse output from a properly operating GPS timing chip - then the negative "Steer" value is an estimate of the free-running CSAC frequency offset, i.e. a calibration of the CSAC.

As long as there is a valid and accurate reference pulse, a series of "Steer" values provide the basis for estimating the frequency offset of the CSAC, as well as its rate of change (drift).

One should note that the PPS output from the GPS chip has substantial noise in the short term. In contrast, the CSAC itself is more stable than the GPS reference over timespans up to 10000 s [16]. As a consequence, the steering corrections calculated by the CSAC are also noisy and that noise primarily due to noise in the reference pulse. In order to provide good estimates of the clock state, it is useful to filter sampled steering data to remove the influence of noise. This will be described next. Figure B.1 show an example telemetry string.

Listing B.1: Example of a telemetry string received from the atomic clock

```
1 0,0x0000,1209CS00909,0x0010
   ,4381,0.86,1.573,17.62,0.996,28.26,-24,---,-1,1,1268126502,586969,1.0
```

Identifier	Description	Notes
Status	Unit Status	See Note 1
Alarm	Pending Unit Alarms	See Note 3
SN	Unit serial number	See Note 2
Mode	Mode of operation	see 6.4.3 for bit definitions.

Figure B.1: Part 1 of table showing telemetry parameters. The table is taken from the SA.45's manual [16]

Contrast	Indication of signal level	Typically ≈ 4000 when locked, and ≈ 0 when unlocked.
LaserI	Laser current (mA)	
TCXO	Tuning Voltage (V)	0-2.5 VDC tuning range $\approx \pm 10$ ppm
HeatP	Physics package Heater Power (mW)	Typical 15mW under NOC
Sig	DC Signal Level (V)	Typical 1.0 V under NOC
Temp	Unit temperature ($^{\circ}\text{C}$)	Absolute accuracy is $\pm 2^{\circ}\text{C}$
Steer	Frequency adjust	In pp10 ¹²
ATune	Analog tuning voltage input	Reads "--" when analog tuning is disabled
Phase	1PPS_CSAC-1PPS_EXT (ns)	Only present when disciplining enabled
DiscOK	Discipline status (0-2)	0=acquiring, 1=locked, 2=holdover
TOD	Time (seconds)	Starts at 0 upon powerup unless set by command
LockT	Time since lock (seconds)	Starts at 0 upon lock
FWver	Firmware version	

Figure B.2: Part 2 of table showing telemetry parameters. The table is taken from the SA.45's manual [16]

B.3 Smoothing of sampled frequency steering data and estimates of the clock state

The frequency steer filter will be based on smoothed values of current and previous "steer" data from the CSAC telemetry string. The method described here is a very simple approach, primarily chosen for its ease of implementation. It is based on the assumption that "Steer" values are sampled at regular intervals. The design of an optimized approach for clock parameter estimation is way beyond the scope of this work. Calculation of smoothed values. Smoothed values of the clock steering correction are calculated using an exponential filter. Input to the calculation is the current sampled "Steer" value, `steer_current`, along with its associated timestamp (in mjd), `t_current`. The exponential filter has a parameter, w , whose inverse is the weight that each new sampled value adds to the existing smooth value:

$$steer_smooth_current = \frac{w - 1}{w} steer_smooth_previous + \frac{1}{w} steer_current$$

Smoothing is initialized by setting `steer_smooth_previous = steer_current`. Since the current smoothed value of the steering parameter is a weighted average of all previous samples as well, it effectively "lags" behind in time. The effect of this lag can be calculated by applying the same exponential smoothing to the associated timestamps

$$t_smooth_current = \frac{w - 1}{w} t_smooth_previous + \frac{1}{w} t_current$$

Again, smoothing is initialized by setting $t_smooth_previous = t_{i_current}$. Updated smoothed values t_smooth and $steer_smooth$ will be computed every time the telemetry string is read (about once per second). Representative data for sampled and smoothed "steer" values are shown in fig XXX. Daily updated clock model Smoothed values will be used to estimate the frequency drift. A simple approach that "works" is to use values at the start of every day and compare to the values from the previous day. Again, this is not an optimized approach. At the start of every new day set the following:

$$t_smooth_today = t_smooth_current \quad steer_smooth_today = steer_smooth_current$$

and label previous day's values $t_smooth_yesterday$ and $steer_smooth_yesterday$. The drift of the steering parameter can be estimated as:

$$steer_drift = \frac{(steer_smooth_today - steer_smooth_yesterday)}{(t_smooth_today - t_smooth_yesterday)}$$

The predicted steering parameter for a given point in time t can now be computed as a simple linear relationship

$$steer_predicted(t) = steer_smooth_today + (t - t_smooth_today) * steer_drift$$

The output of this model is the predicted CSAC steering correction for a given point in time t .

B.4 Phase jump filter – fast timing filter

The "fast" timing filter will check the value of the "Phase" CSAC telemetry phase field. The CSAC disciplining algorithm steers the clock so that the generated PPS output is in sync with the reference PPS input. The reference value for the phase offset filter is therefore 0. Limits for acceptable deviations from the reference value has to be based on normal noise in the phase data. A representative series of phase data is shown in figure B.3.

The filter parameter is a constant $phase_limit$ will be read from a configuration file. Flag for invalid timing will be raised when:

$$abs(phase_current) > phase_limit$$

Tentatively $phase_limit$ is set to 50 ns.




Figure B.3: Figure shows phase offset in nanoseconds as measured by the CSAC

B.5 Frequency correction filter

The frequency correction filter will check the value of the CSAC telemetry "steer" field and compare it to the steering correction predicted by the clock model. Limits for acceptable deviations from the predicted value has to be based on normal noise in the frequency steering data. A representative series of sampled steering data along with model predictions is shown in figure B.4. Filtering of the current frequency steering can now be implemented as follows:

$$Abs(steer_current - steer_predicted(t_current)) > steer_limit$$

The parameter `steer_limit` will be read from a configuration file. The limit is tentatively set to 50 based on the data below. Note that this is a relative frequency correction in units of (10^{12}) , corresponding to rate of change of 0.05 ns/s in the timing output of the clock.




Figure B.4: Sampled clock steering corrections (gray), smoothed values (black solid line) and predicted clock steering values (dashed). Sampled data at 1 s intervals have been smoothed using an exponential filter with filter parameter $w = 100000$.

Appendix C

Data acquisition

C.1 CSAC Logger source code

```
1  '''
2  :Author: Aril Schultzen
3  :Email: aschultzen@gmail.com
4  '''
5
6  import ctypes
7  import fileinput
8  import sys
9  import datetime
10 import time
11 import io
12 import os
13 import serial
14 import jdutil
15
16
17 def get_today_mjd():
18     today = datetime.datetime.utcnow()
19     return jdutil.jd_to_mjd(jdutil.datetime_to_jd(today))
20
21
22 def t_print(message):
23     current_time = datetime.datetime.now().time()
24     complete_message = "[" + str(
25         current_time.isoformat(
26             )) + "]" + "[" + message + "]"
27     print(complete_message)
28
29
30 def main_routine():
31     log_file = open("dp.txt", "a+")
32     t_print("Started CSAC logging script")
33     ser = serial.Serial("/dev/ttyUSB0", 57600, timeout=0.1)
34     sio = io.TextIOWrapper(
35         io.BufferedRWPair(ser,
36             ser),
37         encoding='ascii',
38         newline="\r")
39
```

```

40     while(True):
41         log_file = open("dp.txt", "a+")
42         ser.write(b'^')
43         time.sleep(0.1)
44         telemetry = sio.readline()
45         output = str(get_today_mjd()) + "," + telemetry
46         log_file.write(output)
47         log_file.close()
48         time.sleep(1)
49
50 if __name__ == '__main__':
51     main_routine()

```

C.2 GPS Logger source code

```

1  '''
2  :Author: Aril Schultzen
3  :Email: aschultzen@gmail.com
4  '''
5
6  """
7  GPS Logger requires:
8  - Python v.2.7
9  - python-mysqldb
10
11  EXPECTED TABLE
12  -----
13
14  create table gprmc (
15      id INT NOT NULL AUTO_INCREMENT,
16      sensorID INT ,
17      fix_time TIME,
18      recv_warn VARCHAR(5),
19      latitude DECIMAL(10,5),
20      la_dir VARCHAR(5),
21      longitude DECIMAL(10,5),
22      lo_dir VARCHAR(5),
23      speed DECIMAL(10,5),
24      course DECIMAL(5,2),
25      fix_date DATE,
26      variation DECIMAL(5,2),
27      var_dir VARCHAR(5),
28      faa VARCHAR(5),
29      checksum VARCHAR(5),
30      mjd VARCHAR(50),
31      alt DECIMAL(5,2),
32      PRIMARY KEY (id) );
33  """
34
35  import ctypes
36  import MySQLdb as mdb
37  import ConfigParser
38  import fileinput
39  import sys
40  import datetime
41  import time
42  import io
43  import os
44  import serial

```

```

45 import jdutil
46 from subprocess import call
47
48 config = ConfigParser.ConfigParser()
49
50
51 def dbConnect():
52     con = mdb.connect(config.get('db', 'ip'), config.get('db', 'user'),
53                     config.get('db', 'password'), config.get('db', 'database'))
54     return con
55
56
57 def dbClose(dbConnection):
58     dbConnection.close()
59     t_print("Connection to database closed")
60
61
62 def initConfig():
63     configFile = "config.ini"
64     config.read(configFile)
65
66
67 def t_print(message):
68     current_time = datetime.datetime.now().time()
69     complete_message = "[" + str(
70         current_time.isoformat(
71         )) + "]" + "[" + message + "]"
72     print(complete_message)
73
74
75 def format_date_string(date_s):
76     split = date_s.split(".")
77     split = split[::-1]
78     split = ''.join(split)
79     return split
80
81
82 def insert(con, data):
83     st = data
84     temp = st[12]
85     checksum = temp[1] + temp[2] + temp[3]
86     faa = temp[0]
87     x = con.cursor()
88     date = st[9][4:6] + st[9][2:4] + st[9][0:2]
89     st[9] = date
90
91     try:
92         query = ("INSERT INTO " + config.get('db', 'table') +
93                "(sensorID, fix_time, recv_warn, latitude, la_dir, longitude, lo_dir, ) " +
94                "(speed, course, fix_date, variation, var_dir, faa, checksum, mjd, alt) VALUES " +
95                "(" + config.get('general', 'sensorID') + "," + st[1] + "," + st[2] +
96                "," + st[3] + "," + st[4] + "," + st[5] + "," + st[6] + "," + st[7] +
97                "," + st[8] + "," + st[9] + "," + st[10] + "," + st[11] +
98                "," + faa + "," + st[14] + "," + st[13] + ");")
99         x.execute(query)
100        con.commit()
101    except:
102        con.rollback()
103
104 # Function used to reset the serial configuration
105 # in Linux in case its mangled by something'
106

```

```

107
108 def reset_serial():
109     call("stty -F " + config.get('gps', 'port') + " icanon", shell=True)
110
111
112 def get_today_mjd():
113     today = datetime.datetime.utcnow()
114     return jdutil.jd_to_mjd(jdutil.datetime_to_jd(today))
115
116
117 def main_routine():
118     initConfig()
119     t_print("GPS logger started!")
120     reset_serial()
121     con = dbConnect()
122     counter = 0
123     data = ""
124
125     while(True):
126         ser = serial.Serial(
127             config.get('gps',
128                 'port'),
129             config.get('gps',
130                 'baud'),
131             timeout=0.1)
132         sio = io.TextIOWrapper(io.BufferedRWPair(ser, ser), newline="\r")
133         time.sleep(1)
134         while True:
135             temp = sio.readline()
136             if(temp.find("GNRMC") == 1):
137                 data = temp
138                 data = data.split(",")
139                 sio.readline() # Reading forward manually
140                 temp = sio.readline()
141                 temp = temp.split(",")
142                 data.append(str(temp[9]))
143                 data.append(str(get_today_mjd()))
144                 counter = counter + 1
145                 if(counter == int(config.get('general', 'discard_interval'))):
146                     insert(con, data)
147                     counter = 0
148         dbClose(con)
149
150 if __name__ == '__main__':
151     main_routine()

```

C.3 GPS Logger source code

```

1 '''
2 :Author: Aril Schultzen
3 :Email: aschultzen@gmail.com
4 '''
5 # This script attempts to connect to the
6 # Sensor Server at <ip> : "port" and
7 # IDs itself as <id>. It will then
8 # poll the time solved by the GNSS receiver
9 # connected to Sensor<id> until
10 # terminated.
11

```

```

12 import socket
13 import sys
14 import time
15
16 ip = "10.1.0.46"
17 port = 10001
18 id = 1
19
20 try:
21     s = socket.socket(socket.AF_INET, socket.SOCK_STREAM)
22 except socket.error, msg:
23     print 'Failed to create socket. Error code: ' + str(msg[0]) + ' , Error message : ' + msg[1]
24     sys.exit();
25 try:
26     remote_ip = socket.gethostbyname( ip )
27
28 except socket.gaierror:
29     print 'Could not resolve hostname'
30     sys.exit()
31
32 s.connect((remote_ip , port))
33 s.sendall(b'IDENTIFY -10')
34 recv_buff = s.recv(1024)
35
36 while(1):
37     s.sendall(b'PRINTTIME' + str(id))
38     time.sleep(0.1)
39     recv_buff = s.recv(1024)
40     recv_buff = recv_buff.strip('>\n')
41     print("Sensor " + str(id) + " GNSS solved time: " + recv_buff)
42     time.sleep(0.9)

```

Appendix D

Sensor server software

D.1 Client

sensor_client.c

```
1  #include "sensor_client.h"
2
3  /* CONFIG */
4  #define CONFIG_SERIAL_INTERFACE "serial_interface:"
5  #define CONFIG_CLIENT_ID "client_id:"
6  #define CONFIG_LOG_NAME "log_file_name:"
7  #define CONFIG_LOG_NMEA "log_nmea:"
8  #define CONFIG_FILE_PATH "client_config.ini"
9  #define DEFAULT_SERIAL_INTERFACE "/dev/ttyACMO"
10 #define CONFIG_CONNECTION_ATTEMPTS_MAX "connection_attempts_max:"
11 #define CONFIG_ENTRIES 5
12
13 struct config_map_entry conf_map[1];
14
15 static int identify(int session_fd, int id);
16 static int create_connection(struct sockaddr_in *serv_addr, int *session_fd,
17                             char *ip, int portno);
18 static void receive_nmea(int gps_serial, struct raw_nmea_container *nmea_c);
19 static int format_nmea(struct raw_nmea_container *nmea_c);
20 static void initialize_config(struct config_map_entry *conf_map,
21                              struct config *cfg);
22 static int start_client(int portno, char* ip);
23 static int usage(char *argv[]);
24
25
26 /* Identify the client for the server */
27 static int identify(int session_fd, int id)
28 {
29     /* Converting from int to string */
30     char id_str[5];
31     bzero(id_str, 5);
32     sprintf(id_str, " %d", id); //Notice the space in the second parameter.
33     int read_status = 0;
34
35     /* Declaring message string */
36     char identify_message[sizeof(PROTOCOL_IDENTIFY) + sizeof(id_str) + 1];
```

```

37
38     /* copying */
39     memcpy(identify_message, PROTOCOL_IDENTIFY, sizeof(PROTOCOL_IDENTIFY));
40     memcpy(&identify_message[8], id_str, sizeof(id_str));
41
42     write(session_fd, identify_message, sizeof(identify_message));
43
44     char buffer[100];
45     while ( (read_status = read(session_fd, buffer, sizeof(buffer)-1)) > 0) {
46         if(strstr((char*)buffer, PROTOCOL_OK ) == (buffer)) {
47             /* ID not used. Accepting. */
48             t_print("ID %d accepted by server. \n", id);
49             return 0;
50         } else {
51             /* ID in use. Rejected. */
52             t_print("ID %d rejected by server, already in use. \n", id);
53             return -1;
54         }
55     }
56     /* Something happened during read. read() returns -1 at error */
57     return read_status;
58 }
59
60 /* Create connection to server */
61 static int create_connection(struct sockaddr_in *serv_addr, int *session_fd,
62                             char *ip, int portno)
63 {
64     if((*session_fd = socket(AF_INET, SOCK_STREAM, 0)) < 0) {
65         t_print("Could not create socket \n");
66         return -1;
67     }
68
69     memset(serv_addr, '0', sizeof(*serv_addr));
70
71     serv_addr->sin_family = AF_INET;
72     serv_addr->sin_port = htons(portno);
73
74     if(inet_pton(AF_INET, ip, &(serv_addr->sin_addr))<=0) {
75         t_print("inet_pton error occured! \n");
76         return 1;
77     }
78
79     if( connect(*session_fd, (struct sockaddr *)serv_addr,
80               sizeof(*serv_addr)) < 0) {
81         return 1;
82     }
83
84     return 0;
85 }
86
87 /* Get chosen NMEA from GPS receiver */
88 static void receive_nmea(int gps_serial, struct raw_nmea_container *nmea_c)
89 {
90     char buffer[SENTENCE_LENGTH * 2];
91     int position = 0;
92     memset(buffer, '\0', sizeof(buffer));
93
94     bool rmc = false;
95     bool gga = false;
96
97     /* Get a load of THIS timebomb!! */
98     while(1) {

```



```

99     while(position < 100) {
100         read(gps_serial, buffer+position, 1);
101         if( buffer[position] == '\n' ) break;
102         position++;
103     }
104
105     if(strstr(buffer, RMC ) != NULL) {
106         memcpy(nmea_c->raw_rmc, buffer, position+1);
107         nmea_c->raw_rmc[position + 2] = '\0';
108         rmc = true;
109     }
110
111     if(strstr(buffer, GGA ) != NULL) {
112         memcpy(nmea_c->raw_gga, buffer, position+1);
113         nmea_c->raw_rmc[position + 2] = '\0';
114         gga = true;
115     }
116
117     if(rmc && gga) {
118         break;
119     }
120     position = 0;
121 }
122 }
123
124 /* Send received NMEA data to server */
125 static int format_nmea(struct raw_nmea_container *nmea_c)
126 {
127     int nmea_prefix_length = 6;
128     memcpy(nmea_c->output, "NMEA \n", nmea_prefix_length);
129     int total_length = 0;
130     int newline_length = 1;
131
132     /* RMC */
133     int rmc_length = strlen(nmea_c->raw_rmc);
134     memcpy( nmea_c->output+nmea_prefix_length, nmea_c->raw_rmc, rmc_length );
135     //nmea_c->output[nmea_prefix_length + rmc_length + newline_length] = '\n';
136
137     /* Updating total length */
138     total_length = rmc_length + nmea_prefix_length; //+ newline_length;
139
140     /* GGA */
141     int gga_length = strlen(nmea_c->raw_gga);
142     memcpy( nmea_c->output+total_length, nmea_c->raw_gga, gga_length );
143     nmea_c->output[total_length + gga_length + newline_length] = '\n';
144
145     /* Updating total length */
146     total_length += gga_length + newline_length;
147
148     return total_length;
149 }
150
151 static int make_log(struct raw_nmea_container *nmea_c, int id, char* log_name)
152 {
153     /* Allocating memory for filename buffer */
154     int filename_length = strlen(log_name) + 100;
155     char filename[filename_length];
156
157     /* Clearing buffer */
158     memset(filename, '\0', filename_length);
159
160     /* Copying name from loaded config */

```

```

161     strcpy(filename, log_name);
162
163     /* Casting int to string */
164     char id_string[10];
165     memset(id_string, '\0', 10);
166     sprintf(id_string, "%d", id);
167
168     /* Concating filename and ID */
169     strcat(filename, id_string);
170
171     char log_buffer[SENTENCE_LENGTH * 2];
172     memset(log_buffer, '\0', SENTENCE_LENGTH * 2);
173     strcat(log_buffer, nmea_c->raw_rmc);
174     log_buffer[strlen(log_buffer)-2] = '\0';
175     log_buffer[strlen(log_buffer)-1] = ',';
176
177     strcat(log_buffer, nmea_c->raw_gga);
178
179     return log_to_file(filename, log_buffer, 1);
180 }
181
182 /* Setting up the config structure specific for the server */
183 static void initialize_config(struct config_map_entry *conf_map,
184                             struct config *cfg)
185 {
186     conf_map[0].entry_name = CONFIG_SERIAL_INTERFACE;
187     conf_map[0].modifier = FORMAT_STRING;
188     conf_map[0].destination = &cfg->serial_interface;
189
190     conf_map[1].entry_name = CONFIG_CLIENT_ID;
191     conf_map[1].modifier = FORMAT_INT;
192     conf_map[1].destination = &cfg->client_id;
193
194     conf_map[2].entry_name = CONFIG_LOG_NAME;
195     conf_map[2].modifier = FORMAT_STRING;
196     conf_map[2].destination = &cfg->log_name;
197
198     conf_map[3].entry_name = CONFIG_LOG_NMEA;
199     conf_map[3].modifier = FORMAT_INT;
200     conf_map[3].destination = &cfg->log_nmea;
201
202     conf_map[4].entry_name = CONFIG_CONNECTION_ATTEMPTS_MAX;
203     conf_map[4].modifier = FORMAT_INT;
204     conf_map[4].destination = &cfg->con_attempt_max;
205 }
206
207 static int start_client(int portno, char* ip)
208 {
209     struct termios tty;
210     memset(&tty, 0, sizeof tty);
211
212     struct sockaddr_in serv_addr;
213     int session_fd = 0;
214     int connection_attempts = 1;
215     int con_status;
216
217     struct raw_nmea_container nmea_c;
218     memset(&nmea_c, 0, sizeof(nmea_c));
219
220     struct config cfg;
221
222     initialize_config(conf_map, &cfg);

```

```

223 int load_config_status = load_config(conf_map, CONFIG_FILE_PATH,
224                                   CONFIG_ENTRIES);
225
226 if(!load_config_status) {
227     t_print("Failed to load the config, using default values\n");
228     memcpy(cfg.serial_interface, DEFAULT_SERIAL_INTERFACE,
229            strlen(DEFAULT_SERIAL_INTERFACE)*sizeof(char));
230
231     /* Picking ID number for client at random */
232     cfg.client_id = rand() % ID_MAX;
233     t_print("Picked ID %d at random\n", cfg.client_id);
234
235     /* Disabling logging */
236     cfg.log_nmea = 0;
237
238     /* Setting retry times to 10 */
239     cfg.con_attempt_max = 10;
240 } else {
241     if(cfg.client_id == 0 || cfg.client_id > ID_MAX) {
242         t_print("Client ID can not be less than 1 or more than %d!\n", ID_MAX);
243         exit(0);
244     }
245 }
246
247 /* Establishing connection to GPS receiver */
248 int gps_serial = open_serial(cfg.serial_interface, GPS);
249 if(gps_serial == -1) {
250     t_print("Connection to GPS receiver failed! Exiting...\n");
251     exit(0);
252 } else {
253     t_print("Connection to GPS receiver established!\n");
254 }
255
256 /* Establishing connection to server */
257 while(connection_attempts <= cfg.con_attempt_max) {
258     con_status = create_connection(&serv_addr, &session_fd, ip, portno);
259     if(con_status == 0) {
260         t_print("Connected to server!\n");
261         break;
262     }
263     t_print("Connection attempt %d failed. Code %d\n", connection_attempts,
264            con_status);
265     sleep(1);
266     connection_attempts++;
267 }
268
269 /* Identifying client for server */
270 if( identify(session_fd, cfg.client_id) == -1 ) {
271     exit(0);
272 }
273
274 if(cfg.log_nmea) {
275     t_print("NMEA data logging enabled\n");
276 }
277
278 while (1) {
279     receive_nmea(gps_serial, &nmea_c);
280     int trans_length = format_nmea(&nmea_c);
281     /* Writing to socket (server) */
282     write(session_fd, nmea_c.output, trans_length);
283     if(cfg.log_nmea) {
284         make_log(&nmea_c, cfg.client_id, cfg.log_name);
285     }

```

```

285     }
286     return 0;
287 }
288
289 static int usage(char *argv[])
290 {
291     t_print("Usage: %s -s <SERVER IP> -p <SERVER PORT> \n", argv[0]);
292     return 0;
293 }
294
295 int main(int argc, char *argv[])
296 {
297     char *ip_address = NULL;
298     char *port_number = NULL;
299
300     if(argc < 5) {
301         usage(argv);
302         return 0;
303     }
304
305     while (1) {
306         char c;
307
308         c = getopt (argc, argv, "s:p:");
309         if (c == -1) {
310             break;
311         }
312         switch (c) {
313             case 's':
314                 ip_address = optarg;
315                 break;
316             case 'p':
317                 port_number = optarg;
318                 break;
319             default:
320                 usage(argv);
321         }
322     }
323
324     if(ip_address == NULL || port_number == NULL) {
325         t_print("Missing parameters! \n");
326         exit(0);
327     }
328
329     start_client(atoi(port_number), ip_address);
330     return 0;
331 }

```

sensor_client.h

```

1  #ifndef SENSOR_CLIENT_H
2  #define SENSOR_CLIENT_H
3
4  // Mine
5  #include "net.h"
6  #include "utils.h"
7  #include "protocol.h"
8  #include "nmea.h"
9  #include "utils.h"
10 #include "serial.h"

```

```

11 |
12 | struct config {
13 |     char serial_interface[100];
14 |     int client_id;
15 |     char log_name[100];
16 |     int log_nmea;
17 |     int con_attempt_max;
18 | };
19 |
20 | /* Used by the client */
21 | struct raw_nmea_container {
22 |     /* Raw data */
23 |     char raw_gga[SENTENCE_LENGTH];
24 |     char raw_rmc[SENTENCE_LENGTH];
25 |     char output[SENTENCE_LENGTH * 2];
26 | };
27 |
28 | #endif /* !SENSOR_CLIENT_H */

```

client_config.ini

```

1 | serial_interface: /dev/ttyS0
2 | client_id: 1
3 | log_nmea: 1
4 | log_file_name: log_sensor
5 | connection_attempts_max: 10

```

query_csac.py

```

1 | import ctypes
2 | import fileinput, sys
3 | import datetime
4 | import time
5 | import io
6 | import os
7 | import serial
8 |
9 | def main_routine():
10 |     # Opening serial stream, use ASCII
11 |     ser = serial.Serial("/dev/ttyUSB0",57600, timeout=0.1)
12 |     sio = io.TextIOWrapper(io.BufferedRWPair(ser, ser),encoding='ascii',newline="\r\n")
13 |
14 |     # Open log file, mostly used for debug
15 |     log_file = open("query_csac.txt", "a+")
16 |
17 |     # The query to use
18 |     query = sys.argv[1].strip("\r\n")
19 |
20 |     # How long to sleep between read from serial con.
21 |     sleep_time = 0.2
22 |
23 |     # The minimum length of the answer
24 |     # for the given query.
25 |     minimum_len = 0
26 |
27 |     if(query == '^' or query == '6'):
28 |         minimum_len = 80
29 |     elif(query == 'F'):
30 |         sleep_time = 0.5

```

```

31     minimum_len = 10
32     elif(query == 'M'):
33         minimum_len = 6
34     elif (query == 'S'):
35         sleep_time = 3
36         minimum_len = 2
37     else:
38         minimum_len = 1
39
40     response_len = 0
41
42     if(len(query) > 1):
43         query = "!" + query + "\r\n"
44
45     retry_count = 0
46
47     while (response_len < minimum_len):
48         ser.write(bytes(query))
49         time.sleep(sleep_time)
50         response = sio.readline()
51         response = response.strip("\r\n\x00")
52         response_len = len(response)
53         retry_count = retry_count + 1
54
55     print(response)
56     ser.close()
57     query = query.strip("\r\n")
58     log_string = ("Issued query " + "'" + query + "'" + str(retry_count) + " times\n")
59     log_file.write(log_string)
60 if __name__ == '__main__':
61     main_routine()

```

D.2 Server

sensor_server.c

```

1  #include "sensor_server.h"
2
3  /* VERSION */
4  #define PROGRAM_VERSION "0.8c"
5
6  /* ERRORS */
7  #define ERROR_MAX_CLIENTS_REACHED "CONNECTION REJECTED: MAXIMUM NUMBER OF CLIENTS REACHED\n"
8  #define ERROR_CONFIG_LOAD_FAILED "CONFIG LOAD FAILED: CONFIG FILE CORRUPTED\n"
9  #define ERROR_SEMAPHORE_CREATION_FAILED "SEMAPHORE CREATION FAILED\n"
10 #define ERROR_SOCKET_OPEN_FAILED "ERROR: FAILED TO OPEN SOCKET\n"
11 #define ERROR_SOCKET_BINDING "ERROR: FAILED TO BIND ON %d\n"
12 #define ERROR_CONNECTION_ACCEPT "ERROR: FAILED TO ACCEPT CONNECTION (%d)\n"
13 #define ERROR_FAILED_FORK "ERROR: FORK FAILED (%d)\n"
14 #define ERROR_MISSING_PARAMS "MISSING PARAMETERS!\n"
15
16 /* GENERAL STRINGS */
17 #define PROCESS_REAPED "Process %d reaped. Status: %d Signum: %d\n"
18 #define SIGTERM_RECEIVED "[%d] SIGTERM received!\n"
19 #define SIGINT_RECEIVED "[%d] SIGINT received!\n"
20 #define STOPPING_SERVER "Stopping server...\n"
21 #define CONFIG_LOADED "Config loaded!\n"
22 #define SERVER_RUNNING "Server is running. Accepting connections.\n"

```

```

23 #define WAITING_FOR_CONNECTIONS "Waiting for connections...\n"
24 #define CON_ACCEPTED "Connection accepted\n"
25 #define CLIENT_DISCONNECTED "Client [%d] at [%s] disconnected\n"
26 #define SERVER_STOPPED "Server STOPPED!\n"
27 #define SERVER_STARTING "Sensor server starting...\n"
28 #define CLIENT_KICKED "Client was kicked\n"
29
30 /* USAGE() STRINGS */
31 #define USAGE_DESCRIPTION "Required argument:\n\t -p <PORT NUMBER>\n\n"
32 #define USAGE_PROGRAM_INTRO "Sensor_server: Server part of GPS Jamming/Spoofing system\n\n"
33 #define USAGE_USAGE "Usage: %s [ARGS]\n\n"
34
35 /* CONFIG CONSTANTS*/
36 #define CONFIG_FILE_PATH "config.ini"
37 #define CONFIG_SERVER_MAX_CONNECTIONS "max_clients:"
38 #define CONFIG_SERVER_WARM_UP "warm_up:"
39 #define CONFIG_SERVER_HUMANLY_READABLE "humanly_readable_dumpdata:"
40 #define CONFIG_CSAC_PATH "csac_serial_interface:"
41 #define CONFIG_LOGGING "logging:"
42 #define CONFIG_LOG_PATH "log_path:"
43 #define CONFIG_CSAC_LOG_PATH "csac_log_path:"
44 #define CONFIG_CSAC_LOGGING "csac_logging:"
45 #define SERVER_CONFIG_ENTRIES 8
46
47 /* Server data and stats */
48 struct server_data *s_data;
49
50 /* Shared synchro elements */
51 struct server_synchro *s_synch;
52
53 /* Pointer to shared memory containing the client list */
54 struct client_table_entry *client_list;
55
56 /* Pointer to client list map */
57 struct client_table_entry **client_list_map;
58
59 /* Pointer to shared memory containing config */
60 struct server_config *s_conf;
61
62 /* Pointer to shared CSAC model data */
63 struct csac_model_data *cfd;
64
65 static void remove_client_by_pid(pid_t pid);
66 void remove_client_by_id(int id);
67 static struct client_table_entry* create_client(struct client_table_entry* ptr);
68 static void handle_sigchld();
69 static void handle_sig(int signum);
70 static void initialize_config(struct config_map_entry *conf_map,
71                             struct server_config *s_conf);
72 static int start_server(int port_number);
73 static int usage(char *argv[]);
74 static void setup_session(int session_fd, struct client_table_entry *new_client);
75 static int release_mem_piece(struct client_table_entry* release_me);
76
77 int set_timeout(struct client_table_entry *target,
78               struct timeval h_timeout)
79 {
80     /* setsockopt return -1 on error and 0 on success */
81     target->timeout = h_timeout;
82     if (setsockopt(target->transmission.session_fd, SOL_SOCKET,
83                  SO_RCVTIMEO, (char *)&target->timeout, sizeof(struct timeval)) < 0) {
84         t_print("an error: %s\n", strerror(errno));

```

```

85     return 0;
86 }
87     return 1;
88 }
89
90 /* Prints a formatted string containing server info to monitor */
91 void print_server_data(struct client_table_entry *monitor)
92 {
93     char buffer [1000];
94     int snprintf_status = 0;
95     struct tm *loctime_started;
96     loctime_started = localtime (&s_data->started);
97
98     s_write(&(monitor->transmission), SERVER_TABLE_LABEL,
99           sizeof(SERVER_TABLE_LABEL));
100    s_write(&(monitor->transmission), HORIZONTAL_BAR, sizeof(HORIZONTAL_BAR));
101
102    snprintf_status = snprintf( buffer, 1000,
103                               "PID: %d\n" \
104                               "Number of clients: %d\n" \
105                               "Number of sensors: %d\n" \
106                               "Max clients: %d\n" \
107                               "Sensor Warm-up time: %ds\n" \
108                               "Dump humanly readable data: %d\n" \
109                               "Started: %s" \
110                               "Version: %s\n",
111          s_data->pid,
112          s_data->number_of_clients,
113          s_data->number_of_sensors,
114          s_conf->max_clients,
115          s_conf->warm_up_seconds,
116          s_conf->human_readable_dumpdata,
117          asctime (loctime_started),
118          s_data->version);
119
120    s_write(&(monitor->transmission), buffer, snprintf_status);
121    s_write(&(monitor->transmission), HORIZONTAL_BAR, sizeof(HORIZONTAL_BAR));
122 }
123
124 struct client_table_entry* get_client_by_id(int id)
125 {
126     struct client_table_entry* cli;
127     struct client_table_entry* temp;
128     int found = 0;
129
130     sem_wait(&(s_synch->client_list_sem));
131     list_for_each_entry_safe(cli, temp, &client_list->list, list) {
132         if(cli->client_id == id) {
133             found = 1;
134             break;
135         }
136     }
137     sem_post(&(s_synch->client_list_sem));
138     if(found) {
139         return cli;
140     } else {
141         return NULL;
142     }
143 }
144
145 /* Removes a client with the given PID */
146 static void remove_client_by_pid(pid_t pid)

```



```

147 {
148     struct client_table_entry* cli;
149     struct client_table_entry* temp_remove;
150
151     sem_wait(&(s_synch->client_list_sem));
152     list_for_each_entry_safe(cli, temp_remove,&client_list->list,
153                             list) {
154         if(cli->pid == pid) {
155             /* Decrementing sensor count */
156             if(cli->client_id > 0) {
157                 s_data->number_of_sensors--;
158             }
159             t_print(CLIENT_DISCONNECTED, cli->client_id ,cli->ip);
160             list_del(&cli->list);
161             release_mem_piece(cli);
162         }
163     }
164     /* Decrementing total client count */
165     s_data->number_of_clients--;
166     sem_post(&(s_synch->client_list_sem));
167 }
168
169 /* Removes a client with the given ID */
170 void remove_client_by_id(int id)
171 {
172     struct client_table_entry* cli;
173     struct client_table_entry* temp_remove;
174
175     sem_wait(&(s_synch->client_list_sem));
176     list_for_each_entry_safe(cli, temp_remove,&client_list->list,
177                             list) {
178         if(cli->client_id == id) {
179             list_del(&cli->list);
180             release_mem_piece(cli);
181             s_data->number_of_clients--;
182         }
183     }
184     sem_post(&(s_synch->client_list_sem));
185 }
186
187 static int release_mem_piece(struct client_table_entry* release_me)
188 {
189     int i;
190     for(i = 1; i < s_conf->max_clients; i++){
191         if(client_list_map[i] == NULL){
192             client_list_map[i] = release_me;
193             return 1;
194         }
195         i++;
196     }
197     return 0;
198 }
199
200 static struct client_table_entry* get_mem_piece()
201 {
202     int i;
203     for(i = 1; i < s_conf->max_clients; i++){
204         if(client_list_map[i] != NULL){
205             struct client_table_entry *tmp = client_list_map[i];
206             client_list_map[i] = NULL;
207             return tmp;
208         }

```

```

209     i++;
210 }
211 return NULL;
212 }
213
214 /* Creates an entry in the client list structure and returns a pointer to it*/
215 static struct client_table_entry* create_client(struct client_table_entry* ptr)
216 {
217     sem_wait(&(s_synch->client_list_sem));
218     s_data->number_of_clients++;
219     struct client_table_entry* tmp;
220     tmp = get_mem_piece();
221     list_add_tail( &(tmp->list), &(ptr->list) );
222     sem_post(&(s_synch->client_list_sem));
223     return tmp;
224 }
225
226 /* SIGCHLD Handler */
227 static void handle_sigchld()
228 {
229     pid_t pid;
230     int status;
231     while ((pid = waitpid(-1, &status, WNOHANG)) != -1) {
232         if(pid == 0) {
233             break;
234         }
235
236         if(pid > 0) {
237             remove_client_by_pid(pid);
238         }
239     }
240 }
241
242 /* SIGTERM/INT Handler */
243 static void handle_sig(int signum)
244 {
245     if(signum == 15) {
246         t_print(SIGTERM_RECEIVED, getpid());
247     }
248     if(signum == 2) {
249         t_print(SIGINT_RECEIVED, getpid());
250     }
251     t_print(STOPPING_SERVER, getpid());
252     s_synch->done = 1;
253 }
254
255 /* Setting up the config structure specific for the server */
256 static void initialize_config(struct config_map_entry *conf_map,
257                             struct server_config *s_conf)
258 {
259     conf_map[0].entry_name = CONFIG_SERVER_MAX_CONNECTIONS;
260     conf_map[0].modifier = FORMAT_INT;
261     conf_map[0].destination = &s_conf->max_clients;
262
263     conf_map[1].entry_name = CONFIG_SERVER_WARM_UP;
264     conf_map[1].modifier = FORMAT_INT;
265     conf_map[1].destination = &s_conf->warm_up_seconds;
266
267     conf_map[2].entry_name = CONFIG_SERVER_HUMANLY_READABLE;
268     conf_map[2].modifier = FORMAT_INT;
269     conf_map[2].destination = &s_conf->human_readable_dumpdata;
270

```

```

271     conf_map[3].entry_name = CONFIG_CSAC_PATH;
272     conf_map[3].modifier = FORMAT_STRING;
273     conf_map[3].destination = &s_conf->csac_path;
274
275     conf_map[4].entry_name = CONFIG_LOGGING;
276     conf_map[4].modifier = FORMAT_INT;
277     conf_map[4].destination = &s_conf->logging;
278
279     conf_map[5].entry_name = CONFIG_LOG_PATH;
280     conf_map[5].modifier = FORMAT_STRING;
281     conf_map[5].destination = &s_conf->log_path;
282
283     conf_map[6].entry_name = CONFIG_CSAC_LOG_PATH;
284     conf_map[6].modifier = FORMAT_STRING;
285     conf_map[6].destination = &s_conf->csac_log_path;
286
287     conf_map[7].entry_name = CONFIG_CSAC_LOGGING;
288     conf_map[7].modifier = FORMAT_INT;
289     conf_map[7].destination = &s_conf->csac_logging;
290 }
291
292 /* Sets up the clients structure and initializes data */
293 void setup_session(int session_fd, struct client_table_entry *new_client)
294 {
295     /* Setting the IP address */
296     char ip[INET_ADDRSTRLEN];
297     get_ip_str(session_fd, ip);
298
299     /* Setting the PID */
300     new_client->pid = getpid();
301     new_client->timestamp = time(NULL);
302     strncpy(new_client->ip, ip, INET_ADDRSTRLEN);
303
304     /* Initializing structure, zeroing just to be sure */
305     new_client->client_id = 0;
306     new_client->transmission.session_fd = session_fd;
307
308     /* Zeroing out filters */
309     new_client->fs.ls_f.moved = 0;
310     new_client->fs.ls_f.was_moved = 0;
311
312     new_client->marked_for_kick = 0;
313     new_client->ready = 0;
314
315     /* Setting timeout */
316     struct timeval timeout = {UNIDENTIFIED_TIMEOUT, 0};
317     if(!set_timeout(new_client, timeout)) {
318         t_print("Failed to set timeout for client\n");
319     }
320
321     memset(&new_client->transmission.iobuffer, '0', IO_BUFFER_SIZE*sizeof(char));
322     memset(&new_client->cm.parameter, '0', MAX_PARAMETER_SIZE*sizeof(char));
323
324     /*
325      * Entering child process main loop
326      * (Outer) breaks if server closes.
327      * (Inner) Breaks (disconnects the client) if
328      * respond < 0
329      */
330     while(!s_synch->done) {
331         if(!respond(new_client)) {
332             break;

```

```

333     }
334 }
335 }
336
337 /*
338  * Main loop for the server.
339  * Forks everytime a client connects and calls setup_session()
340  */
341 static int start_server(int port_number)
342 {
343     /* Initializing variables */
344     int server_sockfd;
345     struct sockaddr_in serv_addr;
346     struct config_map_entry conf_map[SERVER_CONFIG_ENTRIES];
347
348     /* Initializing config structure */
349     s_conf = mmap(NULL, sizeof(struct server_config), PROT_READ | PROT_WRITE,
350                 MAP_SHARED | MAP_ANONYMOUS, -1, 0);
351     initialize_config(conf_map, s_conf);
352
353     /* Loading config */
354     int load_config_status = load_config(conf_map, CONFIG_FILE_PATH,
355                                       SERVER_CONFIG_ENTRIES);
356
357     /* Falling back to default if load_config fails */
358     if(load_config_status) {
359         t_print(CONFIG_LOADED);
360         client_list = mmap(NULL,
361                          ((s_conf->max_clients + 1) * sizeof(struct client_table_entry)),
362                          PROT_READ | PROT_WRITE, MAP_SHARED | MAP_ANONYMOUS | MAP_NORESERVE, -1, 0);
363         if(client_list == MAP_FAILED){
364             t_print("Failed to allocate memory for the client list!\n");
365         }
366     } else {
367         t_print(ERROR_CONFIG_LOAD_FAILED);
368         exit(0);
369     }
370
371     client_list_map = malloc((s_conf->max_clients + 1)
372                             * sizeof(struct client_table_entry));
373     int i;
374
375     /* Skip the first entry for some reason */
376     for(i = 1; i < s_conf->max_clients; i++){
377         client_list_map[i] = client_list + i;
378     }
379
380     INIT_LIST_HEAD(&client_list->list);
381
382     /* Create and initialize shared memory for server data */
383     s_data = mmap(NULL, sizeof(struct server_data), PROT_READ | PROT_WRITE,
384                 MAP_SHARED | MAP_ANONYMOUS, -1, 0);
385     if(s_data == MAP_FAILED){
386         t_print("Failed to allocate memory for the server data!\n");
387     }
388
389     bcopy(PROGRAM_VERSION, s_data->version,4);
390     s_data->pid = getpid();
391     s_data->started = time(NULL);
392
393     /* Init shared semaphores and sync elements */
394     s_synch = mmap(NULL, sizeof(struct server_synchro), PROT_READ | PROT_WRITE,

```

```

395         MAP_SHARED | MAP_ANONYMOUS, -1, 0);
396
397     if(s_synch == MAP_FAILED){
398         t_print("Failed to allocate memory for the semaphores!\n");
399     }
400
401     sem_init(&(s_synch->ready_sem), 1, 1);
402     sem_init(&(s_synch->client_list_sem), 1, 1);
403     sem_init(&(s_synch->csac_sem), 1, 1);
404
405     /* Init pointer to shared CSAC_filter data */
406     cfd = mmap(NULL, sizeof(struct csac_model_data), PROT_READ | PROT_WRITE,
407         MAP_SHARED | MAP_ANONYMOUS, -1, 0);
408
409     if(cfd == MAP_FAILED){
410         t_print("Failed to allocate memory for the CSAC filter data!\n");
411     }
412
413     if( &(s_synch->ready_sem) == SEM_FAILED
414         || &(s_synch->client_list_sem) == SEM_FAILED) {
415         t_print(ERROR_SEMAPHORE_CREATION_FAILED);
416         sem_close(&(s_synch->ready_sem));
417         sem_close(&(s_synch->client_list_sem));
418         sem_close(&(s_synch->csac_sem));
419         exit(1);
420     }
421
422
423     pid_t f_pid;
424     f_pid = fork();
425     if(f_pid == 0) {
426         t_print("Forked out CSAC filter [%d]\n", getpid());
427         start_csac_model(cfd);
428         _exit(0);
429     }
430
431     /* Waiting for filter to start */
432     sleep(1);
433     if(s_synch->done)
434         return 1;
435
436     /* Registering the SIGINT handler */
437     struct sigaction sigint_action;
438     memset(&sigint_action, 0, sizeof(struct sigaction));
439     sigint_action.sa_handler = handle_sig;
440     sigaction(SIGINT, &sigint_action, NULL);
441     if (sigaction(SIGCHLD, &sigint_action, 0) == -1) {
442         perror(0);
443         exit(1);
444     }
445
446     /* Registering the SIGTERM handler */
447     struct sigaction sigterm_action;
448     memset(&sigterm_action, 0, sizeof(struct sigaction));
449     sigterm_action.sa_handler = handle_sig;
450     sigaction(SIGTERM, &sigterm_action, NULL);
451     if (sigaction(SIGCHLD, &sigterm_action, 0) == -1) {
452         perror(0);
453         exit(1);
454     }
455
456     /* Registering the SIGCHLD handler */

```

```

457 struct sigaction child_action;
458 child_action.sa_handler = &handle_sigchld;
459 sigemptyset(&child_action.sa_mask);
460 child_action.sa_flags = SA_RESTART | SA_NOCLDSTOP;
461 if (sigaction(SIGCHLD, &child_action, 0) == -1) {
462     perror(0);
463     exit(1);
464 }
465
466 /* Initialize socket */
467 server_sockfd = socket(AF_INET, SOCK_STREAM, 0);
468 if (server_sockfd < 0) {
469     die(62,ERROR_SOCKET_OPEN_FAILED);
470 }
471
472 /*
473  * Initializing the server address struct:
474  * AF_INET = IPV4 Internet protocol
475  * INADDR_ANY = Accept connections to all IPs of the machine
476  * htons(port_number) = Endianness: network to host long(port number).
477  */
478 bzero((char *) &serv_addr, sizeof(serv_addr));
479 serv_addr.sin_family = AF_INET;
480 serv_addr.sin_addr.s_addr = INADDR_ANY;
481 serv_addr.sin_port = htons(port_number);
482
483 /*
484  * Assigns the address (serv_addr) to the socket
485  * referred to by server_sockfd.
486  */
487 if (bind(server_sockfd, (struct sockaddr *) &serv_addr,
488         sizeof(serv_addr)) < 0) {
489     t_print(ERROR_SOCKET_BINDING, port_number);
490     exit(1);
491 }
492
493 /* Marking the connection for listening*/
494 listen(server_sockfd,SOMAXCONN);
495 int session_fd = 0;
496 t_print(SERVER_RUNNING);
497 while (!s_synch->done) {
498     t_print(WAITING_FOR_CONNECTIONS);
499     session_fd = accept(server_sockfd,0,0);
500     if (session_fd!=-1) {
501         if (errno==EINTR) continue;
502         t_print(ERROR_CONNECTION_ACCEPT,errno);
503     }
504     if(s_data->number_of_clients == s_conf->max_clients) {
505         write(session_fd, ERROR_MAX_CLIENTS_REACHED, sizeof(ERROR_MAX_CLIENTS_REACHED));
506         close(session_fd);
507     } else {
508         struct client_table_entry *new_client = create_client(client_list);
509         pid_t pid=fork();
510         if (pid==-1) {
511             t_print(ERROR_FAILED_FORK, errno);
512             /* WHAT HAPPENS WITH THE LIST WHEN FORK FAILS? DEAL WITH IT.*/
513         } else if (pid==0) {
514             close(server_sockfd);
515             setup_session(session_fd, new_client);
516             close(session_fd);
517             if(new_client->marked_for_kick) {
518                 t_print(CLIENT_KICKED, getpid());

```

```

519         }
520         _exit(0);
521     } else {
522         t_print(CON_ACCEPTED);
523         close(session_fd);
524     }
525     }
526 }
527
528 /* Destroying semaphores */
529 sem_destroy(&(s_synch->csac_sem));
530 sem_destroy(&(s_synch->ready_sem));
531 sem_destroy(&(s_synch->client_list_sem));
532
533 /* Freeing */
534 munmap(client_list, sizeof(struct client_table_entry));
535 munmap(s_data, sizeof(struct server_data));
536 munmap(cfd, sizeof(struct csac_model_data));
537 munmap(s_synch, sizeof(struct server_synchro));
538 free(client_list_map);
539
540 /* Closing server FD */
541 close(server_sockfd);
542 t_print(SERVER_STOPPED);
543 return 1;
544 }
545
546 static int usage(char *argv[])
547 {
548     printf(USAGE_USAGE, argv[0]);
549     printf(USAGE_PROGRAM_INTRO);
550     printf(USAGE_DESCRIPTION);
551     return 0;
552 }
553
554 int main(int argc, char *argv[])
555 {
556     char *port_number = NULL;
557
558     /* getopt silent mode set */
559     opterr = 0;
560
561     if(argc < 3) {
562         usage(argv);
563         return 0;
564     }
565
566     while (1) {
567         char c;
568
569         c = getopt (argc, argv, "p:");
570         if (c == -1) {
571             break;
572         }
573
574         switch (c) {
575             case 'p':
576                 port_number = optarg;
577                 break;
578         }
579     }
580

```

```

581     if(port_number == NULL) {
582         printf(ERROR_MISSING_PARAMS);
583     }
584
585     t_print(SERVER_STARTING);
586     start_server(atoi(port_number));
587     exit(0);
588 }

```

sensor_server.h

```

1  /**
2   * @file sensor_server.h
3   * @author Aril Schultzen
4   * @date 13.04.2016
5   * @brief File containing function prototypes, structs and includes for sensor_server.c
6   */
7
8  #ifndef SENSOR_SERVER_H
9  #define SENSOR_SERVER_H
10
11  #define PATH_LENGTH_MAX 1000
12  #define CLIENT_TIMEOUT 5
13  #define MONITOR_TIMEOUT 1000
14  #define UNIDENTIFIED_TIMEOUT 10
15
16  #include <fcntl.h>
17  #include <sys/stat.h>
18  #include "session.h"
19  #include "serial.h"
20  #include "sensor_server_common.h"
21  #include "csac_filter.h"
22
23  /*!@struct*/
24  /*!@brief Contains configuration values for the server
25  */
26  struct server_config {
27     int max_clients;
28     int warm_up_seconds;
29     int human_readable_dumpdata;
30     char csac_path[PATH_LENGTH_MAX];
31     int logging;
32     char log_path[PATH_LENGTH_MAX];
33     int csac_logging;
34     char csac_log_path[PATH_LENGTH_MAX];
35 };
36
37  /*
38   * Made extern because the sessions should
39   * exit if the server is given a SIGINT/TERM
40   */
41  //extern volatile sig_atomic_t done;
42
43  /* Also used by session and action */
44  extern struct client_table_entry *client_list;
45  extern struct server_data *s_data;
46  extern struct server_synchro *s_synch;
47  extern struct server_config *s_conf;
48  extern struct csac_model_data *cfd;

```



```

49
50 /** @brief Removes a client whose ID matches parameter
51 *
52 * Iterates through the linked list and removes the
53 * node containing the client whose ID matches the parameter.
54 * @param id ID for the client
55 * @return Void
56 */
57 void remove_client_by_id(int id);
58
59 /** @brief Returns a client whose ID matches parameter
60 *
61 * Iterates through the linked list and returns
62 * a pointer to the client_table_entry struct in the
63 * list that corresponds with the parameter.
64 * @param id ID for the client
65 * @return client_table_entry *
66 */
67 struct client_table_entry* get_client_by_id(int id);
68
69 /** @brief Prints information about the server.
70 *
71 * Transmits info about the server:
72 * Time when started, PID, number of clients,
73 * number of sensors, max number of clients,
74 * sensor warm-up time and version.
75 *
76 * @param client MONITOR who made the request.
77 * @return Void
78 */
79 void print_server_data(struct client_table_entry *monitor);
80
81 int set_timeout(struct client_table_entry *target,
82                struct timeval h_timeout);
83
84 #endif /* !SENSOR_SERVER_H */

```

config.ini

```

1 humanly_readable_dumpdata: 1
2 max_clients: 100
3 warm_up: 24000
4 csac_serial_interface: /dev/ttyUSB0
5 logging: 1
6 log_path: server_log.txt
7 csac_logging: 1
8 csac_log_path: csac_log.txt

```

sensor_server_common.h

```

1 /**
2 * @file sensor_server_common.h
3 * @author Aril Schultzen
4 * @date 13.04.2016
5 * @brief File containing structs and defines used by session.c, analyzer.c, sensors_server.c and actions.c
6 */
7
8 #ifndef SENSOR_SERVER_COMMON_H

```

```

9  #define SENSOR_SERVER_COMMON_H
10
11 #include <semaphore.h>
12 #include "net.h"
13 #include "colors.h"
14
15 /* General */
16 #define SERVER_TABLE_LABEL "SERVER DATA\n"
17 #define HORIZONTAL_BAR "=====\n"
18 #define ERROR_NO_CLIENT "ERROR: No such client\n"
19 #define ERROR_NO_FILENAME "ERROR: No FILENAME specified\n"
20 #define MAX_FILENAME_SIZE 30
21 #define ID_AS_STRING_MAX 10
22
23 /* Errors */
24 #define ERROR_CODE_NO_FILE -1
25 #define ERROR_CODE_READ_FAILED -2
26 #define ERROR_NO_FILE "ERROR:No such file\n"
27 #define ERROR_READ_FAILED "ERROR:Failed to read file\n"
28
29 /*
30  * command_code struct is used by the parser
31  * to convey an easy to compare command code, as well
32  * as any parameter belonging to that command
33  */
34 struct command_code {
35     int code;
36     char parameter[MAX_PARAMETER_SIZE];
37     int id_parameter;
38 };
39
40 /*!@struct*/
41 /*!@brief Data used by the red_dev_filter.
42  * Read from file.
43  */
44 struct lsf_data {
45     double alt_ref;
46     double lon_ref;
47     double lat_ref;
48     double speed_ref;
49     double alt_dev;
50     double lon_dev;
51     double lat_dev;
52     double speed_dev;
53 };
54
55 struct disturbed_values {
56     int lat_disturbed;
57     int lon_disturbed;
58     int alt_disturbed;
59     int speed_disturbed;
60 };
61
62 struct lsf {
63     struct lsf_data lsf_d;
64     int moved;
65     int was_moved;
66     struct disturbed_values dv;
67 };
68
69 struct filters {
70     struct lsf ls_f;

```

```

71 };
72
73 /*
74  * CLIENT TABLE STRUCT
75  *
76  * list_head list: The head in the list of clients
77  * pid: Process ID for the client connection (See "fork")
78  * session_fd: The file descriptor for the session.
79  * client_id: The connected clients ID
80  * iobuffer: A general purpose buffer for in and output
81  * heartbeat_timeout: Number of seconds of inactivity before disconnect
82  * ip: Clients IP Address.
83  * cm: Command code. Used for quick comparison after commands
84  * are parsed by command parser.
85  */
86
87 /*!@struct*/
88 /*!@brief Contain information about every client that is connected.
89  */
90 struct client_table_entry {
91     struct list_head list; /* The head of the client list */
92     struct transmission_s transmission; /* Everything needed for socket com. */
93     struct timeval timeout; /* Timeout in seconds if not active */
94     struct command_code cm; /* See command code */
95     struct nmea_container nmea; /* All NMEA data associated with the client */
96     pid_t pid; /* The process ID */
97     time_t timestamp; /* When last analyzed */
98     int client_id; /* Clients ID */
99     int client_type; /* Client type, SENSOR or MONITOR */
100    int ready; /* Ready status */
101    int marked_for_kick; /* Marked for kicked at next opportunity */
102    char ip[INET_ADDRSTRLEN]; /* Clients IP address */
103    struct filters fs;
104 };
105
106 /* Server info shared with processes */
107 struct server_data {
108     int number_of_clients; /* Number of clients currently connected */
109     int number_of_sensors; /* Number of sensors, subset of clients */
110     time_t started; /* When the server was started */
111     pid_t pid; /* Servers PID */
112     char version[4]; /* Version of server software */
113 };
114
115 /* Synchronization elements shared with processes */
116 struct server_synchro {
117     sem_t ready_sem;
118     sem_t csac_sem;
119     sem_t client_list_sem;
120     volatile sig_atomic_t done;
121 };
122
123 /*
124  * Roles of client, either SENSOR or MONITOR.
125  * A monitor is only used to monitor the programs state.
126  */
127 enum client_type {
128     SENSOR,
129     MONITOR
130 };
131
132 #endif /* !SENSOR_SERVER_COMMON_H */

```

session.c

```
1  #include "session.h"
2
3  /* ERRORS*/
4  #define ERROR_ILLEGAL_COMMAND "ERROR:Illegal command\n"
5  #define ERROR_NO_ID "ERROR:Client not identified\n"
6  #define ERROR_ID_IN_USE "ERROR:ID in use\n"
7  #define ERROR_ILLEGAL_MESSAGE_SIZE "\rERROR:Illegal message size\n"
8  #define ERROR_WARMUP_NOT_SENSOR "ERROR:Warm-up only applies to sensors\n"
9  #define ERROR_DUMPDATA_FAILED "ERROR:Failed to dump data\n"
10 #define ERROR_LOADDATA_FAILED "ERROR:Failed to load data\n"
11 #define ERROR_NO_COMMAND "ERROR:No command specified\n"
12 #define ERROR_LSD_LOAD_FAILED "ERROR:Failed to load LS data from file\n"
13 #define ERROR_CHECKSUM_FAILED "ERROR: Checksum failed!\n"
14 #define ERROR_ILLEGAL_NMEA "ERROR: Received illegal/corrupt NMEA data\n"
15
16 static int nmea_ready();
17 static int extract_nmea_data(struct client_table_entry *cte);
18 static void calculate_nmea_average(struct client_table_entry *cte);
19 static void calculate_nmea_diff(struct client_table_entry *cte);
20 static int parse_input(struct client_table_entry *cte);
21
22 /*
23  * Used by spawned client processes to "mark" that their NMEA
24  * data is ready for processing. Works as a barrier in a way.
25  */
26 static int nmea_ready()
27 {
28     struct client_table_entry* c_iter;
29     struct client_table_entry* temp;
30     int ready = 0;
31
32     /* iterating through the list of clients */
33     list_for_each_entry_safe(c_iter, temp, &client_list->list, list) {
34         /* Is it a SENSOR?*/
35         if(c_iter->client_type == SENSOR) {
36             /* Is it ready?*/
37             if(c_iter->ready){
38                 ready++;
39             }
40         }
41     }
42     /* if everyone is ready, clear ready flag and carry on */
43     if(ready == s_data->number_of_sensors) {
44         list_for_each_entry_safe(c_iter, temp, &client_list->list, list) {
45             c_iter->ready = 0;
46         }
47         return 1;
48     } else {
49         return 0;
50     }
51 }
52
53 /* Extract position data from NMEA */
54 static int extract_nmea_data(struct client_table_entry *cte)
55 {
56     int buffsize = 100;
57     char buffer[buffsize];
58     memset(&buffer, 0, buffsize);
59
60     /* Extracting latitude */
```

```

61     if(substring_extractor(LATITUDE_START,LATITUDE_START + 1,',',buffer, bufsize,
62                           cte->nmea.raw_rmc, strlen(cte->nmea.raw_rmc)))
63     {
64         cte->nmea.lat_current = atof(buffer);
65     } else {
66         return 0;
67     }
68
69
70     /* Extracting longitude */
71     if(substring_extractor(LONGITUDE_START,LONGITUDE_START + 1,',',buffer, bufsize,
72                           cte->nmea.raw_rmc, strlen(cte->nmea.raw_rmc)))
73     {
74         cte->nmea.lon_current = atof(buffer);
75     } else {
76         return 0;
77     }
78
79     /* Extracting altitude */
80     if(substring_extractor(ALTITUDE_START,ALTITUDE_START + 1,',',buffer, bufsize,
81                           cte->nmea.raw_gga, strlen(cte->nmea.raw_gga)))
82     {
83         cte->nmea.alt_current = atof(buffer);
84     } else {
85         return 0;
86     }
87
88     /* Extracting speed */
89     if(substring_extractor(SPEED_START,SPEED_START + 1,',',buffer, bufsize,
90                           cte->nmea.raw_rmc, strlen(cte->nmea.raw_rmc)))
91     {
92         cte->nmea.speed_current = atof(buffer);
93     } else {
94         return 0;
95     }
96
97     return 1;
98 }
99
100 /* Calculate the average NMEA values */
101 static void calculate_nmea_average(struct client_table_entry *cte)
102 {
103     /* Updating number of samples */
104     cte->nmea.n_samples++;
105
106     /* Updating total */
107     cte->nmea.lat_total = cte->nmea.lat_total + cte->nmea.lat_current;
108     cte->nmea.lon_total = cte->nmea.lon_total + cte->nmea.lon_current;
109     cte->nmea.alt_total = cte->nmea.alt_total + cte->nmea.alt_current;
110     cte->nmea.speed_total = cte->nmea.speed_total + cte->nmea.speed_current;
111
112     cte->nmea.lat_average = ( cte->nmea.lat_total / cte->nmea.n_samples );
113     cte->nmea.lon_average = ( cte->nmea.lon_total / cte->nmea.n_samples );
114     cte->nmea.alt_average = ( cte->nmea.alt_total / cte->nmea.n_samples );
115     cte->nmea.speed_average = ( cte->nmea.speed_total / cte->nmea.n_samples );
116 }
117
118 /*
119  * Calculate the diff between current
120  * NMEA values and the average values.
121  */
122 static void calculate_nmea_diff(struct client_table_entry *cte)

```

```

123 {
124     cte->nmea.lat_avg_diff = (cte->nmea.lat_current - cte->nmea.lat_average);
125     cte->nmea.lon_avg_diff = (cte->nmea.lon_current - cte->nmea.lon_average);
126     cte->nmea.alt_avg_diff = (cte->nmea.alt_current - cte->nmea.alt_average);
127     cte->nmea.speed_avg_diff = (cte->nmea.speed_current - cte->nmea.speed_average);
128 }
129
130 /*
131  * Parses input from clients. Return value indicates status.
132  * Uses the command_code struct to convey parameter and command code.
133  *
134  * Returns -1 if size is wrong
135  * Returns 0 if protocol is not followed
136  * Returns 1 if all is ok
137  */
138
139 static int parse_input(struct client_table_entry *cte)
140 {
141     char *incoming = cte->transmission.iobuffer;
142
143     /* INPUT TO BIG */
144     if(strlen(incoming) > (MAX_PARAMETER_SIZE + MAX_COMMAND_SIZE) + 2) {
145         return -1;
146     }
147
148     /* INPUT TO SMALL */
149     if(strlen(incoming) < (MIN_PARAMETER_SIZE + MIN_COMMAND_SIZE) + 2) {
150         return -1;
151     }
152
153     /* ZEROING COMMAND CODE */
154     cte->cm.code = 0;
155     /* ZEROING ID_PARAMETER */
156     cte->cm.id_parameter = 0;
157
158     /* NMEA */
159     if(strstr((char*)incoming, PROTOCOL_NMEA ) == (incoming)) {
160         cte->cm.code = CODE_NMEA;
161     }
162
163     /* IDENTIFY */
164     else if(strstr((char*)incoming, PROTOCOL_IDENTIFY ) == (incoming)) {
165         int length = (strlen(incoming) - strlen(PROTOCOL_IDENTIFY) );
166         memcpy(cte->cm.parameter, (incoming)+(strlen(PROTOCOL_IDENTIFY)*(sizeof(char))),
167             length);
168         cte->cm.code = CODE_IDENTIFY;
169     }
170
171     /* IDENTIFY SHORT */
172     else if(strstr((char*)incoming, PROTOCOL_IDENTIFY_SHORT ) == (incoming)) {
173         int length = (strlen(incoming) - strlen(PROTOCOL_IDENTIFY_SHORT) );
174         memcpy(cte->cm.parameter,
175             (incoming)+(strlen(PROTOCOL_IDENTIFY_SHORT)*(sizeof(char))), length);
176         cte->cm.code = CODE_IDENTIFY;
177     }
178
179     /* DUMPDATA */
180     else if(strstr((char*)incoming, PROTOCOL_DUMPDATA ) == (incoming)) {
181         int length = (strlen(incoming) - strlen(PROTOCOL_DUMPDATA) );
182         memcpy(cte->cm.parameter, (incoming)+(strlen(PROTOCOL_DUMPDATA)*(sizeof(char))),
183             length);
184         cte->cm.code = CODE_DUMPDATA;

```

```

185     }
186
187     /* DUMPDATA_SHORT */
188     else if(strstr((char*)incoming, PROTOCOL_DUMPDATA_SHORT ) == (incoming)) {
189         int length = (strlen(incoming) - strlen(PROTOCOL_DUMPDATA_SHORT) );
190         memcpy(cte->cm.parameter,
191             (incoming)+(strlen(PROTOCOL_DUMPDATA_SHORT)*(sizeof(char))), length);
192         cte->cm.code = CODE_DUMPDATA;
193     }
194
195     /* PRINT_LOCATION */
196     else if(strstr((char*)incoming, PROTOCOL_PRINT_LOCATION ) == (incoming)) {
197         int length = (strlen(incoming) - strlen(PROTOCOL_PRINT_LOCATION) );
198         memcpy(cte->cm.parameter,
199             (incoming)+(strlen(PROTOCOL_PRINT_LOCATION)*(sizeof(char))), length);
200         cte->cm.code = CODE_PRINT_LOCATION;
201     }
202
203     /* PRINT_LOCATION_SHORT */
204     else if(strstr((char*)incoming, PROTOCOL_PRINT_LOCATION_SHORT ) == (incoming)) {
205         int length = (strlen(incoming) - strlen(PROTOCOL_PRINT_LOCATION_SHORT) );
206         memcpy(cte->cm.parameter,
207             (incoming)+(strlen(PROTOCOL_PRINT_LOCATION_SHORT)*(sizeof(char))), length);
208         cte->cm.code = CODE_PRINT_LOCATION;
209     }
210
211     /* PRINTTIME */
212     else if(strstr((char*)incoming, PROTOCOL_PRINTTIME ) == (incoming)) {
213         int length = (strlen(incoming) - strlen(PROTOCOL_PRINTTIME) );
214         memcpy(cte->cm.parameter,
215             (incoming)+(strlen(PROTOCOL_PRINTTIME)*(sizeof(char))), length);
216         cte->cm.code = CODE_PRINTTIME;
217     }
218
219     /* PRINTCLIENTS */
220     else if(strstr((char*)incoming, PROTOCOL_PRINTCLIENTS ) == (incoming) ||
221         strstr((char*)incoming, PROTOCOL_PRINTCLIENTS_SHORT ) == (incoming)) {
222         cte->cm.code = CODE_PRINTCLIENTS;
223     }
224
225     /* PRINTSERVER */
226     else if(strstr((char*)incoming, PROTOCOL_PRINTSERVER ) == (incoming) ||
227         strstr((char*)incoming, PROTOCOL_PRINTSERVER_SHORT ) == (incoming)) {
228         cte->cm.code = CODE_PRINTSERVER;
229     }
230
231     /* KICK */
232     else if(strstr((char*)incoming, PROTOCOL_KICK ) == (incoming)) {
233         int length = (strlen(incoming) - strlen(PROTOCOL_KICK) );
234         memcpy(cte->cm.parameter, (incoming)+(strlen(PROTOCOL_KICK)*(sizeof(char))),
235             length);
236         cte->cm.code = CODE_KICK;
237     }
238
239     /* EXIT */
240     else if(strstr((char*)incoming, PROTOCOL_EXIT ) == (incoming)) {
241         cte->cm.code = CODE_DISCONNECT;
242     }
243
244     /* DISCONNECT */
245     else if(strstr((char*)incoming, PROTOCOL_DISCONNECT ) == (incoming) ||
246         strstr((char*)incoming, PROTOCOL_DISCONNECT_SHORT ) == (incoming)) {

```

```

247         cte->cm.code = CODE_DISCONNECT;
248     }
249
250     /* HELP */
251     else if(strstr((char*)incoming, PROTOCOL_HELP ) == (incoming) ||
252            strstr((char*)incoming, PROTOCOL_HELP_SHORT ) == (incoming)) {
253         cte->cm.code = CODE_HELP;
254     }
255
256     /* PRINTAVGDIFF */
257     else if(strstr((char*)incoming, PROTOCOL_PRINTAVGDIFF ) == (incoming) ||
258            strstr((char*)incoming, PROTOCOL_PRINTAVGDIFF_SHORT ) == (incoming)) {
259         cte->cm.code = CODE_PRINTAVGDIFF;
260     }
261
262     /* LISTDUMPS */
263     else if(strstr((char*)incoming, PROTOCOL_LISTDUMPS ) == (incoming) ||
264            strstr((char*)incoming, PROTOCOL_LISTDUMPS_SHORT ) == (incoming)) {
265         cte->cm.code = CODE_LISTDUMPS;
266     }
267
268     /* LOADDATA */
269     else if(strstr((char*)incoming, PROTOCOL_LOADDATA ) == (incoming)) {
270         int length = (strlen(incoming) - strlen(PROTOCOL_LOADDATA) );
271         memcpy(cte->cm.parameter, (incoming)+(strlen(PROTOCOL_LOADDATA)*(sizeof(char))),
272              length);
273         cte->cm.code = CODE_LOADDATA;
274     }
275
276     /* LOADDATA_SHORT */
277     else if(strstr((char*)incoming, PROTOCOL_LOADDATA_SHORT ) == (incoming)) {
278         int length = (strlen(incoming) - strlen(PROTOCOL_LOADDATA_SHORT) );
279         memcpy(cte->cm.parameter,
280              (incoming)+(strlen(PROTOCOL_LOADDATA_SHORT)*(sizeof(char))), length);
281         cte->cm.code = CODE_LOADDATA;
282     }
283
284     /* QUERYCSAC */
285     else if(strstr((char*)incoming, PROTOCOL_QUERYCSAC ) == (incoming)) {
286         int length = (strlen(incoming) - strlen(PROTOCOL_QUERYCSAC) );
287         memcpy(cte->cm.parameter,
288              (incoming)+(strlen(PROTOCOL_QUERYCSAC)*(sizeof(char))), length);
289         cte->cm.code = CODE_QUERYCSAC;
290     }
291
292     /* QUERYCSAC_SHORT */
293     else if(strstr((char*)incoming, PROTOCOL_QUERYCSAC_SHORT ) == (incoming)) {
294         int length = (strlen(incoming) - strlen(PROTOCOL_QUERYCSAC_SHORT) );
295         memcpy(cte->cm.parameter,
296              (incoming)+(strlen(PROTOCOL_QUERYCSAC_SHORT)*(sizeof(char))), length);
297         cte->cm.code = CODE_QUERYCSAC;
298     }
299
300     /* PRINT_LOADKRLDATA */
301     else if(strstr((char*)incoming, PROTOCOL_LOADKRLDATA ) == (incoming)) {
302         int length = (strlen(incoming) - strlen(PROTOCOL_LOADKRLDATA) );
303         memcpy(cte->cm.parameter,
304              (incoming)+(strlen(PROTOCOL_LOADKRLDATA)*(sizeof(char))), length);
305         cte->cm.code = CODE_LOADKRLDATA;
306     }
307
308     /* PRINT_LOADKRLDATA_SHORT */

```



```

309     else if(strstr((char*)incoming, PROTOCOL_LOADKRLDATA_SHORT ) == (incoming)) {
310         int length = (strlen(incoming) - strlen(PROTOCOL_LOADKRLDATA_SHORT) );
311         memcpy(cte->cm.parameter,
312                (incoming)+(strlen(PROTOCOL_LOADKRLDATA_SHORT)*(sizeof(char))), length);
313         cte->cm.code = CODE_LOADKRLDATA;
314     }
315
316     /* PROTOCOL_PRINTCFD */
317     else if(strstr((char*)incoming, PROTOCOL_PRINTCFD ) == (incoming)) {
318         int length = (strlen(incoming) - strlen(PROTOCOL_PRINTCFD) );
319         memcpy(cte->cm.parameter, (incoming)+(strlen(PROTOCOL_PRINTCFD)*(sizeof(char))),
320                length);
321         cte->cm.code = CODE_PRINTCFD;
322         printf("PRINTCFD \n");
323     }
324
325     /* PROTOCOL_PRINTCFD_SHORT */
326     else if(strstr((char*)incoming, PROTOCOL_PRINTCFD_SHORT ) == (incoming)) {
327         int length = (strlen(incoming) - strlen(PROTOCOL_PRINTCFD_SHORT) );
328         memcpy(cte->cm.parameter,
329                (incoming)+(strlen(PROTOCOL_PRINTCFD_SHORT)*(sizeof(char))), length);
330         cte->cm.code = CODE_PRINTCFD;
331     }
332
333     else {
334         return 0;
335     }
336
337     /* Attempting to retrieve ID */
338     sscanf(cte->cm.parameter, "%d", &cte->cm.id_parameter);
339
340     return 1;
341 }
342
343 /* Responds to client action */
344 int respond(struct client_table_entry *cte)
345 {
346     bzero(cte->cm.parameter, MAX_PARAMETER_SIZE);
347     /* Only print ">" if client is monitor */
348     if(cte->client_id < 0) {
349         s_write(&(cte->transmission), ">", 1);
350     }
351
352     int read_status = s_read(&(cte->transmission)); /* Blocking */
353     if(read_status == -1) {
354         t_print("[ CLIENT %d ] Read failed or interrupted!\n", cte->client_id);
355         return 0;
356     }
357
358     if(cte->marked_for_kick) {
359         return 0;
360     }
361
362     int parse_status = parse_input(cte);
363
364     if(parse_status == -1) {
365         s_write(&(cte->transmission), ERROR_ILLEGAL_MESSAGE_SIZE,
366                sizeof(ERROR_ILLEGAL_MESSAGE_SIZE));
367     } else if(parse_status == 0) {
368         s_write(&(cte->transmission), ERROR_ILLEGAL_COMMAND,
369                sizeof(ERROR_ILLEGAL_COMMAND));
370     }

```

```

371     /* PARSING OK, CONTINUING */
372     else {
373         /* Comparing CODES to determine the correct action */
374         if(cte->cm.code == CODE_DISCONNECT) {
375             s_write(&(cte->transmission), PROTOCOL_GOODBYE, sizeof(PROTOCOL_GOODBYE));
376             return 0;
377         }
378
379         else if(cte->cm.code == CODE_HELP) {
380             print_help(cte);
381         }
382
383         else if(cte->cm.code == CODE_IDENTIFY) {
384             if(cte->cm.id_parameter == 0) {
385                 s_write(&(cte->transmission), ERROR_ILLEGAL_COMMAND,
386                     sizeof(ERROR_ILLEGAL_COMMAND));
387                 return 0;
388             }
389
390             /* Checking to see if the ID is in use */
391             struct client_table_entry* client_list_iterate;
392             list_for_each_entry(client_list_iterate, &client_list->list, list) {
393                 if(client_list_iterate->client_id == cte->cm.id_parameter) {
394                     cte->client_id = 0;
395                     s_write(&(cte->transmission), "ID in use!\n", 11);
396                     if(cte->cm.id_parameter < 0){
397                         return 1;
398                     }
399                     return 0;
400                 }
401             }
402
403             /* Determining role */
404             if(cte->cm.id_parameter < 0) {
405                 cte->client_type = MONITOR;
406                 struct timeval timeout = {MONITOR_TIMEOUT, 0};
407                 set_timeout(cte, timeout);
408
409             } else {
410                 cte->client_type = SENSOR;
411                 sem_wait(&(s_synch->client_list_sem));
412                 s_data->number_of_sensors++;
413                 sem_post(&(s_synch->client_list_sem));
414             }
415             /* Everything is good, setting id and responding*/
416             s_write(&(cte->transmission), PROTOCOL_OK, sizeof(PROTOCOL_OK));
417             cte->client_id = cte->cm.id_parameter;
418             t_print("[%s] ID set to: %d\n", cte->ip,cte->client_id);
419
420             if(cte->client_type == SENSOR) {
421                 if(load_lsf_data(cte)) {
422                     t_print("Loaded filter data for client %d\n", cte->client_id);
423                 } else {
424                     t_print("Failed to load LS filter data for Sensor %d\n", cte->client_id);
425                     s_write(&(cte->transmission),ERROR_LSD_LOAD_FAILED,
426                         sizeof(ERROR_LSD_LOAD_FAILED));
427                     return 0;
428                 }
429             }
430
431             return 1;
432         }

```

```

433
434     /* Stop here if client is unidentified */
435     else if(cte->client_id == 0) {
436         s_write(&(cte->transmission), ERROR_NO_ID, sizeof(ERROR_NO_ID));
437         return 1;
438     }
439
440     else if(cte->cm.code == CODE_NMEA) {
441         /* Fetching data from buffer */
442         char *rmc_start = strstr(cte->transmission.iobuffer, RMC);
443         char *gga_start = strstr(cte->transmission.iobuffer, GGA);
444
445         if(rmc_start == NULL || gga_start == NULL){
446             s_write(&(cte->transmission), ERROR_ILLEGAL_NMEA, strlen(ERROR_ILLEGAL_NMEA));
447             return 1;
448         }
449
450         memcpy(cte->nmea.raw_rmc, rmc_start, gga_start - rmc_start);
451         memcpy(cte->nmea.raw_gga, gga_start,
452             ( strlen(cte->transmission.iobuffer) - (rmc_start - cte->transmission.iobuffer)
453             - (gga_start - rmc_start)));
454
455         /* Checking NMEA checksum */
456         int rmc_checksum = calculate_nmea_checksum(cte->nmea.raw_rmc);
457         int gga_checksum = calculate_nmea_checksum(cte->nmea.raw_gga);
458
459         /* Continue to filters if ok */
460         if(rmc_checksum && gga_checksum) {
461             s_write(&(cte->transmission), PROTOCOL_OK, strlen(PROTOCOL_OK));
462             cte->timestamp = time(NULL);
463             cte->nmea.checksum_passed = 1;
464
465             if(!extract_nmea_data(cte)){
466                 return 1;
467             }
468
469             calculate_nmea_average(cte);
470             calculate_nmea_diff(cte);
471
472             /* Checksums where OK, client marked ready */
473             cte->ready = 1;
474
475             /* Acquiring ready-lock */
476             sem_wait(&(s_synch->ready_sem));
477
478             /* Checking if the other clients are ready as well*/
479             int ready = nmea_ready();
480
481             /* If everyone is ready, process data */
482             if(ready) {
483                 apply_filters();
484             }
485             /* Releasing ready-lock */
486             sem_post(&(s_synch->ready_sem));
487         } else {
488             cte->nmea.checksum_passed = 0;
489             t_print("RMC and GGA received from %d , checksum failed!\n", cte->client_id);
490             s_write(&(cte->transmission), ERROR_CHECKSUM_FAILED, strlen(ERROR_CHECKSUM_FAILED));
491         }
492     }
493
494     else if(cte->cm.code == CODE_PRINT_LOCATION) {

```

```

495     struct client_table_entry* candidate = get_client_by_id(cte->cm.id_parameter);
496     if(candidate == NULL) {
497         s_write(&(cte->transmission), ERROR_NO_CLIENT, sizeof(ERROR_NO_CLIENT));
498     } else {
499         print_location(cte, candidate);
500     }
501 }
502
503 else if(cte->cm.code == CODE_LOADKRLDATA) {
504     struct client_table_entry* candidate = get_client_by_id(cte->cm.id_parameter);
505     if(candidate == NULL) {
506         s_write(&(cte->transmission), ERROR_NO_CLIENT, sizeof(ERROR_NO_CLIENT));
507     } else {
508         if(load_lsf_data(candidate)) {
509             s_write(&(cte->transmission), PROTOCOL_OK, sizeof(PROTOCOL_OK));
510         } else {
511             s_write(&(cte->transmission), ERROR_LSD_LOAD_FAILED,
512                 sizeof(ERROR_LSD_LOAD_FAILED));
513         }
514     }
515 }
516
517 else if(cte->cm.code == CODE_PRINTCLIENTS) {
518     print_clients(cte);
519 }
520
521 else if(cte->cm.code == CODE_PRINTSERVER) {
522     print_server_data(cte);
523 }
524
525 else if(cte->cm.code == CODE_PRINTTIME) {
526     struct client_table_entry* candidate = get_client_by_id(cte->cm.id_parameter);
527     if(candidate != NULL) {
528         print_client_time(cte, candidate);
529     } else {
530         s_write(&(cte->transmission), ERROR_NO_CLIENT, sizeof(ERROR_NO_CLIENT));
531     }
532 }
533
534 else if(cte->cm.code == CODE_KICK) {
535     struct client_table_entry* candidate = get_client_by_id(cte->cm.id_parameter);
536     if(candidate == NULL) {
537         s_write(&(cte->transmission), ERROR_NO_CLIENT, sizeof(ERROR_NO_CLIENT));
538     } else {
539         kick_client(candidate);
540     }
541 }
542
543 else if(cte->cm.code == CODE_DUMPDATA) {
544     int filename_buffer_size = MAX_FILENAME_SIZE;
545     char filename[filename_buffer_size];
546     int target_id;
547     char id_buffer[ID_AS_STRING_MAX];
548     bzero(id_buffer, ID_AS_STRING_MAX);
549     bzero(filename, filename_buffer_size);
550
551     /* Attempting to extract filename */
552     substring_extractor(2,3, ' ', filename, filename_buffer_size, cte->cm.parameter,
553         MAX_FILENAME_SIZE);
554
555     /* If length of filename = 0 (no filename specified).. */
556     if(strlen(filename) == 0) {

```

```

557         /* ...Cast to int without a care */
558         target_id = atoi(cte->cm.parameter);
559     }
560     /* Else, extract ID */
561     else {
562         substring_extractor(1,2, ' ', id_buffer, ID_AS_STRING_MAX,cte->cm.parameter,
563                             ID_AS_STRING_MAX);
564         target_id = atoi(id_buffer);
565     }
566
567     if(!target_id) {
568         s_write(&(cte->transmission), ERROR_ILLEGAL_COMMAND,
569               sizeof(ERROR_ILLEGAL_COMMAND));
570     } else {
571         struct client_table_entry* candidate = get_client_by_id(target_id);
572         if(candidate != NULL) {
573             if(!datadump(candidate,filename, s_conf->human_readable_dumpdata)) {
574                 s_write(&(cte->transmission), ERROR_DUMPDATA_FAILED,
575                       sizeof(ERROR_DUMPDATA_FAILED));
576             }
577         } else {
578             s_write(&(cte->transmission), ERROR_NO_CLIENT, sizeof(ERROR_NO_CLIENT));
579         }
580     }
581 }
582
583 else if(cte->cm.code == CODE_LOADDATA) {
584     int filename_buffer_size = MAX_FILENAME_SIZE;
585     char filename[filename_buffer_size];
586     int target_id;
587     char id_buffer[ID_AS_STRING_MAX];
588     bzero(id_buffer, ID_AS_STRING_MAX);
589     bzero(filename, filename_buffer_size);
590
591     substring_extractor(2,3, ' ', filename, filename_buffer_size,cte->cm.parameter,
592                       MAX_FILENAME_SIZE);
593
594     /* No filename specified, abort */
595     if(strlen(filename) == 0) {
596         s_write(&(cte->transmission), ERROR_NO_FILENAME, sizeof(ERROR_NO_FILENAME));
597         return 1;
598     }
599     /* Extract target id and move on */
600     else {
601         substring_extractor(1,2, ' ', id_buffer, ID_AS_STRING_MAX,cte->cm.parameter,
602                             ID_AS_STRING_MAX);
603         target_id = atoi(id_buffer);
604     }
605
606     if(!target_id) {
607         s_write(&(cte->transmission), ERROR_ILLEGAL_COMMAND,
608               sizeof(ERROR_ILLEGAL_COMMAND));
609     } else {
610         struct client_table_entry* candidate = get_client_by_id(target_id);
611         if(candidate != NULL) {
612             int load_status = loaddata(candidate,filename);
613             if(load_status == ERROR_CODE_NO_FILE) {
614                 s_write(&(cte->transmission), ERROR_NO_FILE, sizeof(ERROR_NO_FILE));
615             } else if(load_status == ERROR_CODE_READ_FAILED) {
616                 s_write(&(cte->transmission), ERROR_READ_FAILED, sizeof(ERROR_READ_FAILED));
617             }
618         } else {

```

```

619         s_write(&(cte->transmission), ERROR_NO_CLIENT, sizeof(ERROR_NO_CLIENT));
620     }
621 }
622 }
623
624     else if(cte->cm.code == CODE_PRINTAVGDIFF) {
625         print_avg_diff(cte);
626     }
627
628     else if(cte->cm.code == CODE_LISTDUMPS) {
629         listdumps(cte);
630     }
631
632     else if(cte->cm.code == CODE_QUERYCSAC) {
633         if(strlen(cte->cm.parameter) < 3) {
634             s_write(&(cte->transmission), ERROR_NO_COMMAND, sizeof(ERROR_NO_COMMAND));
635             return 1;
636         }
637         client_query_csac(cte, cte->cm.parameter);
638     } else if(cte->cm.code == CODE_PRINTCFD) {
639         print_cfd(cte, cte->cm.id_parameter);
640     }
641
642     else {
643         t_print("No action made for this part of the protocol\n");
644     }
645 }
646 return 1;
647 }

```

session.h

```

1  /**
2  * @file session.h
3  * @author Aril Schultzen
4  * @date 13.04.2016
5  * @brief File containing function prototypes and includes for session.c
6  */
7
8  #ifndef SESSION_H
9  #define SESSION_H
10
11 #include "sensor_server_common.h"
12 #include "filters.h"
13 #include "actions.h"
14 #include "sensor_server.h"
15
16 int respond(struct client_table_entry *cte);
17
18 #endif /* !SESSION_H */

```

actions.c

```

1  #include "actions.h"
2
3  /* GENERAL */
4  #define CLIENT_TABLE_LABEL "CLIENT TABLE\n"
5  #define NEW_LINE "\n"

```

```

6  #define PRINT_LOCATION_HEADER "      CURRENT      MIN      MAX      AVG\n"
7  #define DUMPDATA_HEADER "CURRENT      MIN      MAX      AVERAGE      AVG_DIFF      TOTAL      DISTURBED\n"
8  #define PRINT_AVG_DIFF_HEADER "ID      LAT      LON      ALT      SPEED\n"
9  #define DATADUMP_EXTENSION ".bin"
10 #define DATADUMP_HUMAN_EXTENSION ".txt"
11 #define CSAC_SCRIPT_COMMAND "python query_csac.py "
12
13 /* ERRORS */
14 #define ERROR_APPEND_TOO_LONG "ERROR: TEXT TO APPEND TOO LONG\n"
15 #define ERROR_NO_SENSORS_CONNECTED "NO SENSORS CONNECTED\n"
16 #define ERROR_FCLOSE "Failed to close file, out of space?\n"
17 #define ERROR_FWRITE "Failed to write to file, aborting.\n"
18 #define ERROR_FREAD "Failed to read file, aborting.\n"
19 #define ERROR_FOPEN "Failed to open file, aborting.\n"
20 #define ERROR_UPDATE_WARMUP_ILLEGAL "Warm-up time value has to be greater than 0!\n"
21 #define ERROR_CSAC_FAILED "Communication with CSAC failed!\n"
22
23 /* LOAD_REF_DEV_DATA */
24 #define KRL_FILENAME "ls_data_sensor"
25 #define ALT_REF "alt_ref:"
26 #define LON_REF "lon_ref:"
27 #define LAT_REF "lat_ref:"
28 #define SPEED_REF "speed_ref:"
29 #define ALT_DEV "alt_dev:"
30 #define LON_DEV "lon_dev:"
31 #define LAT_DEV "lat_dev:"
32 #define SPEED_DEV "speed_dev:"
33 #define KRL_DATA_ENTRIES 8
34
35 /* HELP */
36 #define HELP "\n\
37 " COMMAND      | SHORT | PARAM      | DESCRIPTION\n\
38 "-----\n\
39 " HELP          | ?     | NONE      | Prints this table\n\
40 "-----\n\
41 " IDENTIFY      | ID    | ID        | Your ID is set to PARAM ID\n\
42 "-----\n\
43 " DISCONNECT    | EXIT  | NONE      | Disconnects\n\
44 "-----\n\
45 " PRINTCLIENTS | PC    | NONE      | Prints a list of connected clients\n\
46 "-----\n\
47 " PRINTSERVER   | PS    | NONE      | Prints server state and config\n\
48 "-----\n\
49 " PRINTTIME     |      | ID        | Prints time solved from Sensor <ID>\n\
50 "-----\n\
51 " PRINTAVGDIFF | PAD   | NONE      | Prints all average diffs for all clients\n\
52 "-----\n\
53 " PRINTLOC      | PL    | ID        | Prints solved location for Sensor <ID>\n\
54 "-----\n\
55 " LISTDATA      | LSD   | NONE      | Lists all dump files in server directory\n\
56 "-----\n\
57 " DUMPDATA      | DD    | ID & FILE | Dumps state of Sensor <ID> into FILE\n\
58 "-----\n\
59 " LOADDATA      | LD    | ID & FILE | Loads NMEA of FILE into sensor ID\n\
60 "-----\n\
61 " QUERYCSAC     | QC    | COMMAND   | Queries the CSAC with parameter COMMAND\n\
62 "-----\n\
63 " LOADLSFDATA   | LLSFD | ID        | Load reference location data into Sensor<ID>\n\
64 "-----\n\
65 " PRINTCFD      | PFD   |           | Prints CSAC filter data\n\
66 "-----\n\
67

```

```

68  /* SIZES */
69  #define DUMPDATA_TIME_SIZE 13
70  #define MAX_APPEND_LENGTH 20
71
72  void kick_client(struct client_table_entry* client)
73  {
74      sem_wait(&(s_synch->client_list_sem));
75      sem_wait(&(s_synch->ready_sem));
76      client->marked_for_kick = 1;
77      sem_post(&(s_synch->ready_sem));
78      sem_post(&(s_synch->client_list_sem));
79  }
80
81  /* Prints client X's solved time back to monitor */
82  void print_client_time(struct client_table_entry *monitor,
83                       struct client_table_entry* client)
84  {
85      int buffsize = 100;
86      char buffer[buffsize];
87      memset(&buffer, 0, buffsize);
88
89      substring_extractor(RMC_TIME_START,RMC_TIME_START + 1,',',buffer, buffsize,
90                        client->nmea.raw_rmc, strlen(client->nmea.raw_rmc));
91      s_write(&(monitor->transmission), buffer, 12);
92      s_write(&(monitor->transmission), "\n", 1);
93  }
94
95  /* Prints a formatted string containing info about connected clients to monitor */
96  void print_clients(struct client_table_entry *monitor)
97  {
98      char buffer [1000];
99      int snprintf_status = 0;
100     char *c_type = "SENSOR";
101     char *modifier = "";
102
103     struct client_table_entry* client_list_iterate;
104     s_write(&(monitor->transmission), CLIENT_TABLE_LABEL,
105           sizeof(CLIENT_TABLE_LABEL));
106     s_write(&(monitor->transmission), HORIZONTAL_BAR, sizeof(HORIZONTAL_BAR));
107     list_for_each_entry(client_list_iterate,&client_list->list, list) {
108
109         if(client_list_iterate->client_type == MONITOR) {
110             c_type = "MONITOR";
111         } else {
112             c_type = "SENSOR";
113         }
114
115         if(monitor->client_id == client_list_iterate->client_id) {
116             modifier = BOLD_GRN_BLK;
117         } else {
118             modifier = RESET;
119         }
120         snprintf_status = snprintf( buffer, 1000,
121                                   "%sID: %d " \
122                                   "IP:%s, " \
123                                   "PID: %d, " \
124                                   "TYPE: %s, " \
125                                   "NMEA age %d%s\n",
126                                   modifier,
127                                   client_list_iterate->client_id,
128                                   client_list_iterate->ip,
129                                   client_list_iterate->pid,

```



```

130 |                                     c_type,
131 |                                     (int)difftime(time(NULL),client_list_iterate->timestamp),
132 |                                     RESET);
133 |
134 |         s_write(&(monitor->transmission), buffer, snprintf_status);
135 |     }
136 |     s_write(&(monitor->transmission), HORIZONTAL_BAR, sizeof(HORIZONTAL_BAR));
137 | }
138 |
139 | /*
140 | * Prints a string containing simple description
141 | * of the different implemented commands back
142 | * to the monitor.
143 | */
144 | void print_help(struct client_table_entry *monitor)
145 | {
146 |     s_write(&(monitor->transmission), HELP, sizeof(HELP));
147 | }
148 |
149 | /*
150 | * Prints MAX, MIN, CURRENT and AVERAGE position
151 | * for client X back to the monitor
152 | */
153 | void print_location(struct client_table_entry *monitor,
154 |                    struct client_table_entry* client)
155 | {
156 |     char buffer [1000];
157 |     int snprintf_status = 0;
158 |
159 |     char *lat_modifier;
160 |     char *lon_modifier;
161 |     char *alt_modifier;
162 |     char *speed_modifier;
163 |     char *reset = RESET;
164 |
165 |     struct nmea_container nc;
166 |
167 |     nc = client->nmea;
168 |     s_write(&(monitor->transmission), PRINT_LOCATION_HEADER,
169 |            sizeof(PRINT_LOCATION_HEADER));
170 |
171 |     /*Determining colors*/
172 |     if(!nc.lat_disturbed) {
173 |         lat_modifier = BOLD_GRN_BLK;
174 |     } else if(nc.lat_disturbed > 0) {
175 |         lat_modifier = BOLD_RED_BLK;
176 |     } else {
177 |         lat_modifier = BOLD_CYN_BLK;
178 |     }
179 |
180 |     if(!nc.lon_disturbed) {
181 |         lon_modifier = BOLD_GRN_BLK;
182 |     } else if(nc.lon_disturbed > 0) {
183 |         lon_modifier = BOLD_RED_BLK;
184 |     } else {
185 |         lon_modifier = BOLD_CYN_BLK;
186 |     }
187 |
188 |     if(!nc.alt_disturbed) {
189 |         alt_modifier = BOLD_GRN_BLK;
190 |     } else if(nc.alt_disturbed > 0) {
191 |         alt_modifier = BOLD_RED_BLK;

```

```

192     } else {
193         alt_modifier = BOLD_CYN_BLK;
194     }
195
196     if(!nc.speed_disturbed) {
197         speed_modifier = BOLD_GRN_BLK;
198     } else if(nc.speed_disturbed > 0) {
199         speed_modifier = BOLD_RED_BLK;
200     } else {
201         speed_modifier = BOLD_CYN_BLK;
202     }
203
204     snprintf_status = snprintf( buffer, 1000,
205                               "LAT: %s%f%s %f\n" \
206                               "LON: %s%f%s %f\n" \
207                               "ALT: %s %f%s %f\n" \
208                               "SPD: %s %f%s %f\n",
209                               lat_modifier, nc.lat_current,reset, nc.lat_average,
210                               lon_modifier, nc.lon_current,reset, nc.lon_average,
211                               alt_modifier, nc.alt_current,reset, nc.alt_average,
212                               speed_modifier, nc.speed_current,reset, nc.speed_average);
213     s_write(&(monitor->transmission), buffer, snprintf_status);
214 }
215
216 /*
217  * Prints the difference between the calculated
218  * average values for location and the current value
219  */
220 void print_avg_diff(struct client_table_entry *client)
221 {
222     char buffer [1000];
223     int snprintf_status = 0;
224     struct nmea_container nc;
225
226     if(s_data->number_of_sensors > 0) {
227         s_write(&(client->transmission), PRINT_AVG_DIFF_HEADER,
228               sizeof(PRINT_AVG_DIFF_HEADER));
229         struct client_table_entry* client_list_iterate;
230         list_for_each_entry(client_list_iterate,&client_list->list, list) {
231             if(client_list_iterate->client_id > 0) {
232                 nc = client_list_iterate->nmea;
233                 snprintf_status = snprintf( buffer, 1000, "%d %f %f %f %f\n",
234                                             client_list_iterate->client_id, nc.lat_avg_diff, nc.lon_avg_diff,
235                                             nc.alt_avg_diff, nc.speed_avg_diff);
236                 s_write(&(client->transmission), buffer, snprintf_status);
237             }
238         }
239     } else {
240         s_write(&(client->transmission), ERROR_NO_SENSORS_CONNECTED,
241               sizeof(ERROR_NO_SENSORS_CONNECTED));
242     }
243 }
244
245 static int get_pfd_string(char *buffer, int buf_len)
246 {
247     memset(buffer, '\0',buf_len);
248     int snprintf_status = snprintf( buffer, 1000,
249                                   "Phase:                %lf\n\n" \
250                                   "T current:         %lf\n" \
251                                   "T current (smooth): %lf\n" \
252                                   "T previous (smooth): %lf\n" \
253                                   "T today (smooth):  %lf\n" \

```

```

254 |         "T yesterday (smooth):      %lf\n\n" \
255 |         "Steer current:             %lf\n" \
256 |         "Steer current (smooth):    %lf\n" \
257 |         "Steer previous (smooth):   %lf\n\n" \
258 |         "Steer today (smooth):      %lf\n" \
259 |         "Steer yesterday (smooth):  %lf\n\n" \
260 |         "Steer prediction:          %lf\n\n" \
261 |         "MJD today:                 %lf\n" \
262 |         "Days passed since startup: %d\n\n" \
263 |         "Discipline status:         %d\n" \
264 |         "Fast timing filter status  %d\n" \
265 |         "Freq corr. filter status  %d\n\n",
266 |         cfd->phase_current,
267 |         cfd->t_current,
268 |         cfd->t_smooth_current,
269 |         cfd->t_smooth_previous,
270 |         cfd->t_smooth_today,
271 |         cfd->t_smooth_yesterday,
272 |         cfd->steer_current,
273 |         cfd->steer_smooth_current,
274 |         cfd->steer_smooth_previous,
275 |         cfd->steer_smooth_today,
276 |         cfd->steer_smooth_yesterday,
277 |         cfd->steer_prediction,
278 |         cfd->today_mjd,
279 |         cfd->days_passed,
280 |         cfd->discok,
281 |         cfd->ftf_status,
282 |         cfd->fqf_status);
283 |     return snprintf_status;
284 | }
285 |
286 | void print_cfd(struct client_table_entry *monitor, int update_count)
287 | {
288 |     int buf_len = 1000;
289 |     char buffer [buf_len];
290 |     int counter = 0;
291 |
292 |     if(update_count == 0) {
293 |         update_count = 1;
294 |     }
295 |
296 |     while(counter < update_count) {
297 |         get_pfd_string(buffer, buf_len);
298 |         s_write(&(monitor->transmission), buffer, strlen(buffer));
299 |         counter++;
300 |         sleep(1);
301 |     }
302 | }
303 |
304 | int dump_cfd(char *path)
305 | {
306 |     int buf_len = 1000;
307 |     char buffer[buf_len];
308 |
309 |     /* Formatting string with CSAC filter data */
310 |     get_pfd_string(buffer, buf_len);
311 |
312 |     /* Opening and writing to file */
313 |     FILE *cfd_file;
314 |     cfd_file = fopen(path, "w+");
315 |

```

```

316     if(!cfd_file) {
317         t_print("dump_cfd: %s: %s",ERROR_FOPEN, path);
318         return 0;
319     }
320
321     if(!fprintf(cfd_file,"%s", buffer) ) {
322         t_print(ERROR_FWRITE);
323         return 0;
324     }
325
326     if(fcclose(cfd_file)) {
327         t_print(ERROR_FCLOSE);
328     }
329     return 1;
330 }
331
332 /* Dumps data location data for client X into a file */
333 int datadump(struct client_table_entry* client, char *filename,
334             int dump_human_read)
335 {
336     FILE *bin_file;
337     char bin_name[strlen(filename) + strlen(DATADUMP_EXTENSION)];
338     strcpy(bin_name, filename);
339     strcat(bin_name, DATADUMP_EXTENSION);
340
341     bin_file=fopen(bin_name, "wb");
342
343     if(!bin_file) {
344         t_print(ERROR_FOPEN);
345         return 0;
346     }
347
348     if(!fwrite(&client->nmea, sizeof(struct nmea_container), 1, bin_file)) {
349         t_print(ERROR_FWRITE);
350         return 0;
351     }
352
353     if(fcclose(bin_file)) {
354         t_print(ERROR_FCLOSE);
355     }
356
357     if(dump_human_read) {
358         /* Dumping humanly readable data */
359         FILE *h_dump;
360         char h_name[strlen(filename) + strlen(DATADUMP_HUMAN_EXTENSION)];
361         strcpy(h_name, filename);
362         strcat(h_name, DATADUMP_HUMAN_EXTENSION);
363
364         h_dump = fopen(h_name, "wb");
365
366         fprintf(h_dump, "Sensor Server dumpfile created for client %d\n",
367                client->client_id);
368
369         /*
370          * Dumping all from NMEA container
371          * after raw_rmc and including speed_disturbed
372          */
373         int inner_counter = 0;
374         int outer_counter = 0;
375         double *data = &client->nmea.lat_current;
376
377         fprintf(h_dump,DUMPDATA_HEADER);

```

```

378     while(outer_counter < 4) {
379         while(inner_counter < 7) {
380             fprintf(h_dump, "%f ",*data);
381             data++;
382             inner_counter++;
383         }
384         fprintf(h_dump, "%f", *data);
385         inner_counter = 0;
386         outer_counter++;
387     }
388
389     /*
390     * Dumping ref_dev_data
391     */
392     fprintf(h_dump,DUMPDATA_HEADER);
393     inner_counter = 0;
394     double *rdf = &client->fs.ls_f.lsf_d.alt_ref;
395     while(inner_counter < 8) {
396         fprintf(h_dump, "%lf \n",*rdf);
397         rdf++;
398         inner_counter++;
399     }
400
401     if(fcclose(h_dump)) {
402         t_print(ERROR_FCLOSE);
403     }
404 }
405 return 1;
406 }
407
408 /* Print list of dumped data */
409 int listdumps(struct client_table_entry* monitor)
410 {
411     DIR *dp;
412     struct dirent *ep;
413
414     dp = opendir (".");
415     if(dp != NULL) {
416         while ( (ep = readdir(dp)) ) {
417             if(strstr(ep->d_name,DATA_DUMP_EXTENSION) != NULL) {
418                 s_write(&(monitor->transmission),ep->d_name, strlen(ep->d_name));
419                 s_write(&(monitor->transmission),NEW_LINE, sizeof(NEW_LINE));
420             }
421         }
422         closedir (dp);
423     } else {
424         return 0;
425     }
426
427     return 1;
428 }
429
430 /* Load dumped data into the client */
431 int loaddata(struct client_table_entry* target, char *filename)
432 {
433     FILE *dump_file;
434     int file_len = 0;
435
436
437     dump_file=fopen(filename, "rb");
438
439     if(!dump_file) {

```

```

440         t_print(ERROR_FOPEN);
441         return ERROR_CODE_NO_FILE;
442     }
443
444     /* Checking file length */
445     fseek(dump_file, 0, SEEK_END);
446     file_len=ftell(dump_file);
447     fseek(dump_file, 0, SEEK_SET);
448
449     int f_s = fread( &target->nmea,1,sizeof(struct nmea_container), dump_file);
450
451     t_print("Read %d/%d bytes successfully from %s\n", f_s, file_len,filename);
452
453     if(f_s == 0) {
454         t_print(ERROR_FREAD);
455         return ERROR_CODE_READ_FAILED;
456     }
457
458     if(fclose(dump_file)) {
459         t_print(ERROR_FCLOSE);
460     }
461
462     return 1;
463 }
464
465 int query_csac(char *query, char *buffer)
466 {
467     /* Building command */
468     int command_size = MAX_PARAMETER_SIZE + sizeof(CSAC_SCRIPT_COMMAND);
469     char command[command_size];
470     memset(command, '\0', command_size);
471     strcat(command, CSAC_SCRIPT_COMMAND);
472     strcat(command, query);
473
474     /* Acquiring lock*/
475     sem_wait(&(s_synch->csac_sem));
476
477     /* Running command */
478     if(!run_command(command, buffer)) {
479         /* Releasing lock */
480         sem_post(&(s_synch->csac_sem));
481         return 0;
482     }
483
484     /* Releasing lock */
485     sem_post(&(s_synch->csac_sem));
486     return 1;
487 }
488
489
490 int client_query_csac(struct client_table_entry *monitor, char *query)
491 {
492     char buffer[MAX_PARAMETER_SIZE];
493     memset(buffer, '\0', MAX_PARAMETER_SIZE);
494
495     if(!query_csac(query, buffer)) {
496         return 0;
497     }
498
499     if(!s_write(&(monitor->transmission), buffer, strlen(buffer))) {
500         return 0;
501     }

```

```

502     return 1;
503 }
504
505 /*
506  * Load ls filter data into the client struct.
507  * Re-using the config loader.
508  * This whole function needs some work! Magic numbers beware.
509  */
510 int load_lsf_data(struct client_table_entry* target)
511 {
512     /* Request lock */
513     sem_wait(&(s_synch->client_list_sem));
514     sem_wait(&(s_synch->ready_sem));
515     struct config_map_entry conf_map[KRL_DATA_ENTRIES];
516
517     int filename_length = strlen(KRL_FILENAME) + 10;
518     char filename[filename_length];
519     memset(filename, '\0', filename_length);
520     strcpy(filename, KRL_FILENAME);
521
522     /* Way overkill for int to string, but still. */
523     char id[10];
524     memset(id, '\0', 10);
525     sprintf(id, "%d", target->client_id);
526     strcat(filename, id);
527
528     conf_map[0].entry_name = ALT_REF;
529     conf_map[0].modifier = FORMAT_DOUBLE;
530     conf_map[0].destination = &target->fs.ls_f.lsf_d.alt_ref;
531
532     conf_map[1].entry_name = LON_REF;
533     conf_map[1].modifier = FORMAT_DOUBLE;
534     conf_map[1].destination = &target->fs.ls_f.lsf_d.lon_ref;
535
536     conf_map[2].entry_name = LAT_REF;
537     conf_map[2].modifier = FORMAT_DOUBLE;
538     conf_map[2].destination = &target->fs.ls_f.lsf_d.lat_ref;
539
540     conf_map[3].entry_name = SPEED_REF;
541     conf_map[3].modifier = FORMAT_DOUBLE;
542     conf_map[3].destination = &target->fs.ls_f.lsf_d.speed_ref;
543
544     conf_map[4].entry_name = ALT_DEV;
545     conf_map[4].modifier = FORMAT_DOUBLE;
546     conf_map[4].destination = &target->fs.ls_f.lsf_d.alt_dev;
547
548     conf_map[5].entry_name = LON_DEV;
549     conf_map[5].modifier = FORMAT_DOUBLE;
550     conf_map[5].destination = &target->fs.ls_f.lsf_d.lon_dev;
551
552     conf_map[6].entry_name = LAT_DEV;
553     conf_map[6].modifier = FORMAT_DOUBLE;
554     conf_map[6].destination = &target->fs.ls_f.lsf_d.lat_dev;
555
556     conf_map[7].entry_name = SPEED_DEV;
557     conf_map[7].modifier = FORMAT_DOUBLE;
558     conf_map[7].destination = &target->fs.ls_f.lsf_d.speed_dev;
559
560     int load_config_status = load_config(conf_map, filename,
561                                         KRL_DATA_ENTRIES);
562
563     /* releasing lock */

```

```

564     sem_post(&(s_synch->ready_sem));
565     sem_post(&(s_synch->client_list_sem));
566     return load_config_status;
567 }

```

actions.h

```

1  /**
2  * @file actions.h
3  * @brief File containing function prototypes and includes for actions.c
4  *
5  * Function prototypes for functions that implements different
6  * actions that a MONITOR or the system can use to manipulate the
7  * state of the SENSORS or print stats or similar.
8  *
9  * Be advised that any reference to MONITOR in this file means
10 * a client connected to the server who's role is that of a
11 * monitor of the system and not a monitor like a peripheral
12 * connected to a computer. The names of these roles are under
13 * discussion and will probably be changed to avoid misunderstanding.
14 *
15 * @author Aril Schultzen
16 * @date 9.11.2015
17 */
18
19 #ifndef ACTIONS_H
20 #define ACTIONS_H
21
22 #include "sensor_server.h"
23 #include "serial.h"
24 #include <dirent.h>
25
26 /** @brief Kicks a client (both MONITOR or SENSOR)
27 *
28 * Marks the client so respond() in session.c can
29 * disconnect it the next time that client transmits
30 * data. The kick is in other words not instant, this
31 * is however an easy way to gracefully disconnect a
32 * client.
33 *
34 * @param client Pointer to the client_table_entry for the candidate to be kicked.
35 * @return Void
36 */
37 void kick_client(struct client_table_entry* client);
38
39 /** @brief Prints clients solved time to MONITOR
40 *
41 * Extracts the time solved by the GPS receiver, transmitted
42 * via NMEA and stored in the client_table_struct at the server,
43 * and transmits it to the MONITOR that requested it.
44 *
45 * @param monitor Pointer to MONITOR who made the request.
46 * @param client Pointer to SENSOR whose time was requested.
47 * @return Void
48 */
49 void print_client_time(struct client_table_entry *monitor,
50                       struct client_table_entry* client);
51
52 /** @brief Prints a table of clients to the MONITOR
53 *

```



```

54  * Transmits a table of the connected clients to the MONITOR.
55  *
56  * @param monitor Pointer to MONITOR who made the request.
57  * @return Void
58  */
59 void print_clients(struct client_table_entry *monitor);
60
61 /** @brief Prints table of available commands to requesting MONITOR.
62  *
63  * @param monitor Pointer to MONITOR who made the request.
64  * @return Void
65  */
66 void print_help(struct client_table_entry *monitor);
67
68 /** @brief Prints location of SENSOR to requesting MONITOR.
69  *
70  * Prints a overview of current as well as MIN, MAX and AVERAGE
71  * values of LAT, LON, ALT and SPEED recovered from NMEA.
72  *
73  * @param monitor Pointer to MONITOR who made the request.
74  * @param client Pointer to SENSOR whose location is requested.
75  * @return Void
76  */
77 void print_location(struct client_table_entry *monitor,
78                    struct client_table_entry* client);
79
80 /** @brief Prints difference between current position and average.
81  *
82  * Prints the difference between the current position values
83  * recorded from NMEA, and the calculated averages.
84  * Two sensors in close proximity (100m >) should be
85  * subjected to the same noise. If the difference between
86  * sensor A (current-avg) and sensor B (current-avg) changes,
87  * this could mean that one of them is being spoofed.
88  *
89  * @param monitor Pointer to MONITOR who made the request.
90  * @return Void
91  */
92 void print_avg_diff(struct client_table_entry *monitor);
93
94 /** @brief Dumps NMEA data to file for given client
95  *
96  * @param client Pointer to client whose data should be dumped.
97  * @param filename Pointer to filename.
98  * @param human_readable Switch to determine if humanly readable data should be made as well.
99  * @return 1 if success, 0 if fail.
100 */
101 int datadump(struct client_table_entry* client, char *filename,
102             int human_readable);
103
104 /** @brief List dump files in folder
105  *
106  * @param monitor Pointer to requesting monitor
107  * @return 1 if success, 0 if fail.
108  */
109 int listdumps(struct client_table_entry* monitor);
110
111 /** @brief Loads NMEA data into the NMEA struct of a given client (target).
112  *
113  * @param target Pointer to the client whose NMEA data should be loaded
114  * from file.
115  * @param filename Pointer to the filename of the data file.

```

```

116 */
117 int loaddata(struct client_table_entry* target, char *filename);
118
119 /** @brief Uses the query_csac.py to communicate with the CSAC.
120 * Stores the response in a buffer.
121 *
122 * @param buffer Buffer to store the response
123 * @param query Command (query) to send to the CSAC.
124 */
125 int query_csac(char *query, char *buffer);
126
127 /** @brief Uses the query_csac.py to communicate with the CSAC
128 * Prints the response from the CSAC back to the client
129 *
130 * @param monitor Monitor who made the request
131 * @param query Command (query) to send to the CSAC.
132 */
133 int client_query_csac(struct client_table_entry *monitor, char *query);
134
135 /** @brief Loads data for the REF_DEV_FILTER into the client.
136 *
137 * @param target Client to load the data into
138 */
139 int load_lsf_data(struct client_table_entry* target);
140
141 /** @brief Prints the current state of the CSAC filter.
142 *
143 * @param monitor Monitor to print the data to.
144 * @return Status of sprintf() used to build string.
145 */
146 void print_cfd(struct client_table_entry *monitor, int update_count);
147
148 /** @brief Dumps the state of the CSAC filter to file.
149 *
150 * @param Path to desired file to use.
151 * @return 1 if successful, 0 else.
152 */
153 int dump_cfd(char *path);
154 #endif /* !ACTIONS_H */

```

utils.c

```

1 #include "utils.h"
2
3 /* These are also in action.c, duplicates are no solution */
4 #define ERROR_FCLOSE "Failed to close file, out of space?\n"
5 #define ERROR_FWRITE "Failed to write to file, aborting.\n"
6 #define ERROR_FREAD "Failed to read file, aborting.\n"
7 #define ERROR_FOPEN "Failed to open file, aborting.\n"
8
9 #define MJD_SCRIPT_PATH "./get_mjd.py"
10
11 void die (int line_number, const char * format, ...)
12 {
13     va_list vargs;
14     va_start (vargs, format);
15     fprintf (stderr, "%d: ", line_number);
16     vfprintf (stderr, format, vargs);
17     fprintf (stderr, ".\n");
18     exit(1);

```

```

19 }
20
21 /*
22  * Extracts IP address from sockaddr struct.
23  * Supports both IPV4 and IPV6
24  */
25 void extract_ip_str(const struct sockaddr *sa, char *s, size_t maxlen)
26 {
27     switch(sa->sa_family) {
28     case AF_INET:
29         inet_ntop(AF_INET, &(((struct sockaddr_in *)sa)->sin_addr),
30                 s, maxlen);
31         break;
32
33     case AF_INET6:
34         inet_ntop(AF_INET6, &(((struct sockaddr_in6 *)sa)->sin6_addr),
35                 s, maxlen);
36         break;
37
38     default:
39         strncpy(s, "Unknown AF", maxlen);
40     }
41 }
42
43 /*
44  * Extracts IP from session file descriptor
45  */
46 void get_ip_str(int session_fd, char *ip)
47 {
48     struct sockaddr addr;
49     addr.sa_family = AF_INET;
50     socklen_t addr_len = sizeof(addr);
51     if(getpeername(session_fd, (struct sockaddr *) &addr, &addr_len)) {
52         die(44, "getsocketname failed\n");
53     }
54     extract_ip_str(&addr, ip, addr_len);
55 }
56
57 /*
58  * Print with timestamp:
59  * Example : [01.01.01 - 10:10:10] [<Some string>]
60  */
61 void t_print(const char* format, ...)
62 {
63     char buffer[100];
64     time_t rawtime;
65     struct tm *info;
66     time(&rawtime);
67     info = gmtime(&rawtime);
68     strftime(buffer, 80, "[%x - %X] ", info);
69     va_list argptr;
70     va_start(argptr, format);
71     fputs(buffer, stdout);
72     vfprintf(stdout, format, argptr);
73     va_end(argptr);
74 }
75
76 /*
77  * Loads config.
78  * Returns: 0 fail / 1 success
79  */
80 int load_config(struct config_map_entry *cme, char *path, int entries)

```

```

81  {
82      FILE *config_file;
83      int file_len;
84      char *input_buffer;
85
86      int status = 0;
87
88      config_file=fopen(path, "r");
89      if(!config_file) {
90          return 0;
91      }
92
93      fseek(config_file , OL , SEEK_END);
94      file_len = ftell(config_file);
95      rewind(config_file);
96
97      char temp_buffer[file_len];
98
99      /* Allocating memory for the file buffer */
100     input_buffer = calloc( file_len, sizeof(char));
101     if(!input_buffer) {
102         fclose(config_file);
103         t_print("config_loader(): Memory allocation failed, aborting.\n");
104         return 0;
105     }
106
107     /* Get the file into the buffer */
108     if(fread( input_buffer , file_len, 1 , config_file) != 1) {
109         fclose(config_file);
110         free(input_buffer);
111         t_print("config_loader(): Read failed, aborting\n");
112         return 0;
113     }
114
115     int counter = 0;
116     while(counter < entries) {
117         memset(temp_buffer, '\0',file_len);
118         char *search_ptr = strstr(input_buffer,cme->entry_name);
119         if(search_ptr != NULL) {
120             int length = strlen(search_ptr) - strlen(cme->entry_name);
121             memcpy(temp_buffer, search_ptr+(strlen(cme->entry_name)*(sizeof(char))),
122                 length);
123             status = sscanf(temp_buffer, cme->modifier, cme->destination);
124             if(status == EOF || status == 0) {
125                 fclose(config_file);
126                 free(input_buffer);
127                 return -1;
128             }
129         }
130         else{
131             return 0;
132         }
133         counter++;
134         cme++;
135     }
136
137     fclose(config_file);
138     free(input_buffer);
139     return 1;
140 }
141
142 int calculate_nmea_checksum(char *nmea)

```

```

143 {
144     char checksum = 0;
145     int i;
146     int received_checksum = 0;
147     int calculated_checksum = 0;
148
149
150     /* Substring to iterate over */
151     char substring[100] = {0};
152
153     /* Finding end (*) and calculate length */
154     char *substring_end = strstr(nmea, "*");
155     int length = substring_end - (nmea+1);
156
157     /* Copying the substring */
158     memcpy(substring, nmea+1, length);
159
160     /* Calculating checksum */
161     for(i = 0; i < length; i++) {
162         checksum = checksum ^ substring[i];
163     }
164
165     /* Reusing substring buffer */
166     sprintf(substring, "%x\n", checksum);
167
168     /* Converting calculated checksum to int */
169     sscanf(substring, "%d", &calculated_checksum);
170
171     /* Fetching received checksum */
172     memcpy(substring, substring_end+1, strlen(nmea));
173
174     /* Converting received checksum to int*/
175     sscanf(substring, "%d", &received_checksum);
176
177     /* Comparing checksum */
178     if(received_checksum == calculated_checksum) {
179         return 1;
180     } else {
181         return 0;
182     }
183 }
184
185
186 /*
187 * Used to extract words from between two delimiters
188 * delim_num_1 -> The number of the first delimiter, ex.3
189 * delim_num_2 -> The number of the second delimiter, ex.5
190 * delimiter -> The character to be used as a delimiter
191 * string -> Input
192 * buffer -> To transport the string
193 */
194 int substring_extractor(int start, int end, char delimiter, char *buffer,
195                       int buffsize, char *string, int str_len)
196 {
197     int i;
198     int delim_counter = 0;
199     int buffer_index = 0;
200
201     const int carriage_return = 13;
202
203     bzero(buffer, buffsize);
204

```

```

205     for(i = 0; i < str_len; i++) {
206         /* Second delim (end) reached, stopping. */
207         if(delim_counter == end || (int)string[i] == carriage_return) {
208             return 1;
209         }
210
211         if(string[i] == delimiter) {
212             delim_counter++;
213         } else {
214             /* The first delim is reached */
215             if(delim_counter >= start) {
216                 buffer[buffer_index] = string[i];
217                 buffer_index++;
218             }
219         }
220     }
221     /* Reached end of string without encountering end delimit */
222     return 0;
223 }
224
225 int str_len_u(char *buffer, int buf_len)
226 {
227     int i;
228     char prev = 'X';
229     for(i = 0; i < buf_len; i++) {
230         if(buffer[i] == 0x0a && prev == 0x0a) {
231             return i;
232         }
233         prev = buffer[i];
234     }
235     return -1;
236 }
237
238 /* Mega hackish code for getting MJD */
239 int get_today_mjd(char *buffer)
240 {
241     int status = run_command(MJD_SCRIPT_PATH, buffer);
242     /* Removing newline */
243     buffer[strcspn(buffer, "\n")] = 0;
244     return status;
245 }
246
247 int run_command(char *path, char *output)
248 {
249     FILE *fp;
250     int buffer_size = 1000;
251     char buffer[buffer_size];
252     memset(buffer, '\0', buffer_size);
253
254     /* Open the command for reading. */
255     fp = popen(path, "r");
256     if (fp == NULL) {
257         t_print("Failed to run command\n");
258         return 0;
259     }
260
261     /* Read the output a line at a time - output it. */
262     while (fgets(buffer, sizeof(buffer)-1, fp) != NULL) {
263         strcat(output,buffer);
264     }
265
266     /* close */

```

```

267     pclose(fp);
268     return strlen(output);
269 }
270
271 int log_to_file(char *path, char *content, int stamp_switch)
272 {
273     FILE *log_file;
274     log_file = fopen(path, "a+");
275
276     /* Open file */
277     if(!log_file) {
278         t_print(ERROR_FOPEN);
279         return 0;
280     }
281
282     /* Add MJD timestamp */
283     if(stamp_switch == 1) {
284         int timestamp_size = 50;
285         char timestamp[timestamp_size];
286         memset(timestamp, '\0', timestamp_size);
287
288         get_today_mjd(timestamp);
289         if(!fprintf(log_file, "%s", timestamp)) {
290             t_print(ERROR_FWRITE);
291             return 0;
292         }
293     }
294
295     /* Just stamp with regular time */
296     if(stamp_switch == 2){
297         char timestamp[100];
298         memset(timestamp, '\0', 100);
299         time_t rawtime;
300         struct tm *info;
301         time(&rawtime);
302         info = gmtime(&rawtime);
303         strftime(timestamp, 80, "[%x - %X]", info);
304
305         if(!fprintf(log_file, "%s", timestamp)) {
306             t_print(ERROR_FWRITE);
307             return 0;
308         }
309     }
310
311     /* Write content to file */
312     if(!fprintf(log_file, "%s", content)) {
313         t_print(ERROR_FWRITE);
314         return 0;
315     }
316
317     /* Close file */
318     if(fcclose(log_file)) {
319         t_print(ERROR_FCLOSE);
320     }
321     return 1;
322 }

```

utils.h

```
1  /**
2   * @file utils.h
3   * @author Aril Schultzen
4   * @date 13.04.2016
5   * @brief File containing function prototypes and includes for utils.c
6   */
7
8  #ifndef UTILS_H
9  #define UTILS_H
10
11 #include <stdio.h>
12 #include <stdarg.h>
13 #include <stdlib.h>
14 #include <arpa/inet.h>
15 #include <string.h>
16 #include <time.h>
17
18 #include "list.h"
19 #include "config.h"
20
21 /** @brief Terminates program and prints the line
22  *      number where die was called from.
23  *
24  * @param line_number Line number where die() was written
25  * @param format String with error description.
26  * @return Void
27  */
28 void die (int line_number, const char * format, ...);
29
30 /** @brief Extracts IP address from file descriptor
31  *
32  * @param session_fd file descriptor for the session
33  * @param ip Buffer to store the IP address as string.
34  */
35 void get_ip_str(int session_fd, char *ip);
36
37 /** @brief Extracts IP address from sockaddr struct
38  *
39  * Used by get_ip_str() to extract IP address from
40  * sockaddr struct.
41  *
42  * @param session_fd file descriptor for the session
43  * @param ip Buffer to store the IP address as string.
44  * @return Void
45  */
46 void extract_ip_str(const struct sockaddr *sa, char *s, size_t maxlen);
47
48 /** @brief Print function with time-stamp
49  *
50  * Print function like printf() but with time-stamp
51  * in square bracket appended before the String.
52  * Example: [04/13/16 - 08:50:41] Waiting for connections..
53  *
54  * @param format String to print
55  * @return Void
56  */
57 void t_print(const char* format, ...);
58
59 /** @brief Loads config from file using config_map_entry struct
60  *
```



```

61 *   Uses the config_map_entry struct to find the correct entry
62 *   in the config file, cast it to correct type and fill the
63 *   respective memory area (pointer).
64 *
65 *   @param cme Pointer to the config_map_entry struct
66 *   @param path Path to config file
67 *   @param entries Entries in the config file
68 *   @return 1 if success, 0 if fail.
69 */
70 int load_config(struct config_map_entry *cme, char *path, int entries);
71
72 /** @brief Calculates the checksum of a given string of NMEA data.
73 *
74 *   Used to check the integrity of NMEA data from the
75 *   GPS receiver before potential analysis.
76 *
77 *   @param nmea String containing NMEA data to check
78 *   @return 1 if success, 0 if fail.
79 */
80 int calculate_nmea_checksum(char *s);
81
82 /** @brief Extracts words from a String
83 *
84 *   Used to extract a substring from a string by using a
85 *   delimiter. The from and to parameters defines which
86 *   occurrence of the delimiter in the parent string to
87 *   use as start and end for the substring.
88 *
89 *   @param start The delimiter number to start from
90 *   @param end The delimiter number to stop
91 *   @param delimiter Symbol/character to use as delimit
92 *   @param buffer Buffer to store the word(s)
93 *   @param bufsize Size of buffer
94 *   @param string Pointer to parent string
95 *   @param str_len Length of parent string
96 *   @return 1 if success, 0 if no string within the delimits was found.
97 */
98 int substring_extractor(int start, int end, char delimiter, char *buffer,
99                        int bufsize, char *string, int str_len);
100
101 /** @brief Counts bytes from start to first occurrence of null character
102 *
103 *   @param buffer Buffer to search through
104 *   @param buf_len Length of the buffer in bytes
105 */
106 int str_len_u(char *buffer, int buf_len);
107
108 /** @brief Calls a script using run_command to get mjd(now).
109 *
110 *   @param buffer Buffer to store response
111 */
112 int get_today_mjd(char *buffer);
113
114 /** @brief Run a script or a program through the shell
115 *
116 *   @param path Path to program
117 *   @param output Buffer to store response
118 */
119 int run_command(char *path, char *output);
120
121 /** @brief Log to file
122 *

```

```

123 | *   @param content Data to log
124 | *   @param path Path to the log file to log to
125 | *   @param stamp_switch 0 if no timestamp, 1 if MJD.
126 | */
127 | int log_to_file(char *path, char *content, int stamp_switch);
128 | #endif /* !UTILS_H */

```

net.c

```

1 | #include "net.h"
2 |
3 | int s_read(struct transmission_s *tsm)
4 | {
5 |     bzero(tsm->iobuffer, IO_BUFFER_SIZE);
6 |     return read(tsm->session_fd, tsm->iobuffer, IO_BUFFER_SIZE);
7 | }
8 |
9 | int s_write(struct transmission_s *tsm, char *message, int length)
10 | {
11 |     return write(tsm->session_fd, message, length);
12 | }

```

net.h

```

1 | #ifndef NET_H
2 | #define NET_H
3 |
4 | #define _GNU_SOURCE 1
5 | #include <unistd.h>
6 | #include <sys/mman.h>
7 |
8 | #include <stdio.h>
9 | #include <stdlib.h>
10 | #include <string.h>
11 | #include <strings.h>
12 | #include <sys/types.h>
13 | #include <sys/socket.h>
14 | #include <netinet/in.h>
15 | #include <netdb.h>
16 | #include <errno.h>
17 | #include <stdarg.h>
18 | #include <signal.h>
19 | #include <sys/wait.h>
20 | #include <arpa/inet.h>
21 | #include <stdbool.h>
22 |
23 | /* My own header files */
24 | #include "utils.h"
25 | #include "protocol.h"
26 | #include "nmea.h"
27 |
28 | /* GENERAL */
29 | #define IO_BUFFER_SIZE MAX_PARAMETER_SIZE + MAX_COMMAND_SIZE
30 |
31 | struct transmission_s {
32 |     int session_fd;
33 |     char iobuffer[IO_BUFFER_SIZE];
34 | };
35 |

```

```

36 | int s_read(struct transmission_s *tsm);
37 | int s_write(struct transmission_s *tsm, char *message, int length);
38 |
39 | #endif /* !NET_H */

```

csac_filter.c

```

1 | #include "csac_filter.h"
2 |
3 | /* PATH TO CONFIG FILE */
4 | #define CSAC_FILTER_CONFIG_PATH "cfilter_config.ini"
5 |
6 | /* CONFIG CONSTANTS */
7 | #define CONFIG_CFD_PATH "cfd_state_path: "
8 | #define CONFIG_INIT_FROM_FILE "init_cfd_from_file: "
9 | #define CONFIG_TELEMETRY_LOG "telemetry_log: "
10 | #define CONFIG_LOG_PREDICTION "log_predict: "
11 | #define CONFIG_LOG_PRED_PATH "pred_log_path: "
12 | #define CONFIG_INIT_SSC "init_cfd_steer_smooth_current: "
13 | #define CONFIG_INIT_SST "init_cfd_steer_smooth_today: "
14 | #define CONFIG_INIT_SSP "init_cfd_steer_smooth_previous: "
15 | #define CONFIG_INIT_SSY "init_cfd_steer_smooth_yesterday: "
16 | #define CONFIG_PHASE_LIMIT "phase_limit: "
17 | #define CONFIG_STEER_LIMIT "steer_limit: "
18 | #define CONFIG_PRED_LIMIT "pred_limit: "
19 | #define CONFIG_TIME_CONSTANT "time_constant: "
20 | #define CONFIG_WARMUP_DAYS "warmup_days: "
21 | #define CSAC_FILTER_CONFIG_ENTRIES 14
22 |
23 | #define ALARM_FAST_TIMING_FILTER " [ ALARM ] Phase > Limit\n"
24 | #define ALARM_STEER_TO_BIG " [ ALARM ] Steer > limit!\n"
25 | #define ALARM_FREQ_COR_FILTER " [ ALARM ] Steer > predicted!\n"
26 |
27 | static double mjd_diff_day(double mjd_a,
28 |                          double mjd_b)
29 | {
30 |     double diff = mjd_a - mjd_b;
31 |     return diff;
32 | }
33 |
34 | static double get_mjdf()
35 | {
36 |     double mjd_today = 0;
37 |     const int BUFFER_LEN = 100;
38 |     char buffer[BUFFER_LEN];
39 |     memset(buffer, '\0', BUFFER_LEN);
40 |     if(!get_today_mjd(buffer)) {
41 |         t_print("get_mjdf(): Failed to calculate current MJD\n");
42 |         return 0;
43 |     } else {
44 |         if(sscanf(buffer, "%lf", &mjd_today) == EOF) {
45 |             t_print("get_mjdf(): Failed to cast MJD to float\n");
46 |             return 0;
47 |         }
48 |     }
49 |     return mjd_today;
50 | }
51 |
52 | static int load_telemetry(struct csac_model_data
53 |                          *cfd, char *telemetry)

```

```

54 {
55     const int BUFFER_LEN = 100;
56     char buffer[BUFFER_LEN];
57
58     /* Checking discipline mode of the CSAC */
59     if(!substring_extractor(13,14,',',buffer,100,
60         telemetry,strlen(telemetry))) {
61         printf("Failed to extract DiscOK from CSAC data\n");
62         return 0;
63     } else {
64         if(sscanf(buffer, "%d", &cfd->discok) == EOF) {
65             return 0;
66         }
67         /* CSAC is in holdover or acquiring */
68         if(cfd->discok == 2) {
69             return 0;
70         }
71     }
72
73     if(!substring_extractor(12,13,',',buffer,100,
74         telemetry,strlen(telemetry))) {
75         printf("Failed to extract Phase from CSAC data\n");
76         return 0;
77     } else {
78         if(sscanf(buffer, "%lf",
79             &cfd->phase_current) == EOF) {
80             return 0;
81         }
82     }
83
84     if(!substring_extractor(10,11,',',buffer,100,
85         telemetry,strlen(telemetry))) {
86         printf("Failed to extract Steer from CSAC data\n");
87         return 0;
88     } else {
89         if(sscanf(buffer, "%lf",
90             &cfd->steer_current) == EOF) {
91             return 0;
92         }
93     }
94
95     double mjd_today = get_mjdf();
96     if(!mjd_today){
97         return 0;
98     }
99
100     if(mjd_diff_day(mjd_today, cfd->today_mjd) >= 1
101         && cfd->t_current != 0) {
102         cfd->new_day = 1;
103         cfd->today_mjd = mjd_today;
104         cfd->days_passed++;
105     }
106
107     // Initializing today_mjd, only done once at startup
108     if(cfd->today_mjd == 0) {
109         cfd->today_mjd = mjd_today;
110         cfd->days_passed = 0;
111     }
112
113     // Updating running MJD
114     cfd->t_current = mjd_today;
115     return 1;

```

```

116 }
117
118 static void calc_smooth(struct csac_model_data
119                        *cfd)
120 {
121     double W = cfd->cf_conf.time_constant;
122
123     /* Setting previous values */
124     cfd->t_smooth_previous = cfd->t_smooth_current;
125     cfd->steer_smooth_previous =
126         cfd->steer_smooth_current;
127
128     /* Calculating t_smooth_current */
129     cfd->t_smooth_current = (((W-1)/W) *
130                            cfd->t_smooth_previous) + ((1/W) *
131                            cfd->t_current);
132
133     /* Calculating steer_smooth_current */
134     cfd->steer_smooth_current = (((W-1)/W) *
135                                cfd->steer_smooth_previous) + ((1/W) *
136                                cfd->steer_current);
137 }
138
139 /*
140 * Returns 1 if abs(phase_current) is bigger
141 */
142 int fast_timing_filter(int phase_current, int phase_limit)
143 {
144     if(abs(phase_current) > phase_limit) {
145         return 1;
146     }
147     return 0;
148 }
149
150 /*
151 * Returns 1 if abs(cfd->steer_current - cfd->steer_prediction) is bigger
152 */
153 int freq_cor_filter(struct csac_model_data *cfd)
154 {
155     if ( abs(cfd->steer_current -
156            cfd->steer_prediction) >
157          cfd->cf_conf.steer_limit) {
158         return 1;
159     }
160     return 0;
161 }
162
163 static void update_model(struct csac_model_data *cfd)
164 {
165     /* Updating t_smooth */
166     cfd->t_smooth_yesterday = cfd->t_smooth_today;
167     cfd->t_smooth_today = cfd->t_smooth_current;
168
169     /* Updating steer_smooth */
170     cfd->steer_smooth_yesterday =
171         cfd->steer_smooth_today;
172     cfd->steer_smooth_today =
173         cfd->steer_smooth_current;
174
175     /* Updating steer prediction, just for show */
176     get_steer_predict(cfd);
177 }

```

```

178
179 double get_steer_predict(struct csac_model_data *cfd)
180 {
181     if(cfd->days_passed >= cfd->cf_conf.warmup_days) {
182         cfd->steer_prediction = cfd->t_current - cfd->t_smooth_today;
183         cfd->steer_prediction = cfd->steer_prediction *
184             (cfd->steer_smooth_today -
185              cfd->steer_smooth_yesterday);
186         cfd->steer_prediction = cfd->steer_prediction /
187             (cfd->t_smooth_today - cfd->t_smooth_yesterday);
188         cfd->steer_prediction = cfd->steer_prediction
189             +cfd->steer_smooth_today;
190         return cfd->steer_prediction;
191     } else {
192         return -1;
193     }
194 }
195
196 /* Making sure there are no 0 values about */
197 int init_csac_model(struct csac_model_data *cfd,
198                   char *telemetry)
199 {
200
201     if(!load_telemetry(cfd, telemetry)) {
202         return 0;
203     }
204
205     /* Setting preliminary values, don't want to divide by zero */
206     cfd->t_smooth_current = cfd->t_current;
207     cfd->t_smooth_today = cfd->t_smooth_current;
208     cfd->t_smooth_yesterday = cfd->t_smooth_current
209         -0.1;
210
211     /* Setting values from config if preset */
212     if(cfd->cf_conf.init_cfd_from_file) {
213         cfd->steer_smooth_current =
214             cfd->cf_conf.init_cfd_ssc;
215         cfd->steer_smooth_today =
216             cfd->cf_conf.init_cfd_sst;
217         cfd->steer_smooth_previous =
218             cfd->cf_conf.init_cfd_ssp;
219         cfd->steer_smooth_yesterday =
220             cfd->cf_conf.init_cfd_ssy;
221
222         /* Setting preliminary values, don't want to divide by zero */
223     } else {
224         cfd->steer_smooth_current = cfd->steer_current;
225         cfd->steer_smooth_today =
226             cfd->steer_smooth_current;
227         cfd->steer_smooth_previous =
228             cfd->steer_smooth_today;
229     }
230
231     if(cfd->cf_conf.warmup_days == 0) {
232         cfd->new_day = 1;
233     }
234
235     return 1;
236 }
237
238 /* Update the filter with new data */
239 int update_csac_model(struct csac_model_data

```

```

240         *cfd, char *telemetry)
241     {
242         /* Load new telemetry into the filter */
243         if(!load_telemetry(cfd, telemetry) ) {
244             fprintf(stderr, "Telemetry failed to load\n");
245             return 0;
246         }
247
248         /* Calculate smoothed values */
249         calc_smooth(cfd);
250
251         /* Updating prediction if 24 hours has passed since the last update */
252         if(cfd->new_day == 1) {
253
254             /* Update prediction */
255             update_model(cfd);
256
257             /* Updating fast timing filter status */
258             cfd->ftf_status = fast_timing_filter(
259                 cfd->phase_current, cfd->cf_conf.phase_limit);
260
261             /* Updating frequency correction filter status */
262             cfd->fqf_status = freq_cor_filter(cfd);
263
264             /* Clearing new day variable*/
265             cfd->new_day = 0;
266
267             /* If logging is enabled, log steer predicted */
268             if(cfd->cf_conf.pred_logging) {
269                 char log_output[200];
270                 memset(log_output, '\0', 200);
271                 snprintf(log_output, 100, "%lf\n",
272                     cfd->steer_prediction);
273                 log_to_file(cfd->cf_conf.pred_log_path,
274                     log_output, 1);
275             }
276         }
277         return 1;
278     }
279
280     /* Setting up the config structure specific for the server */
281     static void initialize_config(struct
282         config_map_entry *conf_map,
283         struct csac_map_config *cf_conf)
284     {
285         conf_map[0].entry_name = CONFIG_CFD_PATH;
286         conf_map[0].modifier = FORMAT_STRING;
287         conf_map[0].destination = &cf_conf->cfd_state_path;
288
289         conf_map[1].entry_name = CONFIG_INIT_FROM_FILE;
290         conf_map[1].modifier = FORMAT_INT;
291         conf_map[1].destination = &cf_conf->init_cfd_from_file;
292
293         conf_map[2].entry_name = CONFIG_INIT_SSC;
294         conf_map[2].modifier = FORMAT_DOUBLE;
295         conf_map[2].destination = &cf_conf->init_cfd_ssc;
296
297         conf_map[3].entry_name = CONFIG_INIT_SST;
298         conf_map[3].modifier = FORMAT_DOUBLE;
299         conf_map[3].destination = &cf_conf->init_cfd_sst;
300
301         conf_map[4].entry_name = CONFIG_INIT_SSP;

```

```

302     conf_map[4].modifier = FORMAT_DOUBLE;
303     conf_map[4].destination = &cf_conf->init_cfd_ssp;
304
305     conf_map[5].entry_name = CONFIG_PHASE_LIMIT;
306     conf_map[5].modifier = FORMAT_DOUBLE;
307     conf_map[5].destination = &cf_conf->phase_limit;
308
309     conf_map[6].entry_name = CONFIG_STEER_LIMIT;
310     conf_map[6].modifier = FORMAT_DOUBLE;
311     conf_map[6].destination = &cf_conf->steer_limit;
312
313     conf_map[7].entry_name = CONFIG_TIME_CONSTANT;
314     conf_map[7].modifier = FORMAT_DOUBLE;
315     conf_map[7].destination = &cf_conf->time_constant;
316
317     conf_map[8].entry_name = CONFIG_WARMUP_DAYS;
318     conf_map[8].modifier = FORMAT_INT;
319     conf_map[8].destination = &cf_conf->warmup_days;
320
321     conf_map[9].entry_name = CONFIG_INIT_SSY;
322     conf_map[9].modifier = FORMAT_DOUBLE;
323     conf_map[9].destination = &cf_conf->init_cfd_ssy;
324
325     conf_map[10].entry_name = CONFIG_PRED_LIMIT;
326     conf_map[10].modifier = FORMAT_DOUBLE;
327     conf_map[10].destination = &cf_conf->pred_limit;
328
329     conf_map[11].entry_name = CONFIG_TELEMETRY_LOG;
330     conf_map[11].modifier = FORMAT_STRING;
331     conf_map[11].destination = &cf_conf->telemetry_log_path;
332
333     conf_map[12].entry_name = CONFIG_LOG_PREDICTION;
334     conf_map[12].modifier = FORMAT_INT;
335     conf_map[12].destination = &cf_conf->pred_logging;
336
337     conf_map[13].entry_name = CONFIG_LOG_PRED_PATH;
338     conf_map[13].modifier = FORMAT_STRING;
339     conf_map[13].destination = &cf_conf->pred_log_path;
340 }
341
342 void steer_csac(int prediction)
343 {
344     /* Allocating buffer for run_program() */
345     char program_buf[200];
346     memset(program_buf, '\0', 200);
347
348     /* Buffer for the prediction */
349     char pred_string[200];
350     memset(pred_string, '\0', 200);
351     sprintf(pred_string, "%d", prediction);
352
353     /* Buffer for the steer adjust command string */
354     char steer_com_string[200];
355     memset(steer_com_string, '\0', 200);
356
357     /* Building the string */
358     strcat(steer_com_string, "python query_csac.py FA");
359     strcat(steer_com_string, pred_string);
360
361
362     /* Acquiring lock on CSAC serial*/
363     sem_wait(&(s_synch->csac_sem));

```



```

364
365     /* Adjusting frequency according to the models prediction */
366     run_command(steer_com_string, program_buf);
367
368     /* Releasing lock on CSAC serial*/
369     sem_post(&(s_synch->csac_sem));
370
371     char log_buf[200];
372     memset(log_buf, '\0', 200);
373     strcat(log_buf, " Steer: ");
374     strcat(log_buf, pred_string);
375     strcat(log_buf, ", response: ");
376     strcat(log_buf, program_buf);
377
378     /* Logging steer value */
379     log_to_file(s_conf->log_path, log_buf, 2);
380 }
381
382 void disable_csac_disc()
383 {
384     /* Allocating buffer for run_program() */
385     char program_buf[200];
386     memset(program_buf, '\0', 200);
387
388     /* Acquiring lock on CSAC serial*/
389     sem_wait(&(s_synch->csac_sem));
390
391     /* Disabling disciplining */
392     run_command("python query_csac.py Md",
393               program_buf);
394
395     fprintf(stderr, "Disabling CSAC disciplining: %s\n", program_buf);
396
397     /* Releasing lock on CSAC serial*/
398     sem_post(&(s_synch->csac_sem));
399 }
400
401 void enable_csac_disc()
402 {
403     /* Allocating buffer for run_program() */
404     char program_buf[200];
405     memset(program_buf, '\0', 200);
406
407     /* Acquiring lock on CSAC serial*/
408     sem_wait(&(s_synch->csac_sem));
409
410     /* Disabling disciplining */
411     run_command("python query_csac.py MD",
412               program_buf);
413
414     fprintf(stderr, "Enabling CSAC disciplining: %s\n", program_buf);
415
416     /* Releasing lock on CSAC serial*/
417     sem_post(&(s_synch->csac_sem));
418 }
419
420 int check_filters(struct csac_model_data *cmd)
421 {
422     if(freq_cor_filter(cmd)){
423         log_to_file(s_conf->log_path, ALARM_FREQ_COR_FILTER, 2);
424         return 1;
425     }

```

```

426
427  /* If current steer is bigger than the predicted limit */
428  if( abs(cfd->steer_current) > cfd->cf_conf.pred_limit ){
429      log_to_file(s_conf->log_path, ALARM_STEER_TO_BIG, 2);
430      return 1;
431  }
432
433  if(fast_timing_filter(cfd->phase_current, cfd->cf_conf.phase_limit)){
434      log_to_file(s_conf->log_path, ALARM_FAST_TIMING_FILTER, 2);
435      return 1;
436  }
437
438  return 0;
439 }
440
441 int start_csac_model(struct csac_model_data
442                    *cfd)
443 {
444     int raised_alarm = 0;
445     int csac_disc = 1;
446
447     /* Allocating buffer for run_program() */
448     char program_buf[200];
449     memset(program_buf, '\0', 200);
450     int model_init = 0;
451
452     /* csac_filter config */
453     struct config_map_entry
454         conf_map[CSAC_FILTER_CONFIG_ENTRIES];
455
456     /* Initialize config map */
457     initialize_config(conf_map, &cfd->cf_conf);
458
459     /* Load the config */
460     if(!load_config(conf_map, CSAC_FILTER_CONFIG_PATH,
461                   CSAC_FILTER_CONFIG_ENTRIES)){
462         t_print("CSAC model/filter: Failed to load config\n");
463         s_synch->done = 1;
464         return -1;
465     }
466
467     /* Keep going as long as the server is running */
468     while(!s_synch->done) {
469         /* Acquiring lock*/
470         sem_wait(&(s_synch->csac_sem));
471
472         /* Querying CSAC */
473         run_command("python get_telemetry.py",
474                   program_buf);
475
476         /* Releasing lock */
477         sem_post(&(s_synch->csac_sem));
478
479         /* Initialize model if not already initialized */
480         if(!model_init) {
481             model_init = init_csac_model(cfd, program_buf);
482         }
483
484         /* checking alarm */
485         if(cfd->days_passed >= cfd->cf_conf.warmup_days){
486             raised_alarm = check_filters(cfd);
487         }

```

```

488
489     /* If the alarm is raised */
490     if(raised_alarm){
491         if(csac_disc){
492             disable_csac_disc();
493             csac_disc = 0;
494         }
495
496         /* Get mjd to update filter */
497         double mjd_today = get_mjdf();
498
499         /* Calculating MJD */
500         cfd->t_current = mjd_today;
501
502         /* Calc steer predict */
503         int steer_pred = (int)get_steer_predict(cfd);
504         steer_pred = steer_pred * 1000;
505
506         /* Steering CSAC */
507         steer_csac(steer_pred);
508     }
509
510     /* If the alarm is not raised */
511     if(!raised_alarm){
512         if(!csac_disc){
513             enable_csac_disc();
514             csac_disc = 1;
515         }
516         update_csac_model(cfd, program_buf);
517     }
518
519     /* If logging enabled, log all data from the CSAC */
520     if(s_conf->csac_logging) {
521         log_to_file(s_conf->csac_log_path, program_buf,1);
522     }
523
524     /* Dump filter data for every iteration */
525     dump_cfd(cfd->cf_conf.cfd_state_path);
526
527     sleep(0.5);
528     memset(program_buf, '\0', 200);
529 }
530 return 0;
531 }

```

csac_filter.h

```

1  /**
2   * @csac_filter.h
3   * @author Aril Schultzen
4   * @date 05.09.2016
5   * @brief Filter module using CSAC for the sensor_server
6   */
7
8  #ifndef CSAC_FILTER_H
9  #define CSAC_FILTER_H
10
11 #include <stdio.h>
12 #include <stdlib.h>
13 #include <string.h>

```

```

14 #include <stdarg.h>
15 #include <errno.h>
16 #include <unistd.h>
17 #include "utils.h"
18 #include "serial.h"
19
20 #include "sensor_server.h"
21
22 struct csac_map_config {
23     int pred_logging;
24     char cfd_state_path[PATH_LENGTH_MAX];
25     char telemetry_log_path[PATH_LENGTH_MAX];
26     char pred_log_path[PATH_LENGTH_MAX];
27     int init_cfd_from_file;
28     double init_cfd_ssc;
29     double init_cfd_sst;
30     double init_cfd_ssp;
31     double init_cfd_ssy;
32     double phase_limit;
33     double steer_limit;
34     double time_constant;
35     double pred_limit;
36     int warmup_days;
37 };
38
39 struct csac_model_data {
40     /* Phase */
41     double phase_current;
42
43     /* Current */
44     double t_current;
45     double steer_current;
46     double steer_prediction;
47
48     /* Current smooth */
49     double t_smooth_current;
50     double steer_smooth_current;
51
52     /* Previous */
53     double t_smooth_previous;
54     double steer_smooth_previous;
55
56     double t_smooth_today;
57     double steer_smooth_today;
58
59     double t_smooth_yesterday;
60     double steer_smooth_yesterday;
61
62     /* Changes once a day */
63     double today_mjd;
64
65     /* Days passed since startup */
66     int days_passed;
67
68     /* New day, 1 if yes, 0 if no */
69     int new_day;
70
71     /* Discipline mode */
72     int discok;
73
74
75

```

```

76     /* fast timing filter status */
77     int ftf_status;
78
79     /* Frequency correction filter status */
80     int fqf_status;
81
82     /* Config */
83     struct csac_map_config cf_conf;
84 };
85
86 /** @brief Updates the state of the filter from data
87     *     received from the CSAC
88     *
89     * @param cfd State of filter
90     * @param telemetry String of telemetry from the CSAC
91     * @return 0 if error, 1 if success.
92     */
93 int update_csac_model(struct csac_model_data *cfd, char *telemetry);
94
95 /** @brief Initializes the state of the filter by using
96     *     telemetry from the CSAC.
97     *
98     * @param cfd State of filter
99     * @param telemetry String of telemetry from the CSAC
100    * @return 0 if error, 1 if success.
101    */
102 int init_csac_model(struct csac_model_data *cfd, char *telemetry);
103
104 /** @brief Updates the state of the filter from data
105     *     received from the CSAC
106     *
107     * @param cfd State of filter
108     * @return The predicted steer value as double.
109     */
110 double get_steer_predict(struct csac_model_data *cfd);
111
112 /** @brief Starts the csac_filter
113     *
114     * @param cfd State of filter
115     * @return 1 if filter started successfully, 0 if not.
116     */
117 int start_csac_model(struct csac_model_data *cfd);
118
119 #endif /* !CSAC_FILTER_H */

```

cfilter_config.ini

```

1 | cfd_state_path: cfd_state.txt
2 | telemetry_log: telemetry.txt
3 | pred_log_path: pred.txt
4 | log_predict: 0
5 | init_cfd_from_file: 0
6 | init_cfd_steer_smooth_current: -49.134799
7 | init_cfd_steer_smooth_today: -43.508808
8 | init_cfd_steer_smooth_previous: -49.135013
9 | init_cfd_steer_smooth_yesterday: -43.091711
10 | phase_limit: 50
11 | steer_limit: 50
12 | time_constant: 10000
13 | warmup_days: 0

```

14 | pred_limit: 200

get_telemetry.py

```
1 import ctypes
2 import fileinput, sys
3 import datetime
4 import time
5 import io
6 import os
7 import serial
8
9 def main_routine():
10     ser = serial.Serial("/dev/ttyUSB0",57600, timeout=0.1)
11     sio = io.TextIOWrapper(io.BufferedRWPair(ser, ser),encoding='ascii',newline="\r\n")
12
13     log_file = open("telemetry.txt", "a+")
14
15     telemetry_len = 0
16     while (telemetry_len < 60):
17         ser.write(b'!\r\n')
18         time.sleep(0.01)
19         telemetry = sio.readline()
20         telemetry = telemetry.strip("\r\n\x00 ")
21         telemetry_len = len(telemetry)
22
23     print(telemetry)
24     ser.close()
25     log_file.write(telemetry + "\n")
26 if __name__ == '__main__':
27     main_routine()
```

filters.c

```
1 #include "filters.h"
2
3 #define ALARM_RDF "[ ALARM ] Sensor %d triggered LS filter!\n"
4 #define ALARM_RDF_RETURNED "[ ALARM ] Sensor %d cleared LS filter!\n"
5
6 #define LOG_FILE "server_log"
7 #define LOG_STRING_LENGTH 100
8 #define MJD_LENGTH 15
9
10 /** @brief Checks for any "moving" SENSORS
11 *
12 * Checks solved position against known position.
13 * Known position loaded from the config file.
14 * @return Void
15 */
16 static void ls_filter(void);
17
18 /** @brief Checks if a sensor has been marked as moved
19 *
20 * Iterates through client_list and checks for clients marked
21 * as moved. Raises alarm.
22 *
23 * @return Void
24 */
25 static void raise_alarm(void);
```

```

26
27 static int log_alarm(int client_id, char *alarm)
28 {
29     /* allocating memory for string */;
30     char log_string[LOG_STRING_LENGTH];
31     memset(log_string, '\0', LOG_STRING_LENGTH);
32
33     /* Formatting alarm */
34     char alarm_buffer[strlen(alarm) + ID_AS_STRING_MAX];
35     memset(alarm_buffer, '\0', strlen(alarm) + ID_AS_STRING_MAX);
36     snprintf(alarm_buffer, strlen(alarm) + ID_AS_STRING_MAX, alarm, client_id);
37
38     /* Formatting output*/
39     snprintf(log_string, LOG_STRING_LENGTH, "%s", alarm_buffer);
40
41     /* Logging */
42     return log_to_file(s_conf->log_path, log_string, 2);
43 }
44
45
46 void raise_alarm(void)
47 {
48     struct client_table_entry* iterator;
49     struct client_table_entry* safe;
50
51     list_for_each_entry_safe(iterator, safe, &client_list->list, list) {
52         if(iterator->client_id > 0) {
53             /* Checking ls_filter */
54             if(iterator->fs.ls_f.moved == 1) {
55                 iterator->fs.ls_f.was_moved = 1;
56                 iterator->fs.ls_f.moved = 0;
57                 if(s_conf->logging) {
58                     log_alarm(iterator->client_id, ALARM_RDF);
59                 }
60             } else {
61                 if(iterator->fs.ls_f.was_moved) {
62                     iterator->fs.ls_f.was_moved = 0;
63                     if(s_conf->logging) {
64                         log_alarm(iterator->client_id, ALARM_RDF_RETURNED);
65                     }
66                 }
67             }
68         }
69     }
70 }
71
72 void ls_filter(void)
73 {
74     struct client_table_entry* iterator;
75     struct client_table_entry* safe;
76
77     list_for_each_entry_safe(iterator, safe, &client_list->list, list) {
78
79         if(iterator->nmea.lat_current > iterator->fs.ls_f.lsf_d.lat_ref +
80             iterator->fs.ls_f.lsf_d.lat_dev) {
81             iterator->fs.ls_f.moved = 1;
82             iterator->fs.ls_f.dv.lat_disturbed = HIGH;
83         } else if(iterator->nmea.lat_current < iterator->fs.ls_f.lsf_d.lat_ref -
84             iterator->fs.ls_f.lsf_d.lat_dev) {
85             iterator->fs.ls_f.moved = 1;
86             iterator->fs.ls_f.dv.lat_disturbed = LOW;
87         } else {

```

```

88         iterator->fs.ls_f.dv.lat_disturbed = SAFE;
89     }
90
91     if(iterator->nmea.alt_current > iterator->fs.ls_f.lsf_d.alt_ref +
92         iterator->fs.ls_f.lsf_d.alt_dev) {
93         iterator->fs.ls_f.moved = 1;
94         iterator->fs.ls_f.dv.alt_disturbed = HIGH;
95     } else if(iterator->nmea.alt_current < iterator->fs.ls_f.lsf_d.alt_ref -
96         iterator->fs.ls_f.lsf_d.alt_dev) {
97         iterator->fs.ls_f.moved = 1;
98         iterator->fs.ls_f.dv.alt_disturbed = LOW;
99     } else {
100         iterator->fs.ls_f.dv.alt_disturbed = SAFE;
101     }
102
103     if(iterator->nmea.lon_current > iterator->fs.ls_f.lsf_d.lon_ref +
104         iterator->fs.ls_f.lsf_d.lon_dev) {
105         iterator->fs.ls_f.moved = 1;
106         iterator->fs.ls_f.dv.lon_disturbed = HIGH;
107     } else if(iterator->nmea.lon_current < iterator->fs.ls_f.lsf_d.lon_ref -
108         iterator->fs.ls_f.lsf_d.lon_dev) {
109         iterator->fs.ls_f.moved = 1;
110         iterator->fs.ls_f.dv.lon_disturbed = LOW;
111     } else {
112         iterator->fs.ls_f.dv.lon_disturbed = SAFE;
113     }
114
115     if(iterator->nmea.speed_current > iterator->fs.ls_f.lsf_d.speed_ref +
116         iterator->fs.ls_f.lsf_d.speed_dev) {
117         iterator->fs.ls_f.moved = 1;
118         iterator->fs.ls_f.dv.speed_disturbed = HIGH;
119     } else if(iterator->nmea.speed_current < iterator->fs.ls_f.lsf_d.speed_ref -
120         iterator->fs.ls_f.lsf_d.speed_dev) {
121         iterator->fs.ls_f.moved = 1;
122         iterator->fs.ls_f.dv.speed_disturbed = LOW;
123     } else {
124         iterator->fs.ls_f.dv.speed_disturbed = SAFE;
125     }
126 }
127 }
128
129 void apply_filters()
130 {
131     ls_filter();
132     raise_alarm();
133 }

```

filters.h

```

1  /**
2   * @file filters.h
3   * @author Aril Schultzen
4   * @date 13.04.2016
5   * @brief File containing function prototypes and includes for analyzer.h
6   */
7
8  #ifndef ANALYZER_H
9  #define ANALYZER_H
10
11  #include "sensor_server.h"

```



```

12 |
13 | void apply_filters();
14 |
15 | #endif /* !ANALYZER_H */

```

net.c

```

1 | #include "net.h"
2 |
3 | int s_read(struct transmission_s *tsm)
4 | {
5 |     bzero(tsm->iobuffer, IO_BUFFER_SIZE);
6 |     return read(tsm->session_fd, tsm->iobuffer, IO_BUFFER_SIZE);
7 | }
8 |
9 | int s_write(struct transmission_s *tsm, char *message, int length)
10 | {
11 |     return write(tsm->session_fd, message, length);
12 | }

```

net.h

```

1 | #ifndef NET_H
2 | #define NET_H
3 |
4 | #define _GNU_SOURCE 1
5 | #include <unistd.h>
6 | #include <sys/mman.h>
7 |
8 | #include <stdio.h>
9 | #include <stdlib.h>
10 | #include <string.h>
11 | #include <strings.h>
12 | #include <sys/types.h>
13 | #include <sys/socket.h>
14 | #include <netinet/in.h>
15 | #include <netdb.h>
16 | #include <errno.h>
17 | #include <stdarg.h>
18 | #include <signal.h>
19 | #include <sys/wait.h>
20 | #include <arpa/inet.h>
21 | #include <stdbool.h>
22 |
23 | /* My own header files */
24 | #include "utils.h"
25 | #include "protocol.h"
26 | #include "nmea.h"
27 |
28 | /* GENERAL */
29 | #define IO_BUFFER_SIZE MAX_PARAMETER_SIZE + MAX_COMMAND_SIZE
30 |
31 | struct transmission_s {
32 |     int session_fd;
33 |     char iobuffer[IO_BUFFER_SIZE];
34 | };
35 |
36 | int s_read(struct transmission_s *tsm);
37 | int s_write(struct transmission_s *tsm, char *message, int length);

```

```

38 |
39 | #endif /* !NET_H */

```

gps_serial.c

```

1 | #include "serial.h"
2 |
3 | int configure_gps_serial(int fd)
4 | {
5 |     struct termios tty;
6 |     memset (&tty, 0, sizeof tty);
7 |
8 |     if (tcgetattr (fd, &tty) != 0) {
9 |         printf ("error %d from tcgetattr", errno);
10 |         exit(0);
11 |     }
12 |
13 |     cfsetospeed (&tty, B9600);
14 |     cfsetispeed (&tty, B9600);
15 |
16 |     tty.c_cflag &= ~PARENB;
17 |     tty.c_cflag &= ~CSTOPB;
18 |     tty.c_cflag &= ~CSIZE;
19 |     tty.c_cflag |= CS8;
20 |     tty.c_cflag &= ~CRTSCTS;
21 |     tty.c_cflag |= CREAD | CLOCAL;
22 |     tty.c_iflag &= ~(IXON | IXOFF | IXANY);
23 |     tty.c_iflag &= ~(ICANON | ECHO | ECHOE | ISIG);
24 |     tty.c_oflag &= ~OPOST;
25 |     tty.c_cc[VMIN] = 0;
26 |     tty.c_cc[VTIME] = 0;
27 |
28 |     if (tcsetattr (fd, TCSANOW, &tty) != 0) {
29 |         printf ("error %d setting term attributes", errno);
30 |         return -1;
31 |     }
32 |     return 0;
33 | }
34 |
35 | int open_serial(char *portname, serial_device device)
36 | {
37 |     int fd = open (portname, O_RDWR | O_NOCTTY);
38 |     if (fd < 0) {
39 |         t_print ("Error %d opening %s: %s\n", errno, portname, strerror (errno));
40 |     }
41 |
42 |     if(device == GPS) {
43 |         if(configure_gps_serial(fd) < 0) {
44 |             exit(0);
45 |         }
46 |     }
47 |
48 |     return fd;
49 | }

```

serial.h

```

1 | /*
2 | ## CSAC Config #####

```

```

3 | #
4 | # 57600
5 | # 8 bit
6 | # No parity
7 | #
8 | # While CSAC is off:
9 | #
10 | # sudo stty -F /dev/ttyS0 57600
11 | # cat /dev/ttyS0
12 | #
13 | # Turn the CSAC ON
14 | #
15 | # Symmetricom CSAC <- Output
16 | #
17 | #####
18 | */
19 |
20 | #ifndef SERIAL_H
21 | #define SERIAL_H
22 |
23 | #include <errno.h>
24 | #include <termios.h>
25 | #include <unistd.h>
26 | #include <string.h> /* memset */
27 | #include <stdio.h>
28 | #include <stdlib.h>
29 | #include <features.h>
30 | #include <fcntl.h>
31 | #include <signal.h>
32 |
33 | //Mine
34 | #include "utils.h"
35 | #include "protocol.h"
36 |
37 | typedef enum e_serial_device {
38 |     GPS,
39 |     CSAC
40 | } serial_device;
41 |
42 | int open_serial(char *portname, serial_device device);
43 |
44 | /** @brief Queries the CSAC with the command over serial connection
45 | *
46 | * Sends a command to the CSAC and reads buf_len bytes into
47 | * the buffer. Does not deal with formatting in any way.
48 | *
49 | * @param file_descriptor FD for the CSAC serial connection
50 | * @param query Command (query) to send to the CSAC.
51 | * @param buffer Buffer to store the response
52 | * @buf_len buf_len Length of buffer
53 | */
54 | int serial_query(int file_descriptor, char *query, char *buffer, int buf_len);
55 |
56 | #endif /* !SERIAL_H */

```

colors.h

```

1 | #ifndef COLORS_H
2 | #define COLORS_H
3 |

```

```

4  /* RESET */
5  #define RESET "\033[0m"
6
7  /* COLORS */
8  #define BLK_WHT "\033[030;47m"
9
10 /* BOLD */
11 #define BOLD_BLK_WHT "\033[1;30;47m"
12 #define BOLD_GRN_BLK "\033[1;32;40m"
13 #define BOLD_RED_BLK "\033[1;31;40m"
14 #define BOLD_YLW_BLK "\033[1;33;40m"
15 #define BOLD_CYN_BLK "\033[1;36;40m"
16
17 /* BOLD INVERTED*/
18 #define BOLD_BLK_GRN "\033[7;32;40m"
19 #define BOLD_BLK_RED "\033[7;31;40m"
20 #define BOLD_BLK_YLW "\033[7;33;40m"
21 #define BOLD_WHT_CYN "\033[1;37;46m"
22
23 /* UNDERLINED */
24 #define UNDER_RED_BLACK "\033[4;031;40m"
25
26 #endif /* !COLORS_H */

```

config.h

```

1  #ifndef CONFIG_H
2  #define CONFIG_H
3
4  #define FORMAT_INT "%d"
5  #define FORMAT_FLOAT "%f"
6  #define FORMAT_STRING "%s"
7  #define FORMAT_DOUBLE "%lf"
8
9  struct config_map_entry {
10     char *entry_name;
11     char *modifier;
12     void *destination;
13 };
14
15 #endif /* !CONFIG_H */

```

nmea.h

```

1  #ifndef NMEA_H
2  #define NMEA_H
3
4  /* NMEA SENTENCES */
5  #define GGA "$GNGGA"
6  #define RMC "$GNRMC"
7  #define SENTENCE_LENGTH 100
8
9  /* NMEA SENTENCES DELIMITER POSITIONS */
10 #define ALTITUDE_START 9
11 #define LATITUDE_START 3
12 #define LONGITUDE_START 5
13 #define RMC_TIME_START 1
14 #define SPEED_START 7
15

```

```

16 | #define SAFE 0
17 | #define HIGH 1
18 | #define LOW -1
19 |
20 | struct nmea_container {
21 |     /* Raw data */
22 |     char raw_gga[SENTENCE_LENGTH];
23 |     char raw_rmc[SENTENCE_LENGTH];
24 |
25 |     /* Latitude */
26 |     double lat_current;
27 |     double lat_average;
28 |     double lat_avg_diff;
29 |     double lat_total;
30 |     int lat_disturbed;
31 |
32 |     /* Longitude */
33 |     double lon_current;
34 |     double lon_average;
35 |     double lon_avg_diff;
36 |     double lon_total;
37 |     int lon_disturbed;
38 |
39 |     /* Altitude */
40 |     double alt_current;
41 |     double alt_average;
42 |     double alt_avg_diff;
43 |     double alt_total;
44 |     int alt_disturbed;
45 |
46 |     /* Speed */
47 |     double speed_current;
48 |     double speed_average;
49 |     double speed_avg_diff;
50 |     double speed_total;
51 |     int speed_disturbed;
52 |
53 |     /* CHECKSUM */
54 |     int checksum_passed;
55 |
56 |     /* COUNTER FOR AVERAGE */
57 |     int n_samples;
58 | };
59 |
60 | #endif /* !NMEA_H */

```

list.h

```

1 | /**
2 |  * @author kazutomo@mcs.anl.gov
3 |  * @file list.h
4 |  * @brief Linked list implementation from linux kernel source code.
5 |  *
6 |  * This code was lifted from http://www.mcs.anl.gov/~kazutomo/list/.
7 |  * I stumbled upon when writing when writing a Linux clone autumn 15',
8 |  * and tested it in this project. It was planned to be replaced by something
9 |  * smaller.
10 |  * Kazutomo's description:
11 |  *
12 |  * I grab it from linux kernel source code and fix it for user space

```

```

13  * program. Of course, this is a GPL licensed header file.
14  *
15  * Here is a recipe to cook list.h for user space program
16  *
17  * 1. copy list.h from linux/include/list.h
18  * 2. remove
19  *    - #ifdef __KERNE__ and its #endif
20  *    - all #include line
21  *    - prefetch() and rcu related functions
22  * 3. add macro offsetof() and container_of
23  *
24  * - kazutomo@mcs.anl.gov
25  */
26
27 #ifndef _LINUX_LIST_H
28 #define _LINUX_LIST_H
29
30 /**
31  * @name from other kernel headers
32  */
33 /*{*/
34
35 /**
36  * Get offset of a member
37  */
38 #define offsetof(TYPE, MEMBER) ((size_t) &((TYPE *)0)->MEMBER)
39
40 /**
41  * Casts a member of a structure out to the containing structure
42  * @param ptr      the pointer to the member.
43  * @param type     the type of the container struct this is embedded in.
44  * @param member   the name of the member within the struct.
45  *
46  */
47 #define container_of(ptr, type, member) __extension__({
48     const typeof( ((type *)0)->member ) *__mptr = (ptr);
49     (type *) ( (char *)__mptr - offsetof(type,member) );})
50 /*}*/
51
52
53 /*
54  * These are non-NULL pointers that will result in page faults
55  * under normal circumstances, used to verify that nobody uses
56  * non-initialized list entries.
57  */
58 #define LIST_POISON1 ((void *) 0x00100100)
59 #define LIST_POISON2 ((void *) 0x00200200)
60
61 /**
62  * Simple doubly linked list implementation.
63  *
64  * Some of the internal functions ("__xxx") are useful when
65  * manipulating whole lists rather than single entries, as
66  * sometimes we already know the next/prev entries and we can
67  * generate better code by using them directly rather than
68  * using the generic single-entry routines.
69  */
70 struct list_head {
71     struct list_head *next, *prev;
72 };
73
74 #define LIST_HEAD_INIT(name) { &(name), &(name) }

```

```

75
76 #define LIST_HEAD(name) \
77     struct list_head name = LIST_HEAD_INIT(name)
78
79 #define INIT_LIST_HEAD(ptr) do { \
80     (ptr)->next = (ptr); (ptr)->prev = (ptr); \
81 } while (0)
82
83 /*
84  * Insert a new entry between two known consecutive entries.
85  *
86  * This is only for internal list manipulation where we know
87  * the prev/next entries already!
88  */
89 static inline void __list_add(struct list_head *new,
90                             struct list_head *prev,
91                             struct list_head *next)
92 {
93     next->prev = new;
94     new->next = next;
95     new->prev = prev;
96     prev->next = new;
97 }
98
99 /**
100  * list_add - add a new entry
101  * @new: new entry to be added
102  * @head: list head to add it after
103  *
104  * Insert a new entry after the specified head.
105  * This is good for implementing stacks.
106  */
107 static inline void list_add(struct list_head *new, struct list_head *head)
108 {
109     __list_add(new, head, head->next);
110 }
111
112 /**
113  * list_add_tail - add a new entry
114  * @new: new entry to be added
115  * @head: list head to add it before
116  *
117  * Insert a new entry before the specified head.
118  * This is useful for implementing queues.
119  */
120 static inline void list_add_tail(struct list_head *new, struct list_head *head)
121 {
122     __list_add(new, head->prev, head);
123 }
124
125 /*
126  * Delete a list entry by making the prev/next entries
127  * point to each other.
128  *
129  * This is only for internal list manipulation where we know
130  * the prev/next entries already!
131  */
132 static inline void __list_del(struct list_head * prev, struct list_head * next)
133 {
134     next->prev = prev;
135     prev->next = next;
136 }

```

```

137
138 /**
139  * list_del - deletes entry from list.
140  * @entry: the element to delete from the list.
141  * Note: list_empty on entry does not return true after this, the entry is
142  * in an undefined state.
143  */
144 static inline void list_del(struct list_head *entry)
145 {
146     __list_del(entry->prev, entry->next);
147     //entry->next = LIST_POISON1;
148     //entry->prev = LIST_POISON2;
149 }
150
151
152
153 /**
154  * list_del_init - deletes entry from list and reinitialize it.
155  * @entry: the element to delete from the list.
156  */
157 static inline void list_del_init(struct list_head *entry)
158 {
159     __list_del(entry->prev, entry->next);
160     INIT_LIST_HEAD(entry);
161 }
162
163 /**
164  * list_move - delete from one list and add as another's head
165  * @list: the entry to move
166  * @head: the head that will precede our entry
167  */
168 static inline void list_move(struct list_head *list, struct list_head *head)
169 {
170     __list_del(list->prev, list->next);
171     list_add(list, head);
172 }
173
174 /**
175  * list_move_tail - delete from one list and add as another's tail
176  * @list: the entry to move
177  * @head: the head that will follow our entry
178  */
179 static inline void list_move_tail(struct list_head *list,
180                                 struct list_head *head)
181 {
182     __list_del(list->prev, list->next);
183     list_add_tail(list, head);
184 }
185
186 /**
187  * list_empty - tests whether a list is empty
188  * @head: the list to test.
189  */
190 static inline int list_empty(const struct list_head *head)
191 {
192     return head->next == head;
193 }
194
195 static inline void __list_splice(struct list_head *list,
196                                struct list_head *head)
197 {
198     struct list_head *first = list->next;

```



```

199     struct list_head *last = list->prev;
200     struct list_head *at = head->next;
201
202     first->prev = head;
203     head->next = first;
204
205     last->next = at;
206     at->prev = last;
207 }
208
209 /**
210  * list_splice - join two lists
211  * @list: the new list to add.
212  * @head: the place to add it in the first list.
213  */
214 static inline void list_splice(struct list_head *list, struct list_head *head)
215 {
216     if (!list_empty(list))
217         __list_splice(list, head);
218 }
219
220 /**
221  * list_splice_init - join two lists and reinitialise the emptied list.
222  * @list: the new list to add.
223  * @head: the place to add it in the first list.
224  *
225  * The list at @list is reinitialised
226  */
227 static inline void list_splice_init(struct list_head *list,
228                                     struct list_head *head)
229 {
230     if (!list_empty(list)) {
231         __list_splice(list, head);
232         INIT_LIST_HEAD(list);
233     }
234 }
235
236 /**
237  * list_entry - get the struct for this entry
238  * @ptr: the @struct list_head pointer.
239  * @type: the type of the struct this is embedded in.
240  * @member: the name of the list_struct within the struct.
241  */
242 #define list_entry(ptr, type, member) \
243     container_of(ptr, type, member)
244
245 /**
246  * list_for_each - iterate over a list
247  * @pos: the @struct list_head to use as a loop counter.
248  * @head: the head for your list.
249  */
250
251 #define list_for_each(pos, head) \
252     for (pos = (head)->next; pos != (head); \
253          pos = pos->next)
254
255 /**
256  * __list_for_each - iterate over a list
257  * @pos: the @struct list_head to use as a loop counter.
258  * @head: the head for your list.
259  *
260  * This variant differs from list_for_each() in that it's the

```

```

261 * simplest possible list iteration code, no prefetching is done.
262 * Use this for code that knows the list to be very short (empty
263 * or 1 entry) most of the time.
264 */
265 #define __list_for_each(pos, head) \
266     for (pos = (head)->next; pos != (head); pos = pos->next)
267
268 /**
269  * list_for_each_prev - iterate over a list backwards
270  * @pos: the @struct list_head to use as a loop counter.
271  * @head: the head for your list.
272  */
273 #define list_for_each_prev(pos, head) \
274     for (pos = (head)->prev; prefetch(pos->prev), pos != (head); \
275         pos = pos->prev)
276
277 /**
278  * list_for_each_safe - iterate over a list safe against removal of list entry
279  * @pos: the @struct list_head to use as a loop counter.
280  * @n: another @struct list_head to use as temporary storage
281  * @head: the head for your list.
282  */
283 #define list_for_each_safe(pos, n, head) \
284     for (pos = (head)->next, n = pos->next; pos != (head); \
285         pos = n, n = pos->next)
286
287 /**
288  * list_for_each_entry - iterate over list of given type
289  * @pos: the type * to use as a loop counter.
290  * @head: the head for your list.
291  * @member: the name of the list_struct within the struct.
292  */
293 #define list_for_each_entry(pos, head, member) \
294     for (pos = list_entry((head)->next, typeof(*pos), member); \
295         &pos->member != (head); \
296         pos = list_entry(pos->member.next, typeof(*pos), member))
297
298 /**
299  * list_for_each_entry_reverse - iterate backwards over list of given type.
300  * @pos: the type * to use as a loop counter.
301  * @head: the head for your list.
302  * @member: the name of the list_struct within the struct.
303  */
304 #define list_for_each_entry_reverse(pos, head, member) \
305     for (pos = list_entry((head)->prev, typeof(*pos), member); \
306         &pos->member != (head); \
307         pos = list_entry(pos->member.prev, typeof(*pos), member))
308
309 /**
310  * list_prepare_entry - prepare a pos entry for use as a start point in
311  * list_for_each_entry_continue
312  * @pos: the type * to use as a start point
313  * @head: the head of the list
314  * @member: the name of the list_struct within the struct.
315  */
316 #define list_prepare_entry(pos, head, member) \
317     ((pos) ? : list_entry(head, typeof(*pos), member))
318
319 /**
320  * list_for_each_entry_continue - iterate over list of given type
321  * continuing after existing point
322  * @pos: the type * to use as a loop counter.

```

```

323 * @head: the head for your list.
324 * @member: the name of the list_struct within the struct.
325 */
326 #define list_for_each_entry_continue(pos, head, member) \
327     for (pos = list_entry(pos->member.next, typeof(*pos), member); \
328          &pos->member != (head); \
329          pos = list_entry(pos->member.next, typeof(*pos), member))
330
331 /**
332  * list_for_each_entry_safe - iterate over list of given type safe against removal of list entry
333  * @pos: the type * to use as a loop counter.
334  * @n: another type * to use as temporary storage
335  * @head: the head for your list.
336  * @member: the name of the list_struct within the struct.
337  */
338 #define list_for_each_entry_safe(pos, n, head, member) \
339     for (pos = list_entry((head)->next, typeof(*pos), member), \
340          n = list_entry(pos->member.next, typeof(*pos), member); \
341          &pos->member != (head); \
342          pos = n, n = list_entry(n->member.next, typeof(*n), member))
343
344 /**
345  * list_for_each_entry_safe_continue - iterate over list of given type
346  * continuing after existing point safe against removal of list entry
347  * @pos: the type * to use as a loop counter.
348  * @n: another type * to use as temporary storage
349  * @head: the head for your list.
350  * @member: the name of the list_struct within the struct.
351  */
352 #define list_for_each_entry_safe_continue(pos, n, head, member) \
353     for (pos = list_entry(pos->member.next, typeof(*pos), member), \
354          n = list_entry(pos->member.next, typeof(*pos), member); \
355          &pos->member != (head); \
356          pos = n, n = list_entry(n->member.next, typeof(*n), member))
357
358 /**
359  * list_for_each_entry_safe_reverse - iterate backwards over list of given type safe against
360  * removal of list entry
361  * @pos: the type * to use as a loop counter.
362  * @n: another type * to use as temporary storage
363  * @head: the head for your list.
364  * @member: the name of the list_struct within the struct.
365  */
366 #define list_for_each_entry_safe_reverse(pos, n, head, member) \
367     for (pos = list_entry((head)->prev, typeof(*pos), member), \
368          n = list_entry(pos->member.prev, typeof(*pos), member); \
369          &pos->member != (head); \
370          pos = n, n = list_entry(n->member.prev, typeof(*n), member))
371
372
373
374
375 /*
376  * Double linked lists with a single pointer list head.
377  * Mostly useful for hash tables where the two pointer list head is
378  * too wasteful.
379  * You lose the ability to access the tail in O(1).
380  */
381
382 struct hlist_head {
383     struct hlist_node *first;
384 };

```

```

385
386 struct hlist_node {
387     struct hlist_node *next, **pprev;
388 };
389
390 #define HLIST_HEAD_INIT { .first = NULL }
391 #define HLIST_HEAD(name) struct hlist_head name = { .first = NULL }
392 #define INIT_HLIST_HEAD(ptr) ((ptr)->first = NULL)
393 #define INIT_HLIST_NODE(ptr) ((ptr)->next = NULL, (ptr)->pprev = NULL)
394
395 static inline int hlist_unhashed(const struct hlist_node *h)
396 {
397     return !h->pprev;
398 }
399
400 static inline int hlist_empty(const struct hlist_head *h)
401 {
402     return !h->first;
403 }
404
405 static inline void __hlist_del(struct hlist_node *n)
406 {
407     struct hlist_node *next = n->next;
408     struct hlist_node **pprev = n->pprev;
409     *pprev = next;
410     if (next)
411         next->pprev = pprev;
412 }
413
414 static inline void hlist_del(struct hlist_node *n)
415 {
416     __hlist_del(n);
417     n->next = LIST_POISON1;
418     n->pprev = LIST_POISON2;
419 }
420
421
422 static inline void hlist_del_init(struct hlist_node *n)
423 {
424     if (n->pprev) {
425         __hlist_del(n);
426         INIT_HLIST_NODE(n);
427     }
428 }
429
430 static inline void hlist_add_head(struct hlist_node *n, struct hlist_head *h)
431 {
432     struct hlist_node *first = h->first;
433     n->next = first;
434     if (first)
435         first->pprev = &n->next;
436     h->first = n;
437     n->pprev = &h->first;
438 }
439
440
441
442 /* next must be != NULL */
443 static inline void hlist_add_before(struct hlist_node *n,
444                                   struct hlist_node *next)
445 {
446     n->pprev = next->pprev;

```

```

447     n->next = next;
448     next->pprev = &n->next;
449     *(n->pprev) = n;
450 }
451
452 static inline void hlist_add_after(struct hlist_node *n,
453                                 struct hlist_node *next)
454 {
455     next->next = n->next;
456     n->next = next;
457     next->pprev = &n->next;
458
459     if(next->next)
460         next->next->pprev = &next->next;
461 }
462
463
464
465 #define hlist_entry(ptr, type, member) container_of(ptr, type, member)
466
467 #define hlist_for_each(pos, head) \
468     for (pos = (head)->first; pos && ({ prefetch(pos->next); 1; }); \
469         pos = pos->next)
470
471 #define hlist_for_each_safe(pos, n, head) \
472     for (pos = (head)->first; pos && ({ n = pos->next; 1; }); \
473         pos = n)
474
475 /**
476  * hlist_for_each_entry - iterate over list of given type
477  * @tpos:   the type * to use as a loop counter.
478  * @pos:    the @struct hlist_node to use as a loop counter.
479  * @head:   the head for your list.
480  * @member: the name of the hlist_node within the struct.
481  */
482 #define hlist_for_each_entry(tpos, pos, head, member) \
483     for (pos = (head)->first; \
484         pos && ({ prefetch(pos->next); 1;}) && \
485         ({ tpos = hlist_entry(pos, typeof(*tpos), member); 1;}); \
486         pos = pos->next)
487
488 /**
489  * hlist_for_each_entry_continue - iterate over a hlist continuing after existing point
490  * @tpos:   the type * to use as a loop counter.
491  * @pos:    the @struct hlist_node to use as a loop counter.
492  * @member: the name of the hlist_node within the struct.
493  */
494 #define hlist_for_each_entry_continue(tpos, pos, member) \
495     for (pos = (pos)->next; \
496         pos && ({ prefetch(pos->next); 1;}) && \
497         ({ tpos = hlist_entry(pos, typeof(*tpos), member); 1;}); \
498         pos = pos->next)
499
500 /**
501  * hlist_for_each_entry_from - iterate over a hlist continuing from existing point
502  * @tpos:   the type * to use as a loop counter.
503  * @pos:    the @struct hlist_node to use as a loop counter.
504  * @member: the name of the hlist_node within the struct.
505  */
506 #define hlist_for_each_entry_from(tpos, pos, member) \
507     for (; pos && ({ prefetch(pos->next); 1;}) && \
508         ({ tpos = hlist_entry(pos, typeof(*tpos), member); 1;}); \

```

```

509         pos = pos->next)
510
511 /**
512  * hlist_for_each_entry_safe - iterate over list of given type safe against removal of list entry
513  * @tpos:   the type * to use as a loop counter.
514  * @pos:    the @struct hlist_node to use as a loop counter.
515  * @n:      another @struct hlist_node to use as temporary storage
516  * @head:   the head for your list.
517  * @member: the name of the hlist_node within the struct.
518  */
519 #define hlist_for_each_entry_safe(tpos, pos, n, head, member) \
520     for (pos = (head)->first; \
521          pos && ({ n = pos->next; 1; }) && \
522          ({ tpos = hlist_entry(pos, typeof(*tpos), member); 1; }); \
523          pos = n)
524
525
526 #endif

```

protocol.h

```

1  #ifndef PROTOCOL_H
2  #define PROTOCOL_H
3
4  /* CONSTRAINS */
5  #define MAX_COMMAND_SIZE 20
6  #define MAX_PARAMETER_SIZE 2048
7  #define ID_MAX 1000
8  #define MIN_COMMAND_SIZE 2
9  #define MIN_PARAMETER_SIZE 0
10
11 /* COMMANDS TO USE WHEN COMMUNICATING */
12 #define PROTOCOL_DISCONNECT "DISCONNECT"
13 #define PROTOCOL_EXIT "EXIT"
14 #define PROTOCOL_GET_TIME "GETTIME"
15 #define PROTOCOL_IDENTIFY "IDENTIFY"
16 #define PROTOCOL_NMEA "NMEA"
17 #define PROTOCOL_PRINTCLIENTS "PRINTCLIENTS"
18 #define PROTOCOL_PRINTSERVER "PRINTSERVER"
19 #define PROTOCOL_KICK "KICK"
20 #define PROTOCOL_HELP "HELP"
21 #define PROTOCOL_PRINT_LOCATION "PRINTLOC"
22 #define PROTOCOL_PRINTTIME "PRINTTIME"
23 #define PROTOCOL_DUMPDATA "DUMPDATA"
24 #define PROTOCOL_PRINTAVGDIFF "PRINTAVGDIFF"
25 #define PROTOCOL_LISTDUMPS "LISTDATA"
26 #define PROTOCOL_LOADDATA "LOADDATA"
27 #define PROTOCOL_QUERYCSAC "QUERYCSAC"
28 #define PROTOCOL_LOADKRLDATA "LOADLSFDATA"
29 #define PROTOCOL_PRINTCFD "PRINTCFD"
30
31 /* SHORT */
32 #define PROTOCOL_HELP_SHORT "?"
33 #define PROTOCOL_DISCONNECT_SHORT "DC"
34 #define PROTOCOL_DUMPDATA_SHORT "DD"
35 #define PROTOCOL_IDENTIFY_SHORT "ID"
36 #define PROTOCOL_PRINTCLIENTS_SHORT "PC"
37 #define PROTOCOL_PRINTSERVER_SHORT "PS"
38 #define PROTOCOL_PRINT_LOCATION_SHORT "PL"
39 #define PROTOCOL_PRINTAVGDIFF_SHORT "PAD"

```

```

40 #define PROTOCOL_LISTDUMPS_SHORT "LSD"
41 #define PROTOCOL_LOADDATA_SHORT "LD"
42 #define PROTOCOL_QUERYCSAC_SHORT "QC"
43 #define PROTOCOL_LOADKRLDATA_SHORT "LSFD"
44 #define PROTOCOL_PRINTCFD_SHORT "PFD"
45
46 /* RESPONSES */
47 #define PROTOCOL_GOODBYE "Goodbye!\n"
48 #define PROTOCOL_OK "OK!\n\n"
49 #define PROTOCOL_WELCOME "Welcome to the Sensor Server!\n"
50
51 /* COMMAND CODES */
52 /* Used by respond() */
53 #define CODE_DISCONNECT 1
54 #define CODE_GET_TIME 2
55 #define CODE_IDENTIFY 3
56 #define CODE_STORE 4
57 #define CODE_NMEA 5
58 #define CODE_PRINTCLIENTS 6
59 #define CODE_PRINTSERVER 7
60 #define CODE_KICK 8
61 #define CODE_HELP 9
62 #define CODE_PRINT_LOCATION 10
63 #define CODE_WARMUP 11
64 #define CODE_PRINTTIME 12
65 #define CODE_DUMPDATA 13
66 #define CODE_MOVED 14
67 #define CODE_PRINTAVGDIFF 15
68 #define CODE_LISTDUMPS 17
69 #define CODE_LOADDATA 18
70 #define CODE_QUERYCSAC 19
71 #define CODE_LOADKRLDATA 20
72 #define CODE_PRINTCFD 21
73
74 /* SIZES */
75 #define TIME_SIZE 9 /* SIZE OF TIME AS CHARS eg.142546.00, FROM GNRMC */
76
77 #endif /* !PROTOCOL_H */

```

makefile

```

1 SERVER_OBJS = sensor_server.o net.o utils.o session.o filters.o actions.o csac_filter.o
2 CLIENT_OBJS = sensor_client.o net.o utils.o gps_serial.o
3
4 CC = gcc
5 DEBUG = -g
6
7 CFLAGS = -Wall -Wextra -c -g -std=gnu99 -pedantic
8
9 cpu := $(shell uname -m)
10
11 ifeq ($(cpu),armv7l)
12     CFLAGS = -Wall -Wextra -c -std=gnu99 -pedantic -g -march=armv7-a -mtune=arm7 -fsigned-char
13 endif
14
15 LFLAGS = -Wall $(DEBUG)
16
17 server : $(SERVER_OBJS)
18     $(CC) $(LFLAGS) $(SERVER_OBJS) -o server -lpthread
19

```

```
20 client : $(CLIENT_OBJS)
21     $(CC) $(LFLAGS) $(CLIENT_OBJS) -o client
22
23 sensor_server.o : sensor_server.h net.h sensor_server.c
24     $(CC) $(CFLAGS) sensor_server.c
25
26 sensor_client.o : sensor_client.h sensor_client.c
27     $(CC) $(CFLAGS) sensor_client.c
28
29 csac_filter.o : csac_filter.h csac_filter.c utils.h sensor_server.h
30     $(CC) $(CFLAGS) csac_filter.c
31
32 net.o : net.h utils.h net.c
33     $(CC) $(CFLAGS) net.c
34
35 utils.o : utils.h list.h utils.c config.h
36     $(CC) $(CFLAGS) utils.c
37
38 gps_serial.o : serial.h gps_serial.c
39     $(CC) $(CFLAGS) gps_serial.c
40
41 session.o : session.h session.c sensor_server.h
42     $(CC) $(CFLAGS) session.c
43
44 filters.o : filters.h filters.c sensor_server.h
45     $(CC) $(CFLAGS) filters.c
46
47 actions.o : actions.h actions.c sensor_server.h
48     $(CC) $(CFLAGS) actions.c
49
50 clean:
51     \rm *.o
```


Appendix E

Scripts

CSAC query source code

```
1 import ctypes
2 import fileinput, sys
3 import datetime
4 import time
5 import io
6 import os
7 import serial
8
9 def main_routine():
10     # Opening serial stream, use ASCII
11     ser = serial.Serial("/dev/ttyUSB0",57600, timeout=0.1)
12     sio = io.TextIOWrapper(io.BufferedRWPair(ser, ser),encoding='ascii',newline="\r\n")
13
14     # Open log file, mostly used for debug
15     log_file = open("query_csac.txt", "a+")
16
17     # The query to use
18     query = sys.argv[1].strip("\r\n")
19
20     # How long to sleep between read from serial con.
21     sleep_time = 0.2
22
23     # The minimum length of the answer
24     # for the given query.
25     minimum_len = 0
26
27     if(query == '^' or query == '6'):
28         minimum_len = 80
29     elif(query == 'F'):
30         sleep_time = 0.5
31         minimum_len = 10
32     elif(query == 'M'):
33         minimum_len = 6
34     elif (query == 'S'):
35         sleep_time = 3
36         minimum_len = 2
37     else:
38         minimum_len = 1
39
```

```

40     response_len = 0
41
42     if(len(query) > 1):
43         query = "!" + query + "\r\n"
44
45     retry_count = 0
46
47     while (response_len < minimum_len):
48         ser.write(bytes(query))
49         time.sleep(sleep_time)
50         response = ser.readline()
51         response = response.strip("\r\n\x00")
52         response_len = len(response)
53         retry_count = retry_count + 1
54
55     print(response)
56     ser.close()
57     query = query.strip("\r\n")
58     log_string = ("Issued query " + "'" + query + "'" + str(retry_count) + " times\n")
59     log_file.write(log_string)
60 if __name__ == '__main__':
61     main_routine()

```

MJD calculator

```

1  #!/usr/bin/env python
2
3  import datetime
4  import jdutil
5  from dateutil import parser
6
7  today = datetime.datetime.utcnow()
8  print(jdutil.jd_to_mjd(jdutil.datetime_to_jd(today)))

```

Script example

```

1  '''
2  :Author: Aril Schultzen
3  :Email: aschultzen@gmail.com
4  '''
5  # This script attempts to connect to the
6  # Sensor Server at <ip> : "port" and
7  # IDs itself as <id>. It will then
8  # poll the time solved by the GNSS receiver
9  # connected to Sensor<id> until
10 # terminated.
11
12 import socket
13 import sys
14 import time
15
16 ip = "10.1.0.46"
17 port = 10001
18 id = 1
19
20 try:
21     s = socket.socket(socket.AF_INET, socket.SOCK_STREAM)

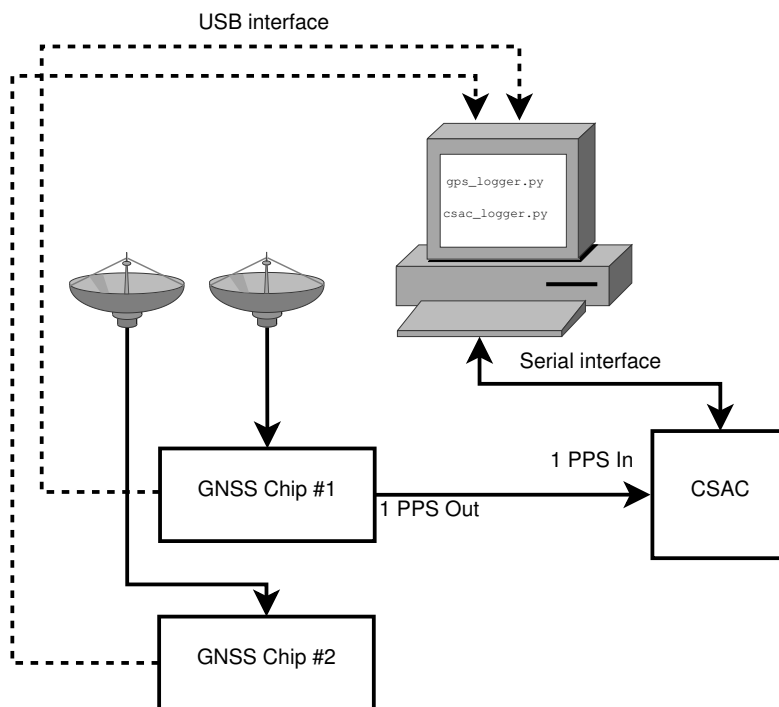
```

```

22 | except socket.error, msg:
23 |     print 'Failed to create socket. Error code: ' + str(msg[0]) + ', Error message : ' + msg[1]
24 |     sys.exit();
25 | try:
26 |     remote_ip = socket.gethostbyname( ip )
27 |
28 | except socket.gaierror:
29 |     print 'Could not resolve hostname'
30 |     sys.exit()
31 |
32 | s.connect((remote_ip , port))
33 | s.sendall(b'IDENTIFY -10')
34 | recv_buff = s.recv(1024)
35 |
36 | while(1):
37 |     s.sendall(b'PRINTTIME' + str(id))
38 |     time.sleep(0.1)
39 |     recv_buff = s.recv(1024)
40 |     recv_buff = recv_buff.strip('>\n')
41 |     print("Sensor " + str(id) + " GNSS solved time: " + recv_buff)
42 |     time.sleep(0.9)

```

E.1 Logger setup schematic



Appendix F

E-mails

F.1 Correspondence with Mr. Davis

Hi Aril,

This would be fine, but you may want to take a look at the Astropy library and see if their time package would meet your needs. It's certain to be more robust and well tested. But if you'd like to use my module, please do. <http://docs.astropy.org/en/stable/time/index.html>

Best,
Matt Davis

On Sun, Oct 23, 2016 at 2:25 PM Aril Schultzen <aschultzen@gmail.com> wrote:

```
> Hi!  
>  
> I am currently writing my master thesis in compsci and I wanted to ask you  
> if it was OK if I used your library for converting dates to/from JD and MJD  
> (https://gist.github.com/jiffyclub/1294443) in my implementation? It  
> will be used to convert time to MJD for a model and also for stamping logs.  
> Your work will of course be acknowledge as your own.  
>  
> Kind regards  
>  
> Aril Schultzen  
>
```

Appendix G

Figures



Figure G.1: Photograph of the system used to measure the 10 MHz output from the atomic clock. Not in the picture is the source of the 10 MHz reference.



Figure G.2: Antenna covered in aluminium foil to simulate a jamming attack.

Complete Bibliography

- [1] Grace Xingxin Gao Liang Heng Daniel Chou. “Reliable GPS-Based Timing for Power Systems: A Multi-Layered, Multi Receiver Architecture”. In: *InsideGNSS* (Nov. 2014), p. 12.
- [2] Yilu Liu et al. “State Estimation and Voltage Security Monitoring Using Synchronized Phasor Measurement”. In: (2001).
- [3] Daniel P. Shepard, Todd E. Humphreys, and Aaron A. Fansler. “Evaluation of the Vulnerability of Phasor Measurement Units to GPS Spoofing Attacks”. In: (2012).
- [4] Todd E. Humphreys Daniel P. Shepard and Aaron A. Fansler. “Going Up Against Time”. In: *GPS World* (2012).
- [5] Navigation National Coordination Office for Space-Based Positioning and Timing. *Timing*. <http://www.gps.gov/applications/timing/>. Accessed: 16-04-2015.
- [6] Steven Johnson. *Where good ideas come from, the natural history of innovation*. New York: Riverhead Books, 2010.
- [7] European Space Agency. *GPS Receivers*. http://www.navipedia.net/index.php/GPS_Receivers. Accessed: 16-04-2015.
- [8] Navigation National Coordination Office for Space-Based Positioning and Timing. *Space Segment*. <http://www.gps.gov/systems/gps/space/>. Accessed: 16-04-2015.
- [9] National Instruments. *GPS Receiver Testing*. http://www.insidegnss.com/special/elib/National_Instruments_GPS_Rx_Testing_tutorial.pdf. Accessed: 16-04-2015.
- [10] Navigation National Coordination Office for Space-Based Positioning and Timing. *Trilateration Exercise*. <http://www.gps.gov/multimedia/tutorials/trilateration/>. Accessed: 21-04-2015.

- [11] Carlene Stephens and Maggie Dennis. “Engineering time: inventing the electronic wristwatch”. In: *British Journal for the History of Science* 33 (2000), pp. 477–497. URL: www.ieee-uffc.org/main/history/step.pdf.
- [12] Carl R. Nave. *HyperPhysics*. <http://hyperphysics.phy-astr.gsu.edu/hbase/acloc.html>. Accessed: 16-04-2015.
- [13] Trimble Navigation Limited. *Trimble GPS Tutorial Getting perfect timing*. http://www.trimble.com/gps_tutorial/howgps-timing2.aspx. Accessed: 18-05-2015.
- [14] Xichen Jiang. “SPOOFING GPS RECEIVER CLOCK OFFSET OF PHASOR MEASUREMENT UNITS”. In: (2012). URL: tcipg.org/sites/default/files/papers/2012_Q2_CPR2.pdf.
- [15] National Marine Electronics Association. *NMEA 0183 standard*. http://www.nmea.org/content/nmea_standards/nmea_0183_v_410.asp. Accessed: 23-10-2016.
- [16] Microsemi Corporation. *SA.45s Chip-Scale Atomic Clock 098-00055-000 User Guide*. URL: http://www.microsemi.com/document-portal/doc_view/133467-sa-45s-chip-scale-atomic-clock-user.
- [17] Symmetricom. *SA.45s CSAC Data sheet*. <http://www.chronos.co.uk/files/pdfs/mic/sa.45s.pdf>. Accessed: 2-7-2015.
- [18] Raspberry Pi Foundation. *Raspberry Pi 3 Model B*. Accessed: 26-9-2016. URL: <https://www.raspberrypi.org/products/raspberrypi-3-model-b/>.
- [19] Mike Thompson and Peter Green. *Raspbian*. Accessed: 26-9-2016. URL: <https://www.raspbian.org/>.
- [20] u-blox. UBX-16000801 - R03. Accessed: 2-7-2015. URL: [https://www.u-blox.com/sites/default/files/products/documents/NEO-LEA-M8T-FW3_ProductSummary_\(UBX-16000801\).pdf](https://www.u-blox.com/sites/default/files/products/documents/NEO-LEA-M8T-FW3_ProductSummary_(UBX-16000801).pdf).
- [21] u-blox. UBX-13003221 - R11. Accessed: 18-10-2016. URL: [https://www.u-blox.com/sites/default/files/products/documents/u-blox8-M8_ReceiverDescrProtSpec_\(UBX-13003221\)_Public.pdf?utm_source=en/images/downloads/Product_Docs/u-bloxM8_ReceiverDescriptionProtocolSpec_\(UBX-13003221\)_Public.pdf](https://www.u-blox.com/sites/default/files/products/documents/u-blox8-M8_ReceiverDescrProtSpec_(UBX-13003221)_Public.pdf?utm_source=en/images/downloads/Product_Docs/u-bloxM8_ReceiverDescriptionProtocolSpec_(UBX-13003221)_Public.pdf).
- [22] Marco Cesati Daniel P.Bovet. *Understanding the Linux Kernel: From I/O ports to process management*. 3rd ed. O’Reilly. ISBN: 0-596-00213-0.

- [23] The GNOME Project. Accessed: 02-10-2015. URL: <https://wiki.gnome.org/Home3>.
- [24] Andre M. Rudoff W. Richard Stevens Bill Fenner. *Unix Network Programming: The sockets Networking API*. 3rd ed. Vol. 1. Addison-Wesley.
- [25] Kazutomo Yoshi. <http://www.mcs.anl.gov/~kazutomo/list/index.html>. Accessed: 15-10-2015.
- [26] Norgeskart.no. Accessed: 30-10-2016. URL: <http://www.norgeskart.no/?sok=Fetveien%2099,%20skedsmo#15/278837/6655264/+hits>.
- [27] M. Kerrisk. *The Linux Programming Interface*. No Starch Press, 2010. ISBN: 9781593272913. URL: <https://books.google.no/books?id=Ps2SH727eCIC>.
- [28] *The Kernel Bug Tracker*. Accessed: 02-03-2016. URL: https://bugzilla.kernel.org/show_bug.cgi?id=8691.
- [29] Inc Free Software Foundation. URL: https://gcc.gnu.org/gcc-4.6/cxx0x_status.html.
- [30] ISO. *Rationale for International Standard Programming Languages C*. URL: <http://www.open-std.org/jtc1/sc22/wg14/www/docs/C99RationaleV5.10.pdf>.
- [31] Inc USB Implementers Forum. *Universal Serial Bus Type-C Connectors and Cable Assemblies Compliance Document*. URL: http://www.usb.org/developers/compliance/usbcpd_testing/USB_Type-C_Compliance_Document_rev_1_1.pdf.
- [32] Matt Davis. *Functions for converting dates to/from JD and MJD*. Accessed: 02-09-2016. URL: <https://gist.github.com/jiffyclub/1294443>.