Stian Kristoffer Endresen

A method for measuring temporal properties of uplink interference in satellite communication

June 2021

NTNU
Norwegian University of
Science and Technology

Master's thesis

Stian Kristoffer Endresen

NTNU
Norwegian University of
Science and Technology
Faculty of Information Technology and Electrical
Engineering
Department of Electronic Systems

Master's thesis

2021

Stian Kristoffer Endresen

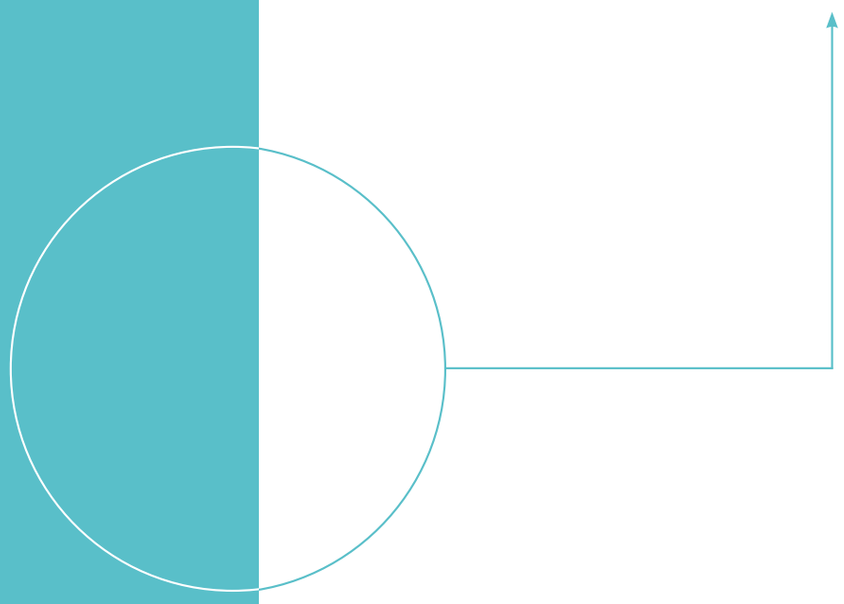

NTNU
Norwegian University of
Science and Technology

NTNU

Norwegian University of
Science and Technology

A method for measuring temporal properties of uplink interference in satellite communication

Stian Kristoffer Endresen

Master in Electronics Systems Design and Innovation
Submission date: June 2021
Supervisor: Nils Torbjörn Ekman
Co-supervisor: Gara Quintana Díaz

Norwegian University of Science and Technology
Department of Electronic Systems

Abstract

When communicating with small satellites in the *Ultra High Frequency* (UHF) band, a high packet loss rate is an issue due to severe radio interference. Thus, in order to design a communication system suited for satellite communication, it is necessary to have measurements of the interference environment. Publicly available previous measurements have shown high levels of interference, but these measurements are limited in scope, and often just display a heat map of the average interference power. In this report, a method is presented for measuring temporal properties of uplink interference, where the aim is to provide information useful for designing a communication system. It is shown that the output produced by the measurement method can be used to estimate certain other probability distributions that in turn can be used to make informed design choices when designing a communication system. The accuracy of these estimates is tested using interference measurements from the Norsat-2 satellite, and it is found that the estimate for one of the two distributions investigated is accurate. A software implementation that utilizes the presented method for measuring interference is developed, and the software is specifically designed to operate on the LUME-1 satellite. Furthermore, a mission is planned to measure interference in the UHF-band using the LUME-1 satellite. There is a distinct focus on the Arctic region in this mission, with the aim of using the interference measurements to make a system for retrieving data from Arctic sensor nodes with small satellites. A measurement configuration for how many measurements to perform and where to perform them is decided with the constraints of the LUME-1 satellite in mind, and the developed software is tested with this configuration to make sure it is mission-ready.

Sammendrag

Ved kommunikasjon med småsatellitter i UHF-båndet er høyt pakketap et problem grunnet sterk interferens. For å designe et kommunikasjonssystem egnet for satellittkommunikasjon er det derfor nødvendig å ha målinger av interferensmiljøet. Offentlig tilgjengelige, tidligere målinger, har vist høyt nivå av interferens. Disse målingene er derimot begrenset i omfang, og viser ofte kun et heatmap over gjennomsnittlig styrke på interferens. I denne rapporten presenteres en metode for å måle tidsmessige egenskaper ved interferens i opplink, der målet er å skaffe informasjon som er nyttig for å designe et kommunikasjonssystem. Det er vist at resultatet som produseres av målemetoden kan brukes til å estimere visse andre sannsynlighetsfordelinger, som igjen kan brukes til å ta informerte designvalg ved design av et kommunikasjonssystem. Nøyaktigheten til disse estimatene er testet ved bruke av interferensmålinger fra satellitten Norsat-2, og det er funnet at estimatet for en av de to fordelingene som er undersøkt er nøyaktig. En softwareimplementasjon som bruker den presenterte målemetoden for å måle interferens er utviklet, og softwaren er designet spesifikt for å operere på satellitten LUME-1. Videre er det planlagt et måleoppdrag for å måle interferens i UHF-båndet med satellitten LUME-1. I dette oppdraget er det et spesielt fokus på det Arktiske området, det målet er å bruke interferensmålingene til å lage et system som kan hente data fra arktiske sensornoder ved bruk av småsatellitter. En målekonfigurasjon for hvor mange målinger som skal utføres og hvor de skal utføres er bestemt med utgangspunkt i egenskapene til LUME-1, og den utviklede softwaren er testet med denne konfigurasjonen for å sikre at den er klar til bruk.

Contents

1. *Introduction*

1.1 Problem statement

The Arctic is an area of growing interest, and EU, NASA, and the Arctic Council state that they believe there will be a growing utilization of resources in this area [1, 2, 3]. This includes both commercial activities and research. However, communication in this region is a challenge due to the lack of infrastructure and satellite coverage [4]. This is particularly a challenge when using scientific sensor nodes that collect data from the Arctic and need to distribute it to researchers. In [5], the possibility of using small satellites to fill this communication gap is investigated and concluded to be a viable option.

There are, however, some challenges when uplinking data to a satellite in the Arctic region. One challenge is the presence of interference from radar systems that transmit powerful radio pulses to determine the speed, attitude, and direction of space objects, like satellites. Particularly, when a satellite crosses the north pole from the *United States* (US) to Russia, the Russian military radars need to verify that the satellite is not an American missile and vice versa. The radar pulses transmitted in these cases typically contain so much power that any communication during the pulses is lost due to a low signal-to-interference ratio. One way to work around this issue is to accept a certain packet loss rate and use a heavy coding scheme in the communication, which allows correcting bit errors and reconstructing lost packets. Energy is, however, a scarce resource in the Arctic due to the low availability of solar power during the winter [5], and one cannot frequently change batteries due to difficulty of access. The communication system must therefore be carefully designed to minimize the effective energy per bit, while still maintaining an acceptable data rate. Important decisions in this design process are the choice of transmit power and the coding scheme. A high transmit power means a high energy per bit, but if it is too low, it will result in loss of communication. Likewise, some coding is necessary to prevent loss of communication, but it increases the effective energy per bit. Selecting a suitable transmit power and coding scheme requires knowledge of the channel and the interference environment. However, measuring the interference with a satellite and downlinking the IQ signal directly is often infeasible due to a strict downlink budget, so another measuring method is needed. This report focuses on such a method.

In this report, we will select a method for characterizing the interference environment in such a way that the measurements can be used as a basis for designing a preliminary communication system adapted to the interference environment. The target satellite for in-orbit interference measurements in this thesis is the LUME-1 satellite. This is a 2U CubeSat developed within the Universidade de Vigo by the Alén Space team that was launched in December 2018. It is available through a collaboration with the University of Vigo and has an onboard *Software Defined Radio* (SDR) which enables uploading new software. In addition, a lab testbed with the same SDR is available for software testing. The measurement method shall utilize the onboard SDR to measure interference. To reduce energy usage and heat generation on the satellite, the method should not require heavy on-board processing, and the size of the data downlinked from the satellite should be minimized to cope with the slow downlink speed of ~75 kB/day. Furthermore, an implementation of the measurement method shall be developed and tested such that it is mission-ready, and can be uploaded to the LUME-1 satellite to perform in-orbit measurements. The UHF-band 435-438 MHz is an amateur band devoted to satellite communication, and we will therefore focus on measuring the interference environment in this band. Thus, the problem statement is given below:

A measurement method for measuring the radio interference properties in a UHF-band satellite communication system shall be developed and implemented, with the intent of using the measurements to provide

statistics that are useful for designing a communication system adapted to the interference environment. The measurement method shall produce measurements that are small in size to accommodate for a limited downlink rate, and the software implementation shall be tested such that it is ready for use in a satellite mission.

1.2 Report outline

A brief outline of what the different chapters in this report contains is given below. First, chapter 2 gives a description of previous work and measurements that are relevant to this report. Chapter 3 describes the UHF-band radio channel in satellite communication and investigates the properties of radar systems, which is a common source of interference. The properties and constraints for the LUME-1 satellite are given in chapter 4, as the measurement software is to be developed for this satellite. In chapter 5, some relevant points about the Norsat-2 satellite are given, since measurements from this satellite will be used for verification later in this report. Chapter 6 describes the measurement method selected for use on the LUME-1 satellite, and chapter 7 describes a software implementation of this method. Chapter 8 contains mathematical derivations of estimates for certain probability distributions based on the measurements that will be produced by the method described in chapter 6. The aim of this is to obtain information that is more directly useful in designing a communication system than the unprocessed measurements themselves. The validity of these estimates is tested using measurements from the Norsat-2 satellite in chapter 9. The effect of changing different measurement parameters is investigated in this chapter, such that suitable parameter choices are made for the measurement mission with the LUME-1 satellite. In chapter 10, the measurements from the Norsat-2 satellite are used to estimate the compression ratio one can expect with the selected measurement method, and a data budget for the LUME-1 satellite is established in order to determine how many measurements can be downlinked within a certain time frame. Different options for how to distribute these measurements in time and space are investigated in chapter 11, and which options to perform with the LUME-1 satellite are decided. In chapter 12, the developed software is tested with a configuration determined by the options selected in chapter 11, using a lab testbed with an identical SDR to the one on the LUME-1 satellite. After this, chapter 13 outlines how the measurements produced by the method described in chapter 6 can be utilized to design a communication system. Lastly, chapter 14 provides some discussion of the results achieved in this report, and chapter 15 lists what remains as future work. A final conclusion is given in chapter 16.

2. *Previous work*

2.1 Estimation of transmission windows

As the radio spectrum is a valuable resource, much work is carried out to better manage it. Users of unlicensed frequency bands need to deal with interference from other radio systems operating within those bands. One solution to this problem is to estimate when transmission opportunities are likely to present themselves and for how long. In [6], a transmitter-side technique is developed for detecting the presence of pulsed interference in 902.11 links, and to estimate temporal statistics of this interference. This is performed by varying the packet transmission duration and observing the corresponding change in packet loss rate to infer information about the timing of the interference pulses. In [7], a hardware demonstrator is described for an opportunistic radio system that adapts to its radio environment to find temporal transmission opportunities and predict the duration of future opportunity windows. This system also aims to minimize harmful interference caused by itself, for the primary users of the spectrum.

The above-mentioned research is focused mostly on creating an adaptive communication system rather than solely characterizing the interference, and the estimation technique, therefore, relies on information obtained by transmitting frames. In [8], however, a method for measuring the distribution of opportunity windows between interference pulses is presented. The method still aims to be used as a transmitter-side technique for estimating when to transmit frames, but it does so without needing to transmit radio packets. It can therefore be used to characterize an interference environment in order to solely increase insight into the channel, without the need for transmitting packets during the characterization. This method does, however, assume that the interference follows an exponential on/off process, which is not the case for periodic interference from radar systems.

2.2 Measuring interference in satellite communication

Some work is already carried out to characterize the interference in satellite communication. In [9], a preliminary noise measurement campaign is carried out to obtain measurements of the in-orbit interference for different geographic locations and carrier frequencies. These measurements are, however, not intended to give a complete overview of the geographic distribution for the interference, and do therefore only cover a small geographic area. In [10], worldwide heat maps of the in-orbit interference is obtained for the VHF band at 145.8 MHz-145.9 MHz, the UHF band at 435.9 MHz-437.3 MHz, the L-band at 1262 MHz-1267 MHz and the S-band at 2401 MHz. However, these measurements are performed with a satellite that has inclination ~55°, which results in no coverage of the Arctic areas. Additionally, the measurements only show the average interference level, with next to no focus on temporal statistics. In [11], the in-orbit interference is measured with the UWE-3 satellite, which has an inclination of 97.6°[12]. Due to this, the satellite measures the interference also over the Arctic region, which could provide useful information for designing a communication system suited for the Arctic. The interference is measured for 16 frequency bands with 200 kHz bandwidth each, in the UHF-band at 435 MHz-438 MHz. These measurements provide some information about what frequency bands contains the most interference but have little focus on the temporal statistics.

In [13, 14], a measurement method for measuring the power and time variability in interference is developed and tested. This method analyzes interference in sub-bands in the frequency domain separately and can estimate some statistics about the stationarity of signals. Therefore, it can identify some time-frequency

characteristics of the interference. The measurement software is uploaded to the LUME-1 satellite, and some measurements have already been carried out with it.

The measurement method described above provides measures that are useful for understanding the interference environment, but it only gives a limited understanding of the temporal characteristics of the interference. In [15], a second measurement method is proposed and investigated, as a continuation on the work carried out in [13], and to complement those results. This method produces a measure that is labeled as the *opportunity distribution*, which obtains more insight into temporal characteristics of the interference within a specific bandwidth. The measurement method is intended to be uplinked to the LUME-1 satellite, similarly to the measurement method in [13]. It is shown in [15] that if one measures the opportunity distribution for a pulsed interference signal, it is possible to derive its pulse repetition frequency and duty cycle from the measurement with precision within 5%, although the precision depends on the resolution used in the opportunity distribution. It can, however, be difficult to use the opportunity distribution to infer information about what kind of interference signal is measured if the signal is composed of several components, or if one has not seen the opportunity distribution for a similar signal before. This report will continue to investigate the opportunity distribution and show how it can be used to provide information that can be directly used to design a communication system, even with no prior knowledge of the type of interference signal measured. This report is a continuation of the work in [15].

3. *Radio Channel*

3.1 Satellite communication in the UHF-band

The UHF-band 435-438 MHz is devoted to satellite communication as an amateur band, and no international regulations limit the allowable bandwidth of a system operating within this band [16, 17]. There are, however, strong recommendations by the *International Amateur Radio Union* (IARU), and the band plan for Region 1 by the IARU recommending a maximum bandwidth of 20 kHz [18]. The Norwegian band plan by the radio amateur organization *Norsk Radio Relæ Liga* (NRRL) also specifies a maximum bandwidth of 20 kHz in this band [19]. If a 20 kHz bandwidth is used in satellite communication, the shortest possible symbol length is 0.05 ms, assuming zero excess bandwidth and no guard band for simplicity. For a communication system to be able to transmit data during an opportunity window, the window should be considerably longer than a single symbol in time. We are most interested in measuring transmission windows that are long enough to allow for transmitting a frame.

3.2 Radar systems

Radio pulses from radar systems are one of the main interference sources in the UHF band for Arctic satellite communication, and it is, therefore, necessary to know what behavior to expect from these systems. A recommendation from the *International Telecommunications Union* (ITU) for some properties of space-tracking ground radars is shown in table 3.1 [20]. As this is merely a recommendation, radar systems might have parameters deviating from this. However, the ITU recommendation still provides a baseline for what type of signals to expect from actual systems.

Radar systems designed for long-range monitoring need to have a slow pulse repetition frequency for the pulses to propagate to the target and for the echo to return to the receiver. Short-range radar systems, on the other hand, can use faster pulse repetition frequencies, as the propagation time is shorter. Additionally, long-range radar systems require higher peak power as more power is lost to free space loss, and in [21], it is estimated that doubling the peak power increases the range by about 25%. Therefore, radar interference with a slow pulse repetition frequency is expected to come from long-range systems, and thus have a higher peak power than interference from systems using faster pulse repetition frequencies.

The time between two consecutive radar pulses is a window that represents a transmission opportunity for a communication system. Thus, it is interesting to calculate the longest and shortest transmission window one can have between two radar pulses in interference from a system that follows the recommendations in table 3.1. The shortest possible window length is found by considering the fastest recommended pulse repetition frequency, which is 3 kHz for radar C. Using the largest duty cycle of 10%, this corresponds to a window duration of $\frac{1-0.10}{3\text{ kHz}} = 0.3$ ms. The longest possible window length is likewise found by considering the slowest recommended repetition frequency. The lowest listed bound on the pulse repetition frequency is 15 Hz for radar B, so this will be used for estimation purposes. Using the smallest duty cycle of 1% results in a window duration of $\frac{1-0.01}{15\text{ Hz}} = 66$ ms. It is worth remembering, however, that since the ITU recommendation states no lower bound on the pulse repetition frequency for type A radar systems, there may exist systems that transmit pulsed interference with even longer windows between pulses.

Parameters	Radar A	Radar B	Radar C
Application	Space object tracking	High altitude surveillance	Surface and air search
Peak RF power output [MW]	1-5	0.3	0.01
Polarization	Circular	Circular	Circular
Pulse duration [ms]	0.25, 0.5, 1, 2, 4, 8	0.01-16	0.001-1
Duty cycle (average) [%]	25	1-25	1-10
Pulse frequency modulation	Search: 100-350 kHz chirp. Track: 1 or 5 MHz linear chirp	2 MHz linear chirp	1 or 0.3 MHz linear chirp
Pulse repetition frequency	up to 41 Hz	15-400 Hz	100-3000 Hz
RX Radiofrequency bandwidth		30 MHz	30 MHz
RX Intermediate frequency bandwidth	1 or 5 MHz depending on chirp width	2 MHz	30 MHz
Antenna beamwidth in azimuth	2.2°	1.8° typical	80°
Antenna beamwidth in elevation	2.2°	1.8° typical	60°

Table 3.1: ITU recommendations for characteristics of ground radars in the 420-450 MHz range [20].

4. *LUME satellite*

Because the target satellite for in-orbit interference measurements in this thesis is the LUME-1 satellite, it is necessary to know the properties and constraints of this satellite. The LUME-1 satellite is a 2U CubeSat developed by Universidade de Vigo and the Alén Space team in collaboration with other organizations, and was launched in December 2018. Its mission ended in 2019, and it is now available to use for other research and experiments as part of a collaboration between *Norwegian University of Science and Technology* (NTNU) and Universidade de Vigo. The satellite is equipped with a SDR, which is called TOTEM. The software in Totem can be updated in-flight, and it will be used for performing interference measurements with the method described in section 6. The mission requirements for this mission, as well as its constraints related to the equipment, are listed below in sections 4.1 and 4.2. These requirements and constraints are modified from [14].

4.1 Mission requirements

The mission requirements for the LUME-1 mission are as follows:

- MR01: Time required to upload the measurement software to the spacecraft shall not exceed 10 days.
- MR02: The spacecraft shall enable updating of the software and its parameters.
- MR03: The spacecraft shall be able to downlink the measured data.
- MR04: Time required to downlink a measurement campaign shall not exceed 14 days.
- MR05: The variation in antenna gain due to pointing mismatch shall not exceed 3 dB within the duration of one measurement.
- MR06: The spacecraft shall provide timestamps for each measurement.

4.2 Available equipment and constraints

4.2.1 Constraints

Some constraints related to the LUME-1 spacecraft are as follows:

- SC01: The spacecraft is tumbling at ~1 rpm.
- SC02: *Attitude Determination And Control System* (ADCS) is not active.
- SC03: Antenna is not omnidirectional, as shown in figure 4.1.
- SC04: The UHF antenna is only circularly polarized when seen from the top and bottom.
- SC05: The satellite has a radio amateur license to transmit in the 437 MHz band.
- SC06: Downlink rate is 4.8 kbps. Considering protocol overhead and retransmission delays, an estimate of 1 kbps will be used.
- SC07: Effective uplink rate is previously estimated to be ~200 bps.

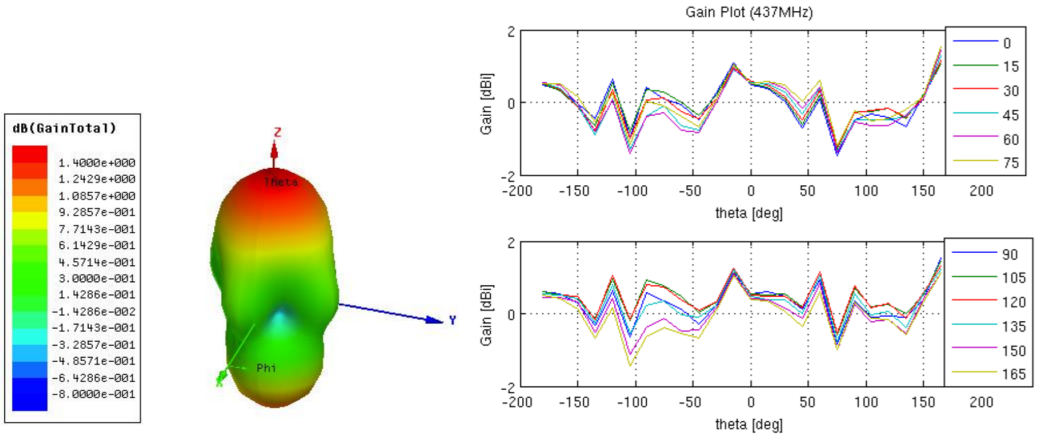

Figure 4.1: Antenna diagram for 2U antenna, simulated by Gomspace [22].

- SC08: Up- and downlink are only available when the spacecraft passes over the ground station in Vigo. Assumed 5 minutes per day.
- SC09: Satellite has an orbital period of 94.4 minutes [23].
- SC10: Measurement software should not run continuously over multiple orbits.

To elaborate on the last constraint, this limits the duration of any single program call to measurement software. Measuring in multiple consecutive orbits is still possible, by scheduling a new call to the measurement software every orbit. This is to reduce power consumption, as too much power is consumed even if the software runs in idle mode between measurements.

Some constraints related to the Totem SDR are:

- TC01: Totem SDR can only be on 12 of the 24 hrs in a day due to power budget. The on-time can be distributed throughout the day.
- TC02: Totem SDR has a 30 sec startup time.
- TC03: Totem SDR can measure in the 435-438 MHz band [24].
- TC04: Totem SDR has a dynamic range of approximately 66 dB.
- TC05: The RF bandwidth cannot exceed 200 kHz-56 MHz [24].
- TC06: Totem SDR enables sample rates of 512 kSps-56 MSps [24]. Software limitations reduce the upper bound.
- TC07: Automatic gain control in Totem cannot be read in real-time.
- TC08: Measurement software must be developed using C, C++, or GNURadio libraries.

Universidade de Vigo owns a ground station that will be used to uplink the software and commands to the satellite, as well as downlink measurement results. Some properties of the ground station are:

- GS01: Radioamateur license in the 437 MHz band.
- GS02: Ground station uses a USRP that can be programmed using GNURadio.

4.2.2 Measurement software size

Using the uplink bit rate estimate from SC07 of 200 bps, available in total 5 minutes each day, we get an average throughput of 7.5 kB/day. As MR01 states that the uplink time shall not exceed 10 days, the size of the measurement program shall thus not exceed 75 kB.

4.2.3 Measurement data size

The downlink rate is 4.8 kbps, but this does not account for protocol overheads or retransmissions. An estimate of 1 kbps is used instead, as stated in SC06. Given the assumption that communication to the satellite is available 5 minutes every day, this provides a daily throughput of 37.5 kB/day. As MR04 states that the time for downlinking a measurement campaign shall not exceed 14 days, the maximum size of a measurement campaign is 525 kB.

5. *Norsat-2 satellite*

In order to investigate the opportunity distribution as a measurement method, IQ signals to analyze are required for testing. It is advantageous if these signals resembles the in-orbit interference that will be found in the LUME-1 mission. Interference signals measured with the Norsat-2 satellite [25] are available for this purpose, and are suited for this use because it is interference measured in-orbit. The Norsat-2 satellite has measured with a carrier frequency of $f_c = 157\,\text{MHz}$, which is different from the $f_c = 435\,\text{MHz}$ carrier that the LUME-1 satellite will use, so the measured interference may differ from that in the UHF-band. However, regardless of this difference, the measured interference can still be used for testing the opportunity distribution as a measurement method.

The Norsat-2 satellite was launched on July 14th, 2017 into a polar sun-synchronous orbit for complementing the Norwegian *Automatic Identification System* (AIS) satellite network that monitors maritime traffic. It is also one of the first satellites to support *VHF Data Exchange System* (VDES)-services. As a collaboration between NTNU and Space Norway in a national project funded by the Norwegian Space Agency, in-orbit radio measurements from Norsat-2 have been obtained. This measurement data will be used to test the measurement method implemented in this report. A description of the Norsat-2 mission is given below. Note that the measurements from the Norsat-2 satellite will only be used as test signals to verify the measurement method, and it is not planned to use results from this report to influence future measurements on the Norsat-2 satellite.

The measurements from the Norsat-2 satellite were obtained While passing over Pisa in Italy, as the satellite measured the received radio signal and downlinked the raw IQ data. This process was repeated seven times, referred to as seven sessions. While the satellite was measuring the radio signal, it was simultaneously receiving VDES signals from a ground station, which was pointing to the satellite. Thus, the measurements contain this VDES signal in addition to other interference.

In each of the seven sessions, the measurements were obtained using a sampling rate of $f_s = 134\,400$ Hz, and a bit depth of 8 bits per I and Q sample. The carrier frequency was set to $f_c = 157\,312\,500 Hz$, which is the same carrier frequency that was used for transmitting the VDES signal. The *Analog to Digital Converter* (ADC) in the satellite was configured to use the maximum gain setting, but with a step attenuator that reduces the gain (default value of 6 dB). The I- and Q part of the measured radio signal was quantized as values between -128 and 127, and the power P that corresponds to a received signal amplitude can be calculated in dBm as

$$P = 10 * \log_{10}\left(\frac{V^2}{resistance}\right) + 30 - gain + attenuation, \tag{5.1}$$

where $gain = 44.65\,\text{dB}$, $attenuation$ is the attenuation of the step attenuator, $resistance = 50\,\Omega$, and V is the signal voltage, given as

$$V = \sqrt{I^2 + Q^2} \cdot step, \tag{5.2}$$

where $step = 6.0316e - 7$, and I and Q are the values of the in-phase and quadrature part of the measured signal. Combining this, the signal effect in dBm is given as

$$
\begin{aligned}
P &= 10 \cdot \log_{10} \left(\frac{I^2 + Q^2}{50\,\Omega} \cdot (6.0316 \cdot 10^{-7})^2 \right) + 30 - 44.65 + 6 \\
&= 10 \cdot \log_{10} \left(I^2 + Q^2 \right) + 10 \cdot \log_{10} \left(\frac{(6.0316 \cdot 10^{-7})^2}{50\,\Omega} \right) - 8.65 \\
&= 10 \cdot \log_{10} \left(I^2 + Q^2 \right) - 150
\end{aligned}
\tag{5.3}
$$

6. *Measurement algorithm*

Chapter 2 states that this report shall continue investigating the measurement method labeled the *opportunity distribution*, which was developed in [15]. This chapter describes how the measurement method works, and how the measurement result can be interpreted.

6.1 Opportunity distribution

The measurement method produces a single, independent measurement for a short period of time. The first step in performing measurement is to sample a radio signal $s[n]$ from the available SDR, with a duration, sampling rate, RF bandwidth, and carrier frequency specified by input parameters. Afterwards, the squared amplitude of the complex baseband signal is calculated as a function of time, $P[n] = |s[n]|^2$. The opportunity distribution of the measured signal is calculated from $P[n]$.

The opportunity distribution is a measure of how much time the signal spends in opportunity windows of different lengths. This measure is essentially the same as what is referred to as the *distribution of opportunity time* in [8]. Given a signal $P[n]$ and a power threshold P_i, we define an opportunity window as an interval where $P[n]$ is continuously below that threshold. We then measure how much time the signal spends in opportunity windows with durations between 0.5 ms and 1 ms, for example. The time spent in windows with duration between 1 ms-2 ms is also measured, and 2 ms-4 ms, and so on up to an upper interval of e.g. 8 ms-16 ms. One would then have measurements from the time bins $\{0.5\,\text{ms} - 1\,\text{ms}, 1\,\text{ms} - 2\,\text{ms}, 2\,\text{ms} - 4\,\text{ms}, 4\,\text{ms} - 8\,\text{ms}, 8\,\text{ms} - 16\,\text{ms}\}$. This procedure is then repeated for a range of power thresholds $\{P_i\}$ for all the time intervals. The opportunity distribution is thus an $M \times N$ grid of measurements, where M is the number of power thresholds and N is the number of time intervals for the opportunity durations. If there are opportunity windows with a duration shorter than the shortest measurement interval, they are grouped into the shortest interval, and vice versa for any opportunities longer than the longest interval. Thus, the time bin 0.5 ms - 1 ms is more precisely everything below 1 ms, and the bin 8 ms-16 ms is more precisely everything above 8 ms. The durations 0.5 ms, 1 ms, 2 ms, 4 ms, 8 ms, and 16 ms are chosen solely for example purposes, and the time bins actually used are discussed in section 6.2, along with what power levels $\{P_i\}$ to use.

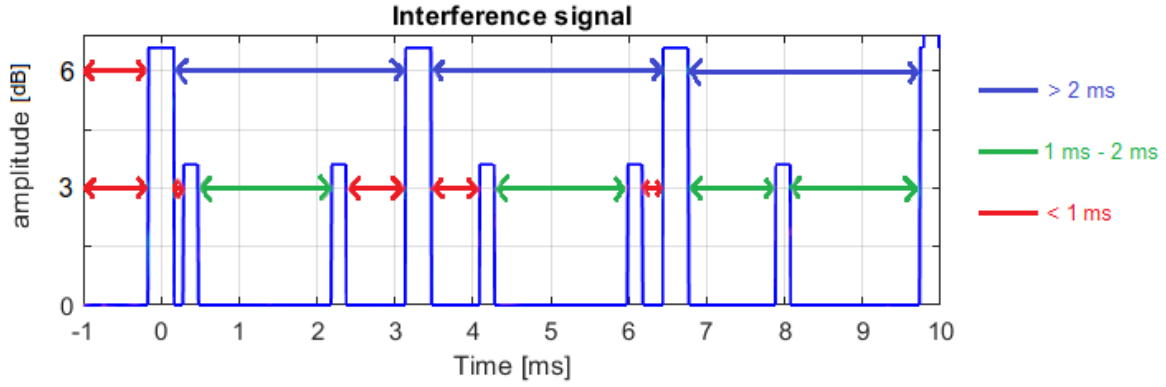

Figure 6.1: Duration of opportunity windows for pulse trains.

Figure 6.1 shows an example signal, where the interference consists of two pulse trains that are measured simultaneously. The opportunity windows for power levels $P = 3$ dB and $P = 6$ dB are marked with colors corresponding to which duration interval they will be categorized into. For power level $P = 6$ dB, the opportunity distribution will show that most of the time is spent in opportunity windows with duration longer than 2 ms, and very little in opportunity windows with duration shorter than 1 ms. Likewise, for the power level $P = 3$ dB, the opportunity distribution will show that only opportunity windows shorter than 2 ms are present and that most of this time is spent in opportunity windows with durations between 1 ms and 2 ms. Another example, with an illustration of the opportunity distribution is shown in section 6.2.1.

To describe the opportunity distribution in a more rigorous way, we define that for a power level P_i and a duration interval $[d_j, d_{j+1}]$, the opportunity distribution denotes how much time in total is spent in opportunity windows with duration in the interval $[d_j, d_{j+1}]$, where an opportunity window is defined as a time interval where the signal is continuously below the power level P_i. The amount of time spent is represented as a fraction of the signal duration. That is, if the measured interference signal lasts for 5 seconds and the time spent in opportunity windows with a duration within the interval $[d_j, d_{j+1}]$ accumulates to a total of 1 second, the opportunity distribution will take the value $\frac{1}{5}$.

An interesting characteristic about the opportunity distribution is that the sum of its value over all time intervals is the *Empirical Cumulative Distribution Function* (ECDF) of the radio signal sampled at the power levels P_i. This is because the sum, at a given power threshold, represents how large fraction of the signal is spent below that threshold, which in turn is equal to the ECDF.

The opportunity distribution as described above only considers the temporal properties of the interference within a given bandwidth, and does not offer any information about spectral properties. This is especially disadvantageous if the interference is concentrated in a narrow frequency band, as it would then be more important to avoid transmitting in that bandwidth rather than knowing the temporal properties of the interference. A possible solution to this would be to calculate a spectrogram of the measured interference and measure the opportunity distribution for each of the bins in the spectrogram. This would provide some spectral resolution, but would greatly increase the data size of the measurements. Furthermore, there may be little difference between bordering frequency bands if the interference stems from powerful military radars that produce so strong interference that saturates the ADC in the radio receiver at the satellite. Therefore, it is decided that the opportunity distribution is to be computed directly from the IQ signal and not the spectrogram, providing no spectral resolution.

6.2 Measurement parameters

When calculating the opportunity distribution, one must decide what time bins to calculate it for. We will refer to the shortest time bin limit as $d_{w_{min}}$ and the longest as $d_{w_{max}}$. Each interval limit in the time bins will be a multiple of the previous one, and the scaling factor between the durations is referred to as s. In the example in section 6.1, the shortest window length was $d_{w_{min}} = 0.5\,\text{ms}$, the longest was $d_{w_{max}} = 16\,\text{ms}$, and the scaling factor between the durations was $s = 2$. When measuring the opportunity distribution, we need to define $d_{w_{min}}$, $d_{w_{max}}$ and the scaling factor s. Some notes regarding the choice of these parameters are given below.

In chapter 3.2, estimates are made showing that the ITU recommendations give rise to systems with opportunity windows of duration between 0.3 ms and 66 ms. Additionally, the length of a symbol used in UHF satellite communication is 0.05 ms at least. Because the duration of a transmission window must be longer than a single symbol in time to allow for transmitting a frame, the shortest transmission window we will measure for, $d_{w_{min}}$, should not be smaller than 0.05 ms. One might even set $d_{w_{min}}$ larger than this, e.g. 0.1 ms. As for the upper limit, $d_{w_{max}}$, the longest duration that can arise between two radar pulses according to the ITU recommendations is estimated to ~66 ms, so $d_{w_{max}}$ should be longer than this to ensure that this duration is covered. The 66 ms is, however, derived from choosing 15 Hz as the lowest possible pulse repetition frequency of a radar system. Because the ITU recommendations in table 3.1 put no lower limit on the recommended pulse repetition frequency for type A radars, systems might exist with slower pulse repetition frequencies than this, and thus longer opportunity durations. It is therefore difficult to determine an upper threshold for d_{max}. Regarding the scaling factor s between the durations

to measure for, it determines how precisely we can estimate the duration of an opportunity window. If we use a scaling factor of e.g. $s = 1.1$, the classification of an opportunity window into one of the chosen window duration intervals will result in a maximum error of e.g. 5% between the duration of the window and the midpoint of the duration interval.

The choice of $d_{w_{min}}$, $d_{w_{max}}$ and s will determine how many time bins is necessary for the duration dimension of the opportunity distribution, which is directly proportional to the size of the opportunity distribution in bytes. Choosing e.g. $d_{w_{min}} = 0.1$ ms, $d_{w_{min}} = 100$ ms and $s = 1.1$ causes the number of windows to be 74, as solving $d_{w_{min}} \cdot s^{num\_bins-1} = d_{w_{max}}$ yields $num\_bins = 74$. The limits for the time intervals to classify the window durations into would then be $\{0.1\,\text{ms}, 0.11\,\text{ms}, 0.121\,\text{ms}, \ldots\ 96\,\text{ms}, 105\,\text{ms}\}$. Note that because all the durations are separated by a scaling factor $s = 1.1$, $d_{w_{max}}$ is not exactly 100 ms, but rather $0.1\,\text{ms} \cdot 1.1^{74-1} = 105$ ms. Further investigation of appropriate choices for $d_{w_{min}}$, $d_{w_{max}}$ and s are provided in chapter 9, and the final choice to use in the LUME-1 measurements is shown in table 11.1 in chapter 11.

Concerning the power thresholds $\{P_i\}$ to use, the dynamic range of the TOTEM SDR is 66 dB from TC04. Thus, it makes sense to use 0 dB as the lowest power threshold and 66dB as the highest. Figure 4.1 shows that the gain of the antenna used on the LUME-1 satellite may vary with $\pm 1.5\ dB$ depending on which direction it is facing, and SC01 states that the spacecraft is tumbling at ~1 rpm. The variation in antenna gain puts a limit on how fine resolution it is sensible to use in the power domain, and since the antenna gain may vary by 3 dB, this is also chosen as the step size. The power thresholds then become the 23 levels {0 dB, 3 dB, 6 dB, ... 66 dB}. The mapping between these power levels and power in dBm depends on the hardware gain on the TOTEM SDR. However no power calibration is performed onboard the satellite, so one would only get relative power values.

6.2.1 Example of measured opportunity distribution

To better understand what the opportunity distribution illustrates, figure 6.2 shows an example of an opportunity distribution that was measured in [15], along with a snippet of the radio signal that the opportunity distribution was calculated from and the ECDF derived from the opportunity distribution. This was obtained by transmitting a radio signal in a lab testbed using an SDR and measuring it with a TOTEM SDR, before calculating the opportunity distribution from the measured signal. The radio signal that was transmitted was a pulse train with a pulse repetition frequency of 20 Hz and a duty cycle of 25%, and it was measured for a duration of 2 s. The measurement parameters used were: $d_{w_{min}} = 0.1$ ms, $d_{w_{max}} = 100$ ms and $s = 1.1$. A more detailed description of the measurement setup can be found in [15].

The opportunity distribution shows a clear vertical line from 24 dB to 57 dB at the time bin 36.8 ms-40.6 ms, where it takes a value of 73.1%. This tells us that there are many opportunity windows with a duration within this interval. The duration matches the duration between two consecutive pulses in the pulse train, which is $\frac{1-0.25}{20\,\text{Hz}} = 37.5$ ms. Furthermore, since the duty cycle of the pulse train is 25%, one would expect that 75% of the time in the radio signal is spent in the windows between the pulses. This can be seen from the opportunity distribution, by the fact that 73.1% of the time is spent in opportunity windows with a duration between 36.8 ms and 40.6 ms. The 1.9% deviation is likely because the two opportunity windows at the beginning and end of the radio signal are truncated to a shorter duration, and thus contribute to another time bin than 36.8 ms-40.6 ms. In fact, one can barely see from the distribution that some time is spent in opportunity windows with duration 11.7 ms-12.9 ms and some windows with duration 25.2 ms-27.7 ms, as there are faint non-zero contributions in the opportunity distribution in these duration bins. This is probably the durations of the two truncated windows at the beginning and end of the radio signal, and the contributions are solely due to these two windows.

As for the power level, we can see from the opportunity distribution that the pulse train has a peak power of 57 dB when the pulse is high. For power levels above this, e.g. 60 dB, the radio signal is always below that power level, and one can therefore consider the entire radio signal as a single 2 s long opportunity window. Thus, the opportunity distribution shows that 100% of the time is spent in that opportunity window, which has a duration longer than 100 ms. One can also see that the noise floor is at 24 dB or lower. At power levels below 24 dB, the duration of the opportunity windows is random and determined by how long the noise stays below a certain threshold. One can see this from the horizontal line at 24 dB, which shows that there are opportunity windows with all kinds of durations at this power threshold. For

the ECDF, we see that the distribution derived from the opportunity distribution overlaps perfectly with the actual ECDF calculated directly from the IQ-samples.

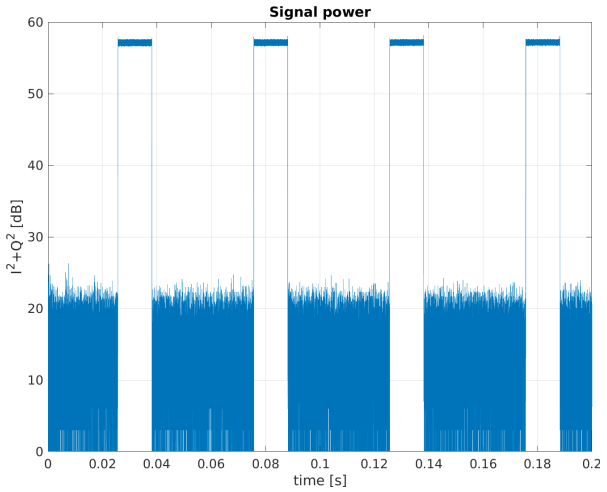

(a) Signal power in time

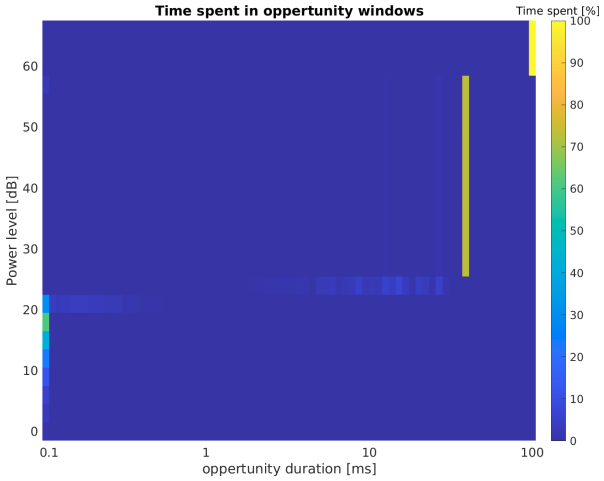

(b) Opportunity distribution

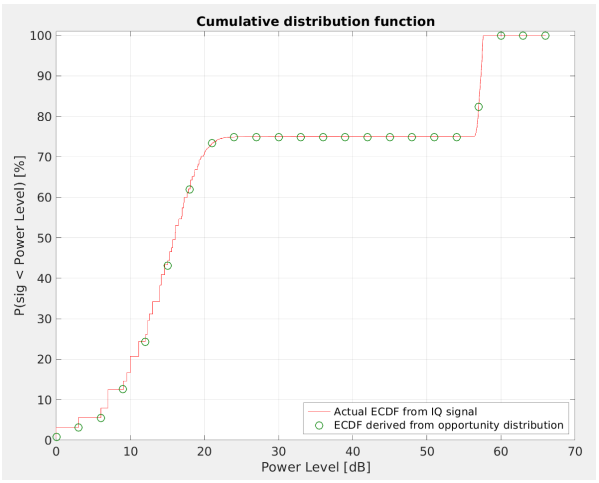

(c) Empirical Cumulative Distribution Function

Figure 6.2: Opportunity distribution measured from a pulse signal with pulse repetition frequency 20 Hz and duty cycle 25%.

7. *Software architecture*

Two software programs are developed to enable measuring the opportunity distribution on the LUME-1 satellite. The first program is called `campaign_manager.sh`, and the second one is called `analyze_signal`. The script `campaign_manager.sh` handles all the top-level scheduling of measurements, interactions with the SDR, and compression of measurements, while `analyze_signal` performs the signal processing needed to calculate the opportunity distribution from a radio signal. Since MR01 requires that the total size of the software uploaded to the satellite is below 75 kB, the programs are developed with this size limitation in mind.

To implement the software, we must decide what options the software should provide for how to distribute the measurements in a campaign in space and time. We define a measurement campaign as a collection of measurements, performed during multiple orbits. Several measurement blocks are computed in each orbit, where a measurement block is defined as a sequence of consecutive measurements with no added delay in between. A measurement is defined as a single opportunity distribution calculated from a radio signal received on the satellite. Figure 7.1 shows an illustration of how the measurements in a campaign may look like if there are four orbits, three measurement blocks per orbit, and ten measurements per measurement block. Each dot in this figure represents one measurement. Note that the figure only provides a visualization of how the measurements in a campaign are positioned in space, and the distance between measurements is not to scale. The spacing between the measurements within a measurement block is equal to the duration of the sampled radio signal that is used to calculate a single measurement.

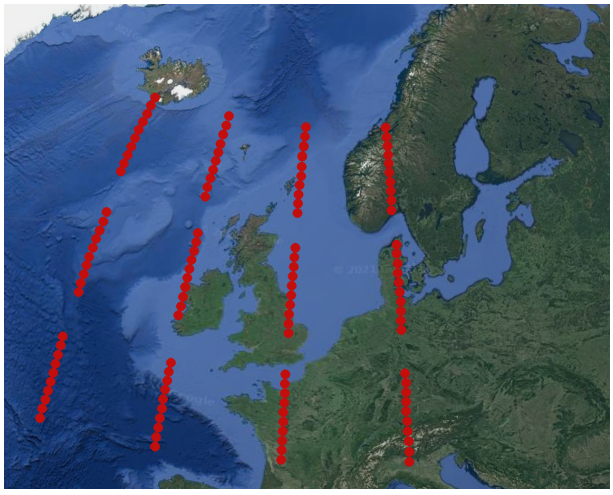

Figure 7.1: Illustration of the measurements in a campaign, where each red dot represents a measured opportunity distribution (not to scale).

As described in section 4.2.1, a program call to the measurement software should not run over multiple orbits. The measurement software is still developed with the option to measure for several orbits such that the software is flexible, and in case this feature is needed in a future mission.

7.1 `campaign_manager.sh`

The `campaign_manager.sh` program is implemented in such a way that a measurement campaign can be initiated with a single call to this program. Such a call is initiated when the user sends a command to the satellite, containing arguments that specify how to perform the measurements, how many measurements should be performed, and how often they should be performed. This command initiates a call to `campaign_manager.sh`, which schedules the measurements, saves the measurement results with timestamps, and compresses the measurements.

Figure 7.2 shows an outline of the steps performed by the `campaign_manager.sh` program. When this program is called, it first configures the Totem SDR by enabling the *Radio Frequency* (RF) front-end, setting the sampling frequency, carrier frequency, RF bandwidth, hardware gain, and reading a *Finite Impulse Response* (FIR) filter into memory. The FIR filter is needed because the SDR is normally limited to a minimum sampling rate of 2.08 MHz, and enabling an FIR filter enables the use of lower sampling rates. This filter must be uploaded to the satellite as a separate file, and must thus be included as a part of the data budget.

After the initial setup, a loop that runs once for each satellite orbit we want to measure in is initiated. This is the loop with the index `num_orbits` in figure 7.2. In each iteration of this loop, the measurement blocks for that orbit are obtained, and at the end of the orbit, the measurement files are compressed and stored on the satellite. Lastly, a delay is introduced to wait until the next orbit. Obtaining the measurement blocks for an orbit is done by running the loop with the index `num_meas_blocks` in figure 7.2. In each iteration of this loop, a measurement block is obtained, and then a delay is introduced to wait for when the next measurement block is to be obtained.

A measurement block is obtained when running the loop with the index `num_meas`. In each iteration in this loop, a single measurement is obtained. This is done by first sampling the radio signal from the SDR, which is performed by using a command-line utility from the libiio package from Analog Devices [26]. This utility outputs the read signal to a binary file, which is used by the `analyze_signal` program to calculate the opportunity distribution of the radio signal. The binary file containing the radio signal is then deleted, and the calculated opportunity distribution is saved with a name determined by the timestamp for when the radio signal was recorded. The signal processing in `analyze_signal` is executed as a background process, which enables the sampling of the radio signal for a measurement to start immediately after the radio signal for the previous measurement was sampled, without having to wait for `analyze_signal` to

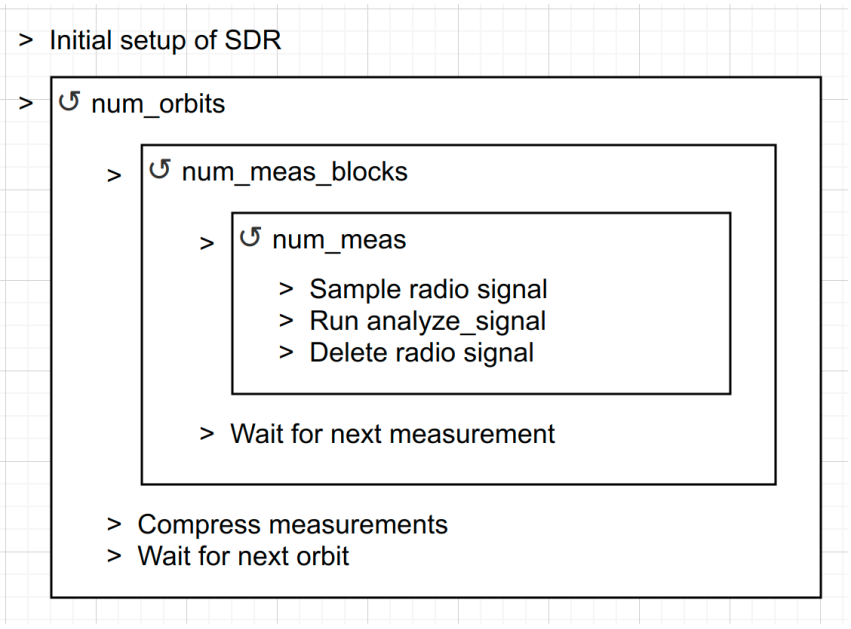

Figure 7.2: Illustration of the layout in `campaign_manager.sh`

finish.

`campaign_manager.sh` takes several input parameters to determine how to execute the measurement campaign. These parameters are listed in table 7.1 and must be specified for every call to `campaign_manager.sh`. Note that for the power levels $\{P_i\}$, they are relative power levels in dB, where 0 dB is the power of the lowest signal strength possible to measure with the Totem SDR.

Input parameter	Description
`f_s`	Sampling frequency for SDR
`f_c`	Carrier frequency for SDR
`bw`	Bandwidth of RF filter for SDR
`gain`	Hardware gain for SDR
`num_samples`	Number of samples to sample the radio signal for
`buffer_size`	Buffer size when sampling the radio signal. Recommended: 32768
`power_min_dB`	Lowest dB-value to calculate opportunity distribution for
`power_max_dB`	Highest dB-value to calculate opportunity distribution for
`power_step_dB`	dB step size between power levels to calculate opportunity distribution for
`time_min`	Shortest window duration to calculate opportunity distribution for, d_{min}
`time_max`	Longest window duration to calculate opportunity distribution for, d_{max}
`time_step`	Multiplicative scaling factor, s, between consecutive window durations to calculate the opportunity distribution for
`num_quantization_levels`	Number of quantization levels to quantize opportunity distribution into
`num_orbits`	Number of satellite orbits to measure in
`orbit_period`	Period of each satellite orbit in seconds
`num_meas_blocks`	Number of measurements to perform each orbit
`meas_block_period`	Period for how frequent to initiate a measurement block in seconds
`num_meas`	Number of measurements per measurement block
`program_path`	Directory to where the `analyze_signal` program and the FIR-filter is located
`output_path`	Directory where measurement results should be saved

Table 7.1: Input parameters to `campaign_manager.sh`

7.2 `analyze_signal`

`analyze_signal` is a program written in C++, which calculates the opportunity distribution of a sampled radio signal as described in section 6. The power levels $\{P_i\}$ and time bins to calculate the distribution for are determined by input parameters to the program. Likewise, the name of the binary file that contains the radio signal to analyze, and the filename to save the opportunity distribution to are also determined by input parameters. All the input parameters to `analyze_signal` are passed from `campaign_manager.sh`, and a comprehensive list of them are `f_s`, `input_file`, `output_file`, `power_min_dB`, `power_max_dB`, `power_step_dB`, `time_min`, `time_max`, `time_step` and `num_quantization_levels`.

When writing the measured opportunity distribution to an output file, 8 bits are used to represent each measurement point if `num_quantization_levels` <= 256, and 16 bits otherwise. Additionally, to reduce the file size, parts of the distribution are cropped out before saving it. Namely, the bins corresponding to power levels where no interference is present are removed. After this, the bins corresponding to durations for which there is no opportunity window with that duration or longer, are removed. For the example in figure 6.2b, the bins that will be removed are everything above 60 dB and everything to the right of 40.6 ms. Note that nothing is cropped away unless it can be perfectly reconstructed in a lossless manner.

7.2.1 Quantization

Before saving the opportunity distribution, it is rounded into discrete values, determined by the input `num_quantization_levels`. This is, however, not done by rounding the opportunity distribution directly. Instead, for each power bin, the *cumulative* opportunity distribution along the duration-axis is calculated from the opportunity distribution, and this cumulative opportunity distribution is quantized. Lastly, the quantized regular opportunity distribution is calculated from the quantized cumulative distribution and saved as an output file. The advantage with this approach to quantization is that the accumulated sum of the quantization errors in the opportunity distribution at a given power threshold will never be larger than $\frac{1}{\texttt{num\_quantization\_levels}}$. This also means that if one decides to derive the ECDF from the measurement, as described in section 6.1, the rounding errors will not accumulate to a greater error, and the error in the CDF due to the quantization will never exceed $\frac{1}{\texttt{num\_quantization\_levels}}$. Furthermore, avoiding the accumulation of rounding errors allows for estimating certain statistics with greater accuracy, as shown in section 9.2.

8. *Deriving packet loss from the opportunity distribution*

This section describes how the opportunity distribution can be used to derive the packet loss rate of a communication system. In addition to this, two other probability distributions that provide useful information for designing a communication system are estimated. The first of these distributions is the probability that a certain percentage of the payload data in a communication packet falls outside the *main opportunity window*. The main opportunity window is here defined as the window containing the header sequence for that communication packet. An illustration of this is shown in figure 8.1a. This figure shows a communication packet with a header sequence in the middle, marked by a hatched area, and gray interference pules. The payload data is defined as all parts of the packet that are not the header. The parts of the payload data that fall outside the main opportunity window are dotted. The second distribution is the probability that a certain percentage of the payload data is lost to interference. An illustration of this is shown in figure 8.1b, where the parts of the payload data that are lost to interference are dotted. This second distribution is the one that is most useful for designing a communication system, but in order to derive it, we need an estimate for the first distribution. We will also see that in order to derive the second distribution, one also needs some additional measurements besides the opportunity distribution.

To calculate the above-mentioned estimates, we will use the opportunity distribution to estimate how frequent opportunity windows with different window durations occur and look at how likely it is that a communication packet falls within each of these windows.

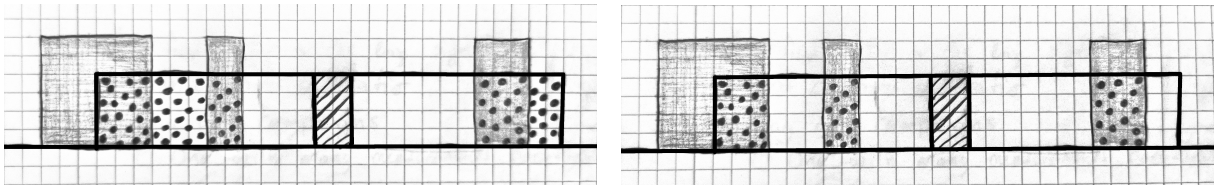

(a) The dotted area is the part of the payload data that falls outside the main opportunity window.

(b) The dotted area is the part of the payload data that is lost to interference.

Figure 8.1: Illustration of communication packet. The hatched area in the middle represent the header sequence, while the gray columns represent interference pulses.

8.1 Assumptions and notation

To calculate the packet loss rate and the above-mentioned distributions, we must first choose a model to use for the communication system. We will assume that a communication packet is transmitted at a random time, and calculate the probability that this packet is hit by interference. As shown in figure 8.2, we will model the communication packet as containing two parts: a header sequence and payload data, where the header sequence is located in the middle of the packet. The header sequence is assumed to be located in the middle because if it is used to estimate channel properties like phase drift, the channel estimation is most accurate close to the header. The duration of the entire packet is denoted as d_c, and the duration of the header as d_h. We define a power threshold P, and say that if any interference occurs

during the duration of the communication packet, and this interference has power equal to or larger than P, we will count it as interfering with the packet. If any interference occurs within the frame header, we count the entire communication packet as being lost, because the receiver is then unable to detect the packet. This is equivalent to the entire packet in figure 8.1 being dotted out.

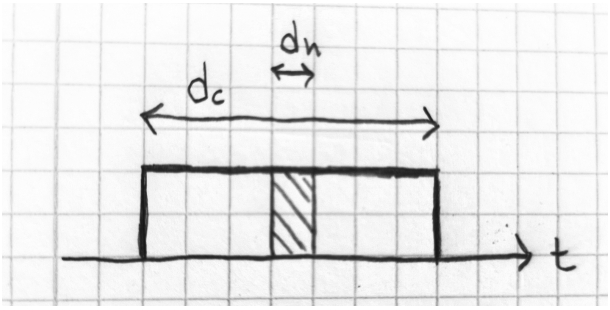

Figure 8.2: A single communication packet in time. The header is marked with diagonal lines.

As described in section 6, the opportunity distribution actually shows how likely the signal is to be in an opportunity window with duration within some interval, $[d_j, d_{j+1}]$, where the size of the interval was defined by a multiplicative scaling factor $s = \frac{d_{j+1}}{d_j}$. In this chapter, however, we will approximate all these opportunity windows as having the same duration, $d_w = \frac{d_j + d_{j+1}}{2}$. The opportunity distribution is thus defined for a discrete sequence of window durations $\{d_w\}$, and we will denote the shortest of these durations as $d_{w_{min}}$, and the longest as $d_{w_{max}}$. We will use the notation $\sum_{d_w=d_{w_{min}}}^{d_{w_{max}}}$ when we want to sum over all durations. We denote the opportunity distribution as $OppDist[P, d_w]$. $OppDist[P, d_w]$ represents how likely the signal is, at any given time, to be in an opportunity window with duration d_w, where an opportunity window is defined as a time interval where the signal is continuously below the power level P.

8.2 Probability that no interference occurs during the communication packet

The probability that there is no interference in the packet at all is equal to the probability that the entire packet falls within a single opportunity window. This can be derived from the opportunity distribution. First, the expected number of opportunity windows with duration d_w that occur per second, is

$$N_w = \frac{OppDist[P, d_w] \cdot 1\,\mathrm{s}}{d_w}. \tag{8.1}$$

For each of these windows, we must find the probability that the entire communication packet falls inside that window. Figure 8.3 shows an illustration of this, and we see that if the entire packet is to fall inside the window, there is an interval with duration $d_w - d_c$ in which the packet has to start. Because there are N_w windows every second, the probability that the packet falls inside one of these windows, is

$$\begin{aligned} p_{fall\_within}(d_c, d_w) =& \frac{(d_w - d_c) \cdot N_w}{1\,\mathrm{s}} \\ =& (d_w - d_c)\ OppDist[P, d_w]/d_w. \end{aligned} \tag{8.2}$$

Since this is the probability that a communication packet falls within an opportunity window with the specific duration d_w, we get the probability that the packet falls within a window with *any* duration by simply adding together $p_{fall\_within}(d_c, d_w)$ for all window durations of interest. The windows of interest are the ones with duration equal to or longer than the duration of the communication packet, as there is no way for the packet to fit within an opportunity window shorter than this. Thus, the probability that a communication packet is received without any interference is

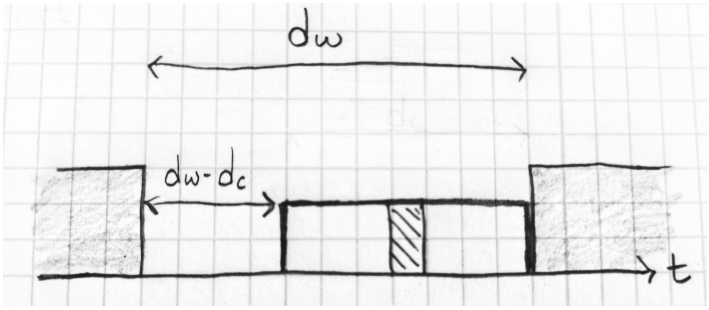

Figure 8.3: A communication packet that falls within an opportunity window.

$$\begin{aligned} p_{no\_interference}(d_c) &= \sum_{d_w=d_c}^{d_{w_{max}}} p_{fall\_within}(d_c, d_w) \\ &= \sum_{d_w=d_c}^{d_{w_{max}}} OppDist[P, d_w] \cdot (d_w - d_c)/d_w. \end{aligned} \tag{8.3}$$

8.3 Packet loss probability

We assume that an entire communication packet is lost if, and only if, any interference occurs in the header sequence. Using the result from equation 8.3, the probability that no interference occurs during the header is $p_{no\_interference}(d_h)$, such that the probability for packet loss is $1 - p_{no\_interference}(d_h)$. This is found by observing that the probability that interference is present in the header is the same as the probability that interference is present in a communication packet with duration d_h.

8.4 Distribution for how large portion of the payload data falls outside the opportunity window

In this section, we will estimate the distribution for how likely it is that a certain percentage, x%, of the payload data falls outside the main opportunity window. To get this, we will first find the cumulative probability that less than x% of the payload data falls outside. This will be estimated by dividing the opportunity windows into four different categories based on their duration: those with duration $d_w > d_c$, those with duration $d_w \in [\frac{d_c+d_h}{2}, d_c]$, those with duration $d_w \in [d_h, \frac{d_c+d_h}{2}]$, and those with duration $d_w < d_h$. The opportunity windows with duration $d_w < d_h$ are ignored because it is not possible to fit a header inside these windows. For each of the three remaining categories, we will find the distribution for how likely it is that a communication packet falls within such a window and that less than x% of the payload data is outside the main opportunity window. This is described in sections 8.4.1, 8.4.2 and 8.4.3, and combined into a single expression in section 8.4.4. In sections 8.4.1, 8.4.2, and 8.4.3, we will not consider the cases where no interference is present in the packet at all, but this will still be taken into account in the final expression.

8.4.1 Windows with duration $d_w > d_c$

For an opportunity window that lasts longer than the communication packet, the only way the packet can fall outside the window is for one of the edges of the packet to fall outside as shown in figure 8.4. From this figure, we see that a packet must start within an interval of duration $\frac{d_c-d_h}{2}$ if the beginning of the packet is to be obstructed by interference. If a given communication packet falls within the window, there is a $d_w - d_h$ interval in which the packet may start, such that the probability that the beginning of it is obstructed by interference is $\frac{(d_c-d_h)/2}{d_w-d_h}$. This probability must be multiplied by 2 to get the probability that either the beginning or the end of the packet is obstructed by interference. Thus, if a packet is

transmitted at a random time, the probability that it falls within an opportunity window with duration d_w, and that the beginning or end of the packet is obstructed by interference is

$$\begin{aligned} & 2 \cdot \frac{(d_c - d_h)/2}{d_w - d_h} \cdot p_{fall\_within}(d_h, d_w) \\ &= \frac{(d_c - d_h)}{d_w - d_h} \cdot (d_w - d_h) OppDist[P, d_w]/d_w \\ &= (d_c - d_h) \cdot OppDist[P, d_w]/d_w. \end{aligned} \tag{8.4}$$

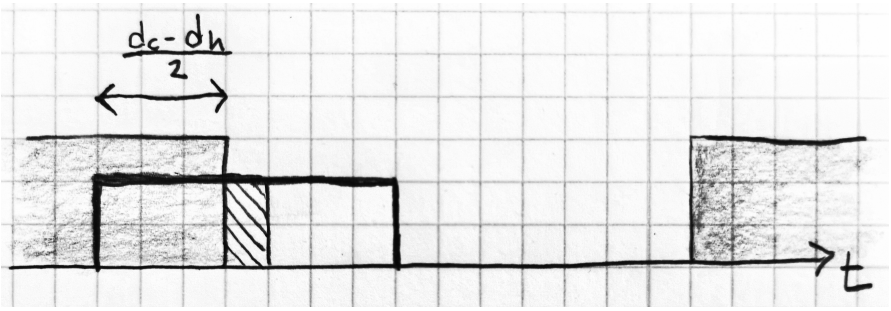

Figure 8.4: A communication packet in an opportunity window with duration $d_w > d_c$, where the beginning of the packet falls outside the opportunity window.

Concerning the portion of the payload data that will fall outside the opportunity window, we see from figure 8.4 that it will vary between 0% and 50% of the payload data, with uniform probability distribution for the different portions. Therefore, the probability that a packet will fall inside any window with duration d_w, and that less than x of the payload data will be obstructed by interference, is given as

$$\begin{cases} (d_c - d_h) \cdot OppDist[P, d_w]/d_w \cdot \frac{x}{0.50}, & x < 0.50 \\ (d_c - d_h) \cdot OppDist[P, d_w]/d_w, & x > 0.50 \end{cases} \tag{8.5}$$

Because this probability only considers opportunity windows with the specific duration d_w, we need to sum it over all window durations of interest, $d_w < d_c$, to get the probability that this happens for *any* opportunity window with duration $d_w < d_c$:

$$p_{less\_than\_x\_outside,\,1}(x) = \sum_{d_w = d_c}^{d_{w_{max}}} \begin{cases} (d_c - d_h) \cdot OppDist[P, d_w]/d_w \cdot \frac{x}{0.50}, & x < 0.50 \\ (d_c - d_h) \cdot OppDist[P, d_w]/d_w, & x > 0.50 \end{cases} \tag{8.6}$$

8.4.2 Windows with duration $d_w \in \left[\frac{d_c + d_h}{2}, d_c\right\rangle$

Figure 8.5 shows an opportunity window that has duration $d_w \in \left[\frac{d_c+d_h}{2}, d_c\right\rangle$. A communication packet that falls within such a window will always have parts of it that fall outside the window because the duration of the window is shorter than the duration of the packet. Furthermore, there are two ways in which parts of it can fall outside the window. The first is that *either* the beginning or the end falls outside, as shown in figure 8.5a, and the second is that *both* the beginning and the end fall outside the window, as would happen with the packet in figure 8.5b if it was moved slightly to the right.

First, the probability that a communication packet falls within a window with duration d_w, where $d_w \in \left[\frac{d_c+d_h}{2}, d_c\right\rangle$, is equal to the probability that the entire header falls within the window. Following the same reasoning as in section 8.3, this probability is

$$p_{fall\_within}(d_h, d_w) = (d_w - d_h) OppDist[P, d_w]/d_w. \tag{8.7}$$

Given that this happens, we see from figure 8.5 that the probability that both the beginning and end of the packet falls outside the window is $\frac{d_c - d_w}{d_w - d_h}$, while the probability that only the beginning or only the

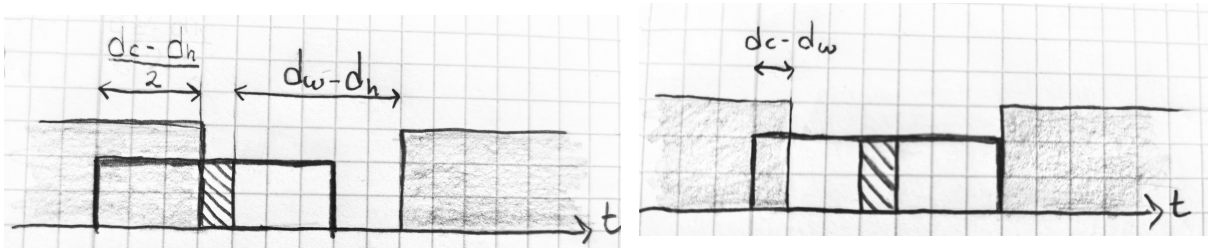

(a) A packet located such that the beginning of the packet falls outside the opportunity window.

(b) If the packet is located at the position in the figure, or more to the right, both the beginning and the end of the packet will fall outside the opportunity window.

Figure 8.5: Two communication packets, each in an opportunity window with duration $d_w \in \left[\frac{d_c+d_h}{2}, d_c\right\rangle$.

end falls outside is $1 - \frac{d_c-d_w}{d_w-d_h} = \frac{2d_w-d_h-d_c}{d_w-d_h}$ Combining this with equation 8.7, we get the probability that the communication packet falls within a window with duration d_w, and both the beginning and end of the packet are outside the window as

$$\begin{aligned}&(d_w - d_h)OppDist[P, d_w]/d_w \cdot \frac{d_c - d_w}{d_w - d_h} \\ =&(d_c - d_w)OppDist[P, d_w]/d_w.\end{aligned} \tag{8.8}$$

Likewise, the probability that the communication packet falls within a window with duration d_w, and that either the beginning or end is outside the window, is

$$(2d_w - d_h - d_c)OppDist[P, d_w]/d_w. \tag{8.9}$$

Furthermore, when both the beginning and end of the packet fall outside the window, the total duration of the payload data that is lost to interference is always equal to $(d_c - d_w)$. Because the duration of the payload data is $(d_c - d_h)$, the portion of the payload data that is outside the window is $\frac{d_c-d_w}{d_c-d_h}$. Therefore, the probability that a packet falls inside a window with duration d_w, and that both the beginning and the end of the packet fall outside the window in such a way that less than x% of the payload data is outside the window, is

$$\begin{cases} 0, & x < \frac{d_c-d_w}{d_c-d_h} \\ (d_c - d_w)OppDist[P, d_w]/d_w, & otherwise \end{cases} \tag{8.10}$$

Because this probability only considers opportunity windows with the specific duration d_w, we need to sum it over all window durations of interest, $d_w \in \left[\frac{d_c+d_h}{2}, d_c\right\rangle$ to get the probability that this happens for any opportunity window with *any* duration $d_w \in \left[\frac{d_c+d_h}{2}, d_c\right\rangle$:

$$p_{less\_than\_x\_outside,\,2}(x) = \sum_{d_w=\frac{d_c+d_h}{2}}^{d_c} \begin{cases} 0, & x < \frac{d_c-d_w}{d_c-d_h} \\ (d_c - d_w)OppDist[P, d_w]/d_w, & otherwise \end{cases} \tag{8.11}$$

When only the beginning *or* end of the packet falls outside the main window, the duration that falls outside will vary between $(d_c - d_w)$ and $\frac{d_c-d_h}{2}$ with uniform distribution. Because the total duration of the payload data is $(d_c - d_h)$, the portion of the payload data that is outside varies between $\frac{d_c-d_w}{d_c-d_h}$ and $\frac{(d_c-d_h)/2}{d_c-d_h} = 0.50$. Therefore, the probability that a packet will fall inside a window with duration d_w and that either the beginning or end of it falls outside the window in such a way that less than x% of the payload data is outside the window, is

$$\begin{cases} 0, & x < \frac{d_c-d_w}{d_c-d_h} \\ (2d_w - d_h - d_c)OppDist[P, d_w]/d_w \cdot \frac{x-\frac{d_c-d_w}{d_c-d_h}}{0.50-\frac{d_c-d_w}{d_c-d_h}}, & x \in \left[\frac{d_c-d_w}{d_c-d_h}, 0.50\right] \\ (2d_w - d_h - d_c)OppDist[P, d_w]/d_w, & x > 0.50. \end{cases} \tag{8.12}$$

As this probability only considers opportunity windows with the specific duration d_w, we must sum it over all window durations of interest, $d_w \in \left[\frac{d_c+d_h}{2}, d_c\right\rangle$ to get the probability that this happens for a window with *any* duration $d_w \in \left[\frac{d_c+d_h}{2}, d_c\right\rangle$:

$$p_{less\_than\_x\_outside,\,3}(x) = \sum_{d_w=\frac{d_c+d_h}{2}}^{d_c} \begin{cases} 0, & x < \frac{d_c-d_w}{d_c-d_h} \\ (2d_w - d_h - d_c)OppDist[P, d_w]/d_w \cdot \frac{x-\frac{d_c-d_w}{d_c-d_h}}{0.50-\frac{d_c-d_w}{d_c-d_h}}, & x \in [\frac{d_c-d_w}{d_c-d_h}, 0.50] \\ (2d_w - d_h - d_c)OppDist[P, d_w]/d_w, & x > 0.50. \end{cases} \tag{8.13}$$

8.4.3 Windows with duration $d_w \in \left[d_h, \frac{d_c+d_h}{2}\right\rangle$

Figure 8.6 shows an opportunity window with duration $d_w \in \left[d_h, \frac{d_c+d_h}{2}\right\rangle$. The probability that a communication packet falls within a window with duration d_w in equation 8.7 is valid also for windows with duration $d_w \in \left[d_h, \frac{d_c+d_h}{2}\right\rangle$. Furthermore, such a window is so narrow that any communication packet that falls within it will always have both its beginning and end outside the opportunity window. The total duration of the payload data that is outside the window, is then always equal to the duration of the packet minus the duration of the window, $(d_c - d_w)$. Because the combined duration of the payload data is $(d_c - d_h)$, the portion of it that falls outside the window, is $\frac{d_c-d_w}{d_c-d_h}$. Thus, the probability that a packet will fall inside a window with duration d_w, and that less than x% of the payload data is outside the window, is

$$\begin{cases} 0, & x < \frac{d_c-d_w}{d_c-d_h} \\ (d_w - d_h)OppDist[P, d_w], & x > \frac{d_c-d_w}{d_c-d_h} \end{cases} \tag{8.14}$$

Because this probability only considers opportunity windows with the specific duration d_w, we must sum it over all window durations of interest, $d_w \in \left[d_h, \frac{d_c+d_h}{2}\right\rangle$ to get the probability that this happens for a window with *any* duration $d_w \in \left[d_h, \frac{d_c+d_h}{2}\right\rangle$:

$$p_{less\_than\_x\_outside,\,4}(x) = \sum_{d_w=d_h}^{\frac{d_c+d_h}{2}} \begin{cases} 0, & x < \frac{d_c-d_w}{d_c-d_h} \\ (d_w - d_h)OppDist[P, d_w], & x > \frac{d_c-d_w}{d_c-d_h} \end{cases} \tag{8.15}$$

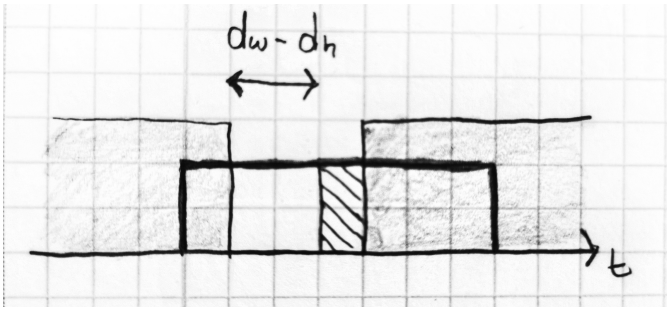

Figure 8.6: A communication packet in an opportunity window with duration $d_w \in \left[d_h, \frac{d_c+d_h}{2}\right\rangle$.

8.4.4 Combined probability that less than x% of the payload data falls outside the opportunity window.

In section 8.4.1, 8.4.2, and 8.4.3, we investigated all possible ways that parts of a communication packet can fall outside the main opportunity window. To get the combined probability that less than x% of the payload data in a packet falls outside the main opportunity window, we must add the probability that this happens for the four cases we calculated in section 8.4.1, 8.4.2, and 8.4.3. In addition to this, we must add the probability that the entire packet falls inside an opportunity window, without any interference in the payload data:

$$
\begin{aligned}
p_{less\_than\_x\_outside}(x) = \quad & p_{no\_interference}(d_c) \\
+ \quad & p_{less\_than\_x\_outside,1}(x) \\
+ \quad & p_{less\_than\_x\_outside,2}(x) \\
+ \quad & p_{less\_than\_x\_outside,3}(x) \\
+ \quad & p_{less\_than\_x\_outside,4}(x) \\
= \quad & \sum_{d_w=d_c}^{d_{wmax}} OppDist[P,d_w](d_w-d_c)/d_w \\
+ \quad & \sum_{d_w=d_c}^{d_{wmax}} \begin{cases} (d_c-d_h)\cdot OppDist[P,d_w]/d_w \cdot \frac{x}{0.50}, & x<0.50 \\ (d_c-d_h)\cdot OppDist[P,d_w]/d_w, & x>0.50 \end{cases} \\
+ \quad & \sum_{d_w=\frac{d_c+d_h}{2}}^{d_c} \begin{cases} 0, & x<\frac{d_c-d_w}{d_c-d_h} \\ (d_c-d_w)OppDist[P,d_w]/d_w, & otherwise \end{cases} \\
+ \quad & \sum_{d_w=\frac{d_c+d_h}{2}}^{d_c} \begin{cases} 0, & x<\frac{d_c-d_w}{d_c-d_h} \\ (2d_w-d_h-d_c)OppDist[P,d_w]/d_w \cdot \frac{x-\frac{d_c-d_w}{d_c-d_h}}{0.50-\frac{d_c-d_w}{d_c-d_h}}, & x\in[\frac{d_c-d_w}{d_c-d_h},0.50] \\ (2d_w-d_h-d_c)OppDist[P,d_w]/d_w, & x>0.50. \end{cases} \\
+ \quad & \sum_{d_w=d_h}^{\frac{d_c+d_h}{2}} \begin{cases} 0, & x<\frac{d_c-d_w}{d_c-d_h} \\ (d_w-d_h)OppDist[P,d_w], & x>\frac{d_c-d_w}{d_c-d_h} \end{cases}
\end{aligned}
\tag{8.16}
$$

Note that in this expression, it is important to not double-count the windows in the sum limits. That is, the longest window included in the sum with upper limit $d_w = \frac{d_c+d_h}{2}$ should not be included in the sums with $d_w = \frac{d_c+d_h}{2}$ as the lower limit, and equivalently for the limit $d_w = d_c$. The distribution for the probability that *exactly* x% of the payload data is lost to interference is found by deriving $p_{less\_than\_x\_outside}(x)$ and adding the probability that the entire packet is lost to interference from section 8.3 at $x = 100\%$. Performing the derivation gives

$$
\begin{aligned}
p_{outside}(x) = \quad & p'_{less\_than\_x\_outside}(x) + \delta(x-100\%)\cdot(1-p_{no\_interference}(d_h)) \\
= \quad & \sum_{d_w=d_c}^{d_{wmax}} OppDist[P,d_w](d_w-d_c)/d_w \cdot \delta(\chi) \\
+ \quad & \sum_{d_w=d_c}^{d_{wmax}} \begin{cases} (d_c-d_h)\cdot OppDist[P,d_w]/d_w \cdot \frac{1}{0.50}, & \chi<0.50 \\ 0, & \chi>0.50 \end{cases} \\
+ \quad & \sum_{d_w=\frac{d_c+d_h}{2}}^{d_c} (d_c-d_w)OppDist[P,d_w]/d_w \cdot \delta\left(\chi-\frac{d_c-d_w}{d_c-d_h}\right) \\
+ \quad & \sum_{d_w=\frac{d_c+d_h}{2}}^{d_c} \begin{cases} 0, & \chi<\frac{d_c-d_w}{d_c-d_h} \\ (2d_w-d_h-d_c)OppDist[P,d_w]/d_w \cdot \frac{1}{0.50-\frac{d_c-d_w}{d_c-d_h}}, & \chi\in[\frac{d_c-d_w}{d_c-d_h},0.50] \\ 0, & \chi>0.50. \end{cases} \\
+ \quad & \sum_{d_w=d_h}^{\frac{d_c+d_h}{2}} (d_w-d_h)OppDist[P,d_w]\cdot\delta\left(\chi-\frac{d_c-d_w}{d_c-d_h}\right) \\
+ \quad & \delta(x-100\%)\cdot\left(1-\sum_{d_w=d_h}^{d_{wmax}} OppDist[P,d_w]\cdot(d_w-d_h)/d_w\right)
\end{aligned}
\tag{8.17}
$$

8.5 Distribution for how large portion of the payload data is lost to interference

The probability in equation 8.16 describes how likely it is that a certain portion of the payload data falls outside the main opportunity window. It is however not guaranteed that everything that falls outside this window is obstructed by interference, since the interference pulse that ends the opportunity window may be short, as shown in figure 8.8. If we know that the total duration of the payload data that falls outside the main window is $d_{outside}$, we need to calculate how large portion of this is actually lost to interference, d_{lost}. We define a distribution $f(d_{lost}; d_{outside})$ that denotes how likely it is that d_{lost} of the packet is lost to interference, given that $d_{outside}$ of it is outside the main opportunity window.

We will make an approximation that will simplify the estimation of $p_{lost}(x)$, by allowing us to use the estimate for $p_{outside}(\chi)$ to estimate $p_{lost}(x)$. To explain this approximation, we consider the case shown in figure 8.7a, where 15% of the payload data in a communication packet falls outside the main opportunity window to the left, and 35% falls outside to the right. The distribution for how large portion of the signal that is lost to interference on the left side is $f\left(d_{lost};\ 0.15(d_c - d_h)\right)$, and likewise, the distribution for how much is lost on the right side, is $f\left(d_{lost};\ 0.35(d_c - d_h)\right)$. The combined duration that is lost to interference, is the sum of these two. Because the distribution for a sum of random variables is the convolution between their distributions, the combined duration that is lost to interference will be distributed as $f\left(d_{lost};\ 0.15(d_c - d_h)\right) * f\left(d_{lost};\ 0.35(d_c - d_h)\right)$. However, $p_{outside}(\chi)$ denotes the probability that $\chi\%$ of the payload data falls outside the main opportunity window, regardless of whether $\chi\%$ is split between the right and left side of the window, or if it is all on one side. There is therefore no way to distinguish between the cases in figure 8.7a and 8.7b, if we are to use $p_{outside}(\chi)$ in the estimation of $p_{lost}(x)$. We, therefore, make the approximation that when a portion of the payload data, $d_{outside}$, is outside the main opportunity window, we assume that the portion of the signal that is lost to interference can be calculated as if that entire portion fell outside the window on one side, and not divided between the left and right. This is essentially approximating the case in figure 8.7a as the case in figure 8.7b. This approximation only affects the cases where part of the communication packet falls outside the main opportunity window to the left, and part of it to the right. Instead of $f\left(d_{lost};\ 0.15(d_c - d_h)\right) * f\left(d_{lost};\ 0.35(d_c - d_h)\right)$, the portion of the payload data that is lost to interference is now estimated as $f\left(d_{lost};\ 0.50(d_c - d_h)\right)$.

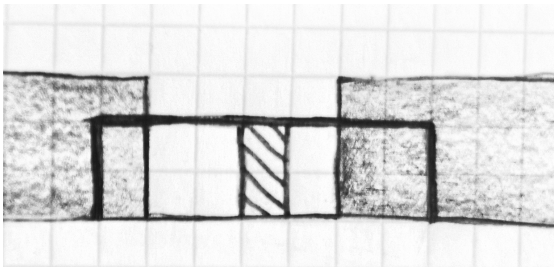

(a) 15% of of the payload data fall outside the main window on the left, and 35% on the right.

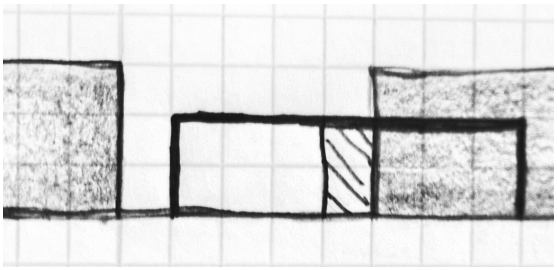

(b) 0% of the payload data fall outside the main window on the left, and 50% on the right.

Figure 8.7: Two communication packets, where part of the packets fall outside the opportunity window on the left or right side of the main window.

It is not possible to estimate $f(d_{lost}; d_{outside})$ based on the opportunity distribution alone because we also need information about how likely a pulse is to have a certain duration. This information is not contained in the opportunity distribution, but can be obtained by measuring what we will refer to as the *pulse distribution*. While the opportunity distribution denotes how much time is spent in opportunity windows with a given length, the pulse distribution is the corresponding measure for how much time is spent in interference pulses of a given length. Because this report focuses on the properties of the opportunity distribution, we will not investigate the pulse distribution in detail. We will, however, still assume that the pulse distribution is known when estimating $f(d_{lost}; d_{outside})$. How precise this estimate is will provide an indication of how useful the pulse distribution is, and can be used to justify whether or not one should focus future work on investigating it. We will denote the pulse distribution as $PulsDist[\cdot, \cdot]$. $PulsDist[p, d_p]$ represents how likely the signal is, at any given time, to be in an interference pulse with

duration d_p, where an interference pulse is defined as a time interval where the signal is continuously at or above the power level P.

In figure 8.8, we have labelled the duration of the first interference pulse preceding the main opportunity window, as d_{p_1}, and the opportunity window preceding this pulse as d_{w_1}. In order to estimate $f(d_{lost}; d_{outside})$, we need to know the distribution for the durations d_{w_1} and d_{p_1}. The number of opportunity windows with duration d_{w_1} that occur during a second, is $\frac{OppDist[P,d_{w_1}]\cdot 1s}{d_{w_1}}$ from equation 8.1. If we choose an opportunity window at random, the probability that it will have duration d_{w_1} is thus

$$\begin{aligned} P_w[d_{w_1}] &= \frac{OppDist[P, d_{w_1}] \cdot 1s/d_{w_1}}{\sum_{d_i=d_{w_{min}}}^{d_{w_{max}}} OppDist[P, d_i] \cdot 1s/d_i} \\ &= \frac{OppDist[P, d_{w_1}]/d_{w_1}}{\sum_{d_i=d_{w_{min}}}^{d_{w_{max}}} OppDist[P, d_i]/d_i} \end{aligned} \tag{8.18}$$

Likewise, if we choose an interference pulse at random, the probability that it has duration d_{p_1}, is

$$P_p[d_{p_1}] = \frac{PulsDist[p, d_{p_1}]/d_{p_1}}{\sum_{d_i=d_{p_{min}}}^{d_{p_{max}}} OppDist[P, d_i]/d_i} \tag{8.19}$$

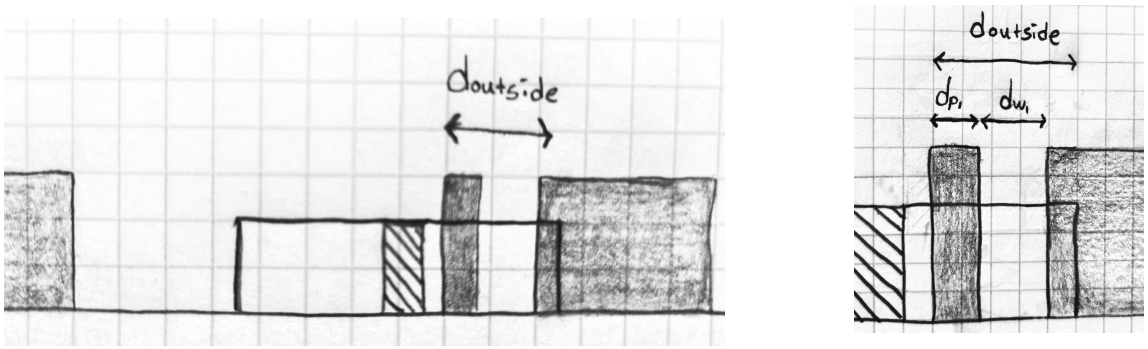

(a) The communication packet

(b) Close-up view of the part that is outside the main opportunity window.

Figure 8.8: A communication packet, where a portion of the packet with duration $d_{outside}$ falls outside the opportunity window.

Depending on the duration d_{p_1}, d_{w_1}, and $d_{outside}$, there are three expressions for how large portion of the communication packet is lost to interference:

- If $d_{outside} < d_{p_1}$, everything that is outside the opportunity window is lost to interference, $d_{lost} = d_{outside}$.
- If $d_{p_1} \leq d_{outside} \leq d_{p_1} + d_{w_1}$, the duration of the signal loss due to interference is is $d_{lost} = d_{p_1}$.
- If $d_{outside} > d_{p_1} + d_{w_1}$, as is the case in figure 8.8, the signal loss due to interference is composed of two parts. The first contribution is due to the interference pulse with duration d_{p_1}, and the second contribution is the interference that occurs after the window with duration d_{w_1}. We will denote the duration of the signal that is lost from this contribution as d_{lost_2}. The duration that falls to the right of the window with duration d_{w_1}, is $d_{outside} - d_{p_1} - d_{w_1}$. Thus, d_{lost_2} is distributed according to $f(d_{lost_2}; d_{outside} - d_{p_1} - d_{w_1})$.

In order to calculate $f(d_{lost}; d_{outside})$, one must consider all possible durations for d_{p_1} and d_{w_1}. For each combination of d_{p_1} and d_{w_1}, the probability that this particular outcome occurs is $P_p[d_{p_1}] \cdot P_w[d_{w_1}]$. This assumes that d_{p_1} and d_{w_1} are distributed independently of each other, and independent of the preceding pulse- and window durations. $f(d_{lost}; d_{outside})$ can be calculated as a weighted sum over all the outcomes, where each term in the sum is the portion of the signal that will be lost for that particular outcome,

weighted by the probability that the outcome occurs. Each term in the sum will correspond to one of the three possibilities in the bullet list above, so the expression for $f(d_{lost}, d_{outside})$ is thus

$$f(d_{lost}; d_{outside}) = \sum_{d_{p_1}=d_{p_{min}}}^{d_{p_{max}}} \sum_{d_{w_1}=d_{w_{min}}}^{d_{w_{max}}} P_p[d_{p_1}] P_w[d_{w_1}] \cdot \begin{cases} \delta[d_{lost} - d_{outside}], & d_{outside} < d_{p_1} \\ \delta[d_{lost} - d_{p_1}], & d_{outside} \in [d_{p_1}, d_{p_1} + d_{w_1}] \\ f(d_{lost} - d_{p_1}; d_{outside} - d_{p_1} - d_{w_1}), & d_{outside} > d_{p_1} + d_{w_1}. \end{cases} \tag{8.20}$$

Note that the expression for $f(d_{lost}; d_{outside})$ contains $f(d_{lost}; d_{outside})$ in itself, which makes it difficult, or impossible, to solve analytically. It is, however, still possible to calculate the expression numerically, by first selecting a discrete list of values for $d_{outside}$ and d_{lost}, for which one wishes to compute $f(d_{lost}; d_{outside})$. f is then calculated for the smallest value for $d_{outside}$, for all values for d_{lost}. After this, f is computed for the second smallest value of $d_{outside}$, then the third smallest value, and so on. One is then guaranteed that when computing f for a given value of $d_{outside}$, all values of f for the shorter values of $d_{outside}$ are already computed, and one can therefore look up $f(d_{lost} - d_{p_1}; d_{outside} - d_{p_1} - d_{w_1})$.

$f(d_{lost}; d_{outside})$ can be used to derive an expression for how likely it is that $x\%$ of the payload data is lost to interference. We remember that $p_{outside}(\chi)$ from equation 8.17 denotes the probability that $\chi\%$ of the payload data falls outside the main opportunity window, which corresponds to a duration $d_{outside} = \chi \cdot (d_c - d_h)$ outside the window. The distribution for how large portion of the signal is lost to interference is then $f(x \cdot (d_c - d_h);\ \chi \cdot (d_c - d_h))$, where x is the percentage of the payload data that is lost to interference. Using this, we can get the probability that x% of the payload data is lost by integrating over all possible values for χ,

$$\int_{0\%}^{100\%} f\left(x \cdot (d_c - d_h);\ \chi \cdot (d_c - d_h)\right) \cdot p_{outside}(\chi)\ \ d\chi. \tag{8.21}$$

However, we must remember that $p_{outside}(\chi)$ includes the cases where the header is hit by interference. In these cases, the entire packet is lost to interference, so the final expression for $p_{lost}(x)$ becomes

$$\begin{aligned} p_{lost}(x) = \int_{0\%}^{100\%} f\left(x \cdot (d_c - d_h);\ \chi \cdot (d_c - d_h)\right) \cdot \left(p_{outside}(\chi) - \delta(\chi - 100\%)(1 - p_{no\_interference}(d_h))\right)\ \ d\chi \\ + \int_{0\%}^{100\%} \delta(x - 100\%) \cdot \delta(\chi - 100\%)(1 - p_{no\_interference}(d_h))\ \ d\chi. \end{aligned} \tag{8.22}$$

From this, we can calculate the probability that *less than* x% of the payload data is lost to interference as

$$p_{less\_than\_x\_lost}(x) = \int_0^x p_{lost}(z)\ dz. \tag{8.23}$$

8.5.1 Improving the precision of the estimates

The opportunity distribution at a power level P might look something like figure 8.9a. At the beginning of this chapter, we said that for the purposes of estimating $p_{less\_than\_x\_outside}(\chi)$ and $p_{lost}(x)$, we would assume that all opportunity windows with duration within an interval $[d_j, d_{j+1}]$ in the opportunity distribution could be approximated as having the same duration, $d_w = \frac{d_j + d_{j+1}}{2}$. This is essentially approximating the opportunity distribution as the one in figure 8.9b. The validity of this approximation is worse the larger the interval $[d_j, d_{j+1}]$ is. This means that the approximation does not hold well when using a large scaling factor s, where s is defined as described in section 6.2. A better approximation would be to assume that the opportunity windows are equally likely to have any duration within the interval $[d_j, d_{j+1}]$. This would, however, make it more difficult to derive expressions for $p_{less\_than\_x\_outside}(\chi)$ and $p_{less\_than\_x\_lost}(x)$. The next best option is to upsample the opportunity distribution in the duration

dimension before computing the estimates. Figure 8.9c shows how the opportunity distribution might look after upsampling with an upsampling factor of 5. When estimating $p_{less\_than\_x\_outside}(\chi)$ and $p_{lost}(x)$, the opportunity distribution would then be approximated as the one in figure 8.9d rather than as the one in figure 8.9b. Note that the probability of each bin in the upsampled opportunity distribution is a fifth of the probability of the corresponding bin in the original opportunity distribution. This is a necessary normalization when upsampling the opportunity distribution, as more bins require the probability per bin to decrease.

The above procedure is advantageous compared to just using a smaller scaling factor s to make the duration intervals shorter to begin with. Using a smaller scaling factor would indeed result in better estimates, but it would also increase the size of the opportunity distribution, which influences the data budget. This is investigated more in section 9.3.

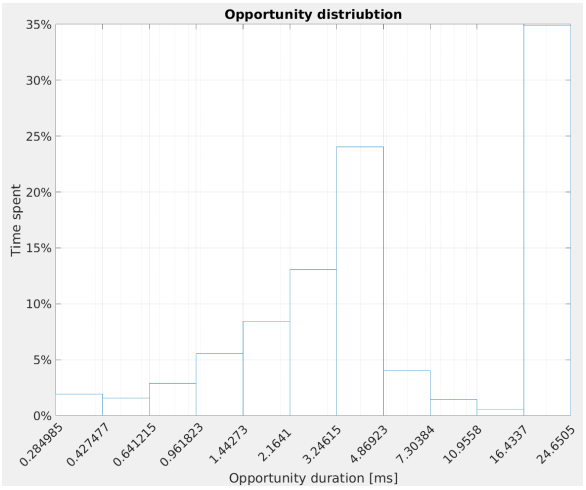

(a) An example of how an opportunity distribution might look like.

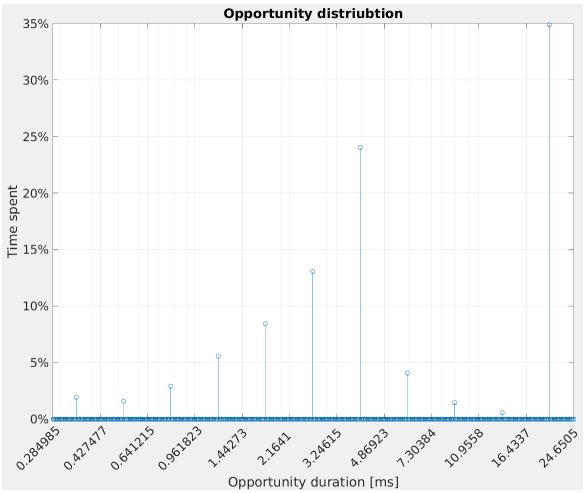

(b) The opportunity distribution is approximated as this while estimating $p_{less\_than\_x\_outside}(\chi)$ and $p_{lost}(x)$.

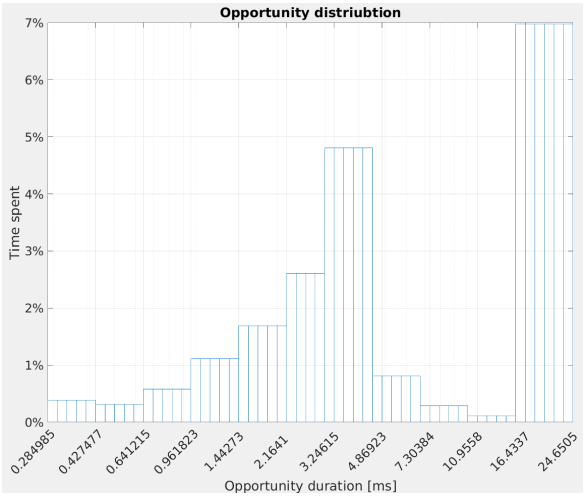

(c) The opportunity distribution upsampled.

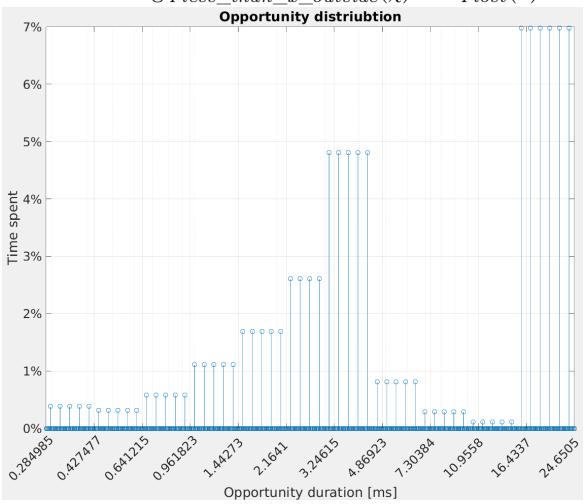

(d) The upsampled opportunity distribution is approximated as this while estimating $p_{less\_than\_x\_outside}(\chi)$ and $p_{lost}(x)$.

Figure 8.9

9. *Validation of estimated expressions*

The estimates derived in chapter 8 must be validated in order to verify their correctness, and how they are affected by changing the resolution and number of quantization levels in the opportunity distribution. Instead of validating the expressions $p_{lost}(x)$ and $p_{outside}(x)$ directly, we will validate the cumulative distributions $p_{less\_than\_x\_outside}(x)$ and $p_{less\_than\_x\_lost}(x)$ in equations 8.16 and 8.23. The reason for this is that the distributions $p_{lost}(x)$ and $p_{outside}(x)$ contain delta pulses, which makes it difficult to plot and compare with the true distributions. Furthermore, the value of the cumulative distributions at $x = 0\%$ and $x = 100\%$ is the probability that no interference is present in the communication packet, and one minus the probability of packet loss respectively. The estimation of these values is therefore also validated when estimating the cumulative distributions. In order to do the verification, we will use an arbitrarily session from the Norsat-2 measurements as a test signal and calculate the estimates based on the opportunity distribution and pulse distribution for this session. Session seven is chosen as the test signal, which is is 12.5 minutes long. A 5-second segment of this signal is used to calculate the opportunity distribution and pulse distribution that the estimates are based on.

9.1 Ideal measurement parameters

In this section, the estimates for $p_{less\_than\_x\_outside}(x)$ and $p_{less\_than\_x\_lost}(x)$ are calculated using very precise measurement parameters for the opportunity distribution. To do this, one must first choose a test duration for the communication packet, d_c and its header, d_h, to use in the validation. The choice for these values is not critical, as one only needs to select some example values to estimate for. In section 3.1, we found that the duration of a symbol used in UHF satellite communication must be 0.05 ms or longer. We set the duration of the header to be ten times this duration, $d_h = 0.5\,\text{ms}$, and the duration of the communication packet to be ten times the duration of the header, $d_c = 5\,\text{ms}$. We will also test with a shorter duration, $d_c = 1\,\text{ms}$ and $d_h = 0.1\,\text{ms}$, just to get a different realization. As for the power threshold P, we will also test two different cases, $P = -129$ dBm and $P = -123$ dBm.

To compute the expressions in equations 8.16 and 8.23, which we are to verify, we require the opportunity distribution and pulse distribution of the signal. The shortest duration to calculate the opportunity distribution for, $d_{w_{min}}$, must be shorter than the duration of the packet header, d_h. Choosing $d_{w_{min}}$ too large will lead to imprecise estimates. Likewise, $d_{w_{max}}$ must be chosen sufficiently large, and the scaling factor s between must be sufficiently small. As we do not know what is sufficient before performing the validation, we set $d_{w_{min}} = \frac{1}{f_s} = 7.4\,\mu\text{s}$, $d_{w_{max}} = 5\,\text{s}$ and $s = 1.01$, as these values are probably far beyond what would be sufficient. Additionally, the opportunity distribution is not quantized. The pulse distribution is calculated using the same parameters as the opportunity distribution.

As describes in section 8.5, estimating $p_{less\_than\_x\_lost}(x)$ involves estimating $f(d_{lost}; d_{outside})$, which requires us to select a discrete list of d_{lost} and $d_{outside}$ to estimate $f(d_{lost}; d_{outside})$ for. When estimating $f(d_{lost}; d_{outside})$, we set both these lists to be 100 values linearly spaced out between 0 and $d_c - d_h$. More values could lead to a more precise estimate, but would increase the time required to compute the estimate.

To validate the estimates for $p_{less\_than\_x\_outside}(x)$ and $p_{less\_than\_x\_lost}(x)$, one needs to know the true value for these distributions, and compare them with the estimates. For the IQ-signal that we are using for the verification, these distributions can be found by placing virtual communication packets at various points in the IQ signal and keeping track of how often a portion of the packet falls outside the main opportunity window, and how often a portion is lost to interference. This is performed, and the resulting

probability is used as a baseline for what we ideally want our estimate to be. $p_{less\_than\_x\_outside}(x)$ and $p_{less\_than\_x\_lost}(x)$ are then calculated from the opportunity distribution and pulse distribution, using the expressions in equations 8.16 and 8.23.

Figure 9.1 shows estimates for $p_{less\_than\_x\_outside}(x)$ and $p_{less\_than\_x\_lost}(x)$, compared with the true distributions for the different values of P, d_c and d_h that we have chosen. We can see that the estimate for $p_{less\_than\_x\_outside}(x)$ appears to be accurate for all the cases, as it aligns well with the true value for the probability. From e.g. the case in figure 9.1b, we can see that there is a 43.5% probability that no interference is present in the packet at all and a 87% probability that no interference is present in the header. Additionally, the probability that e.g. less than 60% of the payload data falls outside the main opportunity window is 79%. As for the estimate of $p_{less\_than\_x\_lost}(x)$, it appears to be accurate in some of the cases, but not all of them. Especially for the case in figure 9.1c, it does not overlap with the true distribution. That is, the estimate appears to be correct for $x > 20\%$, but this is only because the distribution flattens out for those x-values.

Regarding the processing time required to compute the estimates, $p_{less\_than\_x\_outside}(x)$ was estimated in less than 0.3 s for all the four cases in figure 9.1, while the estimate of $p_{less\_than\_x\_lost}(x)$ took between 20 minutes and 40 minutes to compute. This, combined with the fact that $p_{less\_than\_x\_lost}(x)$ requires the pulse distribution in addition to the opportunity distribution, makes it less attractive as an estimate.

In sections 9.2 - 9.6 we investigate how the resolution in the opportunity distribution affects the precision of the estimate of $p_{less\_than\_x\_outside}(x)$. Specifically, quantization of the opportunity distribution and the choice of the parameters d_{max}, d_{min} and the scaling factor s are investigated.

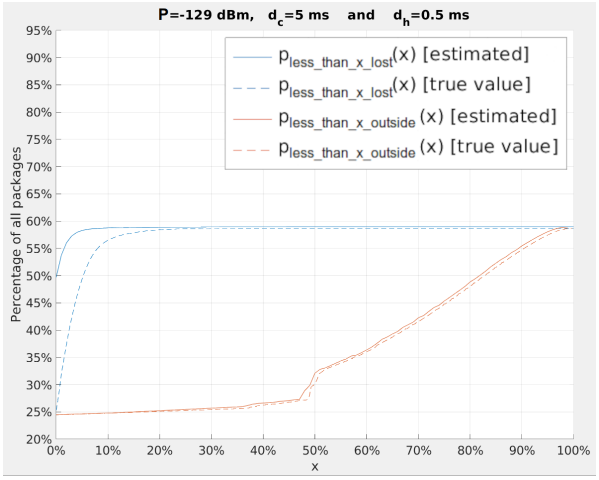

(a) $P = -129dBm$, $d_c = 5\,\text{ms}$ and $d_h = 0.5\,\text{ms}$

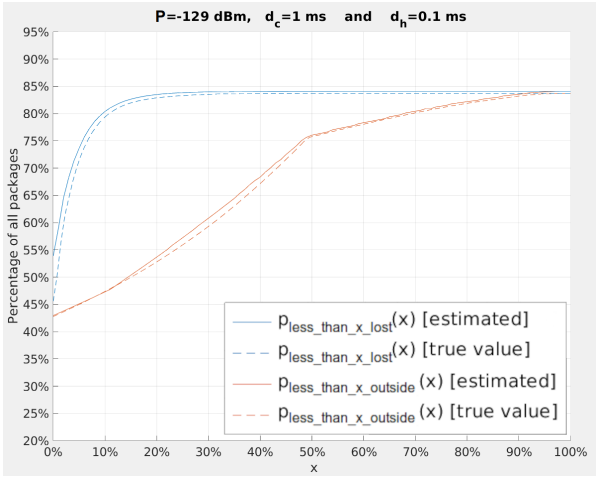

(b) $P = -129dBm$, $d_c = 1\,\text{ms}$ and $d_h = 0.1\,\text{ms}$

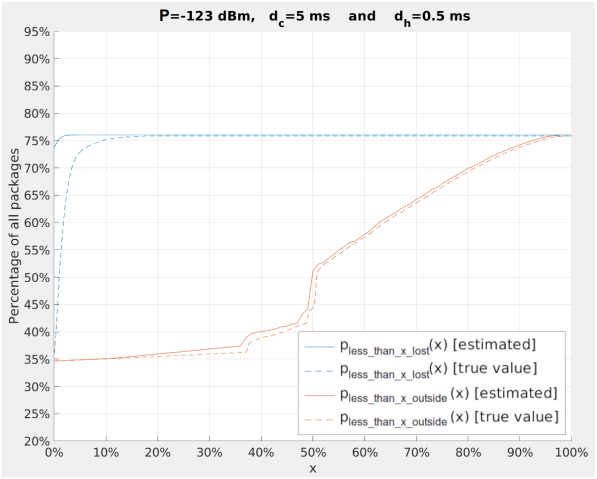

(c) $P = -123dBm$, $d_c = 5\,\text{ms}$ and $d_h = 0.5\,\text{ms}$

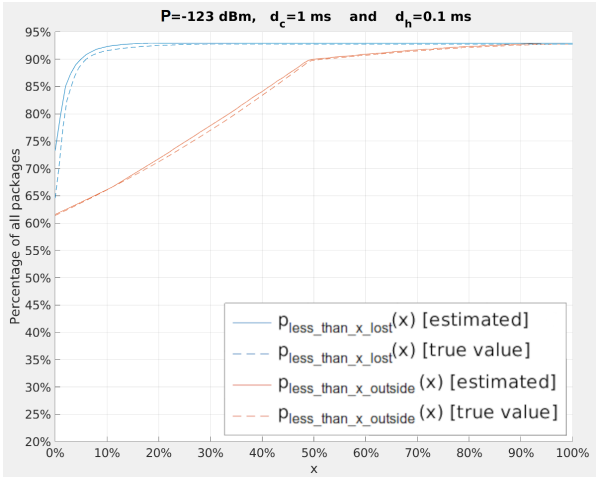

(d) $P = -123dBm$, $d_c = 1\,\text{ms}$ and $d_h = 0.1\,\text{ms}$

Figure 9.1: $p_{less\_than\_x\_outside}(x)$ and $p_{less\_than\_x\_lost}(x)$ estimated from the opportunity distribution versus the true distribution, for different values of P, d_c and d_h.

9.2 Effect of changing the number of quantization levels

In sections 9.1 and 9.3, the opportunity distribution and pulse distribution were calculated without quantizing it into discrete values. Quantizing is necessary in order to reduce the data size of the measurements, and the fewer quantization levels that are used, the smaller the measurement size is in bytes. There should, however, be enough quantization levels so that the measurements are not severely affected by the quantization errors. To find out how many quantization levels are necessary, the test in section 9.1 is repeated, but the opportunity distribution and pulse distribution are quantized into a set number of quantization levels. As described in section 7.2.1, the opportunity distribution is rounded in a special way that prevents rounding errors to accumulate in the cumulative distribution. To see how this affects the estimate of $p_{less\_than\_x\_outside}(x)$, we will test the effect of quantization both by rounding as described in section 7.2.1 and by rounding the opportunity distribution directly, in a regular fashion. Figure 9.2 and 9.3 show the estimates for $p_{less\_than\_x\_outside}(x)$ in these cases.

We see from figures 9.2 and 9.3 that when the opportunity distribution is rounded with the method described in section 7.2.1, the estimate for $p_{less\_than\_x\_outside}(x)$ is clearly more accurate. For this case, it appears that even 128 quantization levels produce quite an accurate estimate, and at 64 quantization levels, there start to be some artifacts in the estimate. Therefore, 128 quantization levels seem like a suitable choice.

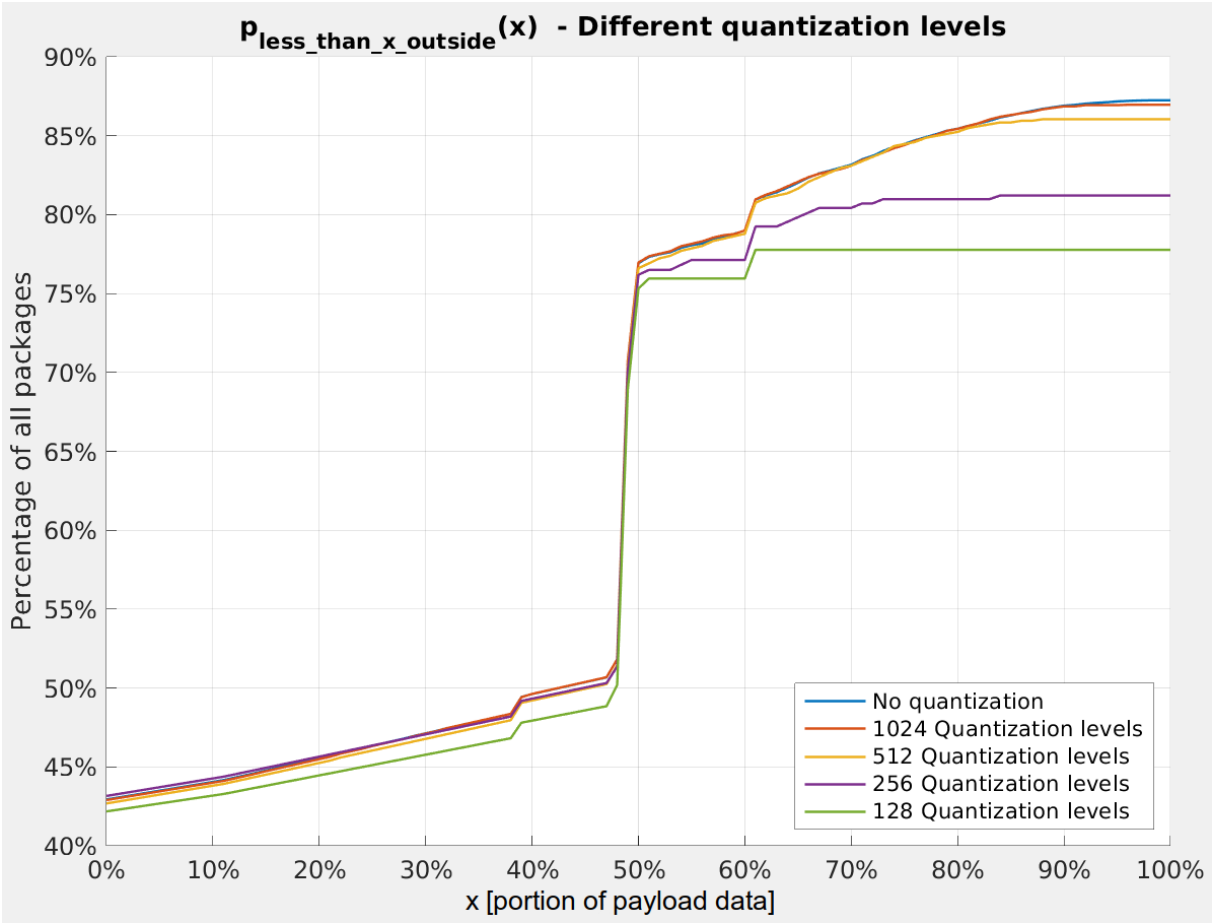

Figure 9.2: Estimates of $p_{less\_than\_x\_outside}(x)$, using an opportunity distribution that is rounded directly.

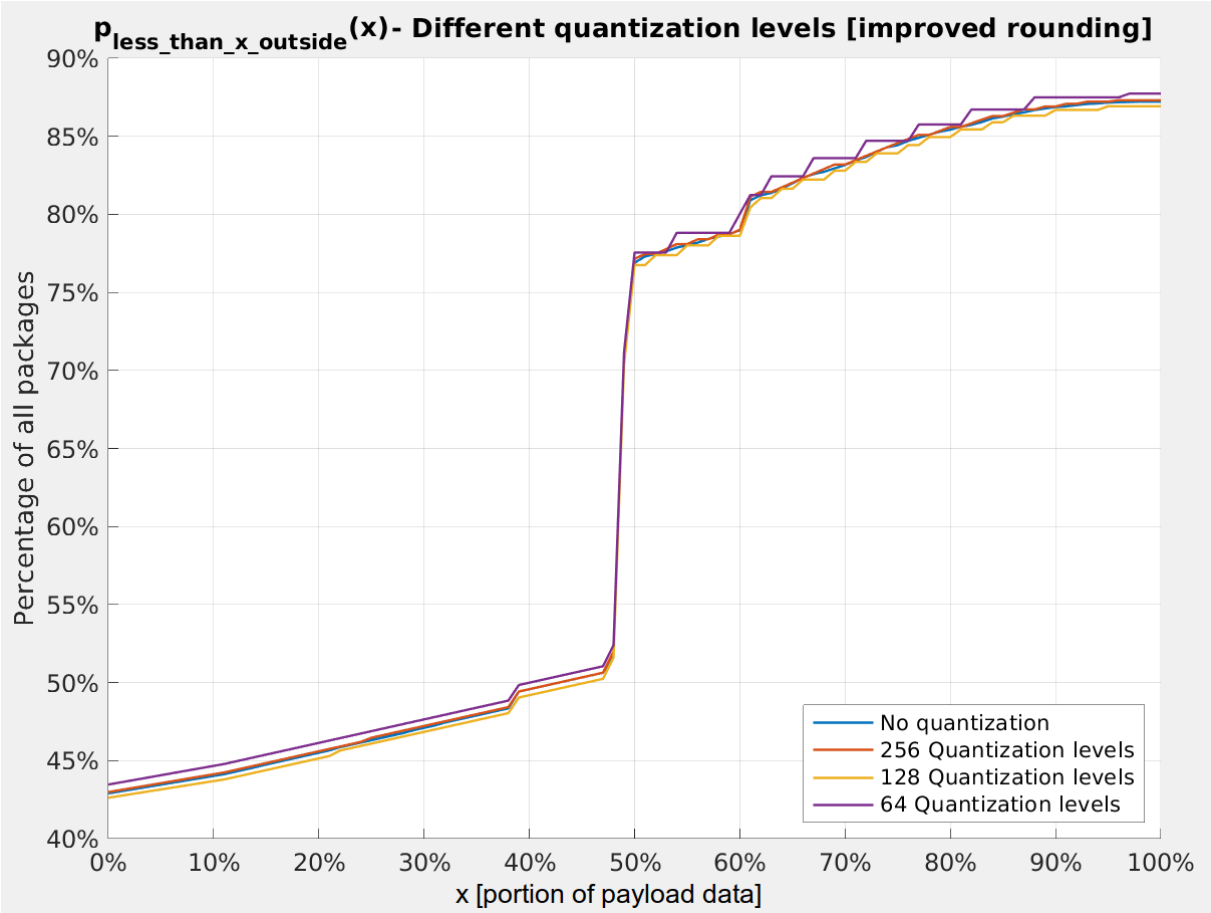

Figure 9.3: Estimates of $p_{less\_than\_x\_outside}(x)$, using an opportunity distribution that is rounded with the procedure described in section 7.2.1.

9.3 Effect of changing the scaling factor s

In the previous estimates for $p_{less\_than\_x\_outside}(x)$, a scaling factor $s = 1.01$ was used as a parameter for computing the opportunity distribution. In other words, the duration intervals for which the opportunity distribution is calculated, are very narrow. This leads to more precise estimates, but it also means that the opportunity distribution is calculated for many bins, which increases the size of the measurement in bytes. It is therefore advantageous to use as low a scaling factor as possible, while still maintaining sufficient precision in the estimates. To investigate how the choice of scaling factor affects the precision, we repeat the procedure described in section 9.1, but with different values for the scaling factor s. The values that will be tested, are $s = \{1.01, 1.015, 1.03, 1.05, 1.10, 1.15, 1.30, 1.50, 2.0\}$. As described in section 8.5.1, upsampling the opportunity distribution before computing the estimates will lead to more accurate estimates when the scaling factor is large. Therefore, the opportunity distribution is upsampled with an upsampling factor such that the upsampled opportunity distribution has the same, or better, resolution as the opportunity distribution one would get if a scaling factor $s = 1.01$ was used. Figures 9.4 - 9.6 show the result of these tests.

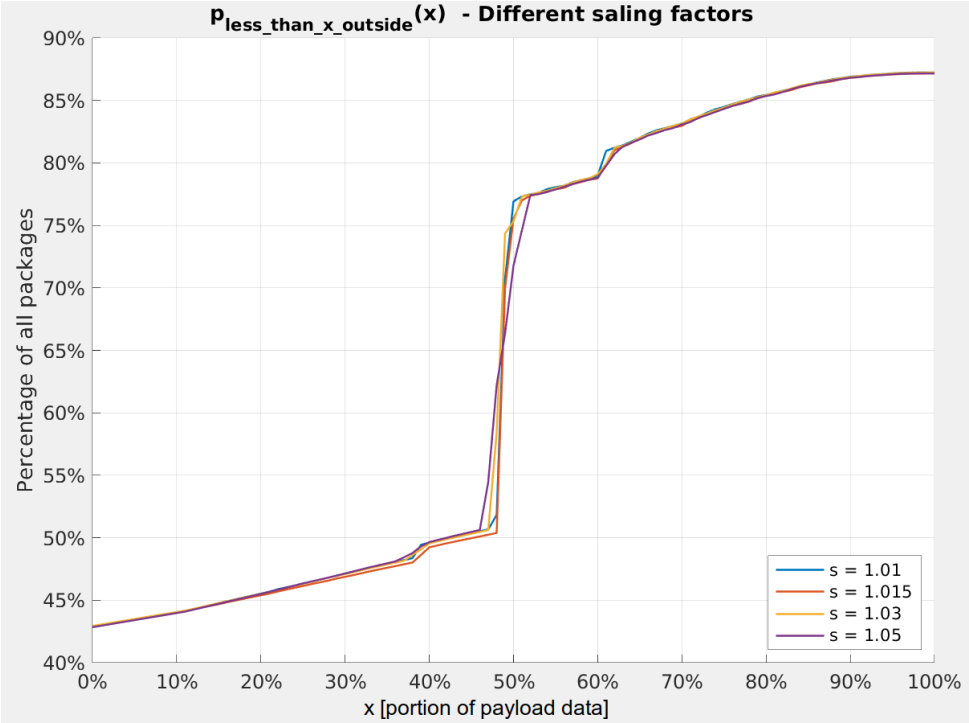

Figure 9.4: Estimates of $p_{less\_than\_x\_outside}(x)$, using scaling factors $s = \{1.01, 1.015, 1.03\}$

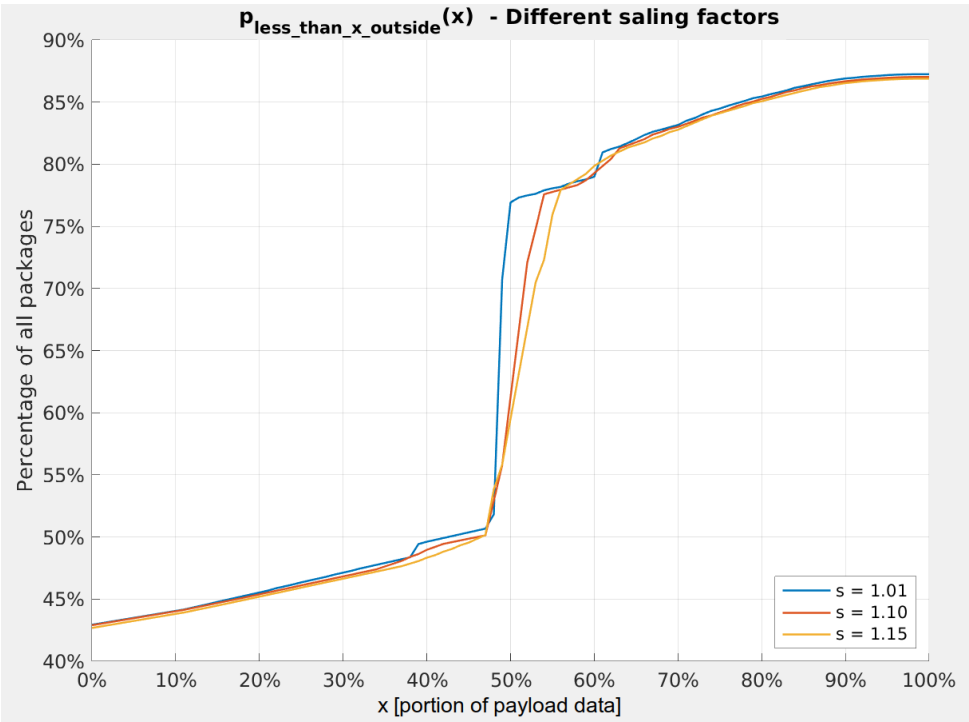

Figure 9.5: Estimates of $p_{less\_than\_x\_outside}(x)$, using scaling factors $s = \{1.10, 1.15\}$

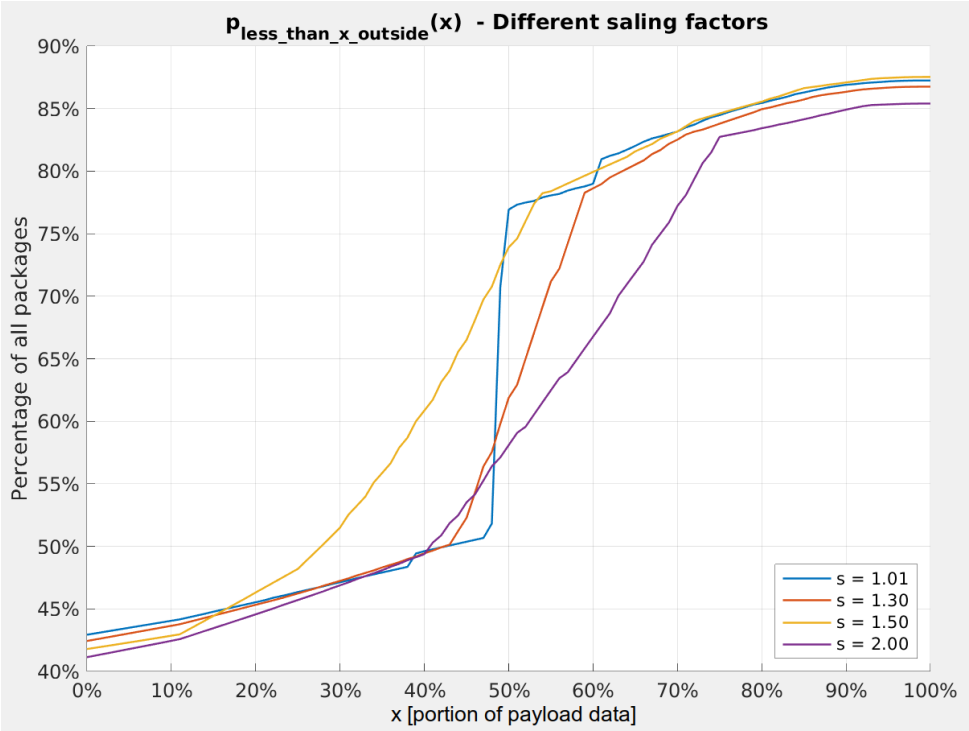

Figure 9.6: Estimates of $p_{less\_than\_x\_outside}(x)$, using scaling factors $s = \{1.30, 1.50, 2.00\}$

To illustrate how important the choice of s is for the data budget, table 9.1 shows how many bins the opportunity distribution has for the duration, for each of the scaling factors tested. The data size of a measured opportunity is also shown, assuming each value is represented with one byte, and that the power bins go from 0 dB to 66 dB with 3 dB steps in between.

Scaling factor	Number of duration bins in the opportunity distribution	Data size
1.01	1419	31.87 kB
1.015	949	21.31 kB
1.03	478	10.74 kB
1.05	290	6.51 kB
1.10	149	3.35 kB
1.15	101	2.27 kB
1.30	54	1.21 kB
1.50	35	805 B
2.00	21	483 B

Table 9.1

We see that the estimate for $p_{less\_than\_x\_outside}(x)$ deviate from the true probability when the scaling factor is chosen large Any scaling factor $s <= 1.05$ seems to have no significant negative effect on the precision of the estimate, while for scaling factors larger than this, some deviations in the estimate start to appear, with deviations increasing with the value of the scaling factor. The deviations also seem to appear mostly for the x-values where there is a rapid change in $p_{less\_than\_x\_outside}(x)$.

Table 9.1 shows that a small change to the scaling factor greatly affects the data budget. One can therefore argue that the scaling factor should be chosen no larger than what is absolutely necessary, and that one may even choose a scaling factor that is so small that it introduces some deviation in the estimate of $p_{less\_than\_x\_outside}(x)$, simply because the reduction in data size is so significant. Because of this, we select $s = 1.30$ as the scaling factor to use. This scaling factor does introduce deviations in the estimate, but those deviations are accepted to reduce the data size.

9.4 Effect of changing the minimum duration in the opportunity distribution, $d_{w_{min}}$

The estimates above were derived from an opportunity distribution where the shortest duration that the opportunity distribution was calculated for, $d_{w_{min}}$, was set to $\frac{1}{f_s} = 7.44\,\mu$s. There was no motivation for this choice, other than that $d_{w_{min}}$ should not be too large. In this section, we will investigate how the choice of $d_{w_{min}}$ affects the estimate for $p_{less\_than\_x\_outside}(x)$.

From equation 8.16, we see that to estimate $p_{less\_than\_x\_outside}(x)$, we only use the bins in the opportunity distribution where the duration d_w is greater than or equal to the duration of the header of the communication packet, d_h. The estimate for $p_{less\_than\_x\_outside}(x)$ is therefore not affected by the choice of $d_{w_{min}}$, as long as it is less than d_h. To see how larger choices for $d_{w_{min}}$ affect the estimate, we repeat the procedure described in section 9.1, but with different values for $d_{w_{min}}$. As the header duration is $d_h = 0.5\,$ms, the values for $d_{w_{min}}$ that will be tested, are $d_{w_{min}} = \{0.50\,\text{ms},\ 0.75\,\text{ms},\ 1.00\,\text{ms}\}$. Figure 9.7 shows the estimates for these parameters.

It is clear from figure 9.7 that when $d_{w_{min}}$ is larger than d_h, we will have a severe deviation in the estimate for $p_{less\_than\_x\_outside}(x)$. The choice of $d_{w_{min}}$ is however not critical for the data budget, as changing it from e.g. 1.0 ms to 0.5 ms only increases the size of the opportunity distribution by 7.6%, given that the other parameters are as described in section 9.1. When choosing $d_{w_{min}}$, the most important consideration is therefore that it is not larger than the smallest value which one might want to use for d_h. In section 6.2, we found that choosing $d_{w_{min}} = 0.1\,$ms would satisfy this, as the header sequence cannot be shorter than this due to restrictions on communication systems in the UHF-band. We will therefore select $d_{w_{min}} = 0.1\,$ms as a suitable value.

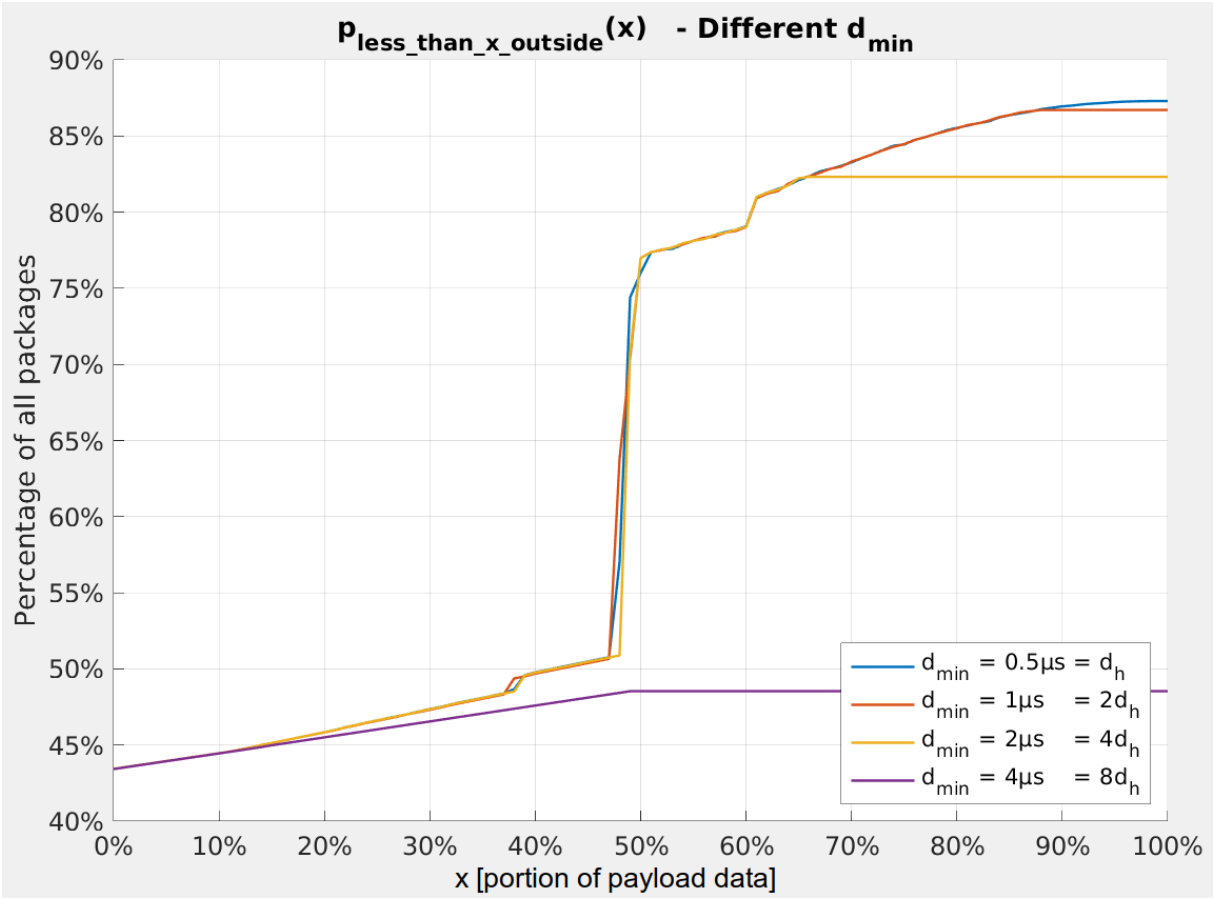

Figure 9.7: Estimate of $p_{less\_than\_x\_outside}(x)$, for different choices of $d_{w_{min}}$, with $d_h = 0.5\,$ms for reference.

9.5 Effect of changing the maximum duration in the opportunity distribution, $d_{w_{max}}$

Similarly to investigating what effect the choice of $d_{w_{min}}$ has on the precision of the estimates, we will also test the choice of $d_{w_{max}}$. We can expect that the precision of the estimate starts to deteriorate when $d_{w_{max}}$ approaches the duration of the communication packet, d_c, which is set to $d_c = 5\,\text{ms}$ in this verification. The values for $d_{w_{max}}$ that we are testing are: $d_{w_{max}} = \{5\,\text{s},\ 200\,\text{ms},\ 50\,\text{ms},\ 10\,\text{ms}\} = \{1000d_c,\ 40d_c,\ 10d_c, 2d_c\}$. Figure 9.8 shows the estimates for $p_{less\_than\_x\_outside}(x)$ with these parameters.

We see from figure 9.8 that choosing $d_{w_{max}}$ less than $\sim 10d_c$ will significantly affect the estimate for $p_{less\_than\_x\_outside}(x)$ for low values of x, while Choosing $d_{w_{max}} \geq 40d_c$ appears to give a precise estimate. $d_{w_{max}} \in [10d_c, 40d_c]$ result in a somewhat precise estimate, but still with some deviations for low values of x. The choice of $d_{w_{max}}$ affects the data budget in a similar manner to $d_{w_{min}}$, and we remember that a large change in $d_{w_{min}}$ causes only a moderate change in the data budget. The data budget, although still important, is therefore not the main concern when choosing $d_{w_{max}}$, and it is more important that $d_{w_{max}}$ is chosen large enough for the estimate of $p_{less\_than\_x\_outside}(x)$ to be precise.

When $d_c = 5\,\mu\text{s}$, we see that 200 ms, or maybe even 100 ms would be a sufficient choice for d_{max}. But if one wishes to estimate $p_{less\_than\_x\_outside}(x)$ for longer communication packets as well, d_{max} should be larger. However, there is another factor limiting how large d_{max} could be chosen as. If the opportunity distribution that $p_{less\_than\_x\_outside}(x)$ is estimated from is calculated from a radio signal that is e.g. 5 seconds long, one cannot accurately measure statistics of windows with a duration close to 5 seconds, because one does not observe enough of those windows to get an accurate estimate of their frequency. As a rule of thumb, we want to measure 100 occurrences of something before we say that we have an accurate

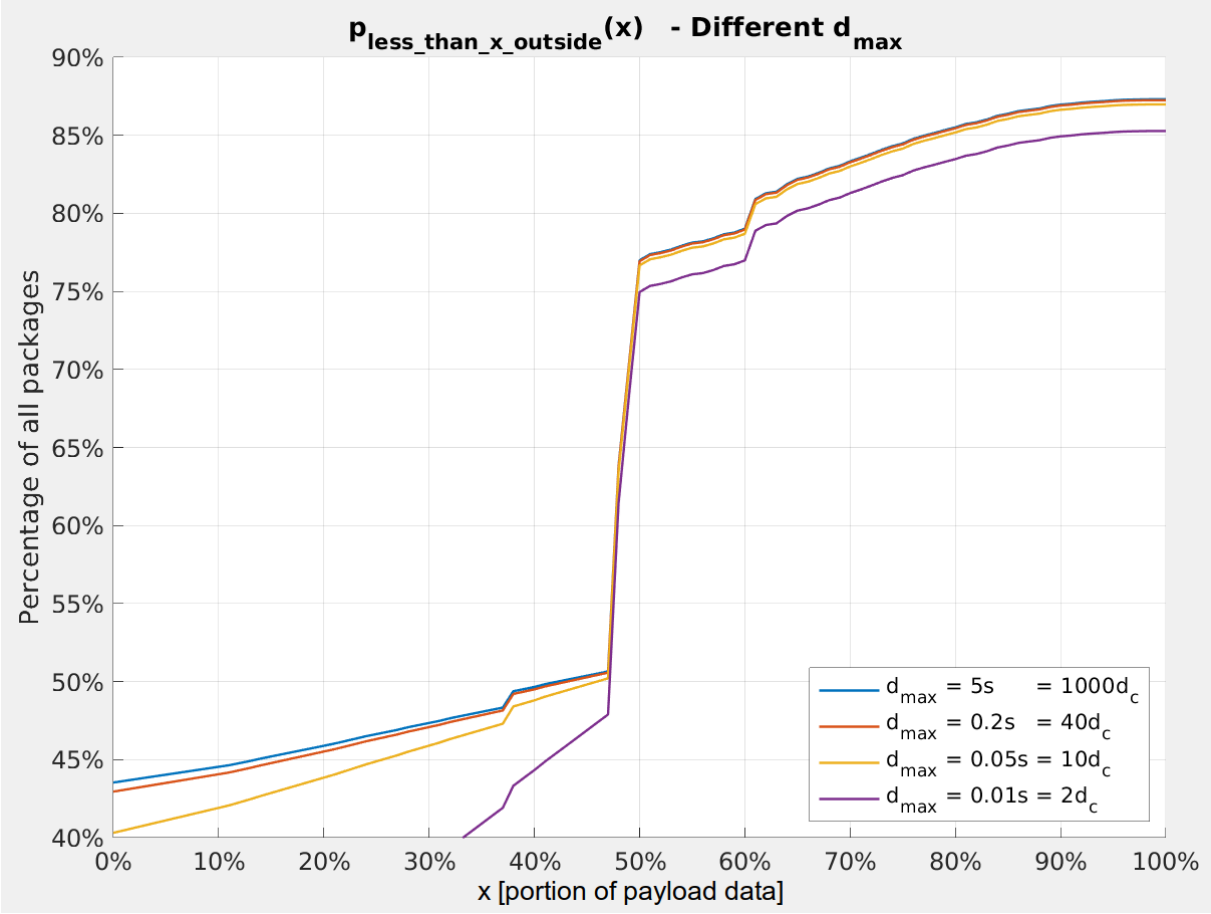

Figure 9.8: Estimate of $p_{less\_than\_x\_outside}(x)$, for different choices of $d_{w_{max}}$, where $d_c = 5\,\text{ms}$ for reference.

estimate of their frequency. If the measured radio signal has a duration of 5 s, this would put d_{max} at 50 ms or less. 50 ms is however shorter than what we require for the estimate of $p_{less\_than\_x\_outside}(x)$ to be precise. We will there, therefore, select $d_{max} = 500$ ms as the maximum opportunity duration to measure for, and accept that the opportunity distribution may not be representative at the bins with the longest duration.

9.6 The combined effect of parameter choices

The effect of changing d_{min}, d_{max}, the scaling factor s, and the number of quantization levels, have now been tested individually. We have concluded that using $d_{min} = 0.1$ ms, $d_{max} = 500$ ms, $s = 1.30$ and $num\_quantization\_levels = 128$ appear to be suitable parameter choices. An estimate of $p_{less\_than\_x\_outside}(x)$, using these parameters for the opportunity distribution, is shown in figure 9.9. It is clear that the estimate for $p_{less\_than\_x\_outside}(x)$ does not overlap as well with the actual distribution as the estimate in figure 9.1a, where more precise parameters were used. The estimate in figure 9.9 is however based on an opportunity distribution that is very small in data size, which is a huge advantage for the LUME-1 mission, due to the low downlink rate. Besides, the average absolute value of the deviation between the estimate and the true distribution is just 1.5 percentage points for the given example, which one can argue is an acceptable difference.

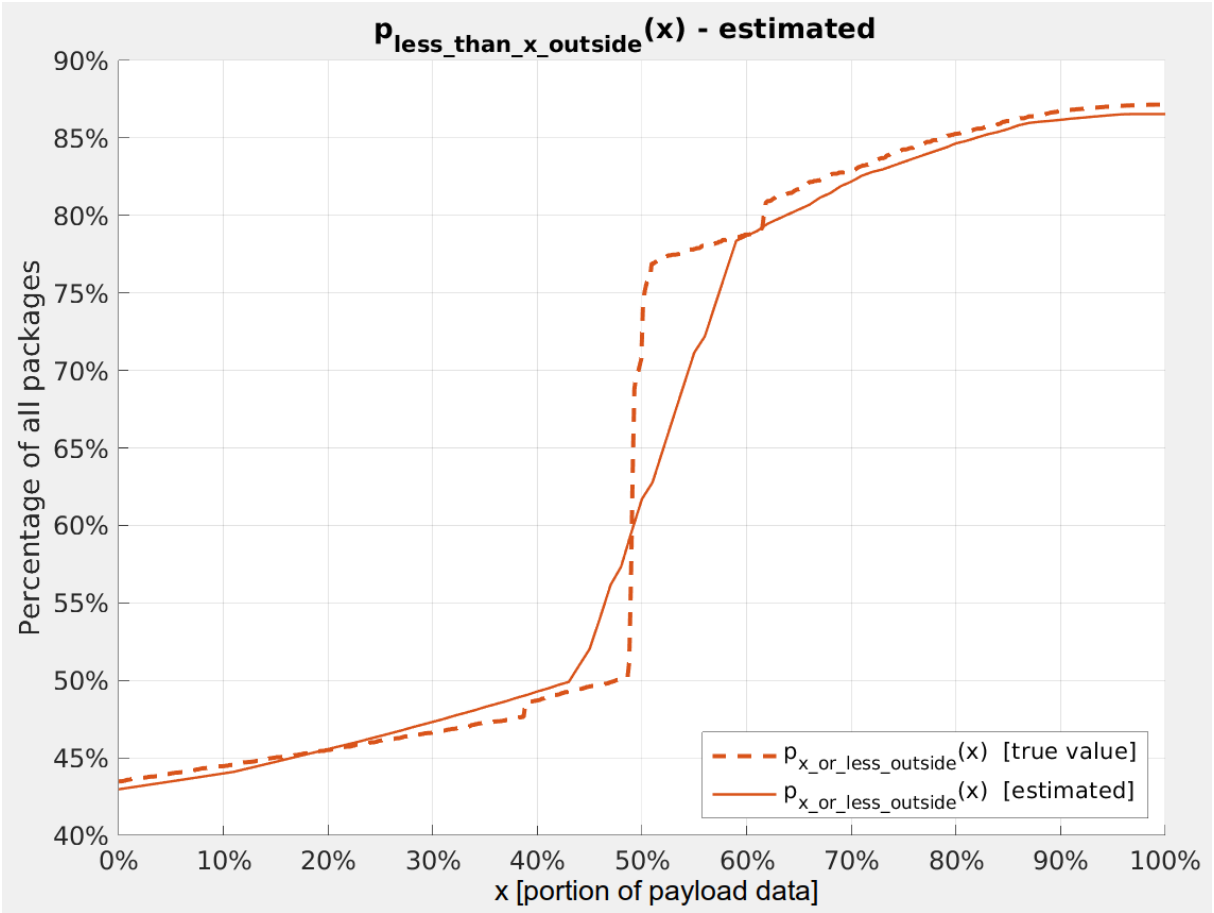

Figure 9.9: Estimate of $p_{less\_than\_x\_outside}(x)$, where $d_{min} = 0.1$ ms, $d_{max} = 500$ ms, $s = 1.30$ and $num\_quantization\_levels = 128$ are used as parameters for the opportunity distribution.

10. *Data budget*

10.1 Estimating compression ratio using Norsat-2 measurements

The opportunity distributions measured on the LUME-1 satellite will be compressed using compression libraries that are pre-installed on the satellite. The specific compression command that is used is `tar -cvf - INPUT_DIR/* | zstd -19 -o OUTPUT_FILE`, which utilizes the tar and zstd compression tools. To find the compression ratio one can expect to get with this, we will use the Norsat-2 data. The Norsat-2 data consists of seven sessions, each lasting ~10 minutes. We divide each session into a series of five-second long snippets, and the opportunity distribution is calculated for each of these snippets. The opportunity distributions are stored on an intermediate format before compression, where one byte, or 8 bits, is used to represent each data point. The measurements in each session are then compressed, and the size of the compressed file is compared to the size of the uncompressed measurements. The opportunity distributions are calculated with the parameters $d_{min} = 0.1\,\text{ms}$, $d_{max} = 500\,\text{ms}$, $s = 1.30$, and $num\_quantization\_levels = 128$. As for the number of power thresholds to calculate the opportunity distribution for, the I- and Q-values go from -128 to 127, such that the maximum value for $I^2 + Q^2$ is $128^2 + 128^2 = 32768 = 45.15$ dB. To cover all possible power levels, the opportunity distribution is calculated for $P = \{0, 3, 6, ..., 45, 48\}$, which is 17 different power thresholds. Figure 10.1 shows the average data size per opportunity distribution for each of the sessions.

The number of bins in the opportunity distribution is 17 in the power dimension, and $1 + \frac{\log\left(\frac{500\,\text{ms}}{0.1\,\text{ms}}\right)}{\log 1.30} = 33$ in the duration dimension. Given that each data point is represented by a value between 0 and 127, one would expect the opportunity distribution to occupy $17 \cdot 33 \cdot \frac{\log_2 128}{8} = 490$ bytes if no compression was implemented other than bit packing. The average data size per opportunity distribution over all the sessions is 105 bytes, which yields an expected compression ratio of $\frac{490}{105} = 4.7$. The compression ratio does however vary a lot across the seven sessions, between 3.7 and 7.8.

10.2 Downlink data budget

As described in section 4.2.3, 525 kB is estimated as an upper bound for how much data can be downlinked from the satellite in a measurement campaign while still fulfilling the mission requirements for downlink time. To find how many measurements can be performed in a measurement campaign, one must know the expected average size of a compressed opportunity distribution in the LUME-1 mission. The measurement parameters used in the LUME-1 mission will be identical to the ones used in the compression test in section 10.1, except for the number of power levels. The number of power levels used in the compression test was 17, while the number of power levels in the LUME-1 campaign will be 23, as described in section 6.2. Assuming the same compression ratio as measured in section 10.1, we can therefore expect a measured opportunity distribution to occupy $104\,\text{B} \cdot \frac{23}{17} = 141\,\text{B}$ in the LUME-1 mission. This means that one can perform $\frac{525 \cdot 1024}{141} = 3802$ measurements during a measurement campaign. One should however aim to produce fewer measurements than this, because the average compression ratio for the particular measurements we get on the LUME-1 satellite may be smaller than expected, and the downlink rate may be slower if there is heavy interference when the satellite passes over the ground station. We will therefore multiply this estimate with a safety factor of 0.8 to decrease the probability that the campaign takes more time to downlink than 14 days. This results in a maximum of $3802 \cdot 0.8 = 3041$ measurements per campaign, or 420 kB.

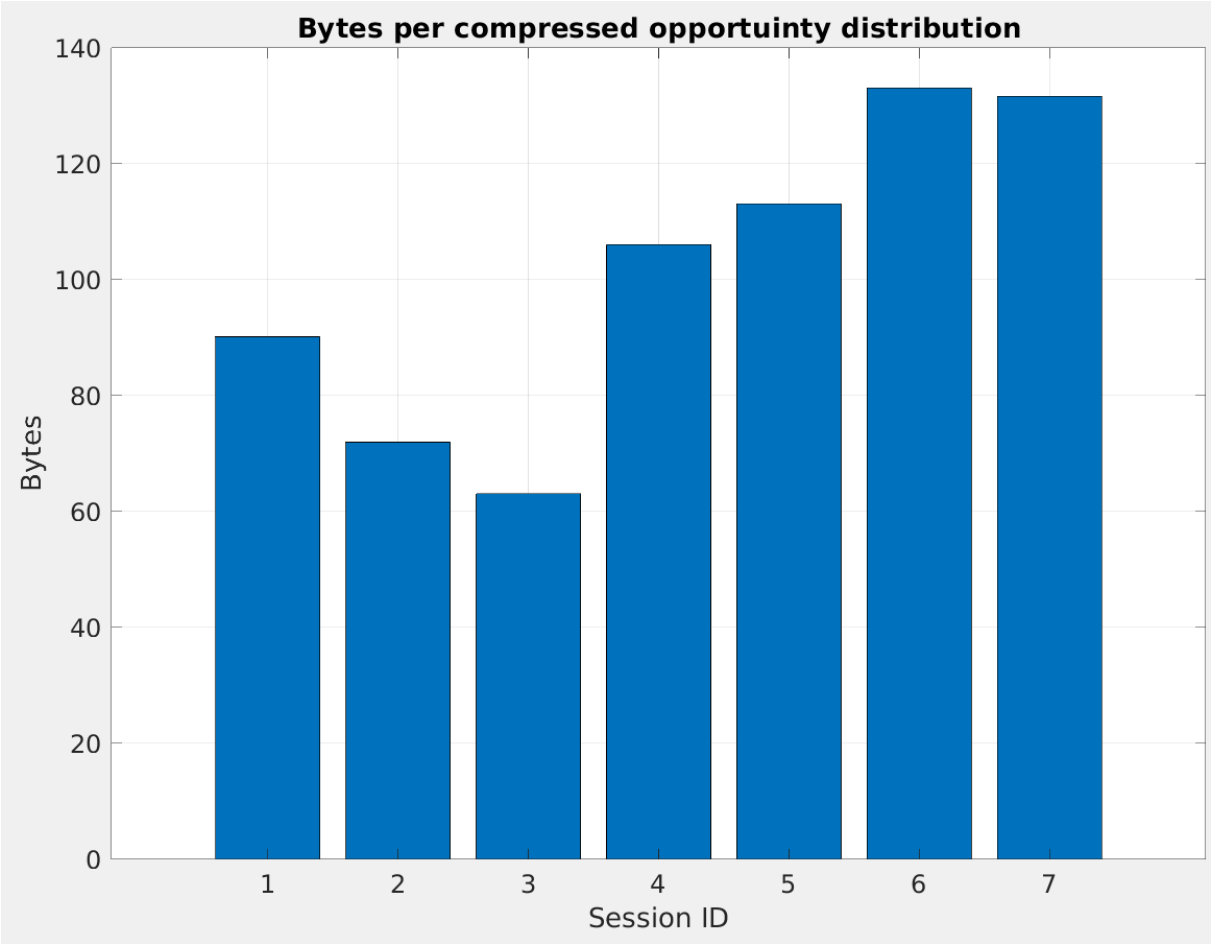

Figure 10.1: Average size of an opportunity distribution after compression

10.3 Uplink data budget

Three files described in section 7 must be uplinked to the satellite in order to perform measurements. These are the `campaign_manager.sh` and `analyze_signal` programs, as well as the FIR-filter `filterConfig.txt`. The sizes of these programs are 5.0 kB, 24.7 kB, and 1.6 kB respectively, which adds to a total of 31.3 kB. Before uplinking the program to the satellite, it will be compressed to reduce its file size. However, we can see that even before compression, the combined file size is well below the 75 kB upper limit imposed by the constraint MR01 in section 4.2.2.

11. *Concept of operations*

Several decisions must be made regarding how to perform a measurement campaign, like how many measurements to perform and how to distribute the measurements in time and space. Sections 11.1-11.4 propose four different options for how to structure a measurement campaign, and in section 11.5, we will discuss the advantages and disadvantages between the different campaign options and select what to do in the actual LUME-1 mission. For each of the options in sections 11.1-11.4, we decide that a measurement campaign only considers one hemisphere. This is due to the constraint TC01 listed in section 4.2.1, which states that the SDR on the satellite can only be on 12 of 24 hours during the day, due to power budget limitations. For all of the four campaign options, every single measurement will be performed by sampling a radio signal that is 5 s long and calculating the opportunity distribution for this signal. 5 s is chosen because previous measurements carried out by G. Q. Diaz [13] use 5 seconds as the measurement duration, and it is easier to compare with those measurements if the same measurement duration is used for this mission.

The measurement parameters to use in the campaign are listed in table 11.1. The slowest available sampling rate on the TOTEM SDR is 521 kHz, which is rounded up to a sampling rate of 600 kHz. As described in section 3.1, 435-438 MHz is devoted to satellite communication as an amateur band, so the carrier frequency is chosen as 435 MHz. This is also used as the carrier frequency in [13], which makes it easier to compare measurements with those. The same applies to the RF filter bandwidth, which

Input parameter	Description
`f_s`	600 kHz
`f_c`	435 MHz
`bw`	200 kHz
`gain`	max gain
`num_samples`	3 MS
`buffer_size`	32768
`power_min_dB`	0
`power_max_dB`	66
`power_step_dB`	3
`time_min`	0.0001 [0.1 ms]
`time_max`	0.5 [500 ms]
`time_step`	1.30
`num_quantization_levels`	128
`num_orbits`	1
`orbit_period`	0 [not used]
`num_meas_blocks`	[depends on campaign option]
`meas_block_period`	[depends on campaign option]
`num_meas`	[depends on campaign option]
`program_path`	[Directory to where the `analyze_signal` program and the FIR-filter is located on the LUME satellite]
`output_path`	[Directory where measurement results should be saved on the LUME satellite]

Table 11.1: Summary of final choice for input parameters in the LUME mission.

is chosen to be 200 kHz. This is the minimum available bandwidth with the TOTEM SDR, and the reason for not using a larger bandwidth is that it should be less than half the sampling rate in order to comply with the Nyquist limit. The hardware gain of the TOTEM SDR should be set as the maximum gain, because signals received by the satellite will be attenuated by the long travel distance. To get a measurement duration of 5 s, the number of samples is set to 3 000 000. The buffer size is left at the default of 32768, but is of less importance, as it does not affect the measurement output. As described in section 6.2, `power_min_dB`, `power_max_dB` and `power_step_dB` are set to 0 dB, 66 dB and 3 dB respectively. The parameters `time_min`, `time_max`, and `time_step` are set to 0.1 ms, 500 ms, and 1.30 respectively, as chosen in chapter 9. Likewise, `num_quantization_levels` is chosen as 128. The `program_path` and `output_path` should be set to the directory on the LUME satellite where the measurement program is located, and the directory where the measurement output is to be placed.

Regarding the `num_orbits` parameter, this must be set to 1 due to the constraint SC10 from section 4.2.1, so to measure for multiple orbits, several calls to `campaign_manager` should be scheduled. This means that the `orbit_period` parameter is unused. As for the remaining parameters, they depend on the campaign configuration and are given for each campaign option in sections 11.1-11.4, in tables 11.2-11.5. In these tables, we will also list the values that the parameters `num_orbits` and `orbit_period` would have if we were to schedule the entire campaign with a single call to `campaign_manager.sh`.

11.1 Option 1

To get a measure of how fast the interference properties in the communication channel changes over time and space, many measurements are performed back-to-back. To compensate for the increased data size, only a small geographic location is considered. Since the Arctic area is the main area of interest, the measurements are performed over this region. Specifically, measurements are performed while the satellite is flying over three southbound tracks originating from the North Pole, as shown in figure 11.1. One track goes in the direction from the US to Russia, another from Russia to the US, and a third one goes over the Greenland Sea between Svalbard and Greenland. The measurements are performed between latitude 90°N to 60°N, which takes the satellite 7.87 min to cover. As each measurement lasts 5 seconds, 94 measurements are performed during this track, such that the total number of measurements in all three tracks combined becomes 282. Given the compression ratios found in chapter 10, this is expected to produce $282 \cdot 141\,\text{B} = 39\,\text{kB}$, which has an expected downlink time of 1.0 days. For comparison, the expected downlink time without compression would be 4.9 days. The input parameters to `campaign_manager.sh` specific to this option are listed in table 11.2.

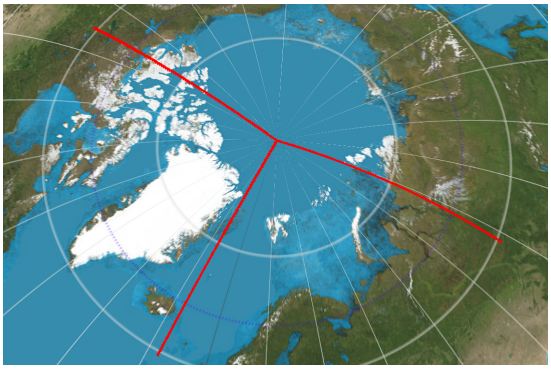

Figure 11.1: Campaign option 1

Input parameter	Description
`num_orbits`	3
`orbit_period`	$94.4 \cdot 60 \cdot 3$
`num_meas_blocks`	1
`meas_block_period`	0 [not used]
`num_meas`	94

Table 11.2: Input parameters for option 1.

11.2 Option 2

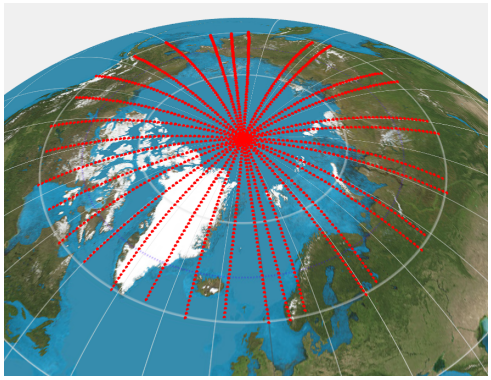

Figure 11.2: Campaign option 2

As the Arctic region is the region of most interest, we will measure the entire time when the satellite is in the geographic area above 60°N. The 60° limit comes from the latitude of the southernmost point of Norway rounded up, and it ensures that all areas that may be of interest are covered. As the Earth rotates 180° in 12 hours, this is how long it takes for the satellite to cover the entire area. This would, however, cover half the area in northbound tracks, and half of it in southbound tracks. The interference from particularly military radar systems is dependent on the satellite's direction of movement, so it is desirable to measure the same location with both a southbound track and a northbound track. Therefore, 24 hours are needed to cover the area with both northbound and southbound tracks. From section 4.2.1, we have that that the LUME-1 satellite's orbit lasts for 94.4 minutes, such that 24 hours corresponds to $\frac{24 \cdot 60}{94.4} = 15.3$ orbits, which is rounded up to 16 orbits in total. A measurement is performed every 10 seconds while above 60°N. Each orbit, the satellite spends $\frac{2 \cdot (90° - 60°)}{360°} \cdot 94.4\,\text{min} = 15.7\,\text{min}$ at latitudes above 60°N, which gives $\frac{15.7\,\text{min}}{10\,\text{s}} = 94$ measurements per orbit. Thus, the campaign contains $94 \cdot 16 = 1504$ measurements in total. Given the compression ratio found in chapter 10, this is expected to produce $1504 \cdot 141\,\text{B} = 207\,\text{kB}$, which has an expected downlink time of 5.5 days. For comparison, the expected downlink time without compression would be 26 days. Figure 11.2 shows an illustration of how measurements are distributed in space for this option, and the input parameters to `campaign_manager.sh` specific to this option are listed in table 11.3.

Input parameter	Description
`num_orbits`	16
`orbit_period`	$94.4 \cdot 60$
`num_meas_blocks`	94
`meas_block_period`	10
`num_meas`	1

Table 11.3: Input parameters for option 2.

11.3 Option 3

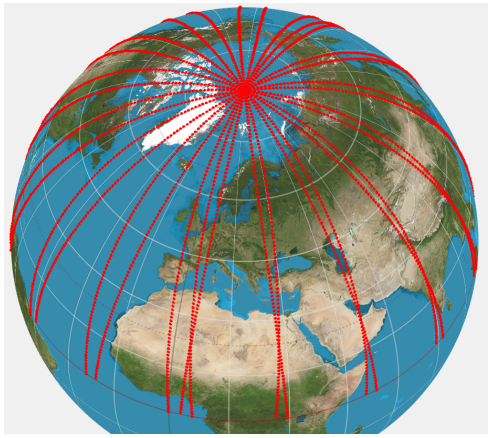

Figure 11.3: Campaign option 3

The Arctic area is obviously the most useful area for obtaining information useful for designing a communication system between Arctic sensor nodes and satellites. However, we also have a motivation for performing measurements outside the Arctic region. The LUME-1 satellite is operated by Alén Space and the Spanish university *Universidade de Vigo*. This gives us an incentive to produce measurements that are also useful for them. Additionally, many ground stations located in Europe will benefit from measurements of uplink interference outside of the Arctic region, as it could be useful for improving their communication strategy. Hence, the geographic area we will cover with measurements is chosen to be an entire hemisphere. Similarly to in option 3, we will measure for 16 orbits in order to cover all areas with both a northbound and southbound track. Due to the increased area to measure the interference over, we will reduce the frequency of the measurements from once every 10 s to once every 15 s. In each orbit, 47.2 minutes are spent in the hemisphere of interest, which gives a total of $\frac{47.2\,\text{min} \cdot 60\,\text{s/min}}{15\,\text{s}} = 189$

measurements per orbit, or $189 \cdot 16 = 3024$ measurements in the entire measurement campaign. Given the compression ratio found in chapter 10, this is expected to produce $3024 \cdot 141\,\mathrm{B} = 416\,\mathrm{kB}$, which has an expected downlink time of 11.1 days. For comparison, the expected downlink time without compression would be 52 days. Figure 11.3 shows an illustration of how measurements are distributed in space for this option, and the input parameters to `campaign_manager.sh` specific to this option are listed in table 11.4.

Input parameter	Description
`num_orbits`	16
`orbit_period`	$94.4 \cdot 60$
`num_meas_blocks`	189
`meas_block_period`	15
`num_meas`	1

Table 11.4: Input parameters for option 3.

11.4 Option 4

This option is a variation of option 3, aimed at having more resolution in the longitude direction. Instead of a 15 s period between measurements, a 60 s period is used, which results in $\frac{47.2\,\mathrm{min} \cdot 60\,\mathrm{s/min}}{60\,\mathrm{s}} = 47$ measurements per orbit. However, instead of measuring over 16 orbits, we will measure over 64 orbits, such that the total number of measurements in the campaign becomes $47 \cdot 64 = 3008$. This is expected to produce $3008 \cdot 141\,\mathrm{B} = 414\,\mathrm{kB}$, which has an expected downlink time of 11.0 days. For comparison, the expected downlink time without compression would be 52 days. Figure 11.4 shows an illustration of how measurements are distributed in space for this option, and the input parameters to `campaign_manager.sh` specific to this option are listed in table 11.5.

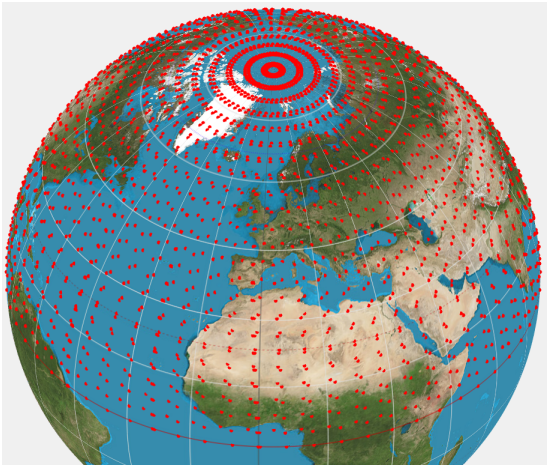

Figure 11.4: Campaign option 4

Input parameter	Description
`num_orbits`	64
`orbit_period`	$94.4 \cdot 60$
`num_meas_blocks`	47
`meas_block_period`	60
`num_meas`	1

Table 11.5: Input parameters for option 4.

11.5 Options for measurements using LUME-1 satellite

Both options 1 and 4 are chosen for use in the LUME-1 measurements. Option 1 will be used as an initial pre-campaign, while the primary measurement campaign will be carried out according to the description of option 4. This section highlights advantages and disadvantages of the different campaign options, and motivates the choice of options 1 and 4.

First, option 1 has the advantage that it provides continuously sampled measurements over the Arctic area, which gives information about the temporal variation of the interference. Additionally, the time to downlink a campaign from option 1 is only one day, which reduces the time between measurements are taken and analyzed. This makes it possible to use option 1 as a complementary campaign to one of the campaigns described in options 2-4.

As for option 2 versus option 3, option 2 provides the most measurements over the Arctic area, which could be very useful for designing a communication system. The shorter expected downlink time of 5.3 days is also an advantage with option 2. However, as described in section 11.3, option 2 does not provide any benefit to the owners of the satellite that is in use. Option 3 is better in this regard, and both Universidade de Vigo and the general SmallSat community could benefit from measuring outside the Arctic region. This motivation is, however, secondary to the mission objective itself, which is to provide information that is useful for designing a communication system for the Arctic region, but it is not the only motivation for choosing option 3 over option 2. Measuring over the entire globe makes it easier to compare with measurements that other people have performed previously, as in [10, 11, 13]. Additionally, the difference in measurement period is only 15 s in option 3 compared to 10 s for option 2, which is only a difference of 50%. And even when one covers the entire globe with measurements, as for option 3, the spatial resolution will still be greatest over the Arctic region because the satellite passes over this region more frequently. The main disadvantage of choosing option 3 over option 2 is the fact that the expected downlink time is doubled. Nevertheless, the expected downlink time for option 3 is still less than the upper limit of 14 days, which we decided in section 4.1.

Option 4 is an alternative variation of option 3, with an almost identical downlink time. The main advantage of option 3 is that its measurements are performed often enough to get an understanding of the temporal variation of the interference along the satellite track. Option 4, however, has a more even sampling along the latitude and longitude far from the poles, which is an advantage in the sense that it provides a good spatial resolution. A higher spatial resolution in the Arctic area is still maintained with option 4, as the satellite passes over this area more frequently than other areas.

One thing to note is that because the measurements in each orbit have to be scheduled separately, it will take some time to uplink the scheduling of all the orbits. This is a disadvantage for option 4, since scheduling 64 software calls will take some time. If it is more practical, however, one always has the option of dividing the campaign into, say, four 16-orbit parts, and perform each part separately. This would result in the measurements being separated in time if the parts were performed with some days in between then. Regardless of this, option 4 is chosen due to the advantages of covering the entire globe with measurements, and the even sampling along the latitude and longitude far from the poles.

12. *Verification and testing*

In order to verify that the software described in section 7 operates as intended and that it is mission-ready, it must be tested in a controlled environment. Some verification is already performed in [15], where it was verified that the `analyze_signal` program produces the opportunity distribution as described in section 6. In sections 12.1 and 12.2, test-as-you-fly testing will be performed for campaign options 1 and 4 from section 11 to verify that the `campaign_manager.sh` script works as intended, and to check that there are no issues in the interface between it and `analyze_signal`.

The measurement setup used to test the software is shown in figure 12.1. The software-defined radio NI USRP-2901 is used as a transmitter to emulate interference signals. This radio is controlled by a

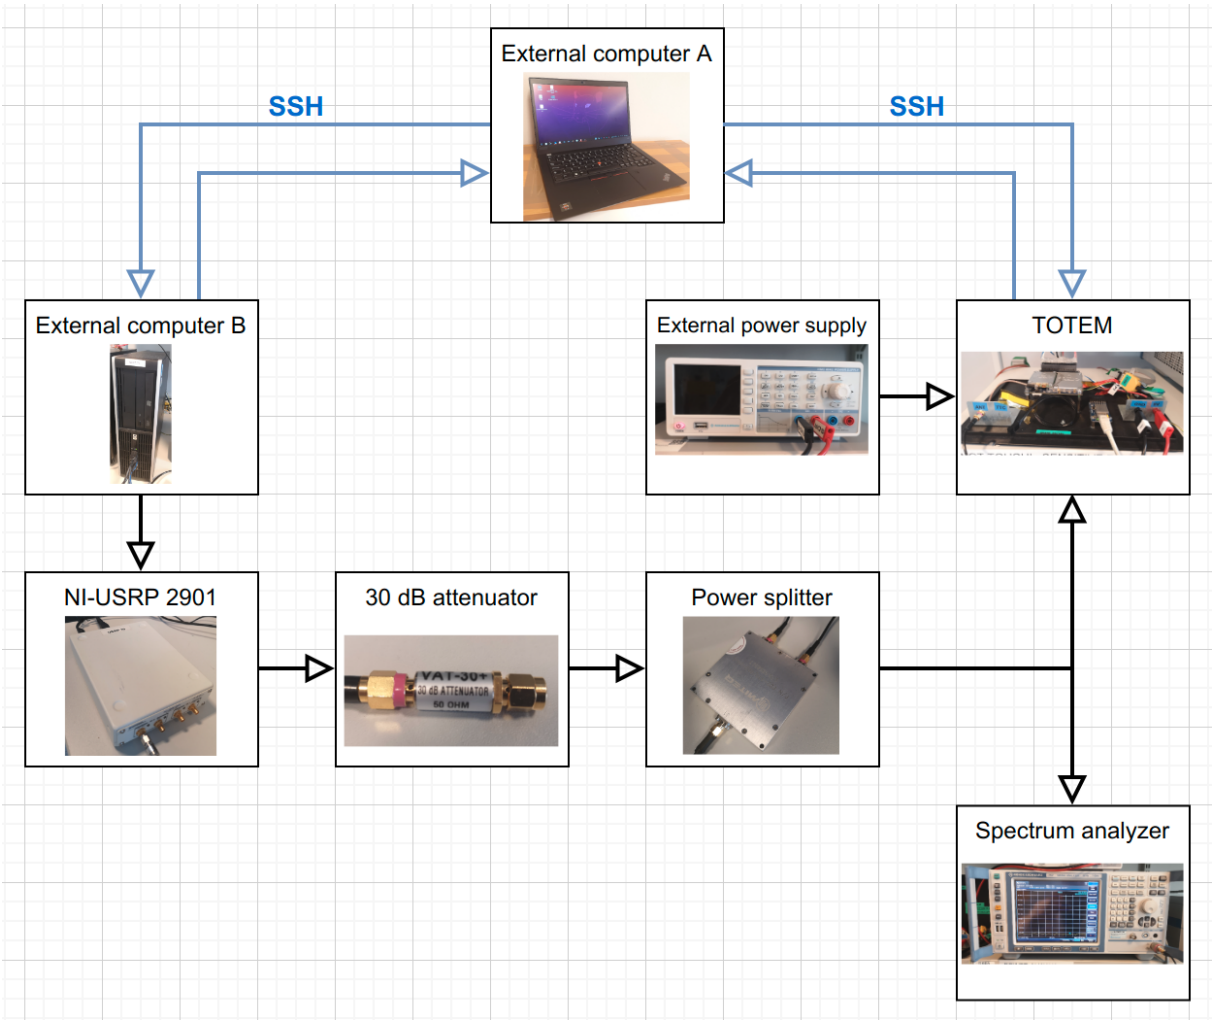

Figure 12.1: Lab setup

computer B, which is again controlled remotely by the external computer A over an SSH connection. The software-defined radio outputs its signal into a power splitter, via a 30 dB attenuator. One of the outputs of the power splitter is fed into a spectrum analyzer, and the other output is fed into a replica of the SDR platform that is used on the LUME-1 satellite. This SDR platform is called TOTEM and is a system on a chip with embedded Linux and programmable logic, plus a wide frequency range transceiver [27]. The 30 dB attenuator is added in order to ensure that the input power to the TOTEM SDR is below its maximum rating, which is 21 dBm. The TOTEM platform is connected to computer A over an SSH connection in order to transfer files to and from the platform and to start the measurement program on the TOTEM board when performing a test. The TOTEM board is powered by an external power supply.

When testing that the developed software works, the NI-USRP 2901 transmits a signal that is configured by computer B, using GNU Radio. The signal is then plotted by the spectrum analyzer to verify that the bandwidth and power spectral density of the transmitted signal is as expected. While the signal is transmitted, computer A sends a single command to the TOTEM platform over SSH, which initiates the `campaign_manager.sh` script with input parameters specified by the command. The signal generated by the NI-USRP 2901 radio is a pulsed complex chirp signal, aimed at emulating a radar of type A from the ITU recommendations in table 3.1. The signal has a pulse duration of 8 ms and a duty cycle of 25%, which gives a pulse repetition frequency of $\frac{1}{8\,\text{ms}/25\%} = 31.25\,\text{Hz}$. The chirp pulse uses a bandwidth of 100 kHz and is centered at 435 MHz. The average spectrum of the transmitted signal over 500 ms is measured with a spectrum analyzer with a resolution bandwidth of 5 kHz, and is shown in figure 12.2.

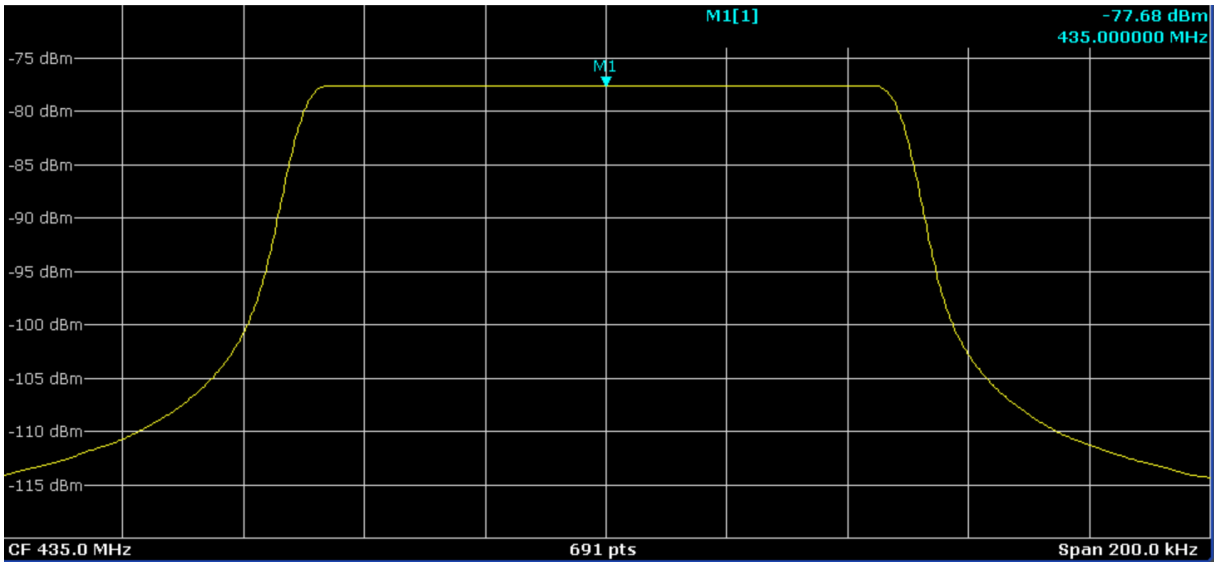

Figure 12.2: Spectrum of the signal generated by the USRP, measured with a resolution bandwidth of 5 kHz and averaged over 500 ms

12.1 Campaign option 1

In this section, the software `campaign_manager.sh` and `analyze_signal` will be tested with the lab setup in figure 12.1, with input parameters corresponding to campaign option 1 from section 11.1. In order to make the software run for multiple orbit times, we will set the input parameters `num_orbits` = 3, and `orbit_time` = 16992, which corresponds to $3 \cdot 94.4$ min.

The campaign runs for approximately 9 hours and 30 minutes and produces three compressed files. Decompressing these files produces 94 measurement files per orbit, or 282 in total. First, to check that the measurements are scheduled at the correct times, we use that the file names indicate when each measurement was performed. Figure 12.3a shows an illustration of the measurement times. Another way to visualize this is to separate the different orbits, and for each orbit, plot when the measurements in that orbit were performed in relation to the orbit start time. This is shown in figure 12.3b, where each rectangle represents one measurement. We see that measurements are only performed in every third orbit

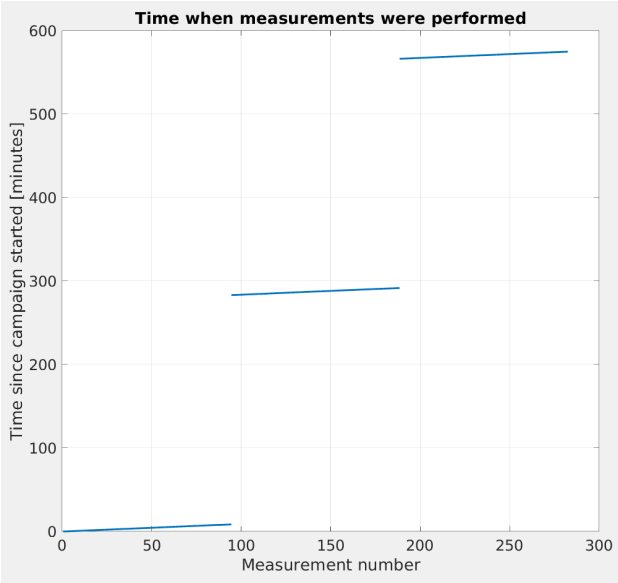

(a) Time when each measurement was performed in the campaign.

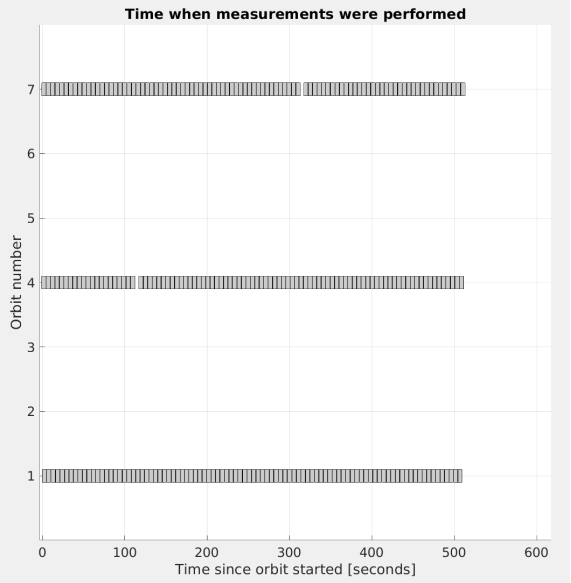

(b) Time when each measurement was performed, for each orbit.

Figure 12.3: Measurement times for campaign option 1.

and that three orbits are measured in total, as expected. As 94 measurements are performed in each orbit, where each measurement is 5 s long, we would expect the last measurement to be finished at time $94 \cdot 5\,\text{s} = 470\,\text{s}$ in every orbit. Instead, we see that the last measurement ends around 510 s, which means that the average time between two measurements is $\frac{510\,\text{s}}{94} = 5.4\,\text{s}$. This would affect the campaign by spacing the measurements slightly more out in time, resulting in a coverage of latitudes 90°N to 57.4°N instead of 90°N to 60°N. The 0.4 s delay between the measurements may be caused by the setup of the SDR before and after fetching the radio signal, or be related to the process of writing the measured signal to a file. An unexpected observation from figure 12.3b, is the extended delay between two measurements at time 110 s in the fourth orbit, and time 315 s in the seventh orbit. It appears that sometimes, the time between two consecutive measurements is significantly longer than the average. This behavior is investigated further in section 12.3.

To verify that the opportunity distribution is calculated correctly in the campaign, one of the opportunity distributions from the measurement campaign is plotted in figure 12.4. The fact that the measured signal is a chirped signal does not affect the opportunity distribution, as it is calculated from the absolute value of the measured radio signal, and the absolute value of a complex chirp is constant. We can therefore see that the measured distribution resembles the example opportunity distribution in figure 6.2 in section 6.2.1, which was also calculated from a pulsed signal. Furthermore, we see from the distribution in figure 12.4 that the opportunity windows between the pulses is measured to be between 19 ms and 25 ms. As the emulated interference signal was configured with a pulse duration of 8 ms and a duty cycle of 25%, this coincides well with the expected duration of $8\,\text{ms} \cdot \frac{1-25\%}{25\%} = 24\,\text{ms}$.

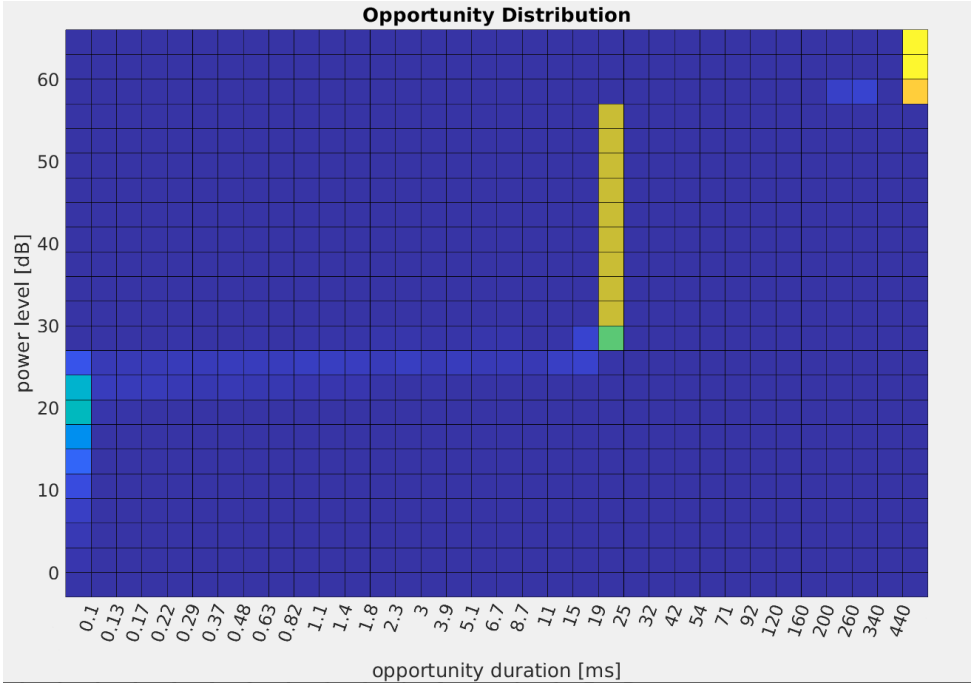

Figure 12.4: An opportunity distribution measured in the test campaign for campaign option 1.

12.2 Campaign option 4

In order to verify that the developed software works also with the parameter configuration for campaign option 4, the test in the section above is repeated, but with input parameters corresponding to those of campaign option 4. The results from these tests are shown in figure 12.5.

As shown in figure 12.5a, the campaign lasts for a total of 100 hours, since 64 orbits lasting 94.4 min each add up to 100 hours in total. The software was configured to perform 47 measurements in each orbit for 64 orbits, with one minute between the measurements. We can see that this is what actually happened from figure 12.5b, as one can observe that 64 orbits were performed and that there are 47 measurements in each orbit, spaced out exactly one minute apart. Unlike for campaign option 1, no timing discrepancies are observed as all measurements appear to be performed at the correct time. It may still be the case that some measurements take longer to perform than the others, as we saw in figure 12.3b for campaign option 1, but because there is a 60-second buffer between each measurement, this does not cause the preceding measurements to be delayed.

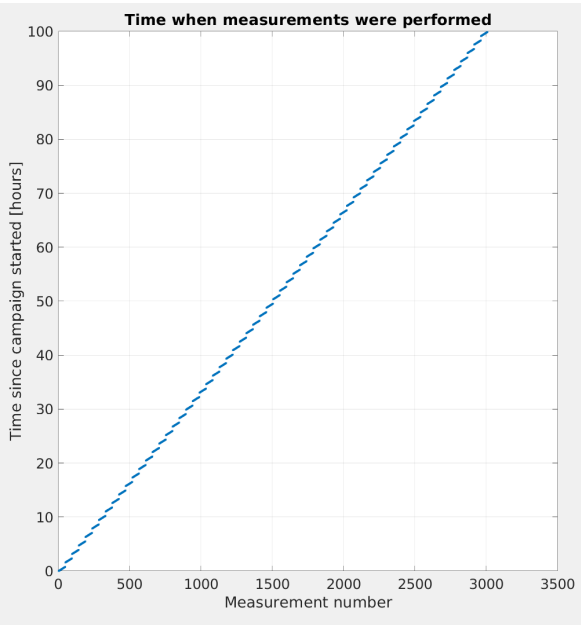

(a) Time when each measurement was performed in the campaign.

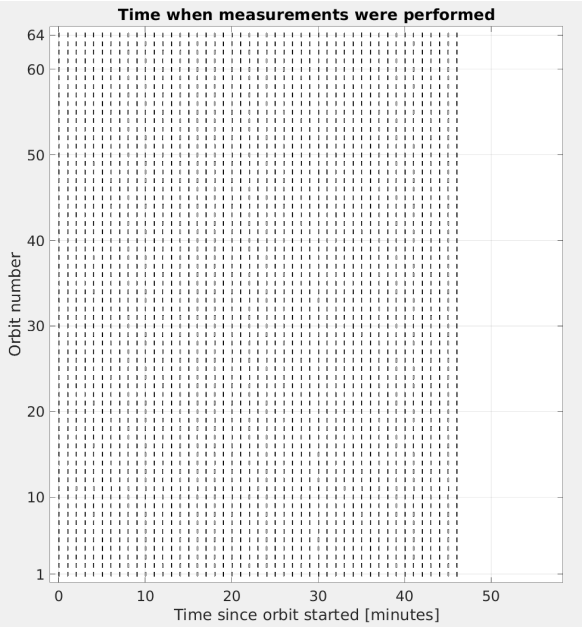

(b) Time when each measurement was performed, for each orbit

Figure 12.5: Measurement times for campaign option 4.

12.3 Delay between radio fetches

When several measurements were performed back-to-back in section 12.1, it was found that in some cases, the time between two measurements is significantly longer than the average duration. This is demonstrated and investigated further by running a script on the TOTEM platform with a loop containing the following three steps:

1. Measure system time
2. Fetch radio signal from TOTEM SDR
3. Delete fetched radio signal

This loop runs for 1000 iterations, and the duration of the radio signal to fetched was 1 second. This is also repeated with a fetch duration of 2 and 4 seconds, to see if the fetch duration affects it. Figure 12.6a shows the time used to perform each fetch, and we can see that the duration between consecutive fetches is usually equal to the duration of the fetched radio signal, plus a constant delay of almost 270 ms. We do, however, also see some rare instances where the fetch takes significantly more time to perform than usual, and this excess delay appears to increase with the duration of the fetched radio signal.

As described in section 7, fetching a radio signal is performed by using the libiio package from Analog Devices [26], which reads the sampled radio signal to a temporary binary file. To investigate whether the excess delay is related to the process of writing the radio signal to the temporary file, the test described in the paragraph above is repeated, but instead of writing the output to a temporary file that is then deleted, it is written to `/dev/null`, which is a null device that immediately discards anything that is written to it. Figure 12.6b shows the time used to perform each fetch in this case, and we can see that the long, sporadic delays are not present anymore. This tells us that those delays occur while the library writes the fetched radio signal to a file. The writing is done continuously while the radio signal is being fetched.

Fetching the radio signal is handled entirely by the libiio library, which does not provide any output other than the fetched radio signal itself. Due to the lack of feedback from the library utility, it is difficult to further determine the cause of the delay. A possible explanation is that the fetching of the radio signal fails mid-fetch, causing it to start over to ensure that the fetched radio signal is continuous. The reason for

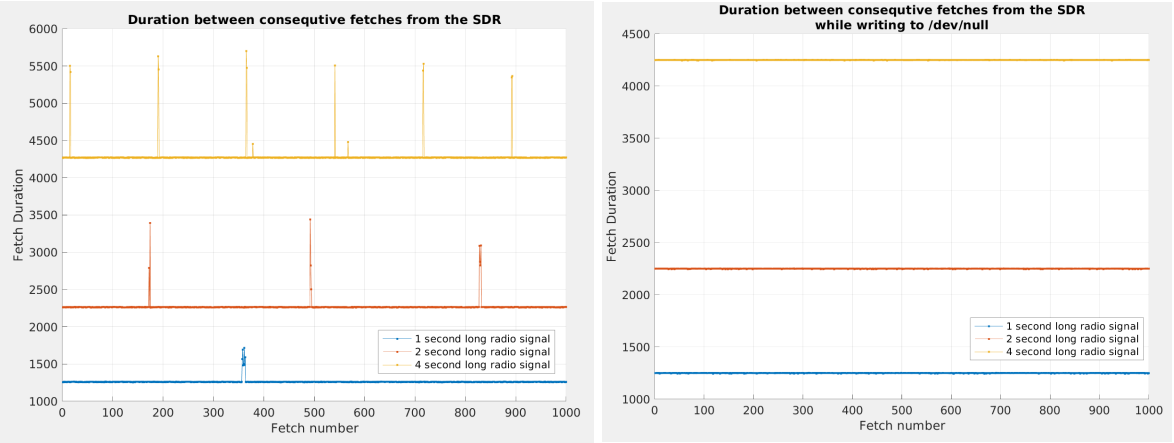

(a) Delay between consecutive fetches from the SDR, when writing the fetch to a temporary file.

(b) Delay between consecutive fetches from the SDR, when writing the fetch to `/dev/null`.

Figure 12.6: Delay between consecutive fetches from the SDR.

the fetch fail could for example be a process running on the TOTEM which interrupts the fetching. This would explain why the excess delay increases with the duration of the fetched signal, as the delay would correspond to the duration of the radio signal that was fetched before the interrupt occurred. It could also explain why the fetches that take more time to perform appear to be clumped together because all the fetches executed while the interrupting process runs would be delayed. It does, however, not explain why the delays do not occur when writing to `/dev/null`. No precise conclusion can therefore be made regarding the cause of these delays.

13. *Proposal for suited communication system*

The end goal when measuring in-orbit interference is to provide information that is useful for designing a communication system fitted to the interference environment. It is appropriate to investigate some options for how such a communication system can be designed. Sections 13.1, 13.2, and 13.3 describe three proposals for how to use the information that can be extracted from the opportunity distribution to design a communication system. These proposals are under the assumption that one has available in-orbit measurements of the opportunity distribution, such that design decisions can be made based on those measurements. A detailed description of how the systems might work is not included here, but rather a description of how the measured opportunity distribution be used in the design process.

Some parameters that should be chosen based on the measured interference are the carrier frequency, transmit power, level of coding, duration to do coding over, and length of a communication packet. The selection of carrier frequency will not be discussed here, because the opportunity distribution focuses on temporal properties rather than spectral properties. As mentioned in section 6.1, there may also be little difference between bordering frequency bands if the interference stems from powerful military radars that produce strong interference that saturates the ADC in the radio receiver at the satellite. The choice of carrier frequency will have to be made based on information obtained by other means than the opportunity distribution, for example by using the measurement technique described by G. Q. Díaz in [13].

13.1 Fixed worst-case design

One of the simplest ways to use the information from the opportunity distribution to design a communication system would be to design a non-adaptive, static system. The system would be static in the sense that it does not change its modulation scheme in order to adapt to a changing interference environment. This could be done by designing the system such that it works when the received interference matches what is measured in the majority of the opportunity distributions. Such a design is the easiest to design, and serves as a simple worst-case design, with outages that are acceptable in regions where we need it to work. When designing such a communication system, the transmit power, the length of a communication packet, and the level of coding used must be determined in unity, as they are dependent on each other. Assuming we have a collection of measured opportunity distributions from the interference environment where the communication system is to operate, a way to determine these parameters is described in the paragraph below.

For a communication system with a given transmit power P_{TX}, communication packet duration d_c, and a forward error-correcting code that can correct $l\%$ loss in the payload data, it is possible to calculate an estimate of the outage probability and the expected effective energy per information bit for the system from the collection of measured opportunity distributions. This can be done for a range of different transmit powers, communication packet durations, and coding schemes, such that one can select the configuration that gives rise to an acceptable trade-off between energy per information bit and outage probability. To calculate the outage probability for a given set of P_{TX}, d_c, and x, one would first have to estimate $p_{outside}(x)$ from each of the measured opportunity distributions, as in chapter 9. The interference power threshold P used in the estimation should be set to P_{RX}, which is the transmit power P_{TX} minus the expected propagation loss and other losses. These estimates for $p_{outside}(x)$ are used as worst-case estimates for $p_{lost}(x)$. The average portion of the payload data that is lost to interference would then be $\int_0^1 x \cdot p_{lost}(x)\ \delta x$. A bit simplified, one can say that the communication link would be operational for the

opportunity distributions where less than $l\%$ of the payload data would be lost to interference on average because the coding scheme would then be able to correct the errors. Note that because this is the average over several packages, this assumes that the coding scheme spans over multiple packets in length, such that it can use the packets with little interference loss to reconstruct the packets with much interference loss. This gives us a rough estimate of the outage probability for the system. To calculate the effective energy per information bit, one can divide the transmit power P_{TX} by the bit rate, and then divide the result by the code rate. The outage probability can also be factored into this if one wants to include the energy spent on trying to establish a connection to the satellite during an outage.

This analysis is a simplification, as the interference is assumed to have an on-off behavior, where there is either no interference or much interference, and nothing in between. Furthermore, the use of $p_{outside}(x)$ as a worst-case estimate for $p_{lost}(x)$ is a very crude approximation, as shown in section 9. Nevertheless, it gives a starting point for the choice of parameters for the communication system.

13.2 Geographically dependent system

While the system in section 13.1 is the easiest to implement, it is not the most efficient scheme possible. In areas where there is little interference, such a system would likely use a higher than necessary transmit power and a lower than necessary code rate, while it would tend to lose the communication link in areas with high interference. An improvement over the static system would be to make it geographically dependent. This could be done by dividing the globe into distinct regions and design a separate communication system for each of the regions using the procedure described in section 13.1. It would then be natural to also design one system for when the satellite travels in a southbound track and another one for when it travels in a northbound track.

When designing a system tailored to a certain geographical region, one has less data to use in the system design because the number of available measurements decreases when one limits the geographic area to look at. Due to this, one should not divide the globe into too many areas, as the measurements would then merely show the particular interference realizations measured, rather than the general behavior. Because the interference might depend on numerous different factors, including the time of day, the time of year, the number of satellites nearby, and the current geopolitical tension, it is difficult to obtain enough measurements to claim that they are representative of the general interference behavior. For the purposes of discussion, however, let's say that as a rule of thumb, at least 100 measurements should be performed per area, and these measurements should be spread out across multiple passes. As described in section 11.4, the number of measurements in a measurement campaign is planned to be 3008, where 1504 are northbound and 1504 are southbound. This results in a maximum of 15 regions if we require there to be 100 measurements per region. One way to divide the globe into 15 sections is shown in figure 13.1,

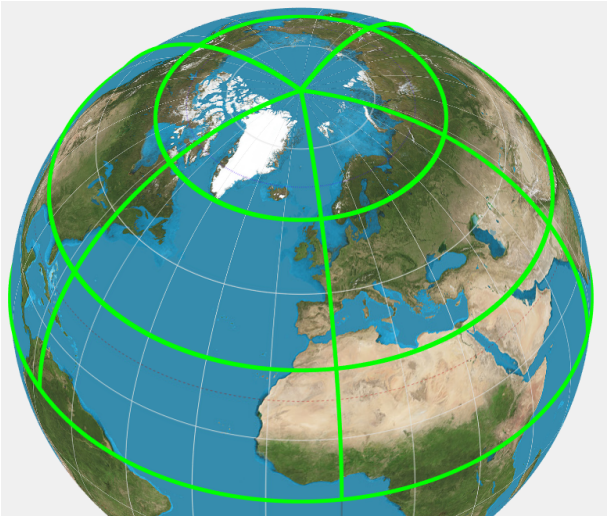

Figure 13.1: A way to divide the globe into 15 regions, such that 100 measurements are performed per region.

and we see that these regions are quite large. If smaller regions are desired, one needs to accept that less than 100 measurements are acquired per region, or perform several measurement campaigns. Designing a communication system based on less than 100 measurements is not recommended, as the system will then be fitted to the specific interference realizations measured rather than the average behavior. Performing two or three measurement campaigns is a better option, as this would also ensure that the spatial variation one has measured for the interference environment is actually spatial variation and not temporal variation. It should also be noted that the regions in figure 13.1 are chosen based on a grid of latitude lines and longitude lines. It would, however, be possible to determine the region boundaries based on the measured interference, such that areas where the interference environment is similar are grouped together, and the optimal parameter settings for a given outage probability are given for this interference environment.

13.3 Adaptive system

The communication system with the best performance would be one that continuously monitors the interference environment, and updates the communication strategy in real-time based on these measurements. Parameters such as the carrier frequency, transmit power, or level of coding could be updated to increase performance. When there is little interference, such a system would adjust accordingly in order to use less energy or get a higher bit rate, and vice versa when there is a lot of interference present.

A possible way to implement an adaptive system would be to measure the opportunity distribution on a regular basis and automatically update the communication system parameters using the procedure described in section 13.1. To simplify the system, one might limit the parameter choices to a predefined set of configurations and select the best suited amongst these. This would require a working two-way communication system, as the satellite needs to downlink the system configurations to the transmitting sensor node or ground station each pass. The advantage of measuring the opportunity distribution in the adaptive system instead of using another measurement technique is that the opportunity distribution is the measure that is used in the measurement campaign, and one does therefore have a good expectation of how the measurements look like when designing the communication system. A disadvantage with implementing an adaptive system this way, however, is that the opportunity distribution as a measurement technique is developed with the constraints of the LUME-1 mission in mind. As the adaptive communication system is not supposed to downlink its measurements, there will be different constraints for the measurement technique used to measure the communication channel, such that other measures than the opportunity distribution might be better suited for this application. If an adaptive system is implemented as described above, one should therefore ideally investigate other options as well, and at the very least, other parameters could be chosen when computing the opportunity distribution, for higher resolution.

14. *Discussion*

In chapter 8, we derived estimates for the distribution for how large portion of the payload data in a communication packet is outside the main opportunity window, $p_{outside}(\chi)$, and the distribution for how large portion of the payload data that will be lost due to interference, $p_{lost}(x)$. As for usefulness, $p_{lost}(x)$ is the most favorable because it directly tells us how much coding is needed in order to cope with the interference. However, in order to estimate this distribution, we need not only the opportunity distribution, but also the *pulse distribution*, which is the corresponding measure for how long time is spent in interference pulses with different durations. This would double the amount of data needed to downlink, and one would have to modify the developed software to be able to calculate the pulse distribution as well. Furthermore, it was found in chapter 9 that the estimate for $p_{lost}(x)$ deviates from the actual distribution. Due to this, we can conclude that no accurate estimate for $p_{lost}(x)$ is found, even if the pulse distribution is known.

In contrast to the estimate for $p_{lost}(x)$, it is shown in chapter 9 that the estimate for $p_{outside}(\chi)$ appears to be accurate. An opportunity distribution that is calculated with measurement parameters chosen to reduce its data size in bytes is sill able to give rise to an estimate of $p_{outside}(\chi)$ whose cumulative distribution deviates only 1.5 percentage points from the true distribution. This, as well as the fact that $p_{outside}(\chi)$ is estimated from only the opportunity distribution, are the advantages of this distribution compared to $p_{lost}(x)$. $p_{outside}(\chi)$ can also be used as a worst-case estimate for $p_{lost}(x)$, as the portion of the payload data that is lost to interference will always be less than or equal to the portion that falls outside the main opportunity window. Furthermore, we will always have that $p_{less\_than\_x\_outside}(0\%) = p_{less\_than\_x\_lost}(0\%)$ and $p_{less\_than\_x\_outside}(100\%) = p_{less\_than\_x\_lost}(100\%)$ for the cumulative distributions, which means that $p_{outside}(\chi)$ alone still provides reliable estimates for the probability that a communication packet has interference in the header, and the probability that there is any interference in the package at all.

When performing verification of the developed software in chapter 12, it should be noted that we set the parameter `num_orbits` to 3 and 64 for CONOPS option 1 and 4 respectively, even though we established in chapter 11 that the measurements performed on the LUME-1 satellite need to use `num_orbits` $= 1$, and that measuring for multiple orbits should rather be performed by scheduling multiple calls to the software. To prepare the software for measuring on the LUME-1 satellite, one would therefore only have to test with `num_orbits` $= 1$, as this is the parameter that will be used on the satellite. However, using `num_orbits` > 1 makes for a more comprehensive test that verifies even the parts of the software which will not be used on the LUME-1 satellite. The way the software is structured, any issues that might arise when using `num_orbits` $= 1$ would also be apparent when using `num_orbits` > 1. Thus, since the software works as expected with `num_orbits` > 1, no additional test is needed with `num_orbits` $= 1$.

In section 12.3, we discovered that two types of delays occur between measurements when measuring back-to-back. First, there is always a short, near-constant delay between consecutive measurements, which was measured to be ~270 ms when no signal processing was performed. This delay is likely to be longer when one computes the opportunity distribution of the measured signal, as the average time between two measurements was then measured to be 5.4 s. The disadvantage with these delays is that it increases the average time between two measurements, such that the measurements are spread out over a longer time interval than they otherwise would be. Secondly, there are delays that occur rarely but last longer than a second. These delays also increase the average time between two measurements, but more important is that they cause loss of even sampling times between consecutive measurements, which is disadvantageous when analyzing the data. This especially affects campaign option 1, where measurements

are performed back-to-back as fast as possible. Campaign option 4 is barely affected by these delays, since measurements in this option are spaced out 60 s apart anyways, and shifting a measurements with a delay on the timescale of 1 s has less of an influence. Also, due to the 55 s buffer between measurements, the delays will not accumulate over time. Therefore, although the presence of these delays is disadvantageous, it is not a critical issue. Because no way is found for getting rid of the delays, one must simply accept that they may also occur in the LUME-1 measurements.

15. *Future work*

This report investigates and prepares the opportunity distribution as a measurement technique for use on the LUME-1 satellite, but the measurements are not actually performed yet. This remains to be done as future work, by measuring on the satellite according to the campaign configurations decided in section 11.5. Included in this is that all the software that is to be uploaded to the LUME-1 satellite needs to be tested by the people who operate the satellite. This is due to their internal testing protocols, and it must be performed regardless of the testing described in chapter 12. For the actual measurement campaigns, the software must be scheduled with one software call per orbit, and then the resulting measurement files must be downlinked.

In addition to performing the measurement campaign, there are ways one can improve the measurement software that is yet to be done. One of these is to investigate the possibility of interfacing the Totem SDR using the libiio C drivers directly instead of the command-line utilities. An advantage of this is that the radio signal would not have to be written to an intermediate binary file, as fetching the radio signal could be performed in the same program as the signal processing. This would reduce both processing requirements and the amount of empty disk space required to perform a measurement. Furthermore, this might get rid of the rare, long delays described in it 12.3, as it was found that these delays may be caused by an error in the process of writing the fetched radio signal to the binary output file.

Another way of improving the measurement software is to modify the `analyze_signal` program such that it can calculate either the opportunity distribution or the pulse distribution. Having both these distributions could be advantageous to get a better understanding of the temporal statistics. The disadvantage with measuring both distributions is that it doubles the amount of data that needs to be downlinked by the satellite. Of the two distributions investigated in this report, only $p_{lost}(x)$ requires the pulse distribution in order to be estimated. It is shown in section 9 that the estimate for this distribution is both less accurate and a lot more time-consuming to compute than the estimate for $p_{outside}(x)$, which only requires the opportunity distribution. If one decides to compute the pulse distribution as well, one should therefore first investigate how it can be used to provide insight that is useful for designing a communication system.

Elaborating on the distribution $p_{lost}(x)$, one might be able to estimate it with a higher accuracy if it is estimated differently. As explained in section 8.5, an approximation had to be done in order to be able to use the estimate for $p_{outside}(x)$ in the estimation of $p_{lost}(x)$, as this simplified the derivation. If one can derive the estimate for $p_{lost}(x)$ without this assumption, this will produce a more accurate estimate. Whether the change in accuracy would be large or insignificant is not possible to say without actually deriving such an estimate and testing it. The imprecision of the estimate for $p_{lost}(x)$, however, might be caused by the assumption that the duration of an interference pulse is independent of the duration of the preceding opportunity windows and interference pulses. This assumption was made because the opportunity- and pulse distribution provide no information about the interdependency between the duration of the pulses and windows. To improve the precision of the estimate for $p_{lost}(x)$, one might have to find a way to measure this interdependency and downlink it as an additional measure, besides the opportunity distribution and pulse distribution. This would, however, require that one finds a way to measure the interdependency on a data format that does not occupy too much disk space and that one can find a better way to estimate $p_{lost}(x)$ using this measure.

16. *Conclusion*

In this report, a method for measuring temporal properties of uplink interference in satellite communication has been presented, and it is shown that the output produced by the measurement method can be used to estimate other probability distributions. The accuracy of these estimates is tested using data from the Norsat-2 satellite for verification, and it is found that the distribution $p_{outside}(x)$ can be accurately estimated. The average deviation between the estimate and the true distribution is found to be 1.5 percentage points with the chosen measurement parameters, which is deemed an acceptable deviation. Furthermore, it is shown that one can use the estimated distribution to make informed design choices when designing a communication system, although some parameters like the carrier frequency should be chosen based on other measurements than the opportunity distribution alone. A measurement campaign configuration is decided with the constraints of the LUME-1 satellite in mind, which achieves an acceptable trade-off between time to downlink measurements and geographic coverage. A software implementation of the chosen measurement method is developed for the LUME-1 satellite, and it is tested with the chosen measurement configuration in a lab to make sure that it is mission-ready. The tests reveal that the software operates as expected, except for some delays between consecutive measurements. It is concluded that these delays are not critical because they only have a small effect on the measurement campaign that covers the entire globe, and a moderate effect on the short pre-campaign.

Bibliography

[1] European Commission, HIGH REPRESENTATIVE OF THE EUROPEAN UNION FOR FOREIGN AFFAIRS AND SECURITY POLICY, *Developing a european union policy towards the arctic region: progress since 2008 and next steps*, June 2013. [Online]. Avaliable at: http://eeas.europa.eu/arctic_region/docs/swd_2012_183.pdf

[2] Joseph C. Casas, *Arctic region communications small satellites (arc-sat)*, Sep. 2012, accessed November 2013. [Online]. Avaliable: http://nix.nasa.gov/search.jsp?R=20120016896&qs=N%3D4294163423%2B4294957324%2B4294587540

[3] Arctic Council, *TELECOMMUNICATIONS INFRASTRUCTURE IN THE ARCTIC*, 2017. [Online]. Avaliable: https://oaarchive.arctic-council.org/bitstream/handle/11374/1924/2017-04-28-ACS_Telecoms_REPORT_WEB-2.pdf?sequence=1

[4] R. Birkeland, *Proposal for a satellite communication payload for sensor networks in the arctic area*, in *Proceedings of the 65th International Astronautical Congress*, 2014

[5] R. Birkeland, *On the Use of Micro Satellites as Communication Nodes*, Trondheim, 2019

[6] Zarikoff, Brad W. *Measuring Pulsed Interference in 802.11 Links*, 2013, [Online]. Available at: https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6221970, Downloaded 20.11.2020

[7] Paulo Marques, *A Hardware Demonstrator of a Cognitive Radio System Using Temporal Opportunities*, 2009, [Online]. Available at: https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5189021, Downloaded 20.11.2020

[8] Tang, P. T., *On the Distribution of Opportunity Time for the Secondary Usage of Spectrum*, 2009, [Online]. Available at: https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=4566074, Downloaded 20.11.2020

[9] G. M. Alberto, *Preliminary noise measurements campaign carried out by HUMSAT-D during 2014*, 2015, [Online]. Available at: https://www.itu.int/en/ITU-R/space/workshops/2015-prague-small-sat/Presentations/Humsat-Praga.pdf, Downloaded: 18.12.2020

[10] M. Buscher, *Investigations on the current and future use of radio frequency allocations for small satellite operations*, 2019, [Online]. Available at: https://depositonce.tu-berlin.de/handle/11303/9161 Downloaded: 18.12.2020

[11] S. Busch, *UWE-3, in-orbit performance and lessons learned of a modular and flexible satellite bus for future pico-satellite formations*, 2015, [Online]. Available at: https://www.researchgate.net/publication/282524791_UWE-3_In-Orbit_Performance_and_Lessons_Learned_of_a_Modular_and_Flexible_Satellite_Bus_for_Future_Pico-Satellite_Formations, Downloaded: 18.12.2020

[12] N2YO, *UWE-3*, 2013, [Online]. Available at: https://www.n2yo.com/satellite/?s=39446. Downloaded: 05.05.2021

[13] G. Quintana-Díaz, T. Ekman and F. Aguado Agelet., *Measurements of Radio Interference in the UHF Amateur Radio Band from a Small Satellite*, Remote Sensing (to be submitted), 2021.

[14] G. Quintana-Díaz., *LUME measurements description*, Internal Document, 1 (Non-published): 1-23, 2021.

[15] Endresen, S. K., *Measurement program for characterizing the interference environment in UHF-band satellite-based communication*, Trondheim, 2020.

[16] Bjørn Myrvold, personal communication, 17.11.2020

[17] IARU, *Amateur satellites*, 2017. [Online]. Available at: https://www.iaru-r1.org/wiki/Amateur_satellites, Downloaded 17.11.2020

[18] IARU, *IARU Region 1 UHF band plan*, 2017. [Online]. Available at: https://www.iaru-r1.org/wp-content/uploads/2020/03/UHF-Bandplan.pdf, Downloaded: 17.11.2020

[19] NRRL, *Norsk båndplan 70cm (435-438 MHz)*, 2016. [Online]. Available at: https://www.nrrl.no/images/bandplaner/BP2016-70cm.pdf, Downloaded: 17.11.2020

[20] ITU, *Characteristics of and protection criteria for radars operating in the radiolocation service in the frequency range 420-450 MHz*, 2019. [Online]. Available at https://www.itu.int/dms_pubrec/itu-r/rec/m/R-REC-M.1462-1-201901-I!!PDF-E.pdf, Downloaded: 22.10.2020

[21] B. Maryland, *Radar Navigation and Maneuvering Board Manual*, 2001. [Online]. Available at: https://books.google.no/books?id=FVzSrhl4MUoC&pg=PA1&lpg=PA1&dq=%22The+word+radar+is+an+acronym+derived+from+the+phrase+RAdio+Detection+And+Ranging+and+applies%22&source=bl&ots=SemWvADABz&sig=ACfU3U1Lw5VCyoSmUkK5HcUHaztbLFJwdg&hl=no&sa=X&ved=2ahUKEwjonILEy_rsAhWJtYsKHY6iCcsQ6AEwBHoECAgQAg#v=onepage&q&f=false, Downloaded 11.11.2020

[22] Gomspace, *NanoCom ANT430 Datasheet / 70 cm band Omnidirectional UHF CubeSat antenna*, 2020. [Online]. Avaliable at: https://gomspace.com/UserFiles/Subsystems/datasheet/gs-ds-nanocom-ant430-41.pdf, Downloaded: 06.10.2020

[23] N2YO, *LUME-1*, 2018, [Online]. Available at: https://www.n2yo.com/satellite/?s=43908. Downloaded: 05.05.2021

[24] Alén Space, *TOTEM-Motherboard datasheet*, 2018. [Online]. Available at: https://public.xeria.es/mindtech2019/es/Company/GetDocFile/2278?idAccount=1675267 Downloaded: 17.12.2020

[25] Bradbury, L. M., *NorSat-2: Enabling advanced maritime communication with VDES*, [Online]. Available at: https://www.sciencedirect.com/science/article/abs/pii/S0094576518303849, Downloaded: 09.04.2020

[26] Analog Devices, *What is libiio?*, 2020. [Online]. Available at: https://wiki.analog.com/resources/tools-software/linux-software/libiio, Downloaded: 10.12.2020

[27] Alén Space, *Software Defined Radio / Totem*, [Online]. Available at: https://info.alen.space/download-totem-sdr-platform-data-sheet Downloaded: 28.04.2021