

# **Thermal sources and timing precision of IceCube's Surface Enhancement DAQ**

**Wärmequellen und Timing-Präzision im Datenerfassungssystem  
des IceCube Surface Enhancement**

Bachelor's Thesis of

Philipp Dittrich

at the KIT Faculty of Physics  
ETP – Institute for Experimental Particle Physics  
IAP – Institute for Astroparticle Physics

Reviewer: Prof. Dr. Frank G. Schröder  
(Associate Professor,  
University of Delaware, USA)

Second reviewer: Dr. Andreas Haungs

Advisor: Dr. Thomas Huber

Second advisor: M.Sc. Megha Venugopal

01. Mai 2023 – 30. September 2023



---

I declare that I have developed and written the enclosed thesis completely by myself. I have submitted neither parts of nor the complete thesis as an examination elsewhere. I have not used any other than the aids that I have mentioned. I have marked all parts of the thesis that I have included from referenced literature, either in their original wording or paraphrasing their contents. This also applies to figures, sketches, images and similar depictions, as well as sources from the internet.

.....  
Philipp Dittrich  
Karlsruhe, September 2023



# Abstract

## Thermal sources and timing precision of IceCube's Surface Enhancement DAQ

The IceCube experiment at the South Pole is about to be equipped with a large set of scintillator panels and radio antennas (Surface Array Enhancement), which are already available and finalised. The readout electronics named TAXI though still show errors, mainly related to their onboard temperature. Especially the antennas' readout data is still highly corrupted. As knowledge about the exact functionality of the electronics has drifted apart over the years, this work is supposed to bring some of that back into one place, giving a detailed overview of clock routings, tick rates and radio data flow within the TAXI board and highlighting stress points where the currently observed errors might arise.

Mainly, the issues can be traced back to temperature related timing problems on the TAXI mainboard, specifically to the communication between the FPGA and the ADC units. Reading out the eight DRS4 sampling channels per antenna requires serialisation of eight 14-bit data streams onto one high-speed FPGA port. Because of the high-frequency nature of this process, inconvenient FPGA coding will quickly lead to errors, when the internal FPGA clock overtakes the processing capabilities, which happens more regularly at higher temperatures.

A closer look has been taken at the temperature problem itself by identifying sources of heat inside the TAXI's metal housing and investigating how the situation could be improved. This partly focuses on the placement of the currently onboard White Rabbit client, the device for inter-TAXI synchronisation, timing and communication, which is argued to be moved out of the TAXI box and into its own housing. The results of the conducted test runs show a 2-3 °C onboard temperature decrease, when using a relocated, standalone White Rabbit client.

Research has also been done regarding the TAXI's power consumption, which varies between 25 and 30 W, depending on the initialisation state of the scintillators. However, the impact of active data acquisition on the onboard temperatures is marginal. The same can be said about the effect of the DRS4 chips' data sampling rate on the power consumption, which has been evaluated with DRS4 evaluation boards. The sampling rate only minimally changes the power consumption, which implies a negligible effect of the DRS4 sampling rate on TAXI temperatures.

Additionally, this thesis contains comparisons between identically set up TAXIs at room temperature, showing onboard temperature differences of about 2 °C and a big difference in the radio sampling quality, with one TAXI showing excessively more corruption than the other.

---

Of the DRS4 evaluation board, the White Rabbit Client and the current TAXI board, IR pictures were taken that are also included in this thesis, alongside results from a TAXI temperature chamber setup evaluating the TAXI's frequency response at different environmental temperatures, where the first artifacts were observed at environmental temperatures of 10 °C.

# Zusammenfassung

## Wärmequellen und Timing-Präzision im Datenerfassungssystem des IceCube Surface Enhancement

Das IceCube Experiment am Südpol soll demnächst mit einer großen Anzahl an Szintillator-Panels und Radioantennen ausgestattet werden (Surface Array Enhancement). Diese sind bereits verfügbar und bereit zum Einsatz. Die Auslese-Elektronik mit dem Namen TAXI hingegen erzeugt immer noch Fehler, die hauptsächlich mit den hohen Temperaturen zu tun haben, die auf dem Mainboard herrschen können. Speziell die Auslesedaten der Radioantennen sind dann stark beschädigt. Da das Wissen über die genaue Funktionsweise der Elektronik über die Jahre verloren ging, soll diese Arbeit nun einiges davon wieder zurück an einen Ort bringen. Dazu gehört eine detaillierte Übersicht der Clocks, die auf dem TAXI Mainboard laufen, und wie sie verteilt werden. Beeinflusst wird davon unter anderem die Tick-Rate der Onboard-Steuereinheit FPGA und der Auslesemechanismus der Radio-Signale, deren Signalweg dargelegt wird. Dabei werden kritische Punkte, an denen Daten beschädigt werden können, aufgezeigt.

Hauptsächlich können die Probleme zurückgeführt werden zu Timingproblemen auf dem TAXI Mainboard, genauer zu der Kommunikation zwischen der FPGA und den ADCs, den Analog-Digital-Konvertern in der Radio-Signalauswertung. Das Auslesen der acht DRS4 Sampling-Kanäle pro Antenne setzt die Serialisierung dieser acht 14-bit Datenflüsse voraus, die dann alle über einen einzigen FPGA-Port eingelesen werden. Wegen der hohen Frequenzen, die bei diesem Prozess auftreten, kann eine ungünstige Programmierung der FPGA schnell zu Fehlern führen. Dies hängt mit der internen FPGA-Tickrate zusammen, die in solchen Fällen zu schnell läuft, um rechtzeitig alle Datenstreams zu verarbeiten. Dies passiert mit erhöhter Wahrscheinlichkeit bei Umgebungstemperaturen über der Raumtemperatur.

Daher wurde ein genauerer Blick auf die Auslöser für die hohen Onboard-Temperaturen geworfen, die derzeit innerhalb des metallenen Gehäuses des TAXI entstehen. Dies bezieht sich teilweise auf die Position des White Rabbit Clients, der sich in der aktuellen Version innerhalb des TAXI-Gehäuses befindet und für Synchronisation, Timing und Kommunikation zwischen einzelnen TAXI-Einheiten zuständig ist. Es wurde diskutiert, diesen außerhalb des TAXI-Gehäuses als eigenständiges Gerät aufzustellen. Die Auswirkungen wurden im Rahmen dieser Arbeit untersucht. Als Resultat zeigt sich eine absolute Temperaturabsenkung um 2-3 °C auf dem Mainboard im Falle eines externen White Rabbit Clients.

Gleichermaßen wurden Untersuchungen durchgeführt, bei denen der Stromverbrauch des TAXI gemessen wurde. Dieser liegt zwischen 25 und 30 W, je nach Initialisierungsstatus der Szintillatoren. Die Datenaufnahme selbst hat einen vernachlässigbaren Einfluss auf die Temperaturen. Gleiches gilt für die Sampling Rate der DRS4 Sampling Chips, die das

---

Radio-Signal aufnehmen. Dies ist das Resultat einer Messung mit DRS4 Evaluation Boards. Die Samplingrate bewirkt nur eine minimale Änderung des Stromverbrauchs, weshalb auf einen vernachlässigbaren Effekt auf die TAXI-Temperaturen geschlossen werden kann.

Zusätzlich dazu enthält diese Arbeit Vergleiche zwischen identisch präparierten TAXIs bei Raumtemperatur. Die beiden Geräte zeigten Onboard-Temperaturdifferenzen von ungefähr 2 °C und große Unterschiede in der Qualität der aufgenommenen Radio-Daten.

Um Hitzequellen sichtbar zu machen, wurden im Rahmen dieser Arbeit ebenfalls Infrarotbilder von einem DRS4 Evaluation Board, dem White Rabbit Client und dem aktuellen TAXI Board aufgenommen, die neben Ergebnissen eines TAXI-Versuchs in einer Kältekammer ebenfalls in diesem Werk enthalten sind. Erste Artefakte waren bereits bei einer Umgebungstemperatur von 10 °C zu sehen.



# Contents

<b>Abstract</b>	<b>i</b>
<b>Zusammenfassung</b>	<b>iii</b>
<b>1. Introduction</b>	<b>1</b>
<b>2. Cosmic Ray Physics</b>	<b>3</b>
2.1. Cosmic Rays . . . . .	3
2.2. Air Showers . . . . .	4
2.3. Detector Types . . . . .	6
2.4. Shower Reconstruction . . . . .	7
<b>3. The IceCube Neutrino Observatory</b>	<b>9</b>
3.1. IceCube In-Ice Detectors - the neutrino observatory . . . . .	9
3.2. IceTop - water Cherenkov tanks . . . . .	10
3.3. The IceCube Surface Array Enhancement . . . . .	10
3.4. IceCube-Gen2 . . . . .	14
<b>4. Studies on TAXI electronics</b>	<b>15</b>
4.1. Radio signal detection and pre-amplification . . . . .	15
4.2. Processing the radio signal - the radioTads . . . . .	15
4.3. Controlling the readout - the FPGA . . . . .	17
4.4. Sampling the radio traces - the DRS4 chips . . . . .	18
4.5. Triggering Mechanisms . . . . .	21
4.6. Readout of the DRS4 chip and analog-digital conversion . . . . .	22
4.7. Bin spikes and timing issues . . . . .	23
4.8. Global synchronisation, drift and data tagging . . . . .	23
<b>5. Studies on temperature dependences of the TAXI</b>	<b>27</b>
5.1. Impact of TAXI data acquisition on the temperatures . . . . .	27
5.2. Different placements of the White Rabbit Client . . . . .	32
5.3. Comparison of two similar TAXI setups . . . . .	36
5.4. Additional heat sources . . . . .	37
<b>6. Conclusion and Outlook</b>	<b>41</b>
<b>A. Appendix</b>	<b>43</b>
A.1. Algorithms used for uncertainty calculation . . . . .	43
A.2. Graphic overview over the TAXI components . . . . .	43

*Contents*

---

A.3. TAXI in the IAP temperature chamber . . . . .	45
A.4. Additional plots: WR-LEN placement . . . . .	46
A.5. IR images . . . . .	49
<b>Bibliography</b>	<b>51</b>

# 1. Introduction

As a strange red glow it was described, when over 500 years B.C., Babylonians first recorded the sighting of an aurora. Since then, much time has passed, and what was inexplicable at that time has now been studied extensively to be induced by charged particles emitted by the Sun. Research showed that its rays not only affect us visibly by spending us light, but also invisibly by darkening our skin and driving photosynthesis in leaves. These findings have augmented our understanding of skin cancer and improved agricultural practices. Amazed by the invisible, scientists went even further and started to look for other sources of energy in the universe that could also affect us on Earth, without us even noticing. This is how Victor Hess first discovered Cosmic Rays in the early 20th century.

From interactions with the atmosphere over traces in water detectors, high-energy particles coming from the universe, called cosmic rays, have since been of big interest for modern astroparticle physics, as they give us detailed information about other celestial bodies that are unobservable with the naked eye, closer elaborated on in chapter 2. For detecting such radiation, usually, large-scale experiments are in use. The reason for this is quite simple - the more area there is, the higher the chance something is detected and the better the results. One of these experiments is called IceCube and is situated at the Earth's South Pole.

With its cubic kilometre in-ice detector array, IceCube has been running and detecting high-energy particles from space for several years already. Now, there are plans for an upgrade, called the IceCube Surface Array Enhancement. It consists of a large amount of detector panels and antennas, spread around the surface of the experiment, with the goal of gaining more detailed information over single events from such high-energy particle interactions, as explained in more detail in chapter 3.

However, merely installing hundreds of detector panels and nearly a hundred antennas does not meet the high requirements the collaboration has for a surface enhancement, A highly robust readout algorithm for these devices is necessary for gaining all the necessary information, so that in the aftermath, scientists can reconstruct the event and understand what has happened. This is why the TAXI was developed, an electronic readout board that is placed at stations distributed all over the experiment's surface. Each TAXI connects to eight panels and three antennas, it communicates with other TAXIs and it sends all its data to the central lab for evaluation.

The system of TAXIs can thus be seen as the brain of the experiment, inter-connected and working as on reflexes. Developing such a reliable readout system is a big challenge though and has been bothering the responsible teams for years already. Currently, the system is closer to finalisation than ever before, however, there are still challenges to be overcome in the data acquisition modes' behaviour, like faulty return values, erroneous readings or occasionally corrupted event data packages. These issues have mostly been narrowed down to problems related to the onboard temperature, in other words, the

board currently runs too hot, causing many of these errors. One goal of this thesis was to investigate possible approaches how to handle the temperature effects on the TAXI DAQ. This includes the analysis of single modules and elements of the TAXI system and how they could be treated in order to improve reliability. These investigations can be found in chapter 5.

With some original developers having transitioned to other projects, there is a noticeable gap as the knowledge they possessed is now dispersed among various individuals. Another goal of this thesis was thus to gather some of this knowledge into one spot again, giving a good overview of the functionality of the antenna readout system, so that the following research can be sped up by new team members having the most necessary information already all in one place. The electronics of the TAXI board are explained in detail in chapter 4.

Additionally, smaller projects are included, partially about the TAXI's power consumption, the sampling rates of specific components, and the data's consistency over different TAXI devices. Overall, this thesis serves as a comprehensive compilation of the current challenges faced in TAXI data acquisition, its data pipeline stability and its clock tree mechanisms, supplemented with potential solutions for the further development of TAXI and future DAQs for large area instrumentations in harsh climatic environments.

## 2. Cosmic Ray Physics

Modern Astroparticle Physics is the research field that includes experiments like the IceCube Neutrino Observatory, which focus strongly on cosmic rays and their sources. Cosmic rays originate from galactic (within our galaxy, the Milky Way) or extragalactic (outside our galaxy) sources and can be categorised into different energy ranges. Generally speaking, cosmic rays are charged particles or nuclei travelling through space at relativistic speeds, as formulated in [1]. That means, they can have all kinds of constitutions and origins. In this chapter, the physics behind those cosmic rays will be briefly summarised, alongside with the rays' interactions with the Earth's atmosphere.

### 2.1. Cosmic Rays

As mentioned before, cosmic rays are less rays of light than particles flying through space. As of today, their origins are not known too well, so measuring and analysing cosmic rays on the Earth still holds big potential for new discoveries. Within our galaxy, the Milky Way, supernova explosions are the primary source of cosmic radiation [2], sending shock waves through the universe and thus accelerating matter in space. Those can explain energies of cosmic radiation particles up to  $3 \times 10^{15}$  eV. Going even higher, there are particles with energies above  $10^{19}$  eV, which are generally considered being from *extra-galactic* sources [3]. Those *Ultra High Energy Cosmic Rays* (UHECR) arrive at our Earth quite rarely and, only get detected once every few weeks. Their sources are expected to be of more impactful origin than supernovae, namely pulsars (systems of multiple neutron stars spinning around their center of mass with a frequency of multiple MHz), jets from supermassive black holes, or gamma ray bursts. Some particles are accelerated so fast that their magnetic Larmor radius surpasses the diameter of the Milky Way, which initially sparked the idea of their extragalactic origin.

Quantitatively, it is quite common to reference the diagram seen in Fig. 2.1. It combines the results of many foregoing experiments into one plot, making three distinct features quite obvious - two so-called "knees" and an "ankle". In other words, there seem to be multiple points in the energy spectrum at which the amount of incoming particles suddenly falls off. So far, solid explanations have not been found. Though there are theories speculating about the reason [4]:

The **first knee** is likely the upper edge of the primary source of galactic radiation, supernovae, marking the beginning of the interval between  $10^{15}$  and  $10^{19}$  eV that connects galactic and extra-galactic radiation. With those primary accelerators falling away, it is to be expected that the flux decreases, with few other galactic sources (yet unknown) still able to give particles enough momentum.




Figure 2.1.: **The flux of high energy particles arriving at the earth as Cosmic Rays as a function of their energy.** For better readability, the y axis is multiplied by  $E^{2.6}$  [4].

The **second knee** is believed to be of the same reason as the first knee, only for heavier nuclei. Those carry higher charges and thus can be accelerated more by the same sources, surpassing the  $10^{15}$  eV barrier.

The **ankle** finally marks the upper limit of detected galactic rays with everything remaining needing to stem from extra-galactic sources. There are some theories surrounding interactions with the cosmic microwave background [5] and other effects of extra-galactic radiation dominating this part of the spectrum [6], but details are yet unknown.

## 2.2. Air Showers

No matter if it is gamma radiation, electrons, hadrons or heavier nuclei, they all interact somehow with our atmosphere, should they enter it. As a result, Cosmic Rays produce so-called air showers within our atmosphere. These usually develop in the stratosphere, at typical heights of 15 to 35 km above sea level and take place over a vertical distance of several kilometres [7, 8]. They are called showers because of their characteristic shape. There usually is a specific primary particle with certain properties (mass, energy, charge), which interacts with air molecules. Because of its high energy, usually, it gets destroyed and fanned out into a range of secondary particles, often lighter mesons and leptons, which interact again themselves, decaying further (see Fig. 2.2). This happens over quite some time, with the amount of particles steadily increasing. Once the primary particle's energy has spread far enough, the secondary particles' energies will be too low to create

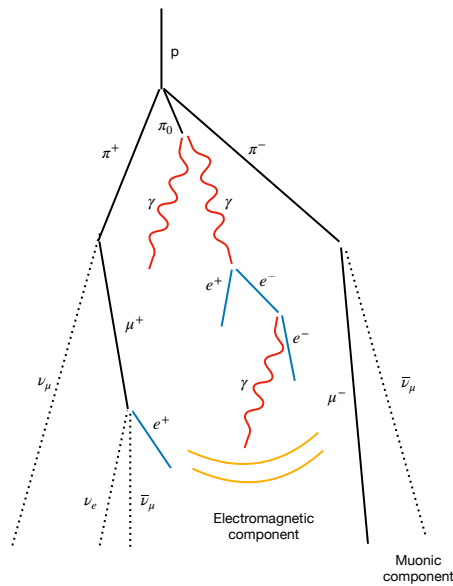


Figure 2.2.: **Schematic of an air shower produced in the Earth's atmosphere.** [1]

new particles instead, so now, some of the positrons and electrons will recombine, and the amount of particles in the shower decreases again. The distance from the primary particle's initial interaction to the point at which the shower has reached its maximum particle count is called the  $X_{\max}$  distance (compare the shower profile in Fig. 2.3). Note this quantity, as it is one of the most used values in shower core reconstruction: It is correlated with the primary particle's mass. The particle's energy, so mass and momentum, influences how the shower will look like once it has reached the Earth's surface - and turned into measurable data for researchers.

### 2.2.1. Radio emission

Because of all the different interactions inside a shower, the resulting particles vary quite widely. Generally, the shower can be split into an electromagnetic, a hadronic and a muonic component. From those three, most have decayed once they have reached the Earth's surface, except the muons, which can still be registered because of their relativistic speeds and the resulting time dilation. The same can be said about any produced neutrinos, which do not decay. Neutrinos and muons can be detected using the IceCube InIce detector array (see chapter 3.1) and scintillation detectors (chapter 3.3.2) respectively. Apart from that, the air shower's electromagnetic component also plays an important role in measurement. It is made up of electrons and positrons, usually resulting from the

$$\pi_0 \rightarrow \gamma\gamma \rightarrow e^+e^+ e^-e^-$$

decay close to the shower's initial interaction. Because of bremsstrahlung and Compton scattering, more high-energy photons are produced which themselves create electron-positron pairs. This chain goes on until the photon's energy goes below the pair production

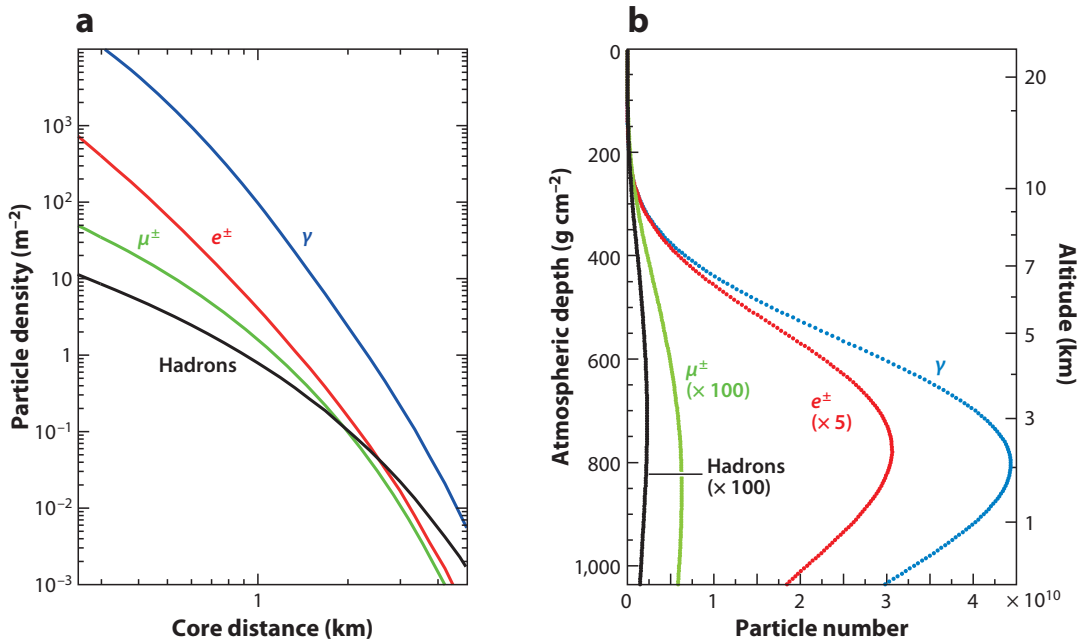


Figure 2.3.: Average (a) lateral and (b) longitudinal shower profiles for vertical, proton-induced showers at  $10^{19}$  eV. The  $X_{\text{max}}$  distance is measured from the top of (b) to the peak of a graph. [7]

threshold - the  $X_{\text{max}}$  distance is reached. Now, a vast amount of electrons and positrons is travelling through the atmosphere - and, based on different effects that are further explained in [9] and [10], they separate spatially and thus create a dipole. Based on how "thick" the shower is at the  $X_{\text{max}}$  distance, different frequencies of radio signals get emitted by this dipole. This emission can range from a few MHz to a few GHz [1, 11].

### 2.3. Detector Types

From the entire shower, what finally arrives on the surface are muons, radio waves and fluorescence light. The latter comes from excitation of air molecules by incoming cosmic rays. By spontaneous relaxation, those molecules emit photons which can be detected by optical fluorescence telescopes like the ones used in the Pierre Auger Observatory [12]. Because of the optical nature of these detectors, this also means that the quality of their data depends strongly on environmental factors like moon phases and clouds.

**Radio waves** are usually measured with antennas. There are certain types that are optimised for specific frequency regions, but they all have in common that they are relatively cheap and undisturbed by environmental factors. However, calibrating antennas can be quite laborious, as they usually cover a wide angle that needs to be accounted for. This can be done with a drone carrying a signal source [1].

**Muons** that arrive on earth usually have energies in the range of a Minimally Interacting Particle [13]. For detecting muons, electromagnetic scintillation detectors are most commonly used. In astroparticle physics experiments, those scintillators usually cover



an area of multiple  $m^2$  in order to increase the event rate<sup>1</sup>. A passing muon will excite the plastic material within the scintillator, which on relaxation will emit photons that get detected. Such a scintillation detector is only used for detecting the number of muons, not their energies.

**Neutrinos** are difficult to observe and are usually reconstructed out of secondary particles they produce. In general, the higher the density of a material, the higher the interaction chance. Covering large volumes, water tanks are well suited for detecting neutrino interactions. Rarely, a neutrino may interact with the water and produce a high energy electron, which then sends out Cherenkov radiation. Because this radiation is emitted in a cone shape, having multiple photodetectors in an array helps in reconstructing the direction of origin and the energy of the particle. Water (or ice) is predestined for this, as it is sufficiently transparent and has a refractive index high enough for electrons to produce Cherenkov light.

## 2.4. Shower Reconstruction

With multiple of the just described methods combined, it has become possible to reconstruct air showers, including information about the primary particle's identity and composition, its energy and, very importantly, the direction where it came from. In experiments like IceCube, which will be described in more detail in the following chapter, an angular resolution of  $1^\circ$  or less is aimed at. For achieving this, large detector arrays with an area of a square kilometre or greater are needed. With such dimensions, even showers from very high energy particles can be fully observed when the shower axis falls into the center of the array. The reconstruction works with the idea that every decaying particle produces new particles that are emitted omnidirectionally. As particles lose energy, the spherically emitted decay products overlap along the shower axis, creating higher energy densities that allow identifying the shower front. With this information travelling near light speed, in order to accurately read out and interpret it, timing resolutions of 1 ns or less are needed in the detector array. The IceCube experiment has been designed to do just that, and will be described now.

---

<sup>1</sup>Actually, the sweetspot must be found between large detector sizes and high resolution of the shower. As scintillator detectors work in a "yes or no" logic, no information is read out about where the muon passed through it - so having large scintillators reduces the accuracy about where the particle hit the surface exactly.



## 3. The IceCube Neutrino Observatory

The IceCube Neutrino Observatory (often called IceCube Experiment or simply IceCube) is a large-scale experiment located at the Earth's South Pole. It consists of several components, all of which are optimised for measuring different kinds of secondary particles coming from air showers induced by cosmic rays. With all its recent and future additions, the IceCube experiment has evolved beyond just being an experiment for detecting neutrinos, it has become a high resolution instrument for understanding and mapping our universe. In the following chapter, the different components of the IceCube experiment are briefly introduced and explained.

### 3.1. IceCube In-Ice Detectors - the neutrino observatory

What gave IceCube its name is a large detector array consisting of a total of 5160 photomultiplier tubes called DOMs (Digital Optical Modules) located on strings, deep in the ice. The detectors are spaced mostly evenly distributed and cover a volume of about  $1 \text{ km}^3$ , with a more dense array right in the middle of the arrangement. As explained before, dense and transparent materials like water are the medium of choice for detecting neutrino interactions, and as this was once the main goal of the IceCube experiment, the South Pole was chosen as the location - because it features a kilometre thick layer of solid water (ice). It has been found that starting from roughly 1 km beneath the surface, the ice becomes sufficiently transparent for visible light. So, 86 holes, 2.5 km deep, were dug in the ice using a hot water drill. Then, 86 strings of approximately the same length were lowered down into those holes - with the afore mentioned DOMs connected to those wires with distances of 17 metres to each other, so that to one string, in total, 60 DOMs are attached [14]. These photomultiplier detectors are sensitive to Cherenkov light emitted by fast-moving electrons in the ice. Incoming neutrinos interact with the ice similarly as with water and thanks to the clarity of the material, the cone shapes of the Cherenkov radiation can clearly be mapped, the particle's path tracked and its energy reconstructed.

Because the DOMs are located so deep in the ice, maintenance is impossible. Shortly after deployment, the ice freezes shut again and seals the DOMs for the rest of the time. Because of glacier movements, the ice on the south pole drifts roughly 10 m per year, causing strain on the strings and cables in the ice [15]. It is unavoidable that at some point, the strings will tear and the array thus become unusable. This is why the IceCube experiment has a limited lifetime.

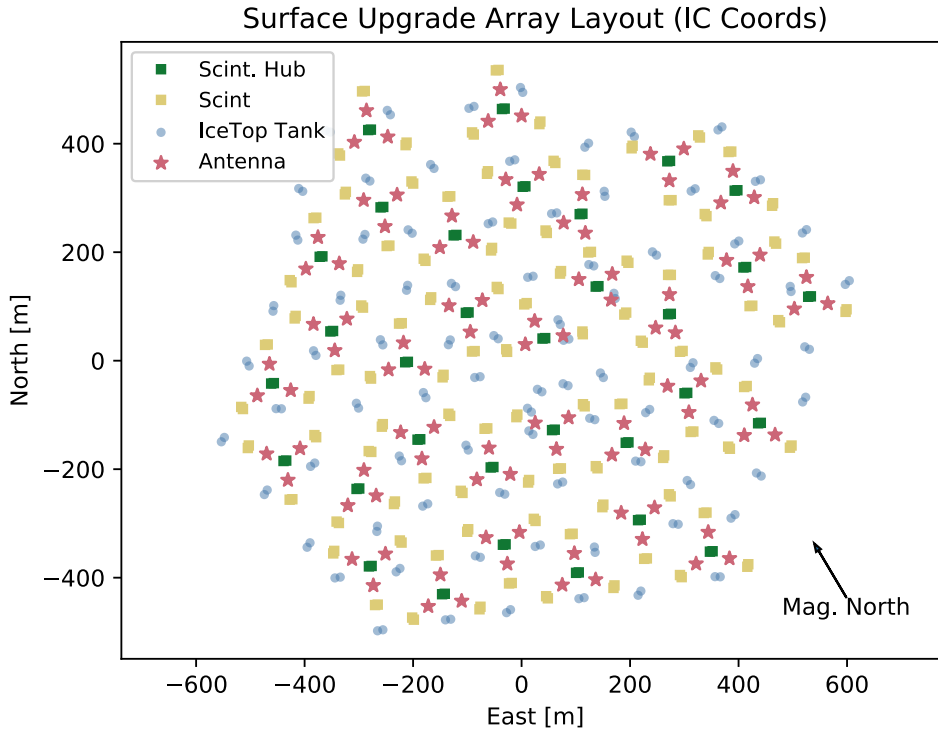


Figure 3.1.: **Planned layout of the IceCube SAE.** The array is organised in 32 groups monitored by one TAXI each (inside the Hub), consisting of eight scintillators each and three antennas in between. [1]

### 3.2. IceTop - water Cherenkov tanks

With just the DOMs in the ice, it is hard to distinguish muons from neutrinos of astrophysical origin. That is why, on the surface, above nearly every string, a pair of water Cherenkov tanks is located with two photomultiplier tubes inside each of them. The main goal of those detectors (often called IceTop, as the tanks were originally positioned on top of the ice<sup>2</sup>) is to get directional information about incoming particles from showers. If a particle is detected first in a water Cherenkov tank on the ice and then beneath the surface by the DOMs, it is highly likely that it initially came from a cosmic ray. The combination of IceTop and the in-ice detectors measuring the high-energy muons in the shower core makes the IceCube experiment a powerful cosmic ray detector, as stated in [14].

### 3.3. The IceCube Surface Array Enhancement

Over the years, studies on radio emission by air showers proved antennas as a very promising tool for reconstructing cosmic rays. That's why the IceCube experiment is

<sup>2</sup>because of accumulating snow, this part of the experiment is likely to soon lose functionality, becoming "in-ice" itself. The substitute version will be positioned on metal poles and can be continuously raised in order to prevent the constant snow accumulation.

planned to have an upgrade in form of radio antennas and scintillators, called Surface Array Enhancement (SAE). Apart from the water Cherenkov tanks on the ice, the experiment is about to get 32 sets of scintillators or antennas, with every single one consisting of a so-called FieldHub, an array of eight scintillators arranged in a triangle shape with one pair in the middle of the triangle, and a total of three antennas located in between two pairs of scintillators (see Fig. 3.1). The scintillators are used for triggering the radio readout, which will be explained in greater detail in chapter 4. The SAE and the in-ice part of IceCube are at some point supposed to communicate, with the SAE being able to veto potential neutrino events, similar to the current IceTop. The communication network used in the IceCube SAE follows the open-source project *White Rabbit*, and plays a key part in synchronising the detectors, delivering a clock and sending data to the central IceCube lab. As the SAE's data acquisition marks the main topic of this thesis, some more details regarding the individual components are explained in the following.

### 3.3.1. The antennas

In the Surface Array Enhancement, for radio signal detection, logarithmic-periodic dipole antennas (LPDA) of the kind *SKALA-v2* are used. They cover a wide angle ( $\pm 45^\circ$  from zenith) and, because of their build form, feature a uniform, omnidirectional gain distribution. The antennas are able to pick up signals in the range from 50 to 650 MHz, but in the IceCube experiment, this interval is limited to the range of 70 to 350 MHz. Below 70 MHz, the galactic background noise increases exponentially, and from times where the antennas were not thought of yet, there exists an active communication channel from the United States Antarctica Program which operates at 360.2 MHz [16], which would severely increase the noise level of the measurements was it not filtered out beforehand, hence the upper limit. The antennas have two polarisation channels for the best omnidirectional coverage, see Fig. 3.2.

### 3.3.2. The scintillators

For detecting charged particles of the shower, scintillator panels with a sensitive area of  $1.5 \text{ m}^2$  are used. They consist of a plastic scintillation material that emits photons after being excited by a charged particle passing through. The plastic fibres inside lead these photons towards a Silicon Photomultiplier (SiPM) while also shifting their wavelength to one at which the SiPMs are the most sensitive. SiPMs feature some major advantages over traditional photomultipliers, as they cannot burn out under too much light - instead, they simply saturate. Each scintillator has electronics onboard that digitise the signal and then send it to the data acquisition electronics [17]. This is of advantage because it outsources some heat-generating elements, while simultaneously creating a digital communication path between the scintillators and the rest of the electronics - an analog signal would be more prone to errors. Because the scintillators in the SAE are only used for triggering antenna readouts if an event occurs above a set intensity threshold, on-device digitisation is a very convenient step for efficiently forwarding only the relevant information.




Figure 3.2.: **The SKALA-v2 antenna.** CAD drawing of one arm with its log-periodic elements [1], and real-world photo for scale (around 3.5 m tall).

#### 3.3.3. The FieldHub

At the core of each of the 32 upcoming Surface Array Enhancement stations sits a FieldHub, a metal container housing the data acquisition electronics, power delivery and a network patchpanel (Fig. 3.3). It is daisy chained to other FieldHubs and insulated against the freezing temperatures at the South Pole<sup>3</sup>. In order to mitigate snow accumulation, it is stood on stackable metal bars that, just like the antennas and scintillators, can be elevated further, should the snow accumulate too high.

#### 3.3.4. The data acquisition electronics: The TAXI

The South Pole is a rough environment for electronics, even more so if these electronics need to detect particles known to interact very weakly. It was thus necessary to develop a special array for reading out the scintillator and radio data and keeping it in sync over the entire area. This is known as TAXI, which stands for Transportable Array for eXtremely large Instrumentation studies. It houses a variety of electronics that process the incoming signal from 8 scintillators and 3 antennas (of which each one has 2 polarisations, so there are 6 ports) and then forward event data to the IceCube Lab. Currently, the electronics show some malfunctions, which is why in the frame of this thesis, further research has been done in order to find possible reasons, with the results discussed in chapters 4 and 5. It is also important to stress that up until now (Mid 2023), only one prototype station with its TAXI has been deployed at the South Pole.

---

<sup>3</sup>As of January 2023, the insulation has been reduced because the electronics inside were suffering from temperature related signal errors in the Southern summer.




Figure 3.3.: **The insides of the FieldHub deployed at the South Pole.** As of end of 2022 [18]. Apart from the TAXI, it houses an eBox for power delivery and a networking patchpanel.

### 3.3.5. Timing and synchronisation: The White Rabbit Network

With 32 active FieldHubs, synchronisation and communication plays a key role for accurate data acquisition. That's why a special, open source network solution is deployed in the IceCube experiment that serves both as a timing foundation as well as a communication channel. The White Rabbit Network's signal is sent over glass fibres, with a central switch *WRS 3/18* splitting the initial signal into up to 18 individual paths that then are daisy chained. Within the frame of the White Rabbit network, it is possible to transmit timing information over up to 80 km of glass fibre, which would allow the daisy chaining of all 32 stations without a loss of precision. Every FieldHub houses a separate White Rabbit client, a node in this network, called *WR-LEN*, a board designed by SEVEN Solutions (now Safran) [19].

By default, these devices are shipped in metal housings, with dedicated cooling fins. In the current iteration of the TAXI though, the internal board of these White Rabbit clients is removed and placed directly inside the TAXI box, on a separate metal layer above the other boards. As the *WR-LEN* heats up significantly and the TAXI's components are sensitive for temperature changes, another goal of this thesis were more detailed investigations as to where this board should best be located.

### 3.3.6. Current issues with the data acquisition

The current iteration of the data acquisition system is still in active development. This seems obvious when looking at the remaining issues that need to be solved before a full deployment:

- **Bin Spikes.** The motivation behind this thesis are frequently occurring spikes in the antenna data. The amount varies among different TAXIs. As this will be looked at in great detail in the following chapters, it is kept short here.

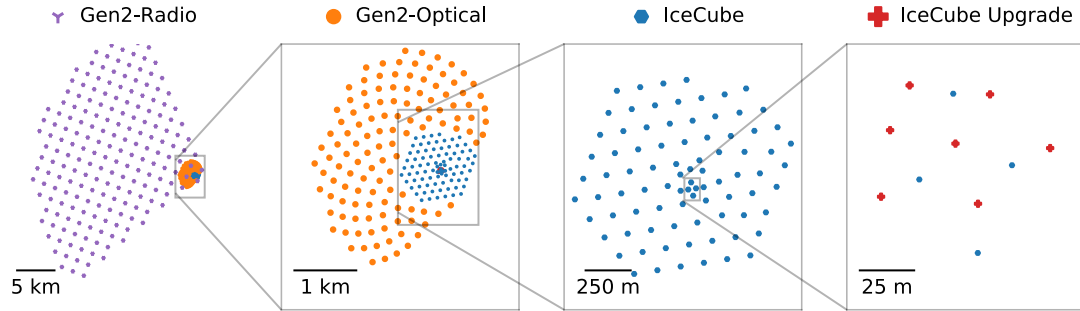


Figure 3.4.: **Schematic of the planned IceCube-Gen2, viewed from the top.** From right to left: A denser core with 7 added strings, the currently existing IceCube array, an upgrade with 120 new strings, and 200 radio stations, arranged in a fan-like shape. [22]

- **Heat sensitivity.** Several of the components of the TAXI board react very sensitively to temperature changes, especially to temperatures above 25 °C [20]. This affects the quality of the readout data (see above) and the communication between the TAXI and the readout electronics in the scintillators, which become unresponsive [21].
- **Deprecated tools for FPGA programming.** The TAXI’s *Spartan-6* FPGA can only be programmed with the proprietary software *Xilinx ISE*. This one is incompatible with all currently used operating systems (Linux, macOS and also Windows 10 and later) and needs a specifically set-up virtual machine to even communicate to a connected FPGA. This, together with mostly limited knowledge about FPGA coding from non-specialists, makes changing parameters in the current setup a challenge.

With all of these issues still to solve and the possibility of an IceCube-Gen2 upgrade in the farther future (see next section), it poses the question of how much the TAXI should be optimised until it needs to be replaced by something entirely new. As no tangible plans exist right now, the TAXI optimisation is still high on the priority list.

### 3.4. IceCube-Gen2

The IceCube experiment’s detector array is vast already, but it has limits. Particularly particles with energies up to EeV ( $10^{18}$  eV) are so rare that even the IceCube array does not catch enough events to get statistically helpful results. For the future, a larger-scale experiment is needed. This is why the idea of IceCube-Gen2 has been formulated (Fig. 3.4) - it would extend the instrumented array in the ice almost by factor 8, add 200 radio stations in a fan-like shape, and make the core array even denser by adding 7 more strings for better angular resolution. The IceCube Surface Array Enhancement is part of this larger project. IceCube-Gen2 would cost around \$350M and is still in an early state of development [22].



## 4. Studies on TAXI electronics

With the TAXI being the central part of the data acquisition electronics for incoming detector and radio signals, it houses a variety of different electronic components optimised for efficient and precise readout of the connected analog or digital measurement instruments. In the current status of the TAXI DAQ, multiple studies have shown that especially the readout of the radio antennas is showcasing some undesired artifacts [1, 20]. The reason for this is not yet known or convincingly proven, so in the frame of this thesis, in-depth research has been done on the functionality, synchronisation and the clock tree timing of the different electronic components on the board with the aim of clarifying the remaining possibilities. In this chapter, the data stream of a radio trace is precisely explained with a focus on synchronisation and timing of the different components. For a graphic overview over all components, their communication and timing paths, please refer to the appendix (chapter A.2).

### 4.1. Radio signal detection and pre-amplification

As mentioned already, the *SKALA-v2* antennas are designed to pick up a wide range of radio signals in a large solid angle. This happens passively - and thus the signal strength is very weak. For this reason, the signal is amplified with a low-noise amplifier (LNA) in the uppermost part of the antenna before further processing. The LNA is of the kind *TQP3M9039* by Qorvo [23] and is powered by 4.5 V DC over a coaxial cable that also serves as a data transmitter. This works, because the constant voltage for powering the LNA is simply a bias shifting the signal's base line by the mentioned 4.5 V. This bias is introduced by the Bias Tee in the next part of the signal chain, the *radioTads*.

### 4.2. Processing the radio signal - the *radioTads*

The analog, pre-amplified signal first passes through a custom front-end signal processing board, commonly referred to as *radioTad*<sup>4</sup> in contemporary usage. This board, of which one exists for every antenna, consists firstly of a so-called *Bias Tee*, a device uniformly known to power readout electronics of antennas. The Bias Tee used in the TAXI is one with the name *TGBT-2R5G+*, designed by Mini-Circuits. It has three ports, of which one is connected to power (4.5 V), one to the further data processing electronics, and one to the LNAs in the antennas. It biases the voltage in the cable from the *RadioTads* to the antennas and thus delivers constant DC power to the LNA. At the same time, the LNA sends its data "on top" of the biased voltage through the same cable, which is split using

---

<sup>4</sup>As told by the creator [24], *radioTad* is an artificial word and has no deeper meaning.




Figure 4.1.: **radioTad frequency response after application of low- and highpass filters.** The upper plot represents the gain distribution, the lower plot shows deviations from its average. The light blue area is the intended frequency response of the antennas. The blue and orange coloured lines represent the two polarisations of one antenna. As the signal is split into 4 differential signals in the radioTads, every polarisation results in a total of 8 channels. [1]

impedance filtering by the Bias Tee [25]. After passing the Bias Tee, the signal is cut by hardware filters. A ULP-340+ by Mini-Circuits acts as a 350 MHz lowpass filter<sup>5</sup>, a SXHP-48+ by the same company cuts away frequencies below 70 MHz. This results in a frequency response graph as seen in Fig. 4.1. An important detail to note is that the gain does not immediately fall off steeply above 350 MHz, but instead linearly diminishes until it reaches a plateau only a little short of 500 MHz. This might appear counterintuitive, as the previously mentioned communication channel at 360 MHz is not eliminated in this phase. In fact, this is done in post by analysis software: It Fourier transforms the signal and puts a 300 MHz lowpass filter over it [1]. The primary purpose of the hardware filters is to prevent ghosting from frequencies above the *Nyquist Limit*, see chapter 4.4.

Once it has passed the filters, the signal is split up into 4 parallel signal paths, each of which is again amplified by THS4508 amplifiers by Texas Instruments and converted to differential pairs [1], before being finally sampled. The post-amps are rated for a voltage supply of 3.75 to 5.25 V [27] and are operated with the same 4.5 V used by the Bias Tee. Until this point, none of the electrical components has needed any kind of synchronisation or clock.

<sup>5</sup>Multiple references state multiple values [1, 20, 26], either 340 or 350 MHz. When considered from a statistical standpoint, there is no hard upper cut, as seen in Fig. 4.1.

### 4.3. Controlling the readout - the FPGA

The main data processing unit of the TAXI DAQ is embedded into its FPGA, a Field-Programmable Gate Array, located on the TAXI mainboard. The used model is a *Spartan-6 XC6SLX45*, developed by Xilinx and can be programmed with a software framework called Xilinx ISE, which includes a coding environment for vhd1 language, and tools for mapping high level code onto the FPGA. This *synthesis* usually involves the software finding a way for programming the FPGA so that its processing power is optimally utilised, with the goal of addressing as many data streams in parallel as possible. The Spartan-6 runs on a 10 MHz reference clock that is synthesised by an *IQD LFTVXO009912* quartz oscillator soldered onto the board. The same frequency is provided by the White Rabbit Network, but as stability is very important, only onboard oscillators are currently in use as reference clocks for the FPGA. There are other oscillators on the TAXI board; one more operates at 10 MHz (currently not in use), and two others run at 25 MHz, feeding the FPGA and the *8720AI* LAN chip by Microchip Technology. This is not explicitly shown in the block diagram in the appendix (chapter A.2).

Being the centre of the data acquisition, among other things, the FPGA is responsible for initialisations of the readout hardware and for processing incoming trigger signals from the scintillators. With its internal electronics, it is capable of taking various clocks as input and turning them into new clocks, running at any multiple of the input frequency [25], using Phased Locked Loops (PLLs). By complicated serialisation of such PLL circuits, even frequencies with decimal values can be produced. In the same way, the external 10 MHz clock is turned into a 118.75 MHz internal clock inside the FPGA which acts as the FPGA's *tick rate*.

The tick rate is a constant value that defines how frequently an FPGA can execute operations. Unlike Central Processing Units (CPUs), FPGAs operate in a First-In-First-Out (FIFO) principle. Once hard-coded, with every tick, the FPGA will request exactly the same information and do exactly the same operations. At a tick rate of 118.75 MHz, the duration such a step can take at most is 8.4 ns. If an operation takes longer than that with no coding measures taken beforehand to use multi-cycle processing, errors will arise in the signal. In general, the FPGA processing is aimed to run as efficiently as possible, which can be achieved by smart vhd1 coding, which then leads to optimal data routing through the FPGA [28].

ISE runs analytics on its synthesised code, comparing its runtime to the given constraints (like desired tick rates and reference clock and jitter values at different temperatures). If the constraints are given correctly, ISE will show well-visible warnings in the synthetisation output, should operations take longer than one tick. The software calculates a *timing score* which, in an ideal scenario, should always be **0**. In the software's simulation of the code, it will add **1** to the timing score for every *picosecond* that is missing for completion of the entire operation within one tick. Incomplete operations can cause input data to be incompletely processed, cut off or mixed up with data of the next tick, shifting its bits. When decoding such corrupted data, this causes disproportionate spikes in the values, which is currently observed in the TAXI radio traces at higher temperatures.

## 4.4. Sampling the radio traces - the DRS4 chips

For effectively working and analysing analog signals with a computer, they must at some point be sampled. Sampling is the process where a continuous signal is transformed into discrete snippets at a consistent frequency. Each of these snippets represents the signal's average over a short duration. As a rule of thumb, when it comes to sampling frequencies, more is always better for an accurate representation of the analog signal. But with higher sampling frequencies, new problems arise - at some point it gets hard designing electronic components that work at such high speeds, and, of course, with the number of samples (essentially data points), the storage size of the entire batch increases linearly with the frequency. So, a sweet spot must be found between an accurate representation of the incoming signal and manageable data size as well as accurate speeds for the chips. Currently, the analog signal is sampled at 950 MHz. This number is justified when looking at the radioTads' frequency response graph (Fig. 4.1).

- - -

There is an important relation in signal processing called the *Nyquist Theorem*. In short, it posits that to sample a continuous signal with a highest frequency of  $f_{\max}$ , the sampling frequency, denoted as  $f_s$ , should be at least double that of  $f_{\max}$  [29], represented by the equation

$$f_s \geq 2 \cdot f_{\max}$$

Every part of the analog signal whose frequency goes above this threshold will create artifacts, also known as aliases, within the wave form. Within the Fourier space, these artifacts appear as "ghosts", which are reflections around the Nyquist frequency  $f_{\max}$ . For a sampling rate of 950 MHz, the Nyquist limit is at 475 MHz, which is slightly but not significantly above the top end of the filtered range. In this current state, a signal with 500 MHz would produce a ghost at 450 MHz.

- - -

Capable of sampling at such elevated frequencies is the DRS4 chip, developed and constructed by the Paul Scherrer Institute [30]. It features 9 channels, of which 8 are in use. The remaining one could be used for sampling an external clock for better timing accuracy [31]. The chip accepts a power supply ranging up to 2.5 V and offers sampling rates between 700 MS/s and 6 GS/s. MS and GS hereby abbreviate Mega-Samples or Giga-Samples, respectively. Each channel of the chip operates separately as a so-called ring sampler: After configuration and its initialisation, the DRS4 chip constantly samples the incoming signal at the designated rate. Once its cache fills up, it begins to replace the oldest data with newer readings. Each channel of the DRS4 chip has a cache of exactly 1024 bins.

Per antenna, one DRS4 chip is used for sampling its data. With two polarisations per antenna, this leads to four DRS4 channels sampling the same antenna signal. However, whether this results in identical data depends on the DRS4 chip's settings. Set by the FPGA, the DRS4 chip can operate in three different *cascading* modes:

- **Cascading.** The four DRS4 channels are provided with samples one after another; only when one channel's buffer is full, the chip will start writing to the next. With

each capable of storing 1024 samples in each ring buffer, the total amount adds to 4096 samples. Once all four are full, the first begins overwriting its oldest bins. At the current 950 MS/s sample rate, this amounts to a coverage time of about 4.3  $\mu$ s once the channels are read out.

- **Semi-Cascading.** Two of the four DRS4 channels operate simultaneously, sampling identical data. This effectively reduces the buffered time interval to (2.2  $\mu$ s), but enhances the statistical accuracy.
- **Non-Cascading.** All four DRS4 channels sample simultaneously. Upon reaching their full capacity of 1024 bins, they begin to replace older data. At the current 950 MS/s sample rate, this mode provides a coverage time of approximately 1.1  $\mu$ s. Although it offers shorter time coverage than other modes, it leads to enhanced statistics and waveform precision, as every DRS4 channel imparts unique nuances that can be mitigated with statistics.

#### 4.4.1. Setting the DRS4 sampling rate

The DRS4 chip will synthesise a sampling frequency on its own with just its input voltage acting as a controller for the sampling speed. This however makes it difficult to precisely control the sampling frequency, in the following named  $f_{\text{DOMINO}}$ . A distinct feature of the DRS4 chip is the capability of running based off an external reference clock, in this case provided by the FPGA. It can be fed into the DRS4 chip through the REFCLK input and acts as the DTAP frequency in the chip's internal logic [31]. This frequency has a period that follows the equation

$$f_{\text{DTAP}} = \frac{1}{2048} \times f_{\text{DOMINO}}.$$

That means, it always switches its state once the domino wave has reached the 512th bin, the one exactly in the middle (between 0 and 1024, the highest bin). It thus takes two entire fillups of the domino ring sampler for generating one period of the DTAP signal. For locking the DRS4 chip to one specific frequency, this DTAP clock must be provided externally - by the FPGA. For generating a sampling frequency of 950 MHz, the refclock must have a frequency of around 460 kHz, a frequency in the output range of the FPGA.

#### 4.4.2. Test Setup: Power draw of the DRS4 chips

There are ongoing discussions within the working group about the idea of lowering the DRS4 chips' sampling rate for the radio traces in the present TAXI and its future developments. One argument is that with a lower sampling rate, the heat generated by the chips might decrease and with it the noise. As this was never tested, a quick setup was developed and carried out in order to find any correlations between the sampling rate and the power draw. For this purpose, a DRS4 evaluation board (labeled v5 #2414) is used. It is a USB-powered, standalone device with 4 inputs that works like an oscilloscope when connected to a PC. On it, 4 individual DRS4 chips work just like on the TAXI, and even though the board's FPGA is an older SPARTAN-3 model, it still retains enough similarity

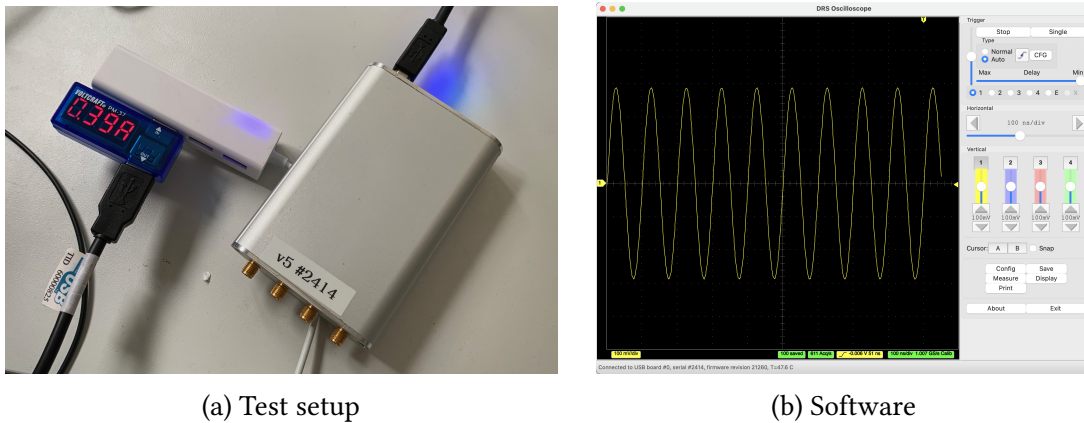


Figure 4.2.: **Setup for measuring the DRS4 Evaluation Board’s power usage.** In (a) from left to right: USB power meter, USB hub, DRS4 Evaluation Board. Image (b) shows an exemplary screenshot of the evaluation software.

to the contemporary version to ensure valid comparisons. After installation, the affiliated software *drsosc* [30] allows easy changes to the board’s sampling rate. For the purpose of this test setup, a USB power meter *PM-37* by Voltcraft is put in between a MacBook Pro and the DRS4 evaluation board, see Fig. 4.2. In the software, the sampling rate is adjusted and the values shown by the power meter are recorded. The power meter alternately displays voltage and current in time intervals of 5 s, from which the power is calculated.

The results are shown in Table 4.1. Very narrow differences are observed. Further investigations on dependencies on the voltage and frequency of the input signal showed correlations so minimal that it was uncertain whether there was any correlation at all. It is interesting though to see the difference between a chip that is not being addressed (i.e. powered by USB, but not communicating, the readout software is not running) and one that is "Paused" (with the trigger set so high that no updates are shown): Even when it is not displaying anything, the sampling is still going on and evaluating the data for anything that might be above the trigger. This is different to the state where it is not addressed at all. Here, there might not even be any sampling happening, thus the lower power consumption. As the difference in power draw was so minimal for all the sampling rates, further research regarding temperature changes on the evaluation board was skipped. Still, some thermal images were captured in order to isolate the primary heat sources on the board, which translate similarly to the TAXI. The images can be found in the appendix (chapter A.5).

The result of this experiment suggests that the state of the antenna readout in fact has only a very minor impact on the overall power consumption of the TAXI and thus on its temperature; it doesn’t matter how high the sampling rates are - or whether it is even sampling at all. Further research has been put into the topic of TAXI workload and its correlation to the operational temperature, whose results are discussed in chapter 5.1.

## 4.5. Triggering Mechanisms

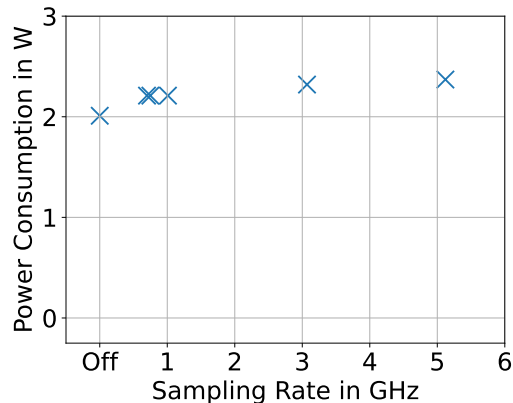
Once fully deployed, the IceCube SAE will be made up of 32 stations which all will feature multiple antennas and scintillators. As described before, each antenna signal is continuously being sampled, the DRS4 chip continuously refreshes its set of samples. At this phase though, no information ever leaves the sampling chips, the state is kept forever, the information is continuously refreshed but never processed or forwarded. And as high-energy events are quite rare and only happen a couple of times per month, most of the data would simply be electrical noise from the amplifiers in the signal chain. Because one second of radio data produces a file several Megabytes large, constantly storing all radio data would quickly overload any available storage capacities. This is why a triggering requirement has been developed. Currently, on the prototype, all three of the following requirements must be fulfilled for the antenna readout to trigger [32]:

- at least 6 scintillators of one station must record
- an event above their individual energy threshold
- within a time frame of 1  $\mu$ s.

Later, once fully deployed, the trigger conditions can and will very likely change. This is easily done by reprogramming the FPGA, as discussed in the subsequent section. But the general idea stays: The scintillators trigger, and then the antennas are read out - all antennas, over the entire surface of the experiment. And because this can mean distances of more than a kilometre, it will take data several microseconds from one end to the other, so by the time the trigger signal reaches the other end of the array, the DRS4 chips there will picture a greatly different time interval than where the trigger initially came from. Because the DRS4 chips constantly overwrite their own information, a uniform image over the entire experiment is currently not possible and a solution is yet to be devised. A short buffer would solve the issue, but such a component is hard to implement, as

Table 4.1.: **Measured power consumption of the used DRS4 Evaluation Board in dependence of changes to the sampling frequency.** "Offline" refers to the state at which the board is connected to power, but not being addressed by the readout software. "Paused" refers to the state of an infinitely high trigger, so that no new data is being sampled.

GS/s	U/V	I/A	P/W
Offline	5.16	0.39	2.01
0.70	5.15	0.43	2.21
0.75	5.15	0.43	2.21
1.01	5.15	0.43	2.21
3.07	5.15	0.45	2.32
5.12	5.15	0.46	2.37
Paused	5.15	0.45	2.32



it is not soldered onto the board. A new DAQ device (possibly for IceCube-Gen2) will supposedly have that feature, but it is still in early stages of development<sup>6</sup>. Still, for the current prototype, the local trigger requirements stated above are used as a criteria for reading out the DRS4 chips. Responsible for this is, once again, the FPGA.

For testing purposes, a so-called software trigger ("soft trigger") can also be activated. It bypasses the scintillator trigger criteria and instead used a fixed time interval at which the antennas are read out. Those readouts have no physical relevance, but are useful for debugging.

## 4.6. Readout of the DRS4 chip and analog-digital conversion

With the trigger requirements met, be it software or hardware trigger, a signal is sent from the FPGA to the DRS4 chips to stop the sampling. Now starts the so-called dead time, in which no new data will be recorded. The dead time is used for reading out all bins in all channels of the DRS4 and, at the current configuration, is approximately 70  $\mu$ s. The readout is facilitated by analog-digital converters (ADC), specifically the LTM9007IY-14 model from Linear Technology. Each DRS4 feeds directly into an ADC. Those have PLLs on their own, and are supplied with a 33 MHz clock from the FPGA. This is the recommended readout frequency by the DRS4's manufacturer [31]. It balances dead time with distortion resulting from excessively high readout frequencies. Both the DRS4 chip and the ADC are supplied with the same 33 MHz clock, but phase-shifted slightly to account for the time the information takes to travel from DRS4 to ADC. The ADC converts the incoming, sampled, analog signal into 14-bit digital samples. This means that between the lowest and highest possible amplitude, there are  $2^{14} = 16\,384$  steps, often referred to as ADC counts. A single ADC count corresponds to

$$\frac{1}{16\,384} \times U_r,$$

with  $U_r$  being the maximum dynamic range of the ADC above which it saturates and loses its linearity.

After the analog-digital conversion, the 8 channels are simultaneously read out by the FPGA. Having 8 free pins each read out 14-bit information at 33 MHz in parallel is challenging from a hardware design perspective, since it would require 8 high-speed inputs on the FPGA. Instead, the 8 pins are serialised to a much higher frequency that is synthesised using an internal PLL on the ADC, so that in the end, the 14 bits of information per channel are read out over one channel with 8 times the frequency,

$$f_{\text{readout}} = 8 \times 33 \text{ MHz} = 264 \text{ MHz}.$$

On the FPGA itself, a deserialiser then converts this information back into separate channels which is then forwarded to the on-board ARM-based Linux system, which then uses the White Rabbit Network to forward this data to the IceCube Lab.

---

<sup>6</sup>This new DAQ will surpass the TAXI and is expected to be completely redesigned among others with the KIT Institute for Data Processing and Electronics (IPE), but is still in discussion about its desired functionality and capabilities to serve large scale detector arrays best.



Currently, the DRS4 readout happens twice in a row. This has been observed to significantly reduce errors in the data (see next paragraph), while, on the other hand, increasing the dead time to the above mentioned  $70\ \mu\text{s}$ . That one ends as soon as all readouts are done. At this point, after indication by the FPGA, the DRS4 chip resumes its sampling again.

## 4.7. Bin spikes and timing issues

What has caused a good number of discussions among team members is the severe corruption of radio traces read out from the DRS4 chips. Currently, almost each recorded radio trace features at least one remarkable spike in the waveform that is clearly artificial and exceeds the antennas' energy range extensively, as can be seen in Fig. 4.3. It has been found that the amount of spikes is correlated to the TAXI's close environmental temperature [20], with readouts above room temperature being especially corrupted. Further research has shown that the reason for the spikes are actually bit shifts and bit flips, the latter occurring from bits of state "1" becoming bits of state "0" [1]. In the past, there have been fixes applied to the firmware like the double readout mentioned above, and the remaining minimal hardware issues should account for roughly one spike in a million samples [28]. Currently, the frequency is closer to one in a thousand, so the spikes must originate from timing issues in the FPGA. Especially at high-speed high-precision timing ports like the serialised readout of the ADCs, the described bit flips and shifts can occur if the entire readout operation is not finished within one tick of the FPGA. Then, bits coming from the serialised readout of the ADCs can bleed into the next tick, shifting the incoming data backwards. As the binary data is later interpreted, having defined start and end signatures of each sample is crucial. The immediate solution would be to probe the current timing constraints in Xilinx ISE for practicability, then test the current FPGA configuration for timing errors and address them with code optimisation and realistic constraints. Learning the FPGA code language and fixing the code unfortunately exceeds the frame of this thesis.

## 4.8. Global synchronisation, drift and data tagging

In the end, once the ADCs are read out, the FPGA passes the filtered, triggered data to the on-board Linux system, which tags and forwards it as an i3 file to the central lab at the South Pole, or in the case of the experiments in this thesis, to a computer located next to the TAXI. This communication happens over the White Rabbit network, which transports the data optically to its receiving end. For this, the WR-LEN client (see chapter 3.3.5) is used.

It has only one input: The glass fibre, which carries all information. On-device, one instance of the signal is looped through for chaining of multiple WR-LENs. Alongside that, the existing fibre signal is decoded, so that in total, three outputs head from the client to the TAXI board: Two SMA cables, labelled 10MHZ and PPS, and one Ethernet cable for communication and data transfer. The two SMA cables yield the necessary timing information for synchronising the TAXIs and tagging recorded data, namely




Figure 4.3.: **Radio Traces from the South Pole.** Taken at the prototype station, using DRS4 cascading sample mode. Depending on the antenna, more or less spikes are visible. The different colours represent different events. [33]

- a) a square-wave clock signal at 10 MHz that is currently only used for the scintillator readout, but is intended to also feed into the TAXI FPGA in the future,
- b) a so-called PPS signal, which provides exactly one clock edge every second to counteract the occurring drift in the 10 MHz signal - with this, it is brought back on time every second anew,
- c) an *IRIG-B* synchronisation clock, running at 100 Hz, transmitted over the same cable as the PPS signal. This channel conveys actual text messages from the central White Rabbit node to the TAXI for synchronising its internal clock to the global one.

The synchronisation between individual TAXIs is just as important as an accurate timing within the TAXI board itself. The White Rabbit network fulfils this role with the PPS signal. As mentioned, at 100 Hz, the IceCube Lab computer sends text data to all receiving WR-LEN clients containing the current date and the time of day in seconds (cumulative, so no minutes or hours). This is the reference used for tagging the individual events that are taken - it is time synchronised per satellite. The FPGA itself though is capable of taking event data at much higher frequencies than 100 Hz - so interpolation is needed. This is possible by a so-called *Real-Time Counter (RTC)* that starts counting upwards as soon as the TAXI is launched. It is linked to the FPGA's internal tick rate, resulting in a frequency of 118.75 MHz. As a result, this number grows significantly in a short span of time. That is why it runs on a 64-bit architecture, so even at this high frequency, it takes this variable almost 5000 years to overflow. With this value constantly counting up in a regular manner, it is useful for interpolating in-between the 100 Hz time stamps sent via the WR network. Basically, a log file is created as soon as there is signal coming from the PPS channel that stores the 100 Hz real-world time stamps next to the FPGA's RTC value

in this exact moment. This comparison chart is then used for interpolating the time for the events recorded by the TAXI. In this configuration, all TAXIs across the entire IceCube surface remain synchronised (with minor shifts in radio data between distant TAXIs, see chapter 4.5).

The need for installing a White Rabbit client in every FieldHub for synchronisation and communication purposes is thus clear. While this technology is effective and offers numerous advantages, there are challenges that come alongside using such high-speed communicators and will be explained in the next chapter.



## 5. Studies on temperature dependences of the TAXI

As shown in reports [20, among others], heat negatively affects the quality of the acquired data by the TAXI. The stability of the communication interface to the scintillator panels also worsens with increasing heat, generating corrupted data packages and increasingly failing to initialise. At the South Pole, the insulation within the FieldHub has already been reduced in order to lower the temperature inside it [33]. The temperature itself was measured to be around 50 °C with environmental temperatures of about –20 °C, the resulting radio data showed significant corruption. Since removing the insulation, the TAXI’s temperature has started to be logged and was measured to reach up to 35 °C during the summer and up to 10 °C during the –80 °C cold winter [21]. For effective data acquisition, this is still too warm. It was therefore motivating to investigate different sources of heat within the TAXI box and how their impact could be reduced.

### 5.1. Impact of TAXI data acquisition on the temperatures

It was not yet known how much the TAXI data acquisition affects the temperature of the device. A short experiment was thus conducted for identifying differences in temperature and power consumption between an idle TAXI and a running one.

#### 5.1.1. Setup

A cold TAXI (unit #22) is powered through a smart plug (*SH-P02* by *Deltaco*) that allows monitoring the power consumption in Watts. The bootup script of the TAXI is modified beforehand to not automatically start the DRS4 chips’ sampling. With this setting, the TAXI is powered up and the temperature of the mainboard is taken at two different locations, as shown in Fig. 5.1) at approximately 60 s time intervals<sup>7</sup>. In the following, the temperatures will be named  $T_{\text{SOC}}$  for the temperature right underneath the onboard Linux SOC (sensor *temp05*) and  $T_{\text{DC}}$  for the temperature next to the DC converter (sensor *temp112*). A pre-programmed USB temperature sensor, *PRO-USB-2* by RS Pro, is fastened to a table at the same height as the TAXI using strong tape (see Fig. 5.9b). The TAXI is then covered with a wooden box to simulate the environment of a FieldHub, described in more detail in chapter 5.2. This setup is allowed to run for several hours. Afterwards, initialisation scripts for activating the DRS4 chips are run. Then, the scintillator panels are initialised one by one. Once initialised, the eventSender is activated for 8 h. During

---

<sup>7</sup>The script is run in a loop with a sleep interval of 60 s, which results in an actual interval between measurements of slightly above 1 minute.




Figure 5.1.: **Top View on the TAXI mainboard.** From top to bottom, it houses connectors to the radioTads, the DRS4 chips with their ADCs, two temperature sensors, and other components, like the LINUX SOC, the FPGA and a DC converter for the correct power delivery. Modified from [20].

the entire process, the TAXI's onboard temperature is recorded and logged. Using the smart plug, in spot checks, the power draw is acquired. As the plug does not support data logging, this was the only viable option.

### 5.1.2. Implementation

It is important to note that the different steps that were described above happened over the time span of several days, meaning the acquired data from different steps are independent from each other. As it proved challenging to initialise the TAXI successfully with onboard temperatures over  $30^{\circ}\text{C}$ , after the radio measurements, the TAXI was powered down in order to allow the mainboard to cool. Even from a cold start though, using the default scripts, after a few minutes, errors started arising while the initialisation was still in progress. Thus, a workaround method was used for initialising the scintillators manually one by one. This approach takes longer, lacks onboard communication regarding full buffers and is currently only used due to issues related to temperature. Once fully deployed, a simpler script is foreseen to initialise the scintillators with just one command.

### 5.1.3. Evaluation

The temperature trends for the different stages can be seen in Figures 5.2, 5.3, 5.4 and 5.5, together with the measured power draw. There is little to no observable difference in power draw between a TAXI with active radio sampling and a completely idle one. Here, the TAXI consistently runs at around 25 W with varying temperatures in the range of  $61$  to  $64^{\circ}\text{C}$ . The temperature measurements with active DRS4 chips were conducted several days later with higher ambient temperatures and a seemingly rising curve in TAXI temperature. Given that the power draw remained consistent, radio sampling is not

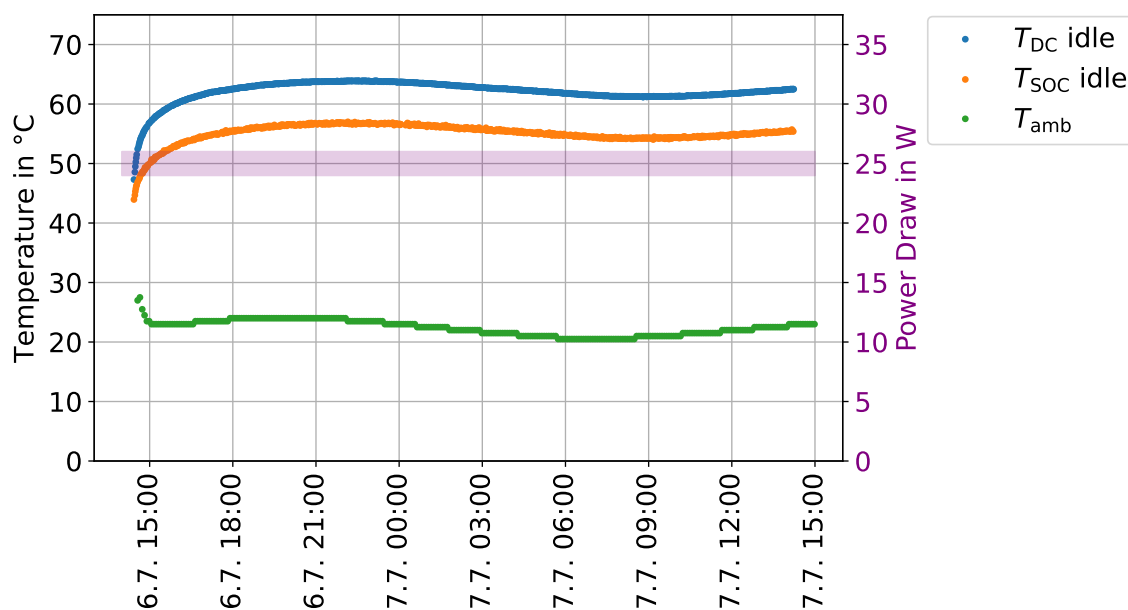


Figure 5.2.: **Temperatures and Power Draw of an idle TAXI.** The power draw is not continuously logged, so it is only schematically drawn in as a purple area. After the cold start, the onboard temperatures follow the ambient temperature trends, never reaching more than 64 °C.

believed to significantly influence the temperatures. It was not logged in detail whether and for how long the TAXI was shut off in-between those measurements, but the rising temperatures could be a response effect of the board heating up after a cold start as it would be the case under normal conditions.

The scintillators were only initialised days later, again starting from a cold TAXI, which is visible in the temperature curve. The initialisation process takes around 5 min to complete, where during this time frame a steady increase in the power draw can be observed. Once initialised, the power draw stays around  $(30 \pm 1)$  W, with temperatures steadily increasing. The measurements ended before the peak was reached. Nevertheless, based on the power consumption data, the missing information can be estimated.

Because of problems with the initialisation, the eventSender could only be started several days later. It did not affect the power draw of the TAXI at all, whose temperature peaked around 66 °C, which leads to believe that this is most likely the point at which the rising curve of Fig 5.4 would have stabilised.

When directly compared to an idle TAXI (Fig. 5.6), only little differences in temperature can be made out. In this graph, different measuring days are combined. In order to make comparisons easier, the starting point of the *idle* dataset has been shifted so that the ambient temperature trends look roughly the same. It is important to mind the influence of ambient temperature on the onboard measurements, as the following experiments in chapter 5.2 will show a strong correlation between them. However, based on the discussed figures alone, it can be inferred that the TAXI's workload influences the onboard temperatures by a maximum of 2-3 °C.




Figure 5.3.: **Temperatures and Power Draw of a TAXI whose DRS4 chips are sampling radio data.** Conducted on a different day than Fig. 5.2. As a cold start was needed and the ambient temperature was high, the temperature increase is not expected to originate from the radio sampling, but instead from the TAXI board still heating up after the cold start.

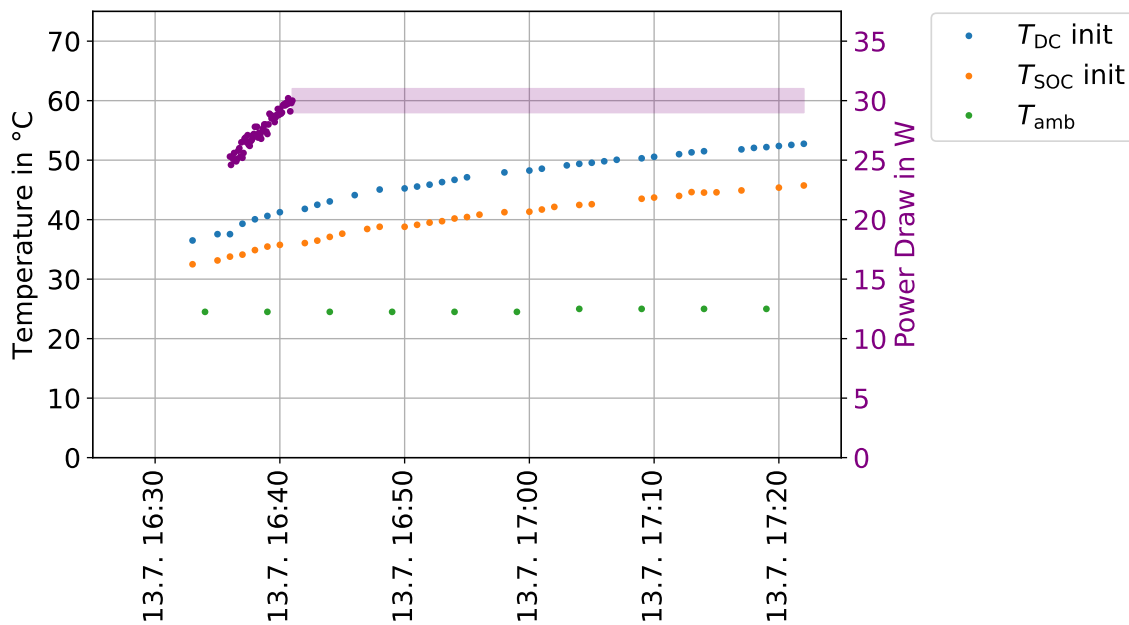


Figure 5.4.: **Temperatures and Power Draw of a TAXI whose connected scintillators are being initialised.** Logging started right after booting the cold TAXI. The initialisation takes around 5 min to finalise, steadily increasing the power consumption. The temperatures are again not meaningful, as the initialisation was needed to be performed right after a cold start.






Figure 5.5.: **Temperatures and Power Draw of a TAXI under full workload.** Over the course of several hours, the temperatures still exhibit the ambient temperature's trends. The temperatures reach a maximum of 66 °C.

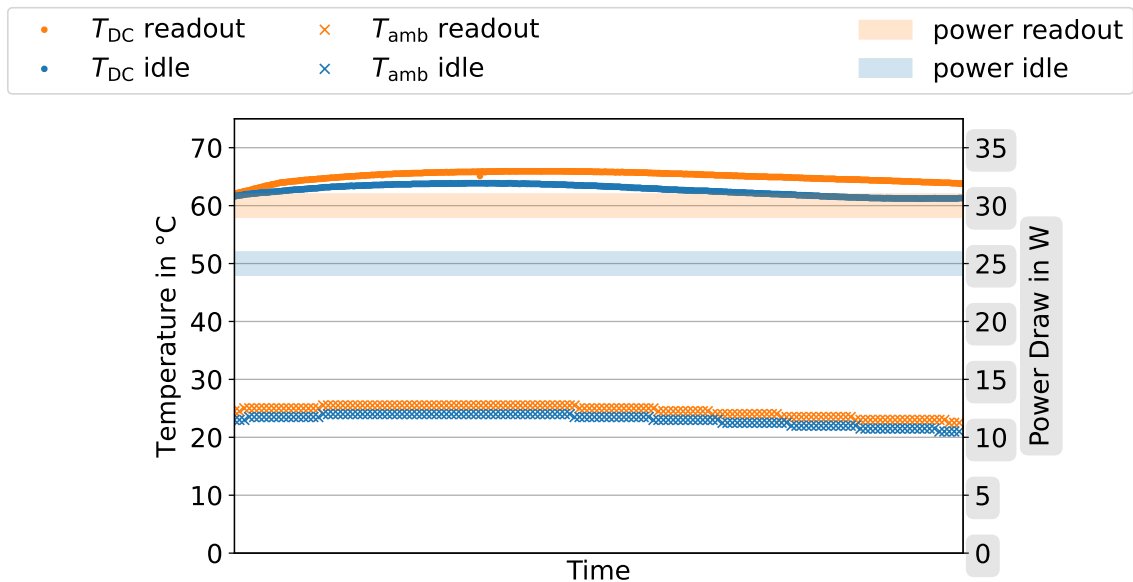


Figure 5.6.: **Comparison of an idle and a full workload TAXI.** As both measurements were taken on different days with different conditions, no marks are given at the x axis. The graph shows the same time interval for both measurements. The colour tinted areas are to be interpreted using the y axis on the right.




Figure 5.7.: **Placement of the WR-LEN standalone unit next to TAXI #09.** Even though the TAXI box is also warm, its physical properties lead to reflected IR radiation instead of black-body emission towards the camera.

## 5.2. Different placements of the White Rabbit Client

One of the main sources of heat inside the TAXI box is the local White Rabbit client (*WR-LEN* by Seven Solutions), which is measured to draw  $(5.9 \pm 1.1)$  W, based on a 5-minute measurement series with the smart plug used in the setups explained in the previous chapter<sup>8</sup>. As the entire TAXI only draws around 30 W of power at most, the idea of moving this board out of the TAXI box sparked multiple discussions within the responsible teams. IR images, as shown in Fig. 5.8, underline this theory further. Adding to the list of advantages is the hereby-involved process of making the power supply of the individual White Rabbit clients independent of the TAXI's, which currently is not the case. Should a single TAXI require maintenance, all other TAXIs daisy chained to the faulty one would go offline the moment the first one is disconnected from power. With the White Rabbit client as a standalone device, the TAXI could be worked on independently from the rest of the chain, keeping the impact on data taking at a minimum.

The entire process of moving the client would not take much time, as the housings for the boards are readily available and the TAXI boxes feature the ethernet port that is needed for such a change already on the outer shell. The WR-LEN would then be placed separately in the FieldHub. Whether the effort is worth the time is the motivation for the following measurements.

### 5.2.1. Setup

The setup for the experiment underwent several iterations. The final one consists of two TAXIs, TAXI#09 and TAXI#22, both placed on wooden panels that initially served

<sup>8</sup>The calculations done for estimating the mean value are briefly explained in the appendix (chapter A.1).

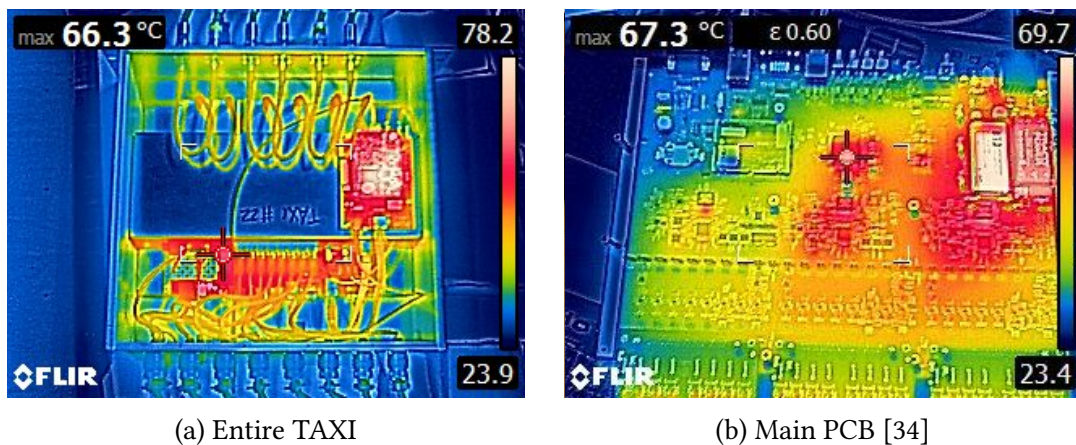


Figure 5.8.: **IR images of the insides of the current TAXI.** Well visible in Fig. (a) is the topmost board, the WR client. Fig. (b) shows an isolated view on only the lowest level PCB, housing the temperature sensors, the DRS4, ADC and FPGA units. More detailed images can be found in the appendix (chapter A.5).

as lids for shipment boxes. TAXI#09 has its internal WR-LEN disconnected and instead gets a standalone unit that is placed next to it, also on the wooden panel (Fig. 5.7). The standalone WR-LEN is connected to the TAXI via two SMA connectors (PPS and 10MHz clock, see chapter 4.8 for more detailed information), as well as via Ethernet (data transfer and communication). It has its own power delivery, whose power draw is separately measured. One TAXI is connected to the antennas, whose effect has been evaluated to be negligible in the previous chapter. Both TAXIs are connected to 4 scintillators each, for similar setups. The cables (antenna and scintillator cables and the glass fibre for the WR network) are routed out and fixated with strong tape. Pieces of polystyrene are used for creating a basic form of insulation around the edges of the wooden plate. Then, the used shipment boxes are equipped with RS-Pro's *PRO-USB-2* temperature sensors (Fig. 5.9a), which are fastened with strong tape and additional stapler pins. The sensors are initialised beforehand and actively measure and log the temperature in a predefined time interval, here set to 5 min.

Then, the crates' tops are put over the setups one by one, so that the TAXIs are now within sufficiently isolated spaces similar to FieldHubs at the South Pole. Additional polystyrene is then used for keeping the lid's pressure from the cables and for filling remaining large air gaps. Another RS-Pro USB temperature sensor is fastened next to the setup close to the identical height for measuring the ambient temperature in the area (Fig. 5.9b). A nearby window is boarded with polystyrene, in order to keep direct sunlight out. Then, the TAXIs are powered up, and on the connected PC, a script is executed which reads out the onboard temperatures for both TAXIs in fixed time intervals of about 60 s and saving that to a log file onto the PC's hard drive. In this manner, the setup is left untouched for around a week. After that, all temperature data is synchronised to a common time base and evaluated.




Figure 5.9.: **Location of the USB temperature sensors.** Both are fastened with strong tape, the FieldHub sensor is additionally held in place with stapler pins (not visible in this image).




Figure 5.10.: **Experimental setup for measuring the WR-LEN's impact on TAXI temperatures.** The TAXI below the table is TAXI#09 and has its WR-LEN outside the metal housing. The wooden boxes simulate the FieldHubs in use at the South Pole.




Figure 5.11.: **Cleaned Temperature Measurements of all sensors in the WR-LEN experiment.** The local trends correlate strongly with the ambient temperature, with the temperature offset showing the actual influence of the WR-LEN client.

### 5.2.2. Evaluation

The measurement run had a duration of 9 days, with a total of 7 sensors constantly taking data. After reading out the USB sensors and fitting all data to their respective time stamps, the result could be visualised as seen in Figure 5.11. The onboard sensors all produced ghosts, with temp05 in TAXI #09 (labelled " $T_{SOC}$  out") being the exception, and are already cleaned up in the plot. For reference, the unedited data is visualised in the appendix. The correlation between all the sensors with the ambient temperature  $T_{amb}$  is well visible. However, besides that, the positioning of the WR-LEN client does shift the curves vertically, making up a difference of several degrees Celsius. Inside the TAXI box, the WR-LEN is positioned directly above the temp05 sensor ( $T_{SOC}$ ), and here, the differences are the most obvious, with an average shift of  $(2.5 \pm 0.2)^\circ\text{C}$ . Uniformly, with the WR-LEN inside the TAXI box, the temperature is higher - even in the FieldHub. A more detailed view, comparing the single curves, can be seen in Figures A.4, A.5 and A.6, in the appendix (chapter A.4).

In order to compare the number of spikes in the radio traces, radio waveforms were taken from both TAXI #22 and TAXI #09 after the measurement run had ended. Exemplary pictures of two random events from this measurement for each TAXI can be seen in Fig. 5.12. No antennas were connected, so the waveforms should only show the baseline of the DRS4 chips and onboard interferences (noise).

Unfortunately, the different DRS4 channels show significantly greater variation among themselves than any temperature difference would cause. Additionally, TAXI #22 (WR-

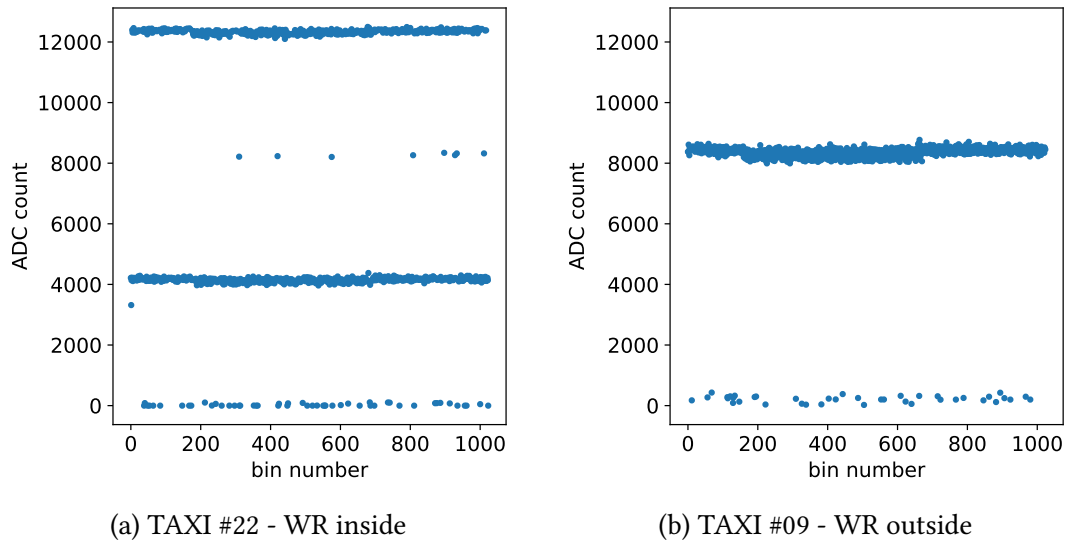


Figure 5.12.: **Exemplary radio traces of both TAXIs.** Non-cascading, taken at full running temperature,  $T_{DC} \approx 62^\circ\text{C}$ . Mind the two baselines in TAXI #22 and the regularly occurring spikes.

LEN inside) shows two baselines, which has been observed with hot TAXIs before [20]. Quantitatively, out of a total of 6408 datasets each containing 1024 bins<sup>9</sup>, an average of  $51 \pm 30$  spikes per trace were recorded, ignoring the doubled baseline. This makes a probability of  $(5 \pm 3)\%$  for detecting a spike in each data bin. The uncertainty is immense, which is due to the different channels' behaviour.

TAXI #09, with the new configuration of a standalone WR-LEN outside the TAXI box, meets the expectations, showing only one baseline, but also showing spikes. The frequency of these again depends heavily on the different ADC channels - the same signal is divided into four separate DRS4 chips after all. Quantitatively, TAXI #09 features on average  $10 \pm 11$  spikes, some channels going as high as 50 and more spikes, some as low as 1 per dataset. The probability of spiking here is  $1 \pm 1\%$  for each of the 1024 bins.

This is unexpected, as while this is the TAXI with a lower onboard temperature, the difference in spike probability is so high that a correlation with the temperature is likely to be overshadowed by a larger, systematical uncertainty which most likely has to do with hardware uncertainties. This is why, afterwards, the two TAXIs were compared in identical configurations.

### 5.3. Comparison of two similar TAXI setups

After finding the greatly different behaviours of TAXIs #09 and #22, the same two devices are compared in more detail. For that, both are set up exactly the same way: The metal boxes are opened, their top cover removed (for reaching lower onboard temperatures).

<sup>9</sup>One event is recorded by 3 antennas with 2 polarisations each, which are each splitting up into 4 DRS4 channels, making a total of 24 datasets per event.

Both onboard WR-LENs are in use. The TAXIs are placed on a table, right next to each other, for similar temperature environments. Then, they are powered up, the DRS4 chips initialised, and then left running for 2 h in this configuration. Then, with both, radio traces are read out, around 250 events. These are then analysed and probed for bin spikes.

A channel-wise investigation of the first 4 events for each TAXI can be seen in Figures 5.13 and 5.14. Even though both TAXIs are placed right next to each other, TAXI #22 runs 2 °C hotter and features disproportionately more spikes. Its baseline also frequently splits up from around 8000 to 4000 and 12000. This is not perceivable in TAXI #09, which not only runs cooler, but also shows way less bin spikes. From this short experiment it can be concluded that different TAXIs react very individually to temperature and therefore, before deployment, must be probed for spikes at different temperatures. This can be done with a measurement setup described in the appendix (chapter A.3).

## **5.4. Additional heat sources**

While equipped with the IR camera, some more photos were taken of the White Rabbit switch located in the central IceCube Lab at the South Pole, and of the DRS4 evaluation board's PCB, with the goal of finding hot spots on their housings and PCBs. These images, alongside others, can be found in the appendix (chapter A.5). The DRS4 boards heat up most where their FPGAs and DRS4 chips are located. The switch heats up more where it is connected to power, which is expected to be the spot at which the internal DC converter is located.

5. Studies on temperature dependences of the TAXI

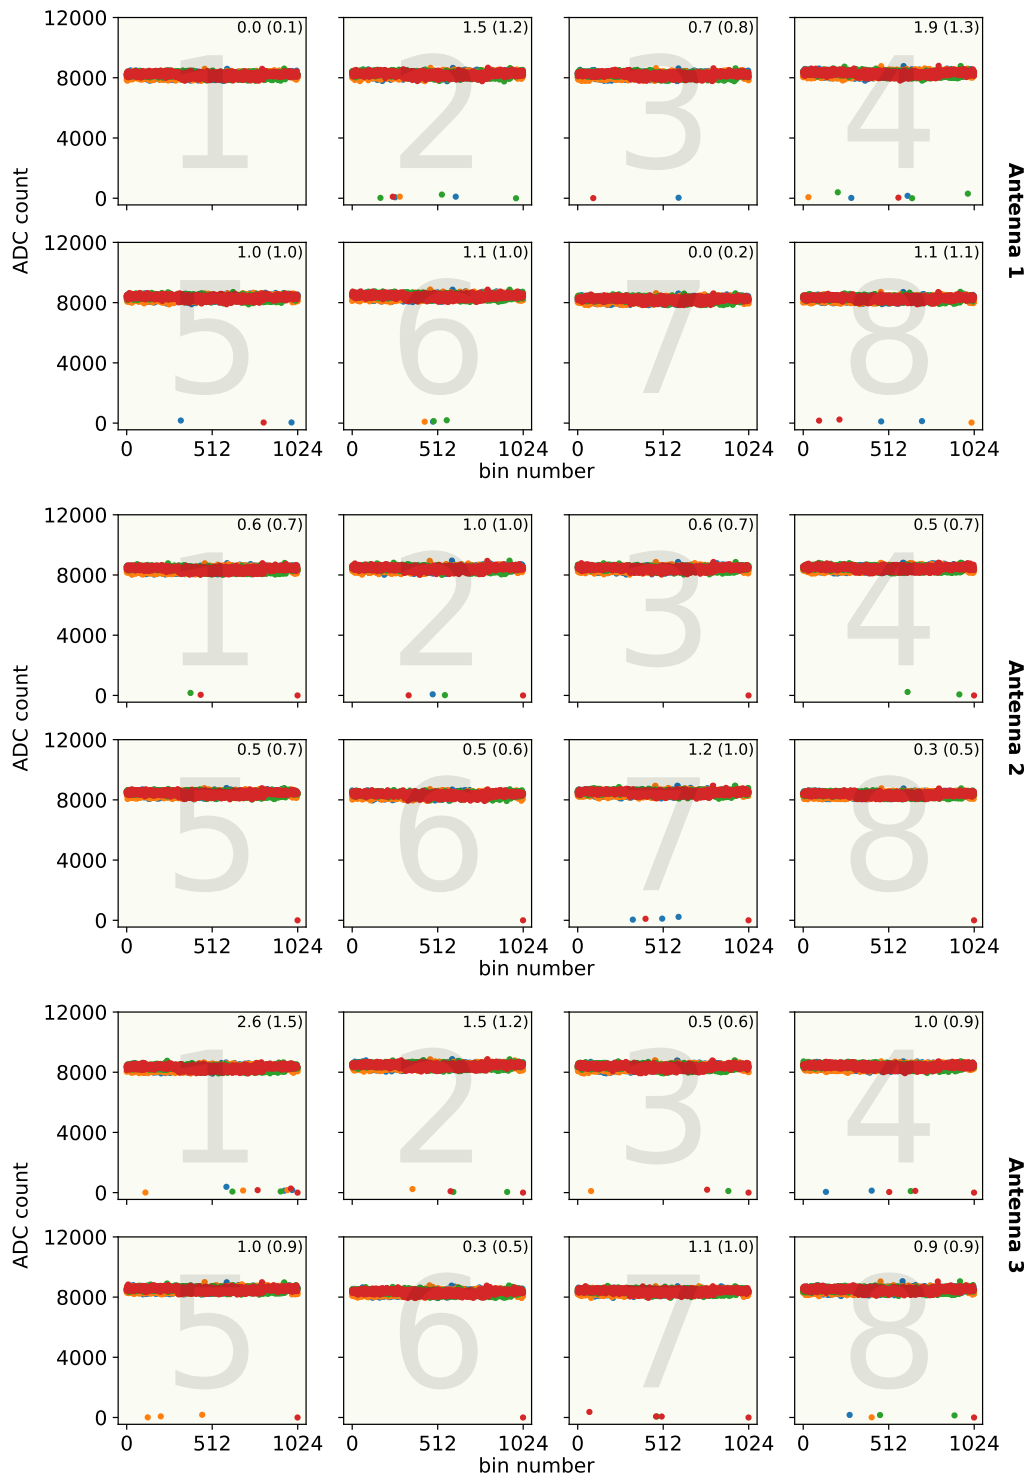


Figure 5.13.: **Radio Traces of TAXI #09 in a standardised setup.** The different dot colours represent 4 different triggered events. The large numbers in the plots' background reference the non-cascading DRS4 channels (1-8). The small numbers in the corner describe the average number of spikes per 1024 bins over all 235 recorded events, with the standard deviation in brackets. No antennas were connected. Onboard temperature:  $T_{DC} = 42.00\text{ }^{\circ}\text{C}$



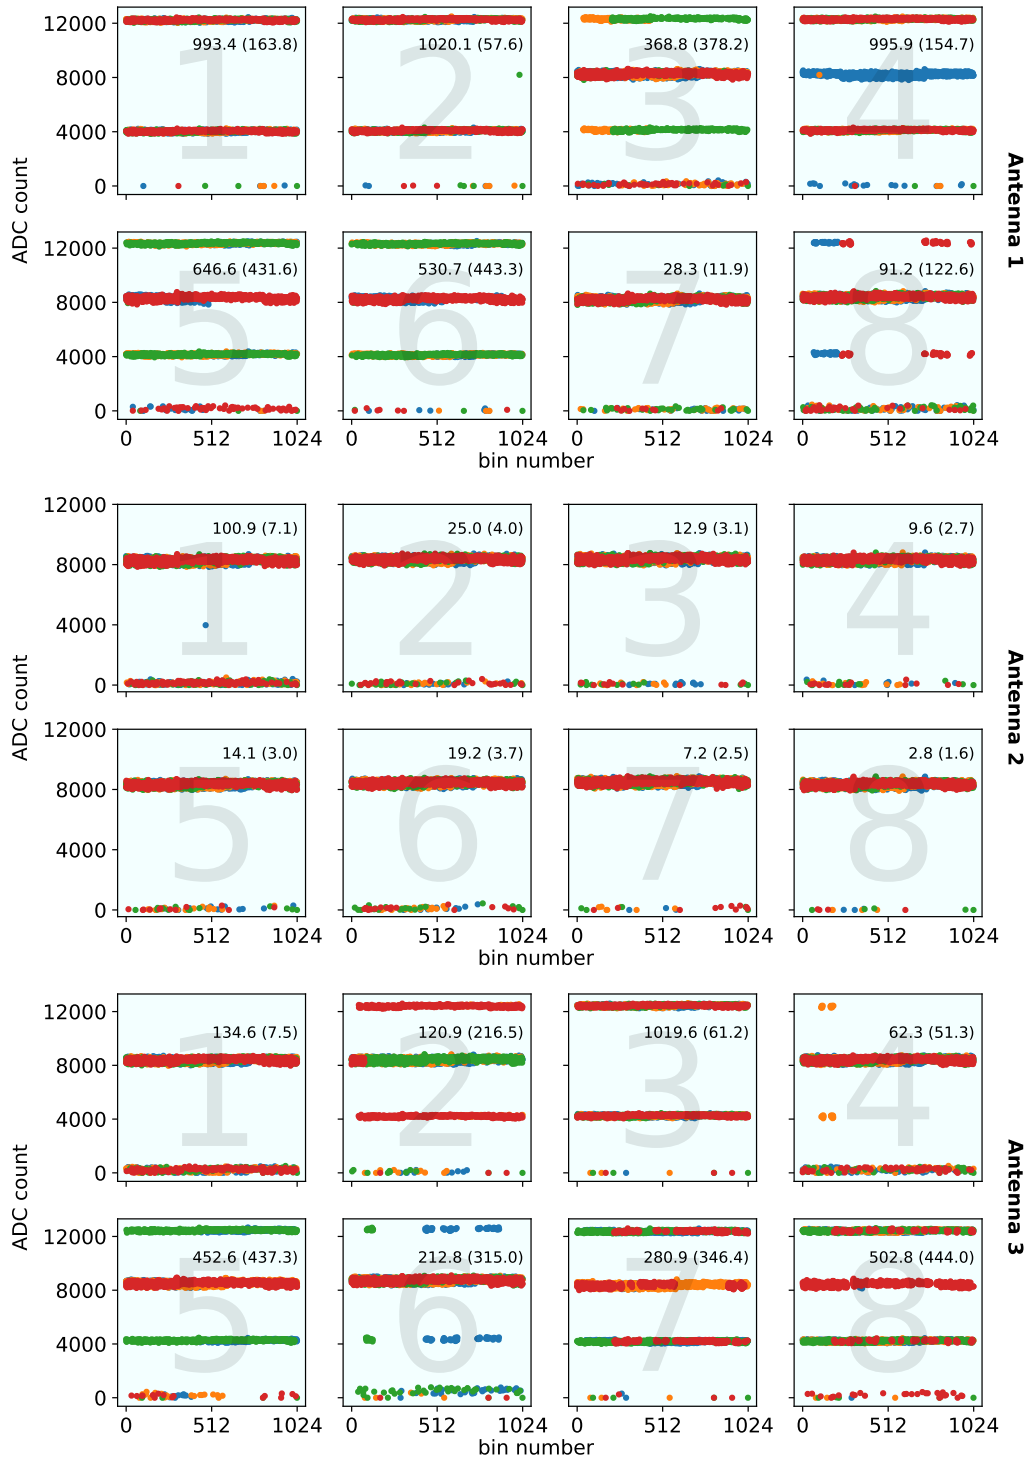


Figure 5.14.: **Radio Traces of TAXI #22 in a standardised setup.** Every bin that has its value not between 6000 and 10000 ADC counts (with the supposed base line being at around 8000), counts as a missed bin, increasing the 1024-bin average in the plot corners over all the 295 events. No antennas were connected, DRS4 set to non-cascading. Onboard temperature:  $T_{DC} = 44.25^\circ\text{C}$



## 6. Conclusion and Outlook

All the research done in this thesis was driven by the need to address the recurrent artificial radio waveform spikes observed in the data acquisition devices of the IceCube Surface Array Enhancement, TAXI. While no modification has been made to the hardware, the results of the conducted research should help better understand the source of these problems and map out plans for clearing them. The issues within the TAXI can be mitigated by optimised FPGA programming taking higher temperatures into account, and a relocalisation of the onboard White Rabbit client.

With the structured and detailed information about all the important components in the signal chain and their communication in chapter 4, further work will hopefully be accelerated with problems easily narrowed down to specific sources. On the TAXI board itself, there are still tasks to fulfill before the full deployment of the Surface Array Enhancement. Namely, the FPGA's code should be looked at in more detail, especially regarding Xilinx ISE's timing score, with the goal of optimising it to stay zero even at higher temperatures. Many of the onboard electronic-related issues are induced by heat. Accordingly, if good thermals can be guaranteed inside the FieldHubs at the South Pole, making the data acquisition electronics themselves more robust to temperature changes will no longer be the primary goal. For better consistency, it should be focused on making TAXIs behave more similarly under temperature fluctuations, based on the results of chapter 5.3.

Other experiments performed within this thesis have shown that there is indeed the possibility of reducing the heat inside the TAXI box by several degrees Celsius by moving the White Rabbit client out of the TAXI housing and into its own (chapter 5.2). This spreads the heat sources further apart, and the better distributed those are, the better the South Pole's environment for passive cooling will dissipate the generated heat. Therefore, once deployed at the Pole, the temperature benefits may very well be even more significant than observed at experiments at KIT.

Reducing the DRS4 chips' sampling rate likely has no effect on the temperatures, as implied by the results of the test run in chapter 4.4.2. The same is true for any reduction in the TAXI workload by initialising fewer scintillators or powering fewer antennas (chapter 5.1), which is also not feasible in practice.

In summary, while significant progress has been made, more research and development is required to further improve the TAXI DAQ in its current state. With the obtained results, the path towards an appropriately operational TAXI is now more clear, but certainly not accessible without additional investigations. Given the ongoing development and expected deployment of the next Surface Array Enhancement stations in the coming years, there is significant potential for refining the TAXI DAQ into a reliable device for analysing the information our Universe sends to us.



# A. Appendix

## A.1. Algorithms used for uncertainty calculation

In chapter 5, on several occasions, values are written in the format  $\mu \pm \sigma$ . This refers to a measurement series' average or mean value  $\mu$  and the respective standard deviation  $\sigma$ . For calculating the mean value, a Python script is used, which executes the calculation

$$\mu = \frac{1}{n} \sum_{i=1}^n x_i$$

over all  $n$  data points  $x_i$ . For calculating the series' standard deviation, another script calculates the result of the formula

$$\sigma = \sqrt{\frac{\sum (x_i - \mu)^2}{n}}.$$

The mean value  $\mu$  hereby represents the expectation value: When assuming a Gaussian probability distribution for all possible values that can be measured,  $\mu$  will have the highest probability. This is a theoretical value though, as in reality the statistics will only converge towards it with more precise measurements and a higher number of overall samples. Still in the Gaussian distribution, the standard deviation  $\sigma$  then states the interval around the mean  $\mu$  in which approximately 68.2% of all measurement values would lay, again in the limit case of an infinite number of measurements.

## A.2. Graphic overview over the TAXI components

As this thesis' research is based on R. Turcotte-Tardif's doctoral thesis [1], the block diagram shown there could be used as a baseline for edits regarding clocks and timing paths of the individual components. The schematic is shown in Fig. A.1. Please refer to chapter 4 for a detailed explanation.

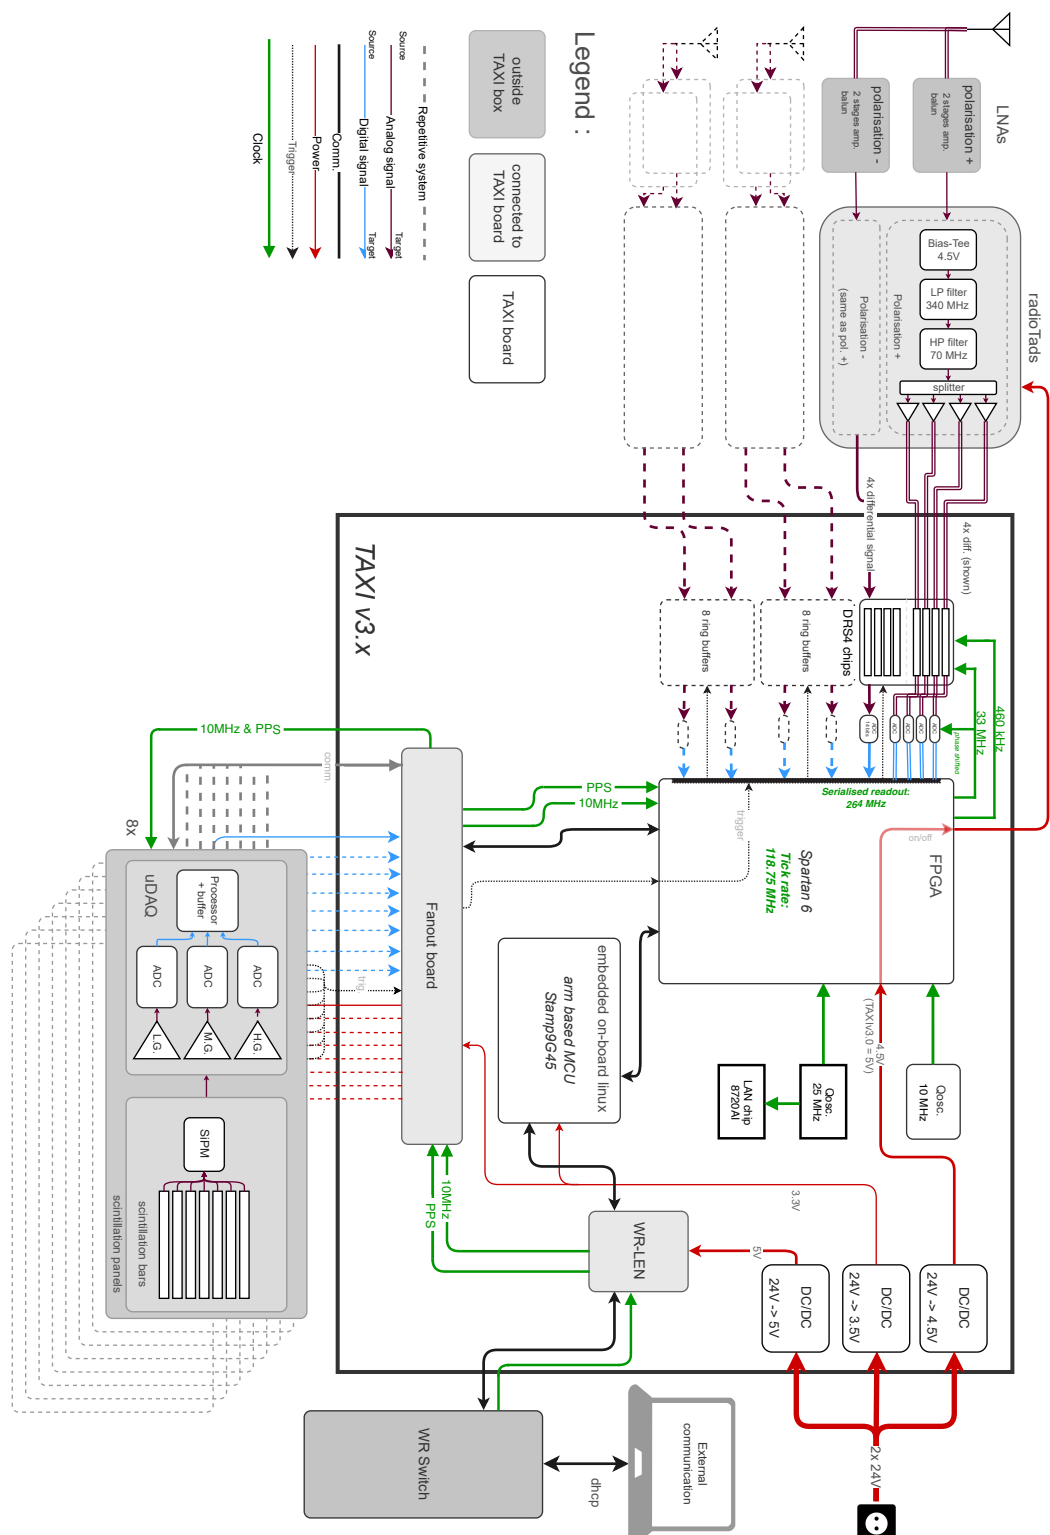


Figure A.1.: Overview over the most important electronic components on the current TAXI board (v3.2) for data acquisition. Edited from [1].

### A.3. TAXI in the IAP temperature chamber

While writing this thesis, A. Weindl conducted measurements of the TAXI radio response in a temperature chamber. As the results are not planned to be shown in any paper, and they fit the topic of this thesis perfectly, with explicit permission, they will be featured here. TAXI #13 was connected to a frequency generator running a frequency sweep from 0 to 500 MHz over a time period of 500 s. The temperature chamber was set to create environments in the range of  $-70$  to  $40$  °C, going up in steps of  $5$  °C. Before each measurement that is shown, the TAXI was shut down and, once "cold" (meaning around ambient temperature), rebooted. The DRS4 chips' radio response was then constantly read out over a period of 520 s, triggered by software at a frequency of around 23 Hz. For each such triggered event, the mean amplitude was stored to a file and then plotted for the entire 520 s duration. The results can be seen in Fig. A.2.

It is obvious how at certain input frequencies, the average gain is lower than at others. This is the hardware bandpass filter described in 4.2 working. Where the gain is low, the input frequency must have been between 400 and 500 MHz or below 50. Interesting to see is that with rising temperature, deviations from the expected gain start first at low gain values (already observable from  $10$  °C ambient temperature). As this is not a logarithmic plot, this is not an effect of distortion. The first explanation that comes to mind is as the graph only shows the mean amplitude, deviations from an amplitude of 0 automatically have a larger effect on the average than bin spikes occurring at the same frequency for higher-gain amplitudes. This could mislead into thinking that at low gain values, there are more bin spikes than at high ones. Should this not be a deception though, what might be an explanation for the just-described effect is the binary configuration of the data that is transmitted at low gains vs high gains. As it is very rare to observe bit flips from 0 to 1, but very common from 1 to 0 [1], having more spikes at low gains might be an effect of that - maybe low-gain signals actually carry a higher number of 1 bits, which would mean some kind of binary inversion, but could lead to more flips happening here.

What is also well visible is the effect of a cold start - most plots don't have spikes at the beginning, but only slowly start developing them mid-range through the entire plot's time interval. This is due to the electronics heating up and with that developing more spikes. From  $25$  °C ambient temperature, the spikes become so apparent that they genuinely corrupt the radio data. This is in line with the findings of [20]. Because of the line-connected data, it is unfortunately not possible to see which channels are more corrupted than others, but looking at the results of chapter 5.3, it is to be expected there are some.

The conclusion of this experiment is once more how important it is to keep the TAXI box cool. Additionally, the theory of the low-gain spikes could be followed up with additional research. Interesting follow-up experiments would be to set the DRS4 chips' baselines once to the lowest and once to the highest possible output level, disconnecting all inputs and simply again sampling the chips' outputs at ambient temperatures between  $0$  and  $40$  °C. It would be interesting to see whether there are actually more bit flips occurring (per channel, as they behave very differently) for high-gain or low-gain data. Looking at the binary characteristics, this could support R. Turcotte's thesis of 1-0 bit flips and further

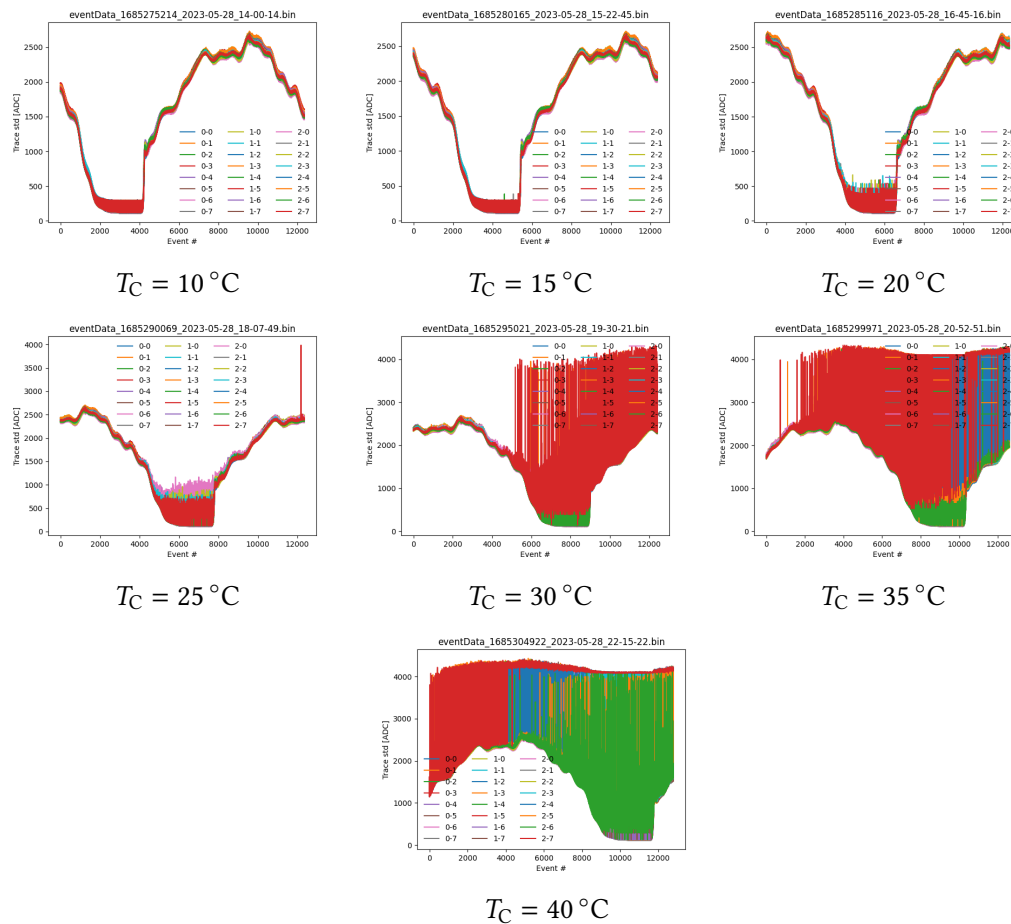


Figure A.2.: **Radio traces of frequency sweeps in a temperature chamber.**  $T_C$  describes the temperature in the chamber. The x axis covers a time range of 520 s, or approx. 12000 events. The y axis shows the average gain per event. Figures by A. Weindl.

ensure that the issue does lay at the serialised communication port between ADC and FPGA.

#### A.4. Additional plots: WR-LEN placement

While evaluating the temperature readings  $T_{SOC}$  and  $T_{DC}$  from the TAXIs' onboard temperature sensors, it was noticed that several readings showed ghosts, shifted values. These were cleaned up in Fig. 5.11, but for debugging reasons, the "messy" plot, tagged Fig. A.3, is shown in the following. The sensors that produced these ghosts are temp112 on both TAXIs (#09 and #22) and temp05 on TAXI #22.






Figure A.3.: **Measured Temperatures of all onboard sensors in the WR-LEN experiment.** It is well visible that there is ghosting, especially in the blue, green and red lines.




Figure A.4.: **Isolated view on the USB sensors in the makeshift FieldHubs.** A temperature difference of several degrees is visible.




Figure A.5.: **Isolated view on the onboard sensors temp05.** The onboard White Rabbit client (located above the sensor) shifts temperatures upward ( $2.5 \pm 0.2$ ) °C. Most false readings are removed for clarity.




Figure A.6.: **Isolated view on the onboard sensors temp112.** Because of its proximity to the DC converter, the baseline is higher. Most false readings are removed for clarity.

## A.5. IR images

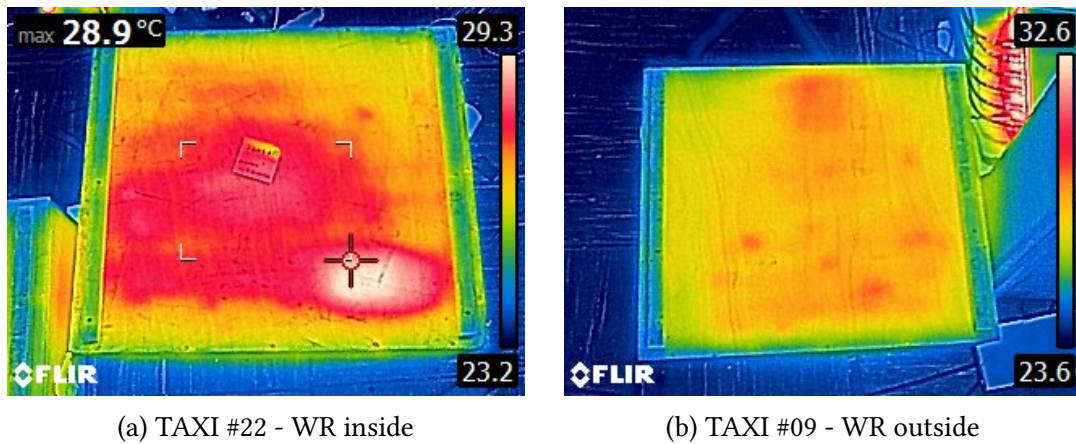


Figure A.7.: **IR images of the makeshift FieldHubs.** In image (a), the hottest spot is right above the DC unit on the PCB. In image (b), the colors are shifted because of disturbances by the other TAXI's ports in the top right corner.

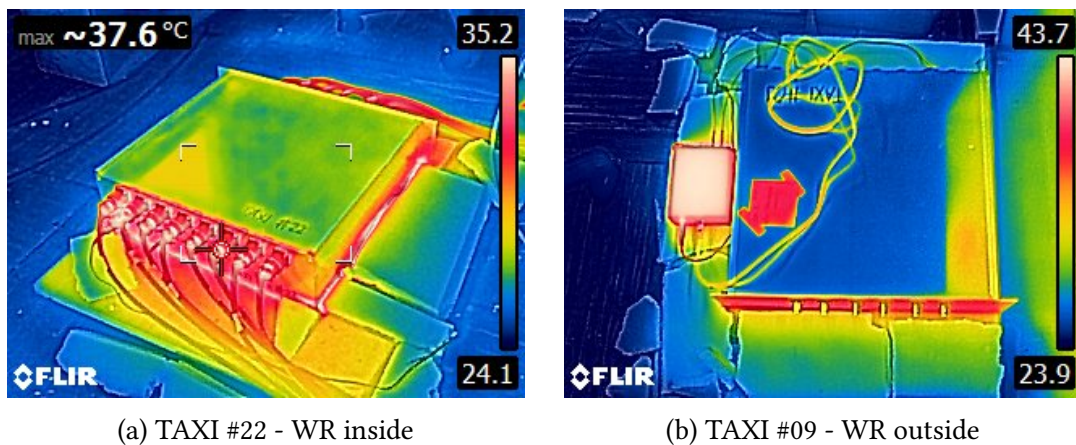


Figure A.8.: **IR image of the closed TAXIs.** Even though the top metal plate is warm, it does not fulfill the properties of a black body and thus only reflects the IR radiation coming from the ceiling. Still, in Fig. (b), the radiated heat by the WR-LEN is well visible.





Figure A.9.: **Close-Up IR images of the onboard White Rabbit Client.** Its hottest points reach way over 80 °C, reinforcing its role as a heat source.



Figure A.10.: **Close-Up IR image of the White Rabbit Switch located at the IceCube Lab.** This is a different device than the passive patchpanel in the FieldHub.



(a) Top

(b) Bottom

Figure A.11.: **IR images of the DRS4 evaluation board's PCB.** After few seconds of operation, the hot spots are already apparent: From left to right: DRS4 chips, FPGA and Power Supply.

# Bibliography

- [1] Roxanne Turcotte-Tardif. “Radio Measurements of Cosmic Rays at the South Pole”. PhD thesis. Karlsruhe, DE: Karlsruhe Institute of Technology, Oct. 2022. URL: <https://doi.org/10.5445/IR/1000160782>.
- [2] Luke O’C. Drury. “Origin of cosmic rays”. In: *Astroparticle Physics* 39-40 (Dec. 2012), pp. 52–60. DOI: 10.1016/j.astropartphys.2012.02.006. URL: <https://doi.org/10.1016%2Fj.astropartphys.2012.02.006>.
- [3] Pasquale Blasi. “Origin of very high- and ultra-high-energy cosmic rays”. In: *Comptes Rendus Physique* 15.4 (Apr. 2014), pp. 329–338. DOI: 10.1016/j.crhy.2014.02.008. URL: <https://doi.org/10.1016%5C%2Fj.crhy.2014.02.008>.
- [4] R. L. Workman et al. “Review of Particle Physics”. In: *PTEP* 2022 (2022), 083C01, pp.525–526. DOI: 10.1093/ptep/ptac097.
- [5] Veniamin Berezhinsky, Askhat Gazizov, and Svetlana Grigorieva. “On astrophysical solution to ultrahigh energy cosmic rays”. In: *Physical Review D* 74.4 (Aug. 2006). DOI: 10.1103/physrevd.74.043005. URL: <https://doi.org/10.1103%5C%2Fphysrevd.74.043005>.
- [6] D. J. Bird et al. (HiRes). “The Cosmic-Ray Energy Spectrum Observed by the Fly’s Eye”. In: *Astrophys. J.* 424 (1994), p. 491. DOI: 10.1086/173906.
- [7] Ralph Engel, Dieter Heck, and Tanguy Pierog. “Extensive Air Showers and Hadronic Interactions at High Energy”. In: *Annual Review of Nuclear and Particle Science* 61.1 (2011), pp. 467–489. DOI: 10.1146/annurev.nucl.012809.104544. eprint: <https://doi.org/10.1146/annurev.nucl.012809.104544>. URL: <https://doi.org/10.1146/annurev.nucl.012809.104544>.
- [8] Frank G. Schröder. “Radio detection of cosmic-ray air showers and high-energy neutrinos”. In: *Progress in Particle and Nuclear Physics* 93 (Mar. 2017), pp. 1–68. DOI: 10.1016/j.pnpnp.2016.12.002. URL: <https://doi.org/10.1016%5C%2Fj.pnpnp.2016.12.002>.
- [9] H. Allan, K. Neat, and J. Jones. “Mechanism of Radio Emission from Extensive Air Showers”. In: *Nature* 215 (1967), pp. 267–268. DOI: 10.1038. URL: <https://doi.org/10.1038/215267a0>.
- [10] G. A. Askar’yan. “Excess negative charge of an electron-photon shower and its coherent radio emission”. In: *J. Exptl. Theoret. Physics (USSR)* 41 (1961), pp. 616–618.
- [11] Tim Huege. “The Renaissance of Radio Detection of Cosmic Rays”. In: *Brazilian Journal of Physics* 44.5 (June 2014), pp. 520–529. DOI: 10.1007/s13538-014-0226-6. URL: <https://doi.org/10.1007%5C%2Fs13538-014-0226-6>.

- [12] “The Pierre Auger Cosmic Ray Observatory”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 798 (2015), pp. 172–213. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2015.06.058>. URL: <https://www.sciencedirect.com/science/article/pii/S0168900215008086>.
- [13] Claus Grupen. “Secondary Cosmic Rays”. In: *Astroparticle Physics*. Cham: Springer International Publishing, 2020, pp. 235–291. ISBN: 978-3-030-27339-2. DOI: 10.1007/978-3-030-27339-2\_7. URL: [https://doi.org/10.1007/978-3-030-27339-2\\_7](https://doi.org/10.1007/978-3-030-27339-2_7).
- [14] Sam De Ridder. “Sensitivity of IceCube cosmic ray measurements to the hadronic interaction models”. PhD thesis. Ghent University, 2019, XI, 210. URL: <https://biblio.ugent.be/publication/8609789>.
- [15] Martin Rongen. “Kalibration des IceCube Neutrino-Observatoriums”. PhD thesis. 2019, p. 2019. DOI: 10.18154/RWTH-2019-09941. URL: <http://publications.rwth-aachen.de/record/771097>.
- [16] *United States Antarctic Activities: Comms Forms, Attachment A*. [https://www.nsf.gov/geo/opp/antarct/treaty/pdf/plans0708/attach\\_08.pdf](https://www.nsf.gov/geo/opp/antarct/treaty/pdf/plans0708/attach_08.pdf). Accessed: 2023-06-27.
- [17] Thomas Huber. “IceScint : A Scintillation Detector Array for the IceCube Ice-Top Enhancement = IceScint: Ein Szintillationsdetektor-Feld für die Oberflächen-Erweiterung von IceCube”. 51.13.04; LK 01. PhD thesis. Karlsruher Institut für Technologie (KIT), 2021. 234 pp. DOI: 10.5445/IR/1000131545.
- [18] T. Huber et al. *IC WOs surface enhancement maintenance instructions Season 2022/23*. PDF presentation, IceCube internal.
- [19] *WR-LEN product page*. [https://safran-navigation-timing.com/product/white-rabbit-len/?model\\_interest\\_\\_c=WR+LEN&product\\_interest\\_\\_\\_single\\_select=Resilient+Timing](https://safran-navigation-timing.com/product/white-rabbit-len/?model_interest__c=WR+LEN&product_interest___single_select=Resilient+Timing). Accessed: 2023-09-14.
- [20] Noah Goehlke. “An IceCube Surface Array Enhancement station for the deployment at Telescope Array”. 51.13.04; LK 01. MA thesis. Karlsruher Institut für Technologie (KIT), 2022. 118 pp. DOI: 10.5445/IR/1000153878.
- [21] A. Weindl. Internal communication, 14 August. 2023.
- [22] M. G. Aartsen et al. “IceCube-Gen2: the window to the extreme Universe”. In: *Journal of Physics G: Nuclear and Particle Physics* 48.6 (Apr. 2021), p. 060501. DOI: 10.1088/1361-6471/abbd48. URL: <https://doi.org/10.1088%5C%2F1361-6471%5C%2Fabbd48>.
- [23] E. de Lera Acedo et al. “Evolution of SKALA (SKALA-2), the log-periodic array antenna for the SKA-low instrument”. In: *2015 International Conference on Electromagnetics in Advanced Applications (ICEAA)*. 2015, pp. 839–843. DOI: 10.1109/ICEAA.2015.7297231.
- [24] Roxanne Turcotte-Tardif. Internal communication, 4 July. 2023.

- 
- [25] Holger Heuermann. *Hochfrequenztechnik : Komponenten für High-Speed- und Hochfrequenzschaltungen*. 3., verbesserte und erweiterte Auflage. SpringerLink. Wiesbaden: Springer Vieweg, 2018, pp. 53–54. ISBN: 978-3-658-23198-9. URL: <https://doi.org/10.1007/978-3-658-23198-9>.
- [26] Max Renschler. “A Prototype Radio Detector for the IceCube Surface Enhancement”. 51.03.04; LK 01. PhD thesis. Karlsruher Institut für Technologie (KIT), 2020. 169 pp. DOI: 10.5445/IR/1000104529.
- [27] *Data Sheet: Wideband, Low Noise, Low Distortion Fully Differential Amplifier (THS4508)*. Accessed: 2023-07-04. Texas Instruments. URL: [https://www.ti.com/lit/ds/symlink/th4508.pdf?ts=1688492488339&ref\\_url=https%253A%252F%252Fwww.ti.com%252Fproduct%252FTHS4508](https://www.ti.com/lit/ds/symlink/th4508.pdf?ts=1688492488339&ref_url=https%253A%252F%252Fwww.ti.com%252Fproduct%252FTHS4508).
- [28] Marko Kossatz. Phone and text conversations, May-September. 2023.
- [29] Mark Owen. *Practical Signal Processing*. Practical Signal Processing. Cambridge University Press, 2007. ISBN: 978-0-521-85478-8. URL: <https://books.google.de/books?id=lx-tqq-MkK0C>.
- [30] *DRS Chip Home Page*. <https://www.psi.ch/en/drs>. Accessed: 2023-06-28.
- [31] *Data Sheet: 9 Channel, 5 GPS Switched Capacitor Array (DRS4)*. Accessed: 2023-07-05. PSI - Paul Scherrer Institut. URL: <https://www.psi.ch/en/drs/documentation>.
- [32] Marie Johanna Oehler. “The Prototype Station for the IceCube Surface Array Enhancement”. 51.13.04; LK 01. PhD thesis. Karlsruher Institut für Technologie (KIT), 2022. 128 pp. DOI: 10.5445/IR/1000142813.
- [33] Roxanne Turcotte-Tardif. *South Pole Deployment 2023*. PDF presentation 2023, IceCube internal.
- [34] Hrvoje Dujmovic. *TAXI thermal images*. PDF presentation 2022, IceCube internal.





# Acknowledgement

Within the scope of this thesis, over a period of 5 months, several experiments were run, some with many frustrating iterations for getting the perfect results. For never letting me down and always keeping up the good mood I want to thank Megha, who constantly cared for what I did and helped me read out those stubborn USB sensors. Also I want to thank Tom for all his ideas, of which I could only process half because there were so many, and Philippe and Jan for all the silly jokes in the office.

Big thanks also to all the colleagues who tirelessly answered all the very specific questions I had for them, even if I had quite the bad timing. Special thanks belong to Marko Kossatz for the hour-long phonecalls at a time, and to Andreas Weindl, who with his knowledge saved me much research and frustration. Thanks also for all your explanations, the plots and for lending me the necessary cables and screwdrivers, even if it got in the way of your own experiments. Also thanks to Frank, who helped me immensely at figuring out what my thesis would actually be about, and to Prof. Ralph Engel, who gave some very warm and constructive feedback to the drafts I presented to him.

I want to thank Andreas Haungs for introducing me to IceCube and welcoming me so warmly from my very first day at the institute up till today, each time I see you, there's a good mood around. Thanks to Jelena, Federico, Shefali, Mark, Megha and all the other group members who were there and recruited me enthusiastically that one day at the work group presentation at Campus South, you had my sympathy immediately, and even though I didn't join your excursions often, I very much enjoy your spirit. Ah and thanks for occasionally WABing with me.

Thanks to Yan for teaching me how to fly and letting me sample his creaking floorboards. You're amazing!

And I want to thank my two TAXIs for holding on and working properly weeks at a time during the many iterations of my WR-LEN measurements. If there is one thing I'll always remember about working with hardware, it's that this is a real blessing.

