Friedrich-Alexander-Universität
Faculty of Sciences

Doctoral thesis

2026

Jakob Henrichs-Schuhegger

Measurement of Ultra-High-Energy Cosmic Rays with RNO-G



Measurement of Ultra-High-Energy Cosmic Rays with RNO-G

Messung von ultrahochenergetischer kosmischer Strahlung mit RNO-G

Der Naturwissenschaftlichen Fakultät

der Friedrich-Alexander-Universität Erlangen-Nürnberg

zur Erlangung des Doktorgrades Dr. rer. nat.

vorgelegt von

Jakob Henrichs-Schuhegger

aus Monschau

Als Dissertation genehmigt

von der Naturwissenschaftlichen Fakultät

der Friedrich-Alexander-Universität Erlangen-Nürnberg

Tag der mündlichen Prüfung: ??.??.2026

Gutachter/in:

Prof. Dr. Anna Nelles

Prof. Dr. Frank G. Schröder

Abstract

A central unanswered question in astroparticle physics is the origin of ultra-high-energy cosmic rays (UHECRs). Neutrinos, due to their unique properties, serve as valuable messengers that help to investigate and possibly answer this question. The radio detection of neutrinos has emerged as a promising technique for observing ultra-high-energy neutrinos (UHENs), as it enables cost-effective monitoring of large volumes of naturally occurring target materials, such as ice. The Radio Neutrino Observatory in Greenland (RNO-G) is an in-ice radio detector currently under construction on the Greenlandic ice sheet, with the primary goal to detect neutrinos with energies exceeding 10 PeV. Upon completion, RNO-G will comprise 35 independently operating stations, collectively observing approximately 100 km^3 of ice, and will be capable of probing the most optimistic model predictions for UHEN. To date, eight stations have been deployed and are collecting data. A critical part of early detector commissioning is the study of detector characteristics and potential backgrounds, for which UHECRs play a crucial role. When they interact within the atmosphere, they generate a particle cascade that produces radio emission detectable by RNO-G antennas. Additionally, muons produced in the air shower can deposit energy in the ice, resulting in signatures that are effectively indistinguishable from UHEN interactions. Therefore, UHECRs represent a critical background that must be characterized and, whenever possible, vetoed. On the other hand, the radio signals from UHECRs occur at the nanosecond timescale, similar to those from UHENs, providing an opportunity to test the detector under realistic conditions.

In this thesis, we present an analysis to identify the first UHECR air showers using the shallow component of RNO-G. We analyze data from three stations collected in 2022 and 2023 and identify 85 UHECR candidate events, for which we reconstruct their directions, polarizations and energies. Both the number of UHECR candidates and their reconstructed properties are consistent with the expectations from simulations. Furthermore, we investigate the influence of the amplification and digitization response on the signal shape, revealing that this response can significantly alter the correlation score. This result underlines the importance of precise hardware characterization, which is achieved through a full-chain time domain measurement. Additionally, we examine modifications to the trigger path and demonstrate that those changes enhance the trigger performance. Moreover, we thoroughly study various systematic effects and identify the differences between refractive index models at shallow depth as the dominant source of systematic uncertainty. This thesis further describes the database that enables time-dependent detector descriptions, investigates unusual station behavior during the period relevant to this analysis, discusses in detail the refractive index at shallow depths, and introduces a method to measure the refractive index using a ground-penetrating radar. All these efforts contribute towards bringing RNO-G a step closer to routine operation.

Overall, this thesis represents an essential first step towards validating the detector and enabling high-fidelity UHEN detection in the future.

Zusammenfassung

Eine zentrale, bislang unbeantwortete Frage in der Astroteilchenphysik ist der Ursprung ultrahochenergetischer kosmischer Strahlung (UHECR). Neutrinos dienen aufgrund ihrer einzigartigen Eigenschaften als wertvolle Botschafter, die dabei helfen, diese Frage zu untersuchen und möglicherweise zu beantworten. Die Radiodetektion von Neutrinos hat sich als vielversprechende Technik zur Beobachtung ultrahochenergetischer Neutrinos (UHENs) herausgestellt, da sie eine kostengünstige Überwachung großer Volumina natürlich vorkommender Target-Materialien wie beispielsweise Eis ermöglicht. Das Radio Neutrino Observatory in Greenland (RNO-G) ist ein Radio-Detektor im Eis, der derzeit auf dem grönländischen Eisschild errichtet wird und dessen primäres Ziel darin besteht, Neutrinos mit Energien von über 10 PeV nachzuweisen. Nach seiner Fertigstellung wird RNO-G aus 35 unabhängig voneinander arbeitenden Stationen bestehen, die gemeinsam etwa 100 km^3 Eis beobachten. RNO-G wird in der Lage sein, die optimistischsten Modellvorhersagen für UHEN zu überprüfen. Bislang wurden acht Stationen installiert, die bereits Daten sammeln. Ein entscheidender Teil der frühen Inbetriebnahme des Detektors ist die Untersuchung der Detektoreigenschaften und potenzieller Hintergründe, wobei UHECRs eine entscheidende Rolle spielen. Wenn UHECRs in der Atmosphäre Wechselwirkungen eingehen, erzeugen sie eine Teilchenkaskade, die eine von den RNO-G-Antennen nachweisbare Radioemission erzeugt. Zudem können in dem Luftschauer erzeugte Myonen Energie im Eis deponieren, was zu Signaturen führt, die von UHEN-Wechselwirkungen praktisch nicht zu unterscheiden sind. Daher stellen UHECRs einen kritischen Hintergrund dar, der charakterisiert und, wann immer möglich, ausgeschlossen werden muss. Andererseits treten die Radiosignale von UHECRs im Nanosekundenbereich auf, ähnlich wie die von UHENs, was die Möglichkeit bietet, den Detektor unter realistischen Bedingungen zu testen.

In dieser Arbeit stellen wir eine Analyse vor, mit der die ersten UHECR-Luftschauer mithilfe der Antennen von RNO-G kurz unterhalb der Oberfläche identifiziert werden können. Wir analysieren Daten von drei Stationen, die in den Jahren 2022 und 2023 erfasst wurden, und identifizieren 85 potenzielle UHECR-Ereignisse, deren Richtungen, Polarisationen und Energien wir rekonstruieren. Sowohl die Anzahl der potenziellen UHECR-Ereignisse als auch deren rekonstruierte Eigenschaften stimmen mit den Erwartungen aus Simulationen überein. Darüber hinaus untersuchen wir den Einfluss von Verstärkung und Digitalisierung auf die Signalform und zeigen, dass sie den Korrelationswert erheblich verändern können. Dieses Ergebnis unterstreicht die Bedeutung einer präzisen Hardware-Charakterisierung, die durch Messung der gesamten Hardware-Kette im Zeitbereich erreicht wird. Zusätzlich untersuchen wir Modifikationen am Triggerpfad und zeigen, dass diese die Triggerleistung verbessern. Darüber hinaus untersuchen wir verschiedene systematische Effekte und identifizieren die Unterschiede zwischen Brechungsindexmodellen in geringer Tiefe als die dominierende Quelle der systematischen Unsicherheit. Diese Arbeit beschreibt ferner die Datenbank, die zeitabhängige Detektorbeschreibungen ermöglicht, untersucht ungewöhnliches Stationsverhalten während des für diese Analyse relevanten Zeitraums, erörtert detail-

liert den Brechungsindex in geringer Tiefe und stellt eine Methode zur Messung des Brechungsindexes mittels eines Bodenradars vor. All diese Bemühungen tragen dazu bei, RNO-G dem Routinebetrieb einen Schritt näherzubringen. Insgesamt ist diese Arbeit ein wesentlicher erster Schritt zur Validierung des Detektors und zur Ermöglichung einer hochauflösenden UHEN-Detektion in Zukunft.

Table of Contents

Abstract v

Zusammenfassung vii

1 Introduction 1

2 Astroparticle Physics 3

2.1 Acceleration and transport of cosmic rays 4

2.2 Detection of cosmic rays 7

2.2.1 Direct detection 7

2.2.2 Indirect detection 8

2.2.3 Extensive air showers 9

2.3 Cosmic rays at the highest energies 12

2.4 Neutrino astronomy 14

2.4.1 Detection of high-energy neutrinos 15

2.4.2 Detection of ultra-high-energy neutrinos 17

3 Radio Detection 21

3.1 Radio emission from particle cascades 21

3.1.1 Geomagnetic emission 21

3.1.2 Askaryan emission 23

3.1.3 Characteristics of the radio signal 24

3.2 Radio detection of UHECR 26

3.3 Radio detection of UHEN 28

4 Radio Neutrino Observatory in Greenland 31

4.1 Radio antennas 33

4.1.1 Log-Periodic Dipole Array (LPDA) 35

4.1.2 Vpol 36

4.1.3 Hpol 36

4.2 Amplification and digitization 37

4.2.1 Deep components 37

4.2.2 Shallow components . . . 38
4.3 Trigger . . . 39
4.4 Backgrounds . . . 41

5 Shallow RNO-G Hardware and Trigger **43**
5.1 Shallow amplification and digitization response . . . 43
5.1.1 Influence on the shallow UHECR analysis . . . 44
5.1.2 Hardware response descriptions . . . 46
5.1.3 Influence of gain and group delay . . . 51
5.2 RNO-G hardware and calibration database . . . 52
5.2.1 Database structure . . . 52
5.2.2 Web interface . . . 60
5.2.3 Access and querying . . . 62
5.2.4 Status and future plans . . . 63
5.3 RADIANT trigger . . . 65
5.3.1 External vs board noise . . . 65
5.3.2 Trigger efficiency measurements . . . 69
5.3.3 Trigger modifications . . . 73
5.3.4 Trigger threshold in UHECR simulations . . . 75
5.4 Unusual station behavior . . . 80
5.5 Conclusions . . . 83

6 Refractive Index of Ice **87**
6.1 Ice models . . . 87
6.2 Shallow refractive index . . . 89
6.3 Effective antenna description . . . 90
6.4 Influence on simulations . . . 92
6.5 Ground-penetrating radar . . . 93
6.5.1 Dataset . . . 96
6.5.2 Stationary measurement . . . 98

6.5.3 Dynamic measurement . . . 100
6.5.4 Advantages and next steps . . . 103
6.6 Conclusions . . . 104

7 Cosmic Ray Simulations **107**
7.1 Simulation setup . . . 107
7.2 Antenna and hardware description . . . 109
7.3 Trigger description . . . 110
7.4 Noise description . . . 112
7.5 Effective area and expected rate . . . 114
7.6 Conclusions . . . 116

8 Cosmic Ray Analysis **119**
8.1 Analysis strategy . . . 119
8.2 Cosmic-ray templates . . . 120
8.3 Dataset . . . 126
8.4 Systematic uncertainties . . . 129
8.4.1 Refractive index . . . 130
8.4.2 Antenna response . . . 131
8.4.3 Timing calibration . . . 132
8.4.4 Amplification and digitization response . . . 135
8.4.5 Additional investigated systematic effects . . . 137
8.5 Analysis of main dataset . . . 146
8.5.1 Time cluster cut . . . 147
8.5.2 Half trace RMS cut . . . 149
8.5.3 Template cut . . . 151
8.5.4 Event selection . . . 152
8.6 Characterization of observed events . . . 158
8.7 Background contamination . . . 160
8.7.1 Methods for background estimate . . . 160

8.7.2 Potential background sources 163
8.7.3 Source-independent cross-checks 170
8.8 Analysis of further stations 173
8.9 Conclusions 176

9 Conclusions **179**

Appendix **183**

A Supplementary Figures **185**

List of Figures **185**

List of Tables **197**

List of Terms **201**

References **205**

Acknowledgements **217**

Erklärung zur Nutzung von (generativer) KI **219**

Parts of this thesis have been published in:

S. Agarwal et al. (RNO-G), "*Validating the performance of the Radio Neutrino Observatory in Greenland using cosmic-ray air showers*", Submitted to EPJ C (2025)

1 Introduction

A major open question in astroparticle physics concerns identifying astrophysical processes capable of accelerating cosmic rays to energies of up to 10^{20} eV. Neutrinos serve as unique messengers for uncovering these sources, as they are expected to be produced when cosmic-ray particles interact with ambient matter in the vicinity of the acceleration site. Due to their properties, neutrinos are effectively not absorbed by intervening matter or deflected by magnetic fields, and thus travel unimpeded across cosmic distances, pointing back to their sources. However, detecting these neutrinos brings unique challenges, as their small interaction cross-section and low flux require large-scale detectors.

For neutrinos with energies below 10 PeV, neutrino astronomy is well established. Notably, in recent years, the IceCube Neutrino Observatory has identified a blazar [1], a nearby active galaxy [2] and the Milky Way Galaxy [3] as sources of astrophysical neutrinos, and has measured the diffuse neutrino flux up to energies of 10 PeV [4]. In the ultra-high-energy regime ($> 10^{16}$ eV), no other than the exceptional neutrino detected by KM3NeT with an estimated energy of 220 PeV [5] has been observed. The most stringent constraints in this energy range are provided by the IceCube Neutrino Observatory and the Pierre Auger Observatory [6, 7], whose results are in mild tension with the KM3NeT event.

Due to the low flux of ultra-high-energy neutrinos (UHENs), detectors with even larger detection volumes than IceCube or KM3NeT are necessary to detect them. A cost-effective approach is to construct neutrino telescopes that measure the radio emission produced through the *Askaryan effect* [8] when an UHEN interacts in ice. This method offers three primary advantages. First, the long attenuation length of radio waves in ice (e.g. $\mathcal{O}$(1 km) in Greenlandic ice [9]) allows sparse instrumentation, enabling the coverage of a large detection volume with relatively few antennas. In contrast, optical detection modules, as used by neutrino telescopes such as IceCube or KM3NeT, must be deployed closer together due to the short attenuation and scattering lengths of optical light in water or ice ($\mathcal{O}$(100 m) [10, 11]). Second, radio antennas are relatively inexpensive to manufacture. Third, optical detectors need to be deployed in the ice at great depth to achieve sufficient optical properties, which requires drilling deep holes. In comparison, radio antennas can be deployed at shallower depths, making deployment less expensive. Consequently, the radio detection technique provides a cost-effective approach to probe the ultra-high-energy regime of neutrinos.

Building on the experience from previous and current in-ice radio neutrino telescopes such as ARIANNA and ARA, which have successfully demonstrated the radio detection technique [12, 13], the Radio Neutrino Observatory in Greenland (RNO-G) is currently under construction on the Greenlandic ice sheet. Upon completion, RNO-G will observe around 100 km^3 of ice and cover an area of around 50 km^2, making it the largest in-ice neutrino detector to date. The primary objective of RNO-G is to detect neutrinos with energies above 10 PeV, and the experiment will be sensitive to the most optimistic model predictions for UHEN [14]. Additionally, RNO-G serves as a pathfinder for an even larger in-ice radio neutrino telescope, planned as a part of IceCube-Gen2 [15].

Ultra-high-energy cosmic rays (UHECRs) are an important background for in-ice neutrino telescopes. When a UHECR particle interacts within the atmosphere, it produces radio emission primarily through the *geomagnetic effect* [16]. This emission can refract into the ice and reach the detector, thereby becoming a background for the UHEN signal [17]. Additionally, muons produced in the extensive air shower (EAS) initiated by the UHECR interaction can penetrate the ice and deposit energy, producing radio signals that are effectively indistinguishable from those generated by neutrinos [18]. Furthermore, an incompletely developed EAS can impact the ice, initiating a secondary particle cascade that emits radio signals, as demonstrated in Ref. [19–21] and recently confirmed experimentally [22]. Therefore, it is crucial to understand the signals from UHECRs and, whenever possible, to veto them. On the other hand, the radio signals from UHECR air showers offer a valuable opportunity: they have a similar characteristic to that of a UHEN signal, as in both cases the produced electric field occurs on the nanosecond timescale. Therefore, UHECR signals can serve as a unique "test beam" to evaluate the detector performance under realistic conditions, particularly regarding the response to short electric field pulses. Furthermore, with the deployed upward-facing antennas, RNO-G can study the properties of UHECR air showers and, due to its high elevation of around 3200 m, detect impacting air showers. Of particular interest is the potential to observe the in-ice signature of these events in coincidence with the radio emission from the air shower. In summary, it is crucial to understand signals from UHECR air showers for future UHEN searches, and they play a key role in early detector commissioning with regard to studying potential backgrounds and the detector characteristics.

The primary scientific goal of this thesis was to develop the first analysis to identify signals from UHECR air showers with RNO-G. For this purpose, we use early data collected during the 2022 and 2023 data-taking periods with the upward-facing antennas deployed at shallow depths. Furthermore, a comprehensive understanding of the ice surrounding the antennas and how the hardware components affect the signal shape is essential for this analysis. Both aspects will be studied in detail in this thesis. We will demonstrate that the number of identified UHECR candidate events and their reconstructed properties agree with expectations within the carefully studied systematic uncertainties of the analysis.

In the first three chapters, we provide an overview of astroparticle physics at the highest energies (see Chapter 2), introduce the radio detection technique for UHECR and UHEN (see Chapter 3), and describe RNO-G (see Chapter 4). Subsequently, we address the two factors with the greatest influence on the analysis. First, we investigate the impact of the hardware response on the analysis, outline its implementation for future analyses, and study trigger thresholds to ensure simulations with realistic triggers (see Chapter 5). Second, we assess the uncertainty in the ice refractive index at shallow depths and its influence on the analysis, and present a new measurement strategy to determine the refractive index, including a proof-of-principle analysis (see Chapter 6). Afterwards, we explain the simulation framework used to generate realistic UHECR simulations for the analysis (see Chapter 7). Then, we describe the analysis strategy to identify signals from UHECR air showers, discuss systematic uncertainties and potential background sources, and present the final results, including UHECR candidate events and their reconstructed properties (see Chapter 8). Finally, we summarize the results of this thesis and provide an outlook on potential improvements for future analyses (see Chapter 9).

2 Astroparticle Physics

More than 100 years ago, Victor Hess initiated the field of cosmic-ray research by identifying a "penetrating radiation" that increased with altitude during balloon flights [23]. Today, this "penetrating radiation" is recognised as non-thermal particles that travel through the universe [24]. Since the discovery by Hess, the field of cosmic-ray research has advanced considerably. Numerous experiments now measure cosmic rays with high precision across an energy range from less than 1 GeV to more than 10^{20} eV. The resulting energy spectrum is presented in Figure 2.1. Most particles contributing to this spectrum are charged, such as protons or heavier nuclei, although γ rays and neutrinos are sometimes included in the definition of cosmic rays. In this work, the term "cosmic ray" refers exclusively to charged particles, while neutrinos and γ rays are discussed separately, as is often done in literature. The energy spectrum decreases steeply with increasing energy; therefore, it is common to multiply the flux by the energy to enhance the visibility of spectral features. The energy spectrum can be described by multiple power laws $E^{-\alpha}$ with different spectral indices α in different energy regions. Four different energy regions can be identified, which are separated by spectral features. The first, known as the *knee*, occurs at approximately 3 PeV, where the spectral index shifts from about 2.7 to 3. Then, the *second knee* appears near 100 PeV, with the spectral index changing to approximately 3.3, and the *ankle* marks a transition to a spectral index of about 2.5 [24]. An increase in the spectral index is described as the spectrum *softening*, while a decrease is termed *hardening*. Despite significant progress in cosmic-ray research over the last 100 years, fundamental questions remain unsolved, such as the nature of the sources that accelerate cosmic rays to these extremely high energies. Furthermore, investigating cosmic rays at the highest energies provides a unique opportunity to explore particle physics beyond the energy regimes accessible on Earth [25].

There are two different approaches to detecting cosmic rays: direct and indirect measurements. Direct measurements involve detecting the cosmic-ray particle itself with a particle detector, whereas indirect measurements observe secondary particles produced when cosmic rays interact with a natural medium, such as the atmosphere. Due to the steep decline in cosmic-ray flux with increasing energy, direct detection is feasible only at lower energies, below approximately 100 TeV. Consequently, direct detection experiments are effectively limited to Galactic cosmic rays, as it is currently assumed that cosmic rays with energies up to a few PeV originate predominantly from within the Galaxy [24]. This conclusion is supported by γ-ray flux measurements from sources such as the Small Magellanic Cloud, which would be higher if most cosmic rays with energies below a few PeV were of extragalactic origin [26]. Cosmic rays with energies above a few EeV are considered to be of extragalactic origin. At these energies, cosmic rays are not expected to diffuse from sources within the Galaxy to Earth, and their arrival directions should approximately point towards their sources. Therefore, an excess of cosmic rays would be expected along the Galactic plane if they were of Galactic origin; the absence of such an excess supports the extragalactic hypothesis [24].

In this chapter, we first examine possible acceleration and transportation mechanisms for cosmic rays (see Section 2.1), then discuss experimental techniques (see Section 2.2), and

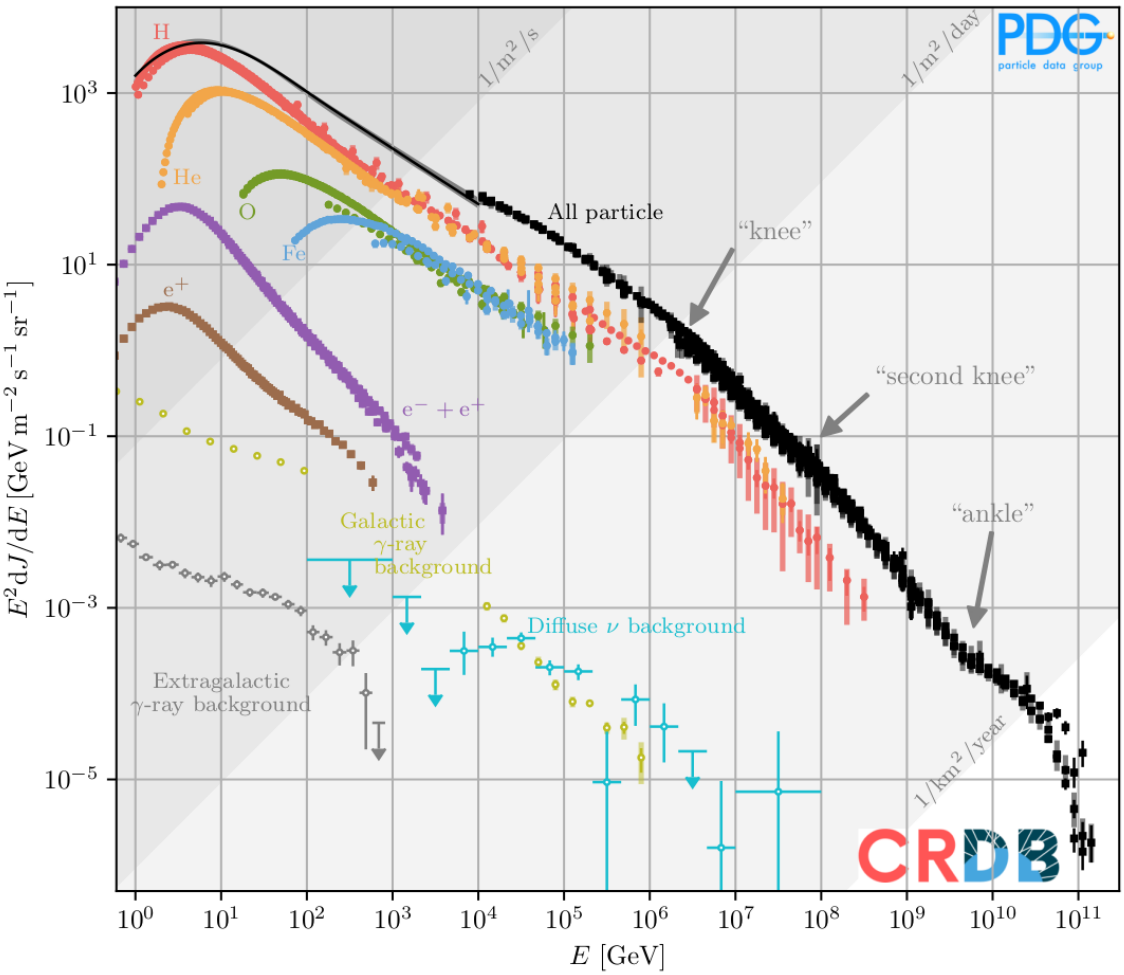


Figure 2.1: The energy spectra of cosmic rays, γ rays, and neutrinos. The flux dJ/dE is multiplied by the energy squared to enhance the visibility of features in the steeply falling spectrum. The plot is taken from [24], and the sources for the various datasets can be found in the original publication.

finally present findings and open questions regarding cosmic rays at the highest energies (see Section 2.3). Lastly, we introduce the role of neutrinos in addressing these unresolved questions, along with methods for detecting these neutrinos (see Section 2.4).

2.1 Acceleration and transport of cosmic rays

The precise mechanism responsible for accelerating cosmic-ray particles at their sources remains unresolved, although several hypotheses have been advanced. Already in 1949, Fermi proposed that particles could be accelerated by repeatedly scattering off moving magnetic fields [27]. Each time the particle scatters[1], it would gain a small amount of energy ΔE proportional to its energy before the scatter E:

$$\Delta E = \alpha E \tag{2.1}$$

[1] The definition of a scattering event depends on the specific process under consideration and will be addressed in detail in subsequent scenarios.

After undergoing n scatterings, the particle would have the following energy E_n:

$$E_n = E_0(1+\alpha)^n, \tag{2.2}$$

where E_0 denotes the initial energy at the beginning of the acceleration. If P_{esc} represents the probability that the particle escapes the acceleration process per encounter, the number of particles produced with energy equal to or greater than E_n is:

$$N(\geq E_n) \propto \sum_{m=n}^{\infty}(1-P_{\text{esc}})^m = \frac{(1-P_{\text{esc}})^n}{P_{\text{esc}}} \tag{2.3}$$

An expression for n can be derived from equation 2.2, leading to the following result.

$$N(>E) \propto \frac{1}{P_{\text{esc}}}\left(\frac{E}{E_0}\right)^{-\gamma}, \ \gamma = \frac{\ln(\frac{1}{1-P_{\text{esc}}})}{\ln(1+\alpha)}, \tag{2.4}$$

This process produces a power law energy spectrum, which is necessary to explain the observed cosmic-ray spectrum. A comprehensive derivation and discussion of this equation is provided in Ref. [25]. The mechanism described above, referred to as *Fermi mechanism*, can be applied to two scenarios[2]. In the first scenario, particles are accelerated at a shock front, a process known as *diffuse shock acceleration*. In this context, particles gain energy by repeatedly crossing the shock front and scattering off turbulent magnetic fields on both sides [28]. Because the media on either side of the shock front move at different velocities, a particle moving from upstream to downstream and back gains energy. Within the context of the derivation of the above equations, scattering in diffuse shock acceleration refers to particles crossing the shock front in both directions. The energy gained per scattering in this scenario is [25]:

$$\Delta E \sim \frac{4}{3}\frac{v}{c}E, \tag{2.5}$$

where c denotes the speed of light and v the velocity of the downstream (shocked) medium relative to the upstream (unshocked) medium. Because the energy gain is linear in velocity, this process is also called *first-order Fermi acceleration*. Further details regarding this mechanism are available in Ref. [29, 30]. The acceleration mechanism described above applies to non-relativistic shocks. For relativistic shocks, an alternative model known as *espresso acceleration* has been proposed, in which a pre-accelerated cosmic ray receives a "one-shot" boost, increasing its energy to above 10^{20} eV [31, 32].

In the second scenario to which the Fermi mechanism can be applied, the accelerated particle is scattered by magnetic fields within a moving plasma cloud [27]. In this scenario, scattering, within the context of the derivation above, refers to a particle entering and exiting the moving cloud. The particle may experience either approaching scatterings (cloud and particle move towards each other before scattering) or overtaking scatterings (cloud and particle move in the same direction before scattering). Depending on the scattering angle and kind of scattering, the particle may gain or lose energy during these encounters. Stochastically, it is more likely that the particle gains energy when scattering

[2] It is important to note that the acceleration process must be collisionless, meaning that no collisions with other particles occur. Otherwise, energy losses during the collisions would render the acceleration inefficient.

off the moving cloud, and thus, on average, the gained particle energy per scattering is [25]:

$$\Delta E \sim \frac{4}{3}\left(\frac{v}{c}\right)^2 E, \tag{2.6}$$

where v is the velocity of the cloud. This process describes interactions with a non-relativistic cloud, and because the energy gain depends on the velocity squared, this mechanism is referred to as *second-order Fermi acceleration*, and it is also known as *stochastic acceleration*. Due to the quadratic dependence on v/c, which is less than one, this process is less efficient than the first-order Fermi acceleration and requires more scatterings to achieve similar particle energies.

Magnetic-field-line reconnection has also been proposed as a mechanism for cosmic-ray acceleration. In this scenario, magnetic field lines within a plasma rearrange, converting magnetic energy into plasma energy, which is used to accelerate charged particles [33]. Lastly, we mention the acceleration of charged particles by electric fields, such as those associated with pulsars in supernova remnants, as another possible acceleration mechanism [34, 35].

The maximum energy to which a cosmic-ray particle can be accelerated at a source is determined by the *Hillas criterion*, which requires that the source size exceeds the Larmor radius[3] of the particle to confine it within the acceleration region during the acceleration process. Consequently, the maximum achievable energy is proportional to the size and the magnetic field strength of the source, as well as the charge Z of the accelerated particle [36]. Therefore, a possible source can accelerate heavier particles (with a larger charge Z) to higher energies than lighter particles.

Lower-energy cosmic rays produced within the Galaxy diffuse towards Earth via a random walk, interacting with the turbulent magnetic field and experiencing deflections along the way. As a result, the arrival directions of the Galactic cosmic rays are smeared, and only small anisotropies are observed in their arrival direction. For very high-energy cosmic rays, the Larmor radius becomes sufficiently large that typical galactic magnetic fields can no longer confine them. For protons, this occurs at energies above $\sim 10^{18}$ EeV. Nevertheless, their arrival directions remain influenced by magnetic field interactions [24]. Multiple processes can cause cosmic-ray particles to lose energy during propagation, such as radiative losses (e.g., Bremsstrahlung or synchrotron radiation, especially for electrons or positrons) or ionization losses. Additionally, the composition of cosmic rays may change during propagation. For instance, heavier nuclei can fragment due to inelastic collisions, or light mesons can be produced in such interactions with matter [24]. Consequently, the energy spectrum and mass composition of cosmic rays differ between their sources and their measurements at Earth, which must be considered when studying the origin of cosmic rays. Cosmic rays are generally classified as *primary* (accelerated at the source) or *secondary* (produced by inelastic collisions of primaries during transport) [24]. The study of secondary cosmic rays offers an opportunity to learn more about the propagation of cosmic rays through the interstellar medium and about the medium itself [25].

[3] The Larmor radius, also known as gyroradius, is the radius of the circular trajectory a charged particle takes under the influence of a magnetic field.

In the ultra-high-energy regime, cosmic rays can inelastically interact with the cosmic microwave background (CMB) or other extragalactic background radiation, thereby limiting their propagation distance. A cutoff in the cosmic-ray spectrum was first predicted in Ref. [37, 38] for protons with an energy above 10^{20} eV, because they are expected to interact with the CMB and produce pions through the Δ resonance:

$$p + \gamma_{\mathrm{CMB}} \rightarrow \Delta^{+} \rightarrow p + \pi^{0} \tag{2.7}$$

$$p + \gamma_{\mathrm{CMB}} \rightarrow \Delta^{+} \rightarrow p + \pi^{+} \tag{2.8}$$

This process, known as the Greisen-Zatsepin-Kuzmin (GZK) effect, is a mechanism for producing so-called *cosmogenic neutrinos* through the decay of charged pions, as postulated in Ref. [39]. Additional information on cosmogenic neutrinos can be found in Section 2.4. However, the observation of a cutoff in the spectrum at 10^{20} eV does not necessarily indicate that it is solely due to the GZK effect; it is also possible that the sources are unable to accelerate particles to higher energies. The mass composition of cosmic rays provides a handle for disentangling both scenarios. According to the Hillas criterion, a source can accelerate heavier nuclei to higher energies than lighter nuclei. As a result, if the cutoff is produced by the source reaching its acceleration limit, the measured mass composition should get heavier before the cutoff. Conversely, if a significant fraction of protons is observed at the highest energies, it implies that the cutoff is likely due to the GZK effect.

The discussion thus far has focused on acceleration scenarios within the standard model of particle physics. If dark matter particles couple to standard model particles, a fraction of cosmic rays could originate from, e.g., dark matter self-annihilation. However, strong constraints on dark matter self-annihilation have been established, e.g., through γ ray observations and CMB measurements [24].

2.2 Detection of cosmic rays

Detecting cosmic rays across the entire energy spectrum requires different experimental approaches. At energies up to several tens of TeV, direct detection is feasible [24], as we discuss in Section 2.2.1. For energies above approximately 1 PeV, cosmic rays are detected indirectly by measuring the products of their interactions in the atmosphere [24], as we outline in Section 2.2.2. These atmospheric interactions generate so-called extensive air showers, which we describe in Section 2.2.3.

2.2.1 Direct detection

Direct detection of cosmic rays requires experiments to be positioned above or high within the atmosphere, as cosmic rays otherwise interact with the atmosphere and thus cannot be measured directly [25]. The primary advantage of direct detection is the ability to identify particle species on an event-by-event basis [24], enabling the measurement of spectra for individual species rather than only an all-particle spectrum. Additionally, analyzing the composition of cosmic rays and comparing it to the chemical composition of astrophysical sources provides information about their potential origin [25]. The study

of unstable isotopes and their ratios relative to stable isotopes can enhance understanding of Galactic propagation processes [24]. However, the need for experiments to operate at high altitudes significantly restricts the possible detector size. Combined with the steeply declining cosmic-ray spectrum, this limitation imposes an upper bound on the energies that these detectors can measure. As a result, direct detection experiments are typically limited to energies of a few tens of TeV. At lower energies, measurements are constrained by the presence of solar energetic particles, which dominate below a few MeV. Given this energy range, these experiments are only sensitive to cosmic rays of Galactic origin [24].

Direct detection measurements can be conducted using instruments deployed on satellites or space stations (e.g., AMS02 [40], PAMELA [41], DAMPE [42], CALET [43], NUCLEON [44]) or via balloon-borne instruments (e.g., CREAM [45], ATIC [46], GAPS [47], BESS [48]). These experiments employ a range of detector types that utilize techniques common to classical particle detectors to achieve high-precision measurements. For instance, magnetic spectrometers are used to determine the charge, charge sign, and rigidity of cosmic-ray particles (see, e.g., Ref. [49, 50]), while calorimeters provide precise energy measurements and help with particle identification (see, e.g., Ref. [51, 52]).

2.2.2 Indirect detection

At energies higher than approximately 1 PeV, the flux of cosmic rays is insufficient for measurements using direct detection experiments deployed in space or on balloons. Consequently, larger detectors must be constructed on Earth, and cosmic rays at these energies can only be detected indirectly. When cosmic rays interact inelastically with the atmosphere, they initiate cascades of secondary particles that travel significant distances, forming a so-called extensive air shower (EAS) (see Section 2.2.3). Indirect detection techniques have enabled measurements of the cosmic-ray spectrum up to energies of 10^{20} eV, where a cutoff is observed [53], and individual events with even higher energies have been reported [54–58]. Due to the low flux, detectors with large surface areas or those capable of observing a large volume of atmosphere are required. Several techniques have been developed to address this challenge. For instance, arrays of particle detectors with spacings from several hundred meters to over 1 km [24], and fluorescence detectors that can observe air distances up to a few 10 km with a single telescope [59]. Cosmic-ray detection methods at these high energies can be categorized into measurements of the particle content and of the electromagnetic signature, the latter of which can be further subdivided.

Arrays of particle detectors, such as scintillator-based detectors, water Cherenkov detectors, and dedicated muon detectors, are deployed to measure the particle content of EAS at the ground level [24]. The energy of the primary cosmic ray can be estimated by integrating over the detected particles, and the shower arrival direction is determined from differences in arrival times [59]. Deploying muon detectors enables measurements of the muon-to-electron ratio, providing insight into the composition of the measured cosmic rays [59]. A key advantage of using particle detectors is their high duty cycle, which allows operation during both day and night. However, the interpretation of these measurements depends on hadronic interaction models, which introduce a systematic uncertainty (see

Section 2.2.3). Additionally, particle detectors exhibit lower accuracy in determining the shower maximum than methods that utilize electromagnetic signatures [59].

The second category includes three types of electromagnetic signatures: fluorescence light, optical Cherenkov light, and radio emission. Experiments measuring these signatures are primarily sensitive to the electromagnetic component of the EAS [59], so complementary measurements of the muonic component are advantageous. Fluorescence light is emitted by nitrogen molecules in the atmosphere, excited by the EAS [59]. In dark and clear nights, this method allows observation of large air volumes. Also, this method enables precise measurements of the air shower energy and direct measurements of the shower maximum[4], which is critical for determining cosmic-ray composition [59]. Electrons and positrons produced in the air shower travel faster than the speed of light in air, and thus generate optical Cherenkov radiation detectable by ground-based optical sensors [59]. This method provides direct energy measurements and allows reconstruction of the position of the shower maximum. However, as with fluorescence detectors, optical Cherenkov light detection is limited to operation during dark, clear nights, which restricts its duty cycle [59]. Lastly, the particle cascade can also be detected via the radio emission produced as it propagates through the atmosphere, a topic we discuss further in Chapter 3. This technique, when used in a radio-quiet environment, can operate day and night, resulting in a high duty cycle. Generally, with the discussed methods, the arrival direction of the air shower can be determined from signal arrival times [59].

Most air-shower experiments employ multiple of the discussed detection techniques to maximize measurement precision and to infer properties of the primary particle that initiated the cascade. Notable examples of hybrid detectors at the highest energies include the Pierre Auger Observatory (fluorescence, water-Cherenkov, scintillator and radio detectors, area approximately 3000 km^2) [60, 61], the Telescope Array (TA) (fluorescence, scintillator detectors, area approximately 700 km^2) [62] and its extension TAx4 [63], and the Tunka experiment (air-Cherenkov, scintillator and radio detectors) [64, 65].

2.2.3 Extensive air showers

When a high-energy cosmic ray[5] interacts with particles in the atmosphere, it initiates a particle cascade containing millions to billions of secondary particles, which is referred to as an extensive air shower (EAS) [59]. Figure 2.2 shows a simplified schematic of such an air shower. Generally, the cascade can be divided into a hadronic, electromagnetic and muonic component.

In the first interaction, secondary particles such as pions and other mesons are produced, initiating a hadronic cascade that forms the core of the air shower [25]. The resulting charged pions may either decay or further interact in the atmosphere, depending on their energy and lifetime. When the charged pions decay, muons are created, which is the largest contribution to the muonic component. Muons with energies above a few GeV typically reach the ground without decaying. Furthermore, neutrinos are produced in the charged pion decay and in a potential subsequent muon decay, which

[4] The shower maximum refers to the point in the atmosphere where the air shower reaches its maximum number of particles (see Section 2.2.3 for more information).

[5] Neutral particles, like e.g., γ rays can initiate an EAS as well [24].

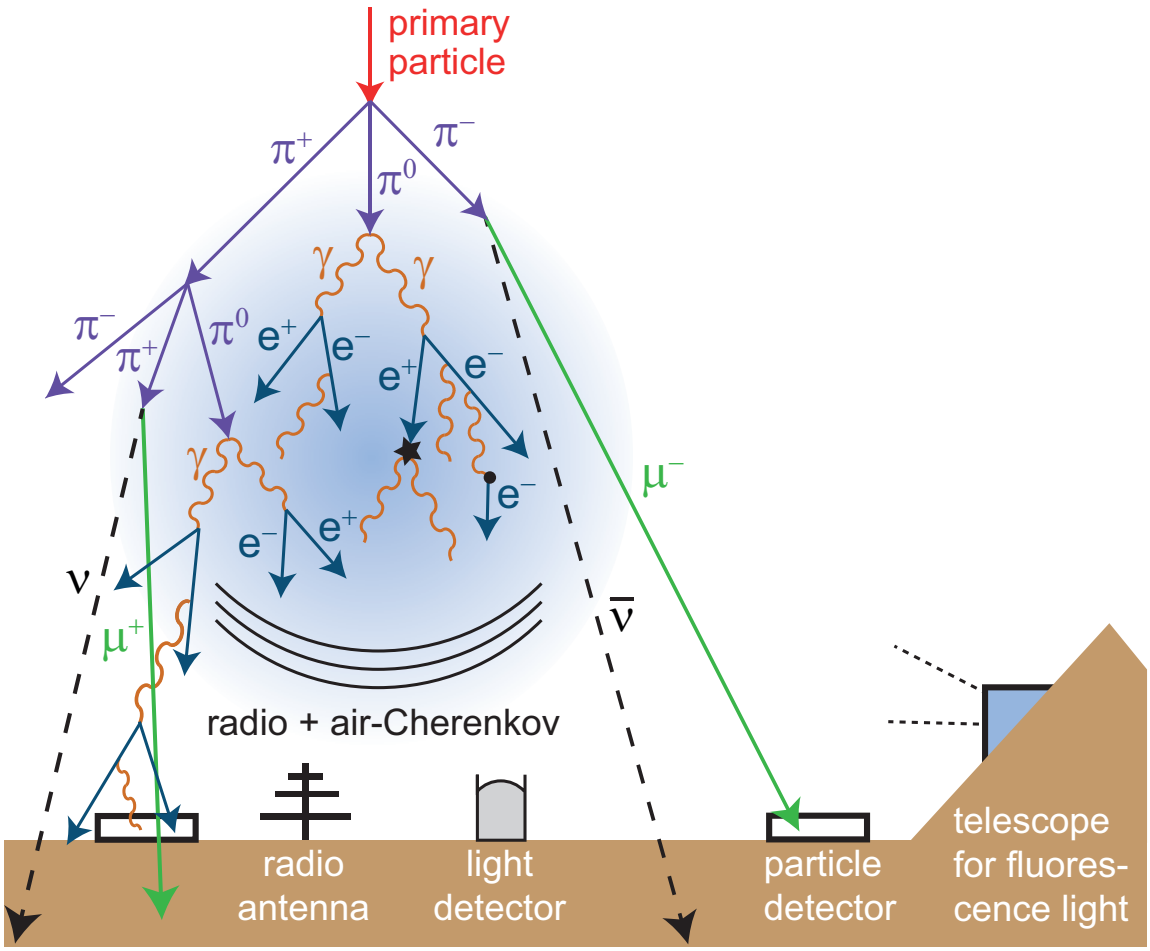


Figure 2.2: A simplified schematic of an extensive air shower, along with different experimental techniques for detecting it. Taken from Ref. [59].

are referred to as *atmospheric neutrinos* [24]. Next, to the charged pions, neutral pions are created, which will immediately decay into two photons. These photons initiate an electromagnetic subshower that develops through pair production, Bremsstrahlung, and inverse Compton scattering [59]. Additionally, muons with sufficiently low energy decay and also contribute to the electromagnetic cascade [24], although it is mainly fed by the hadronic shower core via neutral pion decay [25]. At each hadronic interaction, approximately 1/3 of the energy is transferred to the electromagnetic component. Since most hadrons interact before reaching the ground, the majority of the shower energy is ultimately deposited in the electromagnetic component [25]. Together with the fast development of the electromagnetic cascade, this leads to the majority of particles in the cascade being photons, electrons and positrons. Figure 2.3 shows the longitudinal particle distribution of an air shower. The electrons and positrons in the cascade lose energy and eventually reach the critical energy of around 80 MeV, at which ionization becomes the dominant energy-loss mechanism, leading to a rapid decline in their number [24], as shown in Figure 2.3. In contrast, the muonic component decreases more gradually due to fewer interactions and decays [25]. The point in the atmosphere at which the cascade reaches its maximum particle number, the shower maximum, is a characteristic parameter, often denoted X_{max}, where X represents the atmospheric depth[6].

[6] The atmospheric depth is the integrated atmospheric density a shower traveled through.

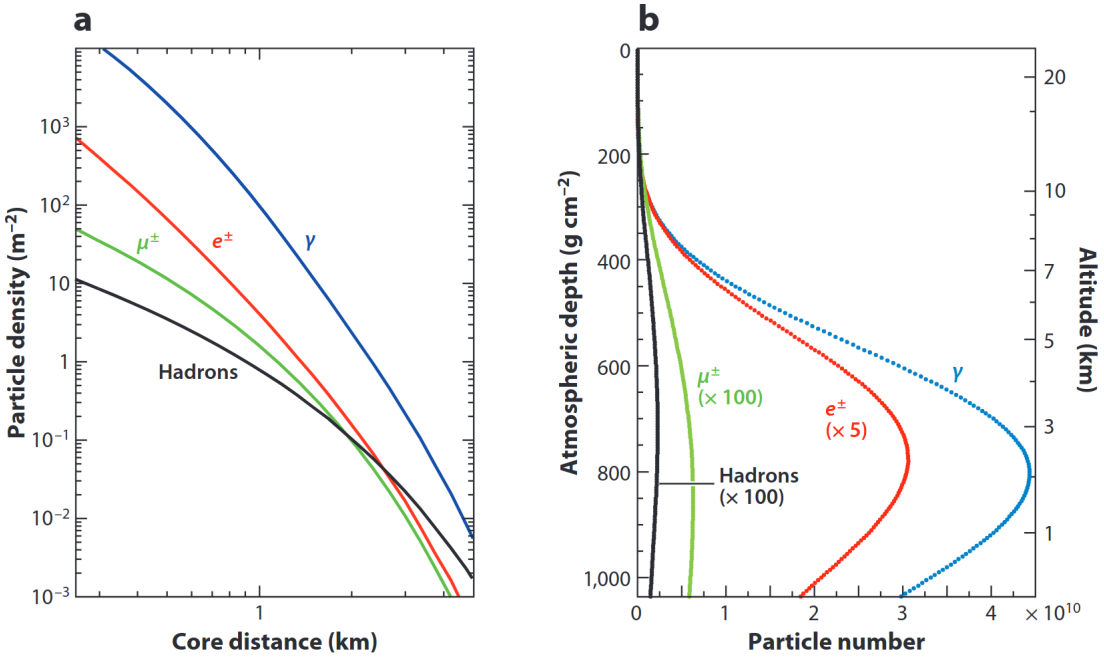


Figure 2.3: Average (a) lateral and (b) longitudinal shower profiles for vertical, proton-induced showers at 10^{19} eV. The lateral distribution of the particles at ground is calculated for 870 g/cm^2, the atmospheric depth of the Pierre Auger Observatory. Taken from Ref. [66].

Most of the shower energy is deposited in the atmosphere. However, if the experiment that measures the air shower is at high elevation, an incompletely developed shower can reach the ground, initiate a new particle cascade, and thus deposit energy there [19–21]. Additionally, a fraction of the energy (10 − 20 % for cosmic rays with an energy above 0.1 EeV) is carried away by muons and neutrinos due to their long travel distances before interacting [24].

Next to the longitudinal distribution, Figure 2.3 also shows the lateral distribution of particles within an air shower, which primarily arises from the transverse momentum of hadronic interactions, resulting in the largest lateral spread for the hadronic component. Multiple Coulomb scatterings of electrons and positrons in the electromagnetic component also contribute to this lateral distribution [24].

In a simplified model, the maximum number of particles N_{max} and the atmospheric depth at which this maximum is reached X_{max} can be related to the energy E of the primary cosmic-ray particle and its mass number A as follows [24]:

$$N_{\mathrm{max}} \propto E \tag{2.9}$$

$$X_{\mathrm{max}} \propto \log(E) \tag{2.10}$$

$$X_{\mathrm{max}} \propto \log\left(A^{-1}\right) \tag{2.11}$$

More complex models exist to describe an EAS, but these are beyond the scope of this introduction. To determine the composition of the cosmic-ray particles, two quantities are important in the simple model: the position of the shower maximum (X_{max}) and the electron-to-muon ratio of the shower. The cross-section for protons or nuclei interacting with the atmosphere is roughly proportional to the number of nucleons. As a result, showers initiated by particles with more nucleons tend to interact at higher altitudes,

thereby affecting the position of the shower maximum. Furthermore, these showers terminate earlier than those initiated by lighter particles and transfer less energy to the electromagnetic component, which impacts the electron-to-muon ratio [59]. Therefore, measuring these quantities enables the extraction of information about the composition of the primary cosmic ray on a statistical basis.

CORSIKA is a simulation tool that enables EAS simulation at the particle level [67]. For these simulations, hadronic interactions at the highest energies need to be described, which is commonly done using one of three models: Sibyll [68], QGSJET [69], or EPOS [70]. These models do not agree on all aspects, such as the predicted number of muons [71, 72], thereby introducing systematic uncertainties. However, studying this energy regime is challenging, as the highest energies correspond to center-of-mass energies exceeding those of the most energetic collisions at the Large Hadron Collider (LHC), and these interactions occur in the forward direction, which is not well studied either [59]. Notably, a collision of a 0.1 EeV cosmic-ray particle within the atmosphere yields approximately a center-of-mass energy comparable to that achieved at the LHC, highlighting the extreme energies involved in the highest-energy EAS [24].

2.3 Cosmic rays at the highest energies

Since the initial detection of cosmic rays, the field has advanced significantly, enabling measurements of cosmic rays to energies of 10^{20} eV. The observation of particles at such high energies implies the existence of astrophysical sources capable of accelerating them to these energies. Additionally, they need to be located within a few hundred megaparsecs, since their propagation horizon is set by energy losses due to interactions with the cosmic background light [28].

The energy spectrum has been measured with high precision. However, discrepancies, such as an energy-dependent difference above 10 EeV, remain between the spectra reported by the two main observatories, the Pierre Auger Observatory and TA. A joint working group determined that the spectra are compatible within systematic uncertainties below energies of 10 EeV, but discrepancies persist at higher energies [73]. Several features have been identified in the energy spectrum at the highest energies, manifesting as deviations from a power law. These features are thought to result from changes in the composition of the cosmic-ray particles [24].

Indirect detection methods do not allow event-by-event measurements of the cosmic-ray composition; instead, the composition is inferred primarily from the observable X_{max} on a statistical basis. For a sample of measured cosmic rays, the mean value $\langle X_{\mathrm{max}} \rangle$ and its spread $\sigma(\langle X_{\mathrm{max}} \rangle)$ are compared with simulations for different primary particles, enabling an estimation of the fractions of protons, helium, carbon-nitrogen-oxygen and iron. However, significant uncertainties in the hadronic interaction models used in these simulations introduce corresponding uncertainties in the mass composition estimates [24]. Figure 2.4 shows $\langle X_{\mathrm{max}} \rangle$ and $\sigma(\langle X_{\mathrm{max}} \rangle)$ as a function of energy, as measured by various experiments. Additionally, the plots display predicted values for pure proton (lightest) and iron (heaviest) compositions, with three lines representing different hadronic interaction models. At energies near 10^{16} eV, the composition is

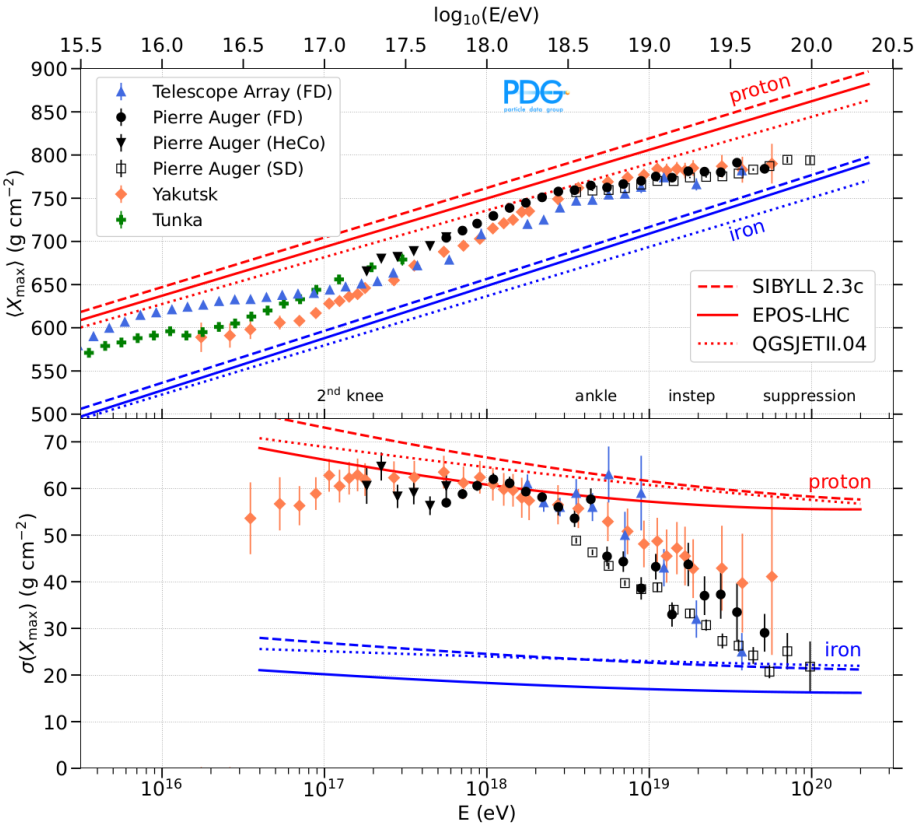


Figure 2.4: Measurements of the mean depth of shower maximum $\langle X_{\mathrm{max}} \rangle$ and its spread $\sigma(\langle X_{\mathrm{max}} \rangle)$ as a function of energy. Additionally, the predictions for a pure proton and iron composition are displayed for the three common hadronic interaction models. The plot is taken from [24], and the sources for the different datasets can be found in the original publication.

dominated by light primaries; it is getting heavier towards energies of approximately 10^{17} eV, and then becomes lighter again near energies of approximately 2×10^{18} eV. Above energies of around 10^{19} eV, the composition is measured to be mixed, with a trend toward heavier elements at higher energies [24]. The trend towards a heavier composition above energies of approximately 10^{19} eV is further supported by measurements of $\sigma(\langle X_{\mathrm{max}} \rangle)$, as shown in Figure 2.4. This trend toward heavier elements at the highest energies has implications for potential source classes, since for some sources, such as active galactic nuclei (AGN) or gamma-ray bursts (GRBs), it has been shown that producing heavy cosmic rays can be difficult [28].

Several potential source classes have been proposed for accelerating cosmic rays to energies up to 10^{20} eV, including GRBs[7], AGN[8], starburst galaxies[9], pulsars or magnetars[10], and tidal disruption events (TDEs)[11]. A brief discussion of these potential source classes is provided in Ref. [28].

[7] Gamma-ray bursts are bursts of γ rays that originate from massive stars collapsing into a black hole or mergers of compact binaries [25].
[8] Active galactic nuclei describe a black hole at the center of a galaxy that is accreting matter. Through the accretion, gravitational energy is ejected in a jet, perpendicular to the accretion disk [74].
[9] Starburst galaxies are galaxies with extremely high star formation rates [28].
[10] Pulsars are rotating neutron stars with a strong magnetic field and magnetars are a subclass with an even stronger magnetic field [25].
[11] Tidal disruption events are events where a black hole accretes a star [75]

As previously discussed, the highest-energy cosmic rays can travel only a limited distance before interacting with the cosmic background light, which restricts the observable volume and the number of potential sources. This limitation makes the identification of sources and the statistical study of their properties possible [24]. An analysis of the arrival direction of cosmic rays with energies above 8 EeV by the Pierre Auger Observatory revealed a large dipole anisotropy with a high significance ($> 6\sigma$). The dipole direction points away from the Galactic center, which supports an extragalactic origin for cosmic rays at these energies [76]. Furthermore, the Pierre Auger Observatory identified an excess (around 4σ) of cosmic rays with energies above $E > 38$ EeV, correlated with the region surrounding the radio galaxy Centaurus A [77]. The Observatory also reports a correlation between detected cosmic rays and a catalogue of starburst galaxies at a significance of around 4σ, likely driven by two starburst galaxies, NGC4945 and M83, within the Centaurus A region, and a third galaxy, NGC253, located in a region for which the observatory reports a mild over-density [77]. TA reports a separate cosmic-ray over-density in the northern sky, referred to as "hotspot", with a significance of 2.8σ [78]. However, the Pierre Auger Observatory does not observe a significant excess for the same region [79].

Despite precise measurements of the energy spectrum, mass composition and arrival direction, it has not yet been possible to identify the acceleration mechanisms or sources responsible for cosmic rays at PeV to EeV energies [24]. Identifying the sources of cosmic rays is challenging because their deflections by galactic and intergalactic magnetic fields prevent them from pointing directly back to their sources. A solution is to use a different messenger, the neutrino. Neutrinos are expected to be produced in the same sources as cosmic rays, mainly through the decay of charged pions and, unlike cosmic rays, are not deflected during their journey to Earth. Consequently, neutrinos point directly back to their source. In Section 2.4, we discuss the advantages and challenges associated with using neutrinos as messengers.

2.4 Neutrino astronomy

The neutrino is a particle of the Standard Model that carries no electromagnetic charge. It interacts only via the weak force or gravity, and is therefore effectively not absorbed by intervening matter or deflected by magnetic fields during its journey from the source to Earth. This property makes the neutrino a unique messenger for investigating the origins of ultra-high-energy cosmic rays (UHECR, $E_{\mathrm{CR}} > 10^{16}$ eV). The weakly interacting nature of neutrinos can be illustrated by the following rule of thumb: approximately 10^{11} neutrinos pass through a human thumbnail each second, yet over an entire lifetime, only one or two will interact within the entire human body [80]. In contrast, γ rays are also neutral particles and thus are not deflected by magnetic fields, but they interact electromagnetically. As a result, the mean free path of PeV-energy γ rays is of the order of galactic scales, since they have a high chance of interacting with background radiation such as the cosmic microwave background (CMB) [59]. This restricts the range of cosmic-ray sources that can be studied using γ rays. However, using neutrinos as the messenger presents unique challenges. Due to their low flux at high energies and

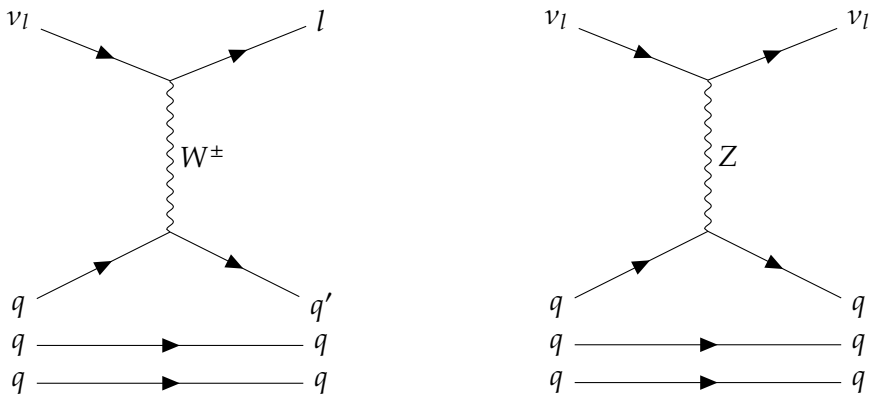


Figure 2.5: Feynman diagram of a charged current (left) and neutral current (right) interaction of a neutrino ν_l with a nucleus (represented as three q). The flavour of the incoming neutrino and outgoing neutrino/charged lepton (l) can be $l = e, \mu, \tau$. In the charged current interaction, one outgoing quark is different from the incoming quark due to charge conservation. The outgoing quarks will produce a hadronic shower.

small cross-section, extremely large detectors are required for their measurement, as we discuss in the next sections.

At the highest energies, neutrinos are classified according to their production site. Neutrinos originating from astrophysical sources are called *astrophysical* neutrinos, while those created by the interaction of cosmic rays with background radiation during transit to Earth are referred to as *cosmogenic* neutrinos. A third category, considered background in neutrino astronomy, is *atmospheric* neutrinos, produced when cosmic rays interact with the atmosphere [24]. In all three cases, the primary production mechanism is the decay of charged pions produced by the interaction of UHECRs with a target[12].

2.4.1 Detection of high-energy neutrinos

Within the high-energy neutrino range, defined here as 20 GeV to 10 PeV, the dominant neutrino interaction is the deep inelastic scattering off a nucleus. In this process, a neutrino scatters off a single quark within the nucleon through the exchange of a Z or $W^\pm$ boson, resulting in the production of a lepton and hadrons. Depending on the exchange boson, the interaction is classified as either neutral current (Z boson) or charged current ($W^\pm$) [81]. Figure 2.5 shows the Feynman diagram for the neutral current (NC) and charged current (CC) interaction. In NC interactions, the final-state lepton is a neutrino with the same flavour as the incoming neutrino. In CC interactions, the final-state lepton is an electron, muon or tau, depending on the flavour of the incoming neutrino. The charged final-state particles from these interactions travel faster than the speed of light in the medium, producing Cherenkov radiation detectable with optical detectors. The various final-state products generate distinct signatures in a detector medium such as ice or water, allowing some sensitivity to neutrino flavour. Neutral current interactions of all flavours and CC interactions of electron neutrinos produce so-called *cascades*: the final-state hadrons generate a hadronic cascade, and the electron

[12] For cosmogenic neutrinos, an additional dominant production channel around EeV energies is the decay of neutrons produced through photodisintegration [25]. In the case of atmospheric neutrinos, there is a subdominant contribution from the decay of charmed mesons, known as *prompt* neutrinos, while the dominant component is called *conventional* neutrinos [25].

in the CC interactions initiates an additional electromagnetic cascade. Consequently, the deposited energy and the resulting Cherenkov radiation are more localized [25]. CC interactions of muon neutrinos produce so-called *tracks*, as the muon travels significant distances due to its long lifetime while emitting Cherenkov light and, above a critical energy, losing energy through stochastic processes. As a result, Cherenkov radiation can be measured along the trajectory of the muon [25]. Lastly, CC interactions of tau neutrinos can produce a so-called *double bang* signature: the final-state hadrons generate a first cascade signature, and the tau decay produces a second cascade signature. If the tau is energetic enough, it travels far enough before decaying such that both cascades are separated, which forms the unique double-bang signature. However, most of the time, the signature is indistinguishable from a normal cascades [25]. In addition to flavour sensitivity, these signatures offer different advantages: Tracks provide the best angular resolution, but some energy may be deposited outside the detection volume. For cascades, all energy is deposited in the detection volume, resulting in good energy resolution, but they provide reduced angular resolution.

Due to the low astrophysical neutrino flux at high energies, large-scale detectors are required. These detectors are made feasible by instrumenting naturally occurring media such as ice and water. The standard approach for high-energy neutrino detection involves deploying arrays of optical modules to measure the Cherenkov radiation. By analyzing the arrival times of Cherenkov photons at the different optical modules, the arrival direction of the neutrino can be reconstructed, while the amount of detected light provides information for energy reconstruction. The first neutrino telescope employing this technique was proposed in the 1970s, the Deep Underwater Muon and Neutrino Detector (DUMAND) project. Although DUMAND was never fully realized, the basic concepts of the current telescopes are based on its studies [25]. The first successful neutrino telescopes were NT200 in Lake Baikal (water), Antarctic Muon And Neutrino Detector Array (AMANDA) at the South Pole (ice) and Astronomy with a Neutrino Telescope and Abyss enviromental RESearch (ANTARES) in the Mediterranean Sea (deep sea) [82]. Presently, the largest operational neutrino telescopes are the cubic-kilometer-scale IceCube Neutrino Observatory [83] at the South Pole, KM3NeT in the Mediterranean Sea (under construction and aiming to match IceCube in scale) [84], and Baikal-GVD in Lake Baikal [85]. Furthermore, the Pacific Ocean Neutrino Telescope (P-ONE), currently in its demonstrator phase, is a planned telescope in the Pacific Ocean off the coast of Canada [86, 87]. Additional proposed large-scale neutrino telescopes include the TRopIcal DEep-sea Neutrino Telescope (TRIDENT) [88] and the Huge Underwater high-energy Neutrino Telescope (HUNT) [89].

In 2014, the IceCube Neutrino Observatory announced the first observation of an astrophysical neutrino flux with a significance exceeding 5σ above the atmospheric neutrino background, which dominates up to around 100 TeV [90]. Neutrino astronomy has since become a well-established field, and the astrophysical neutrino flux is now measured up to energies of 10 PeV. Recently, the IceCube Collaboration reported a deviation from a single power-law description for the astrophysical neutrino flux, with a significance greater than 4σ [4]. Moreover, the KM3NeT Collaboration observed the most energetic neutrino to date, with an estimated neutrino energy of 220 PeV [5]. Investigations into whether this event could be associated with blazars[13], which are

[13] A blazar is a type of AGN characterized by a relativistic jet oriented close to the line of sight of the observer [25].

considered promising astrophysical neutrino source candidates, identified multiple candidates, without being able to conclusively link one of them to the event. An improved position calibration is expected to improve the direction reconstruction and may help confirm or exclude these candidates as the source of the event [91]. In addition, the IceCube Collaboration has identified several astrophysical neutrino sources, including the blazar TXS 0506+056 [1] and the nearby active galaxy (AGN) NGC 1068 [2]. Also, the IceCube Neutrino Observatory has detected neutrinos from the Milky Way Galaxy, produced by cosmic-ray interactions with the interstellar medium [3]. Additionally, coincidences between TDEs and neutrinos detected by the IceCube Neutrino Observatory have been reported [92–94], although improved reconstruction algorithms place the TDEs outside the error contours of the reconstructed neutrino directions [95]. Besides the advances in the field of neutrino astronomy, the IceCube Collaboration reported the first detection of an event consistent with a neutrino interaction at the Glashow resonance[14] [96]. This represents the first such detection, as the resonance energy exceeds the capabilities of current and planned terrestrial particle accelerators [96].

At energies higher than 10^{16} eV, the neutrino flux becomes so low that larger detectors than IceCube or KM3NeT are necessary to detect them. Expanding the optical detection technique to such scales is very expensive, as the optical modules must be placed relatively close together due to the short attenuation and scattering length of $\mathcal{O}(100\,\mathrm{m})$ [10, 11] in ice or seawater. Consequently, the development of cost-effective detection techniques is essential to make these large-scale detectors feasible.

2.4.2 Detection of ultra-high-energy neutrinos

To date, apart from the exceptional ultra-high-energy neutrino (UHEN, $E_\nu > 10^{16}\,\mathrm{eV}$) observed by KM3NeT [5], no other neutrino has been detected in the ultra-high-energy regime. Despite this, theoretical predictions suggest the existence of neutrinos in this energy range, as UHECRs with energies up to around 10^{20} eV have been observed. These neutrinos are produced either through interactions of the UHECRs with matter at their sources or during their propagation to Earth. Figure 2.6 shows the measured γ ray, high-energy neutrino, and UHECR fluxes, alongside model predictions for astrophysical or cosmogenic UHEN fluxes. A comparison shows that all three messengers exhibit similar flux levels, suggesting a common origin[15]. Notably, the predicted UHEN fluxes show significant variation, motivating the development of new detectors capable of constraining this parameter space. Currently, the most stringent limits in this energy regime are provided by IceCube and the Pierre Auger Observatory [6, 7], which are in mild tension with the KM3NeT observation. However, as mentioned before, existing detectors such as IceCube and KM3NeT are not suitable for probing the ultra-high-energy regime because their detection volumes are too small to compensate for the rapidly decreasing neutrino flux at these energies. Additionally, constructing sufficiently large optical neutrino telescopes is not cost-effective. Therefore, alternative techniques for UHEN detection are discussed below.

[14] Resonant production of a W^- boson in the interaction of an electron antineutrino with an electron, leading to an enhanced cross-section [96].

[15] Neutral and charged pions can be created in the interaction of cosmic rays with matter, which subsequently decay into γs and neutrinos, respectively.

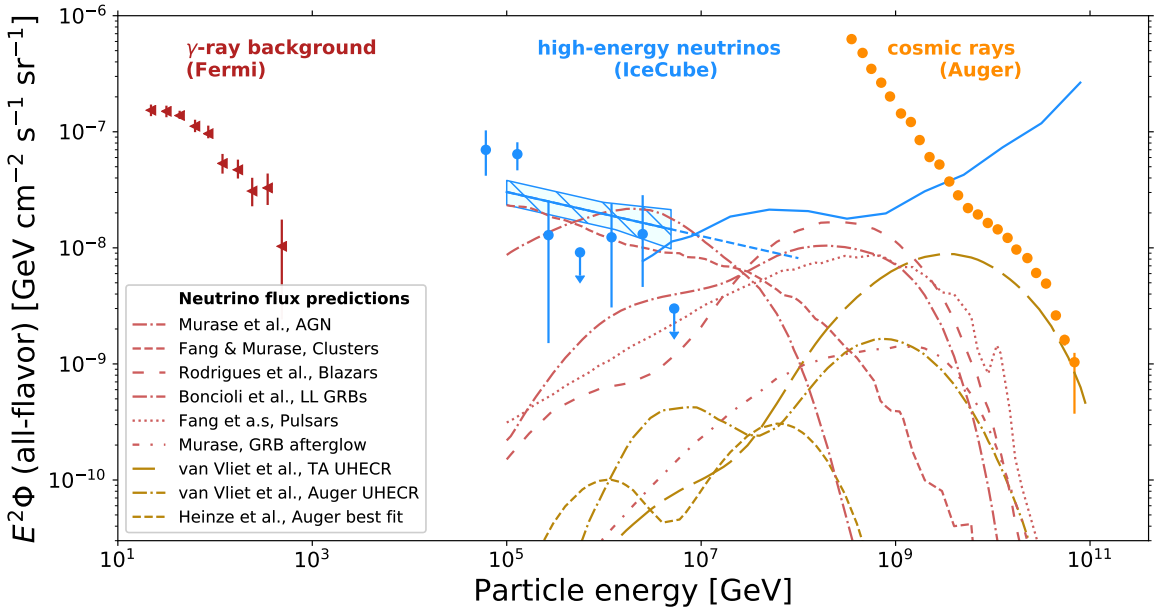


Figure 2.6: The γ ray, high-energy neutrino, and UHECR fluxes as measured by Fermi, the IceCube neutrino observatory, and the Pierre Auger Observatory, respectively. Also, shown are the astrophysical (red lines) and cosmogenic (dark yellow lines) UHEN flux predications from various models. The plot is taken from [14], and the sources for the presented data or models can be found in the original publication.

In the ultra-high-energy regime, the dominant interaction channel for neutrinos is, as in the high-energy regime, deep inelastic scattering with nuclei [81]. As described in the previous section, this interaction can occur via NC or CC processes, both of which produce hadrons in the final state that initiate a hadronic cascade within the medium. Figure 2.5 shows a Feynman diagram of both processes. In CC interactions, a lepton is also created, depositing additional energy in the detector volume as either an electromagnetic[16] or hadronic cascade, depending on the lepton flavour. It should be noted that at energies in the PeV to EeV regime, the neutrino-nucleon cross-section increases sufficiently for the Earth to become opaque to these neutrinos [14]. As a result, detectors in this energy regime are primarily sensitive to UHEN arriving from near the horizon. In the following, we outline three experimental approaches for detecting the particle cascade produced by UHEN interactions with matter.

Acoustic: The energy deposited by the particle cascade locally heats the medium, leading to a rapid expansion. This expansion creates a pressure wave, known as the *acoustic pancake*, that propagates in a cylindrical pattern, perpendicular to the main axis of the cascade [100]. These pressure waves can be detected by hydrophones and can travel significant distances through, e.g. water, due to the large attenuation length of $\mathcal{O}(5\,\mathrm{km})$ at the frequency of 10 kHz [100]. Consequently, the acoustic technique enables cost-effective instrumentation of large volumes of ice or water. However, the energy threshold for neutrino detection using this method is relatively high, around 1 EeV [59]. Although multiple acoustic detectors have been constructed and operated (see Ref. [101], for an overview), current research efforts in this area are limited [59]. Presently, acoustic signals

[16] At high energies, the electromagnetic cascade is influenced by the Landau-Pomeranchuk-Migdal (LPM) [97, 98] effect. The LPM effect elongates the cascade and makes it more irregular, and at very high energies, spatially displaced sub-cascades become visible [99].

are primarily used to calibrate the positions of optical modules in underwater neutrino telescopes, such as KM3NeT.

Radar: The radar technique exploits the property that the plasma created by the particle cascade can refract radio signals with frequencies below the plasma frequency, which depends on the electron density [102]. In this approach, a radio transmitter is deployed within a dense medium, such as ice, to emit radio waves, while receiver antennas are positioned to detect the refracted signal. This method benefits from the large attenuation lengths of radio waves in ice ($\mathcal{O}(1\,\mathrm{km})$ [9]) and is proposed to be sensitive to energies ranging from several PeV to a few EeV [102]. Experimental confirmation of the radar technique has been achieved at accelerators, where particle cascades were initiated in a plastic target [103]. However, confirmation in non-laboratory environments has not yet been demonstrated. A pathfinder experiment, the Radar Echo Telescope for Cosmic Rays (RET-CR), searches for air shower cores that impact and produce a particle cascade in the ice [104], and a neutrino telescope employing this technique has been proposed (Radar Echo Telescope for Neutrinos (RET-N)) [105].

Radio: Particle cascades developing in a medium produce electromagnetic emission in the radio frequency range. This concept has been successfully applied to measure UHECR through radio emission from EAS. The large attenuation length of radio waves in ice of $\mathcal{O}(1\,\mathrm{km})$ [9], combined with the low cost of radio antennas, makes in-ice radio detection particularly suitable for scaling neutrino detectors to large volumes. Another scalable approach for neutrino detection is the search for *earth-skimming* neutrinos, which are mainly ultra-high-energy tau neutrinos. When these neutrinos interact within the Earth near the surface, the final-state tau lepton may escape the Earth and decay in the atmosphere, producing an EAS detectable by radio antennas. A detailed discussion of the emission mechanisms, different detection techniques, and radio neutrino detectors is presented in Chapter 3.

Radio Detection 3

When an ultra-high-energy particle interacts within the atmosphere or a dense medium such as ice, it produces a radio signal. The large attenuation length of radio emission in ice ($\mathcal{O}(1\,\mathrm{km})$) combined with the relatively low cost of radio antennas opens the possibility of constructing neutrino telescopes larger than IceCube or KM3Net to explore the regime of ultra-high-energy neutrinos. In this chapter, we explain the two dominant mechanisms that can produce coherent radio emission from a particle cascade developing within a medium and discuss the resulting signal (see Section 3.1). Furthermore, we describe how experiments utilize the radio emission from air showers to detect UHECR (see Section 3.2) and how radio emission from particle cascades is exploited to enable feasible neutrino telescopes for the ultra-high-energy regime (see Section 3.3).

3.1 Radio emission from particle cascades

Ultra-high-energy particles interacting with the atmosphere or denser media, such as ice, initiate a particle cascade through deep-inelastic scattering. In the atmosphere, this cascade is referred to as an extensive air shower (EAS) (see Section 2.2.3), which typically extends over several kilometers along the shower axis [59]. In contrast, in a denser medium like ice, the particle cascade develops over a much shorter distance, typically on the order of meters [59]. Regardless of the medium, the particle cascade produces coherent electromagnetic emission in the radio-frequency range (MHz - GHz) as it develops. While the underlying emission processes are the same across different media, their relative contributions depend on the density of the medium. The two primary mechanisms responsible for the radio emission are the *geomagnetic* and *Askaryan* effects, which are discussed in Sections 3.1.1 and 3.1.2, respectively. Additionally, a tertiary emission mechanism, known as geo-synchrotron radiation, is suggested to become measurable for very inclined air showers [106]. The emitting region of the particle cascade moves faster than the speed of light in the medium, resulting in a Cherenkov cone structure of the radio emission. Further details regarding the characteristics of the coherent radio signal are provided in Section 3.1.3.

3.1.1 Geomagnetic emission

During the development of a particle cascade in air, electrons and positrons are created as secondary particles, which are mainly concentrated in the shower front [107]. The shower front, also called *pancake*, has a thickness of a few meters because some particles lose energy through interactions with the atmosphere and trail behind, thereby forming an extended shower front [108]. As the shower front propagates toward Earth, electrons and positrons experience the Lorentz force generated by the magnetic field of the Earth and are accelerated in opposite directions, perpendicular to the shower axis. They move away from the shower axis with a constant drift velocity due to multiple collisions, inducing a current in the shower front [107]. The magnitude of this induced current is

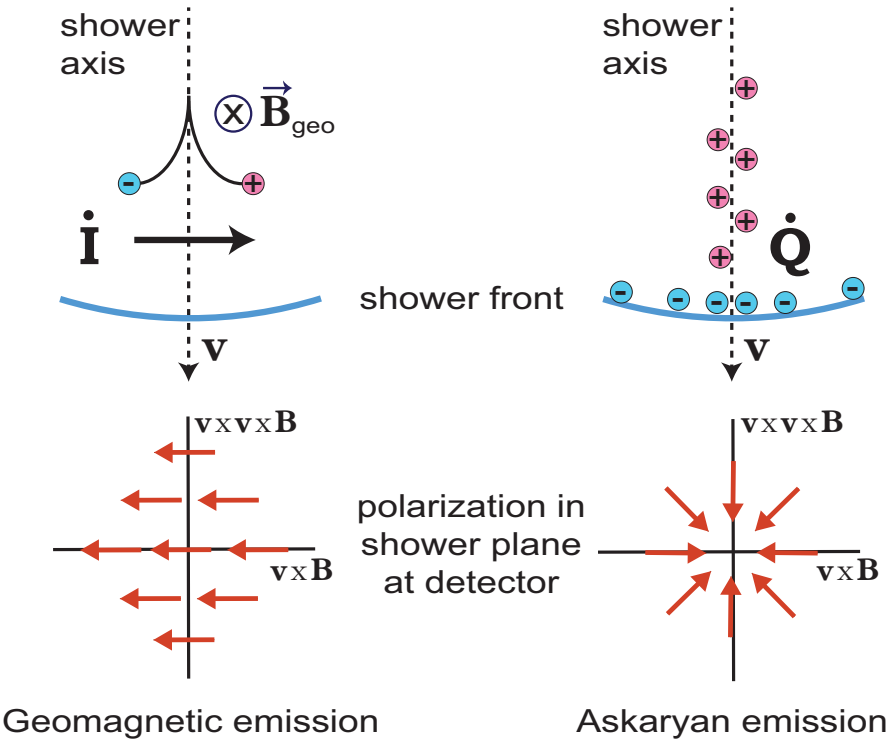


Figure 3.1: Schematics of the geomagnetic (left) and Askaryan (right) mechanisms responsible for radio emission from a particle cascade. Additionally, the resulting signal polarization for both mechanisms is depicted. For a detailed explanation, see the text. Taken from Ref. [111].

approximately proportional to the number of contributing charged particles [109]. During shower development, the number of charged particles increases to a maximum before decreasing due to ionization losses, resulting in a time-varying induced current [110]. In this macroscopic framework, a time-varying current produces electromagnetic radiation, analogous to that of a dipole antenna. Figure 3.1 shows a schematic of this macroscopic model for the geomagnetic emission. This emission mechanism was first described by Kahn and Lerche in 1966 [16] and experimentally confirmed by the LOFAR Prototype Station (LOPES) in 2005 [112]. Alternatively, the emission can be described in a microscopic framework, combining the contributions from individual particles to the full signal, with each emission of a particle described by the endpoint formalism [113]. Simulations have demonstrated the compatibility between the macroscopic and microscopic descriptions of this mechanism [114].

The resulting radio emission has a characteristic linear polarization aligned with the direction of the induced current [110], as illustrated in Figure 3.1. The strength of the emitted signal depends on the induced current, which is proportional to the number of charged particles, and also on the magnitude of the Lorentz force. The Lorentz force $\vec{F}$ is given by [115]

$$\vec{F} = q\vec{v} \times \vec{B}, \tag{3.1}$$

where $\vec{v}$ is the velocity vector of particles with charge q and $\vec{B}$ denotes the magnetic field vector. The magnitude of the Lorentz force depends on the angle between Earth's

magnetic field and the travel direction of the particle cascade:

$$|\vec{F}| = q|\vec{v}||\vec{B}|\sin(\alpha) \tag{3.2}$$

Here, α represents the angle between the magnetic field and the shower direction [59]. The dependence on this angle α can lead to an asymmetry in the distribution of arrival directions for detected events, as air showers with weaker radio signals are more difficult to detect. This asymmetry has been confirmed experimentally [116].

While the same mechanisms can, in principle, be applied to particle cascades in media other than air, in denser media the particle cascade travels a shorter distance ($\mathcal{O}(10\,\mathrm{m})$) compared to those in air. Consequently, the travel distances of the charged particles are insufficient for geomagnetic emission to become relevant, and thus it is negligible in dense media [59].

3.1.2 Askaryan emission

During the development of a particle cascade in a medium, the shower front - where electrons and positrons are mainly concentrated - develops an electron excess, while the ionized medium behind becomes positively charged [59]. This charge excess arises because the medium contains only electrons (not positrons), which can interact with electrons, positrons, and photons from the cascade. Electrons from the medium may be accelerated and join the shower front through Møller ($e^- + e^- \rightarrow e^- + e^-$), Bhabha ($e^+ + e^- \rightarrow e^+ + e^-$) and Compton scattering ($\gamma + e^- \rightarrow \gamma + e^-$). In contrast, positrons can be removed from the shower front by annihilation with electrons from the medium or can be decelerated via Bhabha scattering [117]. The resulting electron excess relative to the total number of electrons and positrons can reach 20-30% [59]. As the cascade progresses, the number of electrons and positrons increases until a maximum is reached, after which it declines as the particles lose energy through processes such as Bremsstrahlung, eventually reaching energies at which ionization becomes the dominant energy-loss mechanism. At this point, particle absorption by interactions with the atmosphere exceeds production, resulting in a net decline in particle number [59]. Consequently, the magnitude of the charge excess varies over time and thus produces electromagnetic radiation in the radio frequency range. This electromagnetic emission is linearly polarized, with the polarization vector oriented toward the shower axis [110]. Figure 3.1 presents a schematic of the emission mechanism and the associated polarization. The emission mechanism was first predicted by G. A. Askaryan [8] and is therefore termed *Askaryan* emission. In the context of radio emission from EAS, this process is also known as *charge-excess* emission.

In dense media such as ice, the Askaryan mechanism is the only relevant process for radio emission, whereas in air, both geomagnetic and Askaryan mechanisms contribute to the overall radio signal. However, the Askaryan emission is subdominant in air, with its relative strength reaching up to around 20% [59]. Askaryan radiation has been detected in laboratory experiments at a particle accelerator [118], as a component of the radio signal from EAS [119], and recently by a radio-neutrino detector from air-shower cores that impacted and continued to develop in ice [22].

3.1.3 Characteristics of the radio signal

Both emission mechanisms described above produce coherent radio emission, with the emitted signal power scaling with N^2, where N is the number of charged particles. In contrast, incoherent radio emission scales only with N. For an EAS initiated by a primary particle with an energy of 10^{18} eV, N can reach $\mathcal{O}(10^8)$ at the shower maximum [108]. Furthermore, the number of electrons and positrons increases linearly with the energy of the primary particle, allowing for an estimation of the primary energy from the measured electric field, as confirmed experimentally [120]. Coherence requires that the emitted wavelength exceed the size of the emitting region. If the wavelength is much smaller than the emitting region, the substructure of the region becomes resolvable, and different parts contribute with varying phases, resulting in a loss of coherence [108]. Based on the coherence condition and the relevant size of the emitting region (the shower front), the frequencies at which the signal remains coherent can be estimated. The relevant length scale ranges from the shower front thickness to the lateral spread of the shower front, depending on the angle at which the shower front is observed. For air showers, the shower front thickness is the most relevant length scale, which is $\mathcal{O}(1\,\mathrm{m})$, and thus they typically have the maximum contributions at frequencies below 100 MHz. Whereas for in-ice showers, the lateral spread is the relevant length scale, which is about 0.1 m, and therefore they have the maximum contributions at frequencies of a couple of 100 MHz [59]. The number of charged particles in the cascade that produce the radio emission increases to a maximum and then decreases until all particles are depleted. Therefore, the particle cascade exists only within a limited temporal and spatial region, and the time integral over the emitted electric field must be zero. The simplest solution that satisfies this condition is a bipolar pulse. A unipolar pulse would not be physical. Indeed, geomagnetic and Askaryan emission create bipolar electric field pulses [108].

Moreover, the particle cascade propagates through the medium at a velocity exceeding the speed of light in that medium. As a result, the radio emission exhibits a Cherenkov cone structure, which means that there is an angle relative to the shower axis where the observer sees all signals arriving simultaneously, regardless of whether the signal was emitted at the beginning or end of the shower. This angle is the Cherenkov angle θ_C and is defined as [117]:

$$\theta_C = \arccos\left(\frac{1}{\beta n}\right) \tag{3.3}$$

Here, β denotes the velocity of the particle cascade relative to the speed of light in vacuum, which is approximately 1, and n is the refractive index of the medium. This Cherenkov angle is approximately 1° for cascades in air and 56° in ice [59]. Due to this geometric effect, the resulting signal structure is a cone with an opening angle θ_C. For an observer at the Cherenkov angle, all signals arrive simultaneously, resulting in a time-compressed, very short electric-field pulse. The emitted signal remains coherent up to several GHz, exceeding what would be possible without time compression of the signal [59]. If the observer is positioned away from the Cherenkov angle, signals from different stages of shower development arrive at different times, causing the pulse to lengthen and the frequency up to which the signal is coherent to decrease. Consequently, events observed at positions other than the Cherenkov angle feature a lower frequency content compared

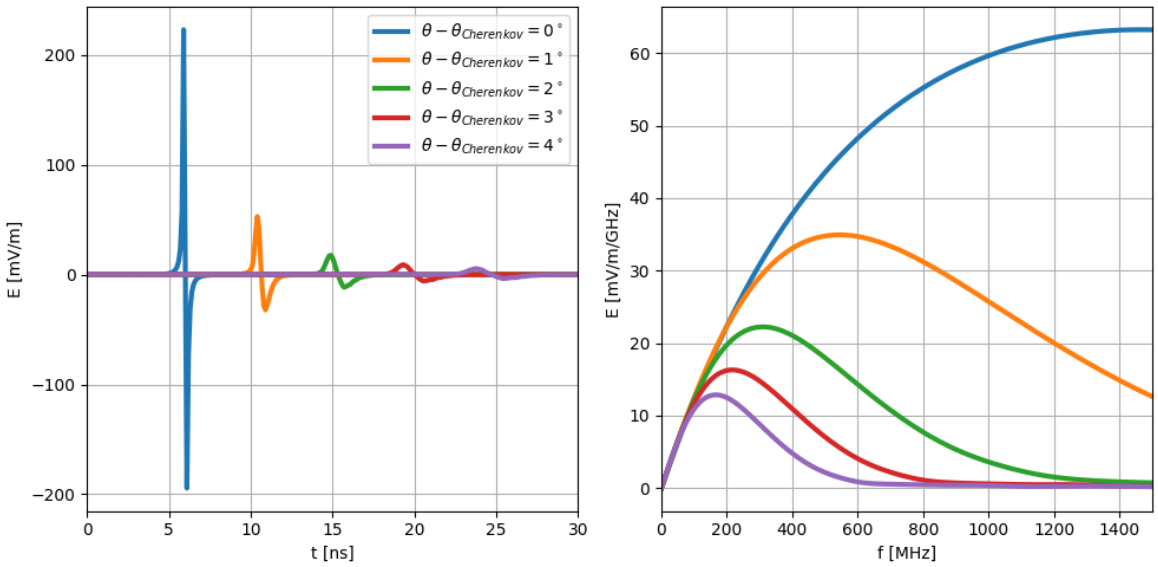


Figure 3.2: The simulated electric field produced by a neutrino interaction in ice (hadronic shower, 1 EeV energy). Shown are the time (left) and frequency (right) domains for different viewing angles relative to the Cherenkov angle. Taken from Ref. [14].

to events observed at the Cherenkov angle, as demonstrated in Figure 3.2 for in-ice Askaryan emission. Although the electric field is time-compressed at the Cherenkov angle, the signal width does not vanish, but it is determined by the shower front thickness for air showers and the lateral spread of the shower front for in-ice showers [117]. It is important to note that the above-discussed Cherenkov cone structure is a geometric effect arising from the speed of the particles relative to the speed of light in the medium. Only the structure and coherent frequency range of the signal are influenced by this effect. The emission mechanisms themselves are independent and would also produce coherent radio emission in a hypothetical medium with a refractive index of $n = 1$ [109].

Additionally, because the particle cascade travels faster than the speed of light in the medium, the arrival-time ordering of the detected signal is affected. For in-ice showers, if the observation angle is smaller than the Cherenkov angle, emission from the late stages of shower development is observed before emission from the initial stages, so the shower appears "bottom-to-top". In contrast, if the observation angle exceeds the Cherenkov angle, the arrival-time ordering matches the emission sequence, and the shower is seen "top-to-bottom" [117]. For air shower emission, the situation is more complex. While in-ice showers are observed under approximately the same angle throughout shower development, air showers have longer longitudinal dimensions, leading to a change in viewing angle during shower development. This leads to the signal from the later stages being "visible" for longer, thereby elongating the second part of the bipolar pulse, making it appear like a unipolar pulse [121].

Because the radio emission forms a forward-beamed Cherenkov cone structure, the emission from an air shower illuminates a so-called *footprint* on the ground. These radio footprints show an asymmetry in the field-strength distribution due to the different polarizations of the geomagnetic and Askaryan emission mechanisms, and if measured in the right frequency range, a visible Cherenkov cone (see Figure 3.3). The footprint size and shape depend on the zenith angle of the air shower, as it is a projection of the

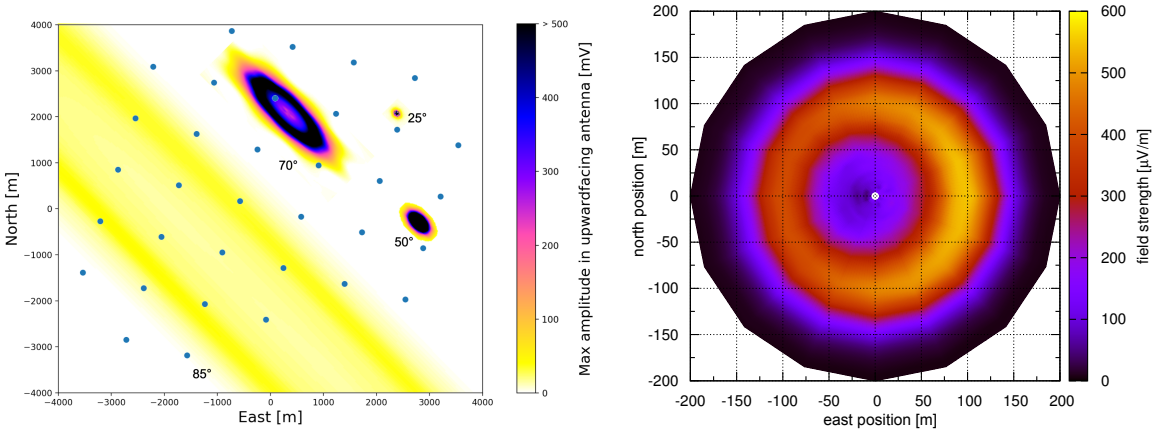


Figure 3.3: Left: The radio footprint of a particle shower with an energy of 3.2×10^{18} eV, simulated at different zenith angles. The dots indicated the potential layout of a detector, with each dot representing the location of antennas. Taken from Ref. [14]. Right: The simulated radio footprint of a vertical air shower with an energy of 10^{17} eV in the frequency band 300 − 1200 MHz. Taken from Ref. [122].

cone onto the ground. Figure 3.3 shows the radio footprint on the ground for four air showers with different zenith angles. The figure illustrates that detector layouts with large spacings between radio antennas are mainly sensitive to inclined air showers. If the footprint is smaller than the distance between the antennas, the probability of triggering an antenna decreases with the size of the footprint. Furthermore, for larger footprints, the radiated energy is distributed over a larger area, so the expected field strength at each point within the footprint is lower, as visible in Figure 3.3.

3.2 Radio detection of UHECR

Detection of ultra-high-energy cosmic rays (UHECRs, $E_{\mathrm{CR}} > 10^{16}$ eV) using the radio technique is done by measuring the radio emission produced by EAS following the interaction of UHECRs with the atmosphere. Multiple radio antennas are distributed over a large area to increase the effective area of the detector and to sample different regions of the radio footprint. The energy threshold for this technique typically ranges from 10^{16} to 10^{18} eV. Key factors influencing the energy threshold include the frequency band of the detector, the antenna array density, and the strength and orientation of Earth's magnetic field [28]. Furthermore, the spacing between individual antennas determines the range of zenith angles to which the detector is sensitive, as the zenith angle affects the size of the radio footprint (see Section 3.1.3).

Radio detection of UHECRs is an active research field, with multiple operational detectors and additional facilities under development. The Auger Engineering Radio Array (AERA), located at the site of the Pierre Auger Observatory, covers approximately 17 km^2 and uses over 150 antennas sensitive in the 30 − 80 MHz frequency band. A key advantage of the AERA site is the ability to cross-calibrate radio measurements with the already established fluorescence (FD) and surface particle (SD) detectors from the Pierre Auger Observatory. Further details about AERA are available in Ref. [60]. Also, AERA paved the way for the radio upgrade of the Pierre Auger Observatory. The Radio

Detector (RD), installed as part of the Auger Prime upgrade and completed in 2024, equips each water-Cherenkov detector with a radio antenna sensitive to the $30 - 80\,\mathrm{MHz}$ frequency range. As a result, the Radio Detector comprises 1660 antennas and covers an area of $3000\,\mathrm{km}^2$ [123].

Another radio instrument is the LOw-Frequency ARray (LOFAR), a radio interferometer with more than 40 stations across Europe, mainly designed for astronomical observations [124]. The dense core of radio antennas in the Netherlands is used for the detection of UHECR air showers, providing data from up to 1000 antennas per recorded air shower. LOFAR uses two types of antennas with different bandwidths: low-band ($10 - 90\,\mathrm{MHz}$[1]) and high-band ($110 - 240\,\mathrm{MHz}$) antennas [125].

Furthermore, LOFAR serves as a pathfinder for the planned Square Kilometer Array (SKA). The low-frequency component of SKA is under construction and, upon completion, will feature an unprecedentedly dense core of nearly $60,000$ antennas ($50 - 350\,\mathrm{MHz}$). SKA is expected to enable the most precise air-shower reconstructions achieved with the radio technique [126].

Lastly, we mention the Tunka Radio extension (Tunka-Rex) experiment, which is the radio extension of the already existing Tunka air-shower experiment, which employs air-Cherenkov detectors and scintillators. Tunka-Rex consists of 44 antenna stations covering approximately $3\,\mathrm{km}^2$, with a denser core and sparser outer stations and is sensitive in the frequency range $30 - 80\,\mathrm{MHz}$ [127].

Radio air-shower detectors have demonstrated strong agreement between measured radio signals and theoretical predictions, including Monte Carlo simulation codes such as CoREAS [122] and ZHAireS [128]. LOFAR has confirmed that the observed polarization aligns with predictions based on the combined geomagnetic and Askaryan effect [129]. Furthermore, theoretical models predict a circularly polarized component in the radio signal due to differences in the effective emission heights and the resulting time delays between these two mechanisms, and LOFAR has confirmed the existence of the circular component [119]. Another predicted feature is the strong time compression of the radio signal along the Cherenkov cone, a feature experimentally confirmed by LOFAR [130]. Regarding the energy scale, the Pierre Auger Observatory demonstrated that the radiation energy from air showers scales quadratically with the energy of the primary particle, which was determined by the surface detectors, consistent with expectations for coherent emission [120]. Also, the measured radiation energy agrees with first-principles Monte-Carlo simulations [120]. In a different study, the Pierre Auger Observatory demonstrated that the radiated energy of an air shower can be related to the energy of the initiating UHECR, and that the radio energy scale is compatible with the well-established fluorescence detector energy scale within systematic uncertainties [131]. Both LOFAR and the Pierre Auger Observatory have demonstrated that radio emission enables precise reconstruction of the atmospheric depth of the shower maximum (X_{max}), achieving precision competitive with established techniques such as fluorescence detection [132, 133]. Additionally, simulated emission profiles fit ground-intensity measurements well, indicating a robust understanding of the emission mechanisms [134]. Finally, the Pierre Auger Observatory has shown that X_{max} measurements from radio

[1] The frequency band can be restricted to $30 - 80\,\mathrm{MHz}$ using filters, which is done for most observations to suppress strong noise outside the frequency band [125].

emission are compatible statistically and on an event-by-event basis with those from fluorescence detectors, further validating the reconstruction method [135]. In conclusion, the experimental results demonstrate that the radio emission mechanisms for air showers are well understood, that simulations and measurements agree well, and that the radio technique has been established as a competitive and complementary tool for measuring EAS from UHECR. Moreover, the knowledge gained from air-shower experiments paved the way for the use of the radio detection technique in neutrino telescopes.

3.3 Radio detection of UHEN

Due to the low flux of neutrinos with energies beyond 10 PeV, no neutrino other than the exceptional event detected by KM3NeT [5] has been observed. Experimental approaches to change this using radio neutrino telescopes can be divided into two categories: in-ice and in-air experiments.

In-ice: Generally, the concept of in-ice experiments is to deploy radio antennas in the ice to measure the radio signal created by the Askaryan effect when a UHEN interacts in the ice [14, 136]. The large attenuation length of ice, such as that in Greenland [9], combined with the relatively low costs of radio antennas, enables the construction of sufficiently large detectors to probe the ultra-high-energy regime cost-effectively. The first operational in-ice neutrino detector was the Radio Ice Cherenkov Experiment (RICE) [137], deployed at the South Pole in 1999 and operated until 2010 [110].

Building on the expertise gained with RICE, the Askaryan Radio Array (ARA) began deployment at the South Pole in 2010 [138]. For ARA, the antennas are deployed in boreholes to depths of 200 m, and their design is driven by the spatial dimensions of the borehole. Two custom-designed antennas are used: one for measuring the horizontal and one for measuring the vertical component of the electromagnetic field [138]. Furthermore, the ARA collaboration developed a phased-array trigger[2] for in-ice neutrino experiments [139], which is also implemented in the next generation of in-ice radio neutrino detectors. ARA has set upper limits on the diffuse neutrino flux [12, 140], as shown in Figure 3.4, and recently reported the first observation of Askaryan radiation in naturally occurring ice. This radiation was measured from partially developed air showers produced by UHECR that penetrated the ice and further developed there [22].

A different approach was pursued by the Antarctic Ross Ice-Shelf ANtenna Neutrino Array (ARIANNA), which began deployment in 2010 at Moore's Bay, Antarctica [99]. For ARIANNA, the antennas are buried in shallow trenches, removing restrictions on the antenna size and allowing the use of off-the-shelf high-gain Log-Periodic Dipole Arrays (LPDAs) with broad bandwidths [13]. The remote location necessitates powering the experiment with renewable energy, namely solar and wind power, which has influenced the detector design to minimize power consumption. On the other hand, the site provides a very low background, and the ice/water boundary of the Ross Ice Shelf enhances the field of view due to the possible detection of reflected signals [13]. ARIANNA has

[2] A trigger that coherently sums signals from multiple antennas to enhance the sensitivity.

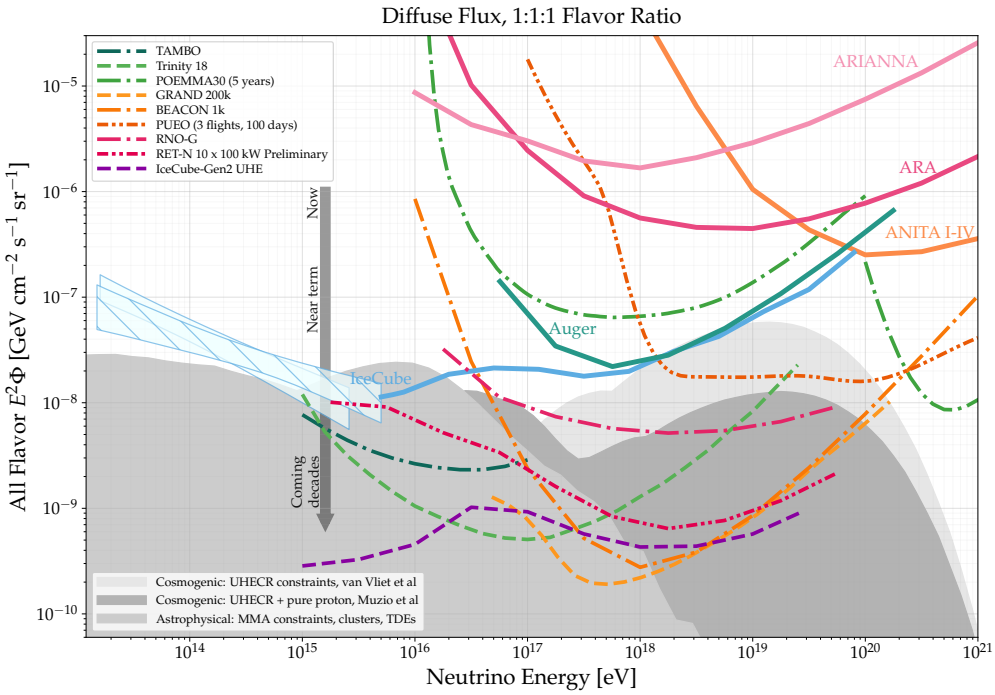


Figure 3.4: The expected differential sensitivities at 90% C.L. for various experiments to an all-flavour diffuse neutrino flux. Unless otherwise specified, an operational lifetime of 10 years is assumed for each experiment. Furthermore, IceCube flux measurements are shown as blue bands, alongside experimental upper limits from IceCube, the Pierre Auger Observatory, ARA, ARIANNA, and ANITA. The figure is taken from Ref.[141], and the sources of the presented data are provided in the original publication. Additional experiments for which sensitivities are presented, but which are not mentioned in the text due to their use of non-radio detection methods, include TAMBO (particle detectors) [142], POEMMA (fluorescence and Cherenkov light) [143], Trinity 18 (Cherenkov light) [144] and RET-N (radar technique) [105].

published upper limits on the diffuse neutrino flux [13], as shown in Figure 3.4, and has detected radio emission from UHECR air showers, including measurements of their polarization [17, 145]. In addition to the Moore's Bay stations, two ARIANNA stations have been installed at the South Pole [13].

Both the ARA and ARIANNA approaches have successfully demonstrated the radio-detection technique and are now combined in the Radio Neutrino Observatory in Greenland (RNO-G), which is currently under construction on the Greenlandic ice sheet. Upon completion, RNO-G will be capable of probing the most optimistic UHEN models and is expected to have discovery potential. The projected sensitivity for a diffuse neutrino flux is presented in Figure 3.4. As will be presented in this thesis, RNO-G can also detect UHECR. Further details about the detector are provided in Chapter 4. Looking ahead, an even larger in-ice radio-neutrino detector is planned as part of the IceCube-Gen2 Neutrino Observatory [15], for which RNO-G serves as a pathfinder. The expected sensitivity for both the optical and radio components of IceCube-Gen2 in the ultra-high-energy regime is shown in Figure 3.4.

In-air: The general approach of the in-air radio neutrino detectors is to detect the radio emission from so-called *earth-skimming* neutrinos, which are mainly ultra-high-energy tau neutrinos that interact in the Earth's crust and produce a tau lepton. If this interaction occurs near the surface, the tau lepton can escape into the atmosphere and decay,

generating a particle cascade and associated radio emission. This technique is primarily sensitive to tau neutrinos, as tau leptons have lifetimes sufficient to exit the Earth but still decay within the detection volume of the potential radio detector [146]. Experiments must either observe neutrinos from the horizon, due to the opacity of the Earth at these energies, or face a mountain ridge where the neutrinos can interact [99].

An experiment searching for earth-skimming neutrinos and UHECRs is the Taiwan Astroparticle Radiowave Observatory for Geo-synchrotron Emissions (TAROGE), with antenna arrays located on a high mountain in Taiwan [147] and in Antarctica [148] that observe the horizon. At both sites, the collaboration has reported the detection of UHECR [147, 148], but an UHEN search has not yet been performed.

Further experiments that are proposed and already operate prototypes are the Beamforming Elevated Array for COsmic Neutrinos (BEACON) [149, 150] and the Giant Radio Array for Neutrino Detection (GRAND) [146, 151]. For BEACON, the antennas would be placed on a high mountain to observe the horizon using beamforming, thereby increasing the sensitivity [149]. For GRAND, it is planned to deploy a large array of antennas on a mountain slope facing another mountain several kilometers away [146]. Both collaborations have reported the detection of UHECR air shower events with their prototypes [150, 151], and their projected sensitivities to a diffuse neutrino flux are shown in Figure 3.4. Recently, BEACON and GRAND have proposed a joint effort to construct the Hybrid Elevated Radio Observatory for Neutrinos (HERON), which would combine both experimental approaches [152].

A different strategy for detecting earth-skimming neutrinos is employed by the Antarctic Impulsive Transient Antenna (ANITA) [153] and its successor, the Payload for Ultrahigh Energy Observations (PUEO) [154]. Both experiments are balloon-borne and survey the Antarctic continent. ANITA has successfully completed four flights [155], while PUEO has recently completed its first. In addition to the tau-lepton-induced air showers, ANITA and PUEO are sensitive to all-flavour neutrino interactions in the Antarctic ice, where the resulting radio emission is refracted out of the ice and detected in the air [153]. A limitation, in addition to the finite lifetime of the balloon-borne detectors, is the large distance to the interaction vertex of possible neutrino interactions, which increases the energy threshold because the signal strength decreases with distance [14, 110]. At the highest energies, the four ANITA flights have set the most competitive limit on the diffuse neutrino flux (see Figure 3.4). The expected sensitivity for multiple PUEO flights is also shown in Figure 3.4. Notably, ANITA has observed so-called *mystery events* [156–158], which are consistent with upward-going showers from tau leptons. However, the directionality of these events creates a significant tension with the standard model cross-section, as tau neutrinos should have been absorbed during propagation through the Earth unless physics beyond the standard model is introduced [159]. The origin of these events has not been resolved.

4 Radio Neutrino Observatory in Greenland

The Radio Neutrino Observatory in Greenland (RNO-G) is an in-ice radio neutrino detector currently under construction and is designed to be sensitive to UHEN with energies above 10 PeV [14]. RNO-G profits from the experience made with previous and ongoing in-ice radio neutrino experiments such as RICE [137], ARA [138], and ARIANNA [13] (see Section 3.3). The detector is located on the Greenlandic ice sheet at an elevation of approximately 3,200 m near Summit Station. Figure 4.1 shows the location of the detector on a map of Greenland. Upon completion, the detector will consist of 35 independently operating stations, resulting in an effective volume of approximately 100 km^3 and covering an area of around 50 km^2. Therefore, RNO-G will be the largest in-ice radio UHEN detector to date [14]. As detection medium, naturally occurring ice is used by instrumenting it with radio antennas, exploiting the large attenuation length of the Greenlandic ice ($\mathcal{O}$(1 km) [9]). Currently, eight stations (labelled with two-digit numbers) are deployed and operational. Figure 4.1 displays the relative location of both deployed and planned stations.

Each station operates as an independent detector, allowing science data to be collected with the already-deployed stations during ongoing construction. Additionally, the stations are powered exclusively by renewable energy and do not depend on a central power plant. The primary power source consists of two solar panels per station, which charge batteries to enable continuous operation during the Arctic summer and shoulder seasons. This approach eliminates the need for power lines, which would bring a logistical overhead for the remote stations. However, exclusive reliance on solar panels limits the stations' lifetimes and excludes data taking during the Arctic winter. To mitigate this limitation, some RNO-G stations are equipped with a wind turbine intended to extend operational periods into the shoulder seasons[1] and potentially during windy intervals in the Arctic winter. The wind turbines are currently in a testing phase, with different models being evaluated for performance under the harsh winter conditions [161, 162]. As of now, five of the eight operating stations are equipped with a wind turbine, although during the period relevant for the analysis presented in this work (2022 and 2023), only stations 11 and 12 were equipped with an early prototype wind turbine. During the Arctic winter, when solar power is unavailable, the stations enter a lower-power hibernation mode, during which no data-taking is possible. Under normal conditions, communication with the stations happens via a private LTE network; in low-power mode, it switches to the Long Range Wide Area Network (LoRaWAN) due to the higher power consumption of the LTE modem [163]. In general, the limited power budget per station drives the design choices made for electronic components.

To increase the effective volume of the detector, the layout was chosen such that the overlap between the effective volumes[2] of the individual stations is minimal by spacing them 1.25 km apart. For UHEN with energies above 1 EeV, the expected coincidence

[1] The time period between the Arctic summer and winter, when the sun rises for only short periods each day.

[2] The effective volume in the context of RNO-G describes the volume in which a UHEN can interact so that the detector is triggering on the event.

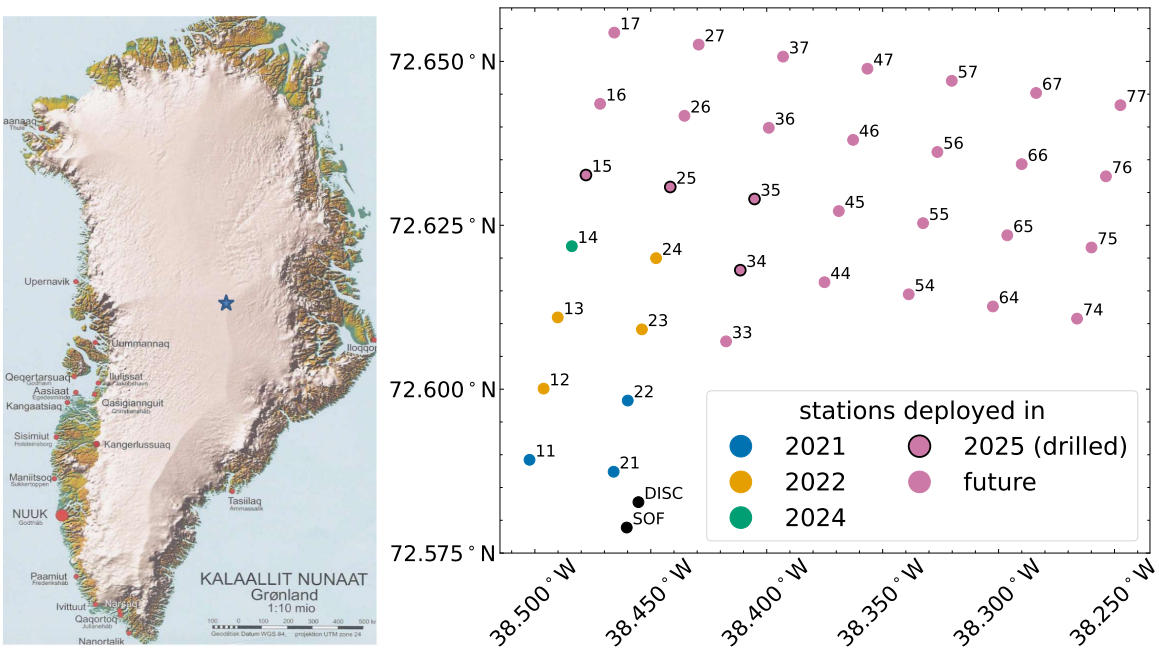


Figure 4.1: Left: Map of Greenland with the location of RNO-G marked by a blue star. The figure is provided by the RNO-G Collaboration. Right: Map of the deployed and planned RNO-G stations. Also shown are non-station objects or locations. The Station Operation Facility (SOF) is the main building at Summit Station, which houses the office of the site manager, the kitchen, and a community space, and is also called "Big House". The Deep Ice Sheet Coring (DISC) [160] is a borehole close to Summit Station, which is used as the origin for the local RNO-G coordinate system.

fraction[3] is approximately 10%, and increases with energy [14]. Because it is most likely that a neutrino will be detected in only a single station, RNO-G is designed to reconstruct the direction of a potential UHEN when it is observed by only a single station [121, 164]. This requirement informed the station layout, as shown in Figure 4.2.

Each station consists of a deep and a shallow component. The deep component includes custom-built antennas (see Section 4.1) that are deployed in three boreholes to a depth of 100 m, serving as the primary part of the instrument for the UHEN detection. Most deep antennas are arranged along the so-called *power string*, with a cluster at the bottom. Four antennas sensitive to the vertical polarization of the signal (Vpol) at the bottom of the string form the so-called *phased array*, which can use a specialized trigger to enhance the sensitivity of the detector (see Section 4.3). Additionally, two antennas sensitive to the horizontal polarization of the signal (Hpol) are positioned directly above the phased array. These six antennas are closely spaced and basically sample the same electric field, which helps with the full reconstruction of the electric field. Furthermore, antennas placed at shallower depths on the string improve the reconstruction of the distance to the interaction vertex [14, 164]. The antennas on the two additional strings, known as *helper-strings*, are essential for accurate azimuthal signal reconstruction. Additionally, two calibration pulsers are co-deployed with the antennas to enable in situ antenna position calibration [165].

The shallow antennas are commercially available Log-Periodic Dipole Arrays (LPDAs) (see Section 4.1), installed in hand-dug trenches near the surface. The shallow component consists of 9 LPDAs, six of which face downwards. These downward-facing antennas

[3] The probability of detecting a neutrino in multiple stations.

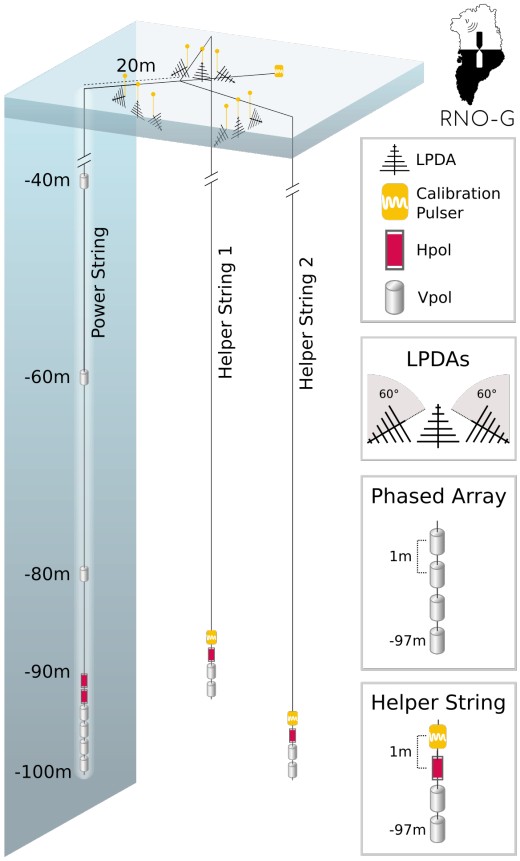


Figure 4.2: Schematic of the first 7 RNO-G stations. 24 antennas of three types (LPDA, Hpol, and Vpol) are distributed in the shallow ice and in three boreholes. Figure from the RNO-G Collaboration.

are oriented so that they have the highest sensitivity to signals from within the ice and thus contribute to the effective volume for UHEN. Moreover, their broadband sensitivity improves the energy and pointing resolution for events detected with these and the deep antennas. The remaining three LPDAs face upwards and are used to detect and veto backgrounds such as UHECR air showers. Consequently, these upward-facing antennas, channels[4] 13, 16, and 19, are the most relevant for the UHECR analysis presented in this work. The newest station (14) features a revised shallow geometry, where only eight LPDAs - four upward and four downward-facing - and an additional dipole antenna at shallow depth are deployed. Studies have shown that this new geometry improves the reconstruction of UHECR events [121], and the design will thus be applied to future stations.

A summary of the detector performance is provided in Ref. [163].

4.1 Radio antennas

Radio antennas convert the electric field produced by an air shower from UHECR or an in-ice shower from UHEN interactions into a measurable voltage waveform. This conversion is characterized by the vector effective length (VEL), which describes the

[4] Each antenna and its associated signal chain is referred to as a channel, numbered from 0 to 23.

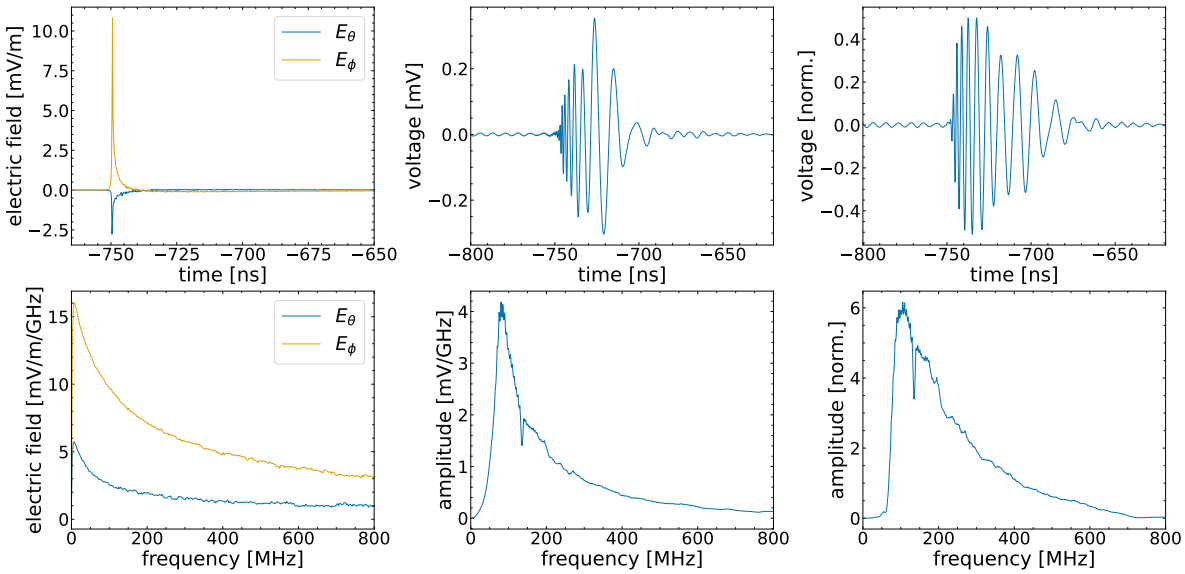


Figure 4.3: Visualization of how the antenna, amplification, and digitization components influence the pulse shape. Left: Time and frequency domain of a simulated electric field from a UHECR air shower. The air shower is simulated with an energy of $E = 10^{18}$ eV and a zenith angle of 55°. Notably, the second part of the characteristic bipolar pulse is barely visible. Middle: Voltage waveform and the corresponding frequency spectrum obtained by applying the antenna response to the electric field. Right: Voltage waveform and corresponding frequency spectrum after the amplification and digitization responses are applied to the signal.

induced voltage in an antenna when subjected to an electric field [166]. Figure 4.3 presents a simulated example electric field from a UHECR air shower, along with the corresponding voltage waveform recorded with an LPDA. The electric field from UHECR air-showers is intrinsically very short in time and has a bipolar structure [122]. However, the antenna response disperses this signal and consequently, the resulting voltage waveform is a pulse with an oscillatory pattern, featuring high-frequency oscillations at the beginning and lower-frequency oscillations toward the end, as shown in the central plots of Figure 4.3. Subsequent amplification and digitization components further disperse the signal (see Section 4.2).

The antennas for RNO-G are installed in ice, and particularly in the upper layer where the LPDAs are located, uncertainties in the index of refraction are significant. The VEL varies for media with different refractive indices, which can be understood with a simple dipole. A dipole is most sensitive when the incoming electric field has a wavelength λ that is a multiple of half the length of the dipole [166], independent of the medium it is placed in. However, while the sensitive wavelength remains unchanged, the propagation speed of the electric field is changed from c, the speed of light in a vacuum, to c/n with n being the index of refraction of the medium. Consequently, the frequency corresponding to the most sensitive wavelength changes.

$$f_{\text{res, med}} = \frac{c_{\text{med}}}{\lambda_{\text{res, med}}} = \frac{c}{n\lambda_{\text{res, med}}} = \frac{c}{n\lambda_{\text{res}}} = \frac{1}{n} f_{\text{res}} \tag{4.1}$$

Here, f denotes the frequency, λ the wavelength, the subscript "res" indicates the resonant frequency or wavelength for which a dipole is most sensitive, and the subscript "med" refers to quantities in the medium, while all other quantities are in vacuum. Therefore, the VEL of the antennas is expected to change by a factor $1/n$ in a medium.

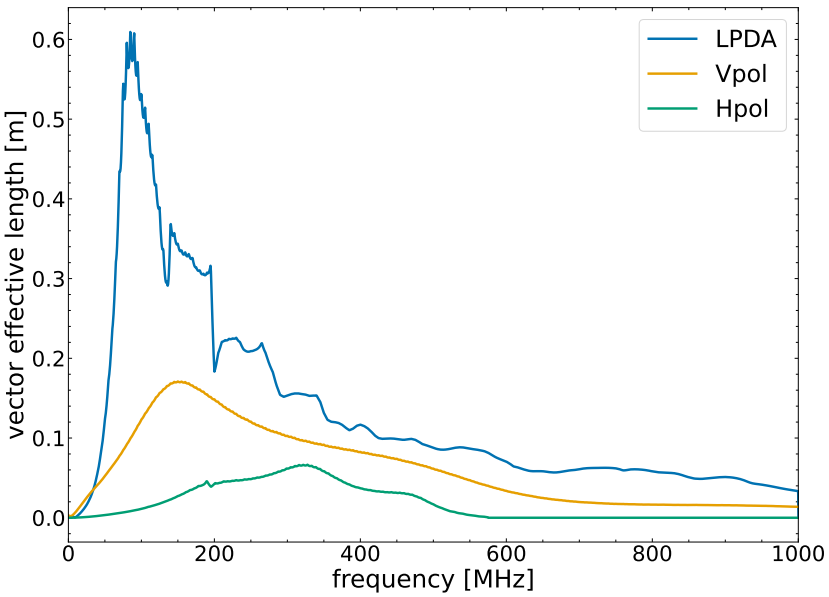


Figure 4.4: The simulated VEL for the LPDA, Hpol, and Vpol antennas. For each antenna, the co-polarized component of the VEL is shown at boresight. For the LPDA, boresight corresponds to a zenith angle of 0°, whereas for the Hpol and Vpol antennas, it is at 90°. The simulations are performed in media with refractive indices of $n = 1.3$ for the LPDA, $n = 1.73$ for the Vpol, and $n = 1.74$ for the Hpol.

This relationship demonstrates that an uncertainty in the index of refraction of the surrounding medium alters the VEL and, consequently, affects the final signal shape. Because the analysis presented in this work is sensitive to the recorded waveform shapes, we will examine the refractive index of the medium surrounding the LPDA and its impact on the analysis in detail in Chapter 6. In the following, we will introduce the three antennas used for RNO-G. Further details and an in-depth discussion of the antennas can be found in Ref. [163].

4.1.1 Log-Periodic Dipole Array (LPDA)

The Log-Periodic Dipole Array (LPDA) is used as the antenna for the shallow component of the detector. Because the antennas are buried directly beneath the surface, their dimensions are not constrained by a borehole, unlike the deep antennas. Consequently, a commercially available antenna from Create [167] (CLP-5130-2N) was chosen, which was previously successfully operated by the ARIANNA experiment [168]. Figure 4.5 shows a photograph of such an LPDA. For the shallow part of the detector, nine LPDAs were previously deployed, but this has now changed to eight LPDAs for the newest and future stations. Three antennas are oriented upwards, while the remaining antennas are directed downwards at approximately 60°. The LPDA measures 1.4 m in length along its spine and 1.45 m in width at its broadest point [121]. When buried in the firn[5], the LPDA is sensitive across the frequency range of 80 to 1000 MHz. The co-polarized component of the VEL for a simulated LPDA at boresight is shown in Figure 4.4. The VEL shown in the figure demonstrates that the LPDA achieves maximum sensitivity at lower frequencies (around 100 MHz), with a steep decline that flattens off towards higher frequencies. As

[5] Name for the upper layer of the ice and snow. For more details see Chapter 6.



Figure 4.5: Images of the science antennas used by RNO-G. From left to right, the antennas depicted are the LPDA (Credit: C. Welling), Vpol, and Hpol.

illustrated in Figure 4.3, the VEL profile will be imprinted on the frequency spectrum of UHECR or UHEN events that may be recorded.

4.1.2 Vpol

The Vpol antennas are constructed to detect the vertically polarized component of the electric field. Their design is constrained by the limited space available within the borehole, and thus, they use a fat-dipole design and are custom-built by the Collaboration. They have a diameter of 12.7 cm and a length of 60 cm. Figure 4.5 shows a picture of a Vpol antenna, while Figure 4.4 displays the simulated VEL. The antenna is sensitive across a frequency range of $\sim 140 - 600$ MHz, and the design ensures rotational symmetry. Nylon caps at both ends facilitate mounting the front-end amplifier (Iglu board) and keep the antenna centered in the borehole, preventing its active element from contacting the borehole walls. A contact with the wall would alter the antenna response and compromise its rotational symmetry. Vpol antennas are the primary science antennas for detecting UHEN and are used for triggering with the deep section of the detector. Moreover, the same antenna design is used for calibration pulsers that are deployed alongside the science antennas. When used as a calibration pulser, the antenna is connected to specialized front-end amplifiers optimized for transmission rather than reception.

4.1.3 Hpol

The Hpol antennas are designed to measure the horizontal component of the electric field. In a classical dipole design, the dipole tines are oriented along the electric field component to be measured. However, since the antennas are deployed inside the borehole, it is not possible to orient dipole tines horizontally for the frequency range of interest, as the required tine length would exceed the borehole diameter. Instead, a slot antenna is used, which couples to the magnetic component of the incoming electromagnetic field and enables the measurement of the horizontal electric component via Booker's extension to Babinet's principle [169]. To maintain rotational symmetry, the antenna

design incorporates four slots. These custom-built antennas were initially developed for the ARA experiment [138] and are subsequently adjusted and further developed for RNO-G. Figure 4.5 shows an image of an Hpol antenna, and Figure 4.4 shows the VEL. The antennas have a diameter of 20.3 cm and a length of 60 cm. When deployed in the ice, their sensitive frequency band ranges from 300 to 850 MHz. Although the antennas are not currently used for triggering, they play a crucial role in the full reconstruction of the electric field. Similar to the Vpol antennas, the Hpol antennas are equipped with nylon caps at both ends to house the front-end amplifier and maintain a central position within the borehole.

4.2 Amplification and digitization

Once the electric field is converted into a voltage waveform by the antenna, the signal is amplified and digitized. Already the signal at the antenna output is no longer a bipolar pulse (see Figure 4.3), and it becomes even more dispersed due to the amplification and digitization. Figure 4.3 shows the final digitized waveform, demonstrating that both the antenna response and the amplification and digitization response affect the waveform shape. Because the UHECR analysis presented in this work is sensitive to the signal shape, a comprehensive understanding of these effects is essential. Therefore, we will provide a detailed discussion of the amplification and digitization response of the shallow component of the detector in Chapter 5. In the following, we introduce the different components comprising the amplification and digitization chain.

4.2.1 Deep components

As previously described, the antennas are positioned in boreholes at depths reaching up to 100 m below the surface. To minimize signal loss during transmission from the antenna to the digitizer board (Radiant) and to maintain a manageable borehole string, the analog signal is converted to an optical signal and transmitted via fiber (Radio-frequency-over-fiber (RFoF)). Figure 4.6 shows a schematic of the signal chain. The first component after the antenna is the Iglu board, which connects via a short coaxial cable and incorporates a low-noise amplifier and a RFoF transmitter. The output of the Iglu board is connected via optical fibers to the Drab board at the surface within the data acquisition (DAQ) box. While the Iglu only contains a single channel, the Drab board accommodates four channels. Each Drab channel receives a signal from an Iglu board, converts it back to analog, and further amplifies and filters it. Two different signal paths follow after the Drab board, depending on the channel type: for channels not part of the phased-array, the signal is routed directly from the Drab board output to the Radiant channel input via a short coaxial cable; for channels part of phased-array, the signal is split after the Drab board, with a fraction directed to the Flower board[6] to form a trigger decision and the remainder continuing to the Radiant board. Further details are provided in Ref. [163, 170].

[6] The Flower board is an additional PCB that implements the trigger decision based on the deep component of the detector.

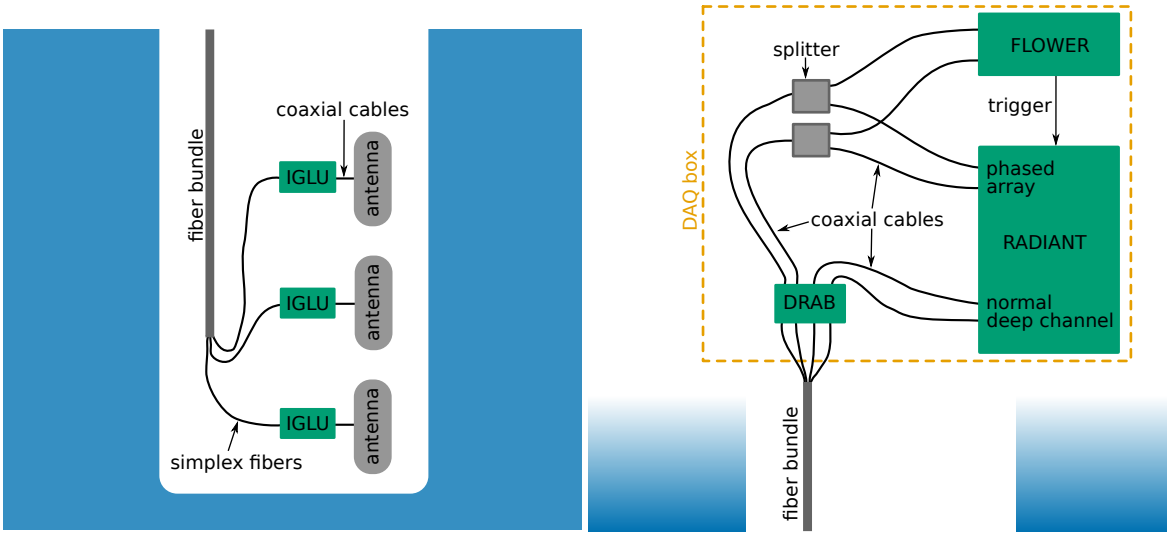


Figure 4.6: The amplification and digitization chain for the deep channels. Left: A schematic depiction of the part in the borehole. Right: A schematic depiction of the part at the surface.

4.2.2 Shallow components

The shallow antennas are deployed directly below the snow surface, enabling signal transmission to the RADIANT via a shorter coaxial cable without large power loss. Moreover, a fully analog system presents a lower risk of errors and is easier to calibrate in advance than an RFoF system. The shallow antennas connect to the SURFACE board, located within the DAQ box on the surface, using 10 m coaxial cables. The SURFACE board incorporates amplification and filtering stages and supports five channels. Signals from the SURFACE board are transmitted directly to the RADIANT input, where they are digitized. A schematic of the shallow chain is shown in Figure 4.7 and further details on the components are available in Refs. [163, 170]. This amplification and digitization chain is implemented in the first seven stations. However, the signal traverses a different cable length compared to the deep channels, which results in a mismatch in the signal arrival times at the RADIANT . This timing issue becomes problematic when a signal that triggers the shallow antennas should also be recorded in the deep antennas, or vice versa. These differences in signal run time can be addressed either by delaying the readout in software or by physically matching the travel times for both components. As a test, the latter approach was adopted for the most recently deployed station 14, where the coaxial cables and SURFACE board were replaced with an IGLU, fiber, and DRAB. The fiber was selected to match the length used for the deep channels, ensuring similar physical time delays in both the shallow and deep channels.

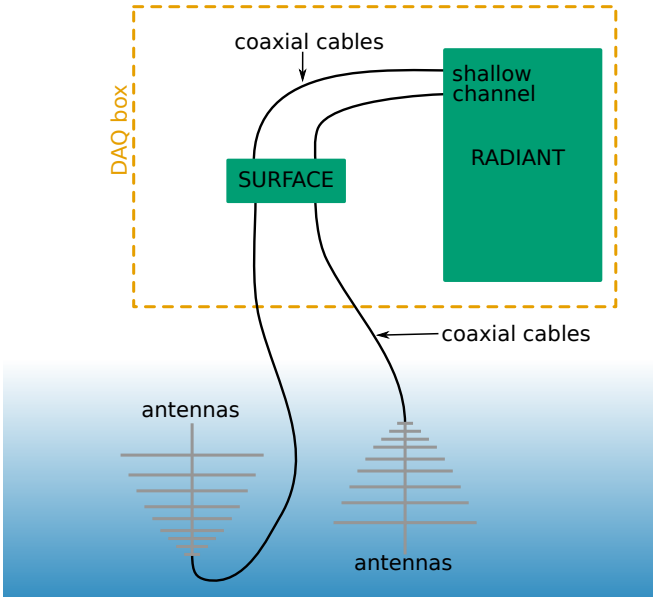


Figure 4.7: A schematic depiction of the shallow channel amplification and digitization chain.

4.3 Trigger

A trigger is a predefined condition that, when met by the signal, prompts the hardware to digitize and save the waveform from the channels. As the current RNO-G hardware is not designed to continuously digitize signals due to power limitations, it supports multiple independent triggers to run simultaneously at each station. In the following section, we will provide a brief overview of each trigger. A comprehensive discussion of the trigger relevant to the analysis presented in this work is presented in Section 5.3.

Forced trigger: This software trigger records waveforms at a specified regular interval when activated. During standard operation, the forced trigger records data at 0.1 Hz. The resulting waveforms primarily consist of thermal noise, as the trigger does not respond to specific waveform features. During periods of elevated noise, some noise events may also be recorded by the forced trigger. These recorded waveforms are used, for example, in absolute gain calibration of the detector [171].

RADIANT trigger: The RADIANT trigger decision is formed by a trigger circuit that is present on every channel, but it is used exclusively for the shallow channels and is therefore sometimes referred to as the shallow trigger. A detailed description of the trigger circuit is available in Refs. [163, 172] and will be further discussed in Section 5.3. To determine whether an event is saved, the signal passes through a diode, and the output must exceed a defined threshold. In science operations, this trigger is always used as a coincidence trigger, with up to two different triggers operable simultaneously. The number and selection of channels required to trigger within a given coincidence window can be specified and modified in software. Currently, the first trigger configuration requires two of the three upward-facing LPDAs to trigger within the coincidence window, which is also called the upward-facing trigger and is most sensitive to UHECR events. The second configuration requires two of the six downward-facing LPDAs to trigger within the coincidence window, also known as the downward-facing trigger. Due to

a firmware issue[7], it is sometimes not possible to determine whether the upward- or downward-facing trigger recorded the event. Consequently, the trigger rates cannot be separated, and only the combined rate will be shown in this work. Additionally, this limitation prevents selecting events based on only one of the two triggers. Furthermore, due to the trigger settings, the signal that triggered the waveform readout is always located in the second half of the recorded waveform.

Deep trigger: The deep trigger uses the four phased-array channels for triggering. The trigger decision is made on the Flower board, which was developed to operate the phased-array trigger (see next paragraph). Prior to the deployment of the advanced phased-array trigger, the Flower board was operated in a simpler mode, requiring each channel to fulfil a high-low threshold within a short time window [163]. Once a trigger decision is made on the Flower board, the Radiant is activated to read out all channels. This trigger was disabled after the phased-array trigger was deployed at the stations.

Phased-array trigger: The phased-array trigger utilizes the four channels corresponding to the Vpol antennas on the power string, which are closely packed and collectively referred to as the phased-array. Similar to the deep trigger, the trigger decision is made on the Flower board. This technique was previously demonstrated by the ARA experiment [139] and was first implemented at the RNO-G stations during the 2024 season [163]. For the phased-array trigger, the fact is utilized that all four antennas effectively measure the same electric field. The waveforms are shifted relative to each other based on travel-time differences across a range of incident angles, and then are coherently summed [14]. As a result, this process enhances the signal-to-noise ratio, thereby increasing the sensitivity of the trigger. Once the trigger decision is made on the Flower board, the Radiant is triggered to read out all channels.

Pulse-per-second trigger: The pulse-per-second (PPS) trigger is a software trigger that enables the readout of waveforms that are timed very precisely relative to the Global Positioning System (GPS) second. During standard operation, this trigger remains disabled.

When an event is triggered with any of the above-described trigger options, waveforms for all 24 channels are recorded. Additionally, during the time periods relevant for this work, the firmware only saves the name of the first trigger. This means, for example, that the upward- and downward-facing triggers are triggering, but only the name of the first one is saved. Moreover, all triggers, except software triggers, have dynamically adjusted trigger thresholds. This was implemented to keep the trigger rate relatively stable even as the noise level changes. For the two Radiant triggers, the adjustments are not applied to the coincidence trigger rate, but instead they are performed at the single-antenna level. As a result, the combined coincidence trigger rate can vary with the amount of correlated noise present, e.g., during high-wind periods [173].

[7] This firmware issue was resolved at the beginning of 2024. However, it remains relevant for this work, as only data collected before 2024 is used.

4.4 Backgrounds

Multiple background types can appear in the RNO-G dataset. An important background for UHEN searches is radio emission from UHECR air showers. This radio emission can refract into the ice and propagate to the radio antennas, thereby being a background to UHEN searches. Moreover, it was demonstrated that incompletely developed air showers may further develop after entering the ice, producing radio emission that can be an additional background source [19–21]. Finally, muons produced in air showers can penetrate the ice and lose energy stochastically, producing particle cascades that are effectively indistinguishable from UHEN signals [18]. Consequently, understanding and, when possible, vetoing signals from air showers is crucial. On the other hand, signals from UHECR air showers share similar characteristics with UHEN signals, as both generate electric fields on nanosecond time scales. These similarities allow the use of UHECR air shower signals to test the detector under realistic conditions, effectively serving as a unique "test beam". Furthermore, the upward-facing LPDAs in RNO-G provide a large effective area for UHECR detection, which enables to measure and study their properties. Additionally, the high elevation of Summit Station places the detector at a vertical depth of $\sim 700\,\mathrm{g/cm^2}$, which is near the shower maximum (X_{max}) for air showers with energies around 10^{17} eV [132]. As a result, RNO-G will be able to observe signatures of air-shower cores impacting the ice and developing particle cascades in the ice, producing Askaryan emission, as reported in Ref. [22].

A significant non-terrestrial background is diffuse Galactic radio emission from the Milky Way. Although the Galactic center is not within the RNO-G field of view, a contribution from the Galaxy has been identified in the noise spectrum [163, 174]. Other UHECR radio detectors have demonstrated that diffuse Galactic emission can be used for absolute gain calibration [175–177], with uncertainties discussed in Ref. [178]. RNO-G employs a similar calibration technique using thermal noise [171], where Galactic noise is considered but remains subdominant for most antennas and frequency ranges. Another non-terrestrial background is the radio emission of solar flares, which can trigger the detector [179]. Efforts are ongoing to incorporate solar flare events into a global antenna position calibration procedure [165]. Additionally, a non-anthropogenic but terrestrial background arises from events during high-wind periods [173]. Further research is underway to clarify the origin of these signals and their correlation with wind speed (see e.g. Ref. [162]). These impulsive background events can trigger the detector and present a significant background for the UHECR analysis presented in this work.

Furthermore, a variety of anthropogenic noise sources are present at Summit Station. Continuous wave (CW) or narrowband signals originate, e.g., from handheld radio devices, aircraft communications, and downlink transmissions from satellites and weather balloons [174]. Those signals can be filtered out using different techniques, and they can be utilized, for example, to determine the index of refraction of bulk ice as done with the transmission from a weather balloon [180]. Additionally, engines from machinery operating at Summit Station can also produce radio emissions, potentially contributing to the background if they are in proximity to the detector. Airplanes, both commercial and those approaching or landing at Summit Station, generate impulsive background events. These signals may originate from transmitters or engines, and it has been demonstrated that they can be used for antenna-position calibration and for

calibrating absolute timing between stations [181]. Other potential background sources are satellites in low Earth orbit. Multiple radio telescopes have reported detecting unintended emission originating from low-Earth-orbit satellites [182–184]. Studies are ongoing to determine whether RNO-G also detects such signals. Finally, additional background can be induced by the stations themselves. For instance, an early version of the power system and a prototype wind turbine caused an increased and highly variable noise level in the shallow antennas [185].

5 Shallow RNO-G Hardware and Trigger

In this chapter, we discuss the shallow amplification and digitization chain, as well as the RADIANT trigger. First, we investigate how different measurements of the amplification and digitization response affect the UHECR simulations and candidate events (see Section 5.1). Next, we introduce the structure and application of the RNO-G hardware and calibration database (see Section 5.2). We then evaluate the status of the RADIANT trigger and identify potential improvements for a redesign of the RADIANT board. Additionally, we describe a method to represent the trigger threshold in UHECR simulations as accurately as possible (see Section 5.3). Finally, we investigate an unusual trigger rate observed during the 2023 data-taking period at station 23 (see Section 5.4).

5.1 Shallow amplification and digitization response

The amplification and digitization chain contains all components that transmit, amplify, or digitize the signal from the antenna. For a shallow channel, this chain includes coaxial cables, the SURFACE board and the RADIANT board. Each component influences the signal by altering its power or introducing a time delay, and both these effects may vary in strength across different signal frequencies. If these changes were uniform across the entire frequency range, they would result in a global scaling factor of the signal amplitude and a uniform delay of the arrival time. However, frequency-dependent changes can significantly alter the signal shape. The way an incoming signal is modified by these components is referred to as the amplification and digitization response, or hardware response. While the antenna also significantly shapes the signal, as described by its VEL, we do not consider its influence in this section.

The hardware response is represented as a frequency-dependent complex number, with the absolute value indicating the change in power per frequency. The frequency-dependent delay, known as the group delay $\tau_g(f)$, is calculated as the derivative of the phase $\phi(f)$ of the complex hardware response, as shown in the following formula [168]:

$$\tau_g(f) = -\frac{1}{2\pi}\frac{\mathrm{d}\phi(f)}{\mathrm{d}f}. \tag{5.1}$$

A constant component of the group delay results in a time shift of the pulse and will be referred to as the time delay, while a frequency-dependent component causes dispersion and alters the pulse shape. In Figure 5.1, we illustrate the effect of the RNO-G hardware response on an incoming signal by presenting an example hardware response and the corresponding change in the time-domain signal of a delta-like pulse. We will discuss the shown hardware response, known as the impulse response, in greater detail in Section 5.1.2. The figure demonstrates that the group delay has a frequency-dependent component, indicating that lower frequencies propagate more slowly through the digitization and amplification components. Also, this effect is evident in the time-domain pulse after the hardware response is applied, where the first oscillations are faster than those at the end of the pulse. For the shown hardware response measurement, the

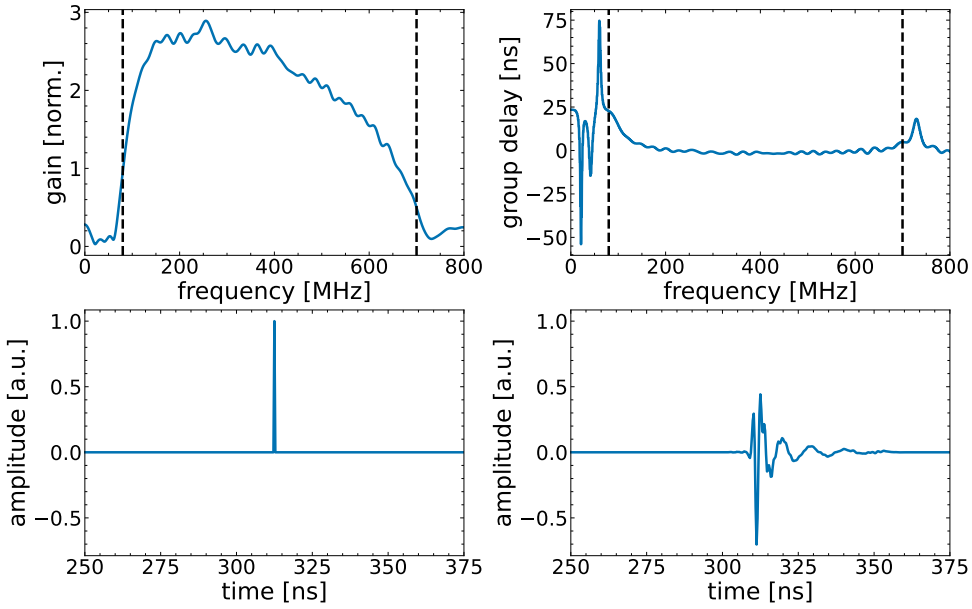


Figure 5.1: Top: The frequency-dependent gain (left) and group delay (right) of the hardware response, referred to as the impulse response, as described in Section 5.1.2. Vertical lines indicate the approximate frequency range over which the hardware is sensitive. The measurement was normalized in the time domain, so the gain does not reflect the actual system gain. Bottom: Illustration of how a delta-like pulse (left) is transformed when the hardware response from the top panels is applied (right).

constant time delay is normalized and therefore very small, as visible in the group delay plot and by comparing the starting position of both time-domain pulses.

A response of a component can be characterized by measuring its scattering parameters (S parameters), which can be done with a Vector Network Analyzer (VNA). S parameters describe the transmission and reflection of a signal in an electric circuit and can be defined for a device with n ports [186]. In this work, we consider only devices with two ports: an input and an output. For a two-port device, four S parameters can be measured, as illustrated in Figure 5.2, and are defined as follows [186]:

$$\begin{aligned} S_{11} &= \left.\frac{b_1}{a_1}\right|_{a_2=0}, \quad S_{12} = \left.\frac{b_1}{a_2}\right|_{a_1=0} \\ S_{21} &= \left.\frac{b_2}{a_1}\right|_{a_2=0}, \quad S_{22} = \left.\frac{b_2}{a_2}\right|_{a_1=0}, \end{aligned} \tag{5.2}$$

The signal fractions a and b are indicated in Figure 5.2. S parameters with identical indices describe a reflection, while those with different indices describe a transmission. In this work, we label the device input with 1 and the output with 2.

5.1.1 Influence on the shallow UHECR analysis

In the shallow cosmic-ray analysis presented in this thesis, we compare template waveforms with recorded waveforms to identify UHECR candidate events (for more information, see Chapter 8). As a result, the analysis is sensitive to variations in the waveform shape. To assess the impact of a slightly inaccurate hardware response

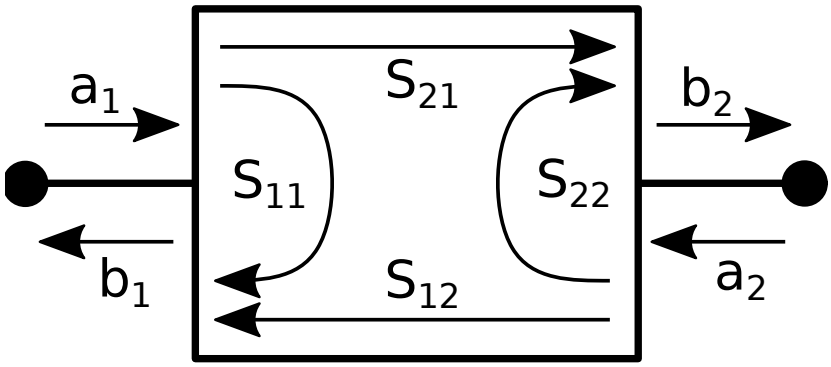


Figure 5.2: Illustration of the S parameter definition.

description on the analysis, we conduct UHECR simulations using the same simulation framework as the main analysis (details regarding the simulation setup are provided in Chapter 7). Two key aspects are relevant for this discussion. First, to create the template waveforms, we apply a hardware response to artificially generated voltage waveforms (a detailed description can be found in Section 8.2). Second, to create the simulated UHECR waveforms, we also apply the hardware response to a simulated voltage waveform. In both cases, the antenna response is already incorporated in the voltage waveforms. After generating the template waveforms and simulated UHECR waveforms, we correlate them, and the correlation value quantifies their similarity. A correlation value of 1 means the two waveforms are identical, and therefore, values close to 1 are optimal for the analysis.

To evaluate the potential impact of inaccurate hardware response descriptions, we perform two simulations. In one simulation, we use the same hardware response description for the template creation and event simulation, while in the other simulation, they differ. All other simulation parameters, including the noise representation, remain identical, and we use an example trigger threshold of 25 mV, corresponding to 9σ[1]. The first simulation serves as a baseline, while the second simulation represents the case where the templates are created with an inaccurate hardware response description. The hardware response descriptions that we use for this test are the *impulse response* and the *database response* (see Section 5.1.2 for detailed descriptions). To compare the two simulations, we calculate the SNR-dependent correlation score cut that we apply in the main analysis. For this cut, we divide the simulation set into multiple SNR bins and determine the 99% quantile for each bin (see Section 8.5 for additional information). The signal-to-noise ratio (SNR) is, in this work, defined as follows:

$$SNR = \frac{\max\left(|\mathrm{wf}|\right)}{\mathrm{RMS}_{\mathrm{noise}}} \tag{5.3}$$

Here, wf denotes the waveform and $\mathrm{RMS}_{\mathrm{noise}}$ is the noise RMS, calculated on the first half of the waveform. The resulting correlation score cuts from both simulations are presented in Figure 5.3, while a comparison of the response descriptions in terms of the gain and group delay is shown in Figure 5.4. When we compare the two hardware response descriptions, we observe differences in the shape of the normalized gain[2].

[1] In this case, σ is the RMS of the simulated noise in the trigger frequency band.
[2] Constant gain factors have no influence because the calculated correlation score is normalized.

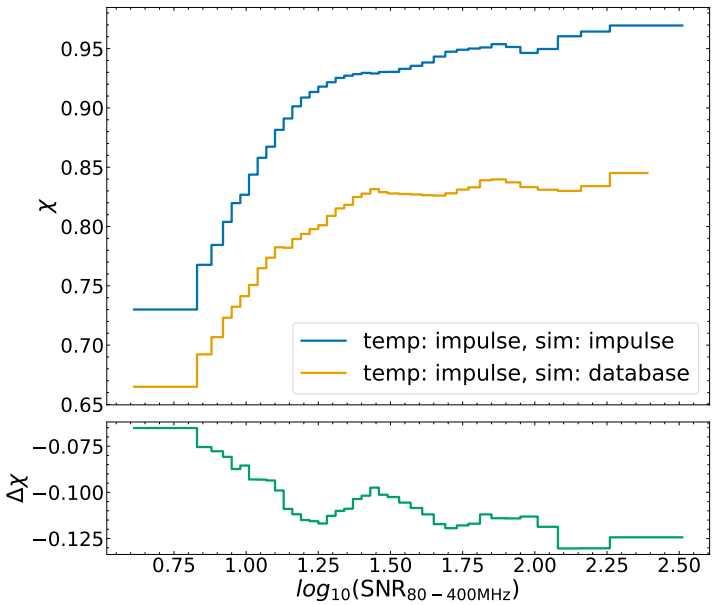


Figure 5.3: The results of simulations performed to investigate the impact of inaccurate hardware response descriptions on the analysis. For the two simulations shown, we used different hardware response descriptions, as specified in the figure. The bottom panel shows the difference between the two SNR-dependent correlation score cuts depicted in the upper panel. In the simulation, before calculating the correlation score and SNR, we apply a bandpass filter with a passband of 80 to 400 MHz to the waveforms, as is done in the UHECR analysis (see Chapter 8).

Additionally, a slight difference in the group delay at low frequencies is visible, although it is difficult to discern due to noise in the database response description. Overall, the differences between the two descriptions seem to be minor. Nevertheless, the simulation demonstrates that this difference can result in a significant change in the correlation score of $\mathcal{O}(0.1)$, as visible in Figure 5.3. Consequently, this finding underscores the importance of accurately characterizing the amplification and digitization response of the detector, particularly for analyses that depend on the waveform shape.

5.1.2 Hardware response descriptions

Three different descriptions of the shallow amplification and digitization response are available from measurements, and their gain and group delay are compared in Figure 5.4. In this figure, we can identify notable differences in both the gain pattern and group delay. The group delay varies in the magnitude of the excess at low frequencies, while for the gain pattern, one of the response descriptions, the *placeholder response*, exhibits a different functional form than the other two descriptions.

Furthermore, to assess the impact of these descriptions on the correlation score between the UHECR templates and candidate events, we generate templates for each hardware response description and correlate them with three reference UHECR candidate waveforms. We selected the three candidate events, which we identified in the main analysis presented in Chapter 8, for their high SNR value (> 30), since the noise influence on the correlation score is minimal at high SNR values. Additionally, all three candidate events exhibit the highest correlation score with the same antenna at the same station, allowing the influence of different antennas or channel-to-channel variations to be neglected. However, the candidate events originate from different zenith and azimuth

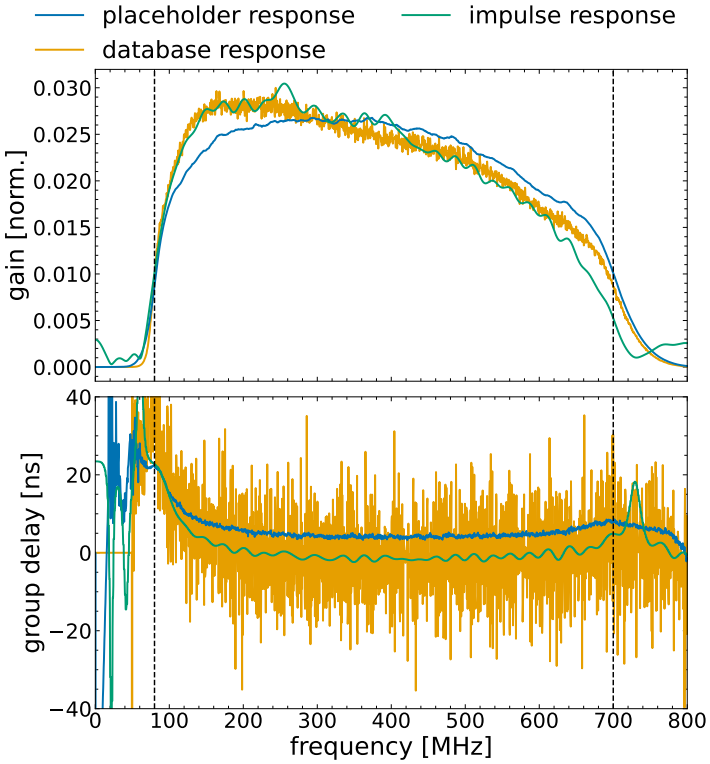


Figure 5.4: Comparison of the gain (top) and group delay (bottom) for different hardware response descriptions. Vertical lines indicate the approximate frequency range in which the hardware is sensitive. The gain pattern is normalized by the integral of the gain over the frequency range 80 − 800 MHz. The shown hardware response descriptions are introduced in Section 5.1.2.

angles, resulting in each being influenced differently by the antenna, as the VEL of the LPDA is not isotropic with respect to the arrival direction of the signal. For the correlation, we consider only the waveform from the channel that provided the highest correlation score when the UHECR candidate event was identified in the main analysis. The templates are created with the same parameters as for the main analysis, only the applied hardware response description is different. Although we generate multiple templates for each hardware response description, we consider only the template with the highest correlation score in this discussion.

In the following paragraphs, we provide a detailed introduction to each hardware response description and evaluate their respective impact on the correlation score with the three reference UHECR candidate events.

Placeholder response: Prior to the availability of more accurate hardware response descriptions, so-called *placeholder* measurements were used in the detector simulation and reconstruction framework, NuRadioReco, by RNO-G to characterize the response. For the shallow channels, this placeholder response is based on an S21 parameter measurement of a Surface amplifier board, whereas for the deep channels, it relies on an S21 parameter measurement of a Drab board, an optical fiber, and an Iglu board combination. As implied by the name, this measurement served as an interim solution and was not intended to provide a comprehensive description.

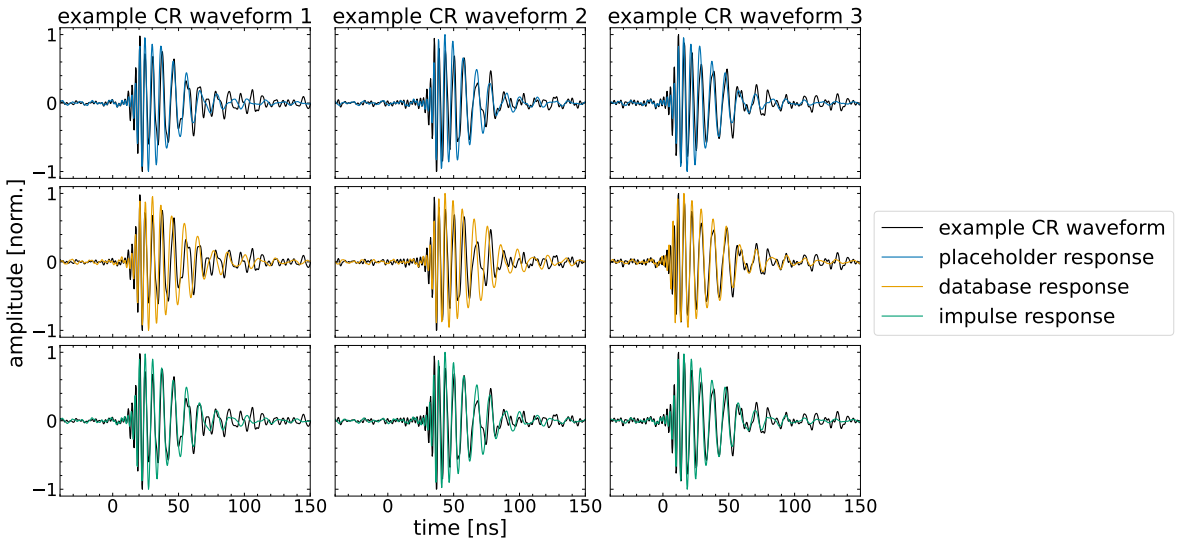


Figure 5.5: Comparison of the template waveforms generated using different hardware response descriptions with the reference cosmic-ray (CR) candidate waveforms. Each column presents the same reference CR candidate waveform, with different colors representing the different hardware response descriptions. All waveforms are normalized to their maximum absolute value.

Figure 5.4 shows that the placeholder response exhibits a different gain pattern compared to the other two responses, likely due to the omission of certain components, such as the RADIANT board. The group delay for this response shows the smallest excess at low frequencies, which may also be due to an incomplete component representation. In Figure 5.5, we compare the templates generated with the placeholder response to the three reference UHECR candidate waveforms. The template waveform aligns reasonably well with the UHECR candidate waveforms, although the most significant discrepancies occur at the beginning of the pulse, where the amplitude is insufficient, and in the middle, where the oscillations are slightly misaligned. The corresponding correlation scores for all three events are shown in Table 5.1. Simulations indicate that UHECR events with such high SNR values should yield correlation scores above 0.9, approaching 1. Therefore, the result suggests a relatively good match but with potential for further improvement. Furthermore, we determined that the group delay, rather than the gain, is the primary factor that influences the correlation score (see Section 5.1.3 for details). Consequently, the observed mismatch in the middle of the pulse, as shown in Figure 5.5, is likely responsible for the suboptimal match between template and candidate waveforms. Although other systematic effects may contribute to the suboptimal correlation score, the hardware response description could be improved by incorporating all relevant components, which are not fully represented in the placeholder response.

Database response: An initial attempt to comprehensively describe all components of the amplification and digitization chain is the so-called *database response*, which we can generate by using the RNO-G hardware and calibration database. In this database, we store S-parameter measurements for each component in the amplification and digitization chain, and, per channel, we can query a combined response for all contributing components. More information about the database is provided in Section 5.2. Unlike the placeholder response, this response includes the SURFACE board response for the exact board deployed for the channel of interest, as well as the coaxial cable from

the antenna to the Surface board, the short coaxial cable from the Surface board to the Radiant board, and the response of the Radiant board.

As an example, the combined response, the combination of the S21 parameter measurements of the components, of channel 13 from station 23, is presented in Figure 5.4. As discussed in the previous paragraph, we can identify a difference in the shape of the gain pattern compared to the placeholder response. Additionally, the group-delay excess at low frequencies is the largest among all responses we consider. In general, we expect changes to appear as we include additional components with their respective group delay and gain characteristics in the description. Moreover, we can see that the database response exhibits a higher noise level than the other response, which is likely due to the VNA not being properly configured for the S-parameter measurements. We anticipate that future measurements will have lower noise levels.

For the comparison with the UHECR candidate waveforms, we generate the templates using the hardware response description, which we query from the database for the specific components associated with the channel that identified the UHECR candidate event. In Figure 5.5, we compare the template waveforms generated with this response description with the UHECR candidate events. For the first two reference UHECR candidate waveforms, the oscillations at the beginning of the pulse are smaller than the template, and towards the end of the pulse, they are out of phase with the template. In the third waveform, the amplitudes at the beginning of the pulse match better, but the waveform is out of phase with the template, although the oscillations towards the end of the pulse match better. These differences in the template-UHECR candidate matching are reflected in the correlation scores, which we provide in Table 5.1, where the first two example waveforms have lower correlation scores than the third.

Overall, the correlation scores for all three UHECR candidate events decreased compared to the placeholder response. This result is counterintuitive, as more actual components of the amplification and digitization chain are described. However, several potential issues may contribute to the low correlation score. First, when we calculate the response from the database, we only multiply the S21 parameters together, neglecting possible reflections and standing waves introduced between components when impedances are not perfectly matched. To address this, we could either measure the S21 parameter of the combined chain in the laboratory or combine the S-parameter measurements, also accounting for the other S parameters and thus taking reflections into account. Work is ongoing to integrate the full S-parameter treatment described in Ref. [187] into the database framework. Second, it is challenging to perform a full S-parameter measurement of the Radiant board, as it has only an input, digitizes and saves the input signal, making it impossible to connect a VNA to the Radiant "output". However, a test point before the analog-to-digital converter (ADC) on the Radiant board allows us to probe the signal, enabling a full S-parameter measurement up to the ADC. As a result, the description of the entire section downstream of the test point is unavailable, which can lead to a difference between the actual and measured Radiant responses.

Impulse response: To address the issues of the incorrect cascading of the S parameters and an incomplete description of the Radiant response, collaborators in Chicago performed an additional measurement of the full digitization and amplification chain.

Table 5.1: The correlation scores for the templates generated with the different hardware response descriptions and the reference UHECR candidate waveforms.

template	waveform 1	waveform 2	waveform 3
placeholder response	0.886	0.888	0.877
database response	0.754	0.762	0.867
impulse response	0.906	0.919	0.956

In this measurement, they injected very short pulses (width: around 200 ps) into the full amplification and digitization chain and recorded them with the RADIANT board. Due to their very short width, the pulses have an effectively flat frequency spectrum within the RNO-G frequency range, making them delta-like pulses. As a result, the Fast Fourier Transform (FFT) of the recorded pulse directly describes the frequency response of the measured system, without requiring division by the FFT of the input pulse. This measurement also includes the full digitization response of the RADIANT. To reduce noise in the frequency domain, we apply a Tukey window [188] around the recorded pulse in the time domain before we calculate the FFT. This measured response is called the impulse response.

It should be noted that the measurement was normalized in the time domain, so the true gain is not reflected in the measurement. For simulations, where the gain is important for trigger decisions, we need to rescale the gain of the response. At the time of the simulations, a full gain calibration of the channels was unavailable to rescale the response. Therefore, we calculated a mean scaling factor from all channels of interest using the stored gain in the database, which is close enough to the true gain of the channels. We then used this mean scaling factor to rescale the impulse response gain for simulation purposes. This rescaling does not affect the correlation score, as the score is normalized and absolute factors do not alter the result. Further details about the scaling of the impulse response are provided in Section 7.2. Additionally, the absolute timing of the measurement was normalized, so the overall time delay may not reflect the true time delay of the entire amplification and digitization chain. Since only the shallow antennas are considered in this work, the overall time delay does not influence the results, as it is consistent across all antennas and does not introduce a systematic error in the direction reconstruction. Again, this does not affect the correlation score because a scan over the entire trace is performed, making the pulse location on the waveform irrelevant. For the correlation score, only the frequency-dependent delay, or group delay, is influential, and this is not altered by the normalization.

The gain and group delay of the impulse response are shown in Figure 5.4. The gain pattern of the impulse response roughly matches the gain pattern of the database response, with the main difference being that the high-frequency cut-off appears earlier than in the database response. However, due to the characteristics of UHECR signals and the sensitivity of the LPDA, the most important region of the response is at low frequencies. Compared to the database response, the group delay at low frequency values is slightly smaller for the impulse response. Figure 5.5 shows a comparison of the template and UHECR candidate waveforms. As with the other two responses, the templates slightly underpredict the pulse amplitude at the beginning of the pulse, but the templates and candidate events are in phase very well. The correlation scores

Table 5.2: The correlation score of the different reference UHECR candidate waveforms with (modified) template waveforms.

template	waveform 1	waveform 2	waveform 3
impulse response	0.906	0.919	0.956
database response	0.754	0.762	0.867
impulse gain & database phase	0.759	0.769	0.867
database gain & impulse phase	0.904	0.915	0.956

between the templates and UHECR candidate waveforms are shown in Table 5.1. All three example events have correlation scores above 0.9, with one event exceeding 0.95, indicating a good match.

These impulse response measurements are the most accurate and complete description of the RNO-G amplification and digitization chain, and we use them as the standard hardware response description for the analysis presented in this thesis. During the 2024 deployment season, the RADIANT board and the entire DAQ box were returned from Summit Station, enabling the retroactive creation of impulse response measurements for the components used during the data-taking periods relevant for this thesis. Furthermore, we found that the differences in the impulse responses across stations and channels for the shallow components are negligible (see Section 8.4 for more information), so we use the same impulse response for all shallow channels and stations.

5.1.3 Influence of gain and group delay

In the previous sections, we demonstrated that minor variations in the description of the amplification and digitization chain can significantly affect the correlation score. Here, we examine whether the gain or the group-delay differences primarily drive changes in the correlation score. To investigate this, we use the difference between the database and impulse response. In addition to the template waveforms generated with the database and impulse response, we construct two additional sets of template waveforms: for one set, we apply a hardware response description that combines the gain from the impulse response and the phase (group delay) from the database response, and for the other set, we reverse these roles. Afterwards, we calculate the correlation scores between these templates and the three reference UHECR candidate waveforms that we already used in the previous section. Although we generate multiple waveforms per template set, we consider only the template that has the highest correlation score with the UHECR candidate waveforms. The resulting correlation scores for the four template sets with the reference UHECR waveforms are presented in Table 5.2. Additionally, the correlation scores for the templates created solely with the database or impulse response are the same as shown in Table 5.1. The Table indicates that varying the gain of the database response to match that of the impulse response results in a minimal change in correlation with the UHECR candidate waveforms. In contrast, exchanging the phase results in a correlation score nearly identical to that obtained with the impulse response template. Therefore, we attribute the observed difference in the correlation scores between the database and impulse response templates primarily to differences in their group delays. These findings suggest that small changes in the gain of the response have a limited

impact on the correlation score, whereas accurate measurements of the group delay are most important.

5.2 RNO-G hardware and calibration database

As we have shown in the previous section, accurate modelling of the hardware response is essential for both simulation and reconstruction. With the RNO-G hardware and calibration database, we address this requirement by storing calibration measurements for all components, thereby enabling the construction of realistic detector responses. This capability is crucial not only for the analysis presented in this thesis but also for all future analyses, especially those sensitive to the shape of the recorded waveforms.

The database features a flexible structure that enables storage of different types of calibration data, such as data collected prior to deployment or in situ measurements obtained in the ice. Furthermore, we store the station and channel positions in the database, ensuring that all information necessary to virtually construct the detector is available. Additionally, we designed the database to track the evolution of the detector over time. For instance, changes to the hardware components can be recorded and are reflected in a time-dependent detector description. Also, we can incorporate improvements in the measurement accuracy in the database, such as updated data of the channel positions.

In this section, we describe the structure of the database (see Section 5.2.1), the two primary methods to interact with the database - the custom web interface (see Section 5.2.2) and the integration into the simulation and reconstruction framework (see Section 5.2.3) - and we conclude with a discussion of the current status and the future plans for the database (see Section 5.2.4).

5.2.1 Database structure

We implemented the database using MongoDB [189], where the data is stored in JSON-like documents that do not require a predefined schema, enabling a flexible and extensible database structure suitable for accommodating, e.g., new types of calibration measurements. For the RNO-G database, two MongoDB instances are hosted on the servers of the DESY computation center, each supporting multiple databases. The first instance, *RNO-G_live_test*, contains two databases, *RNOG_live* and *RNOG_test*, both of which we are actively and solely using for development and testing. The second instance, *RNO-G_live*, serves as the primary instance and contains the production database, *RNOG_hardware_v0*, which is continuously available for user queries. Maintaining multiple databases facilitates ongoing development of both the database and the code that queries it, without the risk of disrupting active simulation or reconstruction code that queries the database.

The database is organized into multiple *collections*, analogous to folders, which we use to organize the different types of data stored. Each collection comprises multiple *documents*, each document representing a JSON-like data structure that stores the actual data. A document may contain diverse data types, such as strings, floats, arrays, or dictionaries,

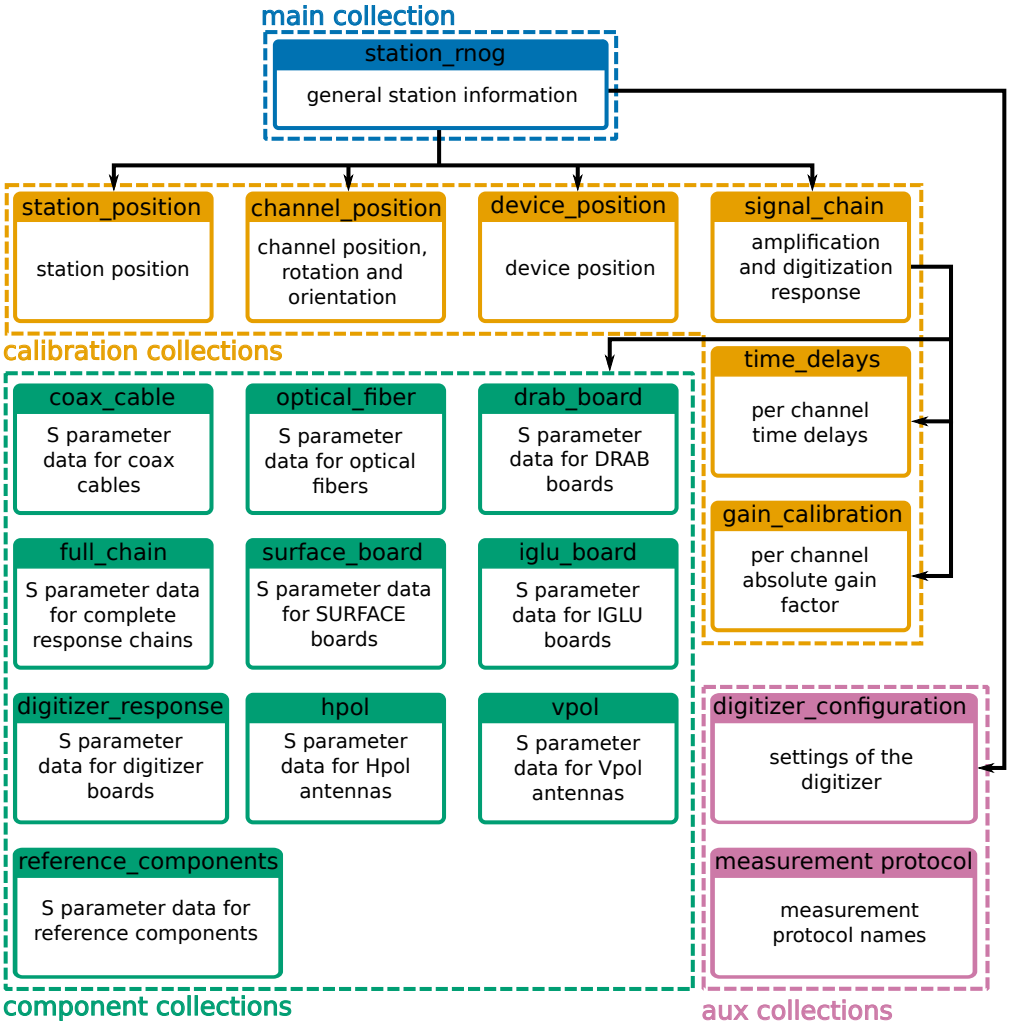


Figure 5.6: Schematic of the RNO-G calibration and hardware database. Different colors represent different types of collections, while the black arrows indicate their relationship, with the originating collection referencing the target collection.

and its structure is not fixed and may vary within a collection. The hardware and calibration database is organized into four types of collections. In Figure 5.6, we provide an overview of the database structure, illustrating all collections and their relationship.

Before we discuss the different types of collections, we need to introduce two more concepts: the *measurement protocol* and the *detector and database time*.

Measurement protocol: RNO-G components are measured by different individuals using different measurement setups. To document the measurement process and identify the responsible individual, we established the measurement protocol. This measurement protocol is a unique string stored with each measurement in the database. Furthermore, the individual performing the measurement creates an entry for the new measurement protocol on a dedicated page within the internal RNO-G wiki. This entry provides detailed information about the measurement, including the setup, measurement date, location, responsible individual, location of the raw data, and the specific components and quantities measured. Therefore, the measurement protocol enables the investigation

of measurements if issues arise, which is particularly important if the original contributors are no longer part of the collaboration. We store all measurement protocol strings in a dedicated collection within the database (*measurement_protocol*).

Detector and database time: To track changes in the hardware components, settings, and knowledge of the detector over time, we implemented two temporal concepts in the database: *detector time* and *database time*.

The detector time enables the user to retrieve the precise detector configuration at any specified time. For instance, the eight stations currently deployed for RNO-G were installed across three different seasons. If the detector time is set to a date when only three stations were operational, the resulting detector description will include only those three stations and their respective components present at that time.

In contrast, the database time, also referred to as *primary time*, tracks the evolution of the knowledge regarding specific measurements. For example, initial position measurements of the stations and channels may lack precision but are sufficient for preliminary use and are stored in the database. When more accurate measurements become available, such as after a subsequent measurement campaign, these updated values replace the previous entries as the primary description and are also stored in the database. The database time documents the time period during which earlier measurements were considered valid, thereby tracking the evolution of the measurements. Selecting a past database time allows for the retrieval of a detector description based on historical measurement knowledge, which is essential for reproducing results from earlier simulations or reconstructions that rely on a now-outdated detector description. We not only track the evolution of the position data but also assign each component measurement a database time range that extends to 2100 if it represents the current best knowledge of the measurement. We made the database detector description available to the collaboration in November 2024, so the earliest selectable database time is the 1st of November 2024. Attempting to select an earlier date will result in an error.

In the following, we provide a more detailed description of the different types of collections and the information they store.

Component collections: The component collections form the foundation of the database structure and store raw data from measured components. Prior to deployment, the responses of the amplifiers, digitizers, cables, and fibers are measured in the laboratory, and the resulting raw data are stored in the database. Certain amplifiers, the Iglu and Drab, have an analog input/output and an optical output/input, which requires measuring both amplifiers together as a pair with a VNA. Consequently, the Iglu (Drab) must always be measured alongside a reference Drab (Iglu). The reference components are also measured together to facilitate subsequent disentanglement of the measurements. These reference component measurements are stored in a dedicated component collection, named *reference_components*. In preparation for the 2024 deployment season, complete hardware chains, consisting of an Iglu, an optical fiber and a Drab, were measured in the laboratory using a VNA. If the hardware chains are deployed in the same sequence, we can use the full-chain measurement to describe the system response, which eliminates

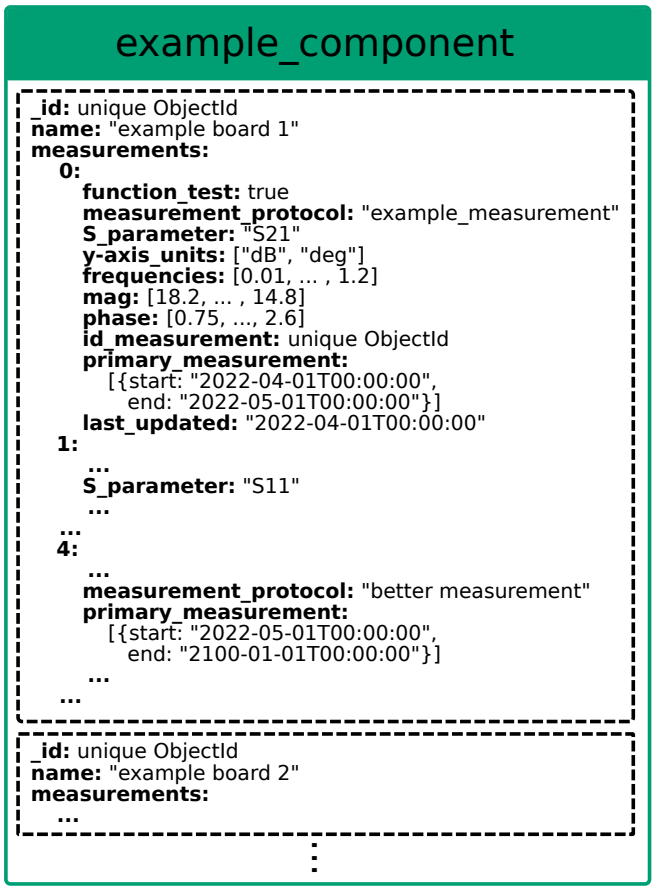


Figure 5.7: Illustration of the structure of an example component collection, showing entries that are consistent across all component collections. Additional details are provided in the accompanying text.

potential uncertainties arising from combining single-component measurements. These measurements are stored in a component collection named *full_chain*. The document structure is intentionally flexible to accommodate, e.g. future in situ measurements of the deployed signal chain in the ice. As described in the previous section, a new measurement of the entire amplification and digitization chain was conducted using a very fast pulse generator, referred to as the *impulse response* measurement. This measurement is also included in the *full_chain* collection. Furthermore, all custom-built Hpol and Vpol antennas undergo testing prior to shipment to Greenland. For this purpose, their S11 parameter is measured in the laboratory and compared with simulations to identify potential issues. Although these S11 measurements currently serve only as a pass/fail criterion, they are also stored in the database.

As we mentioned before, the document structure within the component collections remains flexible to support a wide range of measurement types. However, a common structure is maintained across all documents. In Figure 5.7, we present an example collection with representative documents to illustrate the entries that are consistent throughout all documents. Each document always contains three entries: the name of the measured component, a unique identifier, and an array containing all additional information (*measurements*). The array may include multiple entries, each describing a measurement of an S parameter and providing all necessary information, such as the measurement protocol or the frequency, magnitude and phase data. If the same S parameter is measured multiple times, whether with identical or different setups or

techniques, all measurements are stored within the same document. Selection of the appropriate measurement during data queries is managed using the primary times (database times), as we discussed previously. In addition to the standard entries, optional entries may be included in the entries of the measurement array. Due to the flexible database structure, we can store a variety of supplementary entries, such as the name of a reference component measurement (for Iglu boards, Drab boards or optical fibers), channel ID for multi-channel components, the time delay of a component calculated from the raw phase data, or the attenuation value if an attenuator was used in the measurement setup.

Notably, the magnitude and phase data stored in the database are raw and unprocessed. Any necessary processing, such as time delay subtraction, is performed during data querying.

Calibration collections: There are three different types of collections that we classify as calibration collections. An example of the structure for each collection type is provided in Figure 5.8.

In the first collection type, we store position information, with three separate collections dedicated to channel, device[3] and station positions. Station positions are reported relative to a fixed point at Summit Station[4], which serves as the anchor for the RNO-G coordinate system, while the channel and device positions are measured relative to the station position[5]. All three collections share a common structure: each document contains a unique identifier, a position ID, and an array of measurements. The position ID links each measurement to the corresponding stations, channel or device, as stored in the main collection containing general station information (see the next paragraph for more details). The measurement array includes all position measurements performed for this station, channel or device. Each entry in this array describes a measurement, with the currently active measurement identified by the primary time (database time). In addition to the position, for channels and certain devices, such as pulsers, rotation and orientation angles are also recorded in the measurement entry. Also, a channel or device ID is stored for channels and devices in each entry. Otherwise, measurement entries are consistent across the three different position collections. Additionally, each measurement entry includes the measurement name, time of measurement, a unique ID, the last time the measurement was updated, and time periods during which each measurement serves as the primary measurement.

The second collection type is the signal chain collection, which is a single collection with notable differences from the others. This collection contains, for each channel, all information required to describe the amplification and digitization response. For each channel, the information is stored in a document containing a unique ID, a signal ID and the response descriptions ("measurements"). The signal ID links the response to the relevant channel and station and is also stored in the main collection. The measurement array contains entries describing the response, with the applicable description at any

[3] In the context of the database, a device refers to any component that is not a channel, including pulsers, DAQ boxes, solar panels, and wind turbines.
[4] As the origin of the RNO-G coordinate system, the DISC borehole is used (see Figure 4.1).
[5] The station position is defined as the position of the power string.

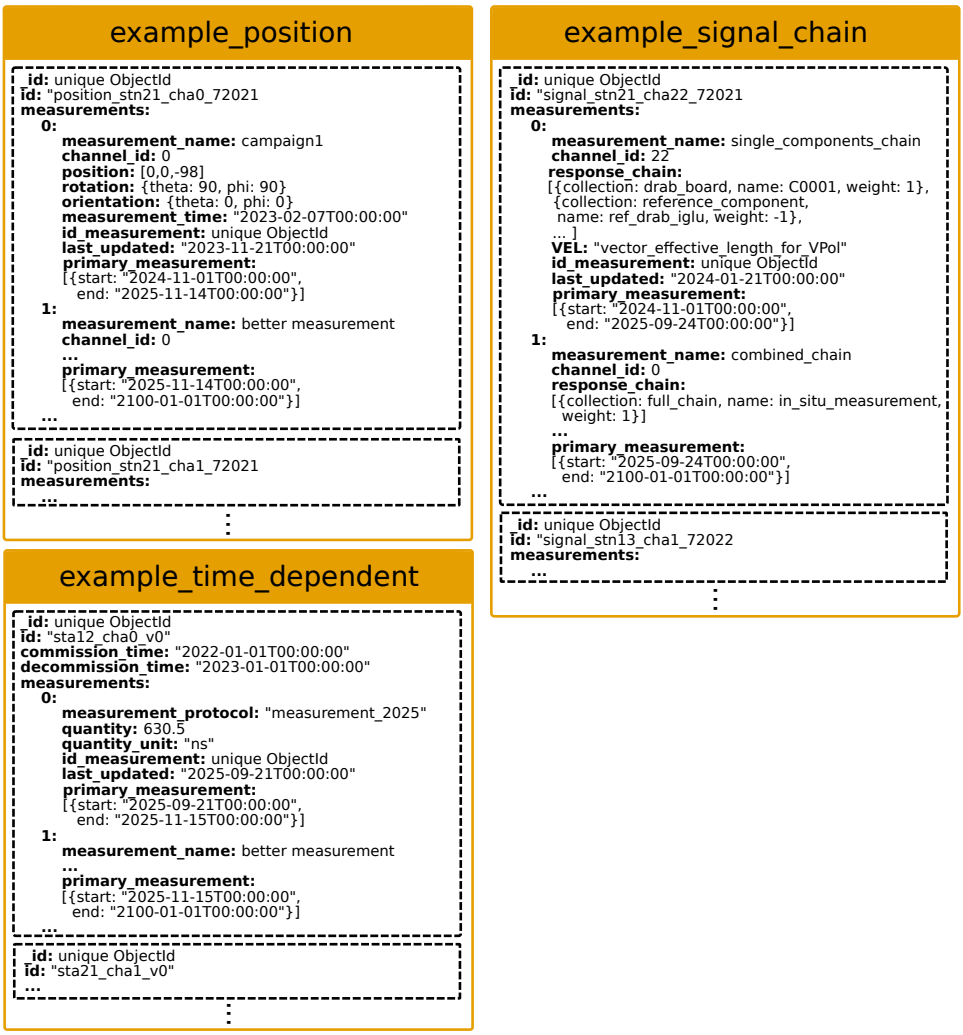


Figure 5.8: Illustration of the structure of example calibration collections for the three types, highlighting all common entries. Additional entries may exist that are not present in every collection and are therefore not depicted. We provide further details in the main text.

given time indicated by the primary time (database time). Each entry also includes a unique measurement ID, channel ID, the VEL name for the antenna response, the last time the description was updated, and a response array. The response array consists of dictionaries that describe the response, which may be multiple if the response is constructed from the S parameters of several components, or a single dictionary for full chain measurements. Each dictionary specifies at least the collection containing the response measurement, its name used to locate it, and a weight, which determines whether the total response is multiplied or divided by this response measurement. Additionally, optional arguments may be included, such as specifying the channel for components with multiple channels. If the channel is part of the phased array, additional entries are required in the response entry to account for the different response of the phased-array trigger. A flag, *is_trigger_chain_absolute*, indicates whether the trigger response is described as a difference from the signal response or as an absolute description. Furthermore, an additional array, *trigger_response_chain*, is added, which mirrors the structure of the signal response array and describes the trigger response. If the trigger chain is described as an absolute response, all necessary components are included in the additional array, even if they are present in the normal signal response.

On the other hand, if described by a difference, only the components not present in the normal signal response, as well as those that need replacement, are included in the trigger response. Currently, all trigger responses in the database are stored as absolute descriptions. Moreover, an optional comment field allows free-text notes for each entry in the measurement array. We did not depict the trigger response and comment field entries in Figure 5.8.

The third type of collection stores time-dependent variables, specifically the absolute gain factor and time delay per channel. These values are essential for making the impulse response description, which we described in Section 5.1, available for simulations and event reconstruction. Each document in the collection type describes the quantity for a channel over a defined time period, which is specified by the commission and decommission times. The appropriate measurement for the desired quantity is determined by the detector time when the quantity is queried. We describe these quantities as time-dependent due to observed temporal variations [190]. Furthermore, each document includes a unique ID linking the measurements to the station and channel, and an array with all available measurements for this quantity and time period, with the currently active measurement identified by the primary time (database time). Each entry in the measurement array contains the measurement protocol, value and unit of the measured quantity, a unique measurement ID and the last time the measurement was updated.

Main collection: The *station_rnog* collection contains general information about each station and is referred to as the main collection. It acts as the starting point for detector description queries and references all necessary measurements across the various collections to enable the construction of a comprehensive detector description. The structure of this collection is illustrated in Figure 5.9. Each station is represented in this collection by a document containing key information, including a unique ID, station name, station ID, commissioning and decommissioning times, comments, IDs linking relevant measurements (position ID and signal ID), and arrays for the channel and device information. In this context, "device" refers to components such as pulsers, solar panels, wind turbines and instrument boxes. When hardware components, such as the RADIANT board, are replaced (as occurred during the 2024 field season) or when the station settings are changed, we reflect these changes in the database by decommissioning the current station document on the date of the modification and creating a new modified document with a commissioning time corresponding to the decommissioning date of the previous document[6]. As a result, the collection can contain multiple documents for the same station. Users can specify the detector time when querying the database to retrieve the appropriate station description. The IDs linking measurements or information to the station correspond to IDs in the relevant collections. For instance, the position of the station is stored in a separate collection, and the associated ID points to the correct position data. Similarly, the signal and trigger digitizer configurations (see the next collection for more information) are linked using this method.

The array for the channel information contains detailed per-channel information, with each entry structured as a dictionary that includes the channel ID, antenna type, comments and independent commissioning and decommissioning times. If only the hardware

[6] In certain instances, modifications may require several hours or days, e.g. when a component is replaced in the field. In these situations, the new commissioning time is set to the date when the modification is complete.



Figure 5.9: Illustration of the structure of the main collection, with the provided values serving as examples. We provide further details in the main text.

of a single channel is updated, the affected channel entry can be decommissioned and a new entry added, thereby reflecting changes at the channel level without requiring a new station document. Each channel entry also contains a list of the installed signal chain components, described by dictionaries that specify the collection, name, and, if applicable, the channel of the deployed component. Additionally, channel entries include IDs that link measured channel positions and the hardware response descriptions, which may differ from the installed components if, for example, in situ measurements are used.

The array for the device information mirrors the structure of the channel array, with each entry describing a deployed device. However, the content of each device entry varies depending on the device type. For example, a pulser entry specifies the amplifier name, which is not required for solar panels. General device information includes the device ID, device name, comments, commissioning and decommissioning times, and, if an amplifier is deployed with the device, its name. The device IDs are fixed to the device type, with the mapping between provided in Table 5.3.

Generally, we do not store measured data in the main collection, but we do store general station information and links to all key measurements. We indicate the collections linked to the main collection in Figure 5.6 by black lines.

Table 5.3: The mapping between device names and IDs, which is used exclusively within the database.

ID	name
0	Helper string C pulser
1	Helper string B pulser
2	Surface pulser
100	DAQ box
101	Solar panel 1
102	Solar panel 2
103	wind-turbine

Aux collections: The auxiliary collections are two collections that do not align with the previous three collection categories: *measurement_protocol* and *digitizer_configuration*. The measurement protocol collection stores the names of all measurement protocols, which are required to populate a drop-down menu on the custom database web interface (see Section 5.2.2 for more information). Each measurement protocol is stored as a document containing the name of the measurement protocol, the insertion timestamp, and a unique ID.

The digitizer configuration collection contains, as the name suggests, configurations for the digitizer boards used, each saved as a separate document. For the Radiant board, the sampling frequency, number of bits, minimum and maximum voltage, and number of samples are stored, along with a unique ID and a human-readable ID. For the Flower board, the same parameters are stored, except that the number of samples is replaced by the noise count. The human-readable ID is referenced in the documents of the main collection to connect a configuration with a specific station. These digitizer configurations are maintained in a separate collection and not included in the documents of the main collection to avoid data duplication, as they are identical across all stations.

5.2.2 Web interface

To facilitate collaboration wide access to the database, we developed a web interface that enables users to add and view measurements. We implemented both the back-end and front-end of the web interface in Python, utilizing the Streamlit package [191] for the graphical user interface. Streamlit is an open-source framework that enables fast website development without requiring knowledge of languages such as JavaScript. It is easy to use, accessible for individuals without prior front-end experience, and it produces clean interfaces suitable for the relatively simple requirements of the database web interface. The web interface code runs within a Docker container on a virtual machine hosted on a DESY server and the website is accessible at `https://radio.zeuthen.desy.de`, which requires authentication via either a GitHub or DESY account. Users start at the home page, which we provide a screenshot of in Figure 5.10. This home page provides general information and primarily serves as a landing page, with navigation to other pages available via the sidebar. These pages can be grouped into three categories: adding S parameter measurements, adding station information, and displaying database content.

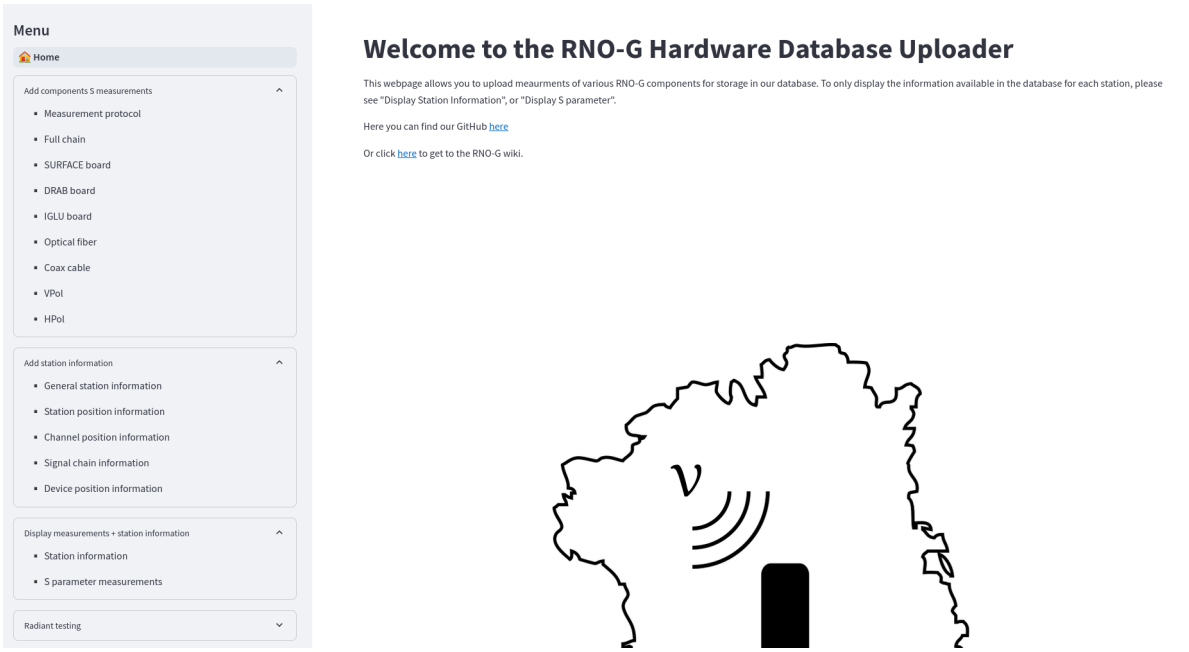


Figure 5.10: Screenshot of the database web interface home page, with the sidebar fully expanded. The last entry in the sidebar, "Radiant testing", serves a different function than the hardware and calibration database and is not relevant to this work.

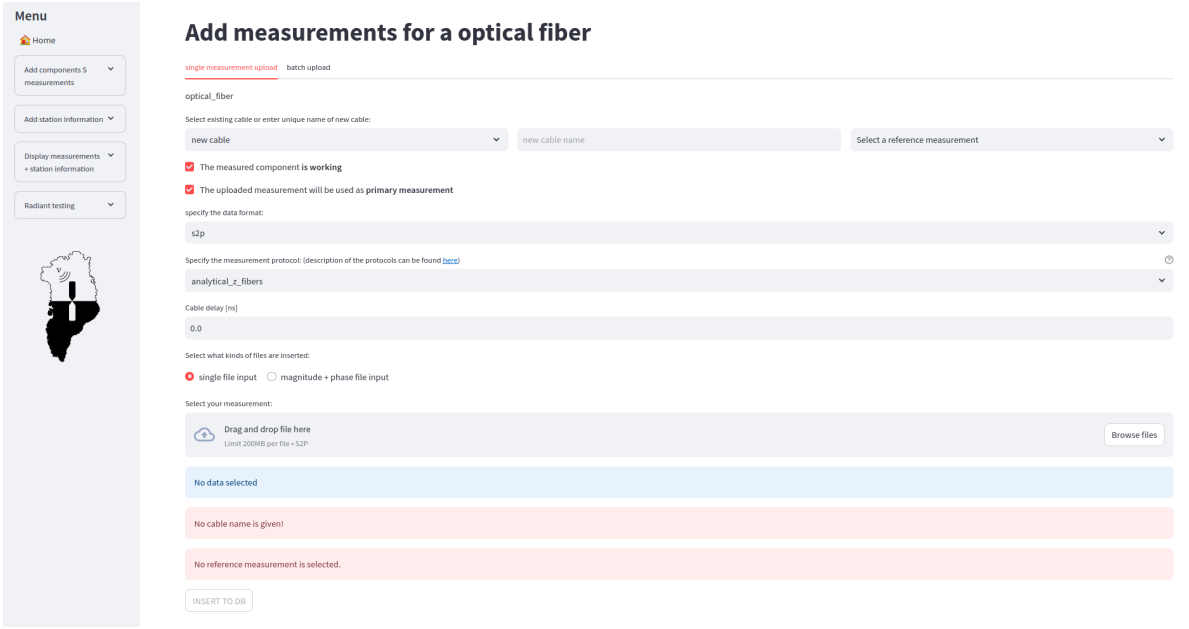


Figure 5.11: Screenshot of a page for adding S-parameter measurements to the database. As an example, we show the single measurement upload page for optical fiber measurements.

Adding S parameters: For each laboratory-measured component, we implemented a dedicated page to insert the measured S parameters into the database. As an example, we show a screenshot of the optical fiber page in Figure 5.11. Generally, the structure of these pages is consistent across the different components. First, users may choose to upload either a single measurement or a batch of measurements. In the single-measurement workflow, users enter all required measurement information, which is stored in the database. Subsequently, users can upload a measurement file, which is automatically plotted at the bottom of the page, allowing them to verify the data before submission. Additionally, automated checks ensure that all fields contain valid values and that the measurement can be parsed correctly. If all checks are successful, the upload button is enabled, and when pressed, the measurement will be written to the appropriate component collection in the database. In the batch upload workflow, multiple data files can be uploaded, plotted and submitted to the database simultaneously. However, for batch uploads, fewer fields are available to enter measurement information and the necessary information must be encoded in the filename, as specified on the respective page. One page, in this category, differs from the others: the upload page for measurement protocols. On this page, users only add the name of the measurement protocol to the database, after which the measurement protocol must also be added to the wiki, including more detailed information about the measurement. Currently, only pages exist for established measurements. If new components are measured or new measurement types become available, we must create additional pages.

Adding station information: Each page in this category serves a different purpose, but all are used to insert station, channel or device information into the database. The first page allows the entry of general station information into the main collection of the database, including the (de)commissioning time, and channel and device information. Additionally, there are pages for entering the station, channel, and device positions into their respective collections. Another page allows adding amplification and digitization chain descriptions to the signal chain collection. In conclusion, to obtain a comprehensive station description from the database, information must be entered across all the pages mentioned here.

Displaying database content: One page allows users to plot the absolute value and phase of the S21 parameter for any measurement and component in the database, which enables manual inspection of the measurements. Another page displays the general station information for a selected station, presenting all available information for the specified detector time from the main, station position, channel position, device position, and signal chain collection.

The web interface is the primary way for most users to add data to the database. While it is possible to add information and measurements via Python scripts, we restrict this functionality to a small group of expert users within the collaboration. Querying information is available to all users through the web interface or via NuRadioReco.

5.2.3 Access and querying

To use the detector description stored in the database for event reconstruction or simulations, we integrated the database into NuRadioReco, enabling its use in both

NuRadioReco and NuRadioMC [192]. We achieved this integration through a database class that provides read and write functions. Only a select group of expert users within the collaboration and the web interface have write access via the database class, while general users are limited to its read functionality. To ensure a seamless integration, a new detector class was developed that utilizes the database class to query a comprehensive detector description and returns a detector object compatible with NuRadioReco and NuRadioMC functions and modules.

When a complete detector description is requested, the query begins by retrieving the general information from the main collection for the relevant stations. Subsequently, the station position is obtained from the *station_position* collection, with the position ID stored in the general information to identify the correct measurement. Afterwards, the per-channel information is queried, starting with the general channel information from the main collection. The channel position ID, stored in the general channel information, is used to retrieve the corresponding channel position from the *channel_position* collection. Then, the amplification and digitization response description is queried from the *signal_chain* collection using the channel signal ID. These descriptions only specify which measurements from which collections are required. Therefore, the appropriate measurements are subsequently queried from the component collections, and, if necessary, the calibrated time delays or absolute gains are queried from the *time_delays* and *gain_calibration* collections. The response description encodes how to combine the retrieved measurements using a *weight*: a weight of 1 indicates multiplying the total response by the measurement, while a weight of -1 indicates dividing the total response by the measurement. At the end, the device information is queried using the device position ID from the main collection to obtain the correct position measurements for the specified devices. Currently, only the device positions are queried, as, for example, response descriptions for pulsers are not yet documented in the database. All collected information is then combined into a single dictionary and made available to the user. When the detector description is accessed via the detector class, all queried information is stored in the returned detector object. We illustrate the interconnections between collections in Figure 5.6, where the black arrows indicate the connections.

In addition to always querying the primary measurement at the specified database time, it is possible to override the primary measurement and query by measurement name, which is useful for studying non-primary detector descriptions. Information about the digitizer configuration, such as the sampling rate, is not currently included in the complete station description query but can be accessed via separate functions. In future updates, we may incorporate this information into the comprehensive querying procedure. Another feature allows the entire detector description to be exported to a file, which the detector class can subsequently use as its information source. This enables the use of the same detector description in multiple independent calculations without requiring repeated database queries, which is particularly advantageous for parallelizing simulation code on computation clusters.

5.2.4 Status and future plans

All eight deployed stations are currently fully described in the database and are available for querying. Additionally, the database accurately reflects all modifications applied to

the stations over the recent years, including changes to the DAQ boxes and sampling rates. Users can retrieve previous detector configurations by setting the detector time to the relevant date.

For the amplification and digitization response description, the database offers two options. The standard description, retrieved when the database time is set to the present (the default), is based on the product of the S21 parameters of all components in the amplification and digitization chain. As we discussed in Section 5.1, this approach neglects potential reflections due to impedance mismatches. A more accurate treatment of the S parameters, as detailed in Ref. [187], is being implemented in the database. Alternatively, users can query the measured impulse response, obtained by collaborators in Chicago through full-chain time-domain measurements. Since this impulse response is normalized, we must scale it to the correct absolute gain for simulation and reconstruction, which we achieve by multiplying the response per channel by an absolute gain factor. The calibration method and initial results are presented in Ref. [171]. The absolute gain factor varies over time [190] and therefore, we store it as a time-dependent quantity in the database. Furthermore, the impulse response measurement does not capture the absolute time delay of the system. Thus, we add an overall time delay, which we calculate from the S-parameter description, as a time-dependent factor. As we demonstrated in Section 5.1, the impulse response description currently provides the most accurate representation of the RNO-G hardware. However, two issues prevent its adoption as the default description. First, we do not have the gain calibration values for the new Radiant board deployed in the 2024 season, which limits the use of this response to periods before the 2024 deployment. Efforts are ongoing to generate these calibration values, and we expect to resolve of this problem soon. Second, there is no impulse response measurement or gain calibration for the Flower board, which precludes an accurate description of the trigger for the phased-array channels. While this does not affect reconstruction code, it, e.g., significantly impacts effective volume simulations. Collaborative efforts are underway to address this issue. As an interim solution, we continue to describe the trigger response for the phased-array channels using S21 measurements, consistent with the default response description. Once these issues are resolved, the impulse response description will become the default. Until then, advanced users can access the impulse response description via its measurement name, *"calibrated_impulse_response_v0"*, when querying the detector description.

Although the database is currently in a stable and functioning state, several tasks remain, as discussed above. Furthermore, additional improvements to the web interface are needed, such as introducing new features and standardizing the code for the different pages. Another potential project involves implementing the full response for the calibration pulser, which is deployed with the strings. Currently, only the positions of these devices are stored in the database. If we can model the signals they produce, it could enable direct comparison with the recorded signals from the pulsers. However, the most challenging upcoming project is preparing the database to store calibrated detector data. The collaboration is advancing towards a full detector calibration utilizing multiple calibration techniques and sources [165], which will yield updated positions and time delays. Additionally, the ice model will be fitted simultaneously, requiring storage of the best-fit parameters. Crucially, the calibration will include correlated uncertainties for all quantities determined by the calibration. To ensure a

proper treatment of these uncertainties, it is essential that we also accurately reflect the correlation of the uncertainties in the database. Up to now, we have not determined the optimal solution for storing all calibrated parameters and their uncertainties in the database.

5.3 Radiant trigger

In this section, we discuss the performance and potential improvements of the Radiant trigger. The Radiant trigger, present on all channels of the Radiant board, is exclusively used for the shallow LPDA antennas, and therefore, sometimes referred to as the shallow trigger. Detailed descriptions of the trigger circuit are available in Refs. [163, 172] and thus, we only provide information relevant to the subsequent section. The trigger logic is based on an analog diode detector circuit, where the diode output is compared to a threshold. The diode output signal exhibits its largest amplitude at negative values below its baseline and is offset by a voltage of 1.2 V [172]. Therefore, the trigger condition is met when the diode output falls below the threshold value. Consequently, the trigger is more sensitive at higher threshold values, which is counterintuitive compared to the conventional interpretation of a trigger threshold. The input signal for the trigger circuit is diverted from the main signal path using a −10 dB directional coupler (TCD-10-1WX+) [163], resulting in only a small fraction of the incoming signal being available on the trigger path. As previously described in Section 4.3, the Radiant trigger thresholds are dynamically adjusted at the channel level. In the data-taking configuration, we set, per channel, a radiant scaler, defined as the number of times the diode output is below the trigger threshold per counted time period (typically one second) [172]. A feedback-based controller maintains the per-channel radiant scaler at a stable value near the predefined goal by adjusting the trigger thresholds. However, even if the radiant scalers for all channels remain stable, variations in the correlated noise rate can cause the coincidence trigger rate to fluctuate.

In the following, we investigate the noise trigger rate of the Radiant board (see Section 5.3.1), explain the procedure to measure trigger efficiency curves (see Section 5.3.2), discuss the trigger performance and evaluate improvements from minor modifications to the trigger circuit (see Section 5.3.3), and present a method to model the trigger threshold in simulations (see Section 5.3.4).

5.3.1 External vs board noise

In addition to the incoming signal, noise originating from the Radiant board itself, which we refer to as board noise, is present on the recorded waveform. Because the input signal on the trigger path is relatively small, due to the −10 dB directional coupler, board noise could dominate the trigger decision rather than external thermal noise. To determine the influence of the board noise on the Radiant trigger, we conducted laboratory measurements. If the trigger is dominated by external thermal noise, the noise trigger rate should increase when we introduce external thermal noise to the Radiant board, while keeping the threshold fixed, compared to when it is absent. To produce external thermal noise for these measurements, we connect a Surface board to

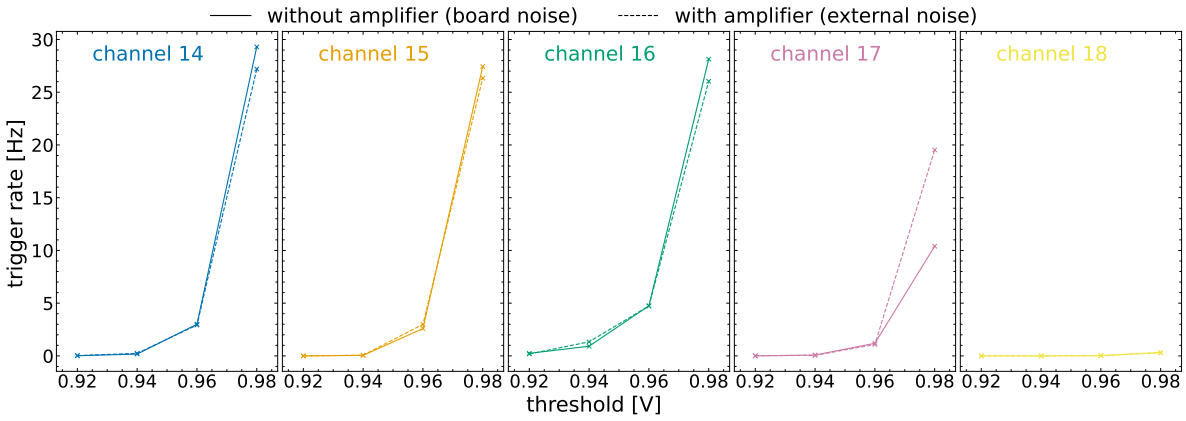


Figure 5.12: The measured noise trigger rate for different channels of a V2 RADIANT board with and without external. Note the counterintuitive trigger thresholds.

the RADIANT channel under test. Rather than connecting an antenna, we terminate the SURFACE board input with a 50 Ω terminator. When no external noise was required, we disconnected the SURFACE board from the RADIANT channel. We used a non-deployed spare RADIANT board (V2)[7] for these laboratory measurements.

Typically, the trigger threshold is dynamically adjusted based on the current noise level to maintain a stable trigger rate. During the laboratory measurements, we fixed the trigger threshold and recorded the noise trigger rate at different threshold values, both with and without the SURFACE board connected. In Figure 5.12, we present the measured noise trigger rates for both configurations and different RADIANT channels as a function of the trigger threshold. Regardless of the measured channel, the trigger rate curves behave similarly whether external noise was present or not. Both curves begin to rise at the same thresholds and exhibit approximately the same noise trigger rate. Although these channels differ in sensitivity, these results indicate that the trigger decision of the RADIANT board is not influenced by external thermal noise. Therefore, the measurements confirm that board noise dominates the RADIANT trigger. If external thermal noise were to dominate the trigger decision, we would expect the measured trigger rate with the amplifier connected to rise at a lower threshold value than without it.

Furthermore, channel 18 exhibits a trigger rate below 1 Hz at a threshold where other channels already exceed 10 Hz. This suggests that either the board noise on channel 18 is lower than on the other channels, allowing a larger trigger threshold value without reaching the noise floor, or that the diode behavior differs slightly, requiring a larger trigger threshold value to achieve the same trigger rate. Further investigations would be necessary to determine the origin of the difference in trigger rate.

Since the trigger is dominated by board noise, waveforms may be triggered even when the feature that triggers is not visible in the recorded waveform. As a result, the recorded waveform can appear to be randomly triggered thermal noise. We show a simplified illustration of this behavior in Figure 5.13. If we insert a waveform with a peak-to-peak voltage A into the signal path, −10 dB of the signal power (around 32% of the signal amplitude) is diverted onto the trigger path at the coupler, while the remainder continues

[7] The V2 version of the RADIANT board was originally deployed with the first seven RNO-G stations and took the data for the UHECR analysis.

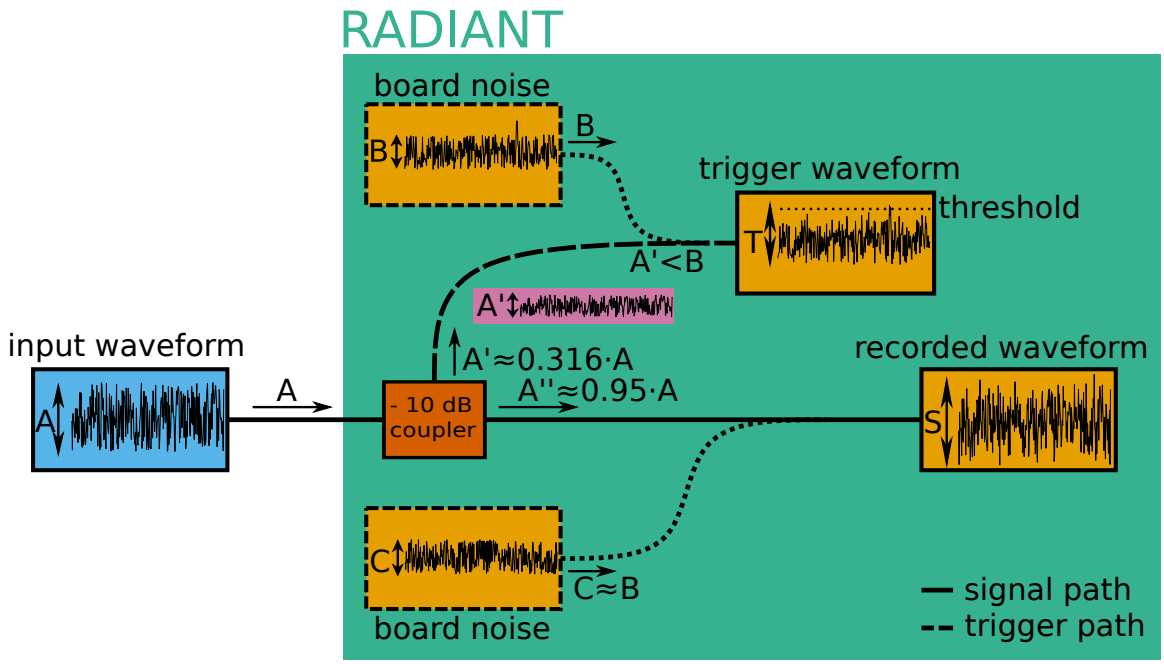


Figure 5.13: A schematic of the signal and trigger paths, which illustrates that the trigger path receives only a small fraction of the input amplitude. Therefore, features responsible for triggering might not be present in the recorded waveforms. For further details, refer to the main text.

along the signal path. Consequently, the peak-to-peak voltage of the signal on the trigger path A' is significantly smaller than the input signal. Before the trigger decision, board noise with a peak-to-peak voltage of B is added to the trigger path. Because the trigger is board-noise-dominated ($B > A'$), an overfluctuation in the board noise can cause the combined signal to exceed the trigger threshold and trigger the channel. As a result, the waveform that continues along the signal path is recorded, and before recording, board noise with a peak-to-peak voltage (C) is added to it. The board noise on the signal path differs slightly from that on the trigger path, however, we can assume that the peak-to-peak voltages are approximately equal ($B \approx C$). Although it does not necessarily exhibit the same overfluctuation. Consequently, the recorded waveform may be influenced by board noise but does not display the feature that led to the trigger decision, making it appear as randomly triggered thermal noise. Also possible, but not essential, is that the board noise on the signal path (C) is smaller than the signal component (A'') that continues along this path after the coupler. If an overfluctuation occurs in the inserted noise waveform, on the trigger path, the overfluctuation might be small compared to the board noise and not result in a trigger. However, if we inject impulsive signals with an amplitude significantly above the board noise into the circuit, they will remain larger than the board noise even after the coupler. Therefore, the RADIANT trigger can respond to those signals, but only until the signal amplitude after the coupler falls below the board noise level. As a result, if we would try to estimate an effective trigger threshold from the recorded waveforms (which are not randomly triggered by board noise), we would find a value higher than the noise level, even if the real trigger threshold sits directly above the noise level on the trigger path.

To address this issue, we can increase the power diverted to the trigger path to ensure that the external thermal noise on the trigger path remains larger than the board noise.

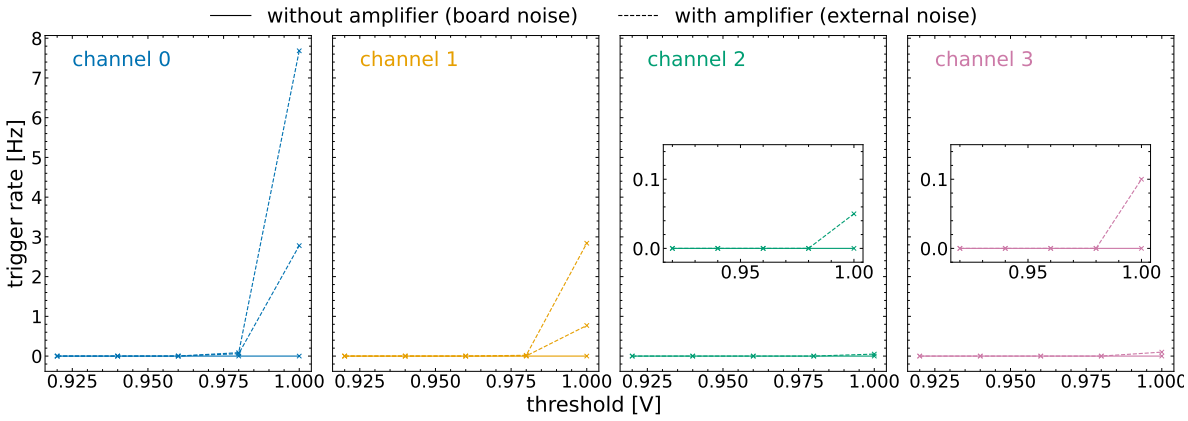


Figure 5.14: The measured noise trigger rate for different channels on a V3 RADIANT board. For channels 0 and 1, we performed the measurement with the amplifier connected twice, with the repeated measurements also shown as dashed lines. Note the counterintuitive trigger thresholds.

In the next-generation RADIANT board (V3), we replaced the directional coupler with a −3 dB power splitter (SBTC-2-10LX+), resulting in significantly greater power in the trigger path, while reducing the power in the signal path. Notably, we made additional changes to the trigger path for the V3 board, as we will discuss in Section 5.3.3. However, replacing the coupler should have the greatest impact on the issue we described above.

To investigate whether the modifications to the trigger path mitigated the problem, we repeated the measurements shown in Figure 5.12 on the V3 RADIANT board. The results for multiple channels are presented in Figure 5.14. Although we measured channels not used for the shallow analysis, they are equipped with the RADIANT trigger circuit and are suitable for testing whether the trigger remains board-noise-dominated after the changes. In all four channels, the trigger rate begins to increase at a lower threshold value in measurements with external noise present. However, the trigger rate difference between the two cases varies considerably across the channels. The observations indicate that the trigger path is now dominated by external noise rather than board noise, demonstrating the effectiveness of the applied changes. For channels 0 and 1, we repeated the measurements with the amplifier connected once. In both cases, the turn-on trigger threshold remains unchanged, but the noise trigger rate values vary significantly. The main difference between the two measurements is the duration, but in both cases, sufficient statistics are available due to the high trigger rate. The reason for this variation remains unknown. Furthermore, the trigger rate for channels 3 and 4 is significantly lower at the highest trigger threshold value than for the other two channels. This difference may indicate that the external noise on these channels is slightly lower, as we measured each channel with a different SURFACE board channel connected or that the diode behavior differs slightly. To determine the origin of this effect, further investigations are necessary. However, neither of these observed effects alters the overall conclusion of the measurement.

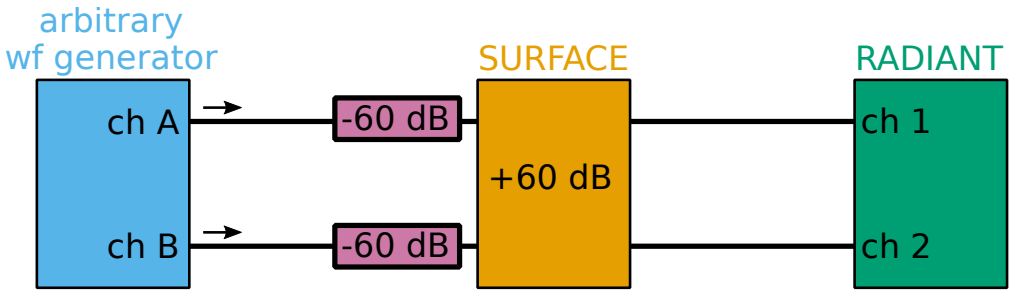


Figure 5.15: Schematic of the setup used to measure trigger efficiency curves for different channels of the RADIANT board. Channel 1 is the channel under test, while channel 2 is used to determine whether triggered events contain a pulse or only thermal noise. Additional information are provided in the main text.

5.3.2 Trigger efficiency measurements

In this thesis, we perform trigger efficiency measurements to evaluate the impact of trigger path modifications on the trigger performance (see Section 5.3.3) and to translate the measured trigger threshold into values usable for realistic UHECR simulations (see Section 5.3.4). In both cases, the measurement setup and analysis strategy are similar to the trigger efficiency measurements described in Ref. [172], and the SNR at which the trigger efficiency curve reaches 50% (SNR($\epsilon_{\text{trig}} = 50\%$)) is used to compare or map measurements. Therefore, in the following, we present the procedure to determine these values.

For the trigger efficiency measurement, we use an arbitrary waveform generator (81160A Pulse Function Arbitrary Noise Generator) to generate UHECR-like waveforms, which we derive from the UHECR analysis templates (for more information on the templates, see Section 8.2). However, due to the limitations of the waveform generator, frequencies above approximately 150 MHz are slightly suppressed. We connected the waveform generator to a SURFACE board, and the SURFACE board to the RADIANT channel under test, using short coaxial cables. This arrangement replicates the field configuration, except that an LPDA is connected to the SURFACE board in the field rather than the waveform generator. Because the waveform generator cannot produce peak-to-peak voltages below 50 mV_{pp} [193] and the RADIANT can only handle input voltages up to 1.2 V_{pp} [172], we placed 60 dB attenuators before the SURFACE board to counteract its approximately 60 dB amplification and ensure that we operate both devices within their specified limits.

Figure 5.15 shows a schematic of the measurement setup, with the channel under test depicted as channel 1 on the RADIANT board, which is connected to channel A on the arbitrary waveform generator. The second channel, channel 2, serves as a "clock" channel. At low amplitudes, distinguishing between a triggered pulse and thermal noise can be challenging. To address this, we generated a second UHECR-like waveform on the time-synchronized channel B of the arbitrary waveform generator, routed through a separate channel of the SURFACE board, and into the "clock" channel on the RADIANT board. We disabled the trigger on channel 2 and set the amplitude of the UHECR-like waveforms so that a clear pulse is always visible in the "clock" channel. When channel 1 triggers, we examine the recorded waveform of channel 2 to determine whether a pulse is present. If we do not observe a pulse in channel 2, we attribute the event to thermal noise and exclude it from the efficiency calculation.

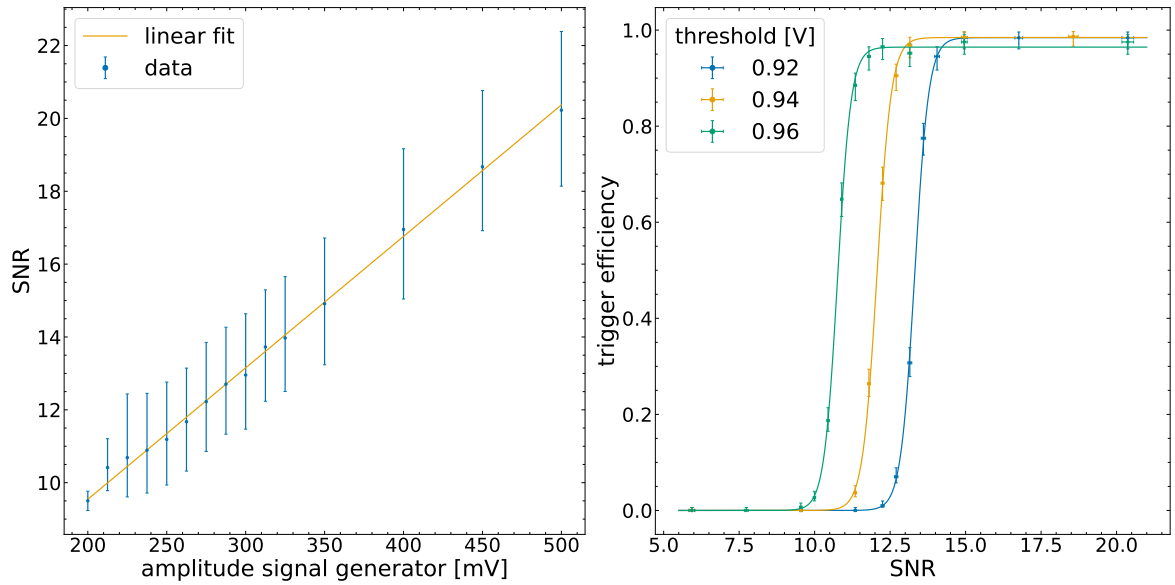


Figure 5.16: Left: The mapping between the arbitrary waveform generator amplitudes and the SNR values for measurements conducted with channel 15 of a RADIANT board. Right: The trigger efficiency curves for channel 15 measured at different trigger thresholds. These measurements were performed prior to any modifications to the channel.

For each trigger efficiency curve measurement, we fix the trigger threshold to a value similar to those recorded during normal data-taking. We could not record trigger efficiencies for thresholds above 1.0 V due to the high noise trigger rate at these thresholds. A trigger efficiency curve contains multiple efficiency measurements at different waveform generator amplitudes. We selected the waveform generator amplitudes so that we have at least one trigger efficiency data point on the rising edge of the trigger efficiency curve. For each fixed threshold and amplitude, we collected data for 300 s, with the arbitrary waveform generator injecting a UHECR-like waveform every second. The measurement duration is a compromise between the number of events per efficiency data point and the total measurement time required to obtain a trigger efficiency curve.

The initial step in analyzing those measurements is to convert the waveform generator amplitude to an SNR value to simplify the comparison between different channels and boards with potentially different noise levels. For this purpose, we group all measurements for a given channel and generator amplitude. Then, we calculate the SNR values for all recorded waveforms in a group, and both the mean and the two-sided 90% quantile of the resulting SNR distribution. These values are shown in Figure 5.16 as a function of the generator amplitude, with the quantile used to define error bars. Subsequently, we fit the SNR-vs-amplitude data with a linear function and use the fit parameters to convert the generator amplitude to SNR values when analyzing the trigger efficiencies. When converting a generator amplitude to an SNR value, we assign an uncertainty to this conversion using the fit covariance matrix (Cov) and standard error propagation, such that the uncertainty on the SNR value is:

$$\Delta\mathrm{SNR} = \sqrt{\mathrm{Cov}_{00} + A^2\mathrm{Cov}_{11} + 2A\mathrm{Cov}_{01}}, \tag{5.4}$$

where A is the amplitude of the waveform generator.

To determine the trigger efficiency for each generator amplitude and trigger threshold, we

divide the number of trigger waveforms by the number of injected waveforms, calculated as the product of the measurement duration and waveform generator rate. Additionally, we apply two quality checks to confirm that the triggered waveform originates from the waveform generator and not thermal noise. First, we examine the clock channel for the presence of a clear waveform, as described above. Second, the maximum amplitude of the triggered waveform must occur in the second half of the waveform, as we expect from the trigger configuration. Only if a waveform passes both checks do we include it in the trigger efficiency calculation. The counting of triggered events follows a binomial process, as each event either triggers the system or does not. Consequently, for each observed (N) and expected number of events (N_{exp}), we can construct a binomial distribution as a function of the probability that an event triggers, which is the trigger efficiency (ϵ_{trig}).

$$B(\epsilon_{\text{trig}}) = \binom{N_{\text{exp}}}{N} \epsilon_{\text{trig}}^{N} (1 - \epsilon_{\text{trig}})^{N_{\text{exp}} - N} \tag{5.5}$$

We use the 68% quantile of this distribution as the uncertainty of the trigger efficiency for each data point. Furthermore, we observed that the expected number of waveforms can vary slightly between measurements, despite constant measurement durations and signal rates from the waveform generator. Up to now, we have not been able to identify the cause of these variations. Therefore, we assume a 2% uncertainty for the expected number of events. To account for this uncertainty, we randomly sample 1000 values for N_{exp} from a uniform distribution within the range $[294, 306]$, corresponding to the 2% uncertainty. If the lower bound of this range is less than the number of triggered events N, we replace it by N. For each sampled N_{exp}, we calculate the trigger efficiency and its 68% quantile. Afterwards, we calculate the mean of these trigger efficiency values, along with the mean and standard deviation of the 68% quantiles. The mean trigger efficiency serves as the final value for a measured amplitude and threshold, while we define the uncertainty as the mean of the 68% quantiles plus or minus the standard deviation. As an example, we show the measured trigger efficiency curves of channel 15 for three different trigger thresholds in Figure 5.16.

To estimate the SNR at which the trigger efficiency reaches 50% ($\text{SNR}(\epsilon_{\text{trig}} = 50\%)$), we fit the following function to the measured efficiencies:

$$f(\text{SNR}) = a \left(\tanh \left(\frac{\text{SNR} - b}{c} \right) + 1 \right), \tag{5.6}$$

In this function, c is a scaling factor that determines the steepness of the increase, while a sets the maximal value of the function. The parameter b describes the point at which the function exhibits its steepest slope, and if a equals 0.5, then b corresponds to $\text{SNR}(\epsilon_{\text{trig}} = 50\%)$. As shown in Figure 5.16, the fitted trigger efficiency curves do not always reach a maximal value of 1, so a does not necessarily equal 0.5. Consequently, we have to calculate $\text{SNR}(\epsilon_{\text{trig}} = 50\%)$ by inverting Equation 5.6, resulting in the following formula:

$$\text{SNR}(\epsilon_{\text{trig}} = 50\%) = \text{arctanh} \left(\frac{0.5}{a} - 1 \right) c + b \tag{5.7}$$

For the calculation of the uncertainty of $\text{SNR}(\epsilon_{\text{trig}} = 50\%)$, we apply standard error

propagation for correlated quantities. This yields the following formula:

$$\begin{aligned}\Delta\mathrm{SNR}(\epsilon_{\mathrm{trig}} = 50\%) = \Bigg[& \left(\frac{2c}{1-4a}\right)^2 \mathrm{cov}_{aa} + \mathrm{cov}_{bb} + \left(\operatorname{arctanh}\left(\frac{0.5}{a} - 1\right)\right)^2 \mathrm{cov}_{cc} \\ & + 2\frac{2c}{1-4a}\mathrm{cov}_{ab} + 2\frac{2c}{1-4a}\operatorname{arctanh}\left(\frac{0.5}{a} - 1\right)\mathrm{cov}_{ac} \\ & + \operatorname{arctanh}\left(\frac{0.5}{a} - 1\right)\mathrm{cov}_{bc}\Bigg]^{\frac{1}{2}},\end{aligned} \tag{5.8}$$

Here, cov_{ij} denotes the covariance matrix of the fit, and a, b, and c are the fit parameters. Calculating the fit for each trigger efficiency curve with both SNR uncertainty and asymmetric trigger efficiency uncertainties is too complex for the purpose of these measurements. Therefore, we reduce complexity by making the trigger efficiency uncertainties symmetric, conservatively using the larger of the two asymmetric uncertainties. Additionally, we constrain the trigger efficiency, including its uncertainty, not to exceed 1. With symmetric uncertainties, the fitting process is simplified, allowing us to use existing Python tools. The uncertainty introduced by this simplification is negligible for the two use cases. In the first case, the primary objective is to qualitatively determine whether modifications to the trigger path improve the trigger performance. In the second case, where we create a mapping for the measured trigger threshold to a simulation trigger threshold, the dominant uncertainty originates from channel-to-channel variations.

As previously noted, the trigger efficiency curves sometimes converge to a value slightly below 1. We observe this effect primarily in lower-threshold measurements, where an increased sensitivity results in a higher trigger rate. One possible explanation is that the Radiant approaches its limit for writing and processing events due to the large number of noise-triggered events, leading to some signal events not being recorded. To reduce the trigger rate at low thresholds, we could enable the coincidence trigger requirement, which enforces the signal and clock channel coincidence on the Radiant board. In contrast, for this analysis, we require the coincidence on the analysis level. This approach also allows us to investigate the hypothesis that the inefficiency of the Radiant trigger is attributable to its high noise trigger rate by comparing measurements with the coincidence requirement at the Radiant and at the analysis level. However, we did not perform such comparisons, as the goal of this work is not to provide a comprehensive characterization of the Radiant trigger.

5.3.3 Trigger modifications

In 2023, a new revision of the Radiant board was developed to address previously identified issues. In the following, we will evaluate the impact of various modifications to the signal and trigger path on the overall trigger performance. As the performance metric, we use trigger efficiency curves measured in the laboratory before and after the modifications. To compare the measured efficiency curves, we calculate the SNR at which the trigger efficiency curve reaches 50% (SNR($\epsilon_{\text{trig}} = 50\%$)).

Multiple modifications to the signal and trigger path were proposed by our collaborator Eric Oberla to improve the performance of the Radiant trigger and subsequently were implemented on a spare Radiant board by the DESY electronic workshop. In the following, we provide a brief overview of these changes. Figure 5.17 shows the circuit diagram of a Radiant channel, with the modified components highlighted.

Capacitors: Three capacitors on the signal path (C1x, C6x and C7x, marked in Figure 5.17 in pink, part number: C0402C221J5RACAUTO) with a capacity of 220 pF were replaced with 1 nF capacitors (part number: CGA2B2C0G1H102J050BA). The impedance of a capacitor is inversely proportional to the product of the angular frequency (ω) and the capacitance (C). Consequently, a capacitor with higher capacitance has lower impedance, thereby reducing the signal attenuation along the path.

Amplifier: An amplifier on the trigger path (U1x, marked in Figure 5.17 in orange, part number: BGA2851,115) is specified to operate at 5 V, with a minimum operation voltage of 4.5 V [194]. However, on the V2 board, it is supplied with only 3 V, which is outside the recommended operation range. Therefore, the amplifier is replaced with a version designed to operate with 3 V (part number: BGA2802,115).

Coupler: As discussed above, the trigger path signal is diverted from the signal path using a directional coupler (T1x, marked in Figure 5.17 in blue, part number: TCD-10-1WX+), resulting in a significantly reduced signal on the trigger path compared to the input signal. To address this, the coupler is replaced with a -3 dB splitter (part number: SBTC-2-10LX+), which is expected to increase the signal on the trigger path.

Prior to any modifications, we conducted baseline measurements of the trigger efficiency on channels 15 and 16 of the spare Radiant board. We analyzed these measurements, as described in Section 5.3.2, and present the resulting values of SNR($\epsilon_{\text{trig}} = 50\%$) in Figure 5.18. Afterwards, capacitors were replaced on channel 15 and the amplifier on channel 16, and we repeated the trigger efficiency measurements for both channels. Again, we present the resulting values of SNR($\epsilon_{\text{trig}} = 50\%$) in Figure 5.18. In both cases, the trigger performance improved because the trigger efficiency reached 50 % at a lower SNR, which indicates a more sensitive trigger. Additionally, we find that the improvement from changing the amplifier is slightly greater than that from changing the capacitors. In the next step, the amplifier on channel 15 was also replaced, combining both modifications on the same trigger circuit. As before, we repeated the trigger efficiency measurement for channel 15, and we present the results in Figure 5.18. The results indicate that the combined modifications enhance the trigger performance more than either change individually. Based on these measurements, it was decided to incorporate the modified amplifier and capacitors into the new Radiant revision.

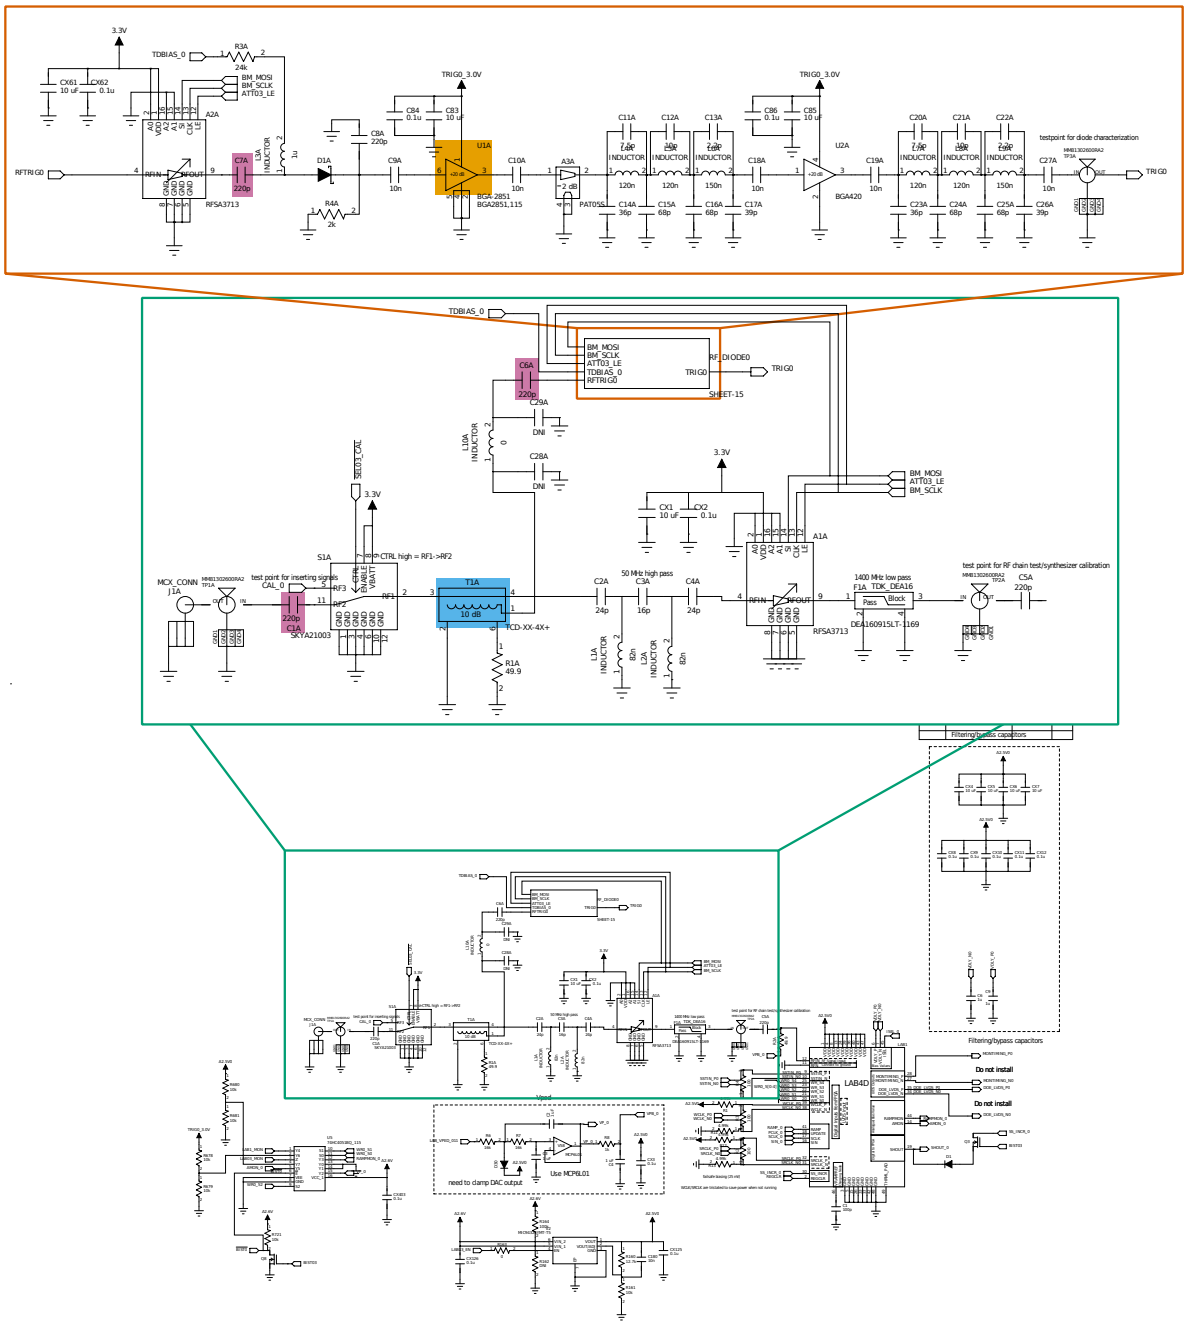


Figure 5.17: Circuit diagram of a RADIANT channel. Two zoomed-in regions provide a detailed view of the trigger circuit, and the highlighted components indicate those that were modified (pink: capacitors, blue: coupler, orange: amplifier). The circuit diagram is provided by the RNO-G Collaboration.

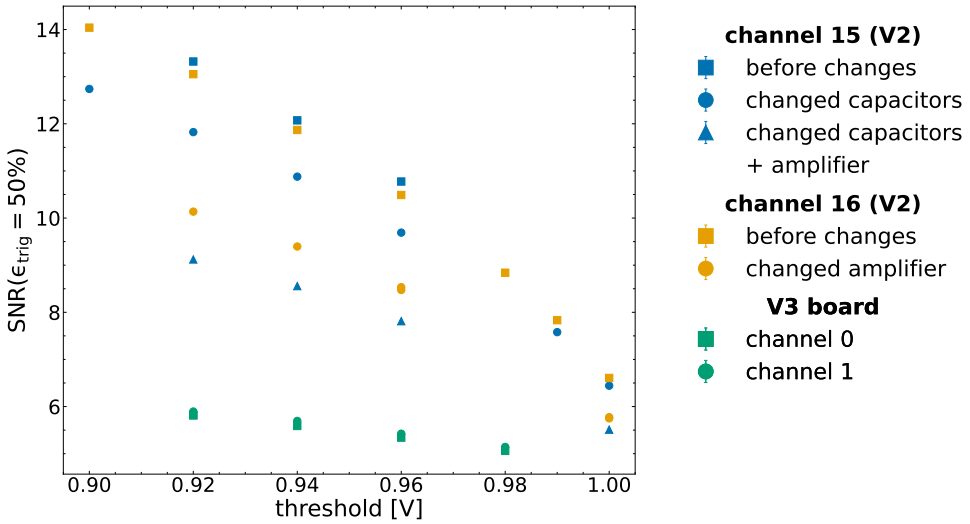


Figure 5.18: The improvement in trigger efficiency resulting from various modifications to the signal and trigger path. The V3 Radiant board includes all modifications (capacitors, amplifier and coupler).

Although we did not measure the effect of changing the coupler before finalizing the new revision, this modification was included. To assess the overall improvements to the trigger path and estimate the impact of the coupler change, we repeated the trigger efficiency measurement on channels from the new revision of the Radiant board (V3). Because we conducted these measurements on a different board, it does not matter for the overall conclusion that we used a different set of channels. The results of this measurement, presented in Figure 5.18, demonstrate that the combined modifications realized on the V3 board yield a significant improvement over the trigger performance of the V2 board. Comparing the V3 board measurements with those obtained after both the amplifier and capacitors were changed allows us to estimate the contribution from the change of the coupler. This comparison indicates that the most substantial single improvement results from the coupler modifications. Notably, the absolute improvement from these changes decreases as the threshold value increases. The reason for this behavior remains unknown.

The V3 Radiant board, including the discussed modifications, was deployed in the field during the 2024 field season. During this period, a new station was commissioned with the updated V3 board, and six of the seven existing stations were retrofitted with the V3 board. Consequently, we expect improved trigger performance at the majority of the operational stations. Further details regarding the Radiant V3 board, its testing, and performance are available in Ref. [172].

5.3.4 Trigger threshold in UHECR simulations

Previous studies have found that accurately modelling the trigger diode, and therefore the change of the waveform along the Radiant trigger path, is challenging [172]. As a result, we approximate the Radiant trigger for all simulations in this work using a bandpass filter and an envelope trigger module available in NuRadioReco [195]. Consequently, the trigger threshold we measure on the Radiant (τ_{measured}) does not correspond directly to the threshold we use for the envelope trigger ($\tau_{\text{simulation}}$). To

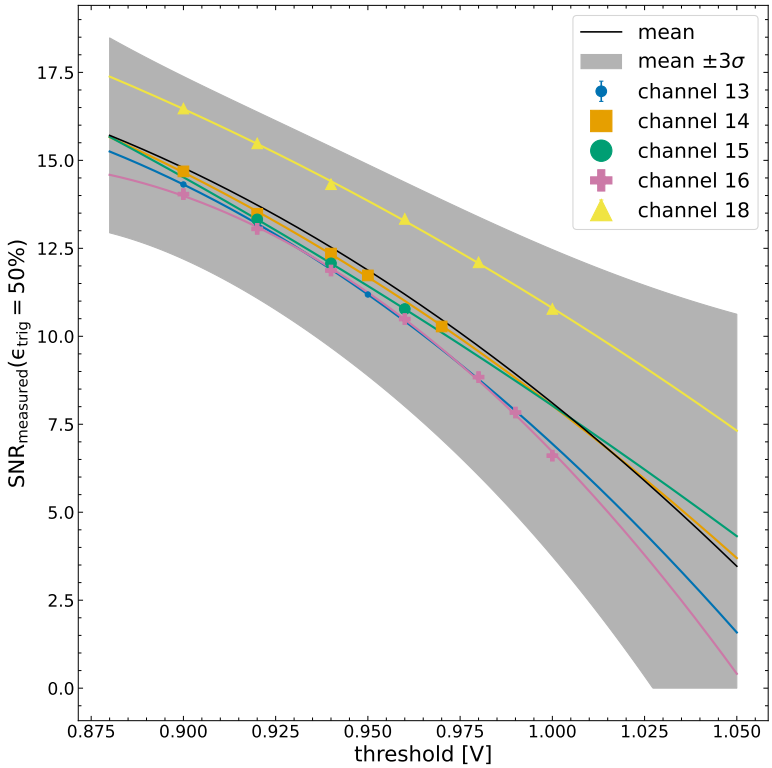


Figure 5.19: The SNR value at which the trigger efficiency reaches 50% for different channels on a spare V2 RADIANT board, as a function of the trigger threshold. The data points are fitted with a quadratic function.

achieve a realistic representation of the RADIANT trigger, we require a mapping between τ_{measured} and $\tau_{\text{simulation}}$. We establish this mapping by relating laboratory measurements of the trigger efficiency to simulated trigger efficiency curves, as we describe in the following.

Laboratory measurements: For different channels on a spare V2 RADIANT board, we performed trigger efficiency measurements at fixed thresholds by injecting UHECR-like waveforms into the RADIANT board, following the setup, procedure, and analysis we described in Section 5.3.2. Examples of the measured trigger efficiency curves are presented in Figure 5.16. From these measurements, we derive, per channel, a mapping between τ_{measured} and the SNR value at which the trigger efficiency reaches 50% ($\text{SNR}_{\text{measured}}(\epsilon_{\text{trig}} = 50\%)$), as shown in Figure 5.19. Each data point includes an error associated with the measured $\text{SNR}_{\text{measured}}(\epsilon_{\text{trig}} = 50\%)$ value, although this error is too small to be visible in the figure. For each channel, we fit the data points to a quadratic function to describe the relationship. The dominant uncertainty of the mapping arises from the channel-to-channel variations, visible in Figure 5.19, rather than from the errors of the quadratic fit. In the following, we consider only the uncertainty due to this channel-to-channel variation.

To combine all channel measurements, we calculate the mean across all channels as a function of the threshold, and we conservatively estimate the uncertainty as $\pm 3\sigma$, where σ is the standard deviation of the measurements. In Figure 5.19, we indicate the mean by a black line, and the uncertainty as a grey area. At higher threshold values, extrapolating the uncertainty region would yield negative SNR values, which do not

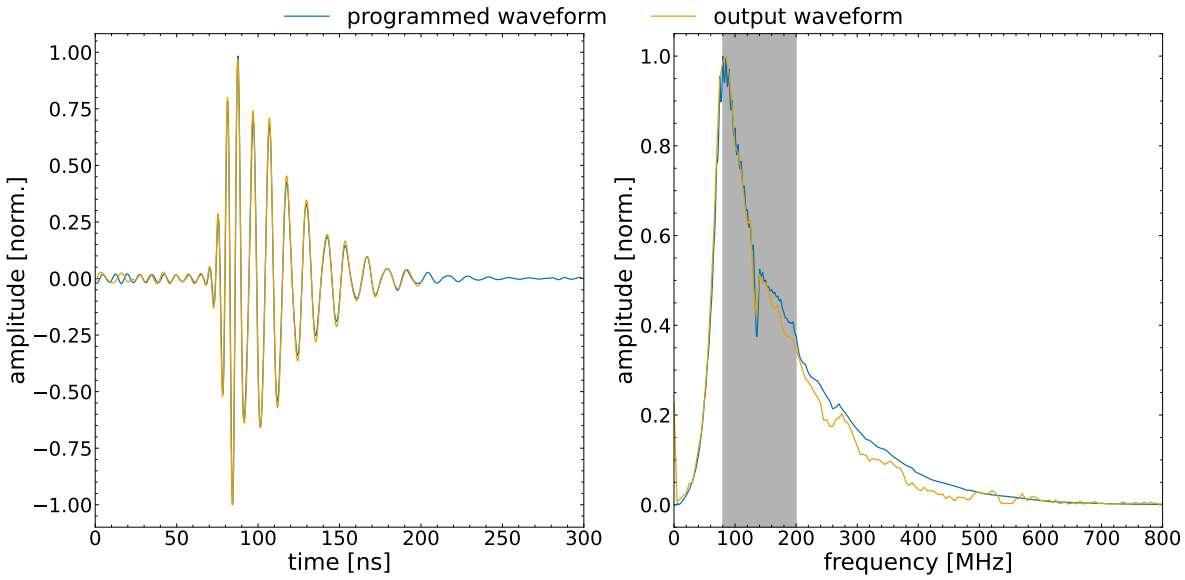


Figure 5.20: Left: A comparison of the programmed and output waveforms of the arbitrary waveform generator, with both waveforms filtered to the trigger-relevant frequency band (80 – 200 MHz). Right: The unfiltered frequency spectra of both waveforms (shown in the left panel). The grey band indicates the trigger-relevant bandwidth.

make sense for our SNR definition[8], and therefore, the minimal possible SNR value is set to zero. Generally, we apply the mapping only to threshold values that result in $\mathrm{SNR}_{\mathrm{measured}}(\epsilon_{\mathrm{trig}} = 50\%)$ values larger than the expected SNR from thermal noise[9]. The arbitrary waveform generator we used for the laboratory measurements has a limited bandwidth. To assess the match of the output waveform relative to the programmed UHECR waveform, we recorded waveforms with an oscilloscope directly connected to the waveform generator output. Then, we compare the recorded waveform with the original UHECR waveform, as shown in Figure 5.20. Within the frequency band relevant for the trigger (80 – 200 MHz [163]), both waveforms exhibit strong similarity, indicating that the limited bandwidth of the waveform generator does not introduce an additional relevant uncertainty compared to the channel-to-channel variations. Furthermore, in Ref. [172], the influence of different UHECR-like input waveforms on the trigger efficiency measurements was examined and found to be sufficiently small to justify the use of a single waveform type for the measurements presented here.

Simulated trigger efficiencies: We reproduce the same relationship we established in the laboratory measurements in simulations, mapping $\tau_{\mathrm{simulation}}$ to $\mathrm{SNR}_{\mathrm{sim}}(\epsilon_{\mathrm{trig}} = 50\%)$ by simulating trigger efficiency curves. To create the simulated trigger efficiency curves, we generate multiple instances of the same UHECR-like waveform with varying thermal and Galactic noise, and apply the envelope trigger. As in the laboratory measurements, we consider a single channel and calculate the trigger efficiency from the number of waveforms that fulfil the trigger condition and that are generated. Furthermore, we calculate the uncertainty of the trigger efficiency similarly to the laboratory measurement (see Section 5.3.2 for more information), with the main difference that we precisely know

[8] The SNR is defined as the maximum of the absolute values of the waveform samples divided by the noise RMS of the waveform (see Equation 5.3).
[9] The thermal noise SNR is around 4, but varies slightly across stations.

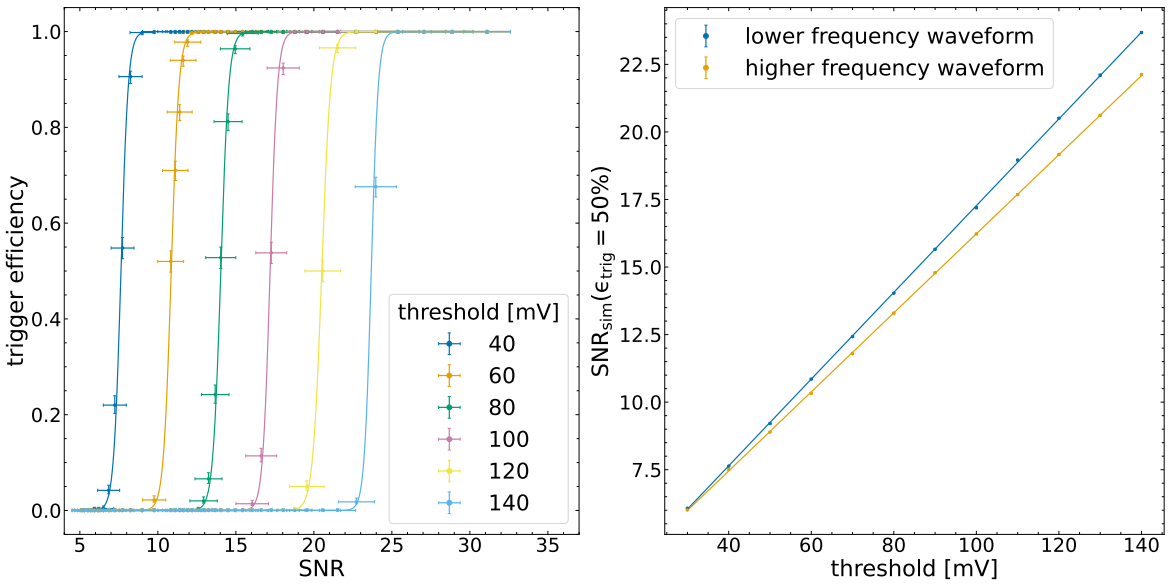


Figure 5.21: Left: Simulated trigger efficiency curves for selected example thresholds. Right: The simulated mapping between the $\tau_{\text{simulation}}$ and the $\text{SNR}_{\text{sim}}(\epsilon_{\text{trig}} = 50\%)$ value. Two mappings are presented, corresponding to the two UHECR-like waveforms used in the simulations.

the number of generated waveforms. A set of 500 simulated UHECR-like waveforms provides the trigger efficiency for a single threshold and SNR value. By varying the amplitude of the UHECR-like waveform while maintaining a constant noise RMS, we simulate the trigger efficiency at different SNR values, to produce a simulated trigger efficiency curve. For each data point, we simulate 500 UHECR-like waveforms, each with a slightly different SNR value due to different noise realizations. Therefore, we calculate the mean SNR value for those waveforms at each data point and use the standard deviation as the uncertainty for the SNR value. Simulated example trigger efficiency curves for different thresholds are presented in Figure 5.21. We fit the data points to the same function as in the laboratory measurements (see Equation 5.6), and the efficiency uncertainty is again conservatively symmetrized to simplify the fitting procedure. Unlike the laboratory measurements, the simulated trigger efficiency curves all reach a maximum value very close to 1. Consequently, we can directly interpret the fit parameter describing the turning point of the function (parameter b in Equation 5.6) as the value of $\text{SNR}_{\text{sim}}(\epsilon_{\text{trig}} = 50\%)$, with the associated uncertainty given by the error of the fit parameter.

Repeating the simulation for different trigger thresholds results in a mapping from $\tau_{\text{simulation}}$ to $\text{SNR}_{\text{sim}}(\epsilon_{\text{trig}} = 50\%)$, as shown in Figure 5.21. Although all data points in the figure have associated uncertainties, these are too small to be visible. To span the entire threshold parameter space, we fit the data points with a linear function. Additionally, we repeated the simulation with a UHECR-like waveform with a different frequency content, and the resulting mapping is also shown in Figure 5.21. The difference between the simulations using different waveforms is the dominant source of uncertainty relative to the uncertainty of the linear fit, and thus is the only uncertainty we consider on the simulation side for the subsequent analysis. For the UHECR-like waveforms, we use two templates (with widths of 0.625 ns and 1.5625 ns) that we developed for the shallow UHECR analysis (see Section 8.2 for more information). The primary difference between the two templates is that one exhibits a higher frequency content than the other.

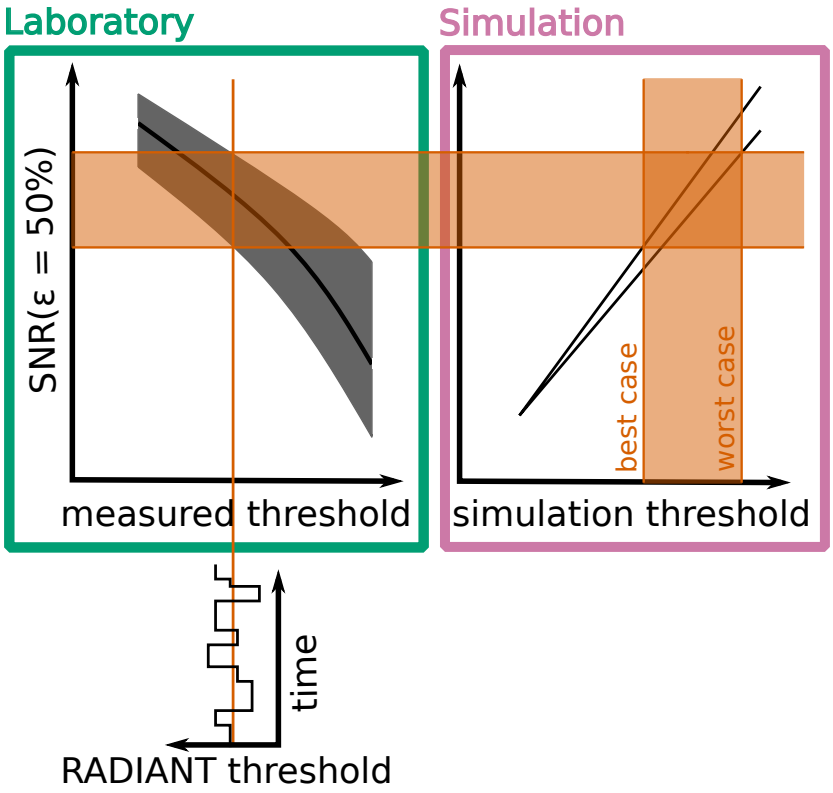


Figure 5.22: Schematic representation of the process by which a measured RADIANT trigger threshold is mapped to a corresponding range of simulation thresholds.

Estimation of the simulation threshold from data: As we already mentioned in Section 4, RNO-G has no fixed trigger threshold, but the threshold is dynamically adjusted to the noise level, which is influenced by both background noise and temperature. Therefore, we map, for each considered time period, the mean τ_{measured} to a mean $\tau_{\text{simulation}}$. In general, we achieve this translation by matching the $\text{SNR}_{\text{measured}}(\epsilon_{\text{trig}} = 50\%)$ to the $\text{SNR}_{\text{sim}}(\epsilon_{\text{trig}} = 50\%)$. A schematic illustrating this mapping is provided in Figure 5.22. First, we translate the mean τ_{measured} measured during the considered data-taking period to the $\text{SNR}_{\text{measured}}(\epsilon_{\text{trig}} = 50\%)$ value, which we then set equal to $\text{SNR}_{\text{sim}}(\epsilon_{\text{trig}} = 50\%)$. Subsequently, we convert the $\text{SNR}_{\text{sim}}(\epsilon_{\text{trig}} = 50\%)$ value into a $\tau_{\text{simulation}}$ value. Due to the significant uncertainties, which can be as high as $\mathcal{O}(50\%)$, we report, for each τ_{measured} both a best- and worst-case simulation trigger threshold. We account for the uncertainties in the measured and simulated mappings by consistently applying the more conservative option, as indicated in Figure 5.22: for the worst-case scenario, we report the larger (less sensitive) threshold, while for the best-case scenario, we report the smaller (more sensitive) threshold.

The procedure we outlined in this section converts a trigger threshold measured on the RADIANT board into a simulation voltage threshold, thereby enabling simulations with realistic trigger thresholds. This is essential for predicting the number of UHECR events expected in a dataset. Given the substantial uncertainties, particularly those associated with the channel-to-channel variations in the laboratory measurements, we report upper (worst-case) and lower (best-case) limits for the simulation threshold for each measured threshold.

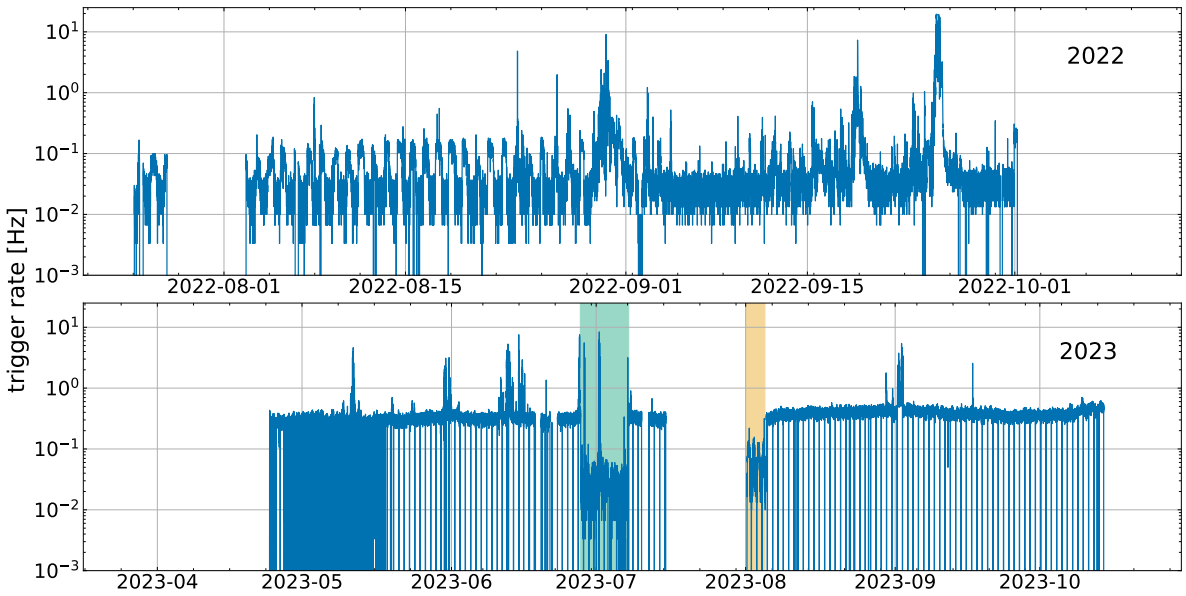


Figure 5.23: The 2022 (top) and 2023 (bottom) trigger rate for both operational RADIANT triggers (one using only the upward-facing antennas and the other using only the downward-facing antennas) combined. The colored regions highlight notable periods in 2023. Further details about these regions are provided in the main text. The regular decreases in the trigger rate observed in 2023 correspond to routine measurements required for voltage calibration, during which data acquisition is temporarily halted.

5.4 Unusual station behavior

During the 2023 data-taking period, we observed anomalous behavior in station 23 compared to the previous data-taking period. Once the station had been activated after the winter break, the shallow trigger operated at an elevated rate. In Figure 5.23, we present the trigger rate of station 23 for both the 2022 and 2023 periods. We display the combined trigger rate for both operational RADIANT trigger, which uses a two antenna coincidence of either the upward or downward-facing antennas. Although RNO-G automatically adjusts the trigger thresholds to stabilize the trigger rate, the coincidence trigger rate of multiple channels may still vary because only the single-antenna trigger rate is stabilized. Consequently, when the amount of correlated noise increases, the antenna-level trigger rate can be stabilized, while the coincidence trigger rate across multiple channels increase. Notably, in 2023, two brief intervals occurred in which the trigger rate reverted to the 2022 level (highlighted in green and yellow in Figure 5.23). The first interval (green) corresponds to the time period where we deactivated the RADIANT trigger on the downward-facing antennas during the investigation of this abnormal behavior. The second period (yellow) occurred under standard settings. A close inspection of the 2022 trigger rate reveals an increase to the 2023 trigger rate level for the last approximately six hours before deactivating the station for the winter, suggesting the phenomenon may have already appeared in 2022. However, due to its short duration, the 2022 period is not treated as a separate configuration in the main analysis.

First, we confirmed that the trigger settings remained unchanged between the time periods before and after the increase in the trigger rate, eliminating a configuration change as the cause. Afterwards, we inspected the recorded trigger thresholds for station

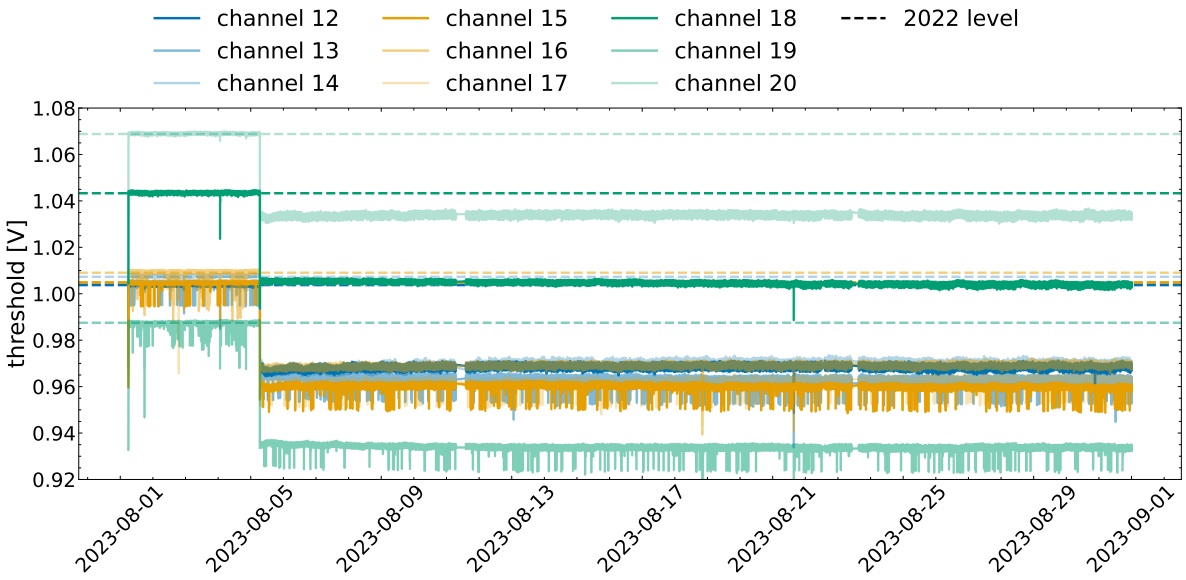


Figure 5.24: The RADIANT thresholds measured at station 23 in August 2023. For reference, the mean RADIANT thresholds measured in August 2022 are indicated by the dashed lines.

23 in August 2023, a period when the trigger rate temporarily returned to the 2022 level. The thresholds are displayed in Figure 5.24, together with the mean threshold per channel from 2022. The trigger thresholds in 2023 are lower than in 2022, resulting in a reduced station sensitivity. Only at the beginning of August, when the trigger rate temporarily returned to the 2022 level, did the thresholds match the 2022 level. Therefore, these results suggest that the increased trigger rate does not originate from an increased sensitivity of the station due to a threshold change. A plausible explanation is an increase in correlated noise in 2023, which led to threshold adjustments to stabilize the single-antenna trigger rate, while the coincidence trigger rate across multiple channels increased due to the enhanced correlation of the noise.

Next, we investigate the time period in August 2023 when the trigger rate abruptly returned to the 2022 level for a short time. Figure 5.25 shows a detailed view of the trigger rate during the August period, around the time it again increased to the 2023 level, and the vertical lines mark the initiation of new runs. In RNO-G, data acquisition is organized into two-hour runs, after which a new run begins, and a configuration file with all instrument settings is loaded. These instrument settings, such as trigger settings, can only be changed at the start of a run and remain constant during the run. The observed change in the trigger rate occurred within a run, as shown in Figure 5.25, indicating that we cannot attribute it to any modification of the data-taking settings. Furthermore, we can exclude calibration measurements or RADIANT calibrations as causes, as these are only performed between runs. In the following, we further analyze the affected run.

We investigated the thresholds during the run and observed the previously described and expected behavior: all threshold values decreased, indicating reduced station sensitivity, as the trigger rate increased again (see Figure A.1). Furthermore, we examined the RADIANT scaler, which measures the number of times the trigger diode output is below the threshold [172]. These scalers are used to adjust the trigger threshold by setting a target, and the RADIANT attempts to achieve that target by adjusting the thresholds. We did not observe a significant change in the scaler values when the trigger rate changed,

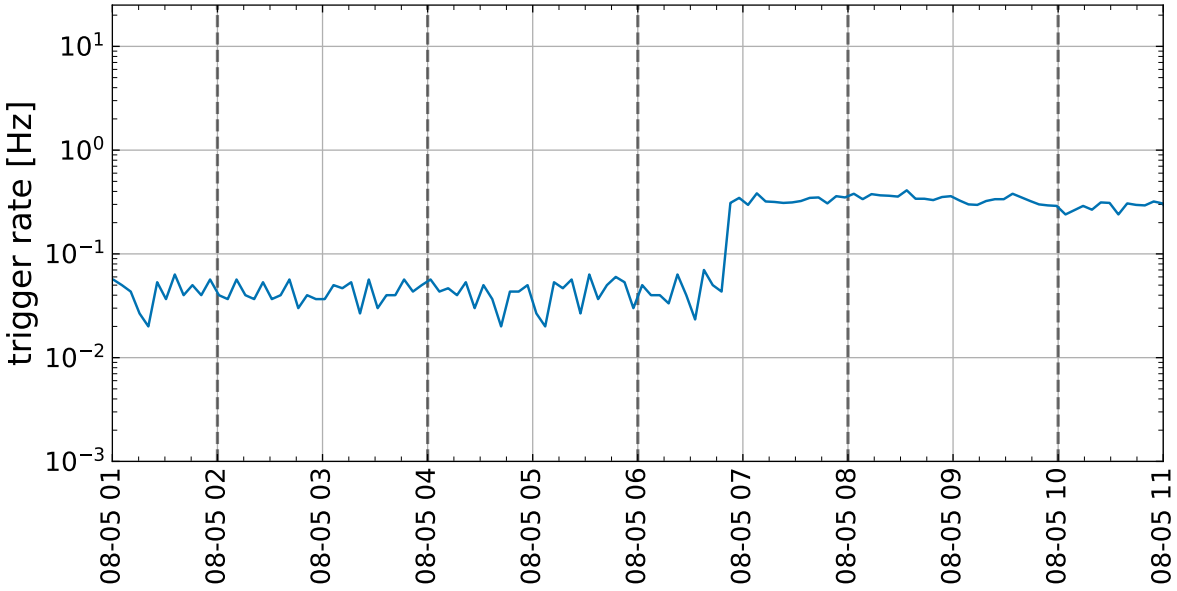


Figure 5.25: The trigger rate for both operational RADIANT triggers (one using only the upward-facing antennas and the other using only the downward-facing antennas) combined. Only a short time period from August 2023 is shown around the time the trigger rate increased from the 2022 to the 2023 level. The vertical lines indicate the start of new data runs.

except for a slight increase in the scaler value spread after the change (see Figure A.2). However, this observation does not allow us to determine the cause of the change. Also, we analyzed the RMS values of the events within the run as a function of the event time (see Figure A.3). All channels, except channel 12, exhibit a normal behavior and do not change when the trigger rate changes. Channel 12 seems to have large block offsets[10], and after correcting for these block offsets, the abnormal behavior of the RMS disappears. Since the block-offsets are introduced during digitization, they do not affect the trigger decision and therefore cannot explain the observed trigger rate behavior. Finally, we visually inspected the average frequency spectrum of the run before and after the change in trigger rate (see Figure A.4). No significant change in the average spectrum is observable that could explain the trigger rate increase.

To identify the channels that contribute most to the increased trigger rate, we systematically removed channels from the trigger condition during data-taking. The resulting trigger rate from this investigation is shown in Figure 5.26. Initially, we switched to a configuration using only downward-facing antennas for triggering, but this did not result in a significant decrease in the trigger rate, as shown in Figure 5.26. This indicates that channels with downward-facing antennas (channels 12, 14, 15, 17, 18, 20) primarily drive the elevated trigger rate. Subsequently, we excluded individual channels from the downward-facing trigger condition, which requires a 2-out-of-6 coincidence. The most substantial changes in the trigger rate occurred when channels 12 and 20 were excluded, suggesting that these channels are the primary contributors to the higher trigger rate. A station diagram (see Figure A.5) shows that channels 12 and 20 are adjacent and located near the second solar panel, which, in principle, could be a possible noise source. However, more studies are necessary to conclusively identify the solar panel as the potential origin of the increased trigger rate.

[10] Artifact of the waveform digitization where blocks of 128 samples have a relative offset to each other and to the baseline [121].

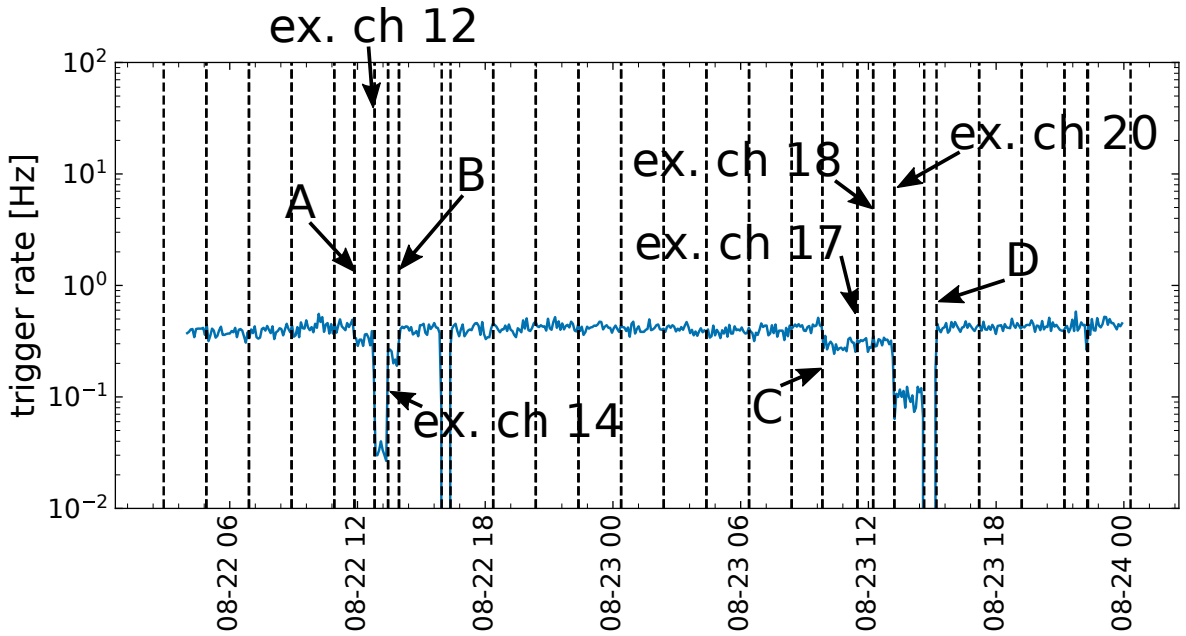


Figure 5.26: The combined trigger rate for both operational RADIANT triggers at station 23 during a 2023 time period when we modified the trigger configuration to investigate the elevated trigger rate. Horizontal lines denote the start of new runs. Line A marks the switch to a pure downward-facing trigger, line B marks the return to the standard configuration, line C marks the switch to a downward-facing trigger with channel 15 excluded from the trigger conditions, and line D marks the return to the standard settings.

Overall, we determined that the change in trigger rate for station 23 cannot be attributed to modifications in the station settings and can occur within a single run. Despite extensive checks, we could not identify the source of this change, except that it appears to be dominated by the channels 12 and 20.

5.5 Conclusions

In this chapter, we provided a detailed examination of the shallow amplification and digitization chain and the RADIANT trigger .

First, we analyzed different measurements of the shallow amplification and digitization response, with particular attention to their impact on the correlation score between template waveforms and simulated or measured UHECR waveforms. Simulation results indicate that imprecise hardware response descriptions can cause significant changes in the correlation score, of $\mathcal{O}(0.1)$. Additionally, we demonstrated that differences in group delay of the measured responses account for the majority of the observed differences in correlation scores. Furthermore, we investigated how different response measurements affect the correlation score of the template with selected measured UHECR candidate waveforms, specifically three high SNR events we identified in the shallow UHECR analysis presented in this thesis. The comparison revealed that the piecewise description, which combines S21 parameter measurements for all components of the amplification and digitization chain into a full response, exhibits shortcomings that result in relatively low correlation scores. This limitation may arise from relying solely on the S21 parameter, which neglects the effects of reflections within the system. To address this issue, efforts

are ongoing to implement a comprehensive treatment of all S parameters. Additionally, the S-parameter measurement of the RADIANT board is limited to a subset of the circuit, which further constraints the description of the hardware response using S-parameter measurements.

Furthermore, we introduced a more comprehensive measurement of the entire amplification and digitization chain: a time-domain impulse response measurement that avoids the limitations of the piecewise description. This impulse response measurement most accurately characterizes the hardware response and currently represents the best available description for the shallow RNO-G hardware response. Consequently, we utilize this response in the shallow UHECR analysis presented in this work.

In summary, an accurate characterization of the amplification and digitization chain is essential for the UHECR analysis, as it relies on the correlation score, which is sensitive to the waveform shape. Further, the thorough understanding we gained of the different measured hardware response descriptions is valuable for all RNO-G analyses that depend on signal-shape-dependent parameters, and constitutes an important step toward future RNO-G neutrino analyses.

In this chapter, we also introduced the RNO-G hardware and calibration database, which serves as a platform for storing and managing calibration measurements for all components or component chains. Additionally, the database stores the channel and station positions, as well as other calibration data required to describe the detector. Overall, it enables a time-dependent description of the detector. Consequently, the database provides significant advantages over the previous method, which relied on storing the information in a single file. Integration of the database with the NURADIORECO simulation and reconstruction framework enables the collaboration to efficiently retrieve comprehensive detector descriptions, including positions and hardware responses, for use in simulation and reconstruction. At present, the most accurate hardware response measurement, the impulse response, is only available in the database in a preliminary form due to incomplete measurements of the trigger component of the phased array channels. However, efforts are ongoing to perform these missing measurements and incorporate them in the database. In the future, a substantial database upgrade will be necessary to store the results of the ongoing full detector calibration, including associated uncertainties, and to make this information broadly available to the collaboration.

Furthermore, we studied the RADIANT trigger in this chapter. Measurements of the initial deployed version of the RADIANT board (V2) revealed that noise originating from the board itself dominates the trigger path. When events are triggered by this board noise, they are indistinguishable from external thermal noise in the waveform recorded on the signal path. Therefore, despite the trigger threshold of the RADIANT board being set above the external thermal noise level, the recorded datasets still contain events resembling thermal noise. Further, we investigated how modifications to components along the signal and trigger path would affect the trigger performance. We measured substantial improvements, leading to the replacement of multiple components during the RADIANT board redesign. The updated RADIANT board (V3), featuring the enhanced RADIANT trigger performance, was deployed during the 2024 campaign in a newly built station and replaced the previous version in nearly all existing stations. Currently, seven out of eight deployed stations operate with the improved RADIANT boards. Additionally,

we established a mapping between the measured trigger threshold in the data and the simulation thresholds for an effective trigger description. This mapping enables the creation of realistic UHECR simulations, which are essential for estimating the expected number of UHECR events in the shallow UHECR analysis presented in this work.

Lastly, we investigated an unusual behavior of station 23 during the 2023 data-taking period, resulting in an elevated trigger rate compared to 2022. We can exclude changes in the trigger settings as the cause and have identified a run in which the trigger rate changed from the 2022 to the 2023 level within it. Additionally, we performed several checks and found that two channels, 12 and 20, primarily contribute to the elevated trigger rate. However, we could not identify the underlying reason for this abnormal behavior. Future work is necessary to improve the understanding of this behavior, which is important because the associated change in the trigger threshold alters the sensitivity of the station and thus the quality of the recorded data.

6 Refractive Index of Ice

In this chapter, we examine the influence of the refractive index of ice on the shallow UHECR analysis and introduce a promising method for measuring it. The antennas we use for the UHECR analysis are positioned at shallow depths, and the analysis focuses on signals originating from above the ice. Consequently, these signals traverse only a short distance through the ice, and the refractive index variation along their path is minimal. Under these conditions, the effect of ray bending is negligible, and ray tracing can be omitted without introducing significant uncertainties. However, as described in Section 4.1, the VEL of an antenna depends on the surrounding medium and its refractive index. In addition, the antenna VEL significantly influences the recorded waveform, and the UHECR analysis depends on a precise understanding of the signal shape. As a consequence, the analysis is sensitive to the refractive index around the antennas.

In the following sections, we introduce different ice models to describe the refractive index profile at the detector site (see Section 6.1), investigate the refractive index at shallow depth (see Section 6.2) and introduce an effective description of the VEL for different refractive indices (see Section 6.3). Furthermore, we examine how small deviations in the refractive index affect the UHECR analysis by simulating different scenarios (see Section 6.4). Lastly, we present a proof-of-principle analysis utilizing ground-penetrating radar data to measure the refractive index profile, which may reduce the uncertainties of the UHECR analysis related to the refractive index in the future (see Section 6.5).

6.1 Ice models

The refractive index profile of glacial ice follows an asymptotic dependency, approaching a constant value at greater depth. This behavior is attributable to its dependence on the ice density, which is assumed to be [9]:

$$n(z) = 1 + 0.854\rho(z) \tag{6.1}$$

Here, $n(z)$ is the refractive index and $\rho(z)$ is the density as a function of depth (z). The glacier density profile changes significantly, beginning at $50 - 70\,\mathrm{kg/m^3}$ for newly fallen snow and asymptotically increasing to densities of $830 - 923\,\mathrm{kg/m^3}$ for glacier ice [196]. The transformation from newly fallen snow to glacier ice occurs over time under the gravitational pressure of the overlying ice and snow, resulting in a density profile that increases with depth. An intermediate stage between fresh snow and glacier ice is the *firn*. Although there is no distinct boundary between snow and firn, firn can be described as snow that has survived at least one summer or, more generally, as altered snow [196]. In contrast, the transition from firn to glacier ice is clearly defined as the point at which the interconnecting passageways between grains are closed due to the pressure [196]. Typically, the refractive index profile is characterized by a single or multiple exponential functions. In the following, we introduce ice models which are used in the context of RNO-G.

Greenland simple: The Greenland simple ice model employs a single exponential function to describe the density profile, as given by the following equation [180]:

$$\rho(z) = \rho_{\mathrm{I}} - (\rho_{\mathrm{I}} - \rho_{\mathrm{S}}) \exp\left(\frac{z}{z_0}\right) \tag{6.2}$$

Here, ρ_{I} is the density of glacier ice, ρ_{S} is the density at the surface, and z_0 is a scale depth. In principle, the densities are known quantities and only the scale depth requires fitting. For the Greenland simple ice model, the density profile is fitted to density data collected near Summit Station. After fitting, the density profile is converted to a refractive index profile using Equation 6.1. The resulting ice model is:

$$n(z) = 1.78 - 0.51 \exp\left(-z/37.25\,\mathrm{m}\right), \tag{6.3}$$

which was fitted by C.Deaconu to data from [197, 198].

Greenland firn: The Greenland firn model applies the Herron-Langway parametrization, which divides the profile into two regions, each described by an exponential function. The boundary between these two regions is defined as the depth at which the critical density of $\rho_{\mathrm{c}} = 550\,\mathrm{kg/m^3}$ is reached. The ice model is fitted to density data collected near Summit Station, as described in Ref. [199], resulting in the following refractive index profile:

$$n(z) = \begin{cases} 1.78 + 0.502 \exp(-z/30.8\,\mathrm{m}), & z \leq 14.9\,\mathrm{m} \\ 1.78 + 0.310 \exp(-z/40.9\,\mathrm{m}), & z > 14.9\,\mathrm{m} \end{cases} \tag{6.4}$$

Greenland poly5: During calibration, the arrival-time differences between two antennas for a signal from a calibration pulser can be measured. These differences can be compared with the time differences expected from simulations, which depend on the ice model used. The Greenland ploy5 model is an empirical model developed to address unexpected mismatches between measured and simulated inter-antenna timings, observed when using the Greenland simple ice model for the simulation [180]. As the name indicates, the model is an exponential polynomial of the fifth degree:

$$\rho(z) = \rho_{\mathrm{I}} + \sum_{i=1}^{5} a_i \exp^i(-z/z_0), \tag{6.5}$$

where ρ_{I} denotes the density of glacier ice, z_0 is a scale depth, and a_i are fitted coefficients. Further information about the model and the resulting fit parameter are available in Ref. [180].

A comprehensive discussion of the different ice models is provided in Ref. [180]. The ice models introduced thus far are all fitted to ice density measurements. However, alternative methods exist to measure the refractive index profile directly, without relying on density data. In recent work, the refractive index was extracted by measuring the S11 parameter[1] of an antenna as it was lowered into the borehole. The fundamental resonance frequency of the antenna is expected to shift by $1/n$, where n is the refractive index,

[1] An explanation of the S parameters is provided in Section 5.1.

enabling independent measurements of the refractive index profile. Two measurement campaigns were conducted near Summit Station, and an ice model was fitted to the extracted data [200], using a three-stage exponential model described in Ref. [165]. A limitation of this technique is the requirement for an empty borehole, which precludes its use near the detector once the antennas are deployed without drilling a new borehole. Another technique is the time-of-flight (TOF) measurement. In these measurements, two antennas are co-lowered in separate boreholes, with one antenna emitting a signal and the other receiving it. The TOF of the signal allows the refractive index to be calculated along the traveled path. During the 2024 deployment campaign, TOF measurements were performed using empty boreholes drilled for the stations. The analysis of these measurements is ongoing, and preliminary results are promising. Again, the resulting data could be fitted by the three-stage exponential model, enabling the determination of the ice model without density data. However, similar to the S11 measurement, this method requires two empty boreholes in close proximity. Due to time constraints, we were not able to include the three-stage exponential model in this work.

6.2 Shallow refractive index

In the context of the UHECR analysis, the refractive index of the firn at shallow depth is of primary importance, as it directly affects the antenna response and thus the recorded signal shape. To estimate the refractive index at the depths where the LPDAs are deployed, we consider the three ice models presented in Section 6.1. Figure 6.1 shows the refractive index profiles for these three models. Based on the deployment depth and potential snow accumulation during the data-taking periods, we estimate that the LPDAs are located at depths of 0.5 to 3 m. The expected LPDA region is indicated as a grey area in the figure. The possible refractive index around the deployed LPDAs spans a broad range, reflecting both the differences among the ice models and the substantial depth interval. From the ice models, we estimate the refractive index around the LPDAs to be between 1.28 and 1.335.

Notably, these ice models assume an undisturbed medium. However, for the deployment of the LPDAs, trenches are excavated, the antennas are inserted, and the trenches are subsequently refilled with snow. Consequently, air pockets may remain around individual antenna tines, or denser clumps of snow may form. Additionally, movement of personnel around the station perimeter during deployment could result in denser regions in the uppermost layers. In principle, these effects could affect the effective refractive index, thereby influencing the antenna response.

We consider the refractive index range presented here a reliable estimate of the medium surrounding the stations and utilize it in Section 6.4 to assess the influence of the refractive index on the UHECR analysis. The development of ice models remains ongoing. For instance, a new three-stage exponential model was developed, and the Greenland poly5 model has been fitted to alternative data sets, producing slightly different results. Additionally, new measurement techniques, such as S11 parameter measurements, are being developed and implemented. Due to time constraints, we were not able to incorporate the most recent data from these techniques and models into this work. Future studies should revisit the refractive index as a source of systematic

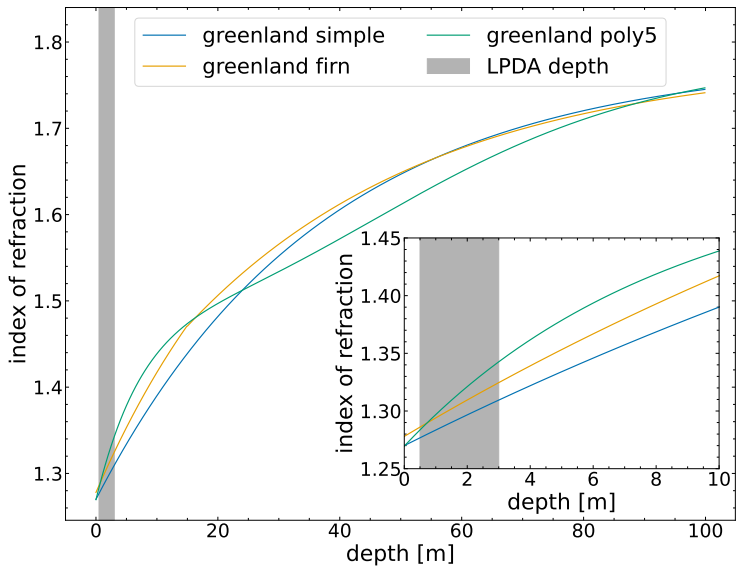


Figure 6.1: The refractive index as a function of depth for different ice models. The grey region denotes the depth range where the LPDAs are expected. The inlet provides a zoomed-in view of the shallow depth region.

uncertainty in the UHECR analysis and incorporate the updated information. Over time, the accumulation of new snow above the antennas is expected to compress the surrounding snow, potentially eliminating small air cavities or denser regions that my persist from the deployment process. Furthermore, due to snow accumulation, the distance between the antennas and the air-snow boundary is increasing, thereby reducing its potential influence on the antenna response.

6.3 Effective antenna description

As we previously discussed in Section 4.1, the VEL varies with the refractive index of the medium surrounding the antenna. To assess how these VEL variations affect the UHECR analysis, we must generate templates for multiple refractive indices, which requires knowledge of the VEL at each refractive index. The VELs are simulated using electromagnetic simulation software programs such as WIPL-D [201] or XFdtd [202]. Currently, we only have access to WIPL-D simulations for refractive indices of $n = 1, 1.3$, and 1.4. Therefore, we need to adapt one of these existing simulations to estimate the VEL at any given refractive index. We achieve this adaptation by transforming the frequency and magnitude of the VEL according to the following relations, which are motivated by Equation 4.1:

$$f_2 = \frac{f_1 n_1}{n_2} \tag{6.6}$$

$$\text{abs}(\text{VEL}_2) = \frac{\text{abs}(\text{VEL}_1)\sqrt{n_2}}{\sqrt{n_1}} \tag{6.7}$$

Here, index 1 refers to the frequency (f) and VEL at the refractive index (n) before the transformation, while index 2 denotes the effective description of the VEL at the new

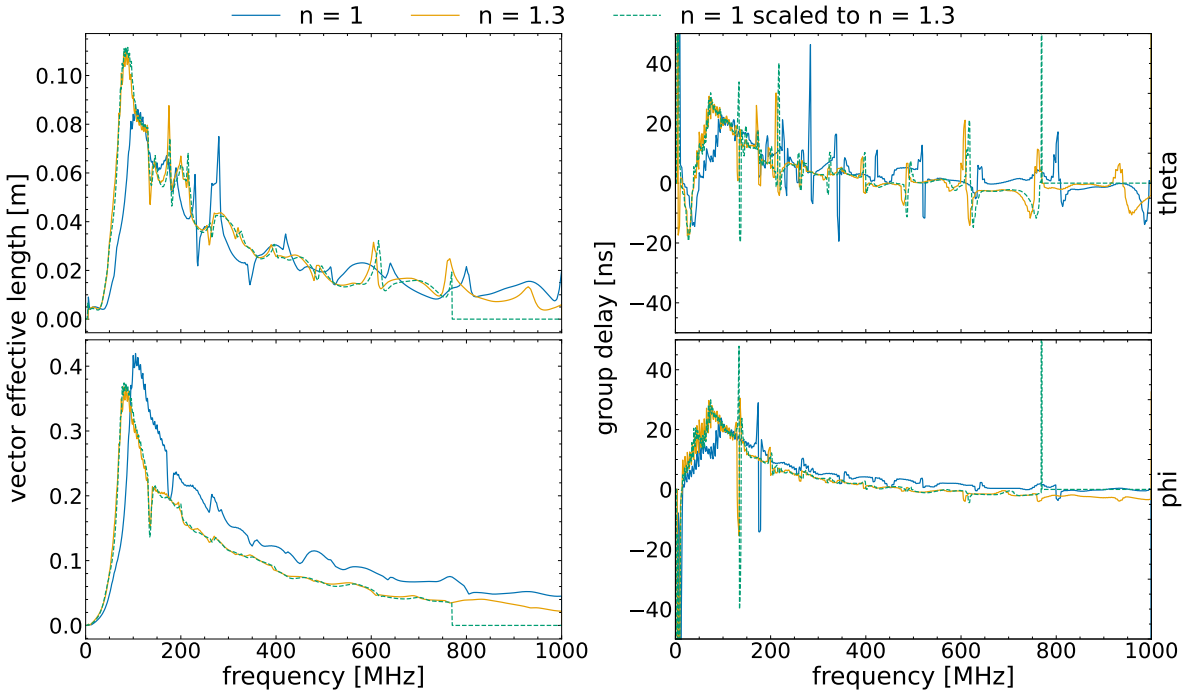


Figure 6.2: The simulated VEL (left) and group delay (right) of an LPDA for two refractive indices $n = 1$ and $n = 1.3$, as well as the VEL and group delay simulated for a refractive index of $n = 1$ scaled to $n = 1.3$. The top panels show the θ component, whereas the bottom panels show the ϕ component of the VEL and group delay.

refractive index. To validate these relations, we scale the VEL in air ($n = 1$) to a refractive index of 1.3, since we have full antenna response simulations for both refractive indices, enabling a direct comparison. The comparison of the scaled and simulated VELs and group delay is shown in Figure 6.2, which demonstrates that we can apply these relations to effectively describe VELs at a given refractive index. We should note that Figure 6.2 displays the VEL for only a single zenith and azimuth angle of a potential incoming signal. To completely transform the VEL to a different refractive index, we must apply the transformation across all possible zenith and azimuth angles.

As an additional verification, we compare templates generated using the VEL , scaled to a refractive index of 1.3, with the VEL simulated for a refractive index of 1.3, focusing on the template with the highest frequency content as a benchmark. The templates are created as we describe in Section 8.2. Correlating these templates yields a correlation score of 0.982, indicating good agreement. To determine whether a scaled VEL at a different refractive index provides a better fit to the template with the simulated VEL at 1.3, we perform a scan by generating templates for refractive indices between 1.2 and 1.4 in increments of 0.01. We achieve the highest correlation score, 0.998, for a template generated with a scaled VEL for a refractive index of 1.32. This suggests that while the VEL scaling generally works, minor discrepancies may affect the correlation score when using the modified VELs. One possible explanation is that the scaling does not perfectly reproduce all features, as visible in Figure 6.2. However, it is uncertain whether these features represent physical effects or simulation artefacts. Overall, the scaling method works, and we will apply it to examine the influence of the refractive index on the UHECR analysis.

6.4 Influence on simulations

In this section, we investigate the influence of the refractive index on the UHECR analysis. To assess the impact of the refractive index, we perform UHECR simulations, using the same framework as in the UHECR analysis. A detailed description of the simulation framework is provided in Chapter 7. For this discussion, it is important to note that when we generate a template or simulated UHECR waveform, we convert the electric field to a voltage waveform using the VEL, which depends on the refractive index. While we have to choose the refractive index for the template and simulation waveforms, the recorded waveform is affected by the actual refractive index surrounding the antenna. To evaluate how the template-data correlation score changes when we select an incorrect refractive index for the template creation, we simulate UHECR waveforms with a fixed refractive index of $n = 1.3$ and correlate them with the template waveforms generated using different refractive indices. For the templates, we select as the refractive index values the upper and lower bound of the range we estimated in Section 6.2, specifically $n = 1.28$ and $n = 1.335$. The refractive index $n = 1.3$ serves as the standard because a complete WIPL-D simulation is available for it. In both cases, we compare the results to a simulation in which $n = 1.3$ is used for both the simulated and template waveforms (standard simulation). The only difference between the standard simulation and the simulation with different refractive indices is the template waveform, while all other simulation parameters remain constant. In all simulations, we set the trigger threshold to 25 mV, corresponding to $9\sigma^2$. To enable an easier comparison, we calculate the SNR-dependent correlation score cut[3] for each simulation. For this, we divide the simulation set into multiple SNR bins and determine the 99% quantile per bin. The resulting correlation score cuts are presented in Figure 6.3. The simulation results demonstrate that using an incorrect refractive index during template creation can significantly affect the correlation score, highlighting the importance of accurately characterizing the ice surrounding the antennas at shallow depths.

The largest difference in correlation score occurs when we use $n = 1.28$ for template creation. This difference in correlation score is SNR-dependent and reaches up to 0.07. Notably, the effect is smaller for the case with the larger refractive index value, even though the increase from $n = 1.3$ to $n = 1.335$ is greater than the decrease to $n = 1.28$. Based on the refractive index differences, we would expect the largest correlation score difference for the mismatch at $n = 1.335$. Additionally, at high SNR values, the difference becomes positive, indicating that the simulation and template with the incorrect refractive index can match better than in the standard case. This behavior may be explained by our findings in Section 6.3, where templates created with the simulated VEL at $n = 1.3$ showed the highest correlation with templates created with the scaled VEL at $n = 1.32$, which is closer to $n = 1.335$ than to 1.28. Another possible explanation for the improved performance with the incorrect refractive index may result from a few simulation outliers, as the 99% quantile per bin is relatively sensitive to such anomalies. Some air shower simulations with minor artefacts may fit better to the incorrect template, distorting the SNR-dependent correlation score cut. The latter explanation is supported by the fact

[2] In this context, σ corresponds to the RMS of the simulated noise in the trigger frequency band.

[3] The SNR-dependent correlation score cut is used in the UHECR analysis to select the UHECR candidate events and is calculated from a UHECR simulation (see Section 8.5.4).

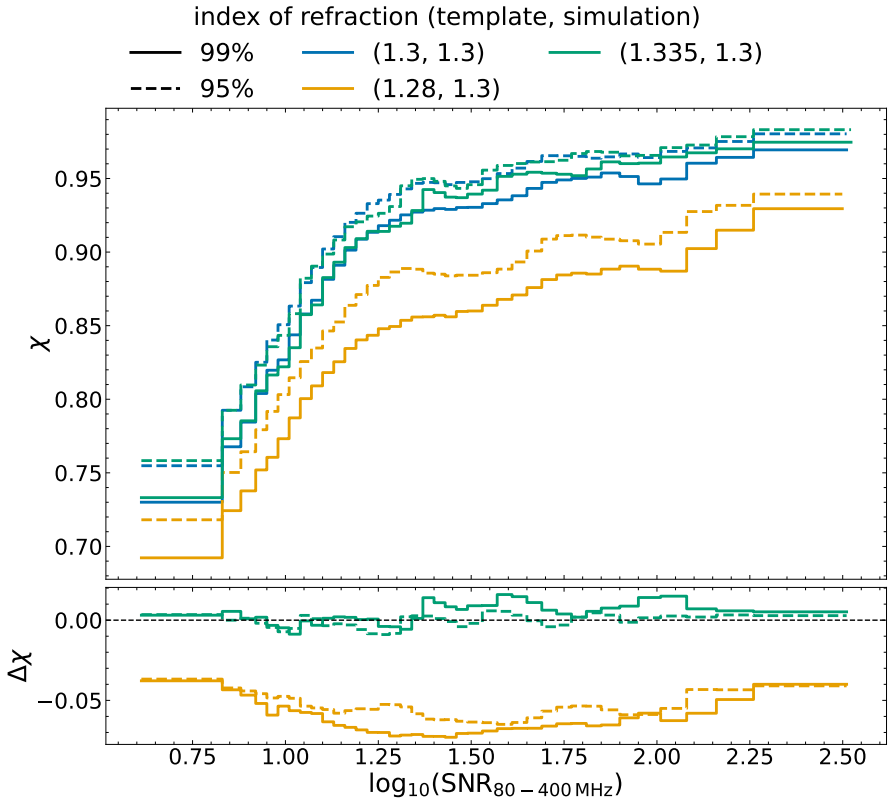


Figure 6.3: Top: The SNR-dependent correlation score cut, calculated from simulations. The colors indicate the different combinations of refractive indices used for the template and simulation. Bottom: The difference in SNR-dependent correlation score cuts relative to the standard combination (n=1.3 for both simulation and template). The colors correspond to those of the non-standard case in the top panel. The dashed black line indicates zero difference.

that the difference between the two curves is smaller for the 95% quantiles, shown as the dotted line in Figure 6.3, which are less sensitive to outliers.

In this investigation, we used refractive index values for the shallow ice around the antennas, derived from different ice models, and demonstrated that varying the index within the allowed range significantly affects the correlation score. Consequently, the refractive index must be treated as a systematic uncertainty in the UHECR analysis. We quantify this uncertainty by considering the difference in SNR-dependent correlation score cuts between the standard simulation and the simulation with $n = 1.28$ for template creation. The choice of $n = 1.28$ provides a conservative estimate of the systematic uncertainty. Furthermore, these findings highlight the need to accurately characterize the media surrounding the antennas and motivate the development or refinement of techniques to understand the ice properties better.

6.5 Ground-penetrating radar

In this section, we examine the feasibility of using a ground-penetrating radar (GPR) to measure the refractive index profile around the RNO-G stations. The GPR is a well-established technique in geophysics that, for example, is employed to investigate concrete structures and detect landmines [203]. It consists of an emitter that transmits short

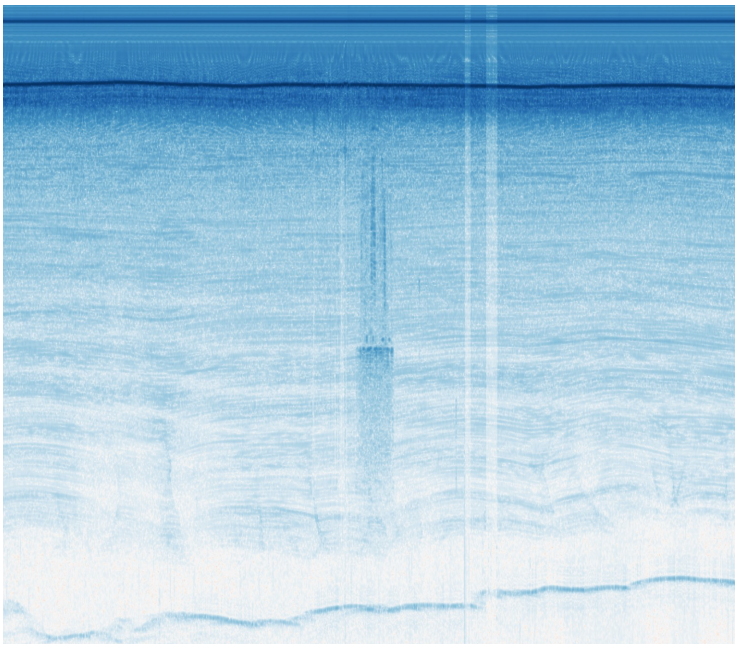


Figure 6.4: Snapshot of a ground-penetrating radar (GPR) measurement from the polarGAP survey [204], with the image provided by Ref. [205]. The central darker blue region is the IceCube detector, and the horizontal blue line at the bottom is the bedrock. The y-axis corresponds to the signal travel time, and the x-axis is either the traveled distance or the number of measurements.

electromagnetic signals into the ground, and a receiver that records any returning signal. Reflections of the transmitted signal can occur at interfaces where there is a sudden change in electromagnetic properties, such as between different media. By analyzing the time difference between signal emission and reception, the GPR can map reflection points in the ground below it. Typically, the emitter and receiver, often integrated into a single device, are moved across the ground during the measurement. The GPR antenna is typically characterized by the frequency at which the bandwidth is centered, its so-called center frequency. A comprehensive introduction to GPRs is provided in Ref. [203].

A GPR does not require direct contact with the ground and can be deployed on platforms such as aircraft. For instance, the polarGAP survey, an airborne campaign over the South Pole, utilized a GPR [204]. Figure 6.4 presents a snapshot from this survey, where the darker blue region in the center corresponds to the IceCube [83] detector, which reflects the radio signal. Therefore, this figure basically shows a picture of the deployed IceCube detector. Additionally, the figure shows the bedrock, visible as a horizontal dark line at the bottom.

The objective of the GPR measurement at RNO-G is to determine the refractive index profile. Because the refractive index is frequency dependent, the center frequency of the GPR device must align with the RNO-G frequency bandwidth. To determine the refractive index with a GPR, we can perform a dynamic measurement by moving the device across the ground during data collection. For simplicity, we assume that signals travel along straight lines to and from a potential reflector in the ice, although the refractive index varies with depth. Accounting for the actual signal paths would require more complex calculations or ray tracing, which may preclude an analytical solution depending on the ice model. Therefore, we will restrict this proof-of-principle analysis to the simplified case. Figure 6.5 shows a schematic illustrating the dynamic measurement

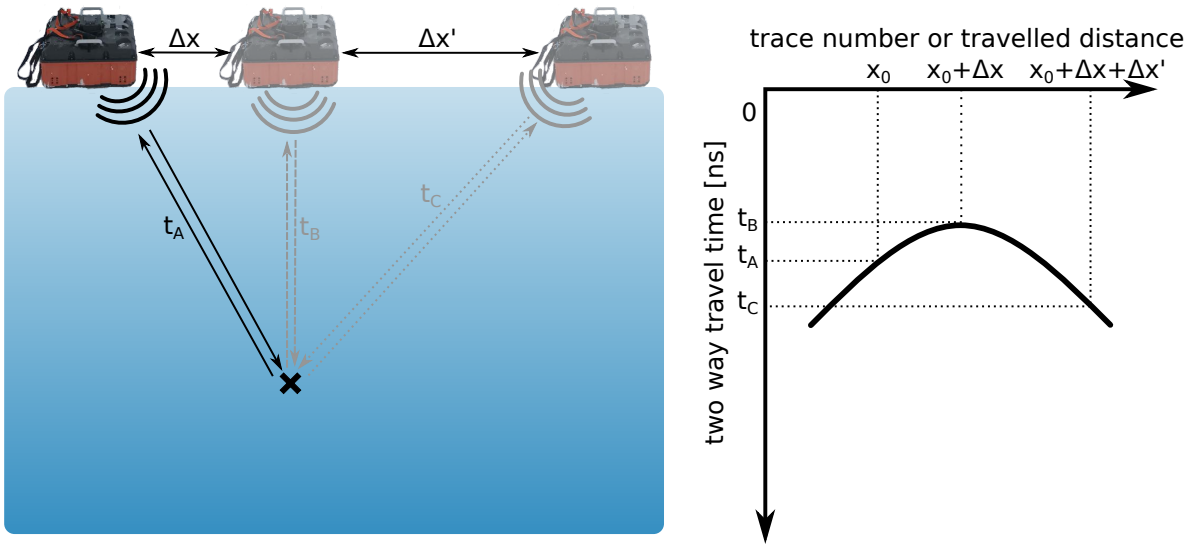


Figure 6.5: Left: Illustration of a dynamic GPR measurement, with a reflector indicated by an x. Right: The expected spectrogram corresponding to such a measurement.

principle. To further simplify the derivation of a formula for determining the refractive index from dynamic measurements, we assume the refractive index profile to be constant. The time t_A required for a signal to travel from an arbitrary GPR position to the reflector and back depends on the distance d_A between the emitter and the reflector and on the constant refractive index n. The relationship can be expressed as follows.

$$t_A = 2\frac{n}{c}d_A, \tag{6.8}$$

Here, c denotes the speed of light in vacuum, and the factor two accounts for the signal traveling from the emitter to the receiver and back. As the GPR moves across the reflector during the dynamic measurement, it will be positioned directly above the reflector after traveling a distance Δx. At this point, the distance to the reflector equals its depth, and the two-way travel time is denoted as t_B. Assuming the GPR moves only horizontally, we can write the traveled distance Δx as a function of the two-way travel times of the signal before and after the GPR moved using Equation 6.8:

$$\Delta x^2 = d_A^2 - d_B^2 = \frac{c^2}{4n^2}(t_A^2 - t_B^2) \tag{6.9}$$

We can solve this formula for the two-way travel time before the GPR was moved (t_A), which leads to the following formula:

$$t_A = \sqrt{4\frac{n^2\Delta x^2}{c^2} + t_B^2} \tag{6.10}$$

This formula describes a hyperbola. When the refractive index varies along the z-axis (i.e., with depth), the analysis becomes more complex. Calculating the signal travel time in this scenario requires integrating the refractive index along the traveled signal path. In this study, we replace the refractive index in Equation 6.10 with an average value $\overline{n}$ experienced by the signal along its path. Further research is necessary to assess the validity of this simplification. Assuming the signal path remains a straight line, we can



Figure 6.6: Image of Christoph Welling and me conducting a dynamic GPR measurement at one of the RNO-G stations in Greenland. The orange-and-black device is the GPR, with a custom GPS antenna attached. The picture is provided by Ref. [207].

calculate the two-way travel time t_A from an arbitrary point A to a reflector and back as follows:

$$t_A = \sqrt{4\frac{\overline{n}^2 \Delta x^2}{c^2} + t_B^2} \tag{6.11}$$

To measure the refractive index with a GPR near the RNO-G detector, we planned to use the antennas installed in the boreholes as radar reflectors for the measurement. Each antenna would then produce a hyperbolic signature, which we could fit using Equation 6.11 to extract the average refractive index. Since the antennas are deployed at different depths, we could determine the average refractive index at different depths. Additionally, we can calculate the same average for different ice models as a function of depth and compare them to the measured values.

In the subsequent sections, we describe the available GPR dataset (see Section 6.5.1) and present a stationary GPR measurement as a test case (see Section 6.5.2). Following this, we analyze a dynamic GPR measurement to extract information about the refractive index profile and to compare the results with different ice models (see Section 6.5.3). This analysis serves as a proof of principle, and additional work is required to address potential uncertainties and systematic effects. In Section 6.5.4, we outline the necessary next steps and discuss the potential advantages of this method.

6.5.1 Dataset

We collected the GPR data for this study during the 2023 maintenance and calibration campaign at Summit Station. For the data acquisition, we used a commercially available GPR with an antenna center frequency of 200 MHz [206], as depicted in Figure 6.6. Although the GPR is equipped with a built-in GPS antenna, we mounted an additional

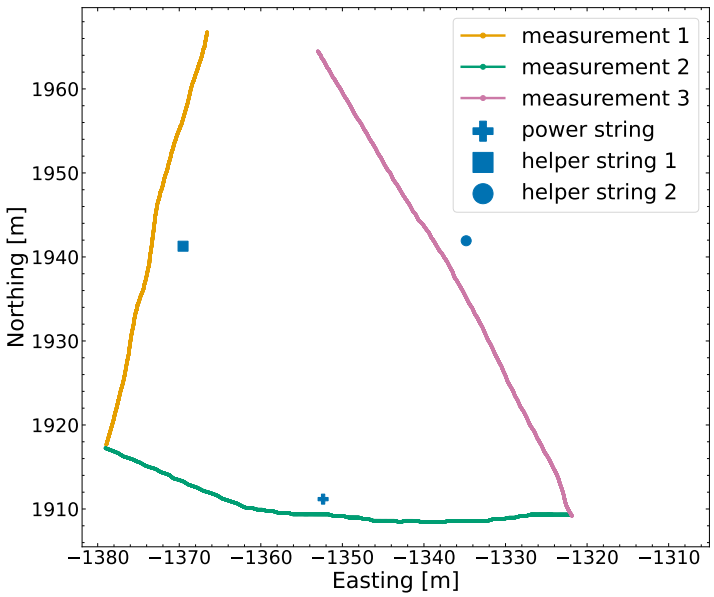


Figure 6.7: A GPS track from a GPR survey conducted at station 12.

custom GPS antenna on a pole on the GPR to enhance the GPS performance.

The general idea of the measurement was to utilize the antennas inside the boreholes as reflectors and use the reflected signals to determine the refractive index at the antenna depths. In addition to several smaller test measurements, which we do not discuss in this thesis, we conducted two main types of measurements: stationary and dynamic. First, we performed a stationary measurement adjacent to an approximately 10 m borehole to evaluate the operation of the GPR, determine whether we can detect a signal reflected from a metal object inside the borehole, and evaluate different GPR settings. Settings we have tested include, for example, the recording time for possible return signals, which alters the maximum depth at which a reflector can be detected, and the number of recorded measurements that are averaged to increase the signal-to-noise ratio. However, the raw data we collected in the field provided limited visibility into a possible signal, making it difficult to assess measurement success. A raw spectrogram from the stationary measurement is presented in Figure 6.9, where only a slight indication of a signal is visible. Consequently, offline data processing is required to extract meaningful information, as we describe in Section 6.5.2.

Despite the challenges with the stationary measurement, we conducted dynamic measurements. As we described previously, dynamic measurements are essential for determining the refractive index. We surveyed all seven deployed stations with the GPR by dragging the device along a triangular path around each station, ensuring that we traversed each antenna string in approximately a straight line. GPS data for a representative survey of a single station is shown in Figure 6.7, which illustrates the survey track.

For each measurement, the GPR records multiple traces, with each trace beginning when a GPR signal is transmitted and continuing for a specified duration to capture all returning signals. Additionally, it can be specified in the settings to average multiple of these traces to enhance the signal-to-noise ratio. Typically, the data is represented as a spectrogram, where the x-axis denotes the trace number, the color indicates the received power, and the y-axis represents the time at which the power is recorded, relative to the

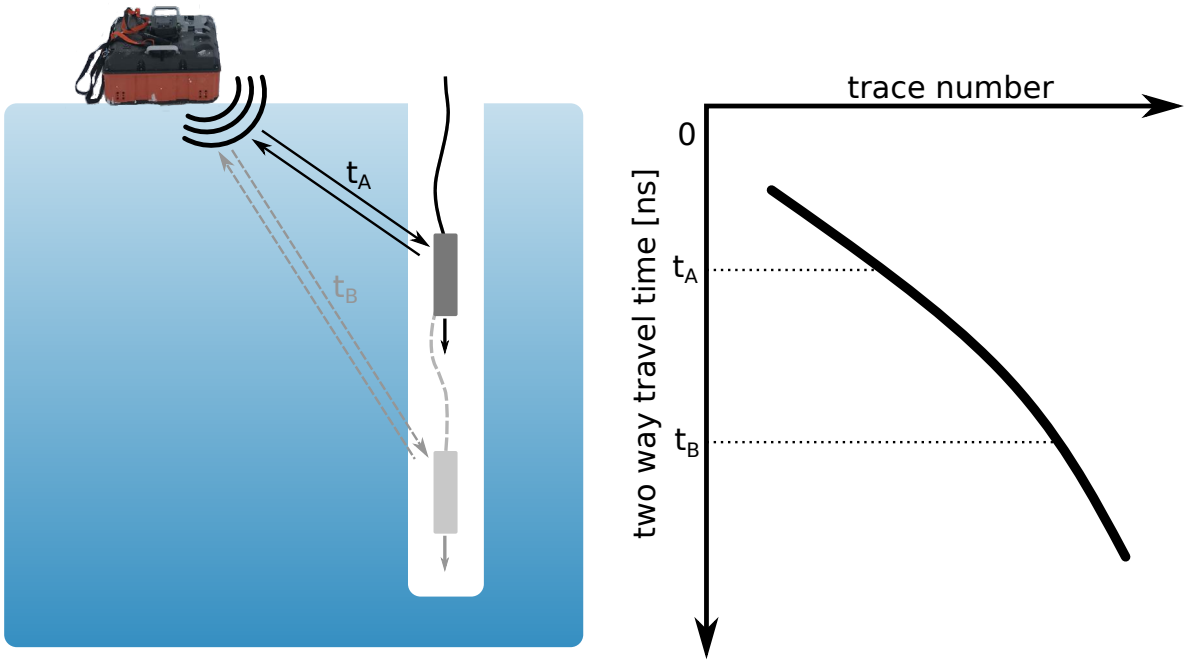


Figure 6.8: Left: Schematic of a stationary GPR measurement conducted adjacent to a borehole, with a metal piece lowered into it. Right: A sketch of the expected spectrogram for such a measurement configuration.

signal transmission time. This time is referred to as the two-way travel time. For dynamic measurements, trace numbers can be converted to a distance using the corresponding GPS data. For the offline reading and processing of the GPR data, we use the Python package ImpDAR [208].

6.5.2 Stationary measurement

As described in the previous section, we conducted a stationary GPR measurement to evaluate whether we can detect reflections from an antenna inside the borehole. In this measurement, we positioned the GPR adjacent to an approximately 10 m deep, accessible, and non-instrumented borehole. To mimic an antenna, we then lowered a metal piece into the borehole at an approximately constant speed while the GPR collected data. We moved the metal piece during the measurement to reduce the likelihood of confusing it with a stationary reflector, such as ice layers. For a stationary measurement with a stationary reflector, the return signal should appear at a consistent time, resulting in a horizontal line in the spectrogram. However, if the reflector is lowered at a constant rate, the travel time increases, producing a linear function with a non-zero slope in the spectrogram. Figure 6.8 presents a schematic of the measurement setup and a sketch of the expected spectrogram. Deviations from this linear behavior may arise due to a non-constant index of refraction and its influence on the ray path. In the presented measurement, the primary source of deviations is likely the non-uniform lowering speed, as we lowered the metal piece manually. We performed several of these measurements with varying settings to optimize them. Here, we only present a single measurement as an example to demonstrate the offline data processing.

Figure 6.9 displays the raw data from the stationary example measurement. The raw data reveals only a faint indication of a signal from a moving reflector, underscoring the need

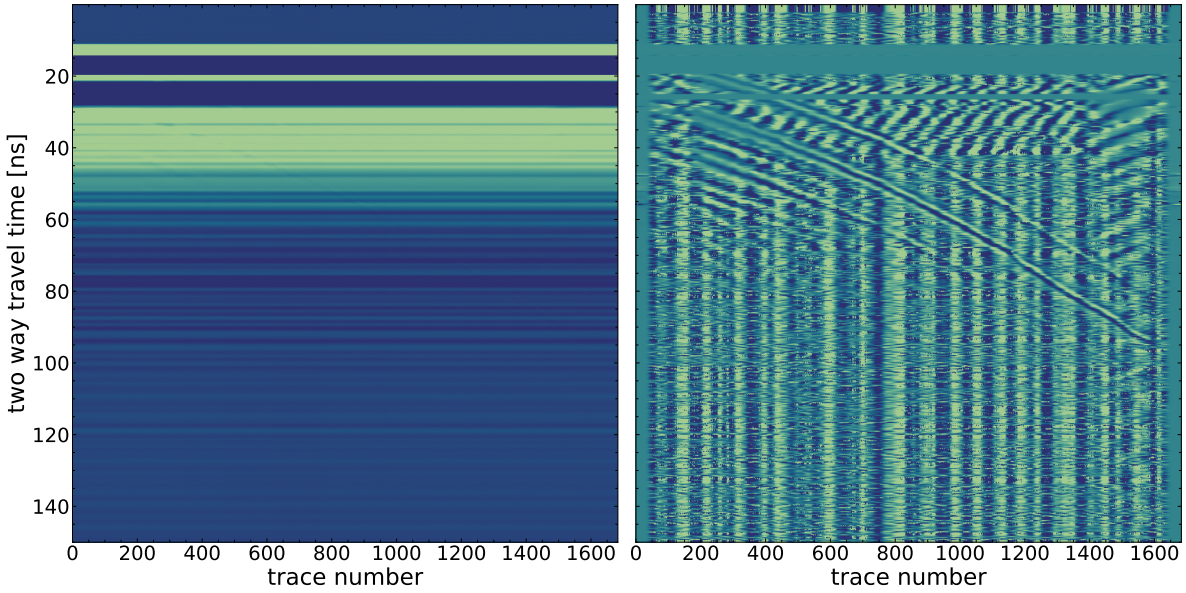


Figure 6.9: Left: Raw GPR data acquired during a stationary measurement near a borehole, with a metal piece lowered into it. Right: Processed version of the data shown in the left panel. In both panels, the color bar is restricted to the two-sided 80% quantile.

for additional offline processing. We evaluated various filters and processing methods, achieving optimal results with a combination of Wiener filters [209]. These filters are implemented in the ImpDAR framework, which enables specifying the number of pixels in the x and y directions that are included in the filter. For the GPR data processing, we first apply a Wiener filter with size (1, 100), then apply three filters with size (1,15). Although alternative filter combinations or sizes may yield improved results, we deemed the mentioned combination sufficient for this proof-of-principle analysis.

The denoised spectrogram shown in Figure 6.9 reveals three approximately linear features. The two-way travel time difference between these lines is approximately 12 to 14 ns, corresponding to a distance of about 1.2 to 1.6 m. Therefore, these lines cannot originate from different sections of the metal piece, as the maximum length of the piece is approximately 1 m. The primary signal, located in the middle, exhibits a maximum two-way travel time of 95 ns, corresponding to depths of 9 to 11 m, depending on the assumed refractive index. This maximum depth aligns with the total depth of the borehole used in the measurements. Additionally, we can estimate the approximate lowering speed of the metal piece from the spectrogram using the estimated initial and final depths, derived from the two-way travel time, and the total measurement duration. This estimate provides only a ballpark number, as it is influenced by factors such as the refractive index we choose when converting the two-way travel times to depths. For the measurement presented, we obtain a value of around 20 cm/s, which is consistent with expected manual lowering speeds and supports the conclusion that we detected a signal reflected by the moving metal piece. However, the origin of the multiple lines remains uncertain. A hypothesis is that they result from an interference pattern, as the spacing between two of the lines is roughly equal to the radar wavelength (1.5 m).

Notably, the detection of a signal in the first place is interesting given the close proximity of the GPR to the borehole ($\mathcal{O}(0.5)$ m), which results in a large incidence angle at the ice-air boundary and would typically lead to total internal reflection. Further investigations are

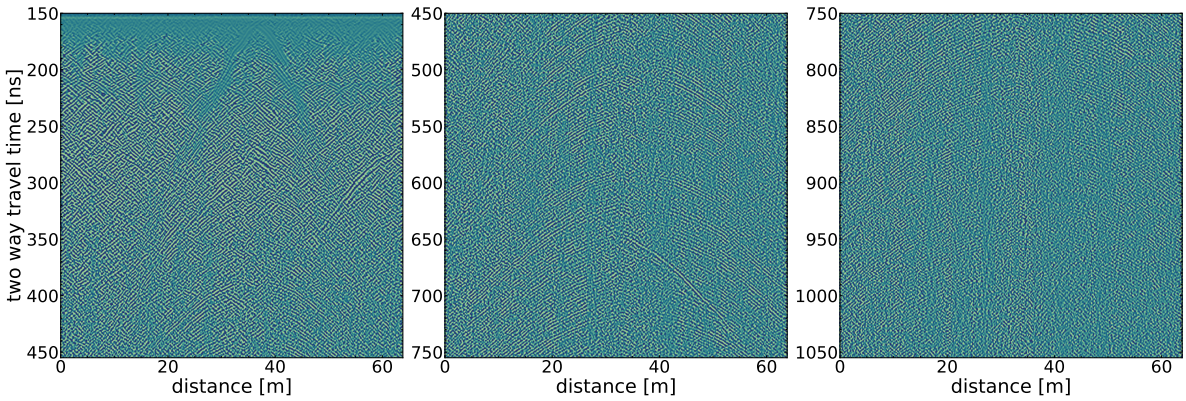


Figure 6.10: Processed spectrogram from a dynamic GPR measurement conducted at the helper string 2 of station 12.

required to understand the behavior at the borehole ice-air boundary and the reflection pattern we have observed. Nevertheless, the results demonstrate that appropriate data processing can reveal faint signals in the collected GPR data.

6.5.3 Dynamic measurement

In this section, we analyze a dynamic GPR measurement to determine whether it is feasible to extract the refractive index from such data. As an example, we consider the measurement we conducted at the helper string 2 of station 12. The GPS track for this measurement is presented in Figure 6.7. As we described previously, data processing is necessary to identify any relevant features in the GPR data. Therefore, we apply the same order and sizes of Wiener filters as for the stationary measurement. Additionally, we apply a $50 - 400\,\mathrm{MHz}$ filter before the Wiener filters to remove the vertical stripes observed in the stationary measurement (see Figure 6.9). The processed data for the dynamical measurement is shown in Figure 6.10, which displays data up to a two-way travel time of 1100 ns, roughly corresponding to the depth of the borehole, beyond which no further reflections are expected.

Since we conducted the measurement at the helper string 2, which contains only antennas at approximately 95 m depth, we expected a strong signal at a two-way travel time of about 1000 ns. However, we do not observe such a signal at these times. Instead, we detect faint hyperbolas across the entire range of two-way travel times, which contradicts the assumption that we should observe the signal reflected from the antennas. One hypothesis is that a small portion of the GPR signal reflects off the borehole ice-air boundary back to the GPR. At the same time, the antennas themselves may be shielded by this air-ice boundary, making them undetectable for the GPR signal.

To further investigate the origin of these signals, we calculate the distance from the GPR to the borehole and display it alongside a spectrogram in Figure 6.11. The crests of the hyperbolas align well with the point of the closest approach to the borehole, supporting the hypothesis that the reflections originate from the borehole. However, the minimum distance to the borehole is only 3.3 m, which may be attributed to the use of the internal GPS of the GPR, which is potentially less accurate than the differential GPS attached to the GPR. Nevertheless, a positional shift of this magnitude does not alter the interpretation that the reflections are from the borehole.

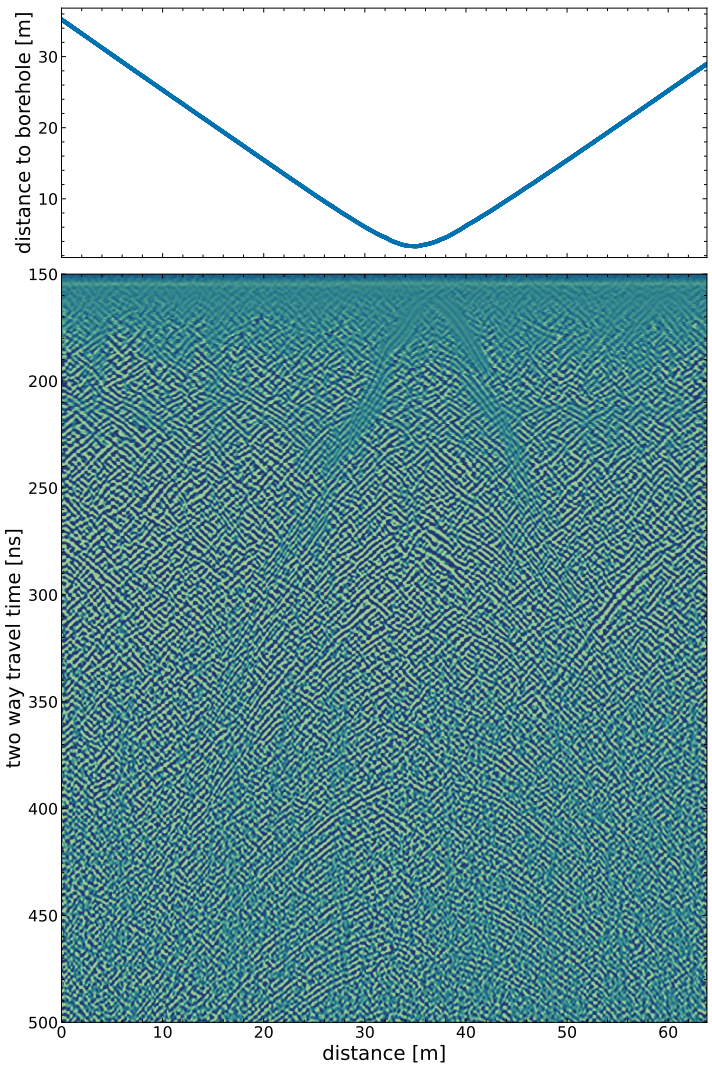


Figure 6.11: A small section of the data presented in Figure 6.10, along with the distance of the GPR from the borehole.

The reflected signal is weak, suggesting that most of the GPR signal is reflected downward, with only minor secondary or higher-order effects returning a small portion to the surface. Alternatively, some reflections may result from subsurface layers that redirect the signal upward. If such a layer intersects the borehole, the GPR signal could be reflected downwards into that layer and subsequently back toward the GPR. Furthermore, the signal vanishes near the crest of the hyperbola, which appears to grow stronger with increasing depth. In summary, the analysis indicates that the reflections originate from the borehole, although further investigations are required to fully understand the reflection mechanism of the signal.

In the subsequent analysis, we extract the average index of refraction from the hyperbolas identified in the dynamic measurements. We manually trace selected hyperbolas at different depths. While this approach is not suitable for large-scale refractive index measurements, it is adequate for an initial proof-of-principle study. To account for potential inaccuracies, we estimate a realistic but conservative error of 3 ns and assign it to all hand-selected two-way travel time data points. Figure 6.12 shows a section of the full spectrogram with the manually selected data for several hyperbolas. The right side of the figure displays these data points fitted to the theoretical model described by

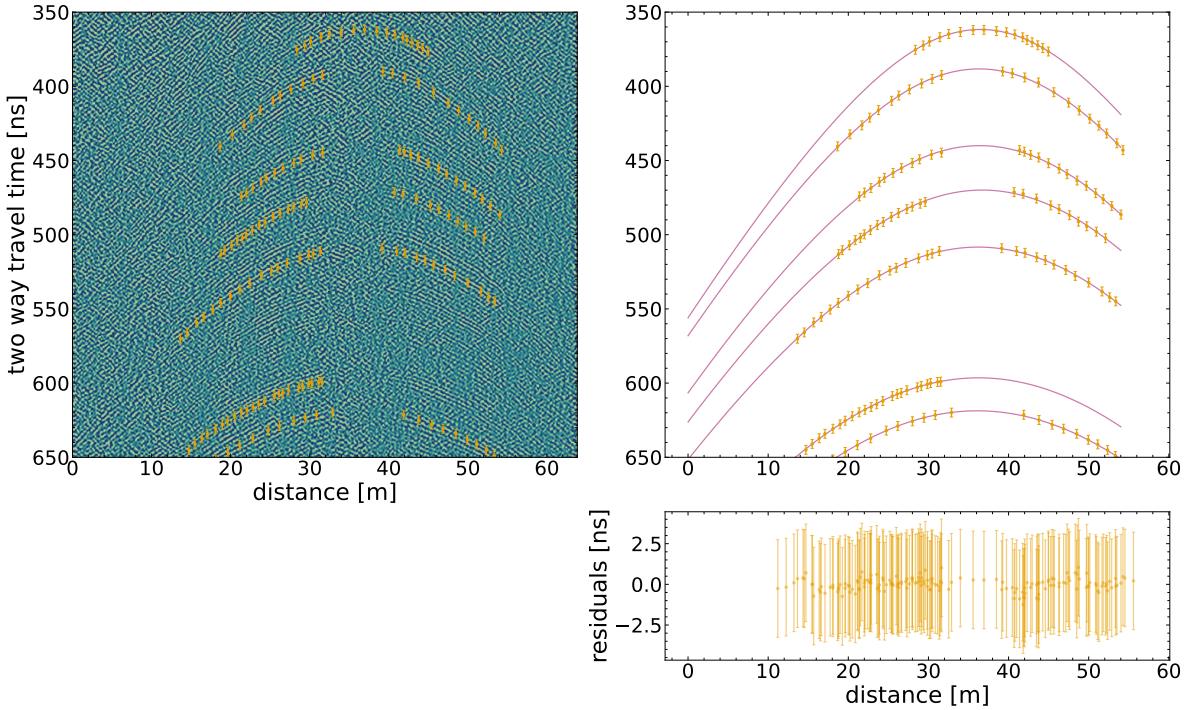


Figure 6.12: Left: A section from the spectrogram shown in Figure 6.10, including data points for the selected hyperbolas. Top right: The hyperbola data from the left panel fitted using the function described in Equation 6.11. Bottom right: The residuals between the hyperbolas and the fit shown in the top right panel.

Equation 6.11. It is important to note that this equation assumes straight-line propagation and that further studies are required to assess the impact of incorporating ray tracing. The lower panel shows the residuals between the fit and the hyperbola data, which are in close agreement. All residuals are consistent with zero within the uncertainties, even if the uncertainties for the two-way travel time would be less conservative.

One of the fit parameters is the average refractive index that the signal encounters on its way to the reflector and back to the GPR, which we can compare with existing ice models. For the comparison, we estimate the depth at which the average refractive index is measured, i.e. the reflector depth. This depth is calculated from the two-way travel time when the GPR is directly above the reflector, which is another fit parameter (t_B) and the measured average refractive index. We calculate the depth uncertainty from the fit uncertainty using standard error propagation. In Figure 6.13, we show the measured average refractive index values, which we derived from data at station 12 helper string 2, along with the theoretical average refractive index profiles from the different ice models introduced in Section 6.1. The data and model values are comparable, suggesting that this method could complement existing techniques for measuring the refractive index. However, we need to be cautious with interpreting the calculated uncertainties. This result represents a proof-of-principle analysis, and a more comprehensive evaluation of the uncertainties and systematic effects is necessary to assess the agreement between the data and ice models.

For the presented study, we assume that the GPR signal propagates along straight lines, although this is not accurate due to the varying refractive index. Consequently, further studies are required to quantify the systematic effect we introduce by this assumption. If the omission of ray tracing introduces significant discrepancies, it may not be possible to

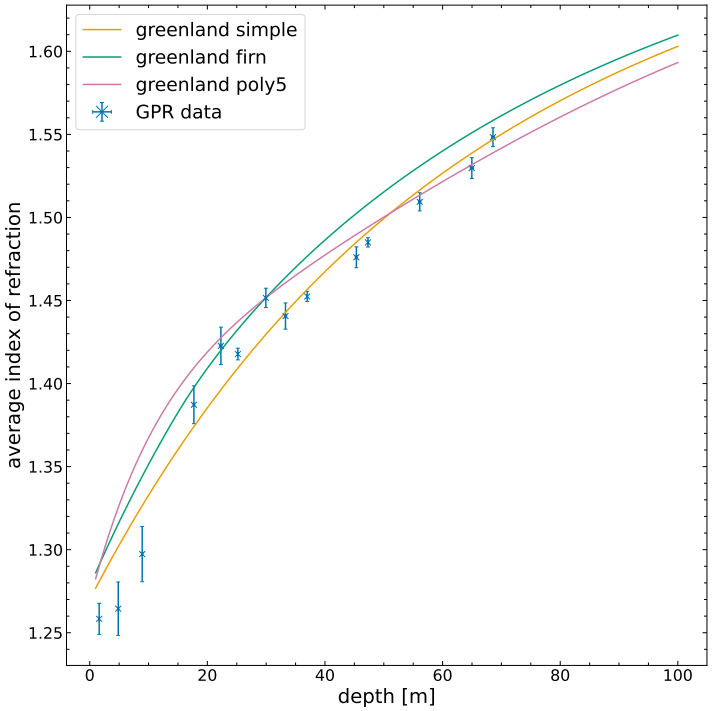


Figure 6.13: The average refractive index as a function of depth, with the data points estimated from a dynamic GPR measurement and lines representing calculations from three different ice models.

derive a general formula for the hyperbolas without assuming a straight line propagation. In this case, we will not be able to fit the general equation to the recorded data to extract a refractive index profile. Nevertheless, it would still be possible to fit ice models to the GPR data by directly comparing simulated hyperbolas, which are influenced by the ice model, with the measured hyperbolas identified in the GPR data.

6.5.4 Advantages and next steps

In the previous sections, we demonstrated the potential of using GPR measurements to determine the refractive index at various depths and compare it with different ice models. The analysis serves as a proof of principle, and further work is required to obtain accurate refractive index measurements. However, multiple factors justify continuing to develop this measurement technique. First, this approach offers a complementary and independent method for determining the refractive index, to methods such as the analysis of ice density data, TOF measurements, or antenna S11 parameter measurements. Second, GPR measurements are easy to conduct and require minimal on-site effort, as the procedure only involves dragging the GPR along at least one of the strings. This simplicity enables repeated measurements at different stations and facilitates the investigation of potential local variations of the refractive index profile. If significant local differences are identified, GPR measurements could be incorporated into the station deployment procedure. While TOF or antenna S11 measurements can also be repeated at each station, these techniques are limited to periods before the antennas are deployed, as an empty borehole is necessary for the measurement. In contrast, GPR measurements can be conducted after the station is fully deployed, allowing for repeated measurements

across, e.g, multiple years. Repeated measurements would enable the study of temporal changes in the refractive index of the upper layers as new snow accumulates.

To establish this technique as a reliable measurement of the refractive index, several key steps are required. First, it is important to assess whether a systematic uncertainty arises from assuming straight-line propagation of the signal in ice. If the effect is substantial, it might only be possible to fit ice models to the GPR data by directly comparing the measured and simulated hyperbolas. Second, automating the selection of the hyperbolas is necessary, as manual identification is impractical for large datasets. The Hough transform, an algorithm for detecting curves in images [210], presents a promising solution. Furthermore, in the current analysis, we simplify the hyperbola calculation by assuming an average refractive index rather than integrating it along the actual signal path. Future work should investigate the impact of this simplification. We also assume in the analysis that the GPR was dragged along a perfectly straight line over the borehole. Simulations are required to assess how deviations from this ideal path affect the measurement. During the 2023 measurement campaign, we did not explicitly monitor the path for straightness or alignment with the borehole. If significant effects are identified, future campaigns can try to improve in this regard, e.g., by marking the intended path in advance. Although a differential GPS was attached to the GPR, we relied solely on the internal GPS in the analysis. For improved accuracy, future analyses should utilize the more precise differential GPS. Furthermore, a better understanding is needed of where the GPR signal is scattering from, to ensure that the physical interpretation of the signal remains valid. Overall, a proper treatment of the uncertainties and systematic effects associated with these measurements must be done.

6.6 Conclusions

In this chapter, we introduced different ice models and examined their impact on the UHECR analysis presented in this thesis. Additionally, we introduced the GPR as a potential new technique for measuring the refractive index.

First, we demonstrated that a significant uncertainty remains in the refractive index at shallow depth, where the LPDAs are located. This uncertainty arises from differences among ice models in how they describe the refractive index profile. Ongoing developments include new measurement techniques for characterizing the ice, such as antenna S11 parameter measurements in the borehole and TOF measurements. Additionally, a new three-stage ice model has been developed and is now primarily used in the context of RNO-G. Due to time constraints, we were not able to incorporate the latest advancements into this work. However, future work should revisit the permissible range of refractive indices at shallow depths to incorporate recent developments, which may help reduce the uncertainty in the refractive index.

Furthermore, we introduced a scaling relation to convert a VEL simulated at a refractive index of n_1 into a VEL at a refractive index of n_2. This scaling relation enabled us to investigate how the permissible range of refractive indices affects the UHECR analysis. The simulations we performed indicate that this refractive index has a substantial impact, potentially resulting in differences of up to 0.07 in correlation score. These

findings underscore that a thorough understanding of the ice is essential not only for understanding the signal path but also for analyses sensitive to the waveform shape.

Lastly, we introduced the GPR as a tool for measuring the refractive index. With this proof-of-principle analysis, we demonstrated that it is feasible to extract the refractive index from GPR data, providing a direct measurement without relying on ice density data. The extracted average refractive index values are comparable to those estimated from existing ice models. However, further research is required to develop a comprehensive analysis that rigorously addresses uncertainties and potential systematic effects. For instance, automating the detection of signals (hyperbolas) in the GPR data and quantifying the influence of the straight-line approximation compared to full ray tracing are necessary steps. We argue that it is advantageous to continue developing this measurement technique, as it offers a complementary method with distinct advantages. First, GPR measurements are easy to conduct and require minimal on-site effort, which enables repeating the measurement at multiple sites to study local variations in the refractive index profile. Second, these measurements can be performed after antenna deployment in the borehole, unlike S11 and TOF measurements, and therefore can be repeated, e.g., annually. This capability may enable investigations of changes in the refractive index of the upper firn layers as new snow accumulates. Generally, the analysis suggests that the GPR offers a valuable complementary method for measuring the refractive index.

In conclusion, in this chapter, we demonstrated that the uncertainty of the shallow refractive index significantly influences the UHECR analysis. Additionally, we introduced an alternative measurement technique to study the ice properties, which may help to better constrain the shallow refractive index in future research.

7 Cosmic Ray Simulations

In this chapter, we present the setup used to produce realistic UHECR simulations for RNO-G. First, we outline the general simulation setup (see Section 7.1). Next, we describe the antenna, amplification, and digitization responses (see Section 7.2) and the Radiant trigger (see Section 7.3). We then discuss noise generation (see Section 7.4) and explain how we calculate the effective area of the detector and the expected UHECR rate (see Section 7.5). In the analysis, we use the simulations to calculate an SNR-dependent correlation score cut to select the UHECR candidate events. The calculation of this cut is covered in Chapter 8.

7.1 Simulation setup

For the simulation, we use a library of 384 air showers generated with CORSIKA (version 7.6400) [67] and CoREAS (version V1.3) [122] using only protons as the primary particle[1]. Most of the other steps in the simulation, such as reading the generated air showers, utilize functions provided by NuRadioReco [211]. The air showers are generated for the conditions at Summit Station and span the zenith angles $5 - 75°$ and energies $10^{16} - 10^{18.5}$ eV. Large zenith angles, above 75°, are excluded, since the signals have to refract into the ice to be detected, which suppresses the sensitivity of RNO-G to near horizontal air showers at the surface. All air showers are simulated for the same azimuth angle of 90°. However, due to the nearly vertical Earth magnetic field at Summit Station (inclination 80.94° downwards) [212], we can mimic different azimuth angles by rotating the detector without introducing a substantial bias [121]. Figure 7.1 shows the parameter distribution for the air-shower library. Since the parameter distributions do not match those expected from the UHECR flux or from geometric arguments, we apply appropriate weighting before calculating the effective area from the simulation (see Section 7.5). To cover all azimuth angles, we rotate the detector in 10° steps and, for each step, we sample each air-shower footprint at 240 points distributed on a star-shaped pattern. Figure 7.1 shows the star-shaped pattern of an example shower with an energy of 10^{18} eV and zenith angle of 35°. Due to the number of air showers and the detector rotation, we sample 3,317,760 realistic events from the generated showers per simulation. For each event that fulfills the assumed trigger condition (see Section 7.3), we save high-level values relevant to the analysis, such as the SNR, rather than the full waveforms, to reduce the size of the stored data. Before we calculate the high-level quantities, we apply a rectangular bandpass filter with a passband of $80 - 400$ MHz to each event. We apply the filter to follow the same steps as in the data processing of the analysis (see Section 8.5).

Although we generate a large number of simulated waveforms, they all stem from a small set of simulated air showers, which can lead to statistical issues. To calculate the effective area (see Section 7.5), we group the simulated air-shower events in energy and zenith-angle bins. For a small number of energy-zenith angle combinations, either none or one air shower exists. A single air shower might show over- or under-fluctuations.

[1] Protons are used as the primary particle as they cover the widest range of $X_{\max}$ values (see Figure 2.4).

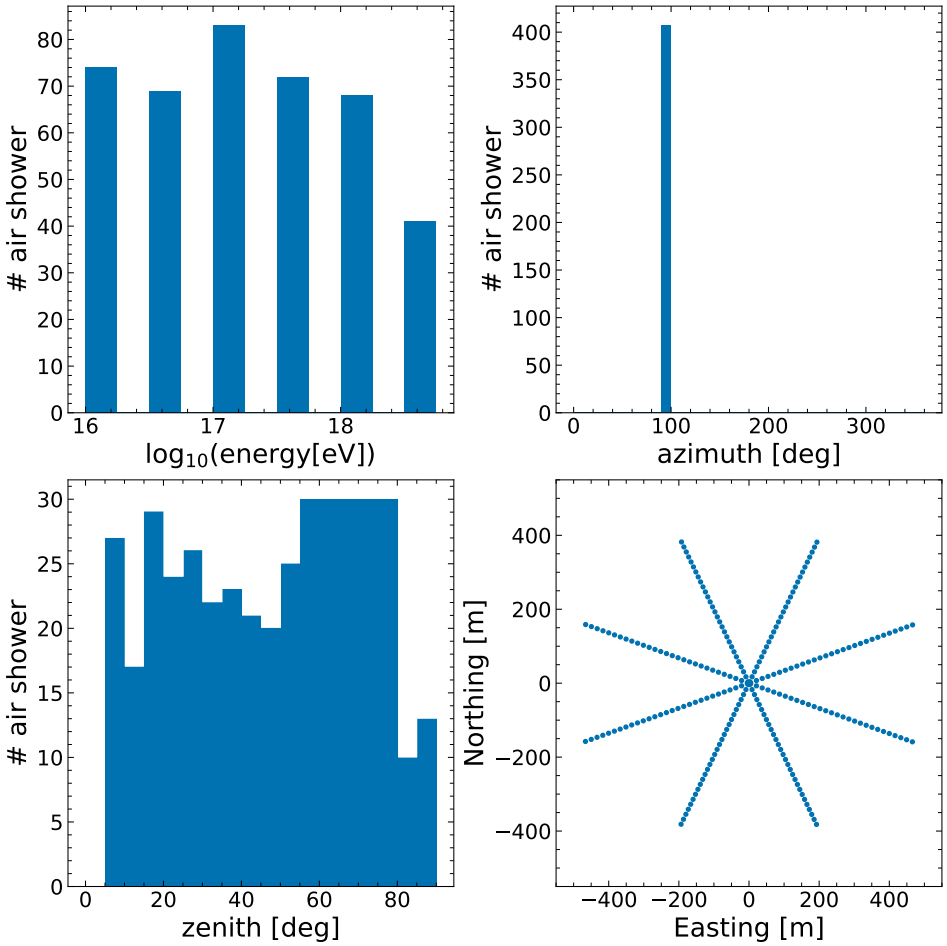


Figure 7.1: The energy, zenith, and azimuth distribution of the air-shower library. Also shown is the footprint of an example air shower, projected onto the ground, with 240 sample positions. The example air shower has an energy of 10^{18} eV and a zenith angle of 35°. The projection onto the ground causes the star-shaped pattern to appear asymmetrical. In a plane perpendicular to the air shower axis, the pattern would be symmetric.

When only one air shower represents a bin, resampling the same air shower may amplify these fluctuations in the final simulation set. The same applies to simulated air showers with simulation artifacts. Thus, increasing the number of air showers in the library is beneficial. A new set of air showers was generated by Philipp Laub, filling the gaps of the air-shower library and increasing the number of air showers per bin. In preliminary tests, we found both shower libraries to be compatible and usable together. Due to time limitations, only the old air-shower library was used for this work.

In these simulations, the detector is modeled as a single station with only shallow antennas. The deep antennas are not simulated because, when a shallow antenna triggers the station, the signal is only visible in the shallow antennas due to a time alignment issue between the deep and shallow readout windows (see Section 4.2.2). Therefore, the deep antennas are not of interest for the analysis and cannot contain additional information. Each station is simulated individually, and we do not account for coincident events across multiple stations. However, coincidence events are rare (about 3%) [172] and occur only at high zenith angles, where the footprint on the ground is large enough to span the 1.25 km distance between stations. Consequently, including coincidence events is expected to only be a small correction.

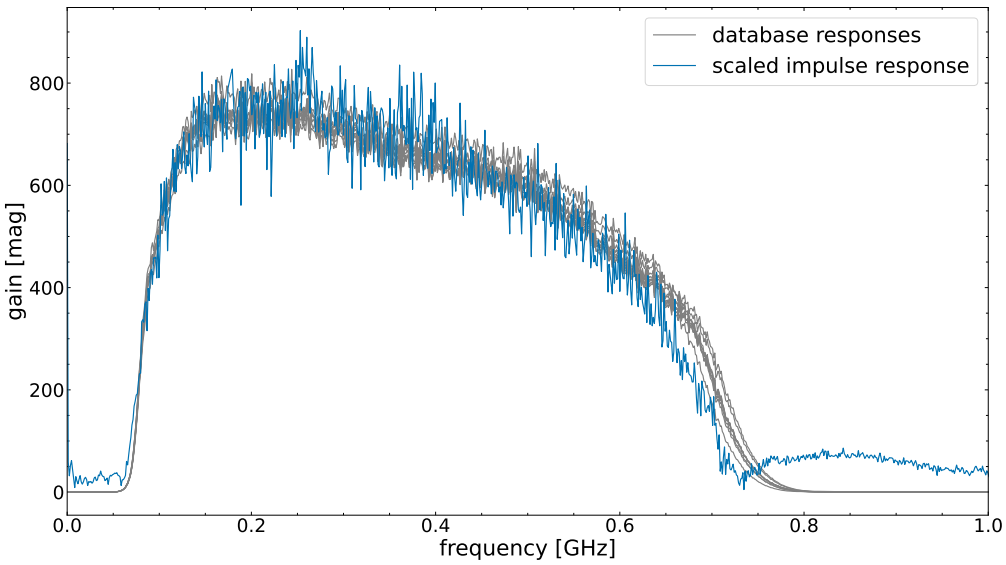


Figure 7.2: Comparison of the amplification and digitization chain gain described by the database response and the scaled impulse response. The impulse response shown here exhibits greater noise than in Section 5.1.2, despite being the same measurement, as we did not apply the same denoising steps for this figure.

7.2 Antenna and hardware description

From the generated air showers and the NuRadioReco framework, we receive an electric field at the antennas of the simulated detector. To convert this electric field into a voltage waveform, we use a vector effective length (VEL) for the LPDAs, simulated with the software WIPL-D [201] and provided via NuRadioReco. In the VEL simulation, the antenna is modeled in an infinite medium with a refractive index of $n = 1.3$, which approximates the expected value for the surrounding ice (see Section 6.2). The simulation yields a complex VEL for both polarizations [17].

For the description of the amplification and digitization response, we use impulse response measurements (for more information, see Section 5.1.2), as these currently provide the most accurate characterization. The impulse response measurements are all normalized, and at the time of the analysis, the absolute gain calibration was a work in progress. Consequently, as a stopgap solution, we scaled the response so that the gain is similar to that of the database response (for more information about the database response, see Section 5.1.2). This scaling ensures that we can trigger properly on the simulated waveforms. To perform this scaling, we calculated a scaling factor SF for all upward-facing shallow channels as described by the database response, which is done as follows:

$$SF = \sqrt{\int_{80\,\mathrm{MHz}}^{800\,\mathrm{MHz}} \left|\mathrm{response}(f)\right|^2 \mathrm{d}f}. \tag{7.1}$$

Here, "response(f)" is the complex response from the database response. Then, we use the mean scaling factor across all considered upward-facing shallow channels to scale the impulse response. It is important to note that the database response was not gain calibrated. However, none of the contributing individual measurements were normalized, so the overall gain should be similar to that of a calibrated response. A

comparison of the resulting impulse response gain with the database response gain is shown in Figure 7.2. The shape of the different responses differs slightly, as expected, but the overall gains are now compatible. With the scaled gain, we can use the impulse response for the simulation and correctly trigger on the simulated waveforms.

Furthermore, the overall time delay of the impulse response is normalized. However, since the time delay is approximately the same across all shallow channels, this normalization does not affect the simulation results and can be disregarded.

At the time of writing, the impulse response is available with proper absolute gain calibration [171], removing the need for this scaling in future analyses.

7.3 Trigger description

To determine whether any simulated event would have been detected by an RNO-G station, we need to apply a trigger to all events. The trigger is designed to resemble the Radiant trigger as closely as possible. Further details about the trigger and its performance can be found in Section 5.3. For the UHECR analysis, only the shallow antennas are of interest, and thus, we model the Radiant trigger exclusively.

Although the downward-facing antennas are less sensitive to UHECR signals coming from above the ice, they can trigger a readout before the upward-facing antennas due to geometrical factors. Therefore, we simulate a trigger using the upward-facing antennas and another using the downward-facing antennas. Both triggers are coincidence triggers with a coincidence time of 50 ns and are modeled accordingly. Due to an oversight, the coincidence time was previously falsely reported as 60 ns in Ref. [163]. For the upward-facing antennas, the coincidence requirement is met if two of the three exceed the threshold, whereas for the downward-facing antennas, two of the six antennas must exceed the threshold within the coincidence time window. Additionally, to enhance the SNR, a bandpass filter is applied prior to the trigger decision. Although it was demonstrated that a passband of $80 - 120$ MHz is optimal [213], this configuration cannot be realized due to hardware limitations. Therefore, a passband of $80 - 200$ MHz is implemented on the Radiant, and for this simulation, a passband of $80 - 180$ MHz is used, which is closer to the theoretical optimal value, making our simulation more conservative. However, the influence of this small difference should be minimal. In the simulation, we apply as the bandpass filter a 10th-order Butterworth filter. The actual Radiant trigger is approximated by an envelope trigger[2], because a direct modeling of the trigger response was not possible for the Radiant (V2) board used in this work [172].

In order to generate realistic simulations and subsequently calculate the expected number of UHECR events, we need realistic trigger thresholds for the simulated trigger. As we approximate the trigger by an envelope trigger, we must map the trigger threshold measured during data acquisition to its simulation counterparts. This mapping, which was derived in Section 5.3.4, can now be applied.

As an illustrative example, we will derive the simulation trigger thresholds for station 13 during the 2022 data-taking period (2022-07-25 to 2022-10-01). First, we extract the

[2] For an envelope trigger, the signal envelope must exceed a given threshold.

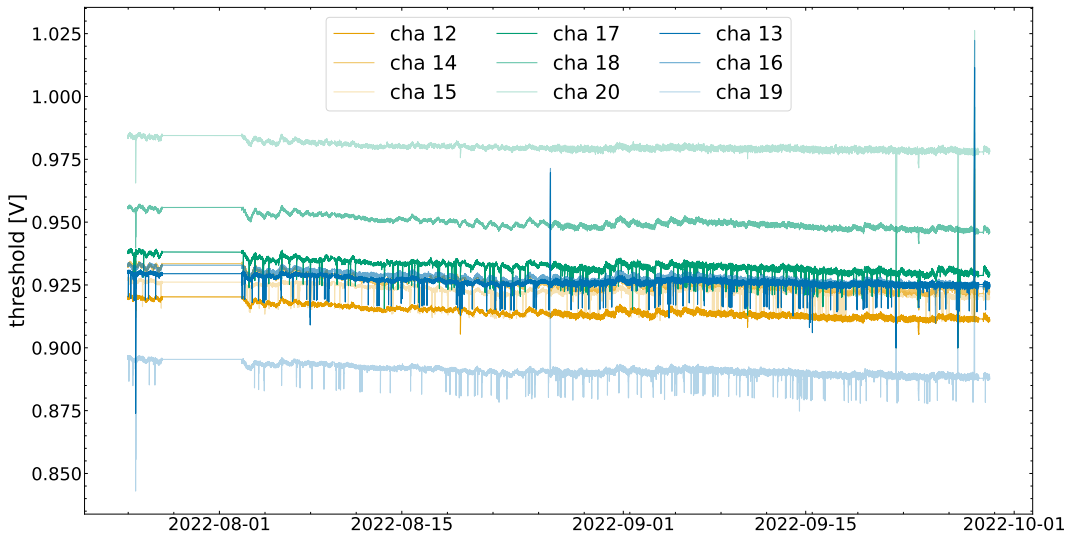


Figure 7.3: Trigger thresholds of station 13 measured during the 2022 data-taking campaign (2022-07-25 to 2022-10-01).

Table 7.1: The mean measured and the simulation thresholds. As an illustrative example, values for station 13 in 2022 are shown.

channel	mean τ_{measured} [mV]	$\tau_{\text{simulation}}$ [mV]	
		best case	worst case
13	926	59	99
16	928	58	98
19	891	71	111
12	914	64	103
14	926	59	99
15	923	60	100
17	933	56	97
18	950	48	91
20	980	30	81

trigger thresholds from the recorded data. Figure 7.3 shows the trigger thresholds of all shallow channels of station 13 in 2022. The thresholds across the different channels differ slightly and show non-constant features. This occurs because we can only specify the single-antenna trigger rate, and its threshold is adjusted accordingly based on the noise level of the channel. Additionally, transient noise sources (e.g., airplanes) can increase noise levels, requiring thresholds to be adjusted to meet their trigger rate goal. As a result, these transient noise sources can lead to spikes or abrupt changes in the threshold. Less abrupt temporal changes can result from factors such as temperature changes that affect the thermal noise, causing the thresholds to be adjusted to maintain trigger rate stability. Next, we calculate the mean threshold per channel over the entire data-taking period. The resulting values for the different channels are presented in Table 7.1. Subsequently, we use the mapping derived in Section 5.3.4 to translate the mean measured thresholds into simulation thresholds. Due to substantial uncertainties in this mapping, we report two thresholds: a best-case and a worst-case value. The best-case value corresponds to a more sensitive trigger and a higher expected UHECR event count, while the worst-case value corresponds to a less sensitive trigger. Both simulation trigger thresholds are displayed in Table 7.1. As a result, each channel of

each station is simulated using a custom threshold estimated from data during data acquisition. For station 23 only, the measured trigger thresholds map to unrealistically low values in the simulation under the best-case scenario due to the large uncertainties from the channel-to-channel variation and extrapolation. Therefore, for station 23, we use a simulation threshold of 25 mV[3] for all channels, rather than the best case scenario. The 25 mV threshold is lower than the best-case thresholds for the other stations, but not unrealistic. For stations 13 and 24, all measured trigger thresholds can be mapped to realistic simulation thresholds. The simulation with realistic trigger thresholds can be used to estimate the expected number of UHECR events per station (for details, see Section 8.5).

7.4 Noise description

To model the noise in the simulations, we consider two sources: thermal and Galactic noise. In the following, we describe the two different noise sources in more detail.

Thermal noise: All sources that contribute to the thermal noise are combined into a single effective temperature T_{noise}, which is subsequently used to calculate the expected noise power P_{noise} using the following formula [166]:

$$P_{\text{noise}} = k_{\text{B}} T_{\text{noise}} \Delta f, \tag{7.2}$$

where Δf is the frequency band for which the noise is calculated, and k_{B} is the Boltzmann constant. The noise temperature can be decomposed into contributions from black-body radiation from the ambient surroundings T_{env} and from the temperature of the receiving electronics T_{sys} [166]. Although an additional potential loss factor may arise from the transmission line between the antenna and receiver, for this work, we assume a lossless transmission line with no contribution. Therefore, the noise power can be written as [166]:

$$P_{\text{noise}} = k_{\text{B}} \left(T_{\text{env}} + T_{\text{sys}} \right) \Delta f. \tag{7.3}$$

The system temperature is estimated from noise temperature measurements to be $T_{\text{sys}} = 80\,\text{K}$ for the electronics relevant to the shallow antennas [163]. For the environmental contribution, the radiated temperature of the surroundings observed by the antenna is calculated as [166]:

$$T_{\text{env}} = \frac{\int_0^{2\pi} \int_0^{\pi} T_{\text{ice}}(\theta, \phi) G(\theta, \phi) \sin\theta d\theta d\phi}{\int_0^{2\pi} \int_0^{\pi} G(\theta, \phi) \sin\theta d\theta d\phi}, \tag{7.4}$$

Here, $T_{\text{ice}}(\theta, \phi)$ represents the temperature from the surrounding ice, and $G(\theta, \phi)$ denotes the gain pattern of the antenna. Both parameters are described as a function of the zenith θ and azimuth ϕ angle. The ice temperature at various angles is calculated by integrating the temperature along the ray paths and accounting for attenuation [171].

[3] The threshold of 25 mV corresponds to 9σ, where σ is the RMS of the simulated noise in the trigger frequency band.

Because the environmental temperature varies with frequency, we use the average over the trigger frequency band 80 − 180 MHz, yielding $T_{\text{env}} \approx 126\,\text{K}$. Combining both noise sources yields a total noise temperature for the shallow antennas of $T_{\text{noise}} \approx 206\,\text{K}$, which we apply in the simulations. Thermal noise is added to the simulated waveforms using a NuRadioReco module, which requires the noise V_{RMS} as an input. Here, V_{RMS} is the RMS of the resulting noise voltage waveforms, which can be calculated using the derived noise temperature. Using the relation $P = V^2/R$ [115], where R is the resistance of the circuit, and P is the power, and applying equation 7.2, we derive for the noise V_{RMS} the following relation:

$$V_{\text{noise}} = \sqrt{P_{\text{noise}} R} = \sqrt{k_{\text{B}} T_{\text{noise}} \Delta f R} \tag{7.5}$$

To calculate the noise V_{RMS}, we assume perfect impedance matching ($R = 50\,\Omega$) and the frequency band $\Delta f = 50 - 800\,\text{MHz}$, which approximately matches the Surface board frequency band. The resulting noise V_{RMS} serves as input to the module, which adds thermal noise to the simulated waveforms. The amplitudes of each frequency bin of the generated noise follow a Rayleigh distribution. The calculated thermal noise temperature is subject to some uncertainty, so we investigated the effect of a ±10 K change on the simulation and found it negligible (see Section 8.4 for more information).

Galactic noise: The second noise source is the diffuse radio emission from our Galaxy. In Ref. [163], it was demonstrated that the diffuse Galactic radio emission can be observed with RNO-G. This is seen primarily in the upward-facing antennas, which are most sensitive to this emission and most relevant for this analysis. Consequently, we must incorporate the Galactic noise as a background in the simulations. The Galactic noise is included using a NuRadioReco module, which takes a radio sky map and calculates the resulting noise on the simulated waveforms. For the simulations presented in this work, we use the 2016 Global Sky Model (GSM) [214] for the Galactic background radiation. Galactic noise varies throughout the day as the Galactic center moves closer or away from the horizon. Since the UHECR events are randomly distributed in time and to account for different Galactic noise contributions, we randomly select the time of day for each simulated waveform. However, the seasonal date for which the Galactic noise is simulated stays the same for all simulations. We evaluated the potential impact of seasonal variations on the simulation and found no significant effect, see Section 8.4 for more information.

Figure 7.4 compares the two noise sources that are considered in the simulations, before any amplification or digitization responses are applied. The figure demonstrates that the Galactic noise is the dominant contribution at low frequencies and becomes subdominant above approximately 200 MHz. Since the LPDAs are most sensitive at low frequencies, this underlines the importance of including the Galactic noise in the simulations.

Future simulations could benefit from exploring whether more realistic noise can be generated using data-driven noise models rather than simulating individual contributions. Furthermore, the impedance matching of components and the validity of assuming $R = 50\,\Omega$ should be further assessed. Finally, in this study, we assume negligible losses and contributions from the transmission line between the antenna and receiver. Future work should further investigate this assumption.

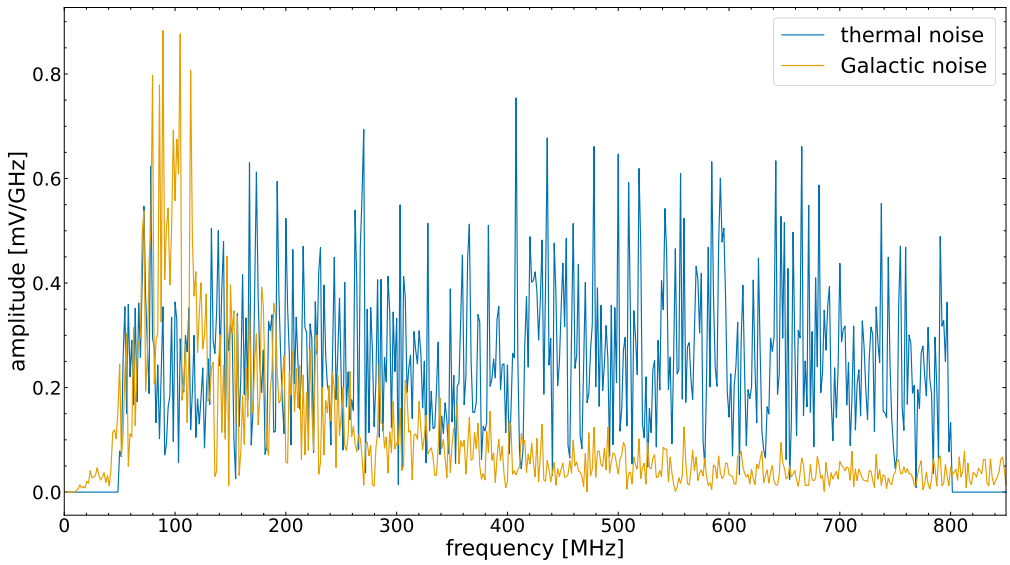


Figure 7.4: The frequency spectrum of simulated thermal and Galactic noise, before any amplification or digitization responses are applied.

7.5 Effective area and expected rate

In this section, we outline the methodology for calculating the effective area and expected event rate for a simulation. The effective area is defined as the region in which a signal from an air shower is sufficiently strong to trigger the detector. This parameter is highly sensitive to the trigger settings and must be determined individually for each simulation. Additionally, the effective area depends on the zenith angle and the energy of air showers. Higher-energy air showers produce signals with larger amplitudes, increasing the likelihood of triggering the detector. The zenith angle influences the projected footprint of the air shower on the ground. For more inclined air showers, which have larger footprints, the radiation energy is dispersed over a wider area, thereby reducing the amplitude of the detectable signal. More vertical showers produce smaller footprints, requiring the air-shower core to be closer to the antenna to be detectable. For the simulations in this work, we calculate the effective area for each simulation as a function of both zenith angle and energy $A_{\text{eff}}(E, \theta)$. In the simulation, we sample each air-shower footprint at 240 positions, as described in Section 7.1, with each position corresponding to a specific area of the footprint [213]. In the following, we refer to this specific area as the area weight, which is associated with the respective sampled position. To determine the effective area, we group all simulated air showers, including those with a different detector rotation, into energy and zenith-angle bins. First, for each air shower within a bin, we sum the area weights of the positions where the signal triggered the detector, yielding the total area weight for that air shower. Then, we take the mean of these total air-shower area weights within each bin as the effective area for the corresponding energy and zenith angle. If none of the sampled positions trigger, we set the total area weight for that air shower to zero and include it in the effective area calculation for that bin.

To obtain the expected rate of UHECR events in the detector, we combine the previously calculated effective area with the UHECR flux $J(E)$. We calculate the expected rate

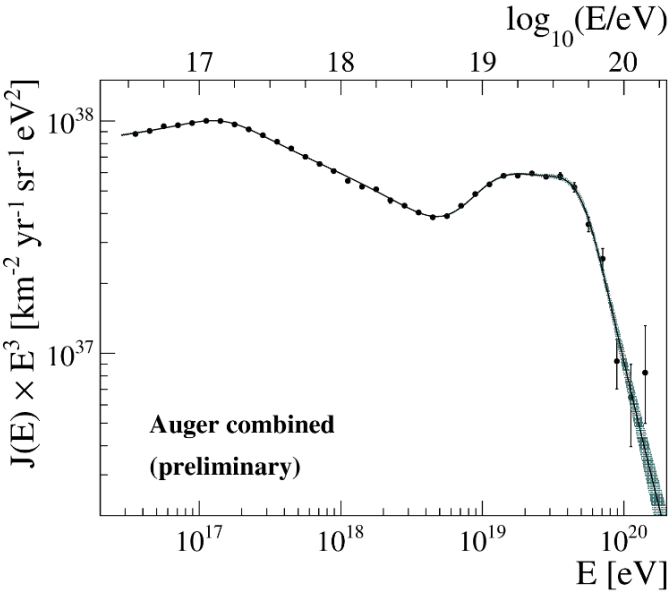


Figure 7.5: UHECR flux published by the Pierre Auger Observatory. It combines measurements of different parts of the detector. Taken from Ref. [215].

(Γ_{UHECR}) per energy and zenith angle bin by evaluating the following integral:

$$\Gamma_{\mathrm{UHECR}} = \int_{E_1}^{E_2} A_{\mathrm{eff}}(E,\theta)\cos(\theta)J(E)\Omega dE, \tag{7.6}$$

Here, E_1 and E_2 denote the boundaries of the energy bin, $\cos(\theta)$ is a projection factor reflecting that only the orthogonal component of the UHECR flux contributes due to the flat geometry of the detector [172], and Ω is the solid angle. The solid angle is calculated as follows.

$$\Omega = \int_0^{2\pi}\int_{\theta_1}^{\theta_2} \sin(\theta)d\theta d\phi, \tag{7.7}$$

In this expression, θ_1 and θ_2 define the boundaries of the zenith angle bin, and ϕ represents the azimuth angle. Since all azimuth angles are within the detector's field of view, we perform the integration over the entire azimuthal range. For the UHECR flux, we take a measurement from the Pierre Auger Observatory [215], as shown in Figure 7.5. This spectrum shows a sharp decrease in flux at higher energies, indicating that although higher-energy events are more likely to trigger the detector, we expect the majority of UHECR events to occur at lower energies around 10^{17} eV. Events with even lower energy are less likely to be detected due to limitations in the trigger and the noise level.

To estimate the expected number of UHECR events (N_{UHECR}) from the simulation, we sum the calculated UHECR rate over all zenith angles and energy bins and integrate over the measurement time period.

$$N_{\mathrm{UHECR}} = \int \sum_{i}^{\mathrm{bins}} \Gamma_{\mathrm{UHECR},i}\, dt. \tag{7.8}$$

Additionally, we must consider analysis efficiencies when calculating the expected number of UHECR events. For example, background rejection cuts may exclude a certain percentage of simulated events, as we discuss further in Section 8.5.

7.6 Conclusions

In this chapter, we presented a framework for conducting realistic UHECR simulations for RNO-G. First, we defined the parameter space for the simulations and described the process of sampling the 384 generated air showers to obtain a representative simulation set. The resulting simulation set contains over 3 million events, each with a waveform for every shallow antenna. To reduce the amount of storage needed, we only save high-level quantities, such as the SNR value, for these events. As previously discussed, sampling from so few generated air showers may introduce some statistical problems. To address this, Philipp Laub generated additional air showers using CORSIKA (version 7.6400) and CoREAS (version V1.3). However, due to time constraints, we were unable to incorporate these new air showers into the current analysis. Preliminary studies indicate that both air-shower libraries are compatible. Future work should more thoroughly assess the compatibility of these two libraries and, if confirmed, combine both to expand the number of available air showers. Moreover, in the current framework, we sample the air-shower footprint at fixed positions. A tool is now available to interpolate the electric field between the footprint positions [216]. Future work should evaluate the impact of this interpolation tool on the simulation.

Furthermore, we describe the antenna, amplification, and digitization responses applied to convert the simulated electric fields into voltage waveforms compatible with measured data. The absolute gain of the used amplification and digitization response was unavailable at the time of the analysis, which required additional scaling to obtain realistic values. Notably, an initial version of the gain calibration is now available, allowing future analyses to use channel-specific gain calibration rather than applying a fixed scaling factor.

Additionally, we introduced the approximated trigger that is implemented in the simulation, along with a discussion of the various settings. As an illustrative example, we presented the process of converting a measured trigger threshold into a simulation threshold by using the mapping derived in Section 5.3.4, for station 13 during the 2022 data-taking period. With this example, we highlight the uncertainty associated with the mapping by displaying the resulting thresholds for both best- and worst- case scenarios.

Furthermore, we explained the process to generate the noise for the simulations, which consists of two contributions: thermal and Galactic noise. The thermal noise is the sum of the black-body radiation of the surrounding ice and the noise generated within the amplifiers. We showed that the combined temperature of both contributions is approximately 206 K. For the conversion from temperature to voltage, we assume an ideal impedance match to 50 Ω. Future work should assess the accuracy of this assumption and determine whether deviations in the real impedance significantly influence the noise. Additionally, we assumed a loss-free cable without any contribution to the noise

temperature between the antenna and receiver, which future work should investigate. For the Galactic noise, we use a module from NuRadioReco that calculates the noise at the antennas based on a radio emission sky map. Moreover, future work could compare a data-driven noise model to simulated noise to further enhance the realism of the simulations.

Finally, we outlined the procedure to calculate the effective area and the expected number of UHECR events. The presented formulas represent an ideal scenario with perfect analysis efficiencies. We address the analysis efficiencies in Section 8.5.

8 Cosmic Ray Analysis

In this chapter, we present the first analysis searching for ultra-high-energy cosmic ray (UHECR, $E_{\mathrm{CR}} > 10^{16}$ eV) events using the shallow antennas of RNO-G. This analysis has been published in Ref. [185] and therefore, substantial portions of this chapter closely follow or are directly adapted from that publication. We begin by outlining the general strategy of the analysis (see Section 8.1) and introducing the method for generating template waveforms that are matched to the data to identify UHECR candidates (see Section 8.2). Subsequently, we specify the dataset used for the analysis (see Section 8.3) and discuss both dominant and sub-dominant systematic effects (see Section 8.4). We then describe the analysis structure and present the results (see Section 8.5). Furthermore, we present the reconstructed properties of all UHECR candidate events (see Section 8.6) and examine whether the identified UHECR candidate events can be attributed to various backgrounds (see Section 8.7). Finally, we analyze data from stations excluded from the main analysis due to their poor data quality (see Section 8.8).

8.1 Analysis strategy

In previous analyses with the ARIANNA experiment, it was shown that a template search can be successfully applied to identify UHECR air shower signals [17, 145]. RNO-G uses the same LPDAs for the shallow detector as ARIANNA, therefore, we can use a similar template search to identify cosmic-ray air showers. In a template search, a set of template waveforms is correlated with the data waveforms and the resulting correlation score is used as the discriminative variable. The correlation score for which an event is considered as UHECR changes as the correlation of data and template improves with signal-to-noise ratio (SNR). Therefore, an SNR-dependent correlation score threshold is used for the UHECR event selection.

However, using a large template set – as in the initial ARIANNA analysis [17], which was not optimized for a small number of templates and includes $\mathcal{O}(200,000)$ template waveforms - makes this method computationally expensive, in particular for the comparatively much larger dataset of RNO-G. By exploiting the fact that the signal shape is predominantly determined by the antenna, amplification, and digitization responses and in particular by the dispersion, in a previous study we have been able to reduce the number of template waveforms to three, while still sufficiently covering the parameter space [217]. To generate the template waveforms, an artificial electric field is created. The two polarization components of this electric field are described by Gaussian functions. Then, the antenna, amplification, and digitization responses are applied to the electric field to obtain the voltage template waveform. Detailed information about the template creation procedure is provided in Section 8.2. In the end, only two template waveforms with different frequency contents are used as the template set for the analysis. Both template waveforms, which are referred to as signal templates, are shown in Figure 8.1. Adding the third template to the template set only marginally increases the coverage by $\mathcal{O}(2\%)$ of all simulated UHECR events, at the cost of allowing more background to pass,

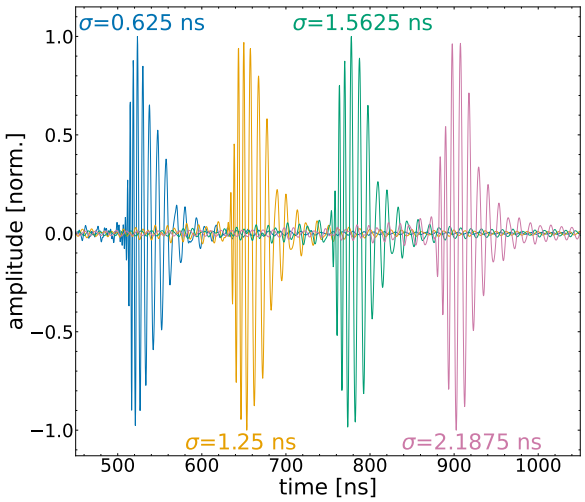


Figure 8.1: Shown are from left to right the waveforms of the two signal (σ = 0.625 ns, σ = 1.25 ns) and the two background templates (σ = 1.5625 ns, σ = 2.1875 ns). The Gaussian width σ is defined in waveform samples and subsequently converted to nanoseconds, which explains the very precise values. Taken from Ref. [185].

as will be shown in Section 8.5.3. Figure 8.1 also shows two background-like template waveforms, which are based on Gaussian functions with a larger width, such that they match low-frequency backgrounds and not UHECR signals.

The template search is based on a cross-correlation of signal and template. The correlation score, χ, is calculated as follows:

$$\chi = \max\left(\left|\frac{\sum_{i=0}^{m}(V_1)_i \cdot (V_2)_{i+\Delta n}}{\sqrt{\sum_{i=0}^{m}(V_1)_i^2} \cdot \sqrt{\sum_{j=\Delta n}^{m+\Delta n}(V_2)_j^2}}\right|\right), \tag{8.1}$$

with V_1, V_2 being the voltage waveforms of the template and data, respectively, Δn the number of samples by which the two waveforms are shifted relative to each other and m being the number of samples of the template waveform. For the calculation of the correlation score, a scan over all possible Δn is performed, while only the maximum of the scan is taken as the final value. Moreover, only a 200 ns window around the pulse of the template waveform is considered for the calculation to reduce the influence of noise. The correlation parameter, χ, lies between 0 and 1, with 1 corresponding to a perfect match. Due to the normalization of the correlation parameter, the overall amplitudes of the data and template waveforms have no effect.

8.2 Cosmic-ray templates

In this section, we describe the creation of the UHECR templates. To generate the template waveforms, we first construct artificial representations of the electric fields. We then apply the antenna, amplification and digitization responses to these artificial electric field to produce the final template waveform. Figure 8.2 illustrates this process for two template waveforms with different frequency content. The following paragraphs provide a more detailed explanation of this procedure.

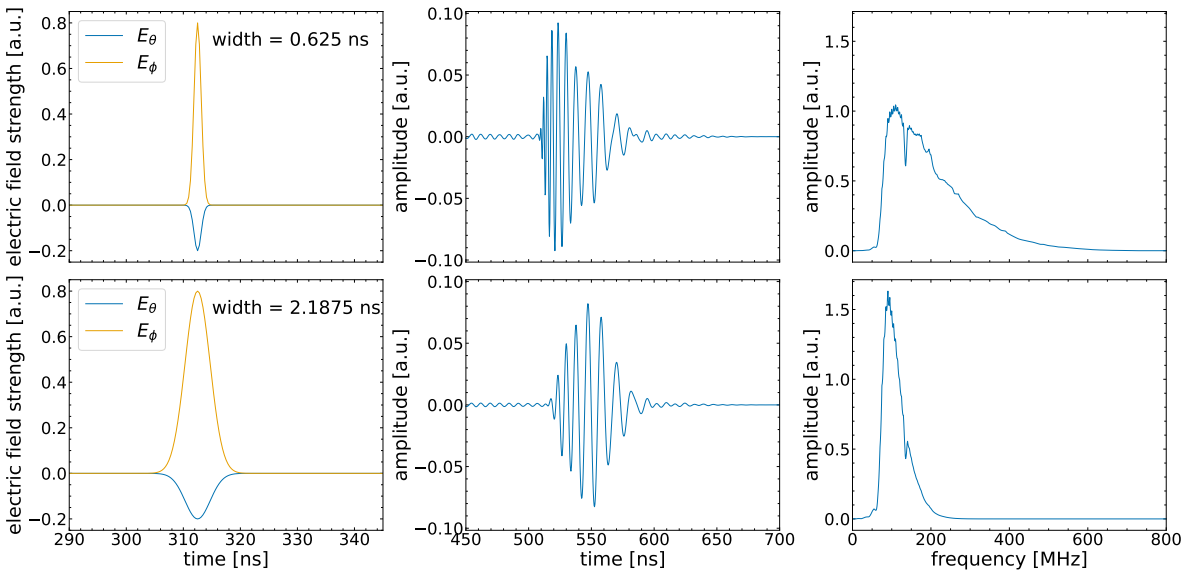


Figure 8.2: Illustration of the generation of two templates with different frequency content. Left: The artificial electric fields. Middle: The resulting template waveforms after applying the antenna, amplification, and digitization responses to the artificial electric field. Right: The frequency spectra corresponding to the waveforms shown in the middle panels.

For template creation, we model both polarization components of the electric field (E_θ, E_ϕ) using Gaussian functions, which is a simplification since the true electric field exhibits a bipolar shape [107]. However, the dominant factors shaping the final template waveform are the antenna, amplification and digitization responses, rather than the precise form of the incoming electric field. Therefore, it is most important for template creation that the electric field used is a nanosecond-scale long pulse. The width of the Gaussian functions determines the frequency content of the resulting template waveform. As we illustrate in Figure 8.2, a template generated with a very short electric-field pulse contains frequencies up to 400 MHz. In contrast, a template generated with a wider pulse exhibits lower frequency content. Templates with varying frequency content are necessary to represent the full range of possible UHECR air shower signals, as the frequency content of an air shower electric field varies with its position in the radio footprint (see Section 3.1.3).

To identify the optimal Gaussian width, we scanned over a range of widths. Since performing a full simulation for each tested Gaussian width would have been computationally expensive, we selected a random air shower from the shower library (see Section 7.1) and sampled it at 30 positions along a single arm of the star-shaped pattern, probing different distances to the shower axis. The distance to the shower axis will also be denoted as "radius" in the following. We processed the sampled electric fields by applying the antenna, amplification, and digitization responses, without including noise and then correlated the resulting waveforms with the templates generated using different Gaussian widths. Figure 8.3 presents the results of this scan, which inform our template selection. In this figure, we express the Gaussian width in samples, which we can convert to nanoseconds by dividing by the 3.2 GHz sampling rate of the RADIANT board. The results indicate that templates with different widths achieve the highest correlation scores at different radii. Templates with smaller widths correlate best at approximately 200 m, while others peak at smaller or larger radii. We attribute this

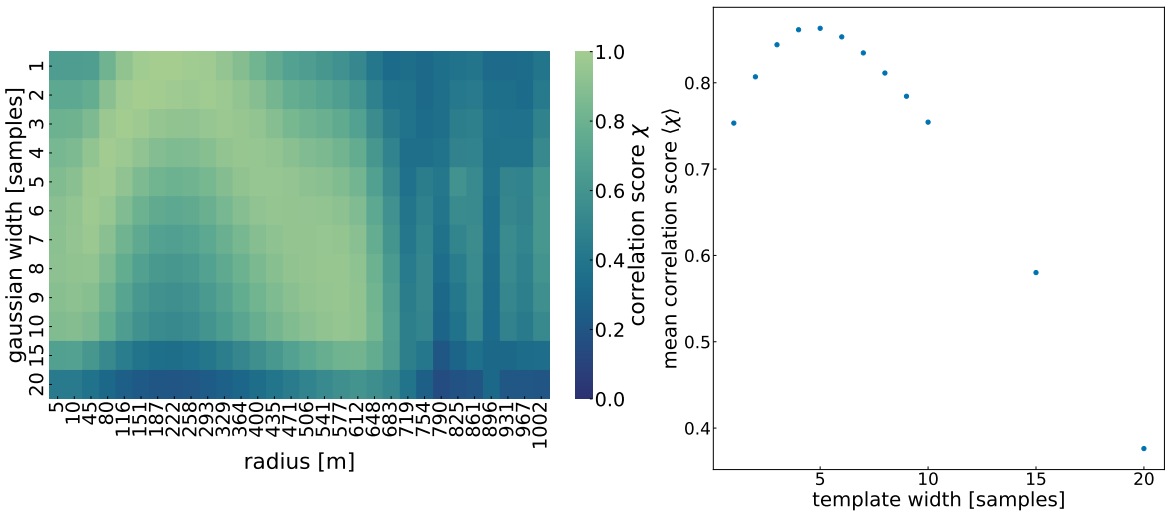


Figure 8.3: The results of a scan over Gaussian widths for template creation. Left: Correlation scores for a template with a waveform from a simulated air shower, as a function of the template width (vertical axis) and the radial distance from the shower axis (horizontal axis). Right: The mean correlation score for each template, excluding the outermost simulated waveforms, for which the simulation noise becomes dominant.

pattern to the time compression of the radio signal, which increases the high-frequency content near the Cherenkov ring. Since we sample the simulated radio footprint along a radial path, we observe only a segment of the Cherenkov ring rather than the full ring. Beyond around 620 m, correlation scores decline and become irregular, as the simulated signal is weaker than the simulation-generated noise at these sampling positions. We exclude these noise-dominated points from further analysis, as they do not contribute any valuable information about the air shower. To select the Gaussian width, we calculate the mean correlation score across all simulated radio waveforms. Figure 8.3 shows the mean correlation scores as a function of the Gaussian width of the template waveform. Based on these results, we proceed with two templates corresponding to the Gaussian widths yielding the highest correlation scores: one with a width of 5 samples (1.5625 ns) and another with a width of 4 samples (1.25 ns). The resulting template waveforms, after applying the antenna, amplification and digitization responses, are shown in Figure 8.1.

The next parameter to specify is the relative amplitude of the Gaussian functions for both polarization components. As an initial estimate, we select a random air shower from the shower library and sample it at three points within the footprint. From these sampled electric fields, we estimate the amplitude ratio of the two polarization components and use it as the amplitude estimate: $E_\theta = -0.2$ a.u. and $E_\phi = 0.8$ a.u.. Since the correlation score is invariant to global scaling, the absolute amplitude is not of interest, and we provide the amplitudes in arbitrary units and normalize them so that they sum to one. To assess whether alternative amplitude ratios improve performance, we conduct a small study using the same simulated waveforms as in the Gaussian width scan and correlate them with templates generated with different amplitude ratios. Figure 8.4 presents the results, showing that the correlation score only varies minimally with amplitude ratio, even when one polarization is set to zero. Global sign changes in the amplitudes are not considered, as they only invert the resulting waveform and do not affect the correlation score. Based on this study, we use the initially estimated amplitudes from simulations, as no other amplitude ratio provides improvement across all radii.

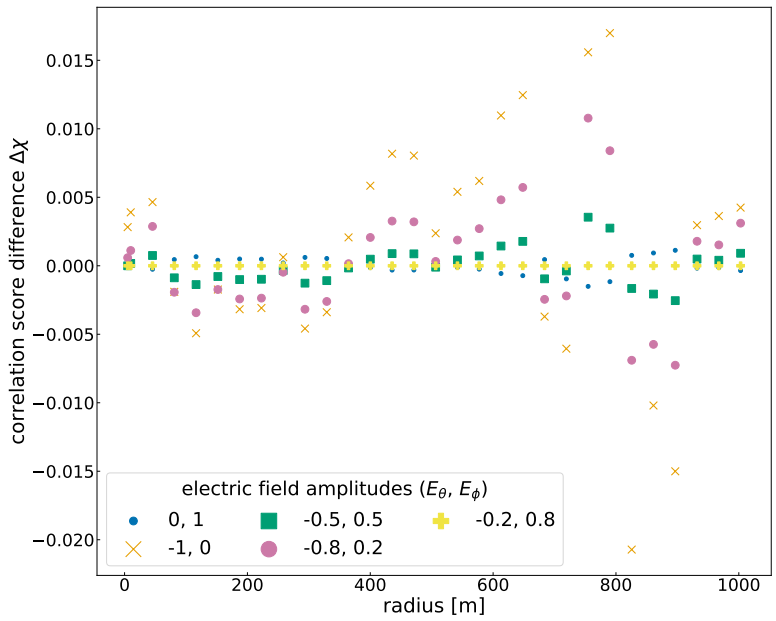


Figure 8.4: The results of a scan over electric-field amplitudes used in template creation. For different electric field amplitudes, templates are generated and correlated with simulated waveforms from a single air shower. Shown is the difference in correlation score relative to the templates created with the amplitudes used in the UHECR analysis, as a function of the distance from the air-shower core. All correlation score differences with a distance greater than approximately 600 m can be neglected, as these events are dominated by simulation noise.

The final parameter for the artificial electric fields is the relative delay between the two polarization components, which is controlled by the position of the Gaussian functions on the electric field waveform (μ_E). For both polarization components, we use the same value, so we do not introduce a relative time delay. Birefringence, an effect that can introduce a relative travel time difference of the two polarizations, should be negligible over the short propagation distance through ice. The absolute location of the two electric field pulses on the waveform is unimportant, since we computed the correlation score by scanning the entire waveform. Table 8.1 summarizes the electric field parameters we discussed in the previous paragraphs, along with all other parameters necessary to generate a template waveform.

To convert the artificial electric field to a voltage waveform, we need to apply the antenna response, which depends on the zenith and azimuth angle of the incoming electric field relative to the antenna. We apply the antenna response using NuRadioReco[211], assuming the artificial electric field originates from an air shower arriving at an antenna with zenith angle θ_{CR}, azimuth angle ϕ_{CR} and antenna rotation angle $\phi_{antenna}$[1]. After being refracted into the ice, θ_{CR} determines the zenith angle for the antenna response (VEL), while the other two angles set the azimuth. For all three, we use fixed values without further optimization, since a single set of angle values sufficiently covers the entire parameter space, as we show later in this section. The specific values are provided in Table 8.1

To estimate the extent to which the two templates cover the parameter space we use in the simulations (see Section 7.1), we conduct a small, noiseless simulation. We employ the same simulation framework as we describe in Chapter 7, but without adding noise

[1] The angle describes the antenna rotation around the axis that goes through the bottom and tip of the antenna and is perpendicular to all tines.

Table 8.1: The parameters employed to generate the two signal and two background templates used in the UHECR analysis.

	signal templates		background templates	
gaussian width [ns]	0.625	1.25	1.5625	2.1875
gaussian width [samples]	2	4	5	7
E_θ [a.u.]	−0.2	−0.2	−0.2	−0.2
E_ϕ [a.u.]	0.8	0.8	0.8	0.8
ϕ_{antenna} [°]	160	160	160	160
θ_{CR} [°]	55	55	55	55
ϕ_{CR} [°]	0	0	0	0
μ_{E} [samples]	1000	1000	1000	1000

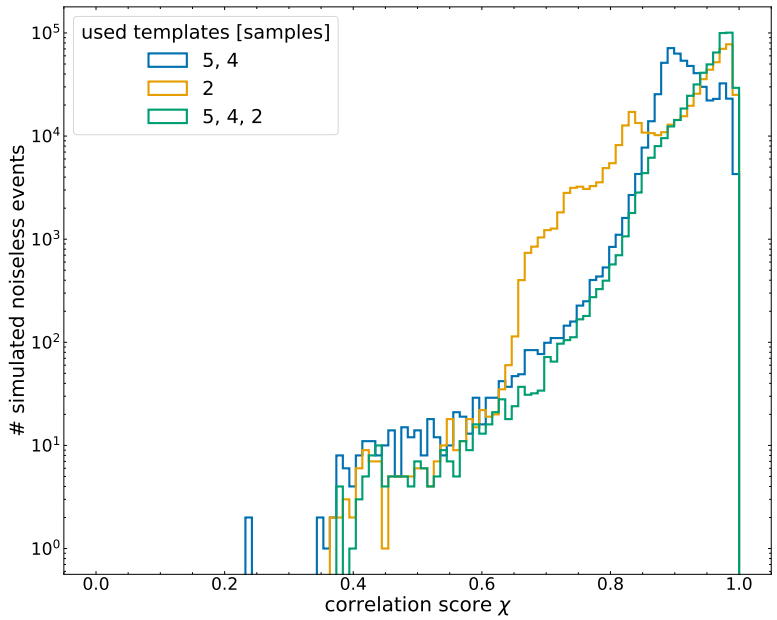


Figure 8.5: The correlation scores for simulated noiseless UHECR waveforms correlated with different template sets.

to the simulated waveforms, which we then correlate with the two created template waveforms. For the detector, we use a single antenna to reduce the number of correlations to calculate, and a trigger threshold of 20 mV, which corresponds to 7.2σ.[2] The trigger threshold is well below the smallest threshold we use in the analysis, 25 mV (9σ), and thus provides a conservative estimate for the considered parameter space. Figure 8.5 shows the distribution of correlation scores when, for each event, we consider only the maximum value between the two templates. Treating each simulated waveform as a point in the parameter space and assuming that the point is covered by templates with a correlation score greater than 0.8, we find that of the 523925 triggered events, only 3371 have a correlation score below 0.8, corresponding to 0.6%. For the UHECR analysis, we consider the parameter-space coverage provided by these templates sufficient. Furthermore, it remains possible that simulation noise-dominated waveforms or simulation artifacts could account for some of the 3371 events. In addition, we used a very conservative trigger threshold, so some simulations with a correlation score below 0.8 may not trigger if we were to apply more realistic thresholds. The sufficient coverage of the parameter

[2] In this context, σ represents the RMS in the trigger frequency band of the simulated noise typically used.

Table 8.2: Mean correlation scores for simulated noiseless UHECR waveforms correlated with different template sets.

used templates [samples]	mean correlation score $\langle\chi\rangle$
5, 4 (standard)	0.9158
5, 4, 1	0.9473
5, 4, 2	0.9531
5, 4, 3	0.9428
5, 4, 6	0.9170
5, 4, 7	0.9172
5, 4, 8	0.9170
5, 4, 2, 1	0.9537

space is also reflected in the full simulations, including noise, which we use to calculate the SNR-dependent analysis cut (see the example in Figure 8.25). These simulations show no outliers with a very low correlation score.

Moreover, Figure 8.5 indicates that most events are clustered around a correlation score of 0.9, with fewer events achieving scores closer to 1. To improve this distribution and shift it toward higher correlation scores, we investigated adding additional templates. We repeated the small simulation described above for different template sets, and we calculated the mean correlation score across all triggered events for each case. As additional templates, we consider those created with Gaussian widths of 1, 2, 3, 6, 7, and 8 samples. First, we added a single template to the template set and the resulting mean correlation scores are presented in Table 8.2. The mean correlation score improves when we add a template with higher-frequency content (lower Gaussian width) to the template set. We observe the largest improvement with the template of width 2 samples (0.625 ns). Additionally, we evaluated the effect of adding two templates to the previous set of two templates on the mean correlation score. As the two additional templates we select those that performed best before: templates with a Gaussian width of 2 and 1 samples. The resulting mean value, also shown in Table 8.2, is only marginally larger than the best-case mean value obtained with a single template. Each additional template increases the computational cost of calculating the correlation score for the dataset or simulation. Therefore, we only included the template with a Gaussian width of 2 samples in the template set, as the improvement from a fourth template is minimal. The distribution of correlation scores for the noiseless simulation with the three-template set is shown in Figure 8.5. With the third template included, only 2281 events have a correlation score below 0.8, corresponding to 0.4%, indicating a slightly increased coverage of the parameter space. Furthermore, we show the correlation score distribution for the noiseless simulation, which is correlated only with the template with a Gaussian width of 2 samples, to demonstrate that this template cannot fully replace the other two templates.

It should be noted that the values presented above may differ from those in Ref. [217] for several reasons. First, compared to Ref. [217], we used an improved description of the amplification and digitization response[3] for template creation. Second, in Ref. [217], we used fewer air showers from the air shower library to cover the parameter space, which is possible but requires scaling the electric fields to allow lower-energy air showers to

[3] The improved description is the impulse response description. More information is provided in Section 5.1.2.

pass the trigger threshold. In this work, we use all air showers in the air shower library for the noiseless simulation to avoid this scaling. Nevertheless, the main conclusion presented here remains consistent with that obtained in Ref. [217].

The situation becomes more complex when we consider background events, such as those produced by the triboelectric effect during high-wind periods.[4] Background events often exhibit only a low-frequency content and therefore correlate better with templates that also have a low-frequency content. Consequently, we developed a background-rejection cut to exclude all events with the highest correlation score with the template referred to as the background template (see Section 8.5). To create the background template, we use the same procedure described previously. We include a dedicated background template with a Gaussian width of 7 samples (2.1875 ns) in the template set. Additionally, we use one of the previously described templates, specifically the one with a Gaussian width of 5 samples (1.5625 ns), as a background template. We found that rejecting all events with the highest correlation score with the background templates results in a loss of only $\mathcal{O}(2\%)$ of the simulated UHECR events, while rejecting a greater proportion of background events. Therefore, we will use templates with Gaussian widths of 4 samples (1.25 ns) and 2 samples (0.625 ns) as signal templates, and those with Gaussian widths of 5 samples (1.5625 ns) and 7 samples (2.1875 ns) as background templates. The waveforms of all four templates are shown in Figure 8.1.

We performed several checks to inform the selection of the template parameters. As the parameter estimates are based on a relatively small number of simulated air showers, they may not be entirely optimal. Further optimization of these values may be pursued in future work. However, this is a complex issue, as the goal is not only to maximize the correlation score but also to maximize the difference between the correlation with UHECR events and that with background events.

8.3 Dataset

For the analysis presented in this work, we consider data from 2022 and 2023, which are treated as separate datasets. Data from 2024 is excluded due to changes in the detector configuration, including modifications to the station setup at the beginning and replacement of the RADIANT board during the data-taking period. During the time periods considered for this work, 7 RNO-G stations were deployed and operational. Figures 8.6 and 8.7 show the trigger rates as a function of time for the shallow antennas for all 7 stations.

Two different triggers contribute to the shallow trigger rate, both require coincident signals in multiple LPDAs. One triggers on the upward-facing LPDAs with a 2 out of 3 coincidence, while the other one triggers on the downward-facing LPDAs with a 2 out of 6 coincidence. Both triggers do not have constant trigger thresholds, but are automatically adjusted to follow the variable noise level. However, these adjustments are done per antenna and thus the resulting coincidence trigger rate can still have a non-constant behavior. This adjustment is done to avoid saturating the data-rate during periods of increased noise. Although generally the upward-facing antennas are more

[4] A discussion of possible background events is provided in Section 4.4.

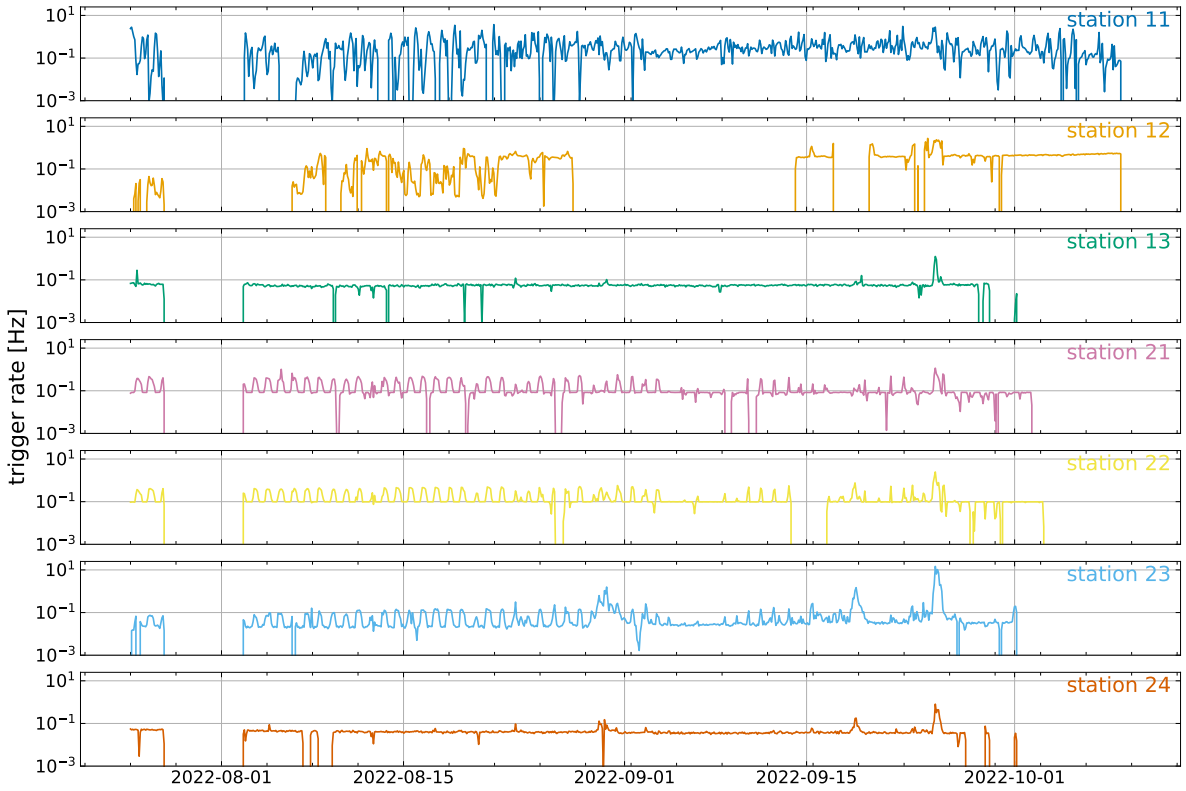


Figure 8.6: The trigger rate (the sum of upward and downward-facing triggers using the LPDAs) as a function of time for all science data of the seven stations in 2022. Taken from Ref. [185].



Figure 8.7: The same as Figure 8.6 for the year 2023. The shaded regions for station 23 show the different configurations of that station. The first configuration is marked by the dark shaded region, while the first light shaded region marks configuration 3 and the second marks configuration 2 (see text for details). Taken from Ref. [185].

sensitive to air shower signals, certain event geometries result in downward-facing antennas fulfilling the trigger condition before the upward-facing ones. Since only the first trigger is recorded, both triggers are relevant for this analysis. Events from other trigger conditions of RNO-G, in particular from the phased array, are not considered for this analysis.

We use the trigger rate to quantify the performance of the stations and find that stations 13 and 24 show the most stable performance (see Figure 8.6 and 8.7). Although station 23 shows small day-night variations in 2022, they occur at a more sensitive trigger threshold compared to the trigger rate variations of stations 21 and 22, making it more feasible to use data from station 23. As a result, only data from the stations 13, 23, and 24 are included in this analysis. In Section 8.8 we will have a closer look at the data quality of the other stations and the findings will support the trigger rate motivated decision. The higher and unstable trigger rates of the other stations can be explained by external factors: During the considered time period, stations 11 and 12 were equipped with a prototype wind turbine, which is thought to be responsible for the increase and large fluctuations in the trigger rates due to a poorly shielded power box [163]. The large day-night variations in the trigger rate of stations 21 and 22 can be attributed to an early version of the power system used by both stations [163], which led to an increased noise level in the shallow antennas during the charging of the batteries from the solar panels. In both cases, we were able to design improvements, and both the wind turbines and power boxes of these stations were exchanged during the 2024 season. As a result, an improvement of the data quality for these stations is expected for the data from mid-2024 onward and will benefit future analyses.

Stations 13, 23, and 24 were installed during the 2022 season and began regular data taking after a short commissioning period. Consequently, the 2022 data-taking period started in July and lasted until October, while the 2023 period lasted from April until October. The exact dates, resulting lifetimes and number of triggered events per station can be found in Table 8.3. We used a month of data from the 2022 data-taking period for the development of the analysis. This month of data is also included in the dataset that is used for the full analysis.

During the 2023 data-taking period, station 23 shows a special behavior compared to the other stations. Without changing the station settings, compared to the 2022 period, the trigger rate went up (the trigger thresholds were raised automatically), leading to a decreased sensitivity of the station. Furthermore, two periods can be identified where the trigger rate drops sharply and stays at a lower level before recovering. In what follows, we separate the 2023 period for station 23 into three different configurations and treat them separately for the analysis and simulation. The first configuration (config 1) covers everything not included in the other configurations, which is the longest period (3119 h lifetime) and can be associated with the standard behavior for 2023. The second configuration (config 2) starts on the 2023-08-01 and lasts the shortest (96 h lifetime). Without changing the run settings, the trigger rate went down (the trigger thresholds were lowered automatically), which leads to a more sensitive configuration, similar to the 2022 period. For the last configuration (config 3), that starts on the 2023-06-27, the downward-facing trigger was turned off and thus the trigger rate is a little lower than normal (220 h lifetime). The downward-facing trigger was turned off to help investigate the sudden changes in the trigger rate, which are described above. A detailed

Table 8.3: The exact start/end dates, lifetime, and number of triggered events for the dataset used (LPDA trigger only), separately for 2022 and 2023. Taken from Ref. [185].

station	dates	lifetime [h]	# triggers
13	2022-07-25 2022-10-01	1410	306,869
23	2022-07-25 2022-10-01	1453	680,902
24	2022-07-25 2022-10-01	1352	217,516
13	2023-04-24 2023-10-12	3929	726,444
23	2023-04-24 2023-10-14	3434	4,476,619
24	2023-04-21 2023-10-11	3604	653,415

investigation of the behavior of station 23 is presented in Section 5.4, but the underlying cause of this behavior remains unidentified.

The main dataset is composed of waveforms from the upward-facing LPDAs, since they are generally more sensitive to air-shower signals than downward-facing antennas. However, the waveforms of the downward-facing antennas are used for a data reduction cut, and additionally thermal noise waveforms from the upward-facing channels are utilized to determine some of the background rejection cut values (see Section 8.5). The thermal noise waveforms are taken using a forced trigger that records a waveform every 10 s. During the data taking of 2022 and 2023 the readout window for the deep and shallow antennas could not be aligned to capture both a signal in the LPDAs and the deep channels, due to analog delays in the system. Thus, the deep antennas do not provide any information for shallow triggered events and are not considered in this work. We do note that the problem with the misaligned readout window has been mitigated for all stations in 2024 and will enable the study of coincidence events with the shallow and deep antennas in future analyses.

8.4 Systematic uncertainties

The effectiveness of the template search depends on how closely the template matches the UHECR waveforms. This is primarily determined by the antenna and hardware (amplification and digitization) responses of the detector. Consequently, an excellent understanding of the antenna and hardware responses is crucial. Simulations indicate that even small differences in the hardware response of the search template compared to the one used for the simulation can lead to a significant reduction of the correlation score, of $\mathcal{O}(0.1)$, relative to the ideal case where both use the same response. We present the relevant simulations in Section 5.1.1, with the observed differences shown in Figure 5.3. These findings emphasize the importance of an accurate description of the hardware response. In this work, we use a high-quality description of the hardware response,

which we refer to as the impulse response, and it is introduced in Section 5.1.2. We find that the channel-to-channel and station-to-station variance of this measurement is negligible, allowing the same template set to be applied to all channels and stations.

To ensure that systematic effects beyond the hardware response are accounted for, we have investigated a variety of factors and their potential impact. A complete list of all performed checks is provided in Table 8.4 together with their final impact on the correlation parameter χ. In the following, we first describe the main systematic uncertainties (see Sections 8.4.1, 8.4.2, 8.4.3, and 8.4.4), and subsequently discuss the systematic effects we found to be negligible (see Section 8.4.5).

8.4.1 Refractive index

The uncertainty in the refractive index n_{firn} of near-surface ice is the dominant systematic uncertainty in this analysis, as well as potentially in all studies using data from shallow antennas. A change in n_{firn} slightly alters the antenna response, which directly translates into an uncertainty on the data-template correlation score if an incorrect n_{firn} is used for the template creation. To quantify this effect, we perform simulations as described in Section 6.4. We conduct, for each station and data-taking period, a custom simulation study to account for the different estimated trigger thresholds when calculating the systematic uncertainty. For the simulations, we use the standard $n_{\text{firn}} = 1.3$ for the simulated UHECR waveform, while generating the templates with a different n_{firn}. The antenna response at different n_{firn} values is obtained by shifting and rescaling the frequency-dependent gain of an antenna simulated in air. Further information regarding this rescaling is available in Section 6.3. In Section 6.2, we estimated the range of possible n_{firn} at LPDA depth to be $1.28 < n_{\text{firn}} < 1.335$, based on different ice models. Two template sets are created, using the extreme values of the estimated n_{firn} range for the antenna response. These modified template sets are correlated with the simulated UHECR waveforms, and SNR-dependent correlation score cut values are calculated for each case. The difference with respect to the standard SNR-dependent correlation score cut is taken as the systematic uncertainty, conservatively choosing the larger difference of both cases. The resulting uncertainty is SNR-dependent and can reach as much as 0.08 in the correlation score. Improvements in the understanding of the near-surface n_{firn} can reduce the systematic uncertainty related to the refractive index and thereby significantly reduce the overall systematic uncertainty.

In addition to the uncertainty in the shallow refractive index across different ice models, further uncertainties may be introduced during deployment. For the LPDA, trenches are excavated and subsequently refilled with snow after the antenna is placed inside. This process may result in small air pockets remaining or denser areas of snow being introduced between the LPDA tines. As a result, the medium surrounding the antenna may become inhomogeneous, introducing additional uncertainty in the refractive index. We are not aware of studies that investigate this effect, but such research would be valuable for future work.

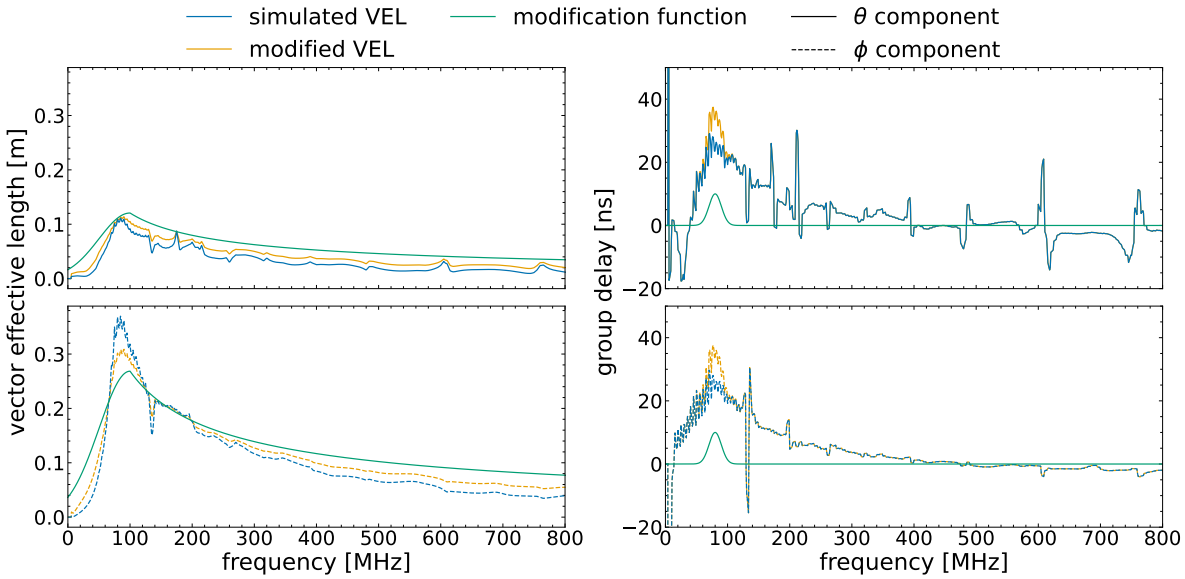


Figure 8.8: Two examples of a modified VEL alongside the original VEL. Left: The θ and ϕ components of a VEL, with only the gain modified. Right: The same simulated VEL as shown in the left panel with the group delay modified. In both cases, the function applied to modify the VEL is shown as well.

8.4.2 Antenna response

As described in Section 7.2, we use a WIPL-D simulation of the LPDA in ice [17] for the antenna response. Currently, no measurements are available that would allow us to confidently quantify variations in the response in situ. Consequently, we estimate a systematic uncertainty due to differences in the antenna response (VEL) shape. To do this, the antenna response is modified by adding or removing power in different parts of the frequency spectrum. In total, we applied three types of variations to the VEL. We modified the group delay by adding a Gaussian function, and altered the gain either by adding a Gaussian function or by multiplying by a $1/f$ function. For the $1/f$ function, we introduced a cutoff to prevent divergence as f approaches zero. At frequencies below the cutoff, we used a Gaussian function peaking at the cutoff. The magnitude of these modifications was informed by VEL measurements conducted in an anechoic chamber [163]. For illustration, Figure 8.8 displays the gain of a modified VEL altered by a $1/f$ function and a group delay modified by a Gaussian function. For each modified response, a new set of templates is generated. Given the large number of created template sets ($\mathcal{O}(30,000)$), it is not computationally feasible to perform simulations with each set. Instead, we correlate the templates with UHECR candidate events, found in a small subset of the data which was selected for the analysis development. To eliminate potential effects arising from the use of waveforms recorded with different hardware and antennas, we consider only UHECR candidate waveforms with the highest correlation score for the same channel and recorded at the same station. The resulting differences in the correlation score compared to the standard template are taken as guidance to estimate a conservative systematic uncertainty. The resulting uncertainty is found to be 0.015 in correlation score, which is non-negligible.

Notably, we altered the VEL only through simple functions and independently modified the gain and group delay. In practice, the discrepancy between simulated and actual

Table 8.4: List of all considered systematic effects and their possible influence on the correlation score χ. Taken from Ref. [185].

systematic effect	impact on χ
refractive index	< 0.08
antenna response	0.015
timing calibration	0.01
hardware response measurement	0.01
amplifier saturation	< 0.01
voltage calibration	< 0.01
statistical fluctuations	< 0.01
station-to-station variation	< 0.01
Galactic noise	< 0.01
noise temperature	< 0.01
antenna depth & rotation (simulation)	< 0.01
total	< 0.083

VEL may be more complex. Future analyses would benefit from a more detailed and quantitative comparison of measured and simulated VEL. All studies utilizing the LPDAs and that are sensitive to the signal shape could benefit from such investigations. Moreover, the LPDA response may be altered due to small antenna deformations during deployment. At present, we are not aware of any detailed studies that have examined the effects introduced by bent antenna tines. Further investigations in this area would be beneficial for future work.

8.4.3 Timing calibration

The waveforms are digitized using a custom chip (ASIC), named LAB4D, exploiting a switch-capacitor-array (SCA) (more information can be found in Ref. [218]). To record meaningful waveforms with the digitizer, it is necessary to tune the charging times of each of the capacitors in the SCA. The tuning is a non-trivial and time-consuming process, and the resulting charging times can be worse than expected from the nominal performance of the chip. However, the templates are created with the assumption of a perfect digitization, which leads to a lower-than-expected correlation score. We investigate the effect on the correlation score by simulating waveforms that are sampled with a poorly tuned digitizer and correlate them with perfectly sampled waveforms. For this toy simulation, we utilize the UHECR template waveforms.

For digitization, the waveform is divided into multiple blocks, each containing 128 samples, which are sampled by two separate SCA units, each with 64 samples. This configuration presents two challenges for the tuning algorithm: tuning of the charging times of each of the 128 capacitors in one of the two sampling cells (referred to as sample-to-sample timing), and ensuring the correct connection between the different blocks of 128 samples (referred to as window-to-window timing). Poor window-to-window timing can introduce artifacts in the recorded waveform and degrade the correlation score. Unfortunately, we could not identify a robust method to quantify window-to-window timing, so we cannot provide a realistic estimate of a potential systematic effect. For the

UHECR analysis, we assume that the windows are well-matched and do not influence the resulting correlation score. In the following paragraphs, we describe how we simulate the SCA to investigate the influence of poor sample-to-sample timing on the correlation score.

After charging for a time Δt, a capacitor on the SCA will have a total charge Q of [115]

$$Q = \int_{\Delta t} CU(t)dt, \tag{8.2}$$

where C denotes the capacitance of the capacitor and $U(t)$ represents the applied voltage. For the digitization of an analog signal, $U(t)$ corresponds to the incoming analog voltage waveform. The LAB4D then calculates the constant voltage U_{const} that must be applied to the capacitor over a nominal time period $\Delta t_{\text{nominal}}$ to yield the same charge Q. To derive the required formula for U_{const}, we first evaluate the integral above for a constant voltage and then solve the result for U_{const}. The result is as follows:

$$U_{\text{const}} = \frac{Q}{C\Delta t_{\text{nominal}}} \tag{8.3}$$

Next, we can substitute the charge Q with the equation for the total collected charge under a time-dependent voltage (Equation 8.2):

$$U_{\text{const}} = \frac{\int_{\Delta t} CU(t)dt}{C\Delta t_{\text{nominal}}} \tag{8.4}$$

This equation describes the digitized voltage of a sample of the SCA as a function of the analog input voltage $U(t)$ and the nominal and actual charging time. We can further expand this equation by a factor of $\Delta t/\Delta t$, resulting in the final form:

$$U_{\text{const}} = \frac{\int_{\Delta t} CU(t)dt}{C\Delta t_{\text{nominal}}} \frac{\Delta t}{\Delta t} = \langle U(t) \rangle \frac{\Delta t}{\Delta t_{\text{nominal}}} \tag{8.5}$$

Here, $\langle U(t) \rangle$ denotes the average analog input voltage over the charging time Δt. The sampling rate of the digitizer determines the nominal time, which is for the data used in the UHECR analysis: $\Delta t_{\text{nominal}} = 312.5\,\text{ps}$. The charging time Δt must be tuned to match the nominal time as closely as possible, a process that is both time-consuming and non-trivial. A difficulty is that when the charging time of a single capacitor is increased by $\alpha\%$, the charging times of the remaining 127 capacitors must be decreased by $\alpha/127\%$ to maintain a constant overall recorded time [218].

We implemented a toy simulation to generate a realistic set of 128 charging times that deviate from the nominal value within a specified range, quantified by the RMS of the difference. With these charging times, we sample a template waveform that exhibits poor sample-to-sample timing using Equation 8.5 and correlate it with a template waveform sampled with ideal timing (where the charging times equal the nominal value for all capacitors). The resulting correlation value estimates the influence of poor sample-to-sample timing on the analysis. We repeat this procedure 100 times to account for statistical fluctuations in the generation of the charging times. Figure 8.9 presents the results of toy simulations using the two signal templates as waveforms and different RMS

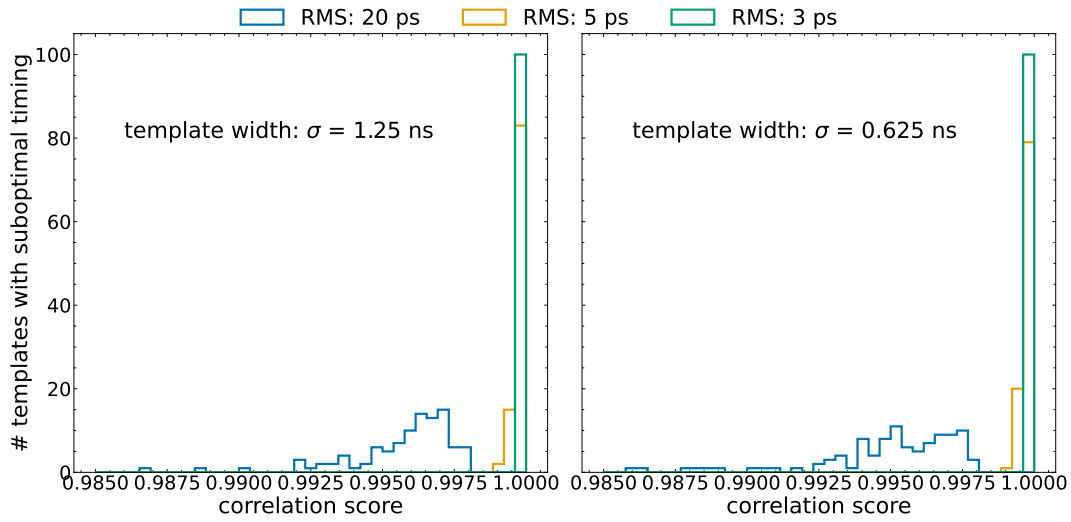


Figure 8.9: Histogram of the correlation scores between a template waveform sampled with ideal and suboptimal sample-to-sample timing. Different colors represent the RMS values of the different charging times used for the suboptimal sampling. For each RMS value, the toy simulation was conducted 100 times. The two panels display the results for the two signal template waveforms.

values for the charging times: 3 ps, 5 ps and 20 ps. The first two values are RMS values obtained in Laboratory tests with the LAB4D [218], representing optimal performance, while the latter reflects a realistic value achieved during Laboratory tests with the fully integrated RADIANT boards. Due to the non-trivial and time-consuming tuning process, it is realistic to assume that the RMS of the charging times during data-taking is in the range of $3-20$ ps. The correlation values for a sample-to-sample timing with RMS values of 3 and 5 ps consistently exceed 0.998, indicating negligible influence on the analysis. However, poorer sample-to-sample timing can reduce the correlation score to 0.99 or slightly lower for some events. Thus, the effect is subdominant but not negligible.

To ensure that the observed differences are not artifacts of the charging-time generation process, we also tested the influence of actual charging times on the correlation score. Additionally, data from the RADIANT tuning indicates that the first and last capacitors in the 128-block can diverge significantly, potentially introducing an even larger bias. The first and last capacitors in the 128-block are also referred to as seam and slow sample, respectively. Unfortunately, for the relevant time period of the analysis, the charging times from the RADIANT boards were not recorded. However, charging-time data from tests of the subsequent RADIANT version (V3), which uses the same LAB4D version, are available, and we expect it to exhibit similar behavior. From a single RADIANT board, we took the charging times of all 24 channels (LAB4Ds) and performed the same resampling and correlation as described above. Figure 8.10 shows the resulting correlation values as a function of the RMS of the measured charging times per channel. The results indicate that the correlation values for realistic charging times with an RMS below 20 ps remain above 0.99, consistent with the simulation outcome. Only one set of charging times at these RMS values produced a lower correlation value for both template waveforms, likely due to the seam and slow sample, which shows a significantly larger charging-time difference than the other samples. Additionally, we observe a decline in the correlation score for RMS values exceeding 20 ps, as anticipated. However, in deployed stations, if the seam or slow sample deviates excessively or if the RMS value is too large, the tuning process is

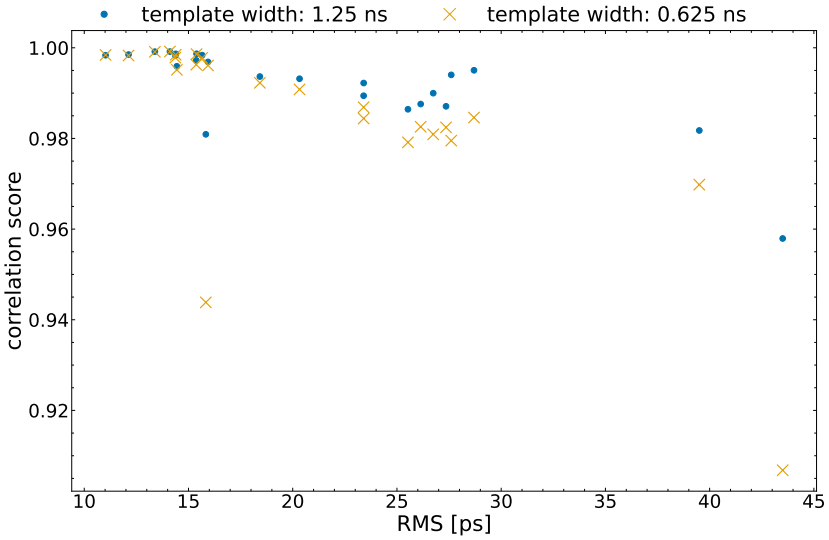


Figure 8.10: The correlation scores between a template waveform sampled with ideal and suboptimal sample-to-sample timing, plotted as a function of the RMS of the charging times used for the suboptimal sampling. The charging times are taken from a measurement conducted with a V3 RADIANT board, where each RADIANT channel provides a set of charging times.

considered unsuccessful and must be repeated, making these extreme cases unrealistic. Furthermore, the correlation scores indicate that the template with higher-frequency content is more strongly affected by poor sample-to-sample timing. Especially for RMS values above 20 ps, the difference between the two template waveforms with different frequency content becomes significant.

We consider a systematic effect of 0.01 in the correlation score realistic, which can occur when the sample-to-sample timing produces an RMS of the charging times of up to 20 ps. Therefore, this time calibration can introduce a subdominant but non-negligible effect, which is accounted for in the UHECR analysis.

8.4.4 Amplification and digitization response

A high-quality characterization of the influence of the hardware response on the signal shape is achieved by measuring the full chain with an impulsive signal. A more detailed explanation of this impulse response measurement is provided in Section 5.1.2. However, factors such as thermal noise, temperature or the overall stability of the measurement setup can have an influence on the precision. Temperature is of particular interest, as the RADIANT board inside the DAQ box experiences daily and seasonal temperature variations. Figure 8.11 displays the temperature inside the DAQ boxes during the 2023 data-taking period for stations 13, 23, and 24. To assess the impact of the external factors on the correlation score, the impulse response measurement was repeated at both the same (room temperature) and different temperatures (0 and −5°C). These repeated measurements were conducted at temperatures lower than those typically reached inside the DAQ boxes. However, components such as the coaxial cables outside the DAQ box, are exposed to lower temperatures. Additionally, if we do not observe a clear temperature effect at these differences, we can exclude an effect for smaller temperature variations within the DAQ boxes.

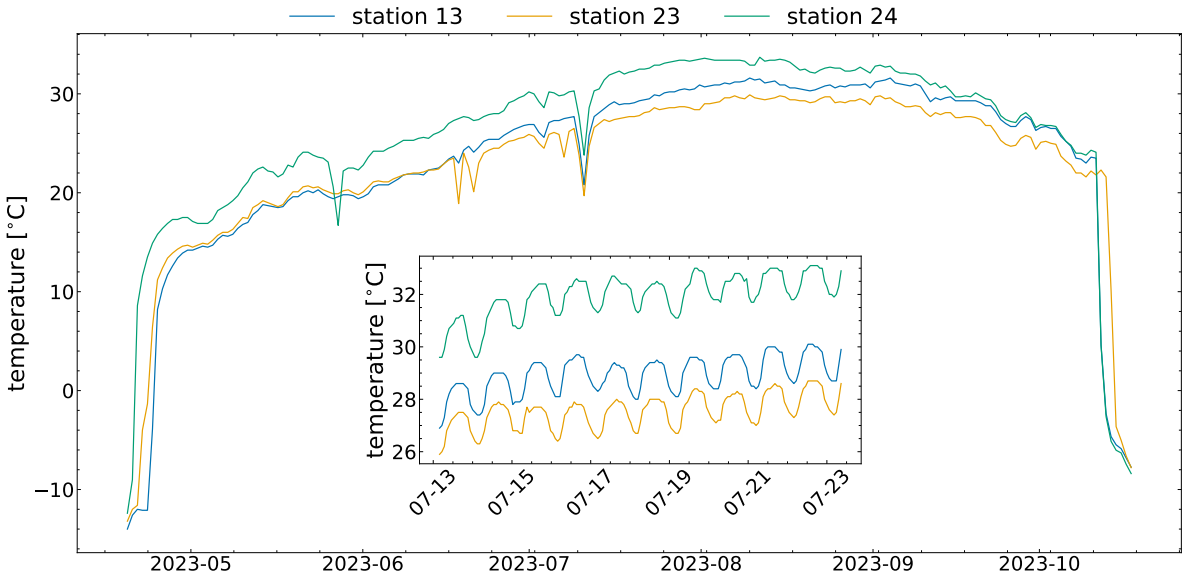


Figure 8.11: Temperatures measured inside the DAQ boxes during the 2023 data-taking period. The inset provides a detailed view of approximately 1.5 weeks. At the beginning and end of the period, the DAQ box cools down because most of its internal electronics are off and do not generate heat.

For each of the measured responses, we generate a template set, which is correlated with UHECR candidate events found in a subset of data that is used for the analysis development. Figure 8.12 shows the difference in correlation score relative to the values we calculated using the standard templates, which we also created from an impulse response measured at room temperature. Notably, the largest difference in correlation scores occurs for the two cases where we used an impulse response measured at room temperature for the template creation. Since all impulse response measurements involved were conducted at room temperature, we can exclude a temperature effect as the source of this variation. This observation suggests that the impulse response measurement setup introduces an inherent uncertainty in the correlation score. The two measurements at lower temperatures exhibit a slightly smaller spread, but the difference in correlation score is mostly positive, indicating that the standard template shows a larger correlation score with the UHECR candidate events. This result suggests the presence of a temperature effect. However, the magnitude of the correlation score difference is comparable to the inherent uncertainty, making it difficult to quantify the strength of a potential temperature effect. Also, the differences we observe at 0 °C and −5 °C are very similar, indicating that any potential temperature effect is not sufficiently strong to be detected at such small temperature differences.

We conclude that the impulse response measurement exhibits an inherent uncertainty, which may arise from thermal noise, instability of the measurement setup, or temperature effects. From the additional measurements, we conservatively estimate the uncertainty in the correlation score to be 0.01, which is subdominant but non-negligible. This uncertainty is represented in Figure 8.12 as the grey band, which contains the majority of data points.

Currently, the impulse response is measured with all amplification and digitization components at room temperature, which we consider an adequate simplification given the inherent measurement uncertainty. For future studies, if the inherent uncertainty

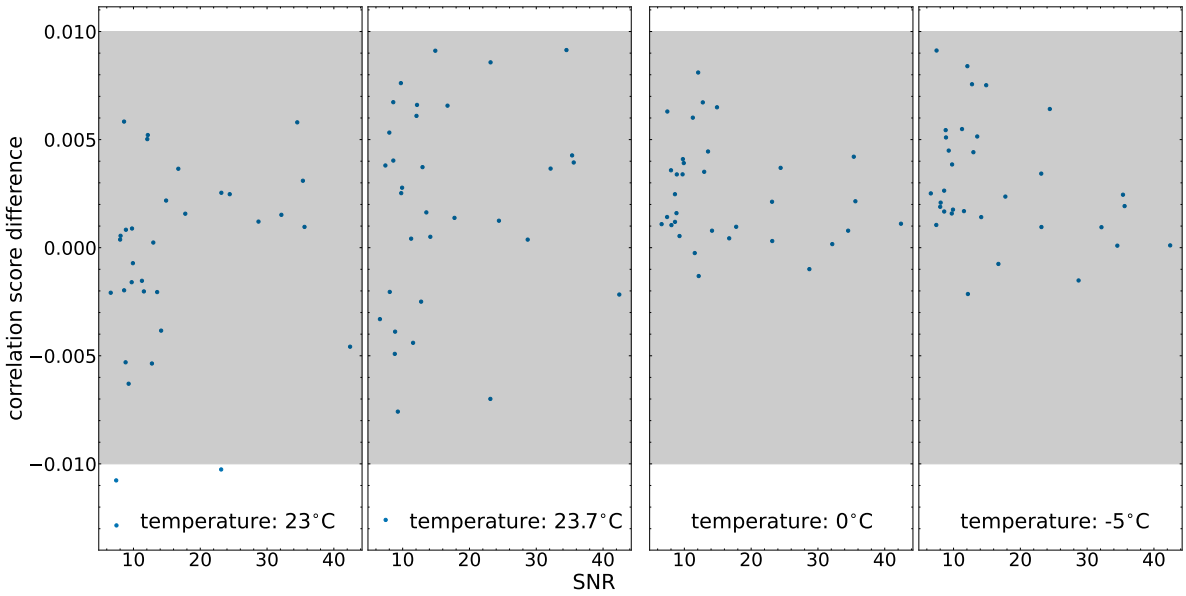


Figure 8.12: The difference in correlation score for template sets created using impulse responses obtained at different temperatures, relative to the standard template set, correlated with UHECR candidate events. The grey band represents the systematic uncertainty derived from these measurements. A positive difference in correlation score indicates that the standard template exhibits a higher correlation score with the UHECR event than the template created at the specified temperature.

is reduced, measuring the impulse response for each component at its operational temperature may be of interest. The influence on the response for shallow channels is likely minimal, as only the coaxial cable is located outside the DAQ box. In contrast, a larger effect may be present in deep channels, where additional components and also amplifiers, such as the IGLU board, are outside the DAQ box and thus at different temperatures.

8.4.5 Additional investigated systematic effects

We find that only the four previously discussed systematic effects contribute with a correlation score of 0.01 or greater. In the following, we describe further potential systematic effects that we found to be negligible. To quantify the impact of a systematic effect on the UHECR simulation, we use the SNR-dependent correlation score cut, calculated after applying the background rejection cuts (see Section 8.5). For both the standard simulation and the simulation including the systematic effect, we determine the SNR-dependent correlation score cut[5] and compare them bin-by-bin. Unless otherwise specified, we use the framework described in Chapter 7 for all UHECR simulations performed in this section, with a trigger threshold of 25 mV, corresponding to 9σ.[6] In general, we consider a systematic effect negligible compared to the previously mentioned effects if it changes the correlation score by less than 0.01. If we compare correlation score cuts, the difference must exceed 0.01 across multiple bins to be considered non-negligible.

[5] The SNR-dependent correlation score cut is set to retain 99% of simulated UHECR events per bin (see Section 8.5.4 for more information), unless otherwise stated.

[6] In this context, σ corresponds to the RMS of the simulated noise in the trigger frequency band.

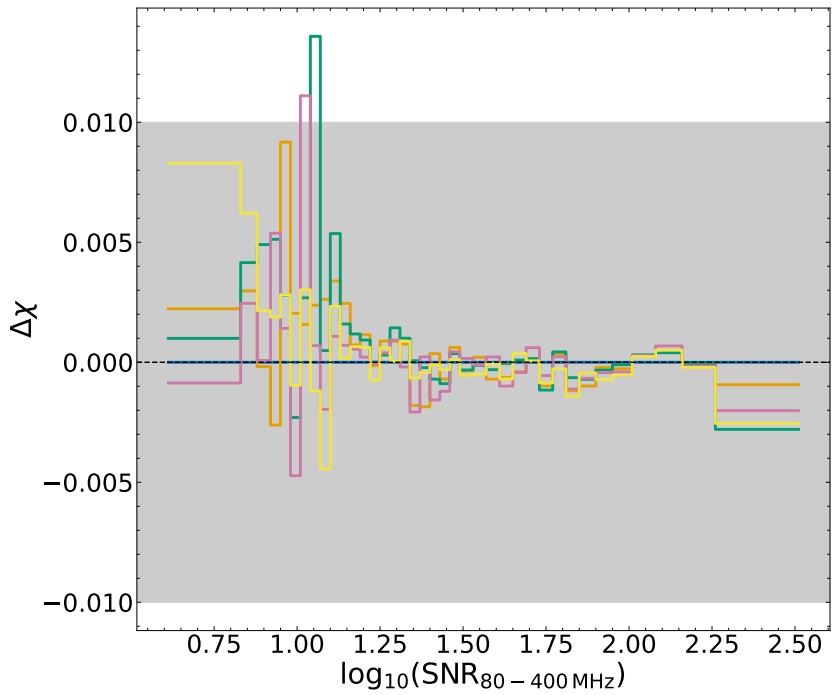


Figure 8.13: The difference in the SNR-dependent correlation score cut, calculated across four simulations with different noise realizations. The difference is always calculated relative to the same simulation, and the four colors represent the results for the different noise realizations. The grey band indicates the region where we consider the difference in the correlation score negligible.

Statistical fluctuations: We analyzed the impact of random noise fluctuations by repeating the simulation four times with different noise realizations for both thermal and Galactic noise. Although four repetitions are insufficient to rigorously quantify the uncertainty arising from noise fluctuations, conducting a sufficient number is computationally expensive. Nevertheless, these four simulations provide an indication of the impact the noise fluctuations have on the analysis. Figure 8.13 shows for all four simulations the difference in the correlation score cut relative to the same reference simulation. As the reference simulation, we selected one of the four simulations. At higher SNR values, the fluctuations remain below a correlation score of 0.003. However, at lower SNR values, all four simulations exhibit a systematic shift towards positive differences in correlation scores, which is unexpected. The systematic bias remains below a correlation score of 0.005 in most bins and exceeds 0.01 in only isolated cases. We expected random noise fluctuations to alter the correlation score of each event, resulting in bin-by-bin variations in the correlation score cut that are randomly distributed around zero. As anticipated, noise has a greater impact at lower SNR values. Overall, the difference due to random noise fluctuations is negligible compared to other systematic effects across all SNR values.

Nevertheless, the systematic shift towards positive differences is unexpected, and we have not yet found a clear cause for this behavior. In a preliminary study, we found that approximately 90% of the simulated events with $\log_{10}(\text{SNR}) < 1.25$, which mainly contribute to the correlation score cut and correspond to the lower tail of the correlation score distribution, vary between simulations. The systematic shift may be related to these findings. Another hypothesis is that the systematic shift arises from dividing the simulation into SNR bins for the calculation of the correlation score cut. Events located

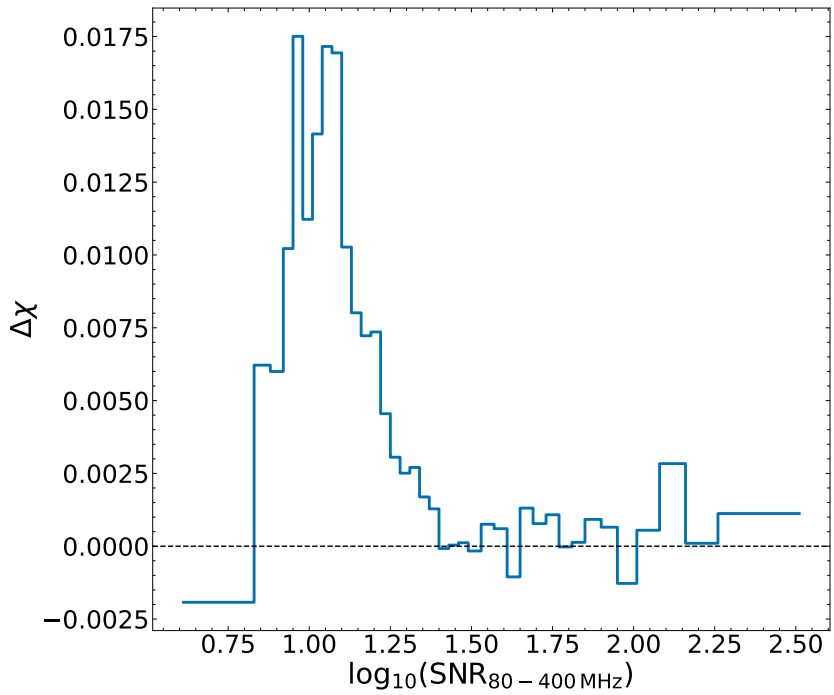


Figure 8.14: Changes in the SNR-dependent correlation score cut when the SNR value of each event is slightly varied.

near the bin edge may move into the neighboring bin without strongly changing the correlation score due to the new noise realization. To investigate this effect, we conducted a toy simulation in which we randomly varied the SNR values slightly and recalculated the analysis cut. Then, we compared the resulting correlation score cut to the original cut, without SNR value scrambling. The difference between the two cuts is shown in Figure 8.14, which demonstrates the same systematic behavior as we observed for low SNR values in Figure 8.13. However, we should approach the interpretation with caution, as the study is highly simplified. We only changed the SNR value, whereas in practice, a change in noise realization would also alter the correlation score. Furthermore, the shown strength of the effect in the toy simulation should be interpreted carefully, as we chose the degree of SNR scrambling to illustrate the effect rather than estimating it from a simulation or data. We conclude that further investigation is required to fully understand the origins of this systematic effect and that, for the UHECR analysis, the introduced systematic effect is negligible compared to other effects.

To study the other systematic effects, a fixed noise realization is used to exclude the influence of the random fluctuations.

Amplifier saturation: We investigated whether a potential nonlinearity in the amplifiers (Surface boards) at high amplitudes could negatively affect the signal shape and, consequently, the correlation score. Nonlinearity effects are expected to affect only very high-amplitude events. To assess the potential nonlinearity, we injected UHECR-like waveforms into the Surface board using an arbitrary waveform generator [193] and recorded the output with an oscilloscope. We conducted the measurement at multiple input amplitudes to span the operational range of the amplifier. For each amplitude, we performed repeated measurements of the same waveform and averaged the results to minimize noise. Then, we correlate the averaged waveforms with a reference waveform,

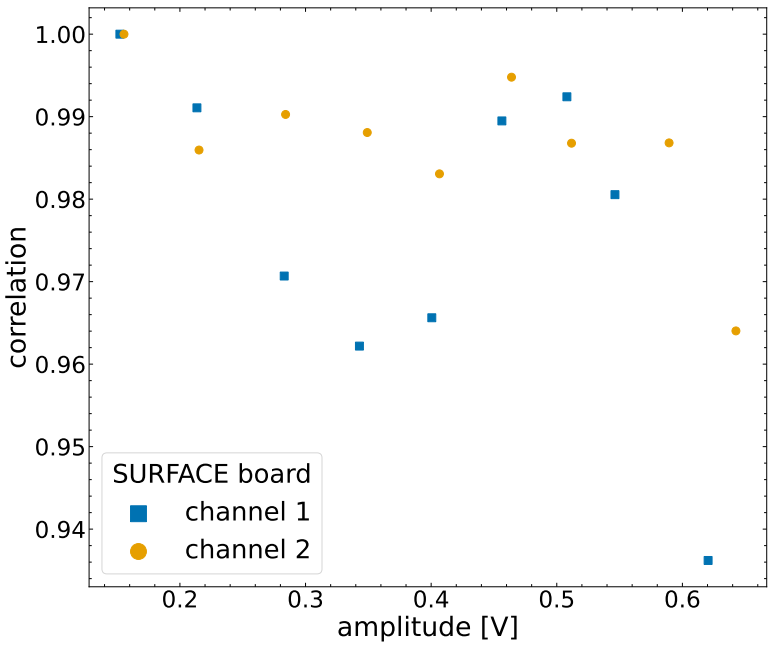


Figure 8.15: Results of amplifier linearity measurements, in which UHECR-like waveforms are passed through a SURFACE board and correlated with a reference measurement. The reference measurement corresponds to the lowest amplitude measurement.

defined as the measurement at the lowest amplitude. Figure 8.15 shows the resulting correlation scores for two different channels of a SURFACE board. Overall, the resulting correlation scores remain high (> 0.96) for all amplitudes except the highest, where we observe a notable change. If the amplifier nonlinearity were responsible, we would expect at the higher amplitudes a progressive decline in the correlation score with increasing amplitude. However, the data exhibit a constant behavior as a function of amplitude but show substantial scatter. Only at the highest amplitude do we find an indication of an amplitude-dependent decrease. Based on these data, we conclude that up to 600 mV, we do not observe an effect attributable to amplifier nonlinearity. For higher amplitudes, preliminary evidence of nonlinearity appears, requiring further investigations. Additionally, future work should aim to reduce the observed scatter in the measurements. The highest energy UHECR candidate event we identified in this analysis has an amplitude of approximately 500 mV. Therefore, even a minor indication of nonlinearity at the highest energy does not significantly impact the present analysis. We conclude that the laboratory measurements revealed no significant effect to this analysis.

Voltage calibration: We investigated whether a potential nonlinearity in the ADC-to-voltage calibration could negatively affect the correlation score. Information about the RNO-G voltage calibration is provided in Ref. [163]. To assess the potential impact on the analysis, we utilized UHECR candidate waveforms that we identified in a subset of data collected for analysis development (see Section 8.3) and calculated the template-data correlation before and after applying the voltage calibration. Because the correlation score is normalized (see Equation 8.1), it remains invariant to scaling factors. However, any potential nonlinearity introduced by the voltage calibration that alters the waveform shape would influence the correlation score. For all identified UHECR candidate events,

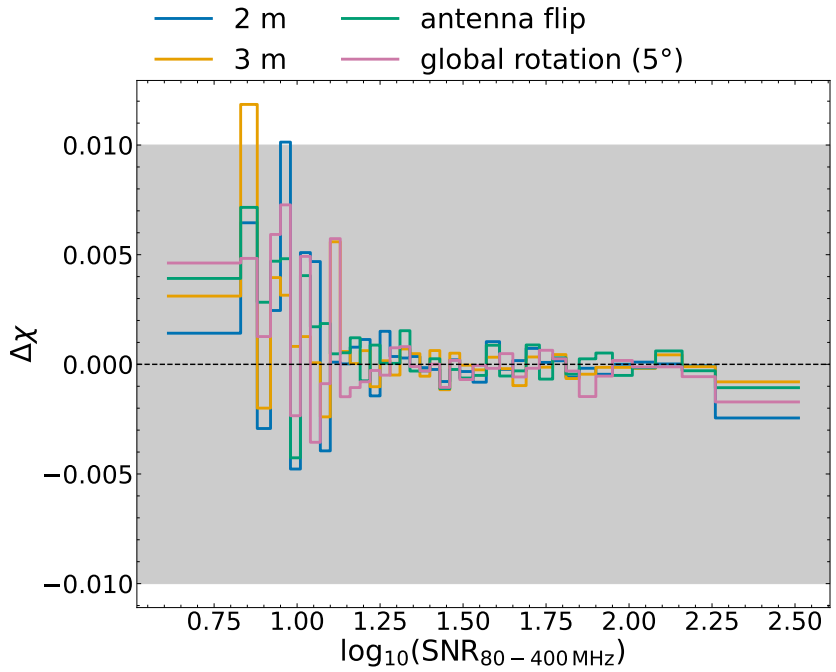


Figure 8.16: Differences in the SNR-dependent correlation score cut, calculated for simulations with modified antenna placement. Each difference is computed relative to the simulation with the standard antenna placement. The grey band indicates the region in which the difference is considered negligible.

the difference in correlation score with and without voltage calibration is less than 10^{-8}. Therefore, the voltage calibration does not significantly influence the correlation score.

Antenna position: We investigated the influence of the antenna depth and orientation on the simulation. In the standard simulation, we place the antenna at a depth of −0.5 m. To assess the impact of varying the antenna depth, we conducted additional simulations with antennas placed at −2 m and −3 m. The differences in the SNR-dependent correlation score cut between the modified and standard depth simulations are shown in Figure 8.16. Across all bins, we find only minor deviations from zero, suggesting that the antenna depth has a negligible effect on the simulation. At higher SNR values, these differences are small and fluctuate randomly around zero. In contrast, at lower SNR values, the fluctuations are more pronounced and exhibit a bias towards positive values. Although the noise realizations remain constant across the simulations, this bias closely resembles the effect previously discussed in the context of statistical noise fluctuations, suggesting a similar origin or mechanism. Overall, for most bins, the difference in correlation scores is below 0.01, and consequently, this systematic effect is not considered further in this analysis.

Moreover, we investigated the impact of the antenna orientation on the simulation. First, we performed a simulation in which all antennas are rotated by 180°. The difference in the correlation score cut calculated relative to the standard simulation is shown in Figure 8.16. Theoretically, a 180° rotation should not affect the simulation outcome due to the symmetry of the LPDA. We only anticipate a sign reversal in the correlation score, which does not alter the result because we consider only the absolute value. At high SNR values, the results show minimal deviations from zero across all bins. These minor

differences are likely attributable to small deviations from symmetry of the simulated VEL, which can slightly modify the simulated waveforms. At lower SNR values, we again observe a bias towards positive values, consistent with previously observed biases. However, the difference in correlation score is less than 0.01 for most bins, and therefore, this effect is not considered for the analysis.

Furthermore, we examined the effect of a global detector rotation on the simulation results. To simulate different azimuth angles, although all cosmic-ray air showers are simulated with the same azimuth angle, we rotate the detector in increments of 10° before converting the incoming electric field into a voltage trace. As a result, we anticipate that a global detector rotation would not alter the simulation outcome. We conducted an additional simulation with the detector rotated by 3°. The resulting difference in the correlation score cut compared to the standard simulation is shown in Figure 8.16. Again, we observe small deviations from zero at high SNR values, which can be attributed to the VEL not being isotropic. When the simulation begins with a slightly different global rotation, different VEL values are sampled at different azimuth angles, leading to minor differences in the simulated waveform. Consistent with previous observations, a bias towards positive values is evident at lower SNR values. Also, for the last simulation, the resulting correlation score cut difference is less than 0.01 in the majority of the bins, indicating that this systematic is negligible.

In summary, variations in the antenna depth, a 180° rotation, or a small global rotation of the detector do not significantly affect the simulation results. Therefore, we will not consider these systematic effects further in this work. While the bias is not examined in detail here, it may be interesting for future work.

Station-to-Station variation: For all simulations, we describe the amplification and digitization response with the same impulse response measurement (see Section 7.2 for more information). Since we use this description to simulate three different stations, each with different hardware components, we examine the impact of using impulse response measurements from different station hardware components on the analysis. In the standard simulation, we use a single impulse response measurement for both template creation and waveform simulation. To assess the impact of station-to-station variations of the impulse response measurements, we generate a second set of templates using an impulse response measurement from a different set of station hardware components. We then correlate the second template set with the UHECR waveforms we simulated using the standard impulse response measurement and calculate the correlation score cut. The difference in the SNR-dependent correlation score cut between this simulation and the standard simulation is shown in Figure 8.17. For most bins, we observe a clear deviation from zero and, additionally, a systematic shift towards negative values, particularly at high SNR values, where the difference reaches -0.015. Such a large difference is considered non-negligible. The negative difference suggests that simulations using different impulse response measurements for template creation and waveform simulation perform better than those using the same measurement for both, which is unexpected. To determine whether this effect persists across all simulated events or only appears in the low-correlation-score events that primarily influence the correlation score cut, we calculate the 90% quantile per SNR bin. This calculation yields a slightly different correlation score cut, which we compare across both simulations. The resulting

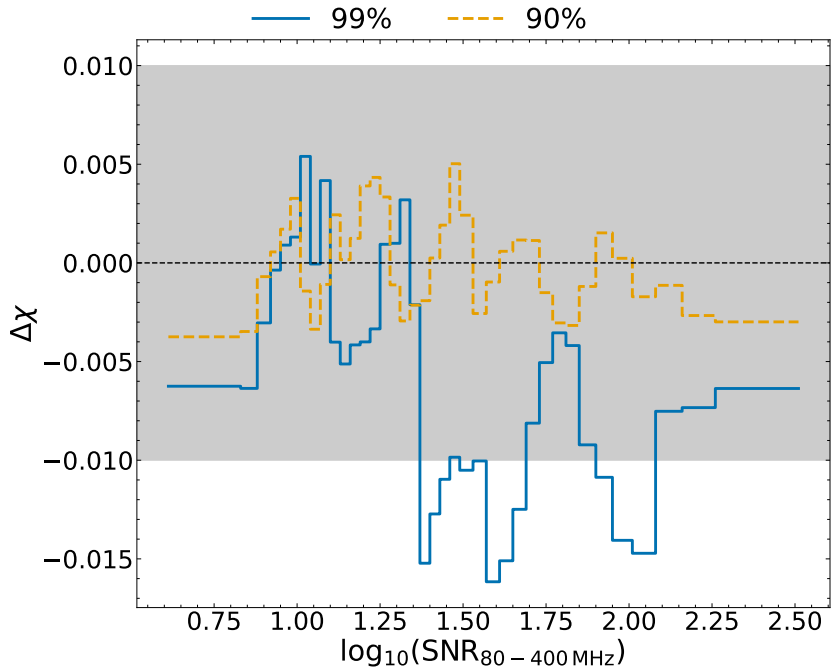


Figure 8.17: The difference in the SNR-dependent correlation score cut, calculated for both the standard simulation and a simulation where the hardware responses for the template and simulated waveforms originate from different stations. The different colors indicate the percentage used to calculate the quantile for each bin in the correlation score cut. The grey band indicates the region in which the difference is considered negligible.

difference, shown in Figure 8.17, indicates that the systematic shift does not persist. As the correlation score cut using the 90% quantile is less sensitive to outliers in the tails of the correlation score distribution, we attribute this effect to outliers in the distribution tails within each bin. Additionally, as discussed in Section 8.4.4, the impulse response measurement exhibits an inherent uncertainty of 0.01 in the correlation score. Therefore, we can likely attribute some of the observed difference to this inherent uncertainty. Although it is not possible to precisely quantify the contribution of the station-to-station variation, it seems not to be significant relative to the inherent uncertainty of the impulse response. Consequently, we use the same impulse response measurement for all templates and simulated waveforms, and we consider the station-to-station variation negligible. Interestingly, for both shown correlation score cut differences, we observe an oscillating behavior as a function of SNR. However, the reason for this pattern remains unknown.

Galactic noise: As discussed in Section 7.4, we consider the diffuse radio emission from the Milky Way Galaxy as a noise source. In the noise simulation, the time of day for which the radio background is estimated is randomly sampled for each simulated waveform. However, the day itself remains the same for all events. To assess whether seasonal variations in the Galactic noise affect the simulations, we repeated the simulations three times using different dates distributed throughout the year. For each simulation, we compare the SNR-dependent correlation score cut with the standard simulation and the resulting differences in correlation score are shown in Figure 8.18. Ideally, if seasonal variations do not influence the correlation score, the difference between the simulations should be zero. The results indicate that the difference remains close to zero and fluctuates around zero for $\log_{10}(\text{SNR})$ values greater than 1.2. In contrast, for

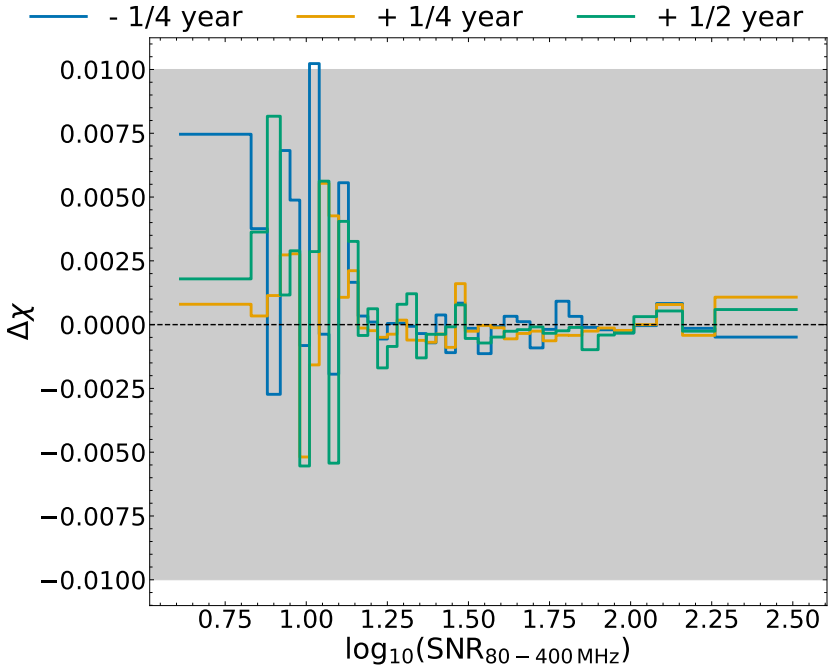


Figure 8.18: Potential effect of seasonal variations in Galactic noise on the SNR-dependent correlation score cut. Shown are the differences in the correlation score cuts calculated for simulations using different dates to calculate the Galactic noise contribution. All differences are calculated and the dates are provided relative to the standard simulation. The grey band indicates the region in which the difference is considered negligible.

$\log_{10}$(SNR) values less than 1.2, the difference exhibits greater fluctuations, with the mean shifting to approximately 0.003 for two of the three dates. For the remaining date (+1/4 year), the fluctuations are smaller, but a bias remains. The noise fluctuations are more pronounced at low SNR values, where noise generally has the greatest impact on the correlation score. Additionally, seasonal variations may introduce a systematic shift if the Galactic noise becomes weaker or stronger, which makes it plausible that the systematic shift originates from seasonal variations. However, this systematic shift occurs in the same region as the effect observed for the statistical noise fluctuations, and similar effects could produce the same systematic behavior here. Consequently, we cannot confidently attribute this behavior solely to Galactic noise. Nevertheless, the difference in the correlation score remains below 0.01 for most bins. Therefore, we consider a possible systematic effect from seasonal variations in Galactic noise negligible for this analysis.

Noise temperature: As described in Section 7.4, we simulate thermal noise by assuming an effective temperature that combines the blackbody radiation temperature of the surrounding ice and the system temperature. We estimated the effective temperature to be approximately 206 K, which we use in the standard simulations. However, this value is an estimate and subject to uncertainty. For example, to determine the environmental blackbody radiation temperature, we averaged over multiple frequency bins. To assess the impact of such uncertainties on the simulations, we assume an error margin of 10 K in the total effective temperature. Then, we conduct two additional simulations: one with the noise temperature increased by 10 K and another with it decreased by 10 K. Figure 8.19 shows the comparison of the SNR-dependent correlation score cuts between

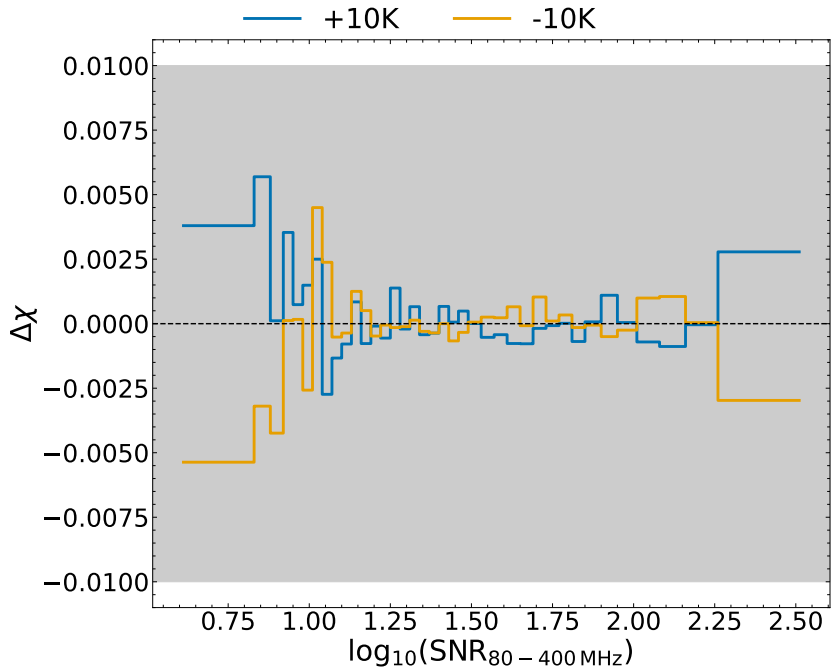


Figure 8.19: Differences in the SNR-dependent correlation score cut, calculated between the standard simulation and two simulations with a varying noise temperatures. The grey band indicates the region in which the difference is considered negligible.

these additional simulations and the standard simulation. The results indicate that the difference fluctuates around zero at higher SNR, while we observe a bias toward positive or negative values at lower SNR. An increase in the noise temperature increases the noise RMS, resulting in a lower SNR value for all simulated waveforms. Additionally, the correlation value decreases as the noise RMS increases. Consequently, simulated events shift toward lower correlation scores and lower SNR values at higher noise temperatures, effectively mimicking a change in the trigger threshold. Therefore, the noise temperature primarily influences the lower SNR cutoff in the simulations and is most relevant for estimating the expected number of UHECR events. Notably, the observed difference coincides with the region where we observed a bias for the statistical noise fluctuations. As a result, it is difficult to determine whether the observed differences are solely due to changes in the noise temperature or are also affected by a similar bias. Overall, the differences remain below 0.01 for all bins, and therefore, we consider this systematic effect negligible for this analysis.

Additional checks: We verified that the reflection and refraction at the ice-air boundary are accurately described in the simulations. Furthermore, birefringence is expected to be negligible due to the short propagation distances $\mathcal{O}(1\,\mathrm{m})$ of the signal in ice. We also examined the impact of varying the bin size of the SNR-dependent correlation score cut, and found it to be negligible.

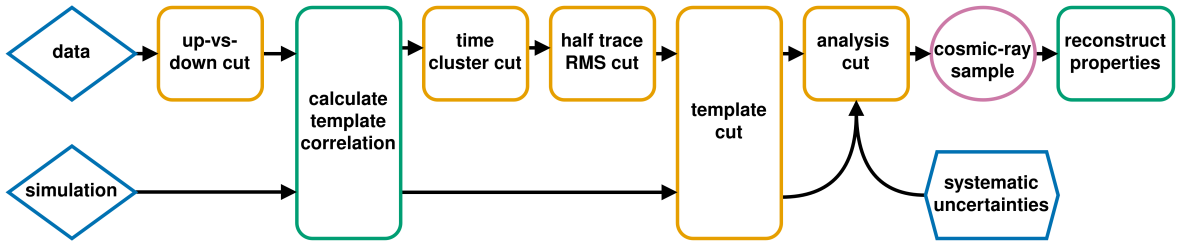


Figure 8.20: A schematic overview of the analysis and simulation pipeline to identify and reconstruct cosmic rays. The large blocks indicate that the same procedure is applied to both the data and simulation pipeline. Taken from Ref. [185].

Table 8.5: Number of events remaining after each step of the analysis for the three stations 13, 23, and 24 in the main dataset, separated by the year 2022 (top) and 2023 (bottom). Taken from Ref. [185].

		13	23	24
2022	before all cuts	306,869	680,902	217,516
	after up-vs-down cut	115,134	547,083	63,377
	after time cluster cut	100,983	103,329	38,606
	after half trace RMS cut	97,374	99,193	37,375
	after template cut	82,366	86,342	32,489
2023	before all cuts	726,444	4,476,619	653,415
	after up-vs-down cut	201,867	1,263,333	375,271
	after time cluster cut	163,287	769,029	65,081
	after half trace RMS cut	151,590	752,959	61,263
	after template cut	129,768	652,425	52,273

8.5 Analysis of main dataset

The datasets of the three stations (13, 23, and 24) are analyzed independently while following the same analysis and simulation procedures as shown in Figure 8.20. In the following, we describe the analysis and simulation procedure step by step.

For each station, the input for the analysis is a dataset as defined in Section 8.3. Table 8.5 shows the number of triggered events per station before any cuts are applied and for those remaining after the cuts that are discussed here.

In the first analysis step, the *up-vs-down cut* is applied, removing all events from the dataset in which the amplitude in the downward-facing antennas is higher than in the upward-facing. This cut ensures that the triggered signal comes from above the ice. Such a signal arrives at the downward-facing antenna in the insensitive back-lobe, resulting in a smaller signal amplitude compared to the upward-facing antennas, where the signal arrives in the sensitive front-lobe. Primarily, thermal noise events are rejected by this cut. The cut is not applied to the simulations used in the analysis, but we confirmed that applying it would not change the simulation results. Its main purpose is to reduce the number of events in the dataset, thereby speeding up the following steps. The number of events surviving the up-vs-down cut are shown in Table 8.5.

In the next step, the data-template correlation is calculated (*template correlation*). For each event, the waveforms from all upward-facing channels are correlated with all

four templates (two signal and two background templates) using Eq. 8.1, resulting in 12 correlation scores per event. Only the maximum of these 12 correlation scores is considered in the next steps of the analysis. Two additional important event quantities are then calculated from the channel waveform with the highest correlation score. The first is the half trace root-mean-square (htRMS) value, defined as the RMS of the first half of the waveform, which is in ADC counts. The second value is the SNR, defined as the ratio of the maximal amplitude and the htRMS value. Prior to these calculations, we apply a bandpass filter to both templates and data/simulation waveforms. Since a study indicated that the bandpass filter slightly increases the data-template correlation score. We performed the study on a small subset of the data selected for analysis development (see Section 8.3 for a description of the data subset). As the filter, we selected a rectangular filter with a passband of $80 - 400\,\mathrm{MHz}$, roughly corresponding to the frequency range covered by the template with the highest frequency content. The bandpass filter also ensures the removal of non-physical signals outside the frequency band of RNO-G.

Subsequently, a series of background rejection cuts are applied to the data, which we elaborate on in the following. First, we apply the time cluster cut (see Section 8.5.1), followed by the half trace RMS cut (see Section 8.5.2), and the template cut (see Section 8.5.3). After these cuts, we select a sample of UHECR candidate events from the background-reduced dataset (see Section 8.5.4) and reconstruct their properties (see Section 8.6).

8.5.1 Time cluster cut

The *time cluster cut* is designed to remove backgrounds that cluster in time and have a high correlation score. This cut is effective because UHECR events are randomly distributed in time rather than clustered. Possible backgrounds are, however, temporally clustered such as signals from airplanes [181] and background signals occurring during high wind periods [173]. While these backgrounds are likely to be identifiable by their signal shape (e.g. [17]), without a model for these backgrounds, we choose the conservative approach to use a time cluster cut. To minimize the number of random coincidences from low correlation score events that are not signal-like, only events that have a correlation score above a certain value (minimal correlation score) are considered for this cut. An event is retained if no other event with a correlation score larger than the minimal correlation score occurs within the time period Δt. If two or more events are within Δt of each other, the whole time period is rejected from the dataset, including those events that did not contribute to the cut due to their small correlation score. A schematic of the time cluster cut is shown in Figure 8.21.

We determined the value for the time period Δt by optimizing it on a small subset of the data that is used for analysis development (see Section 8.3 for a description of the data subset). For all three stations, we use a value of $\Delta t = 60\,\mathrm{s}$. We derive the minimal correlation score from the correlation score distribution of thermal noise events which are taken following a forced trigger. Figure 8.22 shows these distributions for all stations and both data-taking periods (2022 and 2023). The figure shows, for both years, the distribution at the beginning and the end of the data-taking period. All four distributions are stable over time, which makes it possible to choose a constant cut value.

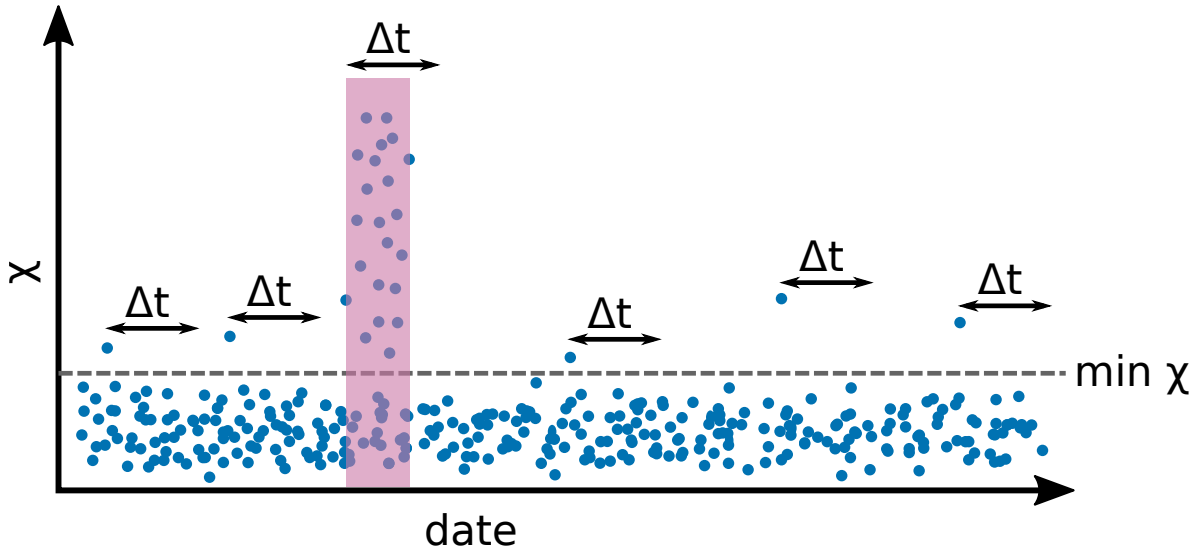


Figure 8.21: Illustration of the time cluster cut. Further details are provided in the main text.

For each station, the value is chosen such that only a few thermal noise events would contribute to the cut and thus reduce the number of possible random coincidences. To further support the decision, a simulation was performed to estimate the expected influence of random coincidence due to thermal noise as a function of the minimal correlation cut value. In the simulation, we generated datasets containing realistic numbers of thermal noise and UHECR events. We assigned both event types a random timestamp and a correlation score exceeding the minimal correlation score. Then, we applied the time cluster cut to these datasets to calculate the expected lifetime loss and UHECR dead time due to random coincidences. The UHECR dead time arises because each thermal noise event with a correlation score higher than the minimal correlation score effectively blocks a time window of $2\Delta t$, during which no UHECR event can occur without being rejected by the time cluster cut. This effective UHECR dead time needs to be accounted for when estimating the expected number of UHECR events. Even when the minimal correlation score is set sufficiently high to exclude thermal noise, random coincidences between different UHECR events may still occur. The resulting loss of UHECR events due to such coincidences is also incorporated into the calculation of the UHECR dead time. Notably, we determine the UHECR dead time exclusively from this simulation and include it in the estimation of the expected number of UHECR events. In contrast, the simulation-based estimate of the lifetime loss due to random coincidences serves only to guide the selection of the minimal correlation score. The lifetime loss that we use to estimate the expected number of UHECR events is derived from actual data and is primarily influenced by non-thermal backgrounds.

For all three stations and both years, we set the minimal correlation score to be 0.55. As a result, less than 0.05% of the events per station and year shown in the distributions in Figure 8.22 would contribute to the time cluster cut. The number of events surviving the time cluster cut can be found in Table 8.5, and the corresponding lifetime losses for both years are provided in Table 8.6. The UHECR dead times values are listed in Table 8.6 as well.

To calculate the UHECR dead time, we estimate the number of thermal noise events contributing to random coincidences. We do this by determining the percentage of

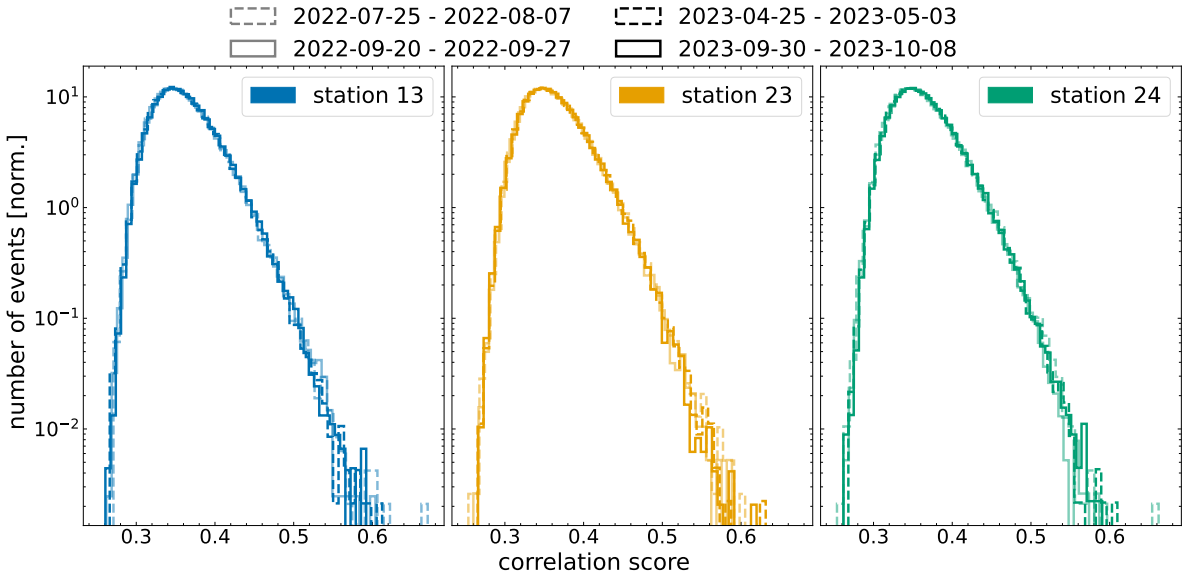


Figure 8.22: Distributions of the correlation scores of forced trigger data for three stations during different time periods in 2022 and 2023. Adapted from Ref. [185].

thermal noise events above the minimal correlation score from a forced-trigger dataset and combining this with the average trigger rate. The simulation further depends on the expected number of UHECR events, for which we obtain an estimate from station-specific simulations. We perform both best- and worst-case trigger scenario simulations (see Section 7.3 for more information) for each station, and the corresponding UHECR dead times are reported in Table 8.6. For each station and trigger scenario, we simulate 30000 days to estimate the UHECR dead time, and repeat these simulations 100 times. As the final UHECR dead time value, we report the mean across the 100 simulations plus 3 times the standard deviation, providing a conservative estimate. Although the mean UHECR dead time is larger for the best case than in the worst case, the standard deviations are smaller, resulting in slightly lower quoted values in the best case than in the worst case. We expect higher mean values in the best-case scenario, as the probability of cosmic-ray-to-cosmic-ray coincidences increases with the number of UHECR events.

In this analysis, we exclude time periods in which two or more events occur in close temporal proximity. Alternatively, we could increase the minimum number of required events to reduce the likelihood of random coincidences and, in turn, the UHECR dead time. However, it may also permit more background events to pass the cut. Future studies could investigate whether adjusting this parameter improves the ratio of UHECR to background event survival probability.

This cut is not applied to the standard UHECR simulations (as indicated in Figure 8.20), as they do not include time stamps or time-clustered backgrounds.

8.5.2 Half trace RMS cut

As the second background rejection cut we apply the *htRMS cut*. The signal is, by construction, in the second half of the recorded waveform given the trigger settings (see Section 4.3). For clean UHECR pulses, we only expect thermal noise on the first

Table 8.6: The lifetime loss and UHECR dead time resulting from the time cluster cut. For the UHECR dead time, we quote the values for both trigger threshold scenarios (worst, best) (see Section 7.3 for more information about the trigger scenarios). Only for station 23, we sometimes only present the UHECR dead time for the worst-case scenario, because no best-case trigger threshold scenario is available.

	station		lifetime loss [%]	UHECR dead time [%]
2022	13		0.3	(0.6, 0.4)
	23		7.6	(1.1, -)
	24		2.2	(0.4, 0.3)
2023	13		1.1	(0.6, 0.3)
	23	cfg1	4.8	(1.8, -)
		cfg2	0.02	(0.4, -)
		cfg3	4.5	(1.3, 0.8)
	24		4.5	(0.5, 0.4)

half of waveforms. Events with a signal other than thermal noise on the first half of the waveform - quantified by the htRMS value - are rejected as background. Possible background events rejected by the htRMS cut are solar flare events [179]. However, this cut is fundamentally agnostic to the origin of the background. The htRMS threshold, above which an event is interpreted as background, is estimated from htRMS thermal noise distributions. Figure 8.23 shows, for each station and year, two thermal noise htRMS distributions: one at the beginning and one at the end of the data-taking period.

For each station, we can see in Figure 8.23 that the htRMS distribution changes during the data-taking period. The changes in 2022 are relatively small compared to the ones in 2023, which likely can be attributed to the longer data-taking period in 2023. Furthermore, for some stations, some distributions have a smaller second peak at higher htRMS values, which can be attributed to narrow band events (e.g., hand-held radio communication or data transmission from a weather balloon) that make it into the forced trigger dataset. The instability of the htRMS distribution makes it necessary to define a value for the cut that is time-dependent.

We calculate the time-dependent threshold by splitting the data into bins of 5 days. For each time bin, we fit a log-normal distribution to the main peak, defined as the region where the distribution remains above 10% of its maximum amplitude. An example of such a distribution for a 5-day bin is shown in Figure 8.24. We then estimate the cut value as the 0.99999 quantile of the fitted distribution. With this approach, it is possible to set a cut value that is not artificially increased by the waveforms with excess noise, such as narrow band events. Occasionally, the fitting procedure yields an unrealistically large cut value due to a strong secondary distribution caused by narrow-band events. Figure 8.24 shows the calculated cut values for station 24, with each value corresponding to a 5-day period. Two cut values clearly deviate from the general trend, which appears to follow a sine function. To address these outlier data points, we employ an iterative approach. Initially, we fit a sine function to all data points and calculate the distance between the fit and each data point. We classify data points with a distance greater than 0.5 ADU as outliers. In subsequent iterations, we fit the sine function again, excluding the identified outliers, and we repeat the process until no further outliers are detected. For outliers, we correct the cut value by replacing it with an estimate from the final sine fit. The fitted sine curves and the corrected outlier data points are also shown in Figure 8.24.

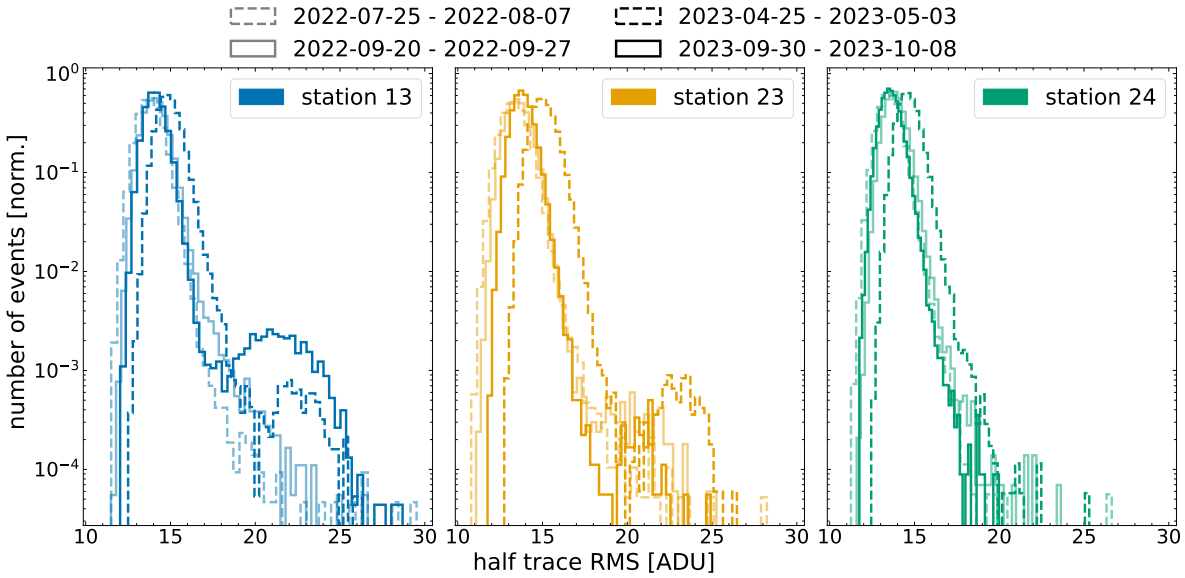


Figure 8.23: The half trace root-mean-square (htRMS) distributions of forced trigger data for three stations and different time periods in 2022 and 2023. The htRMS values are given in ADC counts (ADU). Adapted from Ref. [185].

We did not further investigate the cause of the unstable htRMS values, but it may be related to external factors, such as the outside temperature. Future work should further investigate this unstable behavior. The number of events that survive the htRMS cut can be found in Table 8.5.

This cut is not applied to the standard UHECR simulations (as indicated in Figure 8.20), as these simulations contain only thermal noise in the first half of the waveform and, by definition, the cut has no effect.

8.5.3 Template cut

The last background rejection cut is the *template cut*. A significant number of impulsive backgrounds with a high correlation score have a much lower frequency content compared to what is expected from UHECR events. As a result, we introduced additional background-tagging templates with a low frequency content (σ = 1.5625 ns, σ = 2.1875 ns)[7], shown in Figure 8.1. Events that have the highest correlation score with either of the two background templates are rejected from the dataset. Simulations show that for all stations, less than 2.2% of the expected UHECR events are rejected, which makes this cut very efficient. When we developed the template set, we considered the background template with width σ = 1.5625 ns as a signal template (see Section 8.2). However, retaining it as a signal template results in only a marginal increase of $\mathcal{O}(2\%)$ in the number of simulated UHECR events that pass the cut, while allowing more background events to pass. Consequently, we reclassified it as a background template. For the second background template, we selected a waveform with a lower frequency content than all signal templates. To determine the optimal template width, we performed a small scan correlating templates with varying widths with a subset of background

[7] The Gaussian width σ used to create the templates is defined in waveform samples and subsequently converted to nanoseconds, which explains the very precise values.

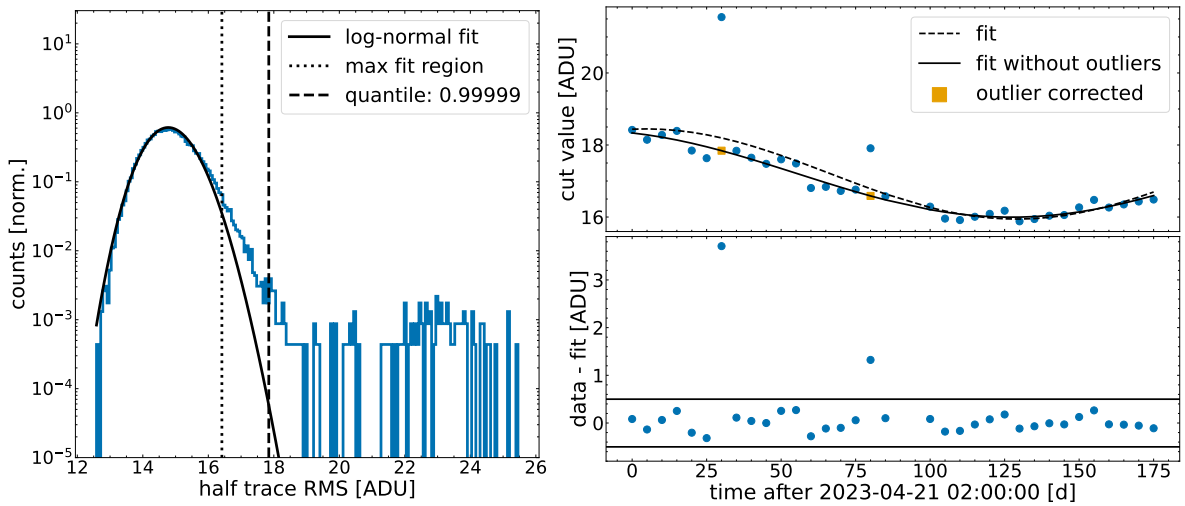


Figure 8.24: Left: Distribution of htRMS values for an example 5-day bin at station 23 in 2023, along with the fitted log-normal distribution. Right: The derived cut values for the htRMS cut for station 23 in 2023 and an illustration of the procedure for correcting outlier values.

and UHECR candidate events from the analysis development dataset. We selected a width of $\sigma = 2.1875\,\text{ns}$ for the second background template, as templates with larger widths allowed more background events to pass. The number of events surviving the last background rejection cut are shown in Table 8.5. This cut is applied to the standard UHECR simulations (as indicated in Figure 8.20), as it affects the shape of the SNR-dependent correlation score cut and the expected number of UHECR events, both of which we derive from the simulations.

8.5.4 Event selection

After applying all background rejection cuts, an analysis cut is applied to select the UHECR sample. We derive this analysis cut from UHECR simulations (more details about the simulation setup are provided in Chapter 7). After generating the simulation, the correlation score, htRMS and SNR are calculated for all triggered events and each event is weighted according to the expected UHECR rate. Subsequently, the background rejection cuts are applied. Since this is pure UHECR simulation, only the template cut has an effect, as the time-cluster and htRMS cut do not remove any simulated events. Furthermore, it should be noted that also the up-vs-down cut is not applied, as it does not affect the pure UHECR simulation. Figure 8.25 shows the correlation score as a function of SNR for a UHECR simulation. For the simulation, we used an example trigger threshold of 25 mV, which corresponds to 9σ.[8] Finally, the simulation is used to calculate the SNR-dependent cut on the correlation score to select the sample of UHECR candidate events. Figure 8.25 shows three lines that are calculated from the shown simulation. The interpretation of each line is that 90%, 95%, and 99% of all simulated UHECR events lie above the line in correlation score per SNR-bin. To determine these lines, the simulation is divided into multiple SNR bins, each containing roughly $10,000$ events to ensure sufficient statistics. For each SNR bin, the correlation cut value is chosen such that the desired fraction of simulated UHECR survives the cut. To ensure that all

[8] In this context, σ corresponds to the RMS of the simulated noise in the trigger frequency band.

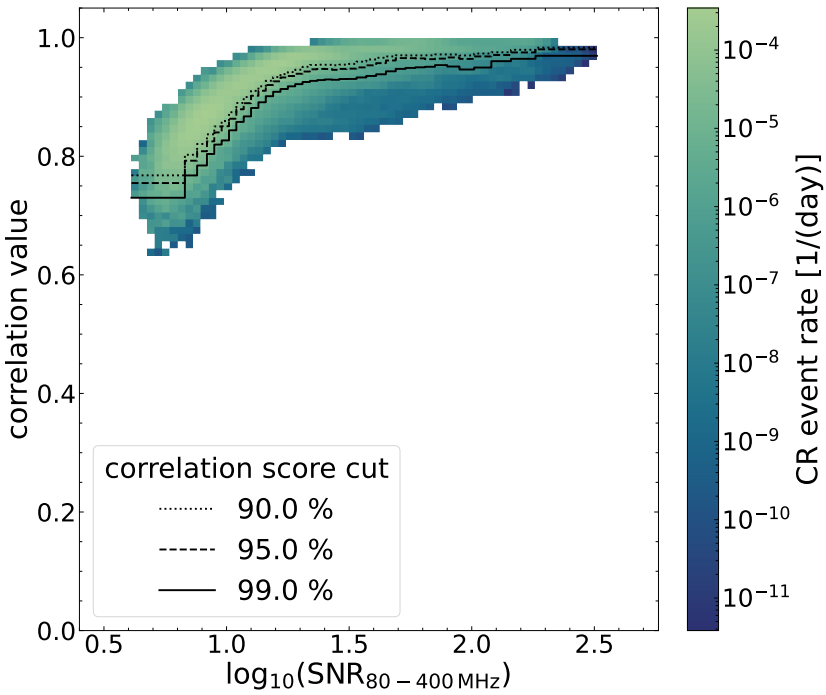


Figure 8.25: The correlation score as function of SNR obtained from an example cosmic-ray (CR) simulation. The colorbar indicates logarithmically the expected event rate per day. Also, shown are three lines that illustrate the correlation score per SNR-bin that retains the 90%, 95% and 99% of the simulated CR events, respectively. The simulation is performed with an example threshold of 25 mV. Taken from Ref. [185].

possible UHECR events are selected, the SNR-dependent correlation score is calculated using the best-case instead of the worst-case simulation[9]. As discussed in Section 8.4, four systematic uncertainties significantly influence the correlation score (larger than 0.01 in correlation score). These systematic uncertainties are added quadratically for each bin. Except for the uncertainty from the hardware response measurement, all systematic uncertainties consistently lower the correlation score per bin. The uncertainty from the hardware response can also increase the correlation score per bin, leading to an asymmetric systematic uncertainty region.

The determined SNR-dependent correlation score cut is applied as the final *analysis cut* to the data. The dataset for all three stations, after applying all background rejection cuts, is shown for 2022 in Figure 8.26 and for 2023 in Figure 8.27. Also shown in Figure 8.26 and 8.27, is the calculated SNR-dependent cut on the correlation score, chosen to retain 99% of the simulated UHECR events (including the systematic uncertainties) and the selected sample of UHECR candidate events.

From the simulation, we obtain a number for the expected rate of UHECR events per day for the best- and worst-case scenario, based of the uncertainties stemming from the realistic simulation threshold. Multiplying this rate by the detector lifetime yields a range for the expected number of UHECR events from the dataset. We also account for the lifetime loss due to the time cluster cut, the UHECR dead time due to random coincidences from the time cluster cut, and the rate reduction due to the template and analysis cut. The expected and found number of UHECR candidate events per station and year are summarized in Table 8.7. Given the conservative best- and worst-case

[9] For station 23, we use a simulation with a 25 mV threshold, as no best-case simulation is available (see Section 7.3).

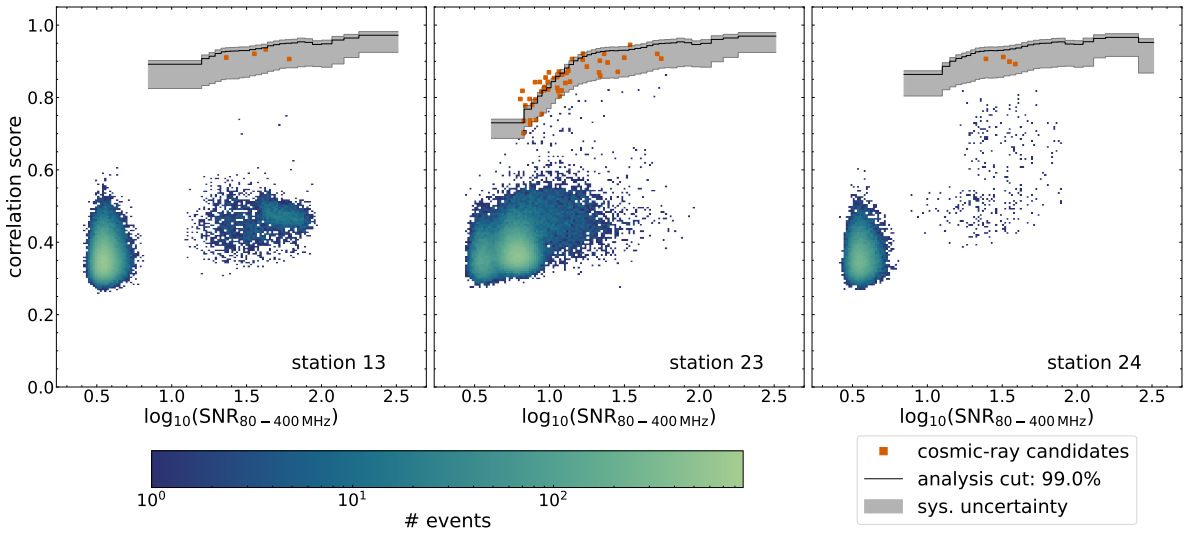


Figure 8.26: The correlation score as function of SNR for the complete 2022 dataset per station after applying all cuts. The selected UHECR candidate events are highlighted in orange. Taken from Ref. [185].

Table 8.7: Found and expected number of UHECR candidate events for all stations and datasets. Taken from Ref. [185].

	station		expected	found
2022	13		1 - 5	4
	23		> 3	47
	24		2 - 8	4
2023	13		4 - 13	8
	23	config 1	> 4	8
		config 2	> 0	2
		config 3	0 - 2	1
	24		4 - 19	11

scenario estimate, we expect the number of all found UHECR candidate events to lie within the expected range for all stations and years. For station 23, only a lower limit is available, as no best-case scenario could be calculated (see Section 7.3). In agreement with expectations, the number of identified UHECR candidate events lies within the predicted range for all stations and years.

Figures 8.26 and 8.27 show that the majority of the UHECR candidate events would not satisfy the correlation cut without accounting for systematic uncertainties, underscoring the importance of their proper treatment. The largest number of UHECR candidate events is found for Station 23 in the year 2022, which is consistent with the station having a lower trigger (and thus cosmic-ray energy) threshold for this year compared to the other stations and the steep increase in the cosmic-ray flux towards lower energies. Due to the lower energy threshold and the cosmic-ray spectrum, we expected for this station to find the majority of the UHECR candidate events at lower SNR (lower energy), which is in agreement with the result shown in Figure 8.26. As discussed in Section 8.3, the trigger threshold for station 23 was increased for most of 2023 relative to the previous year. Consequently, we expect fewer UHECR events and those detected should have a higher SNR value, which is reflected in the data.

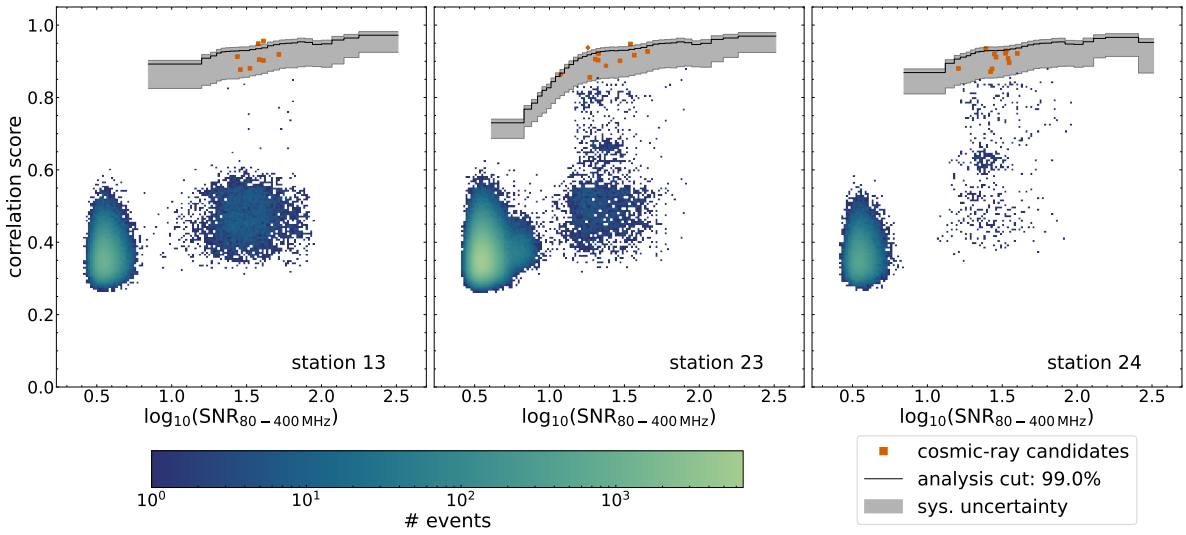


Figure 8.27: Same as Figure 8.26 for 2023. For station 23, only the most conservative SNR-dependent correlation score cut is shown and different markers indicate the candidate events from different configurations. Taken from Ref. [185].

From the left panels of Figures 8.26 and 8.27, it can be shown that two populations of background events – one at lower and the other at higher SNR – survive the background rejection cuts for station 13. Because of the employed noise-riding trigger threshold, only the population at higher SNR ($\log_{10}(\text{SNR}) \gtrsim 1$) is expected, with the lowest SNR value of this population proportional to the station's trigger threshold. The unexpected distribution at low SNR ($\log_{10}(\text{SNR}) \lesssim 1$) is consistent with the SNR and correlation score of thermal noise and its existence can be attributed to the trigger hardware. For the version of the digitizer boards used to collect the presented data, the incoming signal is split between the trigger and signal path for the shallow channels, with the majority of the power directed to the signal path. This leads to the trigger path being dominated by board noise, causing the system to trigger on this noise, which is not visible in the recorded waveforms (i.e. signal path). As a result, the board noise produces false positive triggers that make up the population of thermal noise events (more details are provided in Section 5.3.1). These events pass all background rejection cuts, as no cut is specifically designed to remove thermal noise events. Also, the efficiency for cosmic rays at low SNR values is poor.

The same behavior as for station 13 can be seen for station 24 (right panel, Figures 8.26 and 8.27), while it is less clear in station 23 (center panel, Figures 8.26 and 8.27), especially in 2023 due to its multiple configurations (see Section 8.3). Furthermore, comparing stations 13 and 23 for the year 2022, which have different thresholds, the population at higher SNR values shifts to smaller SNR values for the station with the lower threshold (station 23), which shows the dependence on the trigger threshold for this population. For a small enough trigger threshold, the population at high SNR values and the thermal noise population overlap, as visible in the station 23 data.

The above-described issue with the digitizer board is known, and the next digitizer version (V3) features improved signal splitting to mitigate it (more information is provided in Section 5.3.3). For all stations except 24, the digitizer board was upgraded in 2024. Studies of the trigger efficiency of the new digitizer board in a laboratory set-up have shown that we expect an increased rate of 5^{+5}_{-2} UHECR events per day and per seven

stations [172]. As a result, with the new design, in less than 3000 h of lifetime of a single station, the number of UHECR events is comparable to those found in this work, which uses a total of 15, 182 h of single station lifetime. Future analyses will benefit from the improved trigger performance and the increased UHECR statistics.

To illustrate the events selected, Figure 8.28 and 8.29 display for each year the waveforms of the UHECR candidates with the highest and lowest correlation score. For each candidate event, the waveforms of the three upward-facing channels are overlaid with the best fitting template to illustrate the data-to-template match. Additionally, the corresponding correlation score is shown in the figure. When interpreting these correlation scores, it is important to note that we consider only the maximum correlation score across all three upward-facing channels. In some of these waveforms, we observe a small "ringing" following the main pulse, a phenomenon present in multiple UHECR candidate events. One hypothesis is that this effect originates from reflections caused by impedance mismatches between components, such as the coax-cable-to-Surface-board interface. A preliminary study suggests that reflections originating at the input of either the Surface or Radiant board can produce the necessary time delay if they are again reflected at the antenna output. In this case, the time delay would make the reflected signal appear after the main pulse, yet still be visible in the recorded waveform. However, in this preliminary study, we only considered the signal travel time. Future work should perform a more comprehensive analysis, incorporating the amplitudes of reflected signals to determine if they are sufficient to account for the observed "ringing". Proper cascading of all S-parameters for the different components may provide further insight. Although the impulse response measurement contains all amplification and digitization components, the antenna is not included. Therefore, reflections dependent on the antenna connection are included neither in the response nor in the template waveform.

The signal properties - the zenith angle, azimuth angle, polarization, and energy - for all UHECR candidates are reconstructed and will be discussed in Section 8.6.

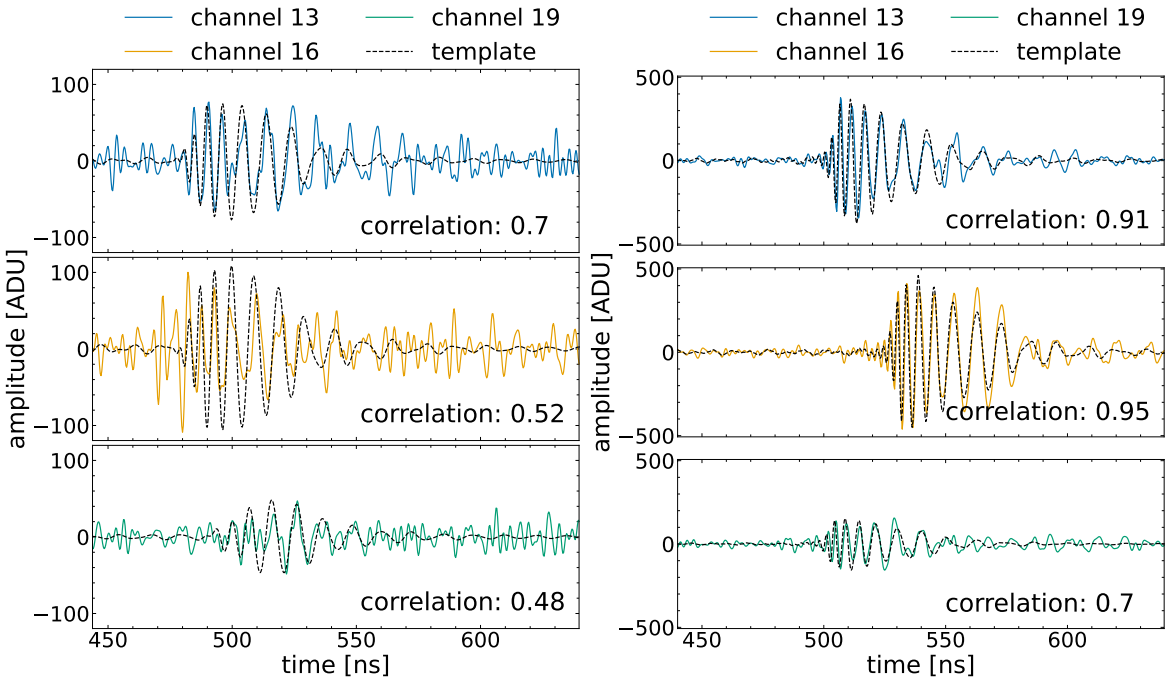


Figure 8.28: Waveforms of the upward-facing LPDAs for the UHECR candidate event with the lowest (left) and highest (right) correlation score across all candidate events in the 2022 data. Adapted from Ref. [185].

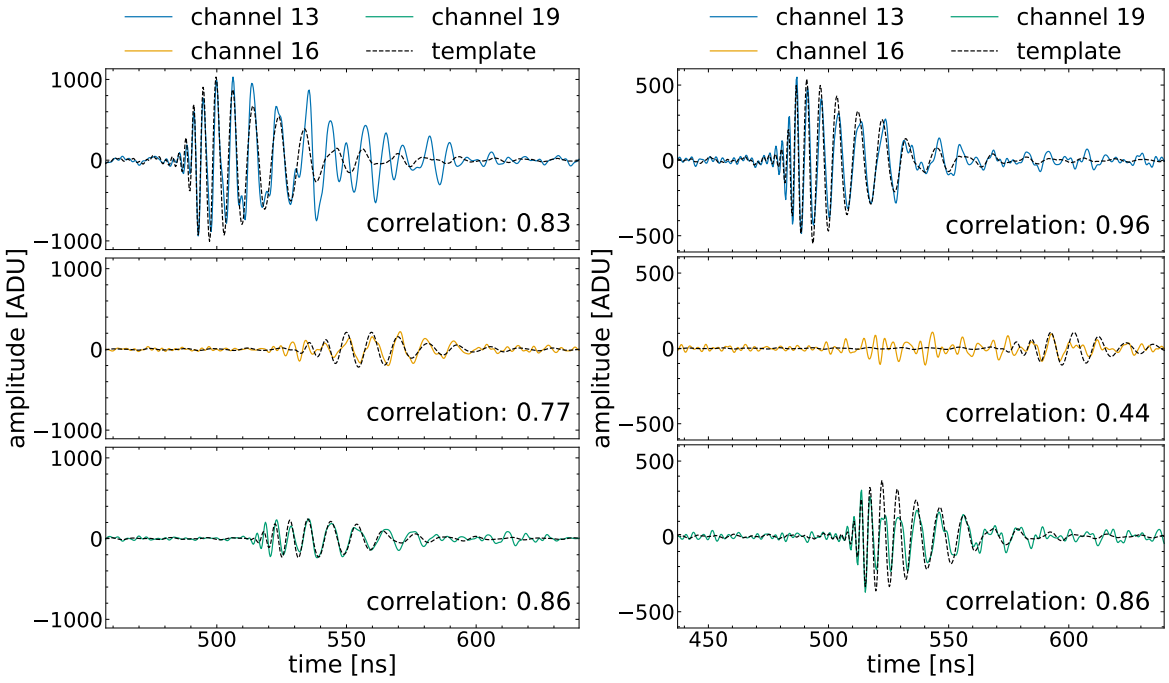


Figure 8.29: Waveforms of the upward-facing LPDAs for the UHECR candidate event with the lowest (left) and highest (right) correlation score across all candidate events in the 2023 data. Adapted from Ref. [185].

8.6 Characterization of observed events

For each UHECR candidate event we reconstruct the zenith angle, azimuth angle, polarization and energy. The reconstruction is done using a forward folding approach, meaning a model for the electric field is fitted to the data waveforms by folding it with the hardware response, which makes it possible to estimate the zenith, azimuth and polarization together. A detailed description of this method can be found in Ref. [121]. In a second step, the energy is reconstructed using the method described in Ref. [219]. The resulting reconstructed parameters for all UHECR candidates can be found in Figure 8.30, overlaid with a simulation that is normalized to the number of found UHECR candidate events. As discussed before, the stations all have a different trigger threshold. So we have chosen a trigger threshold for the shown simulation that is smaller than the thresholds of all stations to show the full possible parameter space. In principle, we can simulate each station with realistic trigger thresholds. However, for the reconstruction, we did not do this due to time constraints. In Figure A.6, the reconstructed values are shown again, separated by station.

For the zenith angle, azimuth angle and energy, the reconstructed values are shown, while for the polarization, the difference ($\Delta_{\text{polarization}}$) to the expected polarization (assuming a purely geomagnetic emission) is shown. The zenith and azimuth angle distributions agree well with the simulations. Due to the large distance of 1.25 km between two stations, only a very small number of events are expected to trigger more than one station. The number of coincidences increases with the area covered by the air-shower radio footprint, which grows with increasing zenith angle. Since some UHECR candidate events are reconstructed at a large zenith angle, we searched for potential coincidence events but found none. This result is consistent with expectations from simulations, given the small number of candidate events per station. However, we do not observe a perfect agreement for the energy. Additionally, some events are reconstructed with a higher energy than expected. Due to the energy dependence of the cosmic-ray flux and the lower threshold of the simulation, more lower-energy events are included in the simulation, which - after normalizing the histogram - effectively suppresses the number of simulated higher energy events. As a result, the simulated energy distribution appears to be shifted toward slightly lower energies compared to the distribution of UHECR candidates. Furthermore, the reconstructed energies still carry sizable uncertainties, since an absolute gain calibration of the detector has not yet been performed. Much larger experiments like the Pierre Auger Observatory have measured the flux of UHECR cutting off beyond energies of 10^{20} eV [53]. Therefore, we expect that the events exceeding 10^{20} eV are either not true cosmic rays or simply mis-reconstructed, as no quality criteria were applied to the reconstructed dataset. Indeed, two of the highest energy events and the event below the trigger threshold are found to likely be background events (see Section 8.7).

Also the polarization difference shows a mismatch, in the form of a bias towards positive values for the UHECR candidate events. This observed bias of $6 - 7°$, is observed exclusively in the events from station 23, which dominate the total event sample. In contrast, the polarization differences for stations 13 and 24, although based on fewer events, agree well with the simulation. The polarization reconstruction was cross-checked by applying the method presented in Ref. [220], which gives overall consistent results,

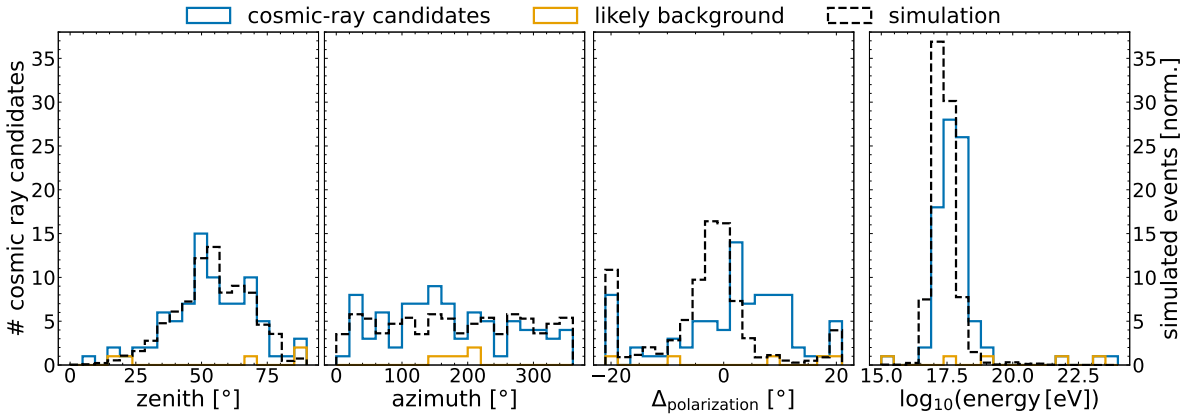


Figure 8.30: Reconstructed quantities for the cosmic-ray candidates (solid blue line), corresponding simulations (dashed line), which are normalized to the number of cosmic-ray candidate events, and likely backgrounds (solid orange line). The likely backgrounds are the five events identified in Section 8.7. Shown are from left to right the zenith angle, azimuth angle, difference to expected polarization angles assuming a purely geomagnetic emission (with the first and last bin being overflow bins), and energy. It should be noted that the shown distributions are consistent in all cases after including systematic uncertainties, which are not shown. Additionally, the simulations use on average a lower trigger threshold than realized in real data. See text for details. Taken from Ref. [185].

Table 8.8: The possible bias in reconstructed polarization and systematic uncertainty for the polarization difference (reconstructed - expected) resulting from systematic uncertainties on the detector description. Taken from Ref. [185].

Systematic uncertainty	σ	Bias [°]	$\Delta\sigma_{68\%}$ [°]
Amplifier gain	5%	0.0	0.3
Antenna position	0.5 m	1.3	2.5
Antenna azimuth	5°	3.3	1.9
Antenna zenith	5°	4.9	6.1

and the same bias is present. To estimate the impact of systematic uncertainties from the detector description and to assess whether they can explain the observed bias, we re-simulated 500 randomly selected simulated events with detector descriptions that vary slightly from the nominal detector description. These variations include geometrical shifts or rotations of the LPDA antennas that have not been accounted for during the installation of the first RNO-G stations. We then subsequently reconstruct these simulated events with the nominal (non-shifted) detector description, and determine the resulting increase in the 68%-containment intervals and bias.

We investigated four sources of systematic uncertainty: an overall scaling of the amplifier gain, a shift of the antenna position in the x-y plane, and a rotation/tilt of the antenna orientation (either in azimuth or in zenith). For each source of systematic uncertainty, we have simulated 25 different detectors by drawing random shifts/rotations from a normal distribution, which were then used to re-simulate the selected simulated events. The shifts/rotations were applied individually to each antenna rather than globally to the entire station. The resulting biases and systematic uncertainties are summarized in Table 8.8.

The results in Table 8.8 show that the observed bias can be explained by systematic uncertainties in the detector description, with the antenna azimuth and zenith angle being the biggest sources of uncertainty. Furthermore, these results emphasize that a more careful antenna position and orientation calibration for the LPDAs is needed, as it

lowers the systematic uncertainties of the reconstruction algorithms significantly.

As a final comparison between simulation and UHECR candidate events, we calculate the 68% quantile of the polarization difference. For the simulation, the value strongly depends on the trigger threshold, as lower SNR events are harder to reconstruct. If we consider a low trigger threshold for the simulation of 4σ[10], we obtain a value of 7.3°, while for a higher trigger threshold (14σ) we obtain a value of 2.9° [121]. To account for the bias on the polarization difference, we fit a normal distribution to the station 23 data and subtract the mean of the distribution from the data. Afterwards, the 68% quantile is calculated to be 8.9° for all UHECR candidate events. The values for simulation and data are compatible within the systematic uncertainties discussed above, which can also lead to a broadening of the polarization difference distribution (quoted as the difference of the 68%-containment in Table 8.8). We note that a better polarization resolution has been reported by the ARIANNA collaboration also using upward-facing LPDAs [145]. A detailed comparison of the resolutions and their limitations requires in-depth study of both detectors (different antenna configuration, trigger, waveform calibration) and differences in the methodological approach (fitting procedure, analysis cuts) and is beyond the scope of this current analysis.

After taking into account the mentioned uncertainties, we can conclude that the reconstructed parameters of the UHECR candidates agree well with the expectation from simulations. Moreover, we want to emphasize that it is highly unlikely for a background source to reproduce all reconstructed parameter distributions. Consequently, the reconstruction supports the interpretation that most of the candidates are events from UHECR air showers.

8.7 Background contamination

Given the discussion of systematic uncertainties and the distributions shown in Figure 8.26 and 8.27, it is unclear whether a perfect separation of signal and background is possible. In the following, we discuss methods for calculating a background estimate (see Section 8.7.1), perform checks to exclude the possibility that the UHECR candidate events originate from known background sources (see Section 8.7.2), and conduct additional analyses to further constrain potential background events (see Section 8.7.3).

8.7.1 Methods for background estimate

Estimating the expected background is challenging because not all background sources and their underlying distributions are known. The background rejection cuts that we apply in the analysis are designed to be agnostic to the origin of the background and instead target its observable characteristics. For instance, a cut targeting the origin of background events occurring during high-wind periods would reject all time periods when the wind exceeds a specified threshold. In contrast, to reject these events, we apply the time cluster cut, which utilizes the fact that they occur in temporal clusters.

[10] Here, σ corresponds to the RMS of the simulated noise in the trigger frequency band (80 − 180 MHz).

Distribution fit: If the underlying background distribution is known, we can fit the distribution or its tail to extrapolate the number of background events in the signal region. However, this requires a comprehensive understanding of the background distribution, which is not available for the impulsive backgrounds in this analysis. Consequently, we will not apply this approach to estimate the background.

ABCD method: Another commonly used technique in particle physics for estimating backgrounds in searches for new physics or for measuring rare Standard Model processes is the ABCD method [221]. The ABCD method requires defining four regions: the signal region (A) and three background regions (B, C, and D). The primary assumption underlying this method is that the number of data points in each region follows a Poisson distribution and that the mean μ for each region satisfies the following relationship:

$$\frac{\mu_{\mathrm{A}}}{\mu_{\mathrm{B}}} = \frac{\mu_{\mathrm{C}}}{\mu_{\mathrm{D}}} = \tau \tag{8.6}$$

Typically, regions C and D contain a high density of background events and can therefore be used to extrapolate the background from region B into the signal region. This is achieved by constructing a likelihood function as the product of the Poisson distributions:

$$\mathscr{L} = \prod_{\mathrm{i=B,C,D}} \mathrm{Pois}(\mathrm{N}_i|\mu_i) = \prod_{\mathrm{i=B,C,D}} \frac{\mu_i^{\mathrm{N}_i} \exp\left(-\mu_i\right)}{\mathrm{N_i}!} \tag{8.7}$$

When the number of events per region is large, computing the factorial directly becomes infeasible. Consequently, the logarithm of the likelihood is maximized:

$$\log \mathscr{L} = \sum_{\mathrm{i=B,C,D}} \log\left(\frac{\mu_i^{\mathrm{N}_i} \exp\left(-\mu_i\right)}{\mathrm{N_i}!}\right) = \sum_{\mathrm{i=B,C,D}} \left[\log(\mu_i^{\mathrm{N}_i} \exp\left(-\mu_i\right)) - \log(\mathrm{N_i}!)\right] \tag{8.8}$$

To calculate the factorial, we apply the following relation:

$$\log(N!) = \log\left(\prod_{i=1}^{N} i\right) = \sum_{i=1}^{N} \log(i) \tag{8.9}$$

With this relation, we simplify the likelihood to the following form.

$$\log \mathscr{L} = \sum_{\mathrm{i=B,C,D}} \left[\mathrm{N}_i \log(\mu_i) - \mu_i - \sum_{i=1}^{N} \log(i)\right] \tag{8.10}$$

Given the number of events per region (N_i), maximizing the likelihood yields estimates for the mean values of the Poisson distributions. This enables the calculation of τ and, consequently, the extrapolation of the mean of the Poisson distribution that describes the number of background events in the signal region (region A). From this extrapolated Poisson distribution, we can determine a confidence interval for the estimated background in the signal region.

In the analysis, we define regions A and B on the dataset after all background rejection

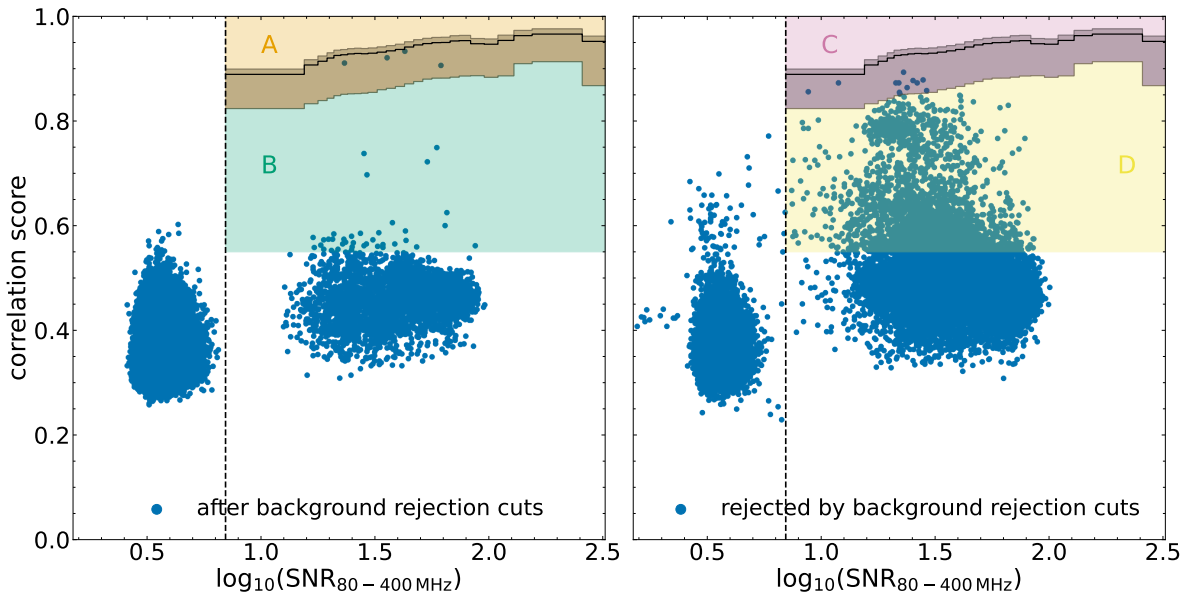


Figure 8.31: Definition of regions A, B, C and D as applied in the ABCD method, using data from station 13 collected in 2022. Left: All data that pass the background rejection cuts. Right: Data that is rejected by the background cuts. The black line represents the SNR-dependent correlation score cut, and the grey region indicates the systematic uncertainties (see Section 8.5). The vertical dashed line marks the lower SNR cutoff for the regions, defined as the lower edge of the last bin of the correlation score cut.

cuts are applied, in the correlation score vs SNR space. Region A comprises all events selected by the SNR-dependent correlation score cut, while region B comprises all other events with a correlation score above a specified threshold. This threshold is set to 0.55, consistent with the minimal correlation score used in the time cluster cut. We define regions C and D on the background sample, which contains all events rejected by the background rejection cuts. Otherwise, we define both regions analogously to regions A and B. Additionally, for all regions, we apply a lower SNR cut-off, corresponding to the start of the SNR-dependent correlation score cut. Figure 8.31 illustrates the four regions using data from station 13 in 2022 as an example.

We maximize the likelihood using the number of events in the regions B, C and D. With the resulting values for μ_B, μ_C and μ_D, we calculate μ_A via Equation 8.6. Subsequently, we use the Poisson distribution for region A to draw the number of events in region A 10000 times to compute the 95% confidence interval for the background estimate. The 95% confidence interval for the background estimate in region A is $[0, 1]$ events.

However, for this method, we defined the lower region cut-off for the correlation score at a somewhat arbitrary value of 0.55. This value is motivated by the time cluster cut, though we could have set it to other values. To assess the sensitivity of the background estimate to this parameter, we repeated the background estimate for different cut-off values. Figure 8.32 shows the 95% confidence interval of the background estimate as a function of the cut-off value. Even a minor change in the cut-off (from 0.55 to 0.54) significantly alters the background estimate (the 95% confidence interval changes from [0, 1] to [0, 2]). This demonstrates that the ABCD method cannot provide a reliable background estimate for this analysis, as small adjustments to the region definition have a substantial impact on the estimate. A potential issue is the selection of the two datasets on which we define regions A, B, C, and D. We designed the background rejection cuts

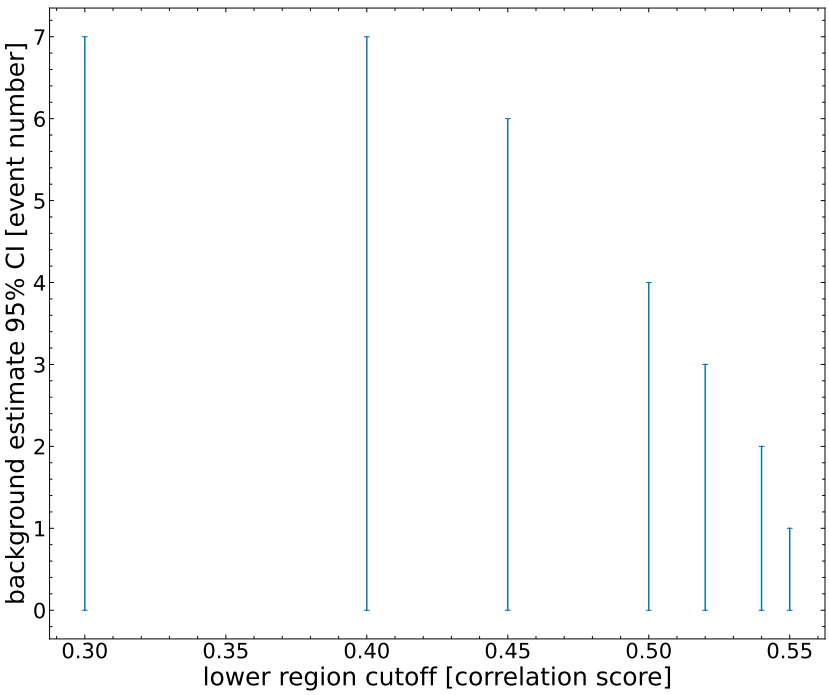


Figure 8.32: Background estimates obtained from the ABCD method as a 95% confidence interval, plotted as a function of the lower correlation score cutoff used to define the regions B and D.

to target background events with high correlation scores. For example, the time cluster cut considers a time cluster only if the correlation score exceeds 0.55. The design of the cuts and the definition of the datasets likely violate the main assumption underlying this method (Equation 8.6). Consequently, we cannot apply this method to estimate the background.

8.7.2 Potential background sources

Given the infeasibility of a comprehensive background estimate, we performed several checks to ensure that the identified UHECR candidate events do not originate from known background sources. In the following paragraph, we describe each of these checks.

Thermal noise: Generally, we do not expect the stations to trigger on thermal noise in the shallow channels due to the high RADIANT trigger thresholds. However, Figures 8.26 and 8.27 reveal a distribution of events at SNR values and correlation scores consistent with thermal noise. To investigate this, we calculate, for each station, the SNR values and correlation scores with UHECR templates for forced trigger events, and overlay the resulting distribution contours in Figure 8.33 with the data from Figures 8.26 and 8.27. The result demonstrates that the distributions of events at low SNR and correlation scores align well with the forced-trigger contours. Since forced-trigger data primarily contain thermal noise events, this result supports the hypothesis that the observed event distribution consists of thermal noise. For station 23, which has the lowest trigger threshold, the impulsive and thermal noise begin to overlap. Nevertheless, a part of the combined distribution still matches the forced-trigger data. In the 2022 data for station 13,

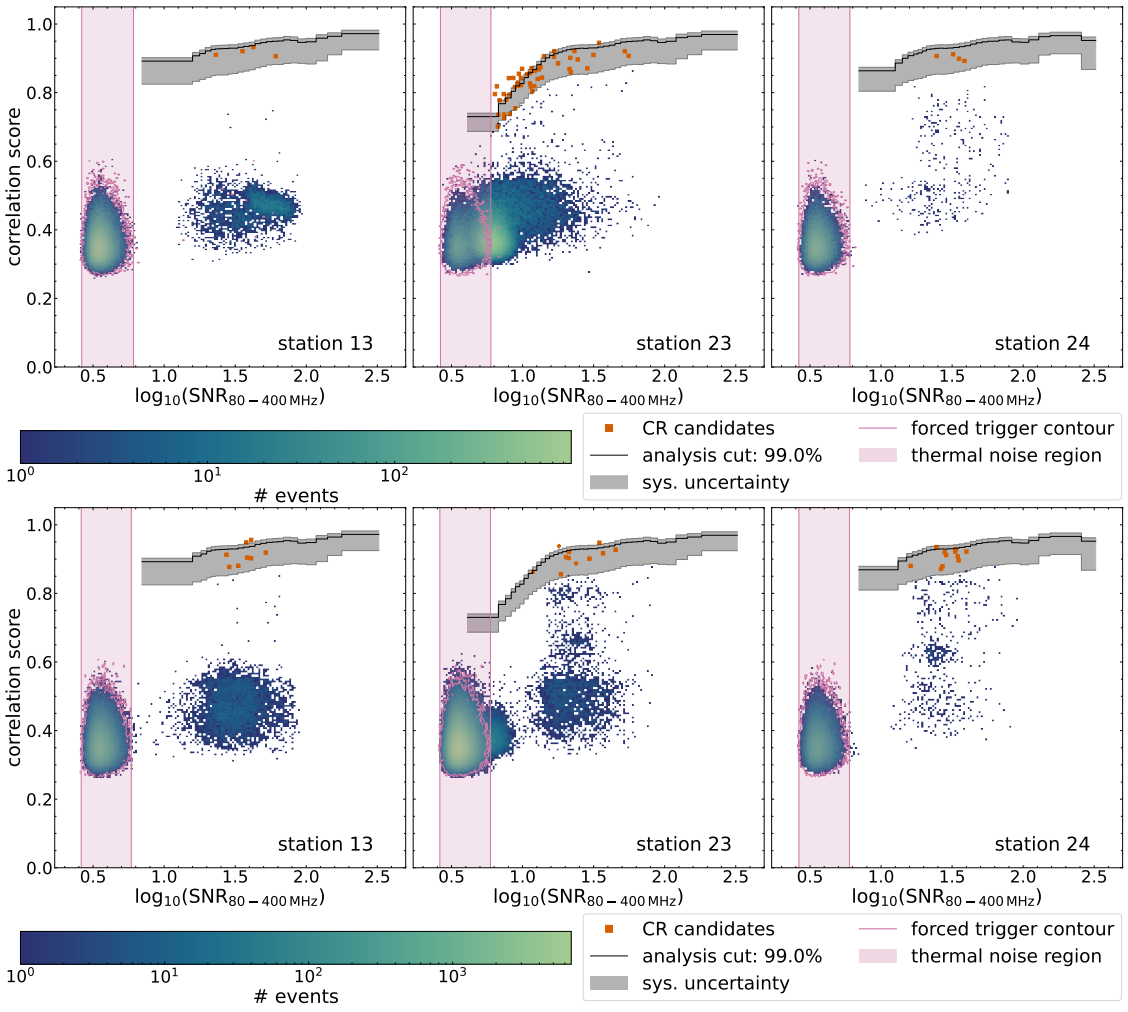


Figure 8.33: The 2022 (top) and 2023 (bottom) data for each station after applying all background rejection cuts, the SNR-dependent correlation score cut including the systematic uncertainties, and the selected cosmic-ray (CR) candidate events. Additionally shown are the contour of the forced trigger data distribution and the region where thermal noise is expected.

a few forced trigger events exhibit high SNR values and fall within the impulsive event distribution. Manual inspection confirmed that these events have impulsive waveforms, suggesting they were triggered by chance during periods of frequent impulsive events, such as high wind periods. Typically, thermal noise should not appear in data triggered by the RADIANT trigger, given the high trigger thresholds. However, they occur because the trigger path is dominated by noise from the RADIANT board itself rather than input noise. A detailed discussion is provided in Section 5.3.1.

In the following, we discuss whether these thermal noise events contribute as a noise source to the signal region. First, we estimate the SNR range where thermal noise is expected using approximately 100, 000 forced-trigger events per station and year. The events correspond to the data set shown in Figure 8.22, which includes 1-2 weeks of forced trigger data from the start and end of the data-taking period. For each station, we use the 99.9968% and 0.0032% quantiles (representing the 4σ region) of the forced trigger SNR distribution to estimate the SNR region in which thermal noise is expected. The resulting regions are indicated in Figure 8.33. For stations 13 and 24, the trigger threshold is sufficiently high that the SNR-dependent correlation score cut does not extend into

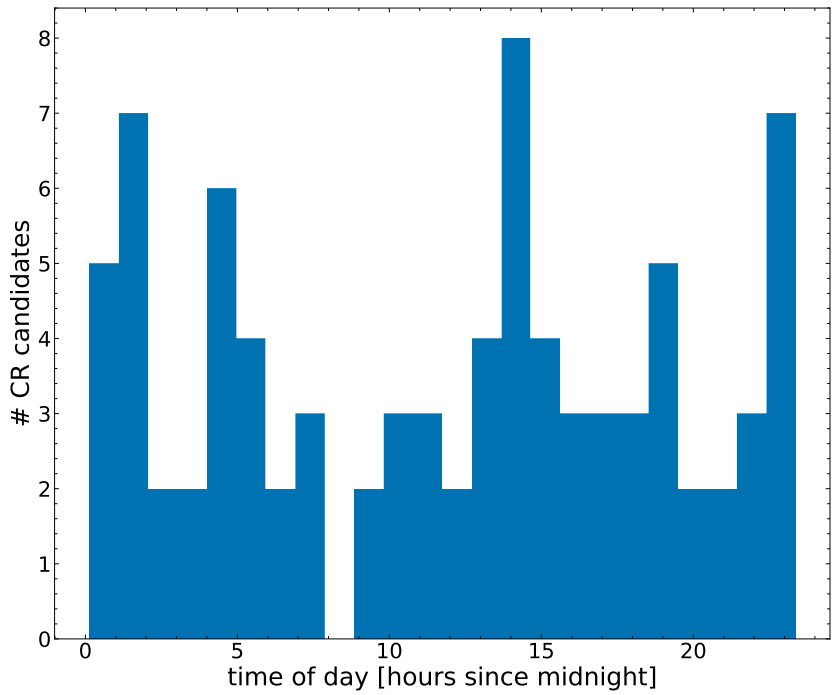


Figure 8.34: Histogram of the time of day at which each cosmic-ray (CR) candidate event was triggered. All CR candidate events identified in this analysis are shown.

the thermal noise region, so no thermal noise events are expected to contribute to the signal region. For station 23, which has the lowest trigger threshold, the correlation score cut does reach into the thermal noise region, and in principle, such events could enter the signal region. However, in both 2022 and 2023, no events in the signal region fall within the thermal noise range. Therefore, no thermal noise events contribute to the signal region of station 23.

Although thermal noise is expected to follow a well-behaved distribution, and a fit to its tail could estimate the expected number of thermal noise events in the signal region, we did not perform this estimation in this analysis. For stations 13 and 24, such an estimation is unnecessary because the signal region does not extend to sufficiently low SNR values. For station 23, the estimation would be challenging due to the difficulty in separating impulsive and thermal noise distributions or in estimating the pure thermal noise trigger rate.

Overall, we conclude that no pure thermal noise events contaminate the UHECR candidate sample.

Time of day: Certain background sources appear at specific times. For example, events associated with heavy machinery at Summit Station or passing snow machines predominantly occur during standard daytime work hours. If the UHECR candidate events originated from such backgrounds, we would expect a temporal clustering of candidate events during these periods. In contrast, UHECR events are not expected to show a time-of-day preference. As a cross-check, we present the distribution of trigger times for UHECR candidate events as a histogram in Figure 8.34. Given the small number of events per bin, the data do not reveal any significant time-of-day clustering. Therefore,

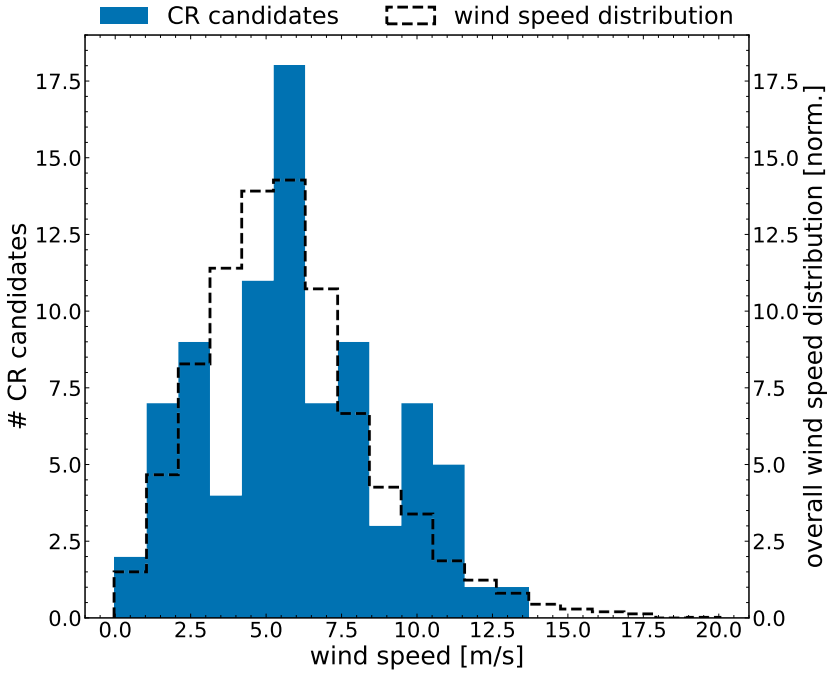


Figure 8.35: Histogram of the wind speed at the trigger time of the cosmic-ray (CR) candidate events (blue) identified in this analysis. Additionally, the overall wind speed distribution for the data-taking period is shown (black dashed), which is normalized to match the total number of candidate events.

most UHECR candidate events likely do not originate from background sources with preferred activity periods. Nevertheless, this does not exclude the presence of some events from such background sources in the sample.

Wind speed: During periods of high wind (wind speed > 10 m/s), impulse events can be produced by the triboelectric effect that trigger the detector [173]. To assess whether a significant fraction of these background events contaminates the UHECR candidate sample, we extract for each candidate event the wind speed at the trigger time from Ref. [222] and show them as a histogram in Figure 8.35. A sample containing only UHECR events should reflect the underlying wind speed distribution during the data-taking period, which is not a flat distribution, as UHECR events do not have a preferred wind speed at which they occur. In contrast, if some events originate from the triboelectric effect, we expect them to cluster at wind speeds above approximately 10 m/s. The UHECR candidate wind speed distribution aligns with the overall wind speed distribution, considering the limited number of events per bin. Notably, none of the candidate events occurred at wind speeds exceeding 13.5 m/s, which can probably be attributed to two factors. First, from the overall wind speed distribution during the considered data-taking periods, we expect only very few UHECR events at these high wind speeds. Second, the time cluster cut is designed to reject periods with substantial time-clustered impulsive noise, such as periods of high wind. Consequently, a UHECR event could have been triggered at higher wind speeds but subsequently rejected by the time cluster cut. In conclusion, we find no evidence for event clustering at high wind speeds, indicating that the UHECR candidate events do not originate from the triboelectric effect. However, since some candidate events are triggered at wind speeds above 10 m/s, we cannot entirely exclude that a few events may result from the triboelectric effect.

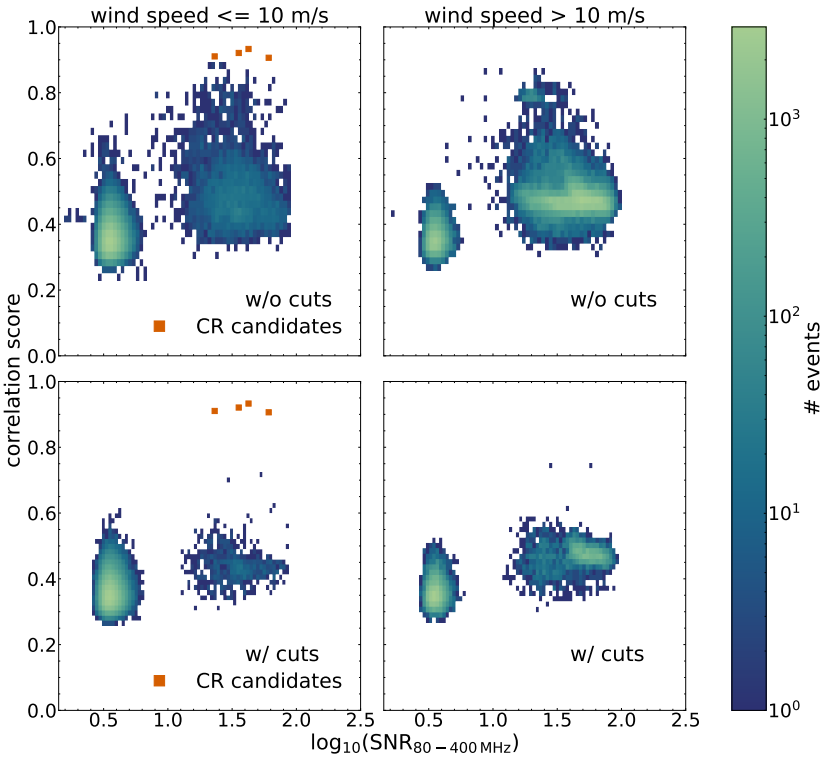


Figure 8.36: Correlation scores for station 13 as a function of the SNR for the 2022 data. The data is shown before (top) and after (bottom) the application of background rejection cuts. Furthermore, the identified cosmic-ray (CR) candidate events for this station and time period are highlighted, and the data recorded at wind speeds $\leq 10\,\mathrm{m/s}$ (left) is separated from the data recorded at wind speeds $> 10\,\mathrm{m/s}$ (right).

In addition to the cross-check described above, we can analyze the entire dataset as a function of wind speed. For this analysis, the dataset for each station is divided into two bins: wind speed values below and above 10 m/s. Figure 8.36 shows, as an example, the dataset for station 13, separated into the two wind speed bins, both before and after the background rejection cuts. The figure indicates that most impulsive events with higher SNR occur at wind speeds above 10 m/s and are thus likely to originate from the triboelectric effect. Comparing the plots before and after the background rejection cuts demonstrates that these cuts are highly effective at rejecting events with high wind speed and large correlation scores. However, many impulsive high-wind events with lower correlation scores survive the cuts. This outcome is expected, as we do not apply a dedicated wind-speed cut, and the time cluster cut targets time clusters with higher correlation values. The remaining events do not compromise the analysis, as they are distant from the signal region. A similar result to the one discussed above is observed when analyzing the datasets from stations 23 and 24 as a function of the wind speed. Overall, these results indicate that the background rejection cuts effectively remove events with high correlation scores during periods of elevated wind speeds.

Airplanes Previous studies have demonstrated that airplanes can generate impulsive signals capable of triggering the detector [181]. As a cross-check to exclude that the UHECR candidate events all originate from airplanes, we investigated the presence of airplanes within 60 km of the station and within ±120 s of the candidate event trigger time. The time interval is selected to correspond to the approximate duration required for

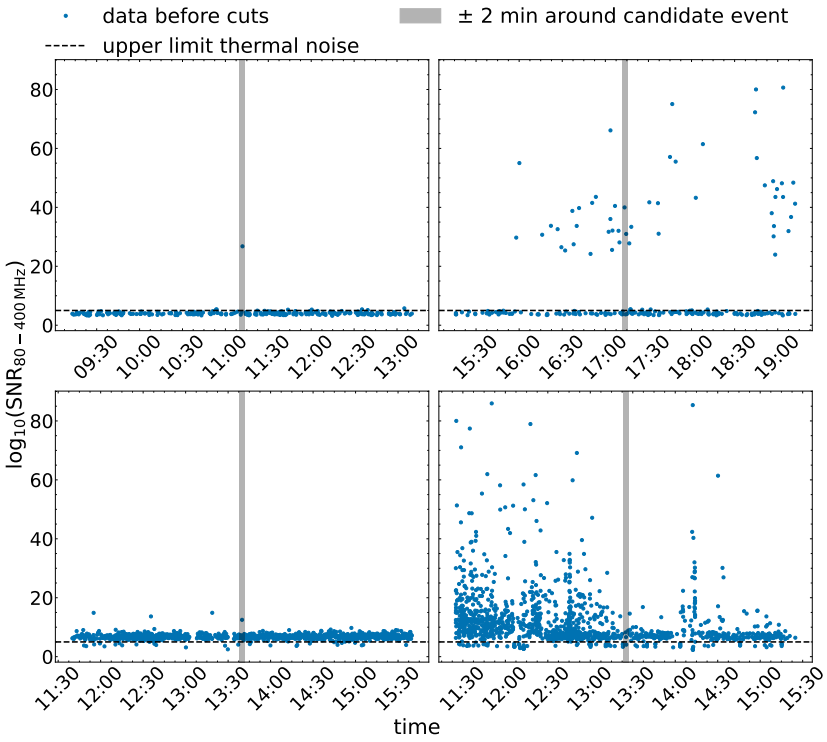


Figure 8.37: The SNR as a function of time for a four-hour period around representative UHECR candidate events with an airplane nearby. For each event, we show the data before any cuts are applied. The shaded region indicates the time period considered for the airplane proximity check, while the dashed line represents an estimate of the upper boundary of the thermal noise SNR distribution. An airplane coincident with a UHECR candidate should lead to multiple high-SNR events in the grey band. Each event was triggered on a different day, but only the time of day is shown in the plot.

an airplane to travel 60 km. Among the 85 UHECR candidate events, 12 meet both criteria, indicating that these events could potentially be attributed to airplanes. Consequently, we examined these 12 events in greater detail.

First, we plotted the SNR of all events within a 4 h window around each candidate as a function of time. If an event originated from an airplane, we would expect that multiple high-SNR signals occur in close temporal proximity to the original event. Figure 8.37 shows the SNR as a function of time for four representative events from the subset of 12 suspicious candidate events. None of these examples exhibits time clustering within a short time (±120 s, time period of the airplane proximity check) around the candidate event. In some cases, we observe additional high-SNR events within the four-hour window, but they do not occur in close temporal proximity to the candidate event. Additionally, for some events, we observe an elevated noise floor (for example, see the lower-left panel of Figure 8.37), which can be attributed to small day-night variations still visible in the station 23 data. Overall, this cross-check suggests that it is highly unlikely that these events originated from nearby airplanes.

Second, we compare the reconstructed directions[11] of the 12 suspicious UHECR candidate events with the airplane directions obtained from a plane tracker. All but three airplanes have a zenith angle close to the horizon, and two of these three events have a zenith angle below the horizon. We manually inspected these two events and found that, although

[11] More information about the reconstruction of the UHECR candidate events is provided in Section 8.6.

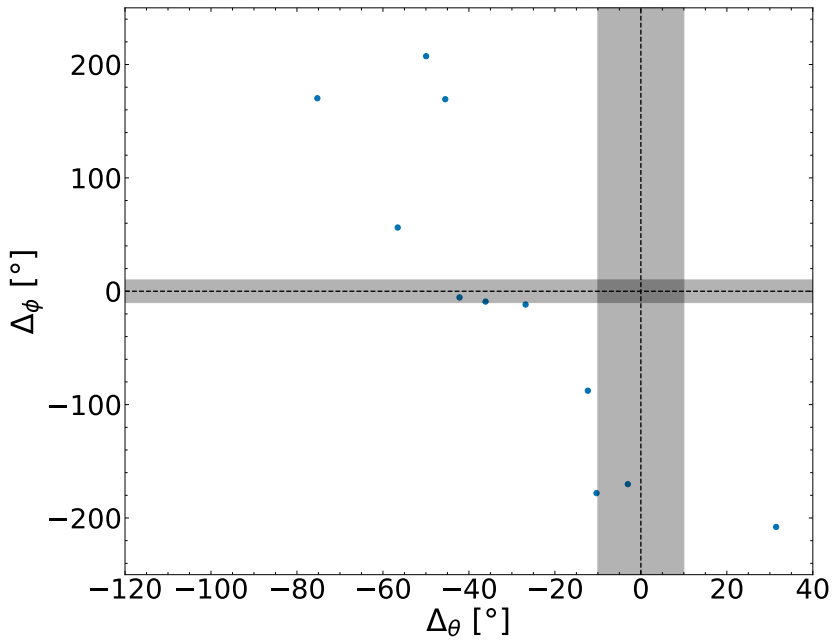


Figure 8.38: Comparison of the reconstructed arrival directions (zenith angle θ and azimuth angle ϕ) of the UHECR candidates with the reported directions of nearby airplanes. Only events with an airplane in proximity to the station in a period of 120 s around the trigger are included. The grey band indicates the region of 10° difference.

reported below the horizon, the two airplanes had landed at Summit Station, indicating that they were transmitting a faulty z position. To compare the reconstructed UHECR zenith and azimuth angle with the airplane directions, we calculate the differences between both, which are shown in Figure 8.38. As a result of the faulty z positions, which are reported for two planes, one of the zenith differences exceeds 90°. For most of the 12 candidate events, the azimuth and zenith directions do not match the airplane directions, as the differences are significantly different from zero. Two events have an azimuth difference less than 10°, but the zenith difference of these events exceeds 35°. Further, a single event has a zenith difference less than 10°, however, the azimuth difference is approximately 170°. Consequently, this cross-check shows that the UHECR arrival direction deviates significantly from the airplane directions.

In summary, with the two cross-checks, we have demonstrated that the UHECR events do not originate from airplanes.

Charge controller The initially used charge controller was found to produce impulsive background noise when the battery reached full charge, and the controller switched off charging. To assess whether these impulsive events significantly contaminate the UHECR candidate event sample, we examined the correlation between the UHECR candidate events and the battery voltage. When the battery voltage attains its maximum, the charge controller switches off charging, potentially generating impulsive noise. If most UHECR candidates originated from the charge controller, we would expect these events to occur when the battery was fully charged. Figure 8.39 shows a histogram of the battery voltage at which each UHECR candidate was triggered. Because the battery voltage distribution is not uniform during data collection due to varying charging and discharging rates, we do not expect an uniform distribution of UHECR candidate events if they are uncorrelated with the battery status. Instead, the candidate event histogram

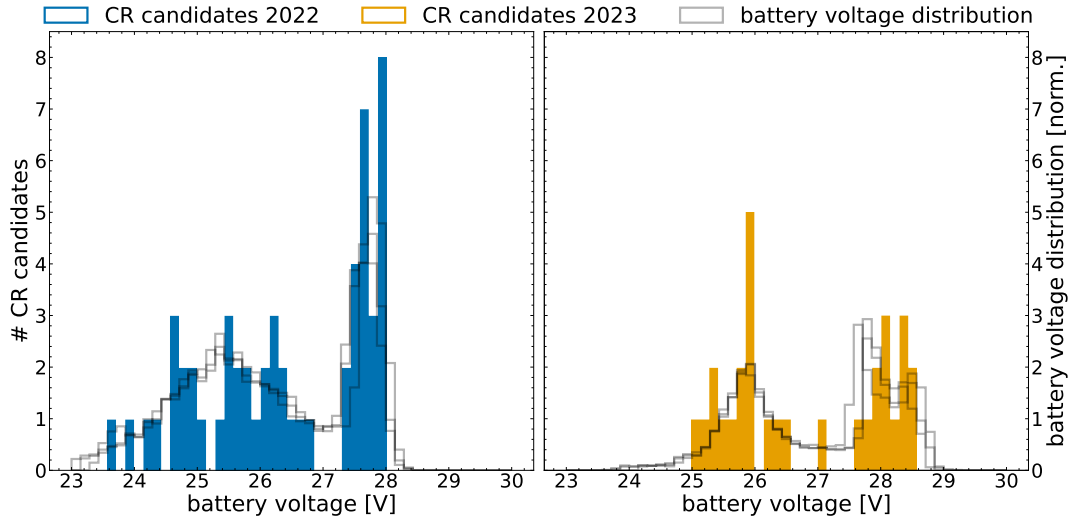


Figure 8.39: Histogram of the battery voltage at which each cosmic-ray (CR) candidate event was recorded. Left: The overall battery voltage distribution and the candidate event distribution for 2022. Right: The same as in the left panel for 2023. In both years, each station shows a slightly different distribution of battery voltages. However, the differences are small enough that we present a combined CR candidate distribution that includes events from all three stations.

should reflect the overall distribution of battery voltages. The overall battery voltage distribution differs slightly between 2022 and 2023. The cause of this difference is not known. Therefore, Figure 8.39 shows the UHECR candidate distribution and the overall battery voltage distribution for each year in separate plots. In both years, the candidate event distribution closely follows the overall distribution, considering the limited number of events per bin. Furthermore, if a majority of UHECR candidate events originated from the charge controller, their reconstructed directions should point towards the device, which is not the case, as we present in Section 8.6. In summary, there is no strong indication of the candidate events originating from the charge controller, although we can not exclude the possibility that individual events are correlated with the charging electronics.

8.7.3 Source-independent cross-checks

In the following, we perform two additional source-independent cross-checks to get a better handle on possible background events.

Background reconstruction As an additional cross-check, we reconstruct background events to determine whether they yield parameter distributions consistent with those expected from the simulation. For the reconstruction, we use the same algorithm as applied for the UHECR candidate reconstruction (see Section 8.6), and as the simulation, we use the one described in Section 8.6. The background events for this cross-check are all events rejected by any of the three background rejection cuts. However, due to computational constraints, we only reconstruct a subset of these events. To select impulsive background events for the reconstruction, we divide the background dataset into three regions based on their correlation scores: less than 0.55, between 0.55 and 0.7, and greater than 0.7. We do not consider background events compatible with thermal

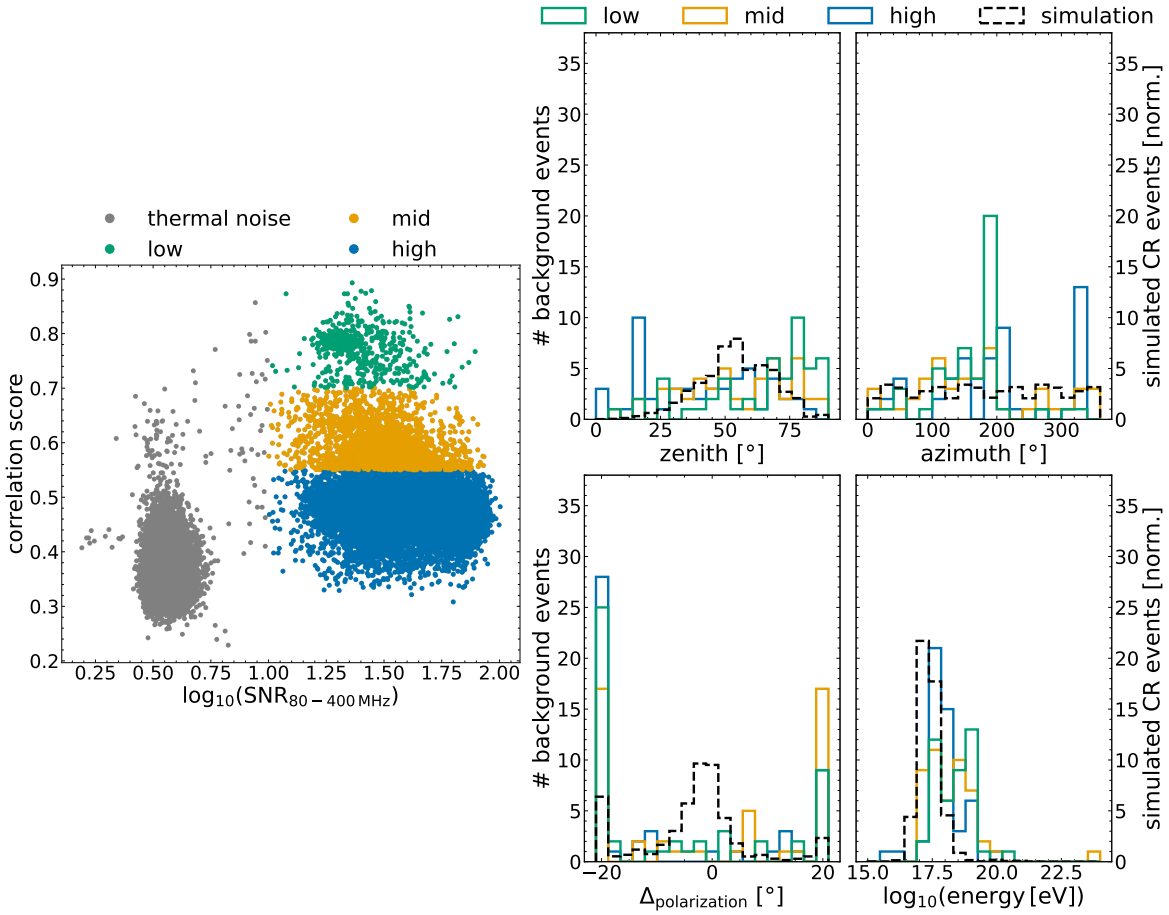


Figure 8.40: Left: Background sample for station 13, with different colors representing the regions we defined for the background reconstruction. Right: The reconstructed values for azimuth angle, zenith angle, polarization difference, and energy of the background events from the three correlation-score regions.

noise, which appear at lower SNR values. Figure 8.40 shows the background data for station 13 with the different regions color-coded, as an example. For each correlation-score region, we randomly select 50 events for the reconstruction to ensure that the sample is not dominated by events with low correlation scores, which could bias the distributions. The resulting reconstructed distributions of zenith angle, azimuth angle, polarization difference (reconstructed - expected polarization), and energy are shown in Figure 8.40 together with the expected distributions from a UHECR simulation. The parameter distributions for zenith angle, azimuth angle, and polarization difference differ from the simulated UHECR distributions for all correlation score regions. Only the reconstructed energy approximately falls within the expected range, which is consistent with the fact that the impulsive background events have amplitudes similar to those of the UHECR candidates. Additionally, in the lowest- and highest-correlation-score regions, we observe clustering in the azimuth angle distribution. Interestingly, events from these two regions cluster around different azimuth angles. We did not investigate the presence of known sources at these azimuth angles, but future analyses could use this information to identify and mitigate additional impulsive background noise sources. The reconstruction of the background events from stations 23 and 24 shows a comparable level of disagreement between the background reconstruction and the simulation (see Figures A.7 and A.8).

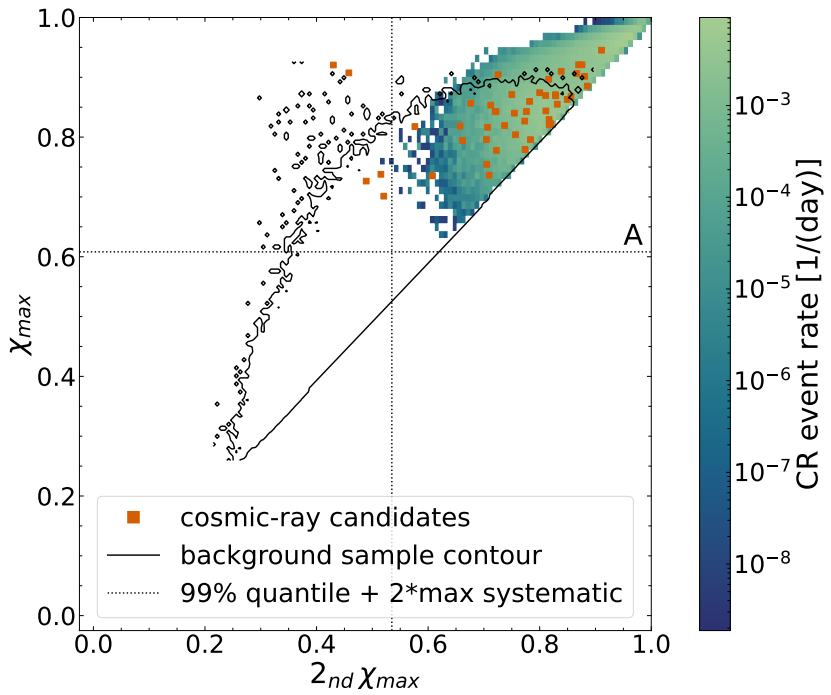


Figure 8.41: Correlation value of the channel with the best match as a function of the correlation value of the second-best channel for station 23. Shown are the contours of a background sample in black, the expected distribution of cosmic rays without systematic uncertainties (color), and the cosmic-ray candidates (squares). Region A indicates the signal region including systematics. Taken from Ref. [185].

This cross-check demonstrates that the reconstructed parameters of pure background events do not follow the distributions expected from simulations. Therefore, it is highly unlikely that the majority of UHECR candidate events are background. However, this cross-check does not allow us to conclude that there are no background events in the UHECR sample.

Second-largest correlation score We investigate the potential of using the channel with the second-largest correlation score for background discrimination. To do so, we consider the distribution of the largest correlation score (χ_{max}) as a function of the second-largest correlation score ($2_{\mathrm{nd}}\chi_{\mathrm{max}}$) in the same event. In a few cases, background events correlate well with the template only in a single channel, which makes the second-largest correlation score a valuable quantity to consider. Also, due to the strong polarization of the UHECR signal and the triangular detector geometry, the UHECR signals are expected to not always correlate equally well for all channels, resulting in a broadened UHECR distribution. Figure 8.41 shows the resulting distributions for station 23 in 2022, including simulation, background data, and UHECR candidate events. By definition, all data points lie above the diagonal line. The black background sample contour represents the overall distribution from events rejected by background rejection cuts, and the simulation is the same as used in Section 8.5 for the SNR-dependent correlation score cut. The background data contour and simulation distribution overlap clearly, making this metric unsuitable as an additional UHECR selection criterion.

However, the second-largest correlation score reveals that five UHECR candidate events are significantly separated from the simulation. To assess their compatibility with the

simulation, we need a distance measure. The best approach would be to separate the simulation into $2_{\text{nd}}\chi_{\text{max}}$ (χ_{max}) bins, calculate the 99% quantile per bin, and evaluate the impact of the different systematic effects on the quantiles. This would provide a decision criterion analogous to the SNR-dependent correlation score cut, including systematic uncertainties (see Section 8.5) for both $2_{\text{nd}}\chi_{\text{max}}$ and χ_{max}. All events that would not satisfy the selection criteria for both quantities would be identified as incompatible with the simulation. Given the time-intensive nature of this method, we adopt a conservative alternative: we calculate the overall 99% quantile for both quantities and apply twice the maximum systematic uncertainty from Section 8.4 as a conservative estimate of the systematic uncertainty. For $2_{\text{nd}}\chi_{\text{max}}$, we assume that the systematic uncertainties are comparable to those estimated for χ_{max}. Even if this assumption is not strictly valid, applying twice the maximum systematic uncertainty ensures a conservative estimate. The resulting 99% quantiles, including the systematic uncertainty, are indicated in Figure 8.41 by a vertical and horizontal dotted line, creating the area *"A"*. We deem events outside the area "A" to be incompatible with the simulation. Consequently, we identify 5 UHECR candidate events from station 23 in 2022 as incompatible with the simulation, and thus, these events are likely background contamination. Notably, two of these background events are also outliers within the background data sample, suggesting they represent a very rare background.

We provide equivalent diagnostic plots for the other stations and data-taking periods in the Appendix (see Figures A.9, A.10, and A.11). For the stations with higher trigger thresholds, the overlap between simulation and background data is substantially reduced. Furthermore, we do not observe additional background contamination in the UHECR sample for these stations and years.

Additionally, we examined the reconstructed properties of the 5 UHECR candidate events that we identified as likely background. The reconstructed properties are shown in Figure 8.30 as *"likely background"*. For three events, the reconstructed energy is either too high ($E > 10^{21}$ eV) compared to the UHECR flux measurements from much larger experiments, which have measured the flux cutting off beyond energies of 10^{20} eV [53] or too low ($E < 10^{15}$ eV) such that it does not fit with the expectations from the trigger thresholds. Additionally, three events exhibit a large polarization difference ($\Delta_{\text{polarization}} > 17°$), and for two events, the reconstructed zenith angle is pointing toward the horizon, which is unlikely for UHECR events. These findings further support classifying these five UHECR candidate events as background.

In summary, using the second-largest correlation score as an additional metric enabled the identification of at least 5 background events within the sample of 85 UHECR candidates.

8.8 Analysis of further stations

As discussed in Section 8.3, the stations 11, 12, 21, and 22 were excluded from the main analysis due to either an experimental wind turbine or an early version of the power system causing an unstable performance of the shallow channels. Nevertheless, we

calculated the data-template correlation from all of these stations. Figures 8.42, A.12, A.13, and A.14 show the correlation score as a function of time for each station and year.

For these stations, the majority of events have a correlation score higher than the expected value for thermal noise (indicated by the grey band in the figures) and, with that, also higher than the minimal correlation score used for the time cluster cut (indicated by a black dashed line in these figures). Using the same cuts as above reduces the station's lifetime drastically. The increase of background events at higher correlation scores suggests that the noise coming from either the early test wind turbine or the early power system is impulsive, relatively similar to UHECR waveforms. This makes the shallow antennas, which are located close to these sources, challenging to use during these periods. Upgrades in the power boxes of the two test wind turbines were implemented in early July 2023 to reduce the noise in the shallow antennas of stations 11 and 12. Figures 8.42 and A.12 show that the noise level begins to decrease around the same time, indicating that the additional shielding has a positive effect.

However, applying the time cluster cut results in lifetime losses of the order of 80% for some of these stations when following the same procedure as defined in Section 8.5. It might be possible to design custom cuts per station to recover some of the data for a full analysis, but we decided against this approach, given that these stations are expected to have higher quality data in the coming years. In 2024, the power system was fully upgraded for the stations 21 and 22 mitigating the known issues. Additionally, the test wind turbine was replaced with a different model in 2024 and preliminary data indicates that they produce less noise in the shallow antennas.

Although these four stations are affected by large amounts of high-correlation background, we are able to identify clear UHECR candidate events by inspecting the waveforms of the highest-correlation score events. For each station and year, the waveform of the channel with the highest correlation score is shown for up to three UHECR candidate events in Figures 8.42, A.12, A.13 and A.14. This demonstrates that, despite the issues discussed above, the stations are able to trigger on high SNR UHECR events when the up-time allows for it.

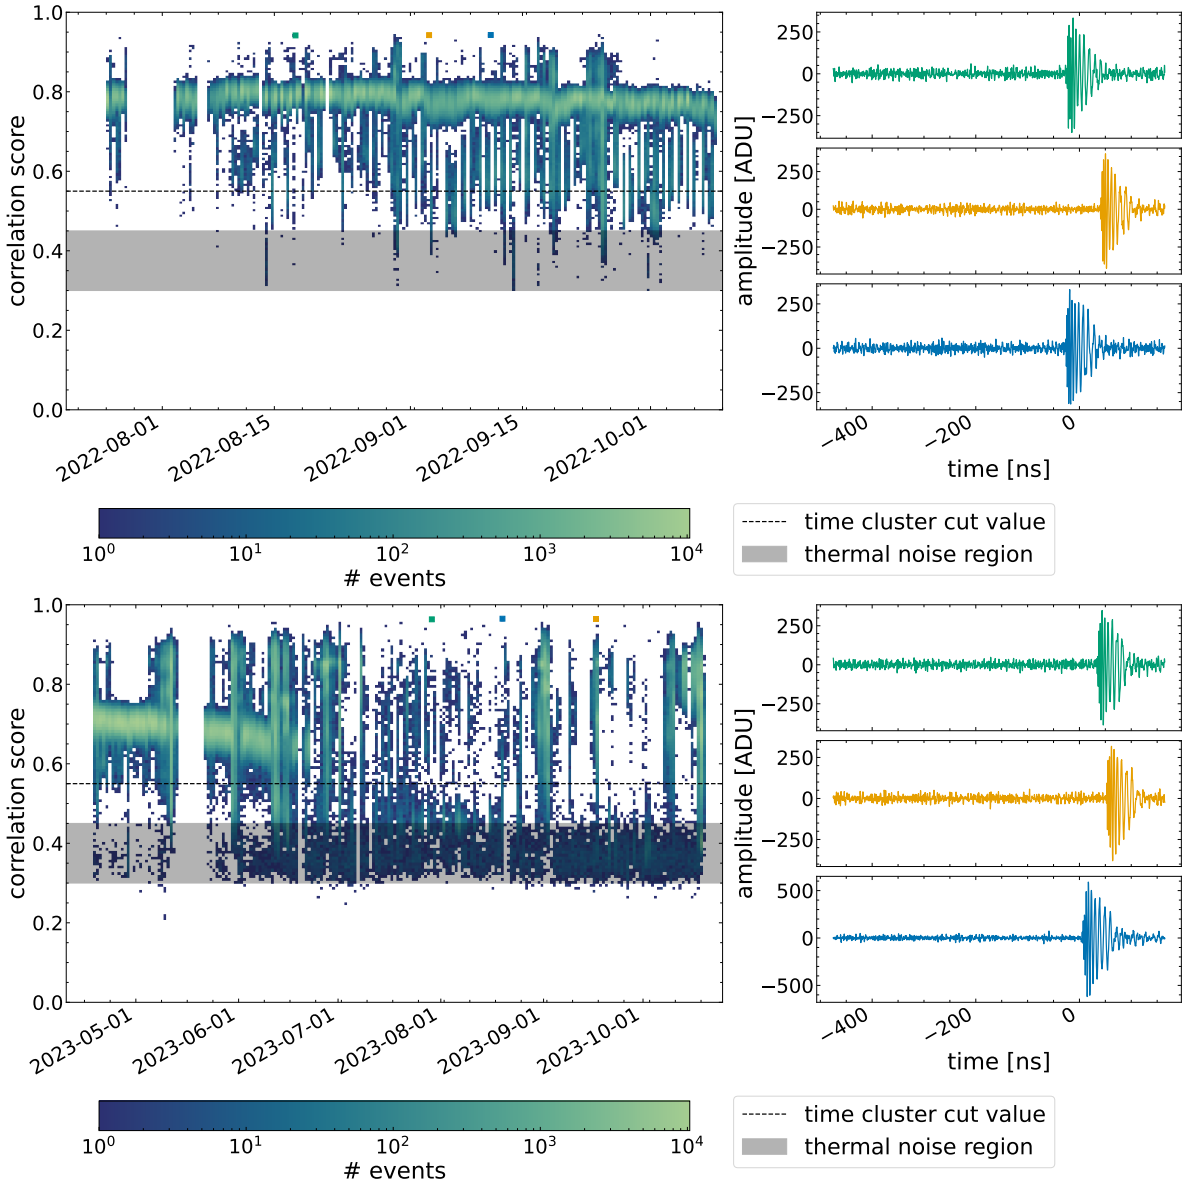


Figure 8.42: The correlation score as a function of time for station 11 and years 2022 (top) and 2023 (bottom) without any cuts applied. Additionally shown are three example waveforms from UHECR candidate events, which are also denoted by squares and colors in the left panel. Each shown waveform corresponds to the channel with the highest correlation score per highlighted event, with the color indicating which event it belongs to. In July 2023, changes to the power system were made. Taken from Ref. [185].

8.9 Conclusions

In this chapter, we provided a detailed description of the first UHECR analysis using the shallow antennas of RNO-G, present the results and discuss potential systematic effects and background contamination.

The UHECR analysis uses a template-matching approach: we generate a set of templates, correlate them with the recorded data, and use the resulting correlation score as the discriminative variable. This method is highly sensitive to the shape of both the recorded and template waveforms. However, we demonstrated that the shape of the recorded waveforms is primarily determined by the antenna, amplification and digitization responses of the detector. Accordingly, we developed a procedure to generate a template by applying the antenna, amplification and digitization responses to artificially generated electric fields. With this process, we produced a set of two signal and two background templates that sufficiently cover the UHECR parameter space. The main distinction between the templates is the width of the artificial electric field, which determines the frequency content in the final waveform. Several parameters can be tuned for generating the templates. We conducted multiple checks and small simulations to inform the selection of these parameters, although we do not claim to have converged on the optimal ones. Further optimization of these parameters may be possible in future work.

The primary dataset for this analysis comprises data from stations 13, 23, and 24 for the years 2022 and 2023. Due to changes in the station configuration and the replacement of key hardware components in 2024, we exclude more recent data-taking periods from the analysis. Further, we do not consider data from other stations because their trigger rates are higher and more unstable during the relevant time periods. In 2024, the issues that caused the elevated and unstable trigger rates at stations 11, 12, 21, and 22 were mitigated. As a result, we expect future analyses to benefit from the improved data quality at these stations, allowing their inclusion in a dataset.

We evaluated several potential systematic effects and identified four as contributing significantly to the analysis. The most substantial contribution arises from the uncertainty in the refractive index of shallow ice. Therefore, improving the understanding of the refractive index is expected to reduce the overall systematic uncertainty of the analysis substantially. Although new methods and ice models to describe the refractive index are being developed or are already available, we could not incorporate them into this analysis due to time constraints. Future work should assess the impact of improved refractive index models on the systematic uncertainty. The other substantial systematic uncertainties are related to the antenna response, the amplification and digitization response, and the timing calibration of the digitizer chip.

In the analysis, we begin with a data-reduction cut, then calculate all analysis quantities and apply three background-rejection cuts. Then, we apply an SNR-dependent correlation score cut to select the UHECR candidate sample. To derive the correlation score cut, we perform custom simulations for each station and data-taking period, accounting for realistic trigger thresholds. Additionally, we apply all relevant cuts to the simulation and incorporate the systematic uncertainty into the correlation score cut. During the development of the background rejection cuts, we observed temporal variations in the noise RMS throughout the data-taking period. Although we did not further investigate

the origin of this variation, understanding its cause is important for evaluating the stability of the detector. The analysis identifies 85 UHECR candidate events across three stations and two years, and the observed number of events agrees with expectations from simulations for all stations and years. Notably, more than half of the events were detected by station 23 in 2022, when it operated in the most sensitive configuration among all stations and data-taking periods. Future UHECR analyses are expected to benefit from the 2024 upgrade of the RADIANT board, which improved the trigger performance. With this upgraded design, we anticipate detecting a comparable number of UHECR candidates with less than 3000 h of single-station lifetime. In contrast, this analysis uses a total of 15182 h of single-station lifetime. As a result, future analyses are expected to detect a larger set of UHECR candidate events.

After selecting the UHECR sample, we reconstruct the direction, polarization, and energy for all candidate events. The distribution of each reconstructed parameter matches expectations from simulations within the systematic uncertainties of the reconstruction. For the polarization, we observe a bias in the events from station 23, which can be attributed to systematic effects, particularly an uncertainty of the antenna orientation. Following proper antenna position and orientation calibration, we expect the bias to be resolved and the polarization resolution (68% confidence interval of the polarization difference) to be improved. Without this calibration, we determine a bias-corrected polarization resolution of 8.9°. It is highly unlikely that a sample of background events would produce the same distribution across all parameters. Therefore, the reconstruction results support the interpretation that the majority of the candidates are UHECR events.

To better assess potential background contamination in the UHECR sample, we performed multiple cross-checks and explored two methods for calculating a background estimate, though without success. We investigated whether the UHECR candidate events originated from known background sources, such as airplanes or high-wind periods. These background source checks indicate that the investigated sources cannot explain the majority of candidate events. However, we cannot exclude a small contamination of background events in the sample based on these cross-checks. Additionally, we conducted two source-independent cross-checks. First, we reconstructed a sample of background events, which showed a distribution different from that in the simulation, supporting the interpretation that most candidate events are UHECR events. Second, we included the second-largest correlation score per event as an additional variable, enabling the identification of five candidate events likely to be background events. The reconstructed properties of these five events further support their classification as likely background events.

Furthermore, we examined the data from the stations excluded from the main analysis. This analysis supports the decision to exclude these stations, as applying the same background-rejection cuts to them would have resulted in a lifetime loss of the order of 80%. Although we could have developed custom cuts for these stations, this approach was not pursued, as we anticipate improvements in data quality in subsequent data-taking periods. Incorporating these stations in future analyses is expected to yield a larger set of UHECR candidate events. Nevertheless, we identified some UHECR candidate events in data from these stations by inspecting the highest-correlation-score events, demonstrating that these stations can trigger on high-SNR UHECR events despite the existing issues.

9 Conclusions

Ultra-high-energy cosmic rays (UHECRs) are an important background for RNO-G and other in-ice radio neutrino telescopes. Consequently, it is essential to characterize and, whenever possible, veto this background. Furthermore, UHECRs offer a valuable opportunity to test the detector under realistic conditions, as their signals occur at the nanosecond timescale, similar to ultra-high-energy neutrino (UHEN) signals. The primary scientific goal of this thesis was to develop an analysis to identify the first UHECR air showers using the shallow antennas of RNO-G. This analysis is an essential first step towards validating the detector and high-fidelity UHEN detection with RNO-G in the future.

In this thesis, we analyzed the full 2022-2023 dataset for three stations and reported the detection of 85 UHECR candidate events, including their reconstructed properties. Both the number of UHECR candidates and their reconstructed properties agree with expectations from simulations. This demonstrates that, despite some necessary modifications - which were identified in this analysis - the detector performs as expected and has successfully passed the commissioning aspect using direct signals from air showers.

For the analysis, we adapted the template-matching method previously developed for the ARIANNA experiment by reducing the number of required signal templates. This was achieved by utilizing the fact that the signal is primarily shaped by the antenna, amplification, and digitization responses, which enables the use of only two signal templates while still sufficiently covering the parameter space.

For each UHECR candidate event, we reconstructed the direction, polarization, and energy. Additionally, we investigated the impact of systematic effects on the reconstruction results. The antenna orientation was identified as having the strongest potential influence on the event reconstruction, and therefore, calibrating the antenna positions and orientation will improve the UHECR reconstruction. It should be noted that it is highly unlikely that background sources would produce UHECR-like distributions for all reconstructed parameters, supporting the interpretation that the majority of the candidate events originate from UHECRs.

To assess potential background contamination in the final UHECR sample, we performed multiple cross-checks to ensure that the events did not originate from known background sources such as airplanes. Additionally, we included the second-largest correlation score per event as an additional variable, enabling the identification of five events that are likely background.

Through the UHECR analysis, we gained a better understanding of the impact that the amplification and digitization chain has on signal-shape-based parameters, such as the correlation score. Simulations indicate that an inaccurate description of the amplification and digitization response can result in differences of $\mathcal{O}(0.1)$ in the correlation score, highlighting the necessity for a precise characterization. Based on these findings, the pre-installation procedure for characterizing the amplification and digitization chain has been changed. We now have an excellent description of the amplification and

digitization chain through full-chain time response measurements, for which we found that channel-to-channel and station-to-station variations are negligible. The improved understanding and description are of immense value for future analyses that rely on the signal shape.

To facilitate an accurate time-dependent description of the detector for simulation and reconstruction, we developed the hardware and calibration database for RNO-G. This database provides the most up-to-date descriptions of hardware components, as well as station and channel positions, while maintaining records of previous measurements to ensure backward compatibility. Integration of the database readout into NuRadioReco allows for seamless use in simulation and reconstruction.

Moreover, we thoroughly investigated the impact of various systematic effects on the correlation score. We found that the uncertainty in the refractive index near the surface is the dominant uncertainty when using shallow antennas. Based on ongoing studies and new techniques (e.g., see Ref. [200]), it is reasonable to expect an improvement of this uncertainty in the near future. The other significant systematic uncertainties are associated with the antenna response, amplification and digitization response, and the time tuning of the digitizer chips.

As another approach to improve understanding of the refractive index, we introduced a ground-penetrating-radar-based measurement technique and conducted a corresponding proof-of-principle analysis using on-site measurements. The results are promising, however, additional work is required to establish this as a robust method for refractive index studies. This technique offers a complementary approach with the advantages that drilling additional holes is not required and that it can be carried out with minimal on-ground effort. Consequently, it could be performed at every station to measure local refractive indices or repeated at regular time intervals to monitor temporal changes in the upper layers of the ice.

The determining factor for the energy threshold of the UHECR in the present data is the trigger on the Radiant (data acquisition) board. Therefore, we investigated potential improvements to the Radiant board to enhance its trigger performance. The identified components were modified in the latest iteration of the Radiant board, resulting in significant improvements in the trigger performance. In the 2024 season, the updated Radiant board was installed in seven of the eight deployed stations, promising an improved sensitivity to UHECRs in future analyses. Additionally, we developed a procedure to translate the measured trigger threshold on the Radiant board into a value suitable for simulations, which use an approximation of the Radiant trigger. This trigger threshold mapping enables realistic UHECR simulations and facilitates estimation of the expected number of events.

During the 2023 data-taking period, station 23 exhibited an unusual behavior. The trigger rate and thresholds changed despite no modifications to the station settings, resulting in three periods with different sensitivities that required consideration in the analysis. We identified that this behavior is likely driven by the downward-facing channels 12 and 20, although the underlying cause remains unidentified. Further investigations into this unstable behavior are suggested, as it affects the detector sensitivity.

Additionally, we examined data from the stations excluded from the UHECR analysis. The results support the trigger-rate-based decision to exclude these stations, as applying the standard background reduction cuts results in substantial lifetime losses. Nevertheless, we identified some UHECR candidates in the data by inspecting the highest-correlation-score events. This demonstrates that, despite the issues, these stations are capable of triggering on high-SNR UHECR events. The increased noise from the power system and wind turbines, which had previously led to the poor data quality of the stations, was meanwhile mitigated.

Future analyses will benefit from the upgraded station hardware and the resulting larger sample of UHECR events. We expect an increase in the number of detected UHECRs by about an order of magnitude. A larger sample size will, e.g., increase the likelihood of detecting multi-station coincidences. In this work, we conducted a search using all three stations included in the main analysis, but found no coincidences, which aligns with expectations given the sample size. Furthermore, the problem with the misaligned readout window for the deep and shallow detector components has been mitigated. Future analyses will be able to detect coincident signals in the shallow and deep antennas caused by down-going UHECRs, and may also observe both the in-air and in-ice components of an impacting air shower.

The collaboration will continue deploying the currently planned RNO-G detector with 35 stations and will work towards their primary scientific goal of searching for UHEN. Due to the low flux of UHEN, achieving a first detection or setting competitive limits will require several years. Throughout this period, UHECRs will remain important for the experiment: initially to characterize the stations under realistic conditions, and later as a background to be considered in UHEN analyses.

Appendix

Supplementary Figures

A

This appendix provides supplementary figures that support statements from different sections of this thesis. It also includes figures from additional stations in cases where only single-station data is presented in the main text.

Figures A.1, A.2, A.3, A.4, and A.5 present results from the investigation of the unusual behavior observed at station 23 during the 2023 data-taking period. Further details regarding this investigation are provided in Section 5.4.

Figure A.6 shows, separately for each station, the reconstructed properties of the UHECR candidate events identified in this thesis. An explanation of the reconstruction method and a discussion of the results are available in Section 8.6.

Figures A.7, A.8, A.9, A.10, and A.11 are from the analysis of potential background contamination in the UHECR sample. These figures are included for completeness, as they show data from the stations not shown in the main text. A discussion of these plots is provided in Section 8.7.

Figures A.12, A.13, and A.14 show the results from the analysis of data from stations excluded from the main analysis. The plot for station 11 is included in the main text. Details about this investigation are provided in Section 8.8.

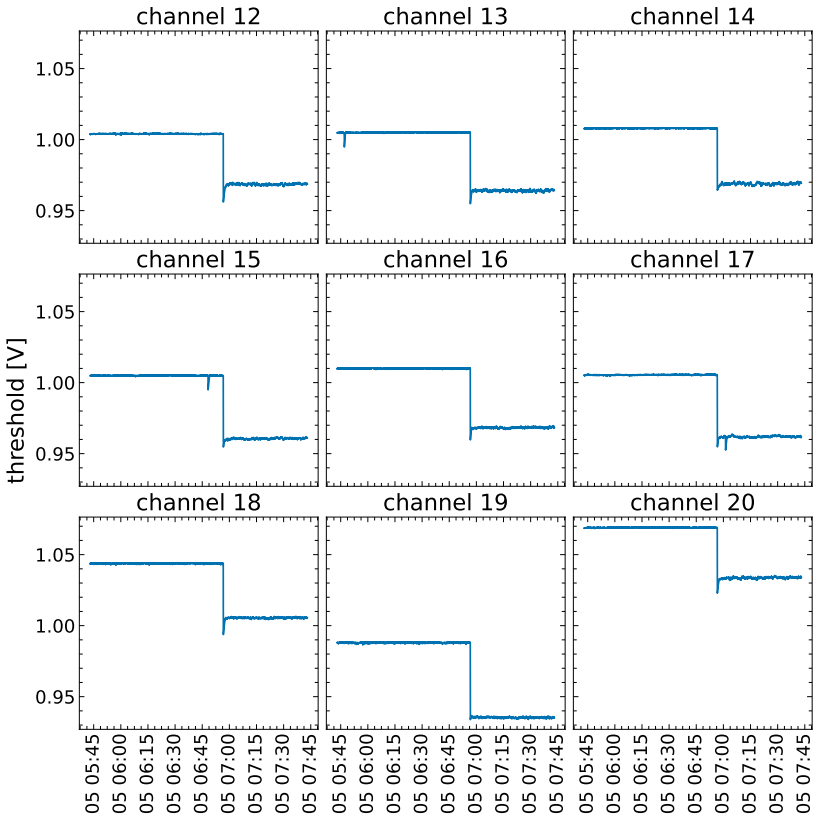


Figure A.1: Radiant thresholds for all shallow channels of station 23 during run 2584. Run 2584 was taken in 2023, and during the run, a sudden increase in the trigger rate was observed.

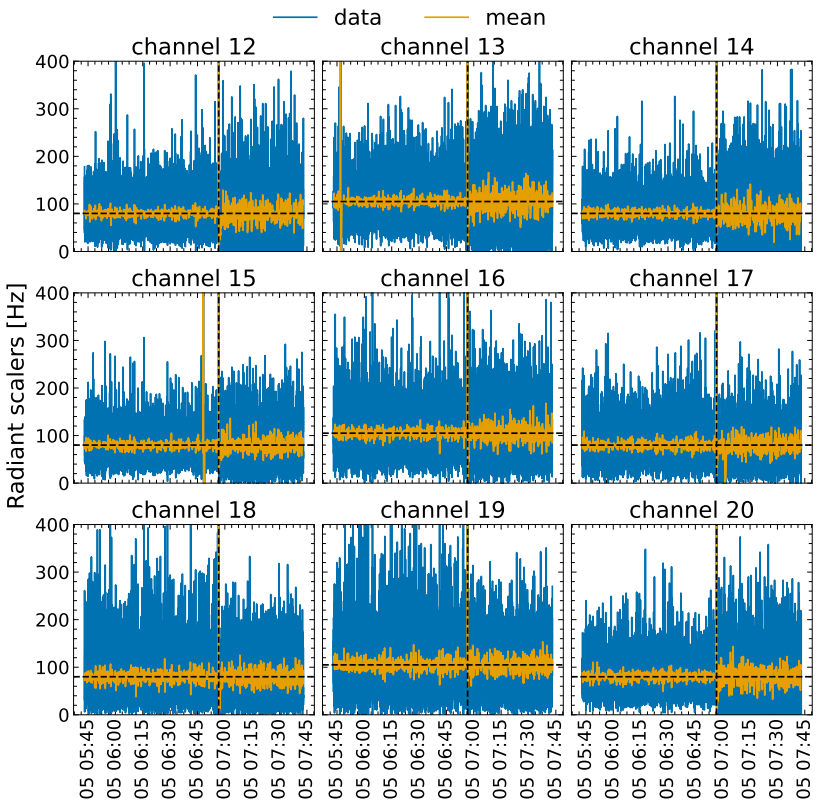


Figure A.2: Radiant scalers for all shallow channels of station 23 during run 2584. Run 2584 was taken in 2023, and during the run, a sudden increase in the trigger rate was observed, marked by the vertical dashed line.

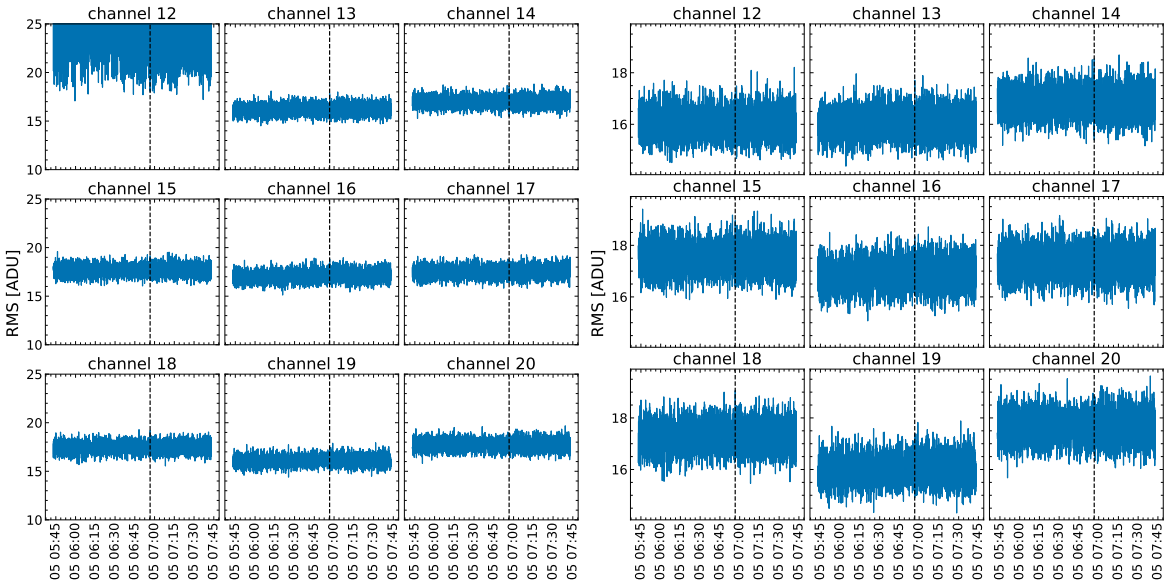


Figure A.3: RMS values for all shallow channels of station 23 during run 2584 before (left) and after (right) block-offset correction. Run 2584 was taken in 2023, and during the run, a sudden increase in the trigger rate was observed, marked by the vertical dashed line.

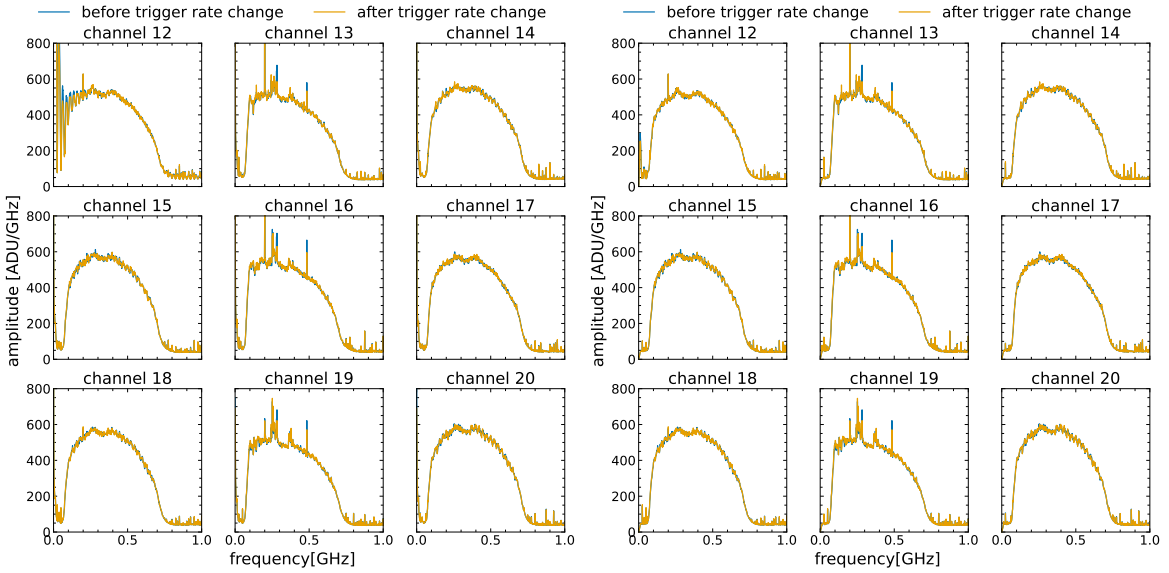


Figure A.4: Average frequency spectrum for all shallow channels of station 23 during run 2584 before (left) and after (right) block-offset correction. Run 2584 was taken in 2023, and during the run, a sudden increase in the trigger rate was observed.

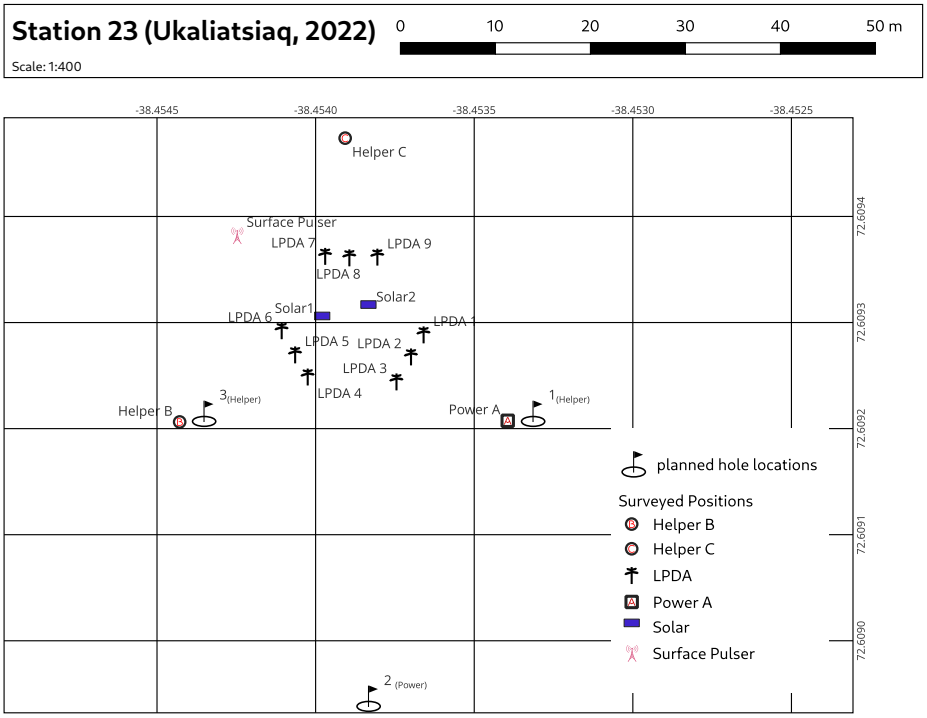


Figure A.5: Position of the different components and shallow channels of station 23. Channels 12 and 20 are denoted as LPDA 9 and 1, respectively. Plot created by Cosmin Deaconu.

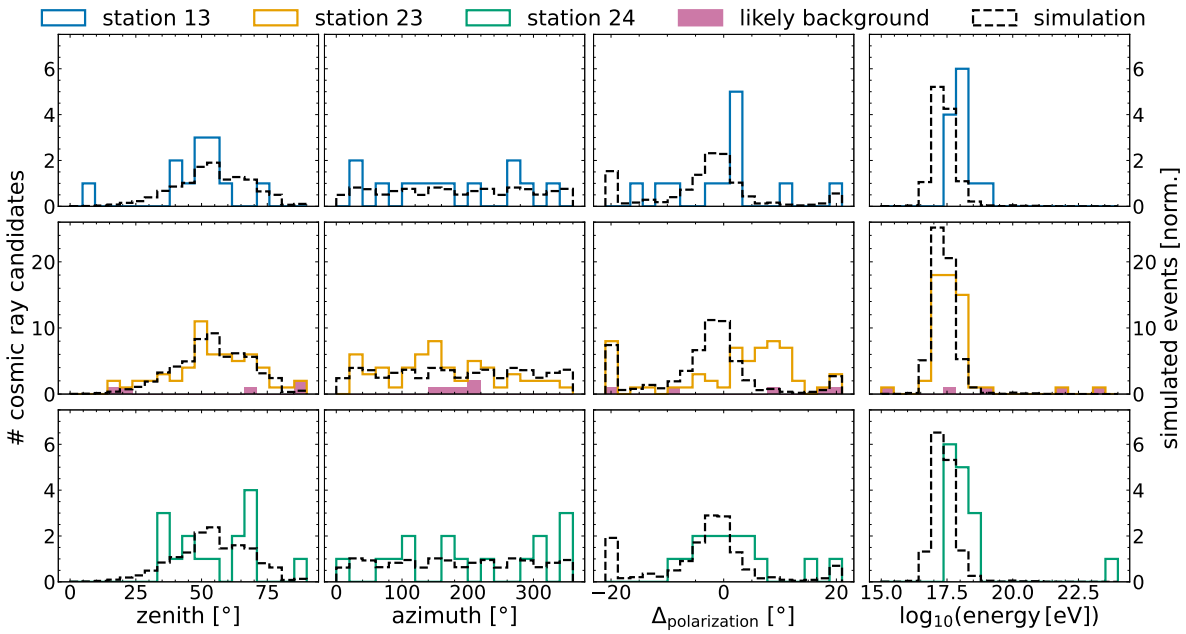


Figure A.6: Reconstructed properties for the UHECR candidates per station, corresponding simulations, which are normalized to the number of UHECR candidate events and likely backgrounds. The likely backgrounds are the five events identified in Section 8.7.3. Shown are from left to right the zenith angle, azimuth angle, difference to expected polarization angles assuming a purely geomagnetic emission (with the first and last bin being overflow bins), and energy. It should be noted that the shown distributions are consistent in all cases after including systematic uncertainties, which are not shown. Additionally, the simulations use on average a lower trigger threshold than realized in real data. See text in Section 8.6 for more details.

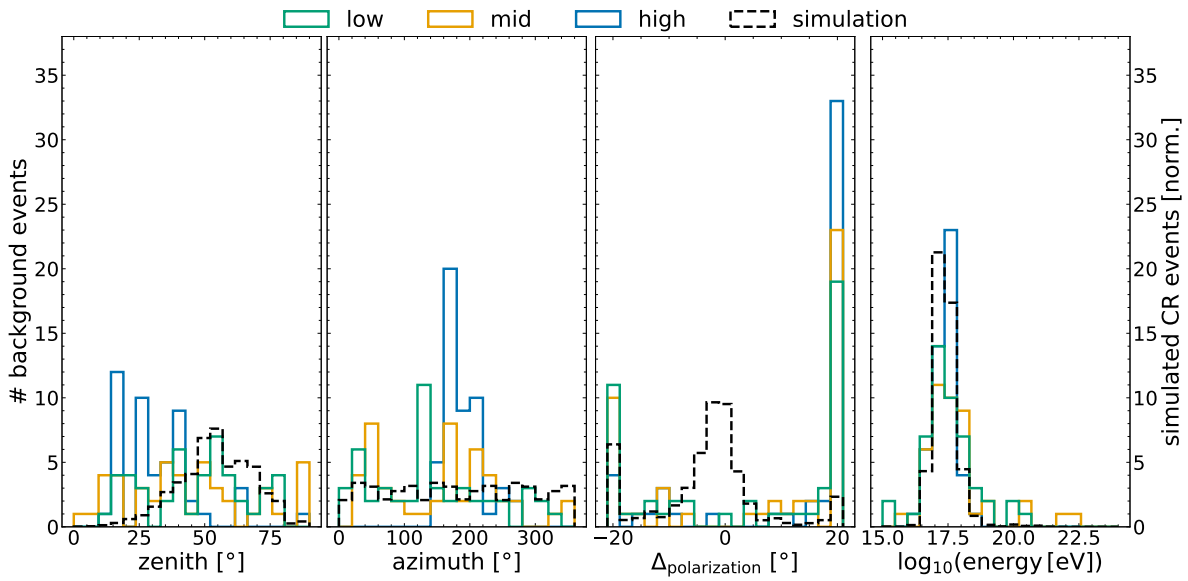


Figure A.7: Reconstructed values for azimuth angle, zenith angle, polarization difference, and energy of the background events from the three correlation-score regions of station 23. See text in Section 8.7.3 for more details.

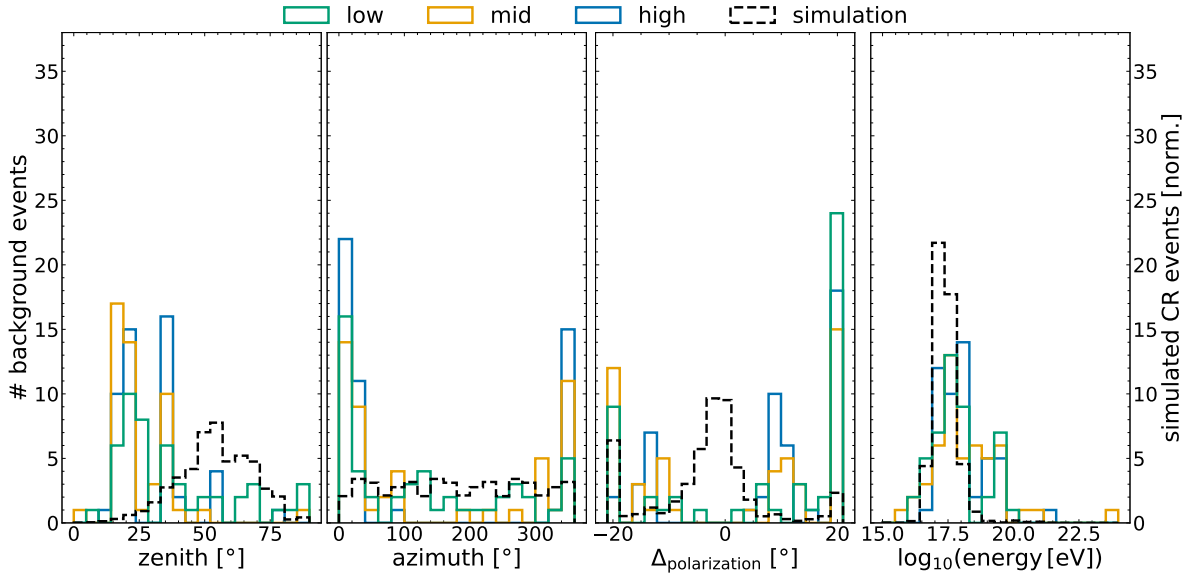


Figure A.8: The same as Figure A.7, but for station 24.

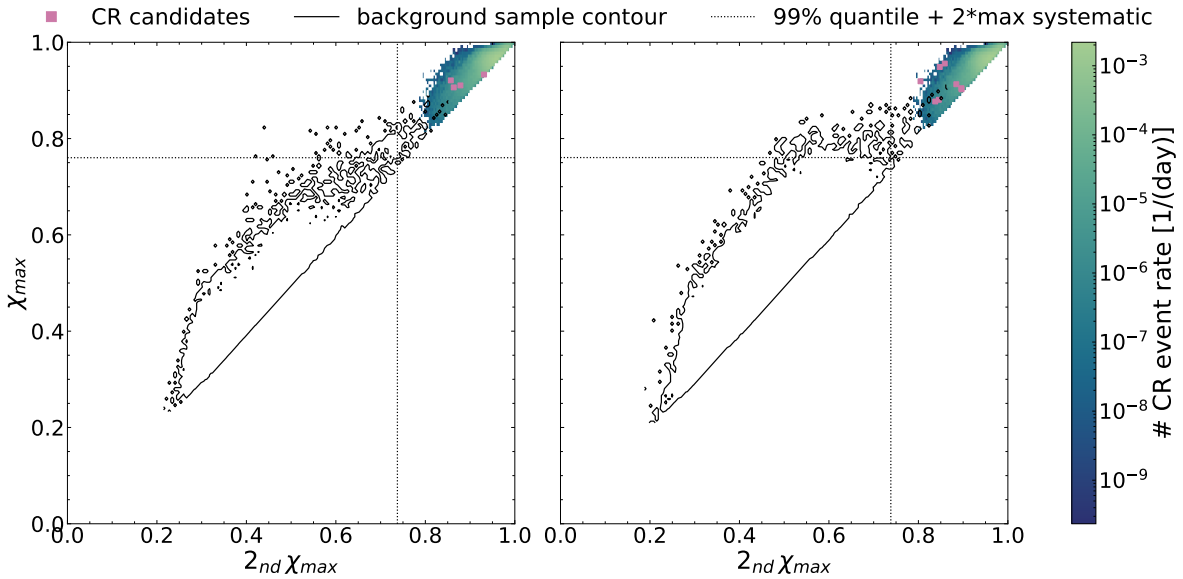


Figure A.9: Correlation value of the channel with the best match as a function of the correlation value of the second-best channel for station 13. Shown are the contours of a background sample in black, the expected distribution of cosmic rays (CRs) without systematic uncertainties (color), and the CR candidates (squares) for the years 2022 (left) and 2023 (right).

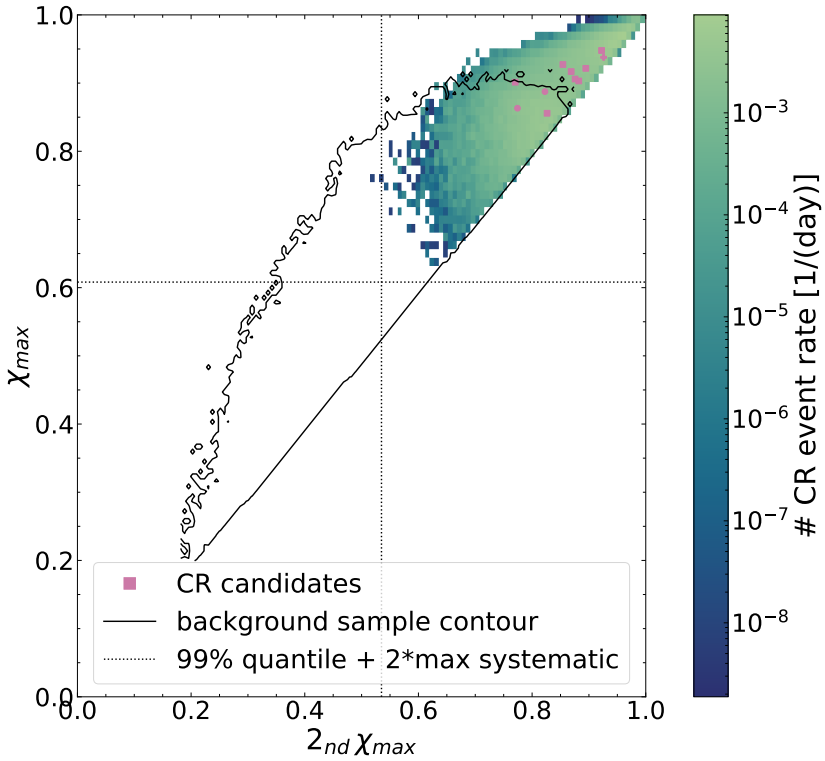


Figure A.10: Correlation value of the channel with the best match as a function of the correlation value of the second-best channel for station 23. Shown are the contours of a background sample in black, the expected distribution of cosmic rays (CRs) without systematic uncertainties (color), and the CR candidates (squares) for the year 2023. Figure 8.41 shows the same for the year 2022.

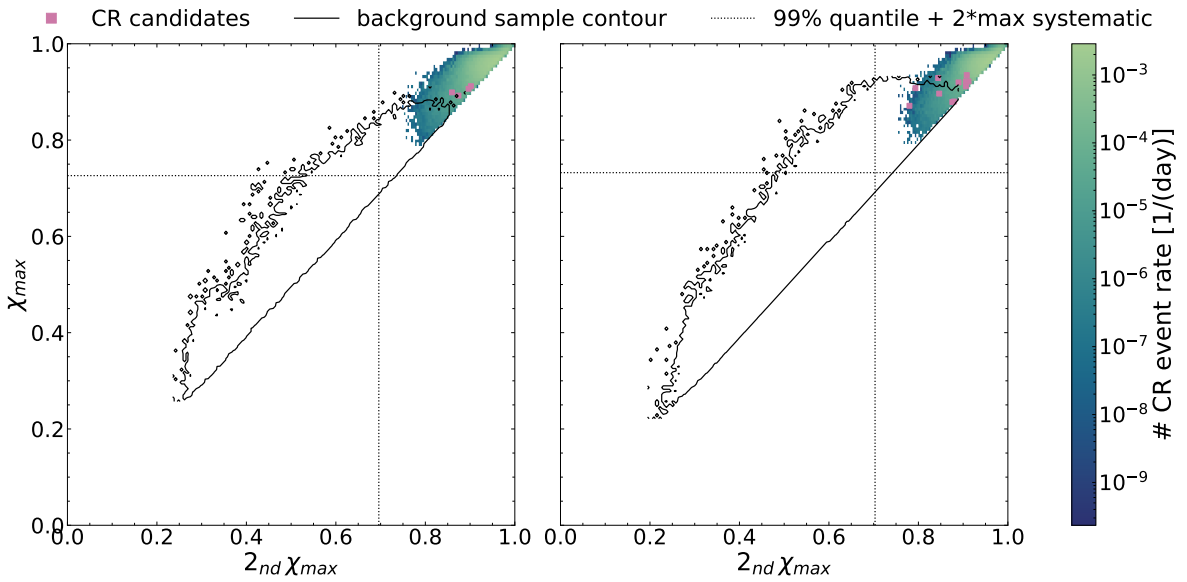


Figure A.11: Correlation value of the channel with the best match as a function of the correlation value of the second-best channel for station 24. Shown are the contours of a background sample in black, the expected distribution of cosmic rays (CRs) without systematic uncertainties (color), and the CR candidates (squares) for the years 2022 (left) and 2023 (right).

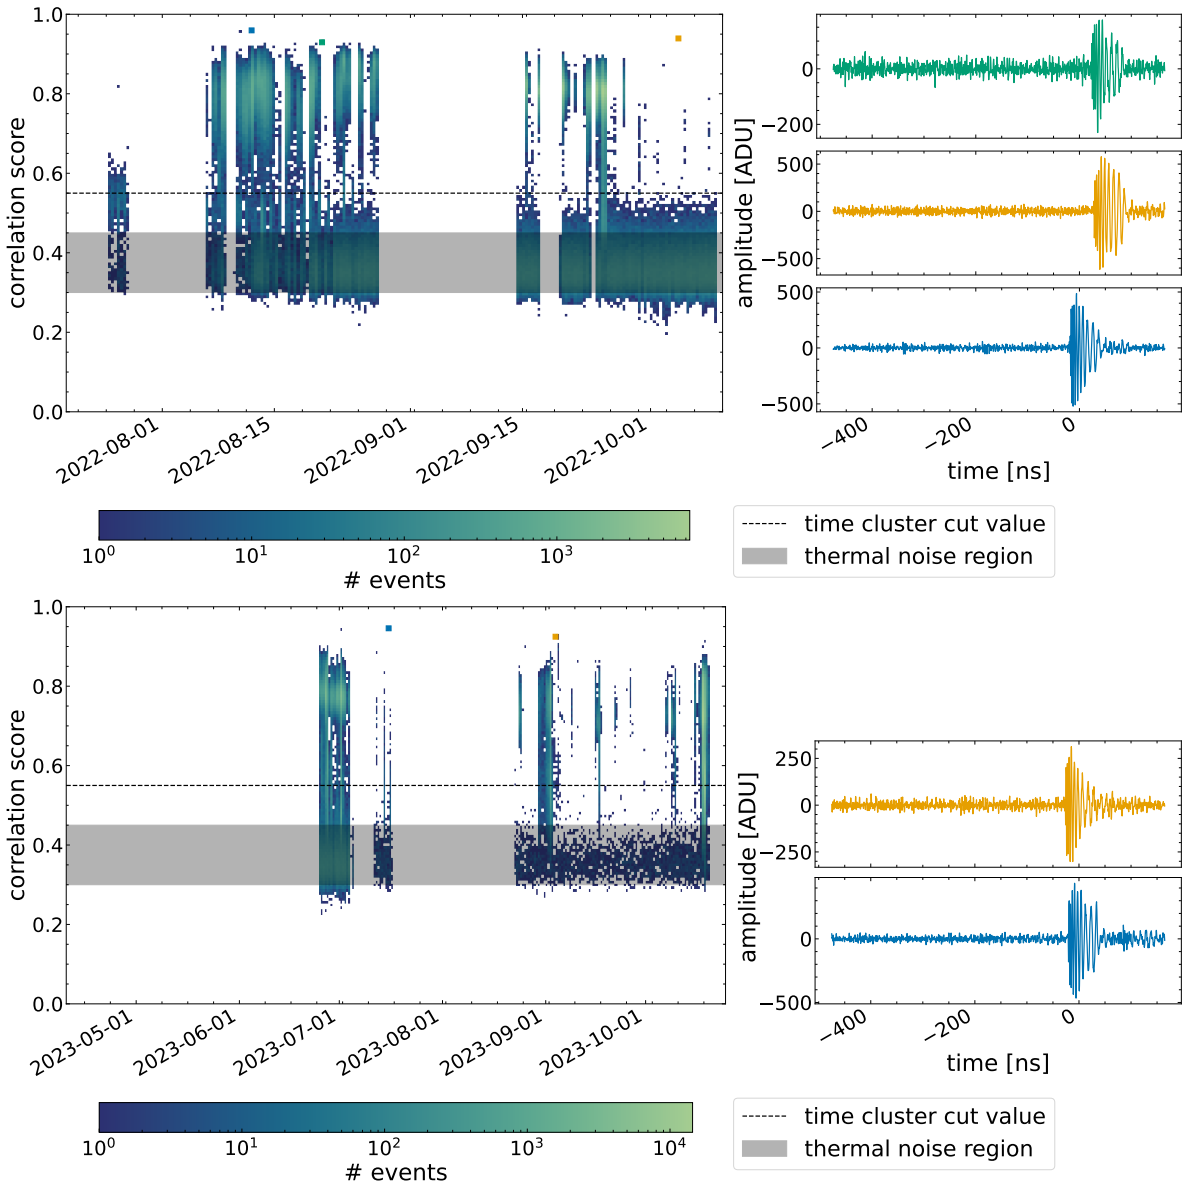


Figure A.12: The correlation score as a function of time for station 12 and years 2022 (top) and 2023 (bottom) without any cuts applied. Additionally shown are three example waveforms from UHECR candidate events, which are also denoted by squares and colors in the left panel. Each shown waveform corresponds to the channel with the highest correlation score per highlighted event, with the color indicating which event it belongs to. In July 2023, changes to the power system were made. Taken from Ref. [185].

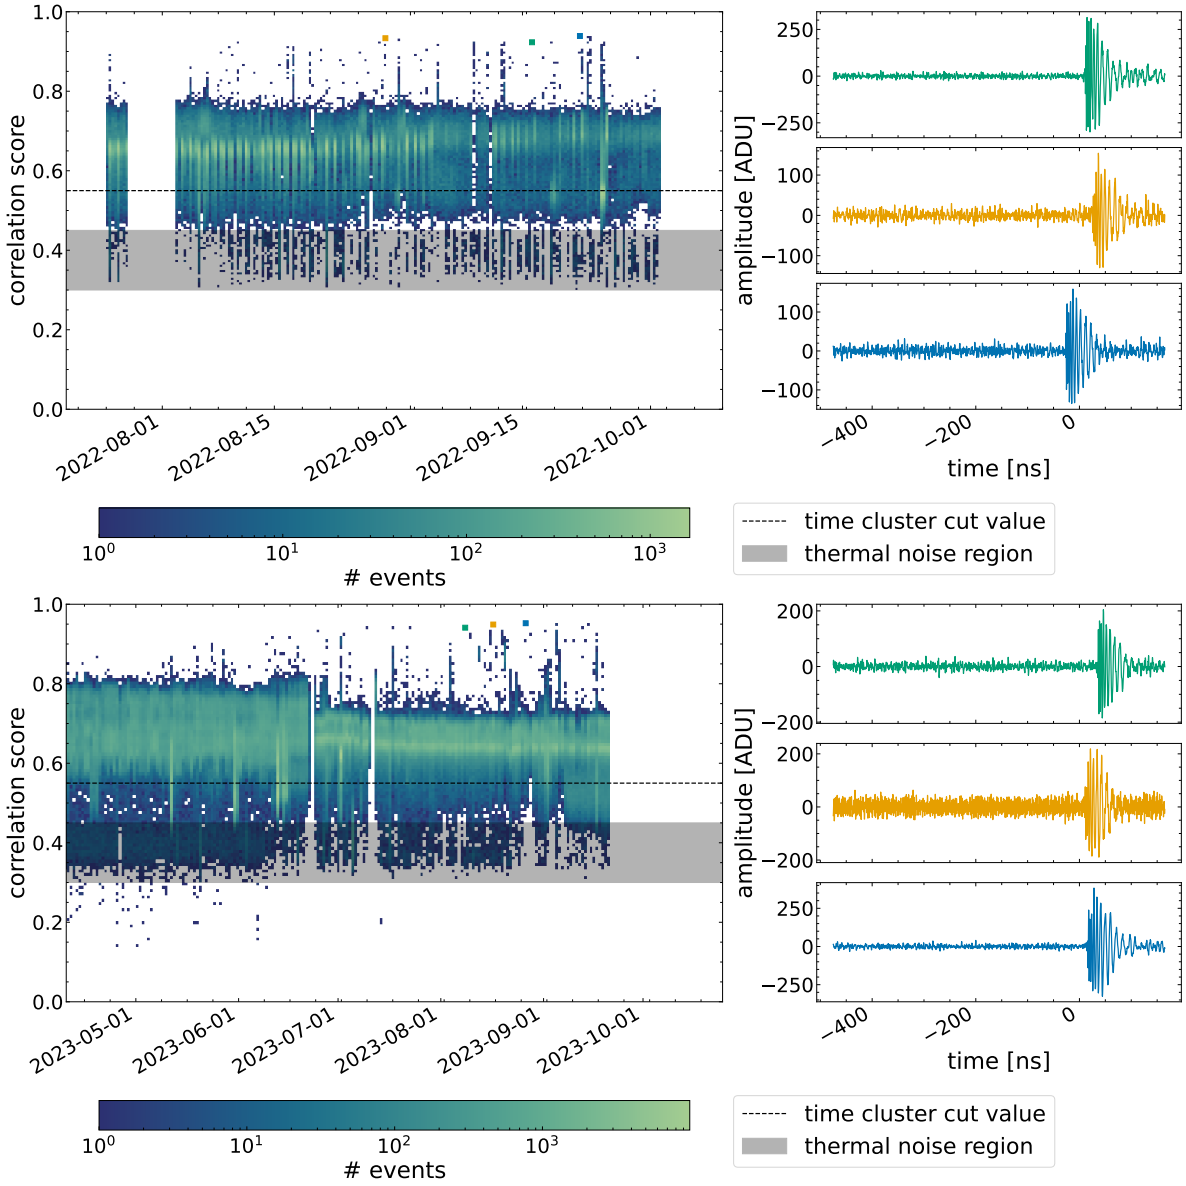


Figure A.13: The same as Figure A.12, but for station 21. Taken from Ref. [185].

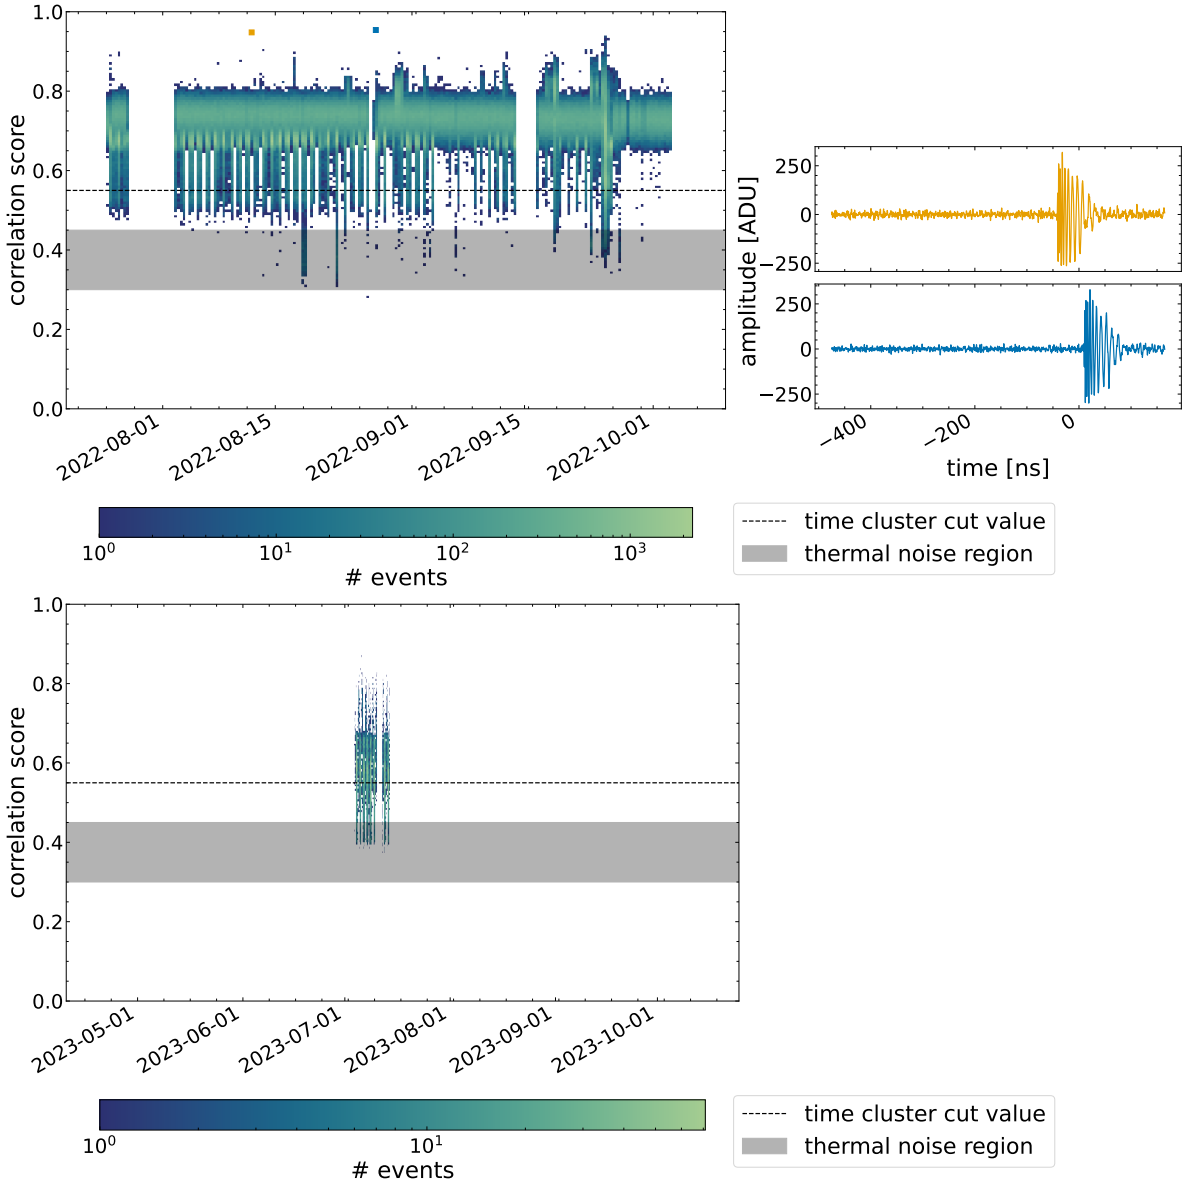


Figure A.14: The same as Figure A.12, but for station 22. Taken from Ref. [185].

List of Figures

2.1 Cosmic-ray energy spectrum. 4
2.2 Schematic of an extensive air shower and possible detection techniques. . . 10
2.3 Lateral and longitudinal profile of an extensive air shower. 11
2.4 Measurements of the shower maximum X_{max}. 13
2.5 Feynman diagram of CC and NC neutrino-nucleus interaction. 15
2.6 Measured γ-ray, neutrino and cosmic-ray flux: the multimessenger picture. 18

3.1 Schematic of the geomagnetic and Askaryan effect. 22
3.2 Simulated Askaryan emission for different viewing angles. 25
3.3 Simulated radio footprints. 26
3.4 Upper limits and sensitivities to a diffuse neutrino flux. 29

4.1 Map of Greenland and RNO-G layout. 32
4.2 Schematic of an RNO-G station. 33
4.3 Simulated signal at different stages of the signal chain. 34
4.4 Simulated VEL for different radio antennas. 35
4.5 Radio antennas used by RNO-G 36
4.6 Schematic of the signal chain used for the deep channels. 38
4.7 Schematic of the signal chain used for the shallow channels. 39

5.1 Impact of the hardware response on a time domain pulse. 44
5.2 Illustration of S parameter definition. 45
5.3 Impact of an inaccurate hardware response on UHECR simulations. 46
5.4 Comparison of different hardware responses. 47
5.5 Template-UHECR waveform comparison for different hardware responses. 48
5.6 Structure of the RNO-G calibration and hardware database. 53
5.7 Structure of the component collections. 55
5.8 Structure of the calibration collections. 57
5.9 Structure of the main collection. 59
5.10 Screenshot of the web interface home page. 61
5.11 Screenshot of a page for adding S-parameters to the database. 61
5.12 Noise trigger rates of a V2 Radiant board. 66
5.13 Schematic to illustrate the board noise-dominated Radiant trigger. 67
5.14 Noise trigger rates of a V3 Radiant board. 68
5.15 Schematic of the setup for trigger efficiency measurements. 69
5.16 Measured trigger efficiency curves and SNR to amplitude mapping. 70
5.17 Radiant circuit diagram. 74
5.18 Trigger performance after trigger path modifications. 75
5.19 Measured trigger threshold to SNR($\epsilon_{\mathrm{trig}} = 50\%$) mapping. 76
5.20 Comparison of the input and output of the arbitrary waveform generator. . 77
5.21 Simulated trigger efficiency curves and threshold to SNR($\epsilon_{\mathrm{trig}} = 50\%$) mapping. 78
5.22 Schematic of the process to estimate the simulation trigger threshold. 79
5.23 Trigger rate of station 23 for the 2022 and 2023 data taking period. 80
5.24 Radiant thresholds of station 23 in August 2023. 81

5.25 Trigger rate of station 23 for a short time period in August 2023. 82
5.26 Investigation of the station 23 behavior by changing the trigger settings. . . 83

6.1 Refractive index for different ice models. 90
6.2 Effective VEL description for different refractive indices. 91
6.3 Impact of the refractive index on UHECR simulations. 93
6.4 GPR image of IceCube. 94
6.5 Schematic of a dynamic GPR measurement. 95
6.6 Image of a dynamic GPR measurement in Greenland. 96
6.7 GPS track of a dynamic GPR measurement. 97
6.8 Schematic of a stationary GPR measurement. 98
6.9 Spectrogram of a stationary GPR measurement. 99
6.10 Spectrogram of a dynamic GPR measurement. 100
6.11 Spectrogram of a dynamic GPR measurement including distance to borehole. 101
6.12 Fit of hyperbolas in GPR spectrogram. 102
6.13 Measured and simulated average refractive index. 103

7.1 UHECR parameter space for simulations. 108
7.2 Comparison of the database response and the scaled impulse response. . . 109
7.3 Trigger thresholds of station 13 for the 2022 data-taking period. 111
7.4 Frequency spectrum of simulated thermal and Galactic noise. 114
7.5 UHECR flux published by the Pierre Auger Observatory. 115

8.1 Signal and background template waveforms. 120
8.2 Illustration of template generation. 121
8.3 Result of a scan over Gaussian widths for template generation. 122
8.4 Result of a scan over E-field amplitudes for template generation. 123
8.5 Noiseless UHECR waveforms correlated with different template sets. 124
8.6 Trigger rates for all stations during the 2022 data-taking period. 127
8.7 Trigger rates for all stations during the 2023 data-taking period. 127
8.8 Comparison of simulated and modified VEL for systematic uncertainty study. 131
8.9 Influence of simulated poor digitizer time tuning on the correlation score. . 134
8.10 Influence of measured poor digitizer time tuning on the correlation score. . 135
8.11 Temperatures inside the DAQ boxes. 136
8.12 Influence of the temperature on the impulse response. 137
8.13 Influence of statistical fluctuations on the UHECR analysis. 138
8.14 Influence of slightly different SNR values on the correlation score cut. . . . 139
8.15 Influence of the amplifier saturation on the UHECR analysis. 140
8.16 Influence of the antenna position on the UHECR analysis. 141
8.17 Influence of station-to-station variations on the UHECR analysis. 143
8.18 Influence of the Galactic noise on the UHECR analysis. 144
8.19 Influence of the noise temperature on the UHECR analysis. 145
8.20 Schematic overview of the analysis and simulation pipeline. 146
8.21 Illustration of the time cluster cut. 148
8.22 Correlation score distributions of forced trigger data. 149
8.23 htRMS distributions of forced trigger data. 151
8.24 Definition of a time-dependent htRMS cut. 152
8.25 Example UHECR simulation with SNR-dependent correlation score cut. . . 153

8.26 Selection of the 2022 UHECR candidate events. 154
8.27 Selection of the 2023 UHECR candidate events. 155
8.28 Example UHECR candidate waveforms (2022). 157
8.29 Example UHECR candidate waveforms (2023). 157
8.30 Reconstructed quantities for the UHECR candidate events. 159
8.31 Regions for the ABCD method. 162
8.32 Background estimate from the ABCD method for different region definitions. 163
8.33 Investigation of thermal noise as possible background contamination. . . . 164
8.34 UHECR candidates as a function of time of day. 165
8.35 UHECR candidates as a function of the wind speed. 166
8.36 Impact of the wind speed on the data. 167
8.37 Airplane cross check with the SNR as a metric. 168
8.38 Comparison of reconstructed UHECR and airplane directions. 169
8.39 UHECR candidates as a function of the battery voltage. 170
8.40 Reconstruction of background events. 171
8.41 Largest vs second-largest correlation score for station 23 (2022). 172
8.42 Correlation score as a function of time for station 11. 175

A.1 Trigger thresholds of station 23 during run 2584. 186
A.2 RADIANT scalers of station 23 during run 2584. 186
A.3 RMS values of station 23 during run 2584. 187
A.4 Average frequency spectrum of station 23 during run 2584. 187
A.5 Map of station 23. 188
A.6 Reconstructed properties for the UHECR candidates per station. 188
A.7 Reconstructed values for background events from station 23. 189
A.8 Reconstructed values for background events from station 24. 189
A.9 Largest vs second-largest correlation score for station 13. 190
A.10 Largest vs second-largest correlation score for station 23 (2023). 190
A.11 Largest vs second-largest correlation score for station 24. 191
A.12 Correlation score as a function of time for station 12. 192
A.13 Correlation score as a function of time for station 21. 193
A.14 Correlation score as a function of time for station 22. 194

List of Tables

5.1 Impact of the hardware responses on the template-UHECR correlation score. 50
5.2 Influence of the gain and group delay on the correlation score. 51
5.3 Mapping between device name and IDs. 60

7.1 Measured and simulated trigger thresholds for station 13. 111

8.1 Final parameters for template generation. 124
8.2 Mean correlation score for simulations with different template sets. 125
8.3 Start/end dates, lifetime and number of triggered events for the dataset. . . 129
8.4 Summary of all considered systematic effects. 132
8.5 Number of events remaining after each step of the UHECR analysis. 146
8.6 Lifetime losses and UHECR dead time from the time cluster cut. 150
8.7 Number of selected UHECR candidate events per station and year. 154
8.8 Impact of systematic uncertainties on the reconstructed polarization. 159

Special Terms

A
ADC analog-to-digital converter.
ADU The ADU is the unit of the ADC counts.
AERA Auger Engineering Radio Array.
AGN active galactic nucleus.
AMANDA Antarctic Muon And Neutrino Detector Array.
ANITA Antarctic Impulsive Transient Antenna.
ANTARES Astronomy with a Neutrino Telescope and Abyss enviromental RESearch.
ARA Askaryan Radio Array.
ARIANNA Antarctic Ross Ice-Shelf ANtenna Neutrino Array.

B
BEACON Beamforming Elevated Array for COsmic Neutrinos.

C
CC charged current.
CMB cosmic microwave background.
CR cosmic ray.
CW continuous wave.

D
DAQ data acquisition.
DISC Deep Ice Sheet Coring.
Drab The Downhole Receiver & Amplifier Board (DRAB) is the amplifier board used in RNO-G to receive and amplify the signal from the IGLU board.
DUMAND Deep Underwater Muon and Neutrino Detector.

E
EAS extensive air shower.

F
FFT Fast Fourier Transform.
Flower The Flexible Octal WavEform Recorder (FLOWER) board is an additional trigger board used in RNO-G, which makes the trigger decision for the four channels of the phased array.

G
GPR ground-penetrating radar.
GPS Global Positioning System.
GRAND Giant Radio Array for Neutrino Detection.
GRB gamma-ray burst.
GZK Greisen-Zatsepin-Kuzmin.

H
HERON Hybrid Elevated Radio Observatory for Neutrinos.
htRMS half trace root-mean-square.

HUNT Huge Underwater high-energy Neutrino Telescope.

I

IGLU The In-ice Gain with Low-power Unit (IGLU) board is the amplifier board used in RNO-G to amplify the signal before sending it to the surface through the optical fiber.

L

LHC Large Hadron Collider.
LOFAR LOw-Frequency ARray.
LOPES LOFAR Prototype Station.
LoRaWAN Long Range Wide Area Network.
LPDA Log-Periodic Dipole Array.
LPM Landau-Pomeranchuk-Migdal.

N

NC neutral current.
NuRadioMC NuRadioMC is a Monte Carlo simulation framework developed for UHECR and UHEN radio detectors.
NuRadioReco NuRadioReco is a detector simulation and reconstruction framework developed for UHECR and UHEN radio detectors.

P

PCB printed circuit board.
peak-to-peak The peak-to-peak value describes the difference between the highest and lowest value of a waveform.
POEMMA Probe Of Extreme Multi-Messenger Astrophysics.
P-ONE Pacific Ocean Neutrino Telescope.
PPS pulse-per-second.
PUEO Payload for Ultrahigh Energy Observations.

R

Radiant The RAdio DIgitizer and Auxiliary Neutrino Trigger (RADIANT) a digitizer and trigger board used for the RNO-G experiment.
Radiant trigger The name "Radiant trigger" refers to the trigger circuit that is integrated on all Radiant channels.
RET-CR Radar Echo Telescope for Cosmic Rays.
RET-N Radar Echo Telescope for Neutrinos.
RFoF Radio-frequency-over-fiber.
RICE Radio Ice Cherenkov Experiment.
RMS root-mean-square.
RNO-G Radio Neutrino Observatory in Greenland.

S

SCA switch-capacitor-array.
shallow trigger The shallow trigger is a name for Radiant trigger if it is used for the shallow LPDAs.
SKA Square Kilometer Array.
SNR signal-to-noise ratio.
SOF Station Operation Facility.

Summit Station Research station on the greenlandic ice sheet at which the RNO-G experiment is located.

Surface The Surface board is the amplifier board used for the shallow channels in RNO-G.

T

TA Telescope Array.

TAMBO Tau Air-shower Mountain-Based Observatory.

TAROGE Taiwan Astroparticle Radiowave Observatory for Geo-synchrotron Emissions.

TDE tidal disruption event.

TOF time-of-flight.

TRIDENT TRopIcal DEep-sea Neutrino Telescope.

Tunka-Rex Tunka Radio extension.

U

UHECR ultra-high-energy cosmic ray.

UHEN ultra-high-energy neutrino.

V

VEL vector effective length.

VNA Vector Network Analyzer.

W

WIPL-D WIPL-D is a electromagnetic simulation software which is used to simulate antenna responses.

X

XFdtd XFdtd is a electromagnetic simulation software which is used to simulate antenna responses.

References

[1] M. G. Aartsen et al. (IceCube), *"Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert"*, Science **361**, 147–151 (2018).

[2] R. Abbasi et al. (IceCube), *"Evidence for neutrino emission from the nearby active galaxy NGC 1068"*, Science **378**, 538–543 (2022).

[3] R. Abbasi et al. (IceCube), *"Observation of high-energy neutrinos from the Galactic plane"*, Science **380**, adc9818 (2023).

[4] R. Abbasi et al. (IceCube), *"Evidence for a Spectral Break or Curvature in the Spectrum of Astrophysical Neutrinos from 5 TeV–10 PeV"*, Accepted at PRL (2025).

[5] S. Aiello et al. (KM3NeT), *"Observation of an ultra-high-energy cosmic neutrino with KM3NeT"*, Nature **638**, [Erratum: Nature 640, E3 (2025)], 376–382 (2025).

[6] R. Abbasi et al. (IceCube), *"Search for Extremely-High-Energy Neutrinos and First Constraints on the Ultrahigh-Energy Cosmic-Ray Proton Fraction with IceCube"*, Phys. Rev. Lett. **135**, 031001 (2025).

[7] A. Aab et al. (Pierre Auger), *"Probing the origin of ultra-high-energy cosmic rays with neutrinos in the EeV energy range using the Pierre Auger Observatory"*, JCAP **10**, 022 (2019).

[8] G. A. Askar'yan, *"Excess Negative Charge of an Electron-Photon Shower and its Coherent Radio Emission"*, Zh. Eksp. Teor. Fiz. **41**, 616–618 (1961).

[9] J. A. Aguilar et al. (RNO-G), *"In situ, broadband measurement of the radio frequency attenuation length at Summit Station, Greenland"*, J. Glaciol. **68**, 1234–1242 (2022).

[10] R. Abbasi et al. (IceCube), *"In situ estimation of ice crystal properties at the South Pole using LED calibration data from the IceCube Neutrino Observatory"*, The Cryosphere **18**, 75–102 (2024).

[11] J. A. Aguilar et al. (ANTARES), *"Transmission of light in deep sea water at the site of the ANTARES Neutrino Telescope"*, Astropart. Phys. **23**, 131–155 (2005).

[12] P. Allison et al. (ARA), *"Low-threshold ultrahigh-energy neutrino search with the Askaryan Radio Array"*, Phys. Rev. D **105**, 122006 (2022).

[13] A. Anker et al. (ARIANNA), *"A search for cosmogenic neutrinos with the ARIANNA test bed using 4.5 years of data"*, JCAP **03**, 053 (2020).

[14] J. Aguilar et al. (RNO-G), *"Design and sensitivity of the Radio Neutrino Observatory in Greenland (RNO-G)"*, JINST **16**, P03025 (2021).

[15] M. G. Aartsen et al. (IceCube-Gen2), *"IceCube-Gen2: the window to the extreme Universe"*, J. Phys. G **48**, 060501 (2021).

[16] F. D. Kahn and I. Lerche, *"Radiation from cosmic ray air showers."*, Proc. R. Soc. Lond. A **289**, 206–213 (1966).

[17] S. W. Barwick et al. (ARIANNA), *"Radio detection of air showers with the ARIANNA experiment on the Ross Ice Shelf"*, Astropart. Phys. **90**, 50–68 (2017).

[18] L. Pyras et al., *"Atmospheric muons at PeV energies in radio neutrino detectors"*, JCAP **10**, 043 (2023).

[19] K. D. de Vries et al., *"The cosmic-ray air-shower signal in Askaryan radio detectors"*, Astropart. Phys. **74**, 96–104 (2016).

[20] A. Coleman et al., *"In-ice Askaryan emission from air showers: Implications for radio neutrino detectors"*, Astropart. Phys. **172**, 103136 (2025).

[21] S. De Kockere et al., *"Simulation of in-ice cosmic ray air shower induced particle cascades"*, Phys. Rev. D **106**, 043023 (2022).

[22] N. Alden et al. (ARA), *"Observation of In-ice Askaryan Radiation from High-Energy Cosmic Rays"*, Phys. Rev. Lett., – (2026).

[23] V. F. Hess, *"Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten"*, Phys. Z. **13**, 1084–1091 (1912).

[24] S. Navas et al. (Particle Data Group), *"Review of particle physics"*, Phys. Rev. D **110**, 030001 (2024).

[25] T. K. Gaisser, R. Engel, and E. Resconi, Cosmic Rays and Particle Physics: 2nd Edition (Cambridge University Press, June 2016).

[26] P. Sreekumar et al., *"Constraints on the cosmic rays in the Small Magellanic Cloud"*, Phys. Rev. Lett. **70**, 127–129 (1993).

[27] E. Fermi, *"On the Origin of the Cosmic Radiation"*, Phys. Rev. **75**, 1169–1174 (1949).

[28] A. Coleman et al., *"Ultra high energy cosmic rays The intersection of the Cosmic and Energy Frontiers"*, Astropart. Phys. **149**, 102819 (2023).

[29] A. R. Bell, *"The acceleration of cosmic rays in shock fronts – I"*, Monthly Notices of the Royal Astronomical Society **182**, 147–156 (1978).

[30] L. O. Drury, *"An introduction to the theory of diffusive shock acceleration of energetic particles in tenuous plasmas"*, Reports on Progress in Physics **46**, 973 (1983).

[31] D. Caprioli, *""Espresso" Acceleration of Ultra-high-energy Cosmic Rays"*, Astrophys. J. Lett. **811**, L38 (2015).

[32] R. Mbarek and D. Caprioli, *"Espresso Acceleration of Ultra-High-Energy Cosmic Rays up to the Hillas Limit in Relativistic MHD Jets"*, PoS **ICRC2019**, 348 (2020).

[33] E. G. Zweibel and M. Yamada, *"Magnetic Reconnection in Astrophysical and Laboratory Plasmas"*, Annual Review of Astronomy and Astrophysics **47**, 291–332 (2009).

[34] P. Goldreich and W. H. Julian, *"Pulsar electrodynamics"*, Astrophys. J. **157**, 869 (1969).

[35] A. R. Bell, *"Cosmic ray acceleration in pulsar-driven supernova remnants"*, Monthly Notices of the Royal Astronomical Society **257**, 493–500 (1992).

[36] A. M. Hillas, *"The Origin of Ultra-High-Energy Cosmic Rays"*, Annual Review of Astronomy and Astrophysics **22**, 425–444 (1984).

[37] K. Greisen, *"End to the Cosmic-Ray Spectrum?"*, Phys. Rev. Lett. **16**, 748–750 (1966).

[38] G. T. Zatsepin and V. A. Kuzmin, *"Upper limit of the spectrum of cosmic rays"*, JETP Lett. **4**, 78–80 (1966).

[39] V. Beresinsky and G. Zatsepin, *"Cosmic rays at ultra high energies (neutrino?)"*, Physics Letters B **28**, 423–424 (1969).

[40] A. Kounine, *"The Alpha Magnetic Spectrometer on the International Space Station"*, Int. J. Mod. Phys. E **21**, 1230005 (2012).

[41] O. Adriani et al. (PAMELA), *"Ten years of PAMELA in space"*, Riv. Nuovo Cim. **40**, 473–522 (2017).

[42] J. Chang et al. (DAMPE), *"The DArk Matter Particle Explorer mission"*, Astropart. Phys. **95**, 6–24 (2017).

[43] S. Torii and P. S. Marrocchesi (CALET), *"The CALorimetric Electron Telescope (CALET) on the International Space Station"*, Adv. Space Res. **64**, 2531–2537 (2019).

[44] E. Atkin et al., *"The NUCLEON experiment. Results of the first year of data acquisition"*, Astropart. Phys. **90**, 69–74 (2017).

[45] E. S. Seo et al. (CREAM), *"CREAM: 70 days of flight from 2 launches in Antarctica"*, Adv. Space Res. **42**, 1656–1663 (2008).

[46] T. G. Guzik et al. (ATIC), *"The ATIC long duration balloon project"*, Advances in Space Research **33**, 1763–1770 (2004).

[47] T. Aramaki et al. (GAPS), *"Antideuteron Sensitivity for the GAPS Experiment"*, Astropart. Phys. **74**, 6–13 (2016).

[48] H. Fuke et al. (BESS), *"Search for cosmic-ray antideuterons"*, Phys. Rev. Lett. **95**, 081101 (2005).

[49] B. Alpat et al. (AMS), *"Charge determination of nuclei with the AMS-02 silicon tracker"*, Nucl. Instrum. Meth. A **540**, 121–130 (2005).

[50] O. Adriani et al., *"The magnetic spectrometer of the PAMELA satellite experiment"*, Nucl. Instrum. Meth. A **511**, 72–75 (2003).

[51] M. Boezio et al., *"A High granularity imaging calorimeter for cosmic ray physics"*, Nucl. Instrum. Meth. A **487**, 407–422 (2002).

[52] C. Adloff et al. (AMS), *"The AMS-02 lead-scintillating fibres Electromagnetic Calorimeter"*, Nucl. Instrum. Meth. A **714**, 147–154 (2013).

[53] A. Aab et al. (Pierre Auger), *"Measurement of the cosmic-ray energy spectrum above* $2.5{\times}10^{18}$ *eV using the Pierre Auger Observatory"*, Phys. Rev. D **102**, 062005 (2020).

[54] A. Abdul Halim et al. (Pierre Auger), *"A Catalog of the Highest-Energy Cosmic Rays Recorded during Phase I of Operation of the Pierre Auger Observatory"*, Astrophys. J. Suppl. **264**, 50 (2023).

[55] D. J. Bird et al. (HIRES), *"Detection of a cosmic ray with measured energy well beyond the expected spectral cutoff due to cosmic microwave radiation"*, Astrophys. J. **441**, 144–150 (1995).

[56] N. Hayashida et al., *"Observation of a very energetic cosmic ray well beyond the predicted 2.7-K cutoff in the primary energy spectrum"*, Phys. Rev. Lett. **73**, 3491–3494 (1994).

[57] M. Chikawa et al., *"Energy spectrum above 3 x* 10^{18} *EV observed with AGASA"*, in International cosmic ray conference, Vol. 1, International Cosmic Ray Conference (Jan. 2001), p. 333.

[58] R. U. Abbasi et al. (Telescope Array), *"An extremely energetic cosmic ray observed by a surface detector array"*, Science **382**, abo5095 (2023).

[59] F. G. Schröder, *"Radio detection of Cosmic-Ray Air Showers and High-Energy Neutrinos"*, Prog. Part. Nucl. Phys. **93**, 1–68 (2017).

[60] A. Aab et al. (Pierre Auger), *"The Pierre Auger Cosmic Ray Observatory"*, Nucl. Instrum. Meth. A **798**, 172–213 (2015).

[61] A. Abdul Halim et al. (Pierre Auger), *"AugerPrime: Status and first results"*, PoS **ICRC2025**, 385 (2025).

[62] T. Abu-Zayyad et al. (Telescope Array), *"The surface detector array of the Telescope Array experiment"*, Nucl. Instrum. Meth. A **689**, 87–97 (2013).

[63] R. U. Abbasi et al. (Telescope Array), *"Surface detectors of the TAx4 experiment"*, Nucl. Instrum. Meth. A **1019**, 165726 (2021).

[64] N. M. Budnev et al. (TAIGA), *"TAIGA the Tunka Advanced Instrument for cosmic ray physics and Gamma Astronomy - present status and perspectives."*, JINST **9**, C09021 (2014).

[65] F. G. Schröder et al., *"Tunka-Rex: a Radio Extension of the Tunka Experiment"*, J. Phys. Conf. Ser. **409**, 012076 (2013).

[66] R. Engel, D. Heck, and T. Pierog, *"Extensive air showers and hadronic interactions at high energy"*, Ann. Rev. Nucl. Part. Sci. **61**, 467–489 (2011).

[67] D. Heck et al., *"CORSIKA: A Monte Carlo code to simulate extensive air showers"*, Wissenschaftliche Berichte **FZKA-6019**, 1–90 (1998).

[68] F. Riehn et al., *"Hadronic interaction model Sibyll 2.3d and extensive air showers"*, Phys. Rev. D **102**, 063002 (2020).

[69] S. Ostapchenko, *"QGSJET-III model: physics and preliminary results"*, EPJ Web Conf. **208**, 11001 (2019).

[70] T. Pierog et al., *"EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider"*, Phys. Rev. C **92**, 034906 (2015).

[71] J. Albrecht et al., *"The Muon Puzzle in cosmic-ray induced air showers and its connection to the Large Hadron Collider"*, Astrophys. Space Sci. **367**, 27 (2022).

[72] A. Aab et al. (Pierre Auger), *"Direct measurement of the muonic content of extensive air showers between* $\mathbf{2 \times 10^{17}}$ *and* $\mathbf{2 \times 10^{18}}$ *eV at the Pierre Auger Observatory"*, Eur. Phys. J. C **80**, 751 (2020).

[73] D. R. Bergman et al. (Pierre Auger, Telescope Array), *"Updated comparison of the UHECR energy spectra measured by the Pierre Auger Observatory and the Telescope Array"*, PoS **ICRC2025**, 381 (2025).

[74] P. Padovani et al., *"Active galactic nuclei: what's in a name?"*, Astron. Astrophys. Rev. **25**, 2 (2017).

[75] G. Lodato et al., *"Recent developments in the theory of tidal disruption events"*, JHEAp **7**, 158–162 (2015).

[76] A. A. Halim et al. (Pierre Auger), *"Large-scale Cosmic-ray Anisotropies with 19 yr of Data from the Pierre Auger Observatory"*, Astrophys. J. **976**, 48 (2024).

[77] A. Abdul Halim et al. (Pierre Auger), *"An update on the arrival direction studies made with data from the Pierre Auger Observatory"*, PoS **ICRC2023**, 252 (2023).

[78] J. Kim et al. (Telescope Array), *"Anisotropies in the arrival direction distribution of ultra-high energy cosmic rays measured by the Telescope Array surface detector"*, PoS **ICRC2023**, 244 (2023).

[79] A. Abdul Halim et al. (Pierre Auger), *"The Distribution of Ultrahigh-energy Cosmic Rays along the Supergalactic Plane Measured at the Pierre Auger Observatory"*, Astrophys. J. **984**, 123 (2025).

[80] `https://neutrinos.fnal.gov/faqs/`, Accessed: 2026-01-22.

[81] J. A. Formaggio and G. P. Zeller, *"From eV to EeV: Neutrino Cross Sections Across Energy Scales"*, Rev. Mod. Phys. **84**, 1307–1341 (2012).

[82] C. Spiering, *"High energy neutrinos and neutrino telescopes"*, in International Conference on History of the Neutrino: 1930-2018 (Mar. 2019).

[83] M. Aartsen et al. (IceCube), *"The IceCube Neutrino Observatory: instrumentation and online systems"*, Journal of Instrumentation **12**, P03012–P03012 (2017).

[84] S. Adrián-Martínez et al. (KM3NeT), *"Letter of intent for KM3NeT 2.0"*, Journal of Physics G: Nuclear and Particle Physics **43**, 084001 (2016).

[85] A. V. Avrorin et al. (Baikal-GVD), *"Baikal-GVD Experiment"*, Phys. Atom. Nucl. **83**, 916–921 (2020).

[86] M. Agostini et al. (P-ONE), *"The Pacific Ocean Neutrino Experiment"*, Nature Astron. **4**, 913–915 (2020).

[87] C. Lagunas Gualda et al. (P-ONE), *"First instrumented line of the Pacific Ocean Neutrino Experiment: status, development and outlook"*, PoS **ICRC2025**, 1087 (2025).

[88] Z. P. Ye et al. (TRIDENT), *"A multi-cubic-kilometre neutrino telescope in the western Pacific Ocean"*, Nature Astron. **7**, 1497–1505 (2023).

[89] T.-Q. Huang et al., *"Proposal for the High Energy Neutrino Telescope"*, PoS **ICRC2023**, 1080 (2023).

[90] M. G. Aartsen et al. (IceCube), *"Observation of High-Energy Astrophysical Neutrinos in Three Years of IceCube Data"*, Phys. Rev. Lett. **113**, 101101 (2014).

[91] O. Adriani et al. (KM3NeT, MessMapp Group, LAT, m Telescope Group, SVOM, Owens Valley Radio Observatory 40-m Telescope Group, Fermi-LAT), *"Characterizing Candidate Blazar Counterparts of the Ultra-High-Energy Event KM3-230213A"*, (2025).

[92] R. Stein et al., *"A tidal disruption event coincident with a high-energy neutrino"*, Nature Astron. **5**, 510–518 (2021).

[93] S. Reusch et al., *"Candidate Tidal Disruption Event AT2019fdr Coincident with a High-Energy Neutrino"*, Phys. Rev. Lett. **128**, 221101 (2022).

[94] S. van Velzen et al., *"Establishing accretion flares from supermassive black holes as a source of high-energy neutrinos"*, Mon. Not. Roy. Astron. Soc. **529**, 2559–2576 (2024).

[95] R. Abbasi et al. (IceCube), *"IceCat-2: Updated IceCube Event Catalog of Alert Tracks"*, PoS **ICRC2025**, 1224 (2025).

[96] M. G. Aartsen et al. (IceCube), *"Detection of a particle shower at the Glashow resonance with IceCube"*, Nature **591**, [Erratum: Nature 592, E11 (2021)], 220–224 (2021).

[97] L. D. Landau and I. Pomeranchuk, *"Limits of applicability of the theory of bremsstrahlung electrons and pair production at high-energies"*, Dokl. Akad. Nauk Ser. Fiz. **92**, 535–536 (1953).

[98] A. B. Migdal, *"Bremsstrahlung and Pair Production at High Energies in Condensed Media"*, Phys. Rev. **103**, 1811–1820 (1956).

[99] S. W. Barwick and C. Glaser, *"Chapter 6: Radio Detection of High Energy Neutrinos in Ice"*, 237–302 (2023).

[100] K. Graf, *"Acoustic neutrino detection in sea water: Technical aspects"*, AIP Conf. Proc. **1535**, 169 (2013).

[101] R. Nahnhauer, *"Acoustic detection of ultra-high energetic neutrinos - a snap shot"*, Nucl. Instrum. Meth. A **692**, 58–64 (2012).

[102] K. D. de Vries, K. Hanson, and T. Meures, *"On the feasibility of RADAR detection of high-energy neutrino-induced showers in ice"*, Astropart. Phys. **60**, 25–31 (2015).

[103] S. Prohira et al., *"Observation of Radar Echoes From High-Energy Particle Cascades"*, Phys. Rev. Lett. **124**, 091101 (2020).

[104] S. Prohira et al. (Radar Echo Telescope), *"The Radar Echo Telescope for Cosmic Rays: Pathfinder experiment for a next-generation neutrino observatory"*, Phys. Rev. D **104**, 102006 (2021).

[105] K. de Vries et al. (Radar Echo Telescope), *"The Radar Echo Telescope for Neutrinos"*, PoS **ICRC2021**, 1195 (2021).

[106] S. Chiche et al., *"Loss of Coherence and Change in Emission Physics for Radio Emission from Very Inclined Cosmic-Ray Air Showers"*, Phys. Rev. Lett. **132**, 231001 (2024).

[107] O. Scholten, K. Werner, and F. Rusydi, *"A Macroscopic Description of Coherent Geo-Magnetic Radiation from Cosmic Ray Air Showers"*, Astropart. Phys. **29**, 94–103 (2008).

[108] O. Scholten, K. D. de Vries, and K. Werner, *"Coherent radiation from extensive air showers"*, Nucl. Instrum. Meth. A **662**, S80–S84 (2012).

[109] K. D. de Vries et al., *"Coherent Cherenkov Radiation from Cosmic-Ray-Induced Air Showers"*, Phys. Rev. Lett. **107**, 061101 (2011).

[110] T. Huege and D. Besson, *"Radio-wave detection of ultra-high-energy neutrinos and cosmic rays"*, PTEP **2017**, 12A106 (2017).

[111] F. G. Schröder, *"Status of the radio technique for cosmic-ray induced air showers"*, Nucl. Part. Phys. Proc. **279-281**, 190–197 (2016).

[112] H. Falcke et al. (LOPES), *"Detection and imaging of atmospheric radio flashes from cosmic ray air showers"*, Nature **435**, 313–316 (2005).

[113] M. Ludwig and T. Huege, *"REAS3: Monte Carlo simulations of radio emission from cosmic ray air showers using an 'end-point' formalism"*, Astropart. Phys. **34**, 438–446 (2011).

[114] T. Huege et al., *"The convergence of EAS radio emission models and a detailed comparison of REAS3 and MGMR simulations"*, Nucl. Instrum. Meth. A **662**, S179–S186 (2012).

[115] W. Demtröder, Elektrizität und Optik, Vol. 2 (Springer, 2009).

[116] D. Ardouin et al., *"Geomagnetic origin of the radio emission from cosmic ray induced air showers observed by CODALEMA"*, Astroparticle Physics **31**, 192–200 (2009).

[117] J. Alvarez-Muñiz, A. Romero-Wolf, and E. Zas, *"Cherenkov radio pulses from electromagnetic showers in the time-domain"*, Phys. Rev. D **81**, 123009 (2010).

[118] K. Bechtol et al., *"SLAC T-510 experiment for radio emission from particle showers: Detailed simulation study and interpretation"*, Phys. Rev. D **105**, 063025 (2022).

[119] O. Scholten et al., *"Measurement of the circular polarization in radio emission from extensive air showers confirms emission mechanisms"*, Phys. Rev. D **94**, 103010 (2016).

[120] A. Aab et al. (Pierre Auger), *"Measurement of the Radiation Energy in the Radio Signal of Extensive Air Showers as a Universal Estimator of Cosmic-Ray Energy"*, Phys. Rev. Lett. **116**, 241101 (2016).

[121] S. Bouma, *"Direction Reconstruction of Radio Signals in Neutrino Detectors in Ice"*, PhD thesis (Friedrich-Alexander-Universität Erlangen-Nürnberg, 2025).

[122] T. Huege, M. Ludwig, and C. W. James, *"Simulating radio emission from air showers with CoREAS"*, AIP Conf. Proc. **1535**, 128 (2013).

[123] B. Pont (Pierre Auger), *"First Data of the 3000 km^2 Radio Detector at the Pierre Auger Observatory"*, in 39th International Cosmic Ray Conference (Sept. 2025).

[124] M. P. van Haarlem et al. (LOFAR), *"LOFAR: The LOw-Frequency ARray"*, Astron. Astrophys. **556**, A2 (2013).

[125] P. Schellart et al. (LOFAR), *"Detecting cosmic rays with the LOFAR radio telescope"*, Astron. Astrophys. **560**, A98 (2013).

[126] S. Buitink et al., *"Progress on air shower observations with the SKA"*, PoS **ICRC2025**, 207 (2025).

[127] P. A. Bezyazeekov et al., *"Measurement of cosmic-ray air showers with the Tunka Radio Extension (Tunka-Rex)"*, Nucl. Instrum. Meth. A **802**, 89–96 (2015).

[128] J. Alvarez-Muñiz et al., *"Coherent Cherenkov radio pulses from hadronic showers up to EeV energies"*, Astropart. Phys. **35**, 287–299 (2012).

[129] P. Schellart et al. (LOFAR), *"Polarized radio emission from extensive air showers measured with LOFAR"*, JCAP **10**, 014 (2014).

[130] A. Nelles et al. (LOFAR), *"Measuring a Cherenkov ring in the radio emission from air showers at 110–190 MHz with LOFAR"*, Astropart. Phys. **65**, 11–21 (2015).

[131] A. Abdul Halim et al. (Pierre Auger), *"Determination of the energy scale of cosmic ray measurements using the Auger Engineering Radio Array"*, PoS **ICRC2025**, 292 (2025).

[132] A. Corstanje et al., *"Depth of shower maximum and mass composition of cosmic rays from 50 PeV to 2 EeV measured with the LOFAR radio telescope"*, Phys. Rev. D **103**, 102006 (2021).

[133] A. Abdul Halim et al. (Pierre Auger), *"Radio Measurements of the Depth of Air-Shower Maximum at the Pierre Auger Observatory"*, Phys. Rev. D **109**, 022002 (2024).

[134] S. Buitink et al. (LOFAR), *"Method for high precision reconstruction of air shower X_{max} using two-dimensional radio intensity profiles"*, Phys. Rev. D **90**, 082003 (2014).

[135] A. Abdul Halim et al. (Pierre Auger), *"Demonstrating Agreement between Radio and Fluorescence Measurements of the Depth of Maximum of Extensive Air Showers at the Pierre Auger Observatory"*, Phys. Rev. Lett. **132**, 021001 (2024).

[136] Z. S. Meyers et al., Radio detection of EeV neutrinos in dielectric media using the Askaryan effect (DESY, Hamburg, 2022).

[137] I. Kravchenko et al. (RICE), *"Updated results from the RICE experiment and future prospects for ultra-high energy neutrino detection at the south pole"*, Phys. Rev. D **85**, 062004 (2012).

[138] P. Allison et al. (ARA), *"Design and initial performance of the Askaryan Radio Array prototype EeV neutrino detector at the South Pole"*, Astroparticle Physics **35**, 457–477 (2012).

[139] P. Allison et al., *"Design and performance of an interferometric trigger array for radio detection of high-energy neutrinos"*, Nucl. Instrum. Meth. A **930**, 112–125 (2019).

[140] P. Allison et al. (ARA), *"Constraints on the diffuse flux of ultrahigh energy neutrinos from four years of Askaryan Radio Array data in two stations"*, Phys. Rev. D **102**, 043021 (2020).

[141] M. Ackermann et al., *"High-energy and ultra-high-energy neutrinos: A Snowmass white paper"*, JHEAp **36**, 55–110 (2022).

[142] C. A. Argüelles et al. (TAMBO), *"TAMBO: A Deep-Valley Neutrino Observatory"*, (2025).

[143] A. V. Olinto et al. (POEMMA), *"The POEMMA (Probe of Extreme Multi-Messenger Astrophysics) observatory"*, JCAP **06**, 007 (2021).

[144] A. N. Otte et al., *"Trinity: An Air-Shower Imaging Instrument to detect Ultrahigh Energy Neutrinos"*, (2019).

[145] A. Anker et al. (ARIANNA), *"Measuring the polarization reconstruction resolution of the ARIANNA neutrino detector with cosmic rays"*, JCAP **04**, 022 (2022).

[146] J. Alvarez-Muñiz et al. (GRAND), *"The Giant Radio Array for Neutrino Detection (GRAND): Science and Design"*, Sci. China Phys. Mech. Astron. **63**, 219501 (2020).

[147] Y. Chen (TAROGE), *"Status and Performance of TAROGE-4 for Radio Detection of Extensive Air Showers"*, PoS **ICRC2025**, 219 (2025).

[148] S.-H. Wang et al. (TAROGE, ARIANNA), *"TAROGE-M: radio antenna array on antarctic high mountain for detecting near-horizontal ultra-high energy air showers"*, JCAP **11**, 022 (2022).

[149] S. Wissel et al., *"Prospects for high-elevation radio detection of >100 PeV tau neutrinos"*, JCAP **11**, 065 (2020).

[150] A. Zeolla and Z. Martin (BEACON), *"Detection of Cosmic Rays at High Elevation by the BEACON Prototype"*, PoS **ICRC2025**, 444 (2025).

[151] O. Martineau-Huynh (GRAND), *"Status of the GRAND project"*, in 39th International Cosmic Ray Conference (July 2025).

[152] K. Kotera et al. (GRAND, BEACON), *"The Hybrid Elevated Radio Observatory for Neutrinos (HERON) Project"*, PoS **ICRC2025**, 1078 (2025).

[153] P. W. Gorham et al. (ANITA), *"The Antarctic Impulsive Transient Antenna Ultra-high Energy Neutrino Detector Design, Performance, and Sensitivity for 2006-2007 Balloon Flight"*, Astropart. Phys. **32**, 10–41 (2009).

[154] Q. Abarr et al. (PUEO), *"The Payload for Ultrahigh Energy Observations (PUEO): a white paper"*, JINST **16**, P08035 (2021).

[155] P. W. Gorham et al. (ANITA), *"Constraints on the ultrahigh-energy cosmic neutrino flux from the fourth flight of ANITA"*, Phys. Rev. D **99**, 122001 (2019).

[156] P. W. Gorham et al. (ANITA), *"Unusual Near-Horizon Cosmic-Ray-like Events Observed by ANITA-IV"*, Phys. Rev. Lett. **126**, 071103 (2021).

[157] P. W. Gorham et al. (ANITA), *"Characteristics of Four Upward-pointing Cosmic-ray-like Events Observed with ANITA"*, Phys. Rev. Lett. **117**, 071101 (2016).

[158] P. W. Gorham et al. (ANITA), *"Observation of an Unusual Upward-going Cosmic-ray-like Event in the Third Flight of ANITA"*, Phys. Rev. Lett. **121**, 161102 (2018).

[159] A. Romero-Wolf et al., *"Comprehensive analysis of anomalous ANITA events disfavors a diffuse tau-neutrino flux origin"*, Phys. Rev. D **99**, 063011 (2019).

[160] A. J. Shturmakov et al., *"A new 122 mm electromechanical drill for deep ice-sheet coring (DISC): 1. Design concepts"*, Annals of Glaciology **47**, 28–34 (2007).

[161] S. Agarwal et al. (RNO-G), *"Construction, commissioning and operation of the Radio Neutrino Observatory in Greenland (RNO-G)"*, PoS **ICRC2025**, 1168 (2025).

[162] P. Schriefer, *"RNO-G's Perspective on Wind - Power Generation and Background Tagging"*, MSc thesis (Friedrich-Alexander-Universität Erlangen-Nürnberg, 2025).

[163] S. Agarwal et al. (RNO-G), *"Instrument design and performance of the first seven stations of RNO-G"*, JINST **20**, P04015 (2025).

[164] I. Plaisier, *"Reconstructing the Arrival Direction of Cosmic Neutrinos with the Radio Neutrino Observatory Greenland (RNO-G)"*, PhD thesis (Friedrich-Alexander-Universität Erlangen-Nürnberg, 2022).

[165] P. Windischhofer (RNO-G), *"Calibrating the Radio Neutrino Observatory in Greenland"*, PoS **ARENA2024**, 003 (2024).

[166] C. A. Balanis, Antenna Theory: Analysis and Design (John Wiley & Sons, 2016).

[167] `https://www.cd-corp.com/english/`, Accessed: 2026-01-07.

[168] S. W. Barwick et al. (ARIANNA), *"Time Domain Response of the ARIANNA Detector"*, Astropart. Phys. **62**, 139–151 (2015).

[169] H. Booker, *"Slot aerials and their relation to complementary wire aerials (Babinet's principle)"*, Journal of the Institution of Electrical Engineers - Part IIIA: Radiolocation **93**, 620–626 (1946).

[170] E. Oberla et al. (RNO-G), *"Low-Power Radiofrequency Systems for the RNO-G Project"*, PoS **ICRC2023**, 1171 (2023).

[171] S. Agarwal et al. (RNO-G), *"Calibration of the Radio Neutrino Observatory in Greenland using thermal noise"*, PoS **ICRC2025**, 1003 (2025).

[172] L. Pyras, *"Cosmic Rays and the Radio Neutrino Observatory Greenland (RNO-G)"*, PhD thesis (Friedrich-Alexander-Universität Erlangen-Nürnberg, 2025).

[173] J. A. Aguilar et al., *"Triboelectric backgrounds to radio-based polar ultra-high energy neutrino (UHEN) experiments"*, Astropart. Phys. **145**, 102790 (2023).

[174] M. Reichert, *"The diffuse Galactic signal in RNO-G"*, BSc thesis (Friedrich-Alexander-Universität Erlangen-Nürnberg, 2023).

[175] A. Nelles et al. (LOFAR), *"Calibrating the absolute amplitude scale for air showers measured at LOFAR"*, JINST **10**, P11005 (2015).

[176] K. Mulrey et al., *"Calibration of the LOFAR low-band antennas using the Galaxy and a model of the signal chain"*, Astropart. Phys. **111**, 1–11 (2019).

[177] A. A. Halim et al. (Pierre Auger), *"Long-term calibration and validation of stability of the Auger Engineering Radio Array using the diffuse Galactic radio emission"*, JINST **20**, P12017 (2025).

[178] M. Büsken, T. Fodran, and T. Huege, *"Uncertainties of the 30–408 MHz Galactic emission as a calibration source for radio detectors in astroparticle physics"*, Astron. Astrophys. **679**, A50 (2023).

[179] S. Agarwal et al. (RNO-G), *"Solar flare observations with the Radio Neutrino Observatory Greenland (RNO-G)"*, Astropart. Phys. **164**, 103024 (2025).

[180] B. Oeyen, *"An empirical ice model for the Radio Neutrino Observatory in Greenland (RNO-G)"*, PhD thesis (Ghent University, 2025).

[181] S. Agarwal et al. (RNO-G), *"Radio emission from airplanes as observed with RNO-G"*, JINST **20**, P11015 (2025).

[182] F. Di Vruno et al., *"Unintended electromagnetic radiation from Starlink satellites detected with LOFAR between 110 and 188 MHz"*, A&A **676**, A75 (2023).

[183] X. Zhang et al., *"Broadband polarized radio emission detected from Starlink satellites below 100 MHz with NenuFAR"*, A&A **698**, A244 (2025).

[184] D. Grigg et al., *"Detection of intended and unintended emissions from Starlink satellites in the SKA-Low frequency range, at the SKA-Low site, with an SKA-Low station analogue"*, A&A **678**, L6 (2023).

[185] S. Agarwal et al. (RNO-G), *"Validating the performance of the Radio Neutrino Observatory in Greenland using cosmic-ray air showers"*, Submitted to EPJ C (2025).

[186] J. B. Hagen, Radio-Frequency Electronics: Circuits and Applications, 2nd ed. (Cambridge University Press, 2009).

[187] T. Fodran, *"Commissioning the Auger Prime Radio Detector. Paving the Way for Mass Separation"*, PhD thesis (Nijmegen U., 2024).

[188] F. Harris, *"On the use of windows for harmonic analysis with the discrete Fourier transform"*, Proceedings of the IEEE **66**, 51–83 (1978).

[189] `https://www.mongodb.com/`, Accessed: 2025-11-23.

[190] S. Agarwal et al. (RNO-G), *"Absolute amplitude calibration of the Radio Neutrino Observatory using system noise"*, In preparation.

[191] `https://streamlit.io/`, Accessed: 2025-11-26.

[192] C. Glaser et al., *"NuRadioMC: Simulating the radio emission of neutrinos from interaction to detector"*, Eur. Phys. J. C **80**, 77 (2020).

[193] `https://www.keysight.com/de/de/assets/7018-01537/data-sheets/5989-6433.pdf`, Accessed: 2026-03-11.

[194] `https://www.mouser.de/datasheet/3/118/1/BGA2851.pdf`, Accessed: 2025-11-10.

[195] C. Glaser et al., *"NuRadioReco: A reconstruction framework for radio neutrino detectors"*, Eur. Phys. J. C **79**, 464 (2019).

[196] K. M. Cuffey and W. S. B. Paterson, The Physics of Glaciers (Academic Press, 2010).

[197] R. L. Hawley, E. M. Morris, and J. R. McConnell, *"Rapid techniques for determining annual accumulation applied at Summit, Greenland"*, Journal of Glaciology **54**, 839–845 (2008).

[198] R. Alley and B. Koci, *"Ice-Core Analysis at Site A, Greenland: Preliminary Results"*, Annals of Glaciology **10**, 1–4 (1988).

[199] C. Deaconu et al., *"Measurements and modeling of near-surface radio propagation in glacial ice and implications for neutrino experiments"*, Phys. Rev. D **98**, 043010 (2018).

[200] S. Agarwal et al. (RNO-G), *"Probing the Firn Refractive Index Profile and Borehole Closure Using Antenna Response"*, Submitted to J. Glaciol. (2025).

[201] `https://wipl-d.com/`, Accessed: 2025-12-13.

[202] `https://www.remcom.com/xfdtd-3d-em-simulation-software`, Accessed: 2025-12-13.

[203] H. M. Jol, Ground Penetrating Radar Theory and Applications (Elsevier, 2008).

[204] F. Ferraccioli et al., Processed airborne radio-echo sounding data from the POLARGAP survey covering the South Pole, and Foundation and Recovery Glaciers, East Antarctica (2015/2016), version 1.0 (NERC EDS UK Polar Data Centre, 2021).

[205] M. Rongen, Private communication, 2025.

[206] `https://www.geophysical.com/products/200-hs-antenna`, Accessed: 2025-11-26.

[207] S. De Kockere, Private communication, 2025.

[208] `https://impdar.readthedocs.io/en/latest/`, Accessed: 2025-12-08.

[209] N. Wiener, Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications (The MIT press, 1949).

[210] R. O. Duda and P. E. Hart, *"Use of the Hough transformation to detect lines and curves in pictures"*, Commun. ACM **15**, 11–15 (1972).

[211] C. Glaser et al., *"NuRadioReco: A reconstruction framework for radio neutrino detectors"*, Eur. Phys. J. C **79**, 464 (2019).

[212] NOAA NCEI Geomagnetic Modeling Team, World Magnetic Model 2025, `https://doi.org/10.25921/aqfd-sd83`, Accessed 2025-08-27.

[213] L. Pyras, *"Optimizing the triggering strategy for the detection of cosmic rays with the Radio Neutrino Observatory Greenland (RNO-G)"*, MSc thesis (Humboldt-Universität zu Berlin, 2020).

[214] H. Zheng et al., *"An improved model of diffuse galactic radio emission from 10 MHz to 5 THz"*, Mon. Not. Roy. Astron. Soc. **464**, 3486–3497 (2017).

[215] V. Verzi (Pierre Auger), *"Measurement of the energy spectrum of ultra-high energy cosmic rays using the Pierre Auger Observatory"*, PoS **ICRC2019**, 450 (2020).

[216] A. Corstanje et al., *"A high-precision interpolation method for pulsed radio signals from cosmic-ray air showers"*, JINST **18**, P09005 (2023).

[217] J. Henrichs et al. (RNO-G), *"Searching for air showers with RNO-G"*, PoS **ARENA2022**, 007 (2023).

[218] J. M. Roberts et al., *"LAB4D: A Low Power, Multi-GSa/s, Transient Digitizer with Sampling Timebase Trimming Capabilities"*, Nucl. Instrum. Meth. A **925**, 92–100 (2019).

[219] C. Welling, C. Glaser, and A. Nelles, *"Reconstructing the cosmic-ray energy from the radio signal measured in one single station"*, JCAP **10**, 075 (2019).

[220] M. Ravn et al., *"Likelihood reconstruction of radio signals of neutrinos and cosmic rays"*, in 10th International Workshop on Acoustic and Radio EeV Neutrino Detection Activities (Sept. 2024).

[221] `https://indico.cern.ch/event/1122790/contributions/4713580/attachments/2381493/4270369/ABCDGuide_draft18Oct18.pdf`, Accessed: 2025-10-21.

[222] `https://gml.noaa.gov/aftp/data/meteorology/in-situ/sum/`, Accessed: 2025-10-16.

Acknowledgements

Writing this thesis and doing all the underlying work would not have been possible without the help and support of many people over the last four years.

First, I would like to thank my supervisor, Anna Nelles, for giving me the opportunity to join DESY as a PhD student and contribute to the amazing RNO-G experiment. Also, I am thankful for her great supervision and continuous support throughout the whole period.

Further, I would like to thank the members of the RNO-G collaboration for sharing their expertise, helping me understand all the details of the experiment, and for providing valuable feedback on my work over the past few years. I would particularly like to thank Cosmin and Eric, to whom I have asked countless questions and who have always been a great help.

Also, I would like to thank all current and former members of the radio group at ECAP and the IceCube group at DESY for making the four years very enjoyable. In particular, I would like to thank my office mates Steffen, Christoph, Ilse, Zack, Navid, Lilly, Enrique, Zeynep, Annanay, and Hannes for turning the office into a fun and productive space. Additionally, I would like to thank my official co-founders of BBB, Navid and Zack, who were always up for a fun idea and for making the time especially memorable.

Moreover, I would like to thank Philipp, Felix, Sjoerd, and Steffen for many fruitful discussions about my analysis and the database. Thank you, Felix, for taking on the database with me, thank you, Timo, for being a great hardware support and for always answering my questions about the RNO-G hardware, and thank you, Lilly, for making the time in the laboratory fun and productive. Additionally, I would like to thank Eric, Cosmin, Christoph, Simon, and Kaeli for an amazing maintenance and calibration season at Summit Station.

I would like to thank Simon, Hannes, Pascal, Cris, Sjoerd, Lilly, Zeynep, Cristina, Annanay, Despina, Teresa, Navid, Enrique, Thijs, Philipp, and Sarah for helping with proofreading my thesis.

Finally, I would like to thank my wife Sarah and my two children for always cheering me up when I came home and for being patient with me. Especially during the final months of this project, Sarah made it possible for me to focus on completing this thesis. Thank you, Sarah, it would not have been possible without you! Also, I would like to thank my mother and sister who always supported me on my way.

Thank you all for making this thesis possible!

Erklärung zur Nutzung von (generativer) KI

Ich, Jakob Henrichs-Schuhegger, erkläre hiermit, in welchem Umfang ich Künstliche Intelligenz (KI), einschließlich Tools wie ChatGPT, Co-Pilot, HAWKI, DeepL Write, Grammarly oder vergleichbare Werkzeuge, bei der Erstellung der vorliegenden Prüfungsleistung genutzt habe. Ich habe KI-gestützte Werkzeuge wie folgt eingesetzt:

- Zur Formulierungshilfe (z. B. Verbesserung von Satzbau, Stil) - KI wurde zur Verbesserung von Stil und Lesbarkeit genutzt, jedoch nicht zur vollständigen Textgenerierung.
- Zur Korrektur von Grammatik und Rechtschreibung - KI wurde zur sprachlichen Optimierung genutzt, jedoch nicht zur inhaltlichen Veränderung oder Textgenerierung.
- Zur Übersetzung von Texten - KI wurde für Übersetzungen genutzt, wobei alle übersetzten Passagen überarbeitet und auf inhaltliche Richtigkeit geprüft wurden.

Ich versichere, alle Nutzungen vollständig und wahrheitsgemäß angegeben zu haben. Mir ist bewusst, dass das Verschweigen oder falsche Angaben zur Nutzung von KI-Tools als Täuschungsversuch gewertet werden können und prüfungsrechtliche Konsequenzen haben.

Ort, Datum	Unterschrift