

Realization of a Self-triggered Detector for the Radio Emission of Cosmic Rays

Zur Erlangung des akademischen Grades eines
DOKTORS DER NATURWISSENSCHAFTEN
von der Fakultät für Physik des
Karlsruher Instituts für Technologie (KIT)

genehmigte

DISSERTATION

von

Dipl. Phys. Adrian Schmidt
aus Lörrach

Karlsruhe
Oktober 2012

Tag der mündlichen Prüfung: 15. Juni 2011

Referent: Prof. Dr. J. Blümer
Institut für Kernphysik
Korreferent: Prof. Dr. H. Gemmeke
Institut für Prozessdatenverarbeitung und Elektronik

Abstract

Realization of a self-triggered detector for the radio emission of cosmic rays

When a high-energetic primary cosmic ray particle penetrates the Earth's atmosphere, a cascade of secondary particles is created, which can be used for an indirect detection. So far, arrays of particle detectors on the surface were employed for this purpose or fluorescence telescopes, which observe the shower development in air. A rather new method uses the radio signal, which is emitted by the charged secondary particles, when they are deflected in the geomagnetic field. Radio detection has the potential to combine the advantages of both established methods, namely the precise, almost calorimetric energy determination of the fluorescence measurement, and the independence from weather and daytime of the particle detectors.

In the past, all measurements of the radio emission were achieved with an external trigger, the experiments were mostly designed as extensions of an existing surface detector. To establish radio detection as an adequate alternative and to enable stand-alone operation, the development of a reliable self-trigger is however essential. The design of a suitable trigger strategy, its realization as an embedded digital system, as well as commissioning and first measurements are the central topics of this thesis. The main challenge arises from the signal background, which is strongly contaminated with artificial interferences. Because of the setup as a widespread array of solar-powered, remote controlled antenna stations, the available electric power budget is very limited, just as the data rate and thus the allowed false trigger rate.

The trigger developed in the context of this thesis is divided into three stages: The first step consists of a dynamic suppression of radio frequency interferences using a Fourier transform into frequency domain, where the interfering signal can be identified and removed autonomously, before transforming back to time domain. Compared to conventional digital filters, this method is far superior, however it is computationally very intensive. To make a real-time implementation possible in the first place, despite the limited performance of the station electronics, a multitude of numerical tricks and optimizations is necessary. With this signal conditioning, the detection threshold is typically lowered by 40 %.

In the second stage, pulses in the conditioned signal are detected with a variable threshold. Additional cuts on the pulse shape reduce the false trigger rate. Moreover, noise transients from the power grid are identified on the basis of their periodicity with line frequency. They are predicted by means of a phase-locked loop and rejected synchronously.

The last trigger stage merges the event candidates from multiple antenna stations, so a direction of origin can be reconstructed. For this purpose, a simple triangulation based on antenna triplets is used, which scales well for a larger antenna array, due to its fast computability. Instead of rejecting events solely based on their zenith angle to remove artificial background from horizon, the method proposed here applies a dynamic suppression of conspicuous directions. In doing so, the background can be reduced by several orders of magnitude down to a few millihertz, while the effective detector aperture is hardly decreased.

Initially, the trigger approach was implemented as far as possible using available hardware. The prototype setup was then tested in a 15 month measuring campaign at the Pierre Auger Observatory. On the basis of these experiences, a new electronics was developed, made-to-measure for radio detection. It was commissioned in the frame of the first full-scale radio experiment AERA, and was already able to detect the first cosmic ray events.

Zusammenfassung

Umsetzung eines selbstgetriggerten Detektors der Radioemission kosmischer Strahlung

Dringt ein hochenergetisches Primärteilchen der kosmischen Strahlung in die Erdatmosphäre ein, entsteht eine Kaskade von Sekundärteilchen, die für den indirekten Nachweis verwendet werden kann. Bisher wurden zu diesem Zweck Netze von Teilchendetektoren am Erdboden eingesetzt, oder auch Fluoreszenzteleskope, die die Schauerentwicklung in Luft beobachten. Eine eher neue Methode nutzt das Radiosignal, das von den geladenen Sekundärteilchen ausgesendet wird, wenn sie im Erdmagnetfeld abgelenkt werden. Die Radiodetektion hat das Potential, die Vorteile beider etablierten Methoden zu verbinden, namentlich die präzise, fast kalorimetrische Energiebestimmung der Fluoreszenzmessung und die Wetter- und Tageszeitunabhängigkeit der Teilchendetektoren.

In der Vergangenheit wurden sämtliche Messungen der Radioemission mithilfe eines externen Triggers erzielt, meist waren die Experimente als Erweiterung eines bestehenden Oberflächendetektors angelegt. Um sich als vollwertige Alternative durchzusetzen und einen selbstständigen Betrieb zu ermöglichen, ist jedoch die Entwicklung eines zuverlässigen Selbsttriggers unerlässlich. Der Entwurf einer geeigneten Triggerstrategie, ihre Umsetzung als eingebettetes digitales System sowie Inbetriebnahme und erste Messungen sind das zentrale Thema dieser Arbeit. Die Schwierigkeit ergibt sich dabei vor allem aus dem Signaluntergrund, der durch künstliche Störquellen stark belastet wird. Wegen des vorgesehenen Aufbaus in Form eines großflächigen Feldes von solarbetriebenen, drahtlos verbundenen Antennenstationen, ist die verfügbare elektrische Leistung stark begrenzt, ebenso die Datenübertragungsrate und somit die erlaubte Falschtriggerrate.

Der im Rahmen dieser Arbeit entwickelte Trigger ist in drei Stufen unterteilt: Der erste Schritt besteht aus einer dynamischen Unterdrückung monofrequenter Störsignale mithilfe einer Fourier-Transformation in den Frequenzraum, wo die Störungen selbstständig identifiziert und entfernt werden, bevor das Signal zurück in den Zeitraum transformiert wird. Verglichen mit konventionellen digitalen Filtern ist diese Methode deutlich überlegen, sie ist jedoch sehr rechenaufwändig. Um eine Echtzeit-Implementierung trotz der begrenzten Leistungsfähigkeit der Stationselektronik überhaupt zu ermöglichen, sind eine Vielzahl von numerischen Kniffen und Optimierungen notwendig. Durch diese Signalaufbereitung wird die Nachweisschwelle um 40 % gesenkt.

In der zweiten Stufe werden Pulse im aufbereiteten Signal durch einen variablen Schwellwert detektiert, zusätzliche Schnitte auf die Pulsform sollen die Falschtriggerrate verringern. Außerdem werden Störpulse vom Stromnetz anhand ihrer Periodizität mit Netzfrequenz erkannt, mithilfe einer Phasenregelschleife vorhergesagt und synchron verworfen.

Die letzte Triggerstufe führt die Ereigniskandidaten von mehreren Antennenstation zusammen, so dass die Ursprungsrichtung des Signals rekonstruiert werden kann. Dazu wird eine einfache Triangulation innerhalb von Antennendreiecken benutzt, die wegen ihrer schnellen Berechenbarkeit auch gut für ein großes Antennenfeld skaliert. Anstatt lediglich Ereignisse anhand ihres Zenitwinkels zu verwerfen, um künstlichen Untergrund vom Horizont zu entfernen, wendet die hier vorgeschlagene Methode eine dynamische Unterdrückung von auffälligen Richtungen an. Dadurch kann der Untergrund um mehrere Größenordnungen auf wenige Millihertz gesenkt werden, während die effektive Detektorapertur kaum reduziert wird.

Der Triggeransatz wurde zunächst im Rahmen der Möglichkeiten als Prototyp mit vorhandener Hardware aufgebaut und in einer 15-monatigen Messkampagne am Pierre Auger Observatorium getestet. Auf Grundlage dieser Erfahrungen wurde eine neue, für die Radiodetektion maßgefertigte Elektronik entwickelt. Sie wurde im Rahmen des ersten großen Radio Experiments AERA in Betrieb genommen und konnte bereits die ersten Teilchen der kosmischen Strahlung detektieren.

Contents

Abstract	I
1. Introduction	1
2. Cosmic rays	3
2.1. The source of cosmic rays	3
2.2. Extensive air showers	5
2.3. Detection of cosmic rays	7
2.4. Radio emission from cosmic ray air showers	10
2.4.1. The radio emission model	10
2.4.2. Properties of detected radio pulses	12
2.4.3. Radio detection experiments	14
3. System development of a radio self-trigger	17
3.1. Measurements with LOPES ^{STAR}	17
3.1.1. Antenna details	19
3.1.2. The signal chain	22
3.1.3. Analog self-trigger investigations	23
3.1.4. Offline self-trigger development	23
3.2. The prototype for digital self-triggering	28
3.2.1. Properties of a Field Programmable Gate Array (FPGA)	29
3.2.2. The prototype hardware	31
3.2.3. The prototype data acquisition system	32
4. Measurements with the prototype self-trigger	35
4.1. Setup at the BLS	35
4.2. Background conditions at the BLS	37
4.3. Search for cosmic rays	45
4.4. Survey for future site	49
4.5. Conclusion	51
5. The AERA experiment	53
5.1. Key features of AERA	53
5.2. Hardware development for AERA	56
5.3. Self trigger firmware	59
5.3.1. The Trigger module	60
5.3.2. The data management unit	60
6. Suppression of radio frequency interferences (RFI)	63
6.1. FPGA Implementation of the RFI-suppression	64
6.2. Choice of the FFT-engine	66
6.3. Effect of the window function	69
6.4. Median filter implementation	72

6.5.	Results of the RFI-suppression	75
6.5.1.	Effect on the random trigger rate	76
6.5.2.	Results of the envelope calculation	77
6.6.	Alternative RFI-suppression methods	79
6.6.1.	Alternatives to the median filter	79
6.6.2.	Matched filter	80
6.6.3.	Notch filter	82
6.7.	Conclusion	83
7.	Pulse detection	85
7.1.	First measurements at AERA	85
7.2.	Pulse parameter cuts	90
7.3.	Periodic background	93
7.4.	Constant fraction discriminator	98
7.5.	Coincidence trigger	100
7.6.	Conclusion	104
8.	Summary and Outlook	107
A.	Appendix	111
A.1.	Characteristics of the Fourier transform	111
A.1.1.	Fast Fourier transform	111
A.1.2.	Fourier transform for real signals	112
A.1.3.	Upsampling in time domain	112
A.1.4.	Hilbert transform	113
A.2.	Digital filters	114
A.2.1.	FIR filter	114
A.2.2.	IIR filter	115
A.3.	Hardware details	115
A.3.1.	The FPGA choice	115
A.3.2.	Interfaces of the AERA board	116
A.4.	Self-triggered cosmic ray events	118
	Bibliography	119
	List of Figures	127
	List of Tables	129

1. Introduction

When Victor Hess took off to his first balloon flights in 1911, he wanted to prove that the natural radioactivity originates from minerals in the Earth's crust, and diminishes with rising height. Instead, he found an increase of the radiation level, discovered the cosmic rays, and in retrospective constituted the topic of astroparticle physics. With ongoing direct measurements of the cosmic rays, the coincident detection of particles in several distant detectors was found in 1938 by Pierre Auger. He deduced that the coincident events were caused by secondary particles, created from a high-energetic cosmic ray in an extensive air shower. This indirect measurement of secondary particles with an array of surface detectors became the standard method for the detection of ultra high-energetic cosmic rays. In 1964 the radio emission from cosmic rays was discovered, followed by the fluorescence detection in 1976. While fluorescence telescopes could be established as high-precision instruments for the air shower development, the radio technique was almost forgotten. In the meantime the cosmic ray experiments grew bigger, the measured spectrum was expanded to higher and higher energies, and the theoretical background of cosmic ray physics was developed (see chapter 2).

Starting with the LOPES experiment in 2003, the method of radio detection was rediscovered, as the fast technological progress made a digital broadband measurement feasible. LOPES is co-located with the surface array KASCADE-Grande, which provides an external trigger and additional shower information about the detected events. When the first measurements and simulations showed that radio detection could combine the good energy reconstruction of fluorescence detectors with the high duty cycle of surface arrays, the interest in this approach dramatically increased and several new experiments were planned. To establish radio detection as a full-fledged alternative, a reliable self-trigger system is however mandatory. Amongst others, the LOPES^{STAR} extension was launched, to develop a suitable analog chain and a self-trigger strategy in a small scale. The obtained trigger algorithm was then implemented in available digital electronics as a prototype (see chapter 3).

To escape the radio noise of the industrial environment at the KIT, a measuring campaign with the digital self-trigger was initiated at the Pierre Auger Observatory in Argentina. After installation of a small antenna triangle in 2008, the measurement continued remote controlled for 15 months. Despite a system upgrade with a fourth antenna in the middle of the operating span, no cosmic rays could be detected. The noise level on site was still too high. Nevertheless, the campaign improved the understanding of the occurring radio noise, and led to requirements for a future system (see chapter 4).

With a survey measurement, a better suited location was found in the North of the Pierre Auger Observatory to install the first full scale radio experiment, the Auger Engineering Radio Array (AERA). Based on the experience of the prototype measurement, a made-to-measure digitizer system was developed. The new electronics is designed for scalability with a large number of antenna stations, including remote operation with wireless communications and a low power consumption for solar supply. Because the radio signal is known to be close to the detection limit, each station is equipped with a large ring buffer, which allows to read out untriggered, surrounding stations of a weak event. This substantially improves the data quality of low-energetic events, which make

up most of the statistics. It also provides the fall-back solution of external triggering (see chapter 5).

To get the most out of the environment conditions, a multi-stage trigger approach is used. At first, the digitized signal is pre-processed to suppress interferences from radio transmitters. Because of the background variability, a self-adjusting filter is implemented, consisting of a Fourier transform to frequency domain, a median filter to remove monofrequent contributions, and a transform back to time domain. To cope with the high data rate of the live data stream, the algorithm is optimized for maximum performance with the limited resources (see chapter 6).

After detecting radio pulses with a constant fraction discriminator, unwanted noise transients are rejected based on their pulse shape and pulse surrounding. Additionally, periodic noise can be predicted and rejected with the help of a custom-made phase-locked loop. When the trigger rate is low enough to comply with the limited capacity of the wireless communication system, the triggers are transmitted to the central data acquisition system. There, the coincidence trigger rejects noise pulses which stem from the horizon. The false trigger rate can be further reduced by dynamically suppressing hot spots in the sky (see chapter 7). The first cosmic ray events are finally identified in comparison with the surface detector of the Pierre Auger Observatory.

2. Cosmic rays

Beyond the classic visual observation of stars, which is as old as mankind itself, astronomy was expanded during the last century to the study of radiation with energies from below 10^{-3} eV like cosmic microwave background[1] up to ultra-high-energy cosmic rays (UHECR) at 10^{20} eV[2]. For radiation with energies above 10^9 eV usually the term *cosmic rays* is used. At this energy the sun fades out as the dominant radiation source, while radiation of galactic and extra-galactic origin is able to penetrate the solar system against the shielding solar wind. Cosmic rays consist mainly of protons, but also other particles like atomic nuclei, electrons, positrons, anti-protons, gamma rays, and neutrinos are included. In a strict sense, the term cosmic rays does not include photons and neutrinos, as they are believed to be only secondary products of the charged cosmic ray particles emitted by the cosmic accelerators. An even closer definition limits cosmic rays to the hadronic atomic nuclei, including the proton as hydrogen nucleus, as the hadrons account for the largest fraction of cosmic rays.

This chapter gives an overview of the current knowledge about the cosmic ray generation and propagation process, followed by the interaction in the Earth's atmosphere and development of an extensive air shower. Finally detection methods and cosmic ray experiments will be exemplified with an emphasis on the radio emission and detection of cosmic rays.

2.1. The source of cosmic rays

The origin of cosmic rays is under debate up to the present day. Unfortunately, the arrival direction of cosmic ray particles does not simply point towards their source, as they are deflected by magnetic fields in space due to their electric charge. In fact, particles with energies up to $5 \cdot 10^{18}$ eV are confined within the galaxy by the galactic magnetic field like in a giant particle storage ring [3], which leads to an isotropic direction distribution. Particle tracks at higher energies are less affected by the magnetic field and should reveal their origin. However, the particles of the highest energies are extremely rare, making the identification of the source location impossible due to the lack of statistics to date.

As a more indirect approach to the question of origin, theoretical emission models are compared with the spectrum of cosmic rays, which has been measured in detail over the last decades in numerous experiments. The spectrum reaches from an energy of 10^9 eV and a flux of hundreds of particles per m^2 and per second to very rare UHECRs at 10^{20} eV occurring only once per km^2 and century. The Spectrum matches a power law

$$\frac{dN}{dE} \propto E^{-\gamma}$$

with an almost constant spectral index γ between -2.5 and -2.7 over several orders of magnitudes (s. fig. 2.1), suggesting a similar acceleration mechanism for all energies.

The most convincing model is the repeated acceleration at relativistic shock fronts, also known as first order Fermi acceleration [5]: When a relativistic particle is reflected




Figure 2.1.: Energy spectrum of the cosmic rays, assembled from several exemplary experiments, showing a steepening of the spectrum at $E = 10^{15.7}$ eV ('knee') and a flattening at $E = 10^{18.3}$ eV ('ankle') [4].

by an approaching shock front, it gains energy similar to a tennis ball hit by a racket. Although the cosmic ray particle does not rebound from a solid surface, but is bent gradually by the moving magnetic field upstream the shock front, the result is the same. The energy gain is $\Delta E = (v/c)^2 \cdot E$ with the shock front velocity v and the relativistic particle energy E . With a certain probability the particle is magnetically reflected back and hits the shock front again. To reach its final energy the cosmic ray particle must bounce off the shock front several times, each time with a certain escape probability. This directly causes a power law spectrum with a negative spectral index compatible with measurements for plausible assumptions[5]. The gyro-radius $r_g \propto E/(Z \cdot B)$ of the accelerated particle must lie within the scale of the shock front, thus particle charge Z , magnetic field B and size of the shock front limit the reachable energy.

Appropriate shock fronts are created by supernovae, when during a stellar explosion large amounts of the star mass are ejected at relativistic speed. The material front expands for some hundred years and is visible as a supernova remnant. A smoking gun evidence for this theory was found in the form of gamma rays at energies up to 10^{14} eV from supernova remnants[6]. Cosmic rays might be accelerated to even higher energies, as the gamma rays are supposed to be secondary particles from the interaction of accelerated charged particle.

However, the scale of supernova remnants and their magnetic fields according to models, can only explain cosmic rays with energies up to 10^{15} eV – 10^{16} eV. This is supported by the fact that around this energy the flux of protons and other light nuclei decreases, the primary composition becomes heavier. This induces a steepening of the energy spectrum, the so-called 'knee', referring to the similarity with a human leg's shape [7]. A second, less distinct knee at $E = 10^{17}$ eV may be caused when the




Figure 2.2.: Sky-map in galactic coordinates of AGNs of the Swift-BAT catalogue smoothed with a resolution of 5° plotted in *reddish color*, and cosmic rays with energies above $5.5 \cdot 10^{19} \text{ eV}$ plotted as *black points*. An accumulation of events from the Centaurus A region is visible [9].

acceleration in supernova remnants reaches its maximum energy for particles of higher proton number. Another possible explanation is the escape of cosmic rays from the galactic confinement at higher energies.

At higher energies, promising candidates for violent shock fronts are active galactic nuclei (AGN): super-massive black holes in the center of most galaxies, which attract matter through a rotating accretion disk. Perpendicular to the disk, some matter is ejected at highly relativistic speed in so-called jets. Theory explains the jets with twisted magnetic fields, compressed in the center of the accretion disk. When the high magnetic pressure blows out, it accelerates a shock front over distances of many kpc. At high energies the models and theories become vague. The flattening 'ankle' of the spectrum at $E = 10^{18.3} \text{ eV}$ might be explained by the transition from galactic to extra-galactic cosmic rays. At least cosmic rays above $E = 10^{19.6} \text{ eV}$ are of extra-galactic origin, since their arrival direction was found to be correlated with neighboring galaxies (s. fig. 2.2), although this discovery decreased in significance lately [8, 9].

Cosmic rays at these energies are affected by GZK-suppression, an effect already predicted in 1966 by Greisen, Zatsepin, and Kuzmin [10] [11]: Due to the proton's highly relativistic energy, the interaction with a photon from the 2.7 K cosmic microwave background provides enough center of mass energy for an excitation of the $\Delta_{(1232)}$ -resonance, which decays again into a pion and a proton with 20 % less energy. Due to the steeply rising cross section with energy, the range of protons reaching the Earth with energies above $E = 10^{19.7} \text{ eV}$ is reduced below $\approx 100 \text{ Mpc}$. For heavier nuclei different mechanisms like photodissociation exist, providing similar energy limits [12]. This cuts the catalogue of possible sources down to relatively nearby objects, and should thereby strongly suppress the flux above the GZK energy threshold, which was finally confirmed by measurements [13] after being a controversial issue for almost a decade [14].

2.2. Extensive air showers

When a high-energetic particle enters the Earth's atmosphere, it collides with a nuclei of the thin air. This inelastic scattering creates secondary particles, which are still high-energetic, scatter again, and create more particles by themselves. The number of particles rises exponentially in a cascade of interactions and generates a so-called 'extensive air shower' [15]. Because of the large momentum asymmetry all particles




Figure 2.3.: Illustration of an extensive air shower developing in three components (*left*), and forming a 'pancake'-shaped shower front (*right*) [17].

are strongly boosted forwards and have similar speed and direction. This results in the concentration of secondary particles in a shower front with a thickness of only a few meters. The transversal speed is distributed randomly, causing a lateral shower extension of up to many kilometers. When after many generations of interactions the mean particle energy is too low for further particle production, their number drops again exponentially due to decay and absorption. Only for primary energies above $E_0 = 10^{14}$ eV the remains of the particle cascade reach ground level and can be measured by surface detectors.

Because of the atmospheric density variation with altitude according to the barometric formula, the mean free path of particles varies strongly. This is usually compensated by using the traversed matter in g/cm^2 as a length unit along the shower track. As the number of shower particles multiplies exponentially until reaching a critical energy of 85 MeV per particle, the maximum number of particles is proportional to the primary energy $N_{\text{max}} \propto E_0$, while the depth of the shower maximum grows logarithmically with energy $X_{\text{max}} \propto \ln(E_0)$. The shower development can be described empirically with the Gaisser-Hillas function [16], with the shower parameters X_0 for the first interaction position and an interaction length λ :

$$N(x) = N_{\text{max}} \cdot \left(\frac{X - X_0}{X_{\text{max}} - X_0} \right)^{\frac{X_{\text{max}} - X_0}{\lambda}} \cdot e^{\frac{X_{\text{max}} - X}{\lambda}}$$

The particles in an extensive air shower can be divided into three fractions (s. fig. 2.3):

The hadronic component originates from strong interactions in the early shower, the first interaction typically results in already more than 50 secondary particles, mainly hadrons. Although all kinds of heavy flavors are created due to the high energy, the production of light K- and π -mesons dominates. Besides causing an hadronic cascade, neutral pions decay into $\gamma\gamma$ -pairs starting the

electromagnetic cascade while charged pions and kaons generate the muonic component. Because of the losses into both other fractions, the hadronic component contains just 1 % of the shower energy at ground level and plays no role for shower detection. Nevertheless the early hadronic interactions have an important impact as a source of the further shower development and are the main cause for shower fluctuations and model uncertainties.

Muons are created together with μ -neutrinos in the decay of charged $\pi^\pm$ and $K^\pm$ mesons. Due to the penetration power of high-energetic muons, they hardly interact in the atmosphere and lose only little energy. Because of their relativistic speed, most of them reach the Earth's surface, despite their short lifetime of $2.2\,\mu\text{s}$. The muonic component carries about 5 % of the showers energy and is important for shower detection, as muons are a direct messenger from the shower maximum. In fact, extensive air showers can be measured by their muons only in underground experiments not even designed for cosmic ray detection.

The electromagnetic component incorporates about 90 % of the shower energy, and is fed by $\gamma\gamma$ -pair production of π^0 decay, muons decaying into electrons and different semileptonic meson decay channels. However, the strongest contribution is the self-preserving combination of $\gamma \rightarrow e^+e^-$ pair production and $e^\pm \rightarrow e^\pm + \gamma$ bremsstrahlung. This continues until reaching a critical electron energy of 85 MeV, when the energy loss due to ionization becomes dominant, and the slow electrons are absorbed by the air. Because the electromagnetic cascade virtually contains all shower energy, it is the most important component for all indirect detection methods.

The remaining few percent of the primary energy is lost in neutrinos, which cannot be detected efficiently. A proper estimate of the fraction of lost energy in the shower models is crucial for correct energy measurement.

For a detailed understanding of the air shower development, the cascade of particle interaction can be simulated with a Monte-Carlo approach, using established simulation engines like CORSIKA [18] or AIRES [19]. The simulation uses known cross-sections and branching ratios from measurements at particle colliders. Because the highest energies in the early air shower cannot be reproduced, extrapolations are used, which may introduce significant deviations [20]. The simulation of billions of secondary particles for each shower needs a lot of computing time, which can be greatly reduced with a hybrid approach by only simulating the first few interactions individually, and using empiric distribution functions for the low energy particles with high statistics [21].

Simulations show for example that the mass of the primary particle has a strong influence on the shower maximum X_{max} : Due to the higher cross-section of heavy nuclei, their first interaction happens earlier and the shower develops faster, which enables a primary mass sensitivity of air showers.

2.3. Detection of cosmic rays

During the history of cosmic ray physics, different detection methods have been developed. The discovery of cosmic rays by Victor Hess in 1911 was achieved with direct measurements of the ionizing radiation during balloon flights [22]. By observing the rise of the radiation level with altitude, he concluded that the radiation source must be extraterrestrial.

Modern direct measurements are conducted with multi-stage particle detectors, including track detection in a strong magnetic field and a calorimeter. This provides not only precise measurements of primary direction and energy, but also particle mass and charge. To avoid absorption or energy loss of the primary particles, the detector must be located above the atmosphere. This is either achieved by satellite experiments like AMS [23], or by mounting the detector on a high altitude balloon, typically drifting for several days in an altitude of 30–40 km like CREAM [24]. Satellite detectors are very limited in size and weight, need a long preparation and construction process and are very expensive. Balloons can lift a larger detector, and may be maintained and upgraded between flights, but have a large disadvantage in respect of observation time. Anyhow, the small detector size in the order of m^2 , prohibits direct measurements to energies below 10^{15} eV , as the particle flux becomes too low at higher energy.

Fortunately cosmic rays above 10^{14} eV can be measured indirectly by their remaining secondary particles at ground level. This method was discovered by Pierre Auger in 1938, when he observed coincidences of multiple particle detectors with distances of hundreds of meters [25]. He concluded, that the detected particles were caused by a single event, an extensive air shower.

Until today the standard method for the measurement of air showers is the use of a surface array of particle detectors. Each particle detector can consist either of scintillator material or a Cherenkov water tank read out by photomultiplier tubes. Typically electrons and muons are counted separately, for example by using an absorber for $e^\pm$ in front of the muon counter. Surface arrays are designed for certain energy ranges: For lower energies, the detector density is important, because the lateral distribution of the shower is only a few dozen meters wide, and the shower must be detected in several detectors to ensure proper reconstruction. For high primary energies the density is less important, a single shower spreads over several kilometers, but the total size of the array must be as large as possible for sufficient statistics of the rare events. For reconstruction of the shower position the barycenter of the measured signals is used, the direction can be determined by triangulating the arrival times of the shower front in each detector. The primary energy is reconstructed based on the number of detected secondary particles. By counting the number of electrons and muons separately, even the shower maximum X_{max} is estimated, as the number of muons stays nearly constant after the maximum, while the electron number drops rapidly. The main disadvantage of surface detectors is that they only measure the footprint of the air shower, a snapshot at the end of its development. Depending on the first interaction at X_0 , the snapshot may be earlier or later along the shower development, which creates a large energy uncertainty for a single shower. Although these shower-to-shower fluctuations are reduced in statistics, the energy and particularly the primary mass reconstruction suffer from relatively large uncertainties in the shower interaction models.

Another method is the detection of air showers using the ultraviolet fluorescence light emitted by nitrogen molecules, which are excited by the electromagnetic component. The light intensity is proportional to the electron density. Special fluorescence telescopes record the temporal development of the light track and thereby the trace of the air shower in real time. Because the whole shower development is measured, the primary energy reconstruction is much less model dependent and has low uncertainties, the shower maximum X_{max} can be read off directly. To receive a high light sensitivity at low exposure time, the pixels made of photomultipliers must be relatively large, the number of pixels is limited to some hundreds. This reduces both the angular and position-resolution, particularly for far away showers only visible in a handful of pixels. A single telescope can cover a large area with an observation angle of 30° and




Figure 2.4.: Map of the Pierre Auger Observatory: Red dots stand for surface detector stations, the fluorescence telescopes and their fields of view are plotted in green. The radio detection experiment AERA and its predecessor setups at the BLS¹ and CLF² are marked blue.

a detection range up to 50 km under optimal conditions. However, the duty cycle of a fluorescence detector is reduced to typically 10–15 %, as it can only measure during dark, moon-less nights with little cloud coverage.

The leading cosmic ray experiment for the highest energies above 10^{18} eV is currently the Pierre Auger Observatory [26], which is located in the province of Mendoza in Argentina. It combines a total of 27 fluorescence telescopes with an array of 1600 water Cherenkov detectors in a hybrid detector, covering an area of 3000 km^2 (s. fig. 2.4). This setup joins the advantages of both approaches: The large surface array provides enough statistics for the rare high-energetic events, while the fluorescence technology delivers detailed information about selected cosmic rays. The hybrid data improves the geometric resolution of the fluorescence detector and allows for the calibration of the surface detector, improving the whole experiment’s accuracy. Still, high quality data is only available for a small subset of events, and the precision of surface-detector-only data is limited by shower to shower fluctuations.

Other, less common approaches are air Cherenkov detectors like Tunka [27], using a photomultiplier surface array to measure the Cherenkov light of air showers, or to use existing neutrino detectors for air shower muon measurements besides their actual purpose, as done with IceCube [28]. Another promising approach to improve air shower measurements is the detection by means of their radio emission, as described in the following section.

¹ The Balloon Launching Station is used to start weather balloons.

² The Central Laser Facility is used for the calibration of the fluorescence telescopes.

2.4. Radio emission from cosmic ray air showers

Radio radiation from air showers was predicted [29, 30] and discovered [31] already in the 1960ies, but due to problems with the historic analogue recording technology and the fast success of the competing surface arrays, interest ceased almost completely in the following years. Today however, with the fast developing digital electronics, and the step towards higher primary energies, radio detection becomes an attractive field of observation again.

2.4.1. The radio emission model

The radio emission of an extensive air shower is produced according to the Maxwell's equations by the different accelerations of all its charged particles due to decays, interactions in matter and deflection in electromagnetic fields. Of course the contributions from electrons and positrons are absolutely dominant, as they are the most numerous particles, and they receive the strongest accelerations due to their low mass. Because the particles are concentrated in the thin shower front travelling at speed of light, the signals of all particles are integrated to a radio pulse in forward direction. The most straightforward and consistent description is achieved by adding up the electromagnetic contributions from each moving particle in a Monte-Carlo approach.

This is done with REAS³ for example, a simulation program expanding the existing air shower simulation CORSIKA⁴ by radio emission. In the latest version REAS 3 [34, 35] the particle tracks are discretized, and the end point contributions of all track segment are summed up. The result is a simulated E-field pulse at surface level, as shown in figure 2.5.

Another, more phenomenological approach, as followed by the simulation framework MGMR⁵, is to separate the different emission mechanisms into macroscopic contributions. This leads to a more graphic description, allowing for a better understanding of particular emission features. The predictions of REAS 3 and MGMR are compatible with each other [36], however MGMR uses some free parameters to weigh different




Figure 2.5.: REAS3-simulation of the radio pulse from an air shower. *Left*: East-west polarization of the time-dependent E-Field at different distances. *Right*: Signal spectrum comparing REAS3 (bold lines) and MGMR (thin lines) simulations [36].

³ Radio Emission from extensive Air Showers [32, 33]

⁴ COsmic Ray Simulations for KAScade [18]

⁵ Macroscopic Geo-Magnetic Radiation [37]

contributions, which reduces the significance of the agreement. The most important macroscopic signal contributions are assembled in the following list.

- The Earth’s magnetic field deflects the air shower electrons and positrons to curved tracks, an acceleration resulting in radio emission similar to the well known synchrotron radiation. However, current simulations suggest that this so-called **geosynchrotron effect** [38] only plays a minor role, contrary to older models [36].
- The geomagnetic field not only accelerates, but first of all separates electrons and positrons. This results in a varying electric dipole perpendicular to both the shower axis and the magnetic field. The dipole moment varies because each single electron is driven apart from its counterpart positron far and farther along the shower track, which increases their personal dipole moment. The flow of all electrons sums up to a macroscopic **transverse current**, and is the most important contribution to radio emission according to simulations [36].
- Additional **variation of the dipole moment** is introduced by the changing number of particles during air shower development. This augments the dipole growth in the early shower, and makes it fade away at the end. Although the charge of electrons and positrons absorbed in the atmosphere does not exactly disappear, it falls out of the moving shower front, and ends to contribute to the integrated radio pulse. The three above-mentioned contributions are caused by the geomagnetic field, and result in linear polarized emission perpendicular to both magnetic field and shower axis.
- Because of the higher cross section of positrons in air due to antiparticle annihilation, the shower front contains more electrons than positrons. This negative net charge excess varies along the shower axis and leads to a radial polarized radio emission, also known as **Askaryan effect** [29].
- Charged particles passing through a dielectric medium at a speed above the phase velocity of light in that medium emit **Cherenkov radiation**. The electrons and positrons of an air shower, emit most of their Cherenkov radiation at wavelengths much shorter than radio e.g. as visible light used in the field of gamma astronomy. The small fraction of Cherenkov radiation in the radio range plays no vital role for cosmic showers in air [39]. However, as shown by measurements with accelerator generated showers, Cherenkov radiation is a major contribution for particle showers induced by neutrinos or muons in dense media like ice, salt or lunar regolith [40]. Sometimes the term ‘Askaryan effect’ is mistakenly used for this radio Cherenkov radiation.
- An extensive air shower loses much energy in ionizing air and thus leaves behind a trace of tenuous plasma. Besides exciting fluorescence light emission, the low energy electrons in the plasma also cause **molecular bremsstrahlung** in the GHz frequency range. Unlike high-energetic bremsstrahlung, this microwave emission is emitted isotropically and partially coherent. This recently gained a lot of interest, because it would be possible to record the shower development from the side, similar to a ‘radio-fluorescence’ detector. Although the microwave emission mechanism was proven at accelerator-induced air showers [41], cosmic air showers were not detected yet.

The different mechanisms cannot be separated strictly from each other, neither in simulation nor in measurements. They influence each other, some of them overlap. But the general composition can be deduced due to different polarization or directivity patterns of the contributions.

Static electric fields in the atmosphere also have a significant influence on the radio emission. Although no effect was found for regular field strengths around 100 V/m, the increased fields above 10 kV/m during thunderstorms can greatly amplify the radio signal due to additional particle acceleration [42, 43]. This effect is currently under closer investigation, the static electric field is monitored, conspicuous periods of time are analyzed separately.

2.4.2. Properties of detected radio pulses

Due to the predominantly coherent emission mechanisms, the contributions of all electrons and positrons add up to a total radio signal amplitude proportional to the particle number, and thus proportional to the primary energy. This means the radio signal power is rising quadratically with shower energy, making radio detection particularly interesting at higher energies.

To ensure coherent emission, the radio source must lie in a region smaller than half a wavelength. With a typical shower front thickness of ≈ 1 m, containing all the electrons and positrons, this puts an upper limit to the radio frequency at ≈ 100 MHz. As the shower front typically gets wider at the outer edges, the radio signal should tend to lower frequencies at higher lateral distances. For measurements the relatively quiet frequency range between the noisy short wave band at 30 MHz and below the FM band at 80 MHz is particularly interesting. Although the first contributions to the radio signal stem from an altitude of several kilometers, and the last fraction is emitted at surface level, the radio signal is compressed by the rapidly moving shower front to a narrow and high pulse. Based on the shower front thickness, the pulse width can be estimated to be in the order of 30 ns. At lateral positions, the compression and thus the coherence of the pulse is less perfect, because the air shower propagation direction and the radio signal path differ. This leads to a larger pulse width and an exponential drop in amplitude with increasing distance to the shower core.

The properties of radio emission lead to a number of assets and drawbacks when comparing radio detection with other methods:

- As radio waves are not absorbed in air, the complete integrated radio signal reaches the surface and can be detected. This **calorimetric energy measurement** is a large advantage over a surface array, which measures only remaining particles from the shower maximum. In particular model dependencies and shower-to-shower fluctuations are reduced, the energy reconstruction is potentially on a par with fluorescence detection.
- The air shower **core position** is estimated with the barycenter of the measured signal amplitudes at each antenna, with their pulse arrival times the primary direction is reconstructed. Close enough to the shower core the resolution can be further increased with interferometry, as all antennas record the same wave train from the same origin [44]. The angular and position resolution exceeds the resolution of a fluorescence telescope and is similar to a surface detector.
- Because radio detection is independent from the time of day or season, its **duty cycle** knows no immanent limitation. Maybe data taken during thunderstorms

must be analyzed separately due to the large influence of strong electric fields, but the duty cycle would still stay above 90 %, rendering a huge benefit in statistics over fluorescence detection. Strongly inclined air showers have to traverse a much longer distance until reaching surface. As only few secondary particles survive, they are hardly detectable with surface detectors. The missing absorption of radio waves potentially enables radio detectors to greatly increase the accepted solid angle. Even the detection of horizontal air showers, induced by high-energetic neutrinos grazing the Earth, may be possible.

- To deduce the **mass of the primary particle**, an exact measurement of the shower maximum X_{max} is necessary. Unlike secondary particles, the radio contribution from the shower maximum is not absorbed in the atmosphere, but still included in the radio signal. Unfortunately the amplitudes from different shower ages are not recorded separately like for a fluorescence telescope, but as an integrated pulse, meaning that the depth of the shower maximum is accessible only indirectly. Recent analysis suggests however, that the shape of the radio front, measured as the lateral pulse arrival distribution, heavily depends on X_{max} , maybe increasing the primary sensitivity to the level of fluorescence detection [45].
- Another advantage of a radio detector is its construction with relatively simple and **cost effective** components: Besides the digital readout electronics, needed for any detector concept, it requires only an antenna with preamplifier, in contrast to more expensive particle detectors. The potential cost advantage still must be proven for a large scale array, but it might be crucial for future experiments.
- A disadvantage of the radio detection method is however, that as parts of the emission mechanisms are based on interaction with the geomagnetic field, the signal strength is **depending on the primary direction**. Parallel to the magnetic field for example, the detector has a blind spot, or at least an area of increased threshold. The shape of the sensitive solid angle is more complicated than just losing efficiency towards horizon, which complicates a proper determination of the detector exposure. This problem also overlaps with angular gain variations of the used antennas: Unlike particle detectors for example, the receiving pattern of any antenna is shaped unevenly, divided into lobes and nulls. Also variable reflections from the ground and the surroundings may have an influence, even for a specially selected antenna design. In summary, the response function of the full radio detector is rather complicated, a reliable amplitude measurement needs careful calibration, monitoring, and reconstruction, otherwise the above mentioned advantages in energy resolution are quickly lost.
- A big problem for radio measurements is always the **noise level**: Because only a tiny fraction of the primary air shower energy is finally contained in the radio pulse, already the inevitable galactic noise limits the observation of radio pulses from cosmic rays to energies above $5 \cdot 10^{16} \text{eV}$. Additional interference is caused by broadcast stations, but also by unintended noise from electrostatic discharge at power lines or any electric device. This background further decreases the signal to noise ratio, which makes event reconstruction difficult, and even worse, hinders self-triggered radio measurements. To establish radio detection as a real alternative to existing methods, the development of an effective self-trigger is however imperative, which is the major topic of this thesis.

As the information obtained from the radio signal, which is dominated by the shower maximum, and the data from a surface detector, which describes the shower remains, are complementary to each other, radio and particle detectors are particularly suited for hybrid measurements. Additionally the radio self-trigger difficulties can be avoided with an external trigger from the surface detector, which is the reason why most of the radio experiments to date follow this approach.

2.4.3. Radio detection experiments

The first measurement of the radio emission of air showers was already achieved in 1964, at the Jodrell Bank observatory in Great Britain [31]. At that time, the radio pulses were recorded with analogue oscilloscope photography and externally triggered by a triangle of scintillators. The measurement used a relatively narrow bandwidth of 2.75 MHz at 44 MHz, using the frequency of BBC radio during nightly sign-off. Although the measurements obviously benefited from the historic low background level at night, radio detection could not be widely established, as the photographic film development and analysis stayed relatively expensive. Radio measurements kept a marginal existence as minor extensions of surface detector experiments in the following decades, although achieving some remarkable results [46]. For example the presumed highest energetic radio event so far, with an estimated primary energy of $2 \cdot 10^{20}$ eV, was detected at the Yakutsk air shower array [47] in the 1980ies. Measurements from that time all suffer from measuring in a relatively narrow frequency range, reducing the recorded amplitude of the wideband radio signal. With advances in digital electronics within the last two decades, wideband measurements became feasible, leading to the modern cosmic ray radio experiments. Some of them some will be described in the following.

Starting in 2003, the **LOPES** experiment was one of the first modern radio detectors, using the frequency range from 40 to 80 MHz. In the final expansion stage LOPES consists of 30 dipole antennas measuring together with and triggered by the well established surface array KASCADE-Grande [48] at the northern campus of the Karlsruhe Institute of Technology (KIT). With the advantage of high quality shower data from KASCADE-Grande to each radio event, the LOPES collaboration strongly contributed to the progress of understanding shower features like 3-dimensional polarization, energy dependence, and lateral distribution [49, 50, 51, 45]. Also the Monte-Carlo simulation tool REAS [35] was developed, and the radio reconstruction with interferometric beamforming was proven at the LOPES experiment [44]. However, the data suffers from the high background level of the industrial environment. The low signal to noise ratio makes precision measurements difficult. The extension LOPES^{STAR} (LOPES Self Triggered Array) is used to develop electronics and methods for following self-triggered measurements [52], leading amongst others to the Auger Engineering Radio Array (AERA) experiment, which will be specified in detail in chapter 5. As the LOPES^{STAR} setup was the starting point of this work, it will be described in closer detail in the next chapter.

LOPES (LOFAR prototype station) itself is, as the name suggests, just a pilot test, proving the feasibility of cosmic ray detection with digital beamforming, for the **LOFAR** (Low Frequency Array) radio telescope [53], which is currently in its last stage of construction. LOFAR uses digital interferometry of hundreds of omnidirectional antennas at stations spread over many kilometers in the Netherlands and neighboring countries. Amongst classic radio astronomy one of LOFARs key tasks is the radio detection of cosmic ray air showers with support of a small particle detector array. With the **NuMoon** experiment [54] the LOFAR observatory is also used for an alternative

approach, looking for radio pulses from cosmic ray and neutrino interactions in the lunar regolith, when the primary particles skim the surface of the moon.

At the same time like LOPES, **CODALEMA** (Cosmic ray Detection Array with Logarithmic Electro-Magnetic Antennas) started measuring at the radio observatory Nançay in France. Profiting from the local low noise environment, and measuring with a large bandwidth from 1 to 200 MHz, CODALEMA made interesting precision measurements [55, 56], and became one of the leading groups in radio detection. A disadvantage, however, is the relatively simple scintillator detector used for triggering, which can only deliver coarse shower information. The CODALEMA group also extended their measurements to the Pierre Auger Observatory, and co-founded the AERA experiment.

The success of LOPES and CODALEMA caused an increasing interest in radio detection, leading for example to retrofits of existing detectors with radio antennas, like at the **Tunka** detector [27] with the goal to improve the X_{max} -resolution. Another example is the project **Radio at IceCube** [57], where the particle detector array IceTop is upgraded with radio antennas: IceTop is intended as Veto for the cubic-kilometer neutrino detector, in order to distinguish unwanted cosmic ray muons from neutrinos, cosmic ray measurements are a welcome by-product. The low-cost radio extension of IceTop to a larger area can improve both the cosmic ray reconstruction, and the Veto for inclined showers which do not trigger the smaller IceTop detector.

Besides cosmic ray measurements, radio detection is also an interesting approach to detect neutrino-induced particle showers. Because of the low neutrino cross section the first interaction and at least the early part of the particle cascade takes place in dense media like rock or ice, where the radio emission is dominated by the Askaryan effect instead of geomagnetic emissions. Although this mechanism was proven in accelerator laboratories, no neutrinos were radio-detected so far. An example for the neutrino detection efforts are the **RICE** experiment [58], where together with the regular optical IceCube detector modules, also radio modules were buried in the ice shelf, and its planned successor **ARA** (Askaryan Radio Array) [59]. Another approach is to detect the radio emission, instead of inside the ice, from above the ice-shelf with the balloon-borne detector **ANITA** (Antarctic Impulse Transient Antenna) [60]. However, no up-going neutrinos were detected to date, but only down-going cosmic rays, whose radio emission was reflected by the ice. In contrast, the **TREND** experiment (TIANSHAN Radio Experiment for Neutrino Detection), located in a remote valley in XinJiang, China, follows the construction of regular radio arrays. It was constructed to measure horizontal air showers, created by neutrinos penetrating the surrounding mountains. The key problem of this self-triggered measurement is the reliable discrimination of neutrino-induced radio pulses from noise pulses, which also predominantly originate from the horizon. Until now, only cosmic ray showers with higher elevation were detected [61].

3. System development of a radio self-trigger

Most of the radio measurements so far were made with relatively small antenna fields, using an external trigger from local particle detectors. The next step towards fully establishing radio detection as an equivalent alternative, is a radio array with an area of several square-kilometers, measuring self-triggered and potentially stand-alone. The self-trigger feature is important on the one hand to profit from the nonexistent radio absorption by measuring highly inclined cosmic rays, which are hardly detectable for a surface detector, and on the other hand for the prospective construction of a cheap radio-only detector. In 2005, the Institute for Data Processing and Electronics (IPE) started the extension $\text{LOPES}^{\text{STAR}}$ of the LOPES experiment, to develop a complete system for self-triggering, from suitable antennas over the analogue chain to the actual trigger strategy. This research is the starting point of the thesis at hand, and is described in detail in the following chapter. It led to a digital trigger system, first tested as a prototype, then enhanced to a final system deployed at the AERA experiment.

3.1. Measurements with $\text{LOPES}^{\text{STAR}}$

The $\text{LOPES}^{\text{STAR}}$ experiment consists of 11 antennas, located south of the actual LOPES experiment (s. fig. 3.1). The location outside of the dense KASCADE array improves the noise background compared to LOPES, where the local photomultipliers cause many transients due to their high voltage recharge after each air shower. $\text{LOPES}^{\text{STAR}}$ is divided into three clusters, each one arranging its antennas in almost equilateral triangles with an edge length of 65 m to 70 m. The regular triangle configuration improves later direction reconstruction and simplifies the setup of a coincidence trigger. The electronics of each antenna field is housed inside the nearby existing container from the Grande detector. The subsequent numbering of the Grande containers is also used for the naming of the antenna fields, the clusters D17, D19, and D30 are respectively close to container 17, 19, and 30. All antennas are dual polarized in east-west and north-south direction, thus providing a total of 22 channels. The $\text{LOPES}^{\text{STAR}}$ experiment uses the free areas between KIT buildings as a guest, therefore it is limited in space and sometimes suffers restrictions by local constructions. For example, the D19 cluster was measuring with two antennas instead of three until 2008 because of a close-by construction site. The D30 field was dismantled at the end of 2010 to make room for a new building.

The KASCADE Grande experiment not only provides housing for the electronics, but more importantly a fast external trigger and information about the detected air showers including primary energy and direction. KASCADE Grande covers the energy range between 10^{16}eV and 10^{18}eV , while $\text{LOPES}^{\text{STAR}}$ becomes fully effective at $10^{17.5}\text{eV}$ [62]. The Grande trigger is generated centrally, when neighboring Grande stations detect secondary particles coincidentally. It is organized in 18 trigger clusters, each consisting of 6 stations in hexagonal shape plus one central station [63]. To trigger $\text{LOPES}^{\text{STAR}}$, all 7 stations from hexagon 13 (s. fig. 3.1) are required. This guarantees that a major




Figure 3.1.: Map of the cosmic ray detectors at the Karlsruhe Institute of Technology. The KASCADE air shower detector covers the square in the northeast corner, its upgrade Grande is marked with blue squares. The LOPES experiment is indicated by red triangles, where triangles pointing upwards stand for east-west polarized antennas, downward triangles for north-south polarization, and stars for dual-polarized antennas. The green hexagon shows the time calibration beacon. The antennas of the LOPES^{STAR} extension are marked by black stars, the used Grande-trigger hexagon 13 is indicated by the dashed blue lines.

part of the air shower falls within the array, ensuring a proper reconstruction quality. The trigger is delivered to each antenna field via glass fiber as a simple strobe signal. The resulting trigger rate is around 0.05 Hz, however, due to the small overlap of the energy ranges of LOPES^{STAR} and Grande, most of the recorded events do not contain a clear radio signal, as they are too low-energetic.

3.1.1. Antenna details

The first component in the signal chain is the antenna, which has to fulfil a several requirements to properly detect the radio signal of cosmic rays.

- The radio signal induced by a cosmic ray consists of a short pulse, which is thus wide in frequency domain. Due to its coherence, the recorded pulse height is achieved by integrating the amplitudes spectrum over the used frequency range, which means that the pulse amplitude rises linearly with the band-width for a flat spectrum. In contrast, the uncorrelated background noise is determined by integrating over its power spectrum, thus the noise level only increases with the square root of the band-width. Therefore broadband recording improves the signal to noise level. Because of the large background contributions from the short wave band below 30 MHz and the FM band above 80 MHz, radio detection is usually limited to the frequencies in between. The antenna therefore should be **broadband**, and to achieve a high angular shower acceptance, also have a large beam-width. Antenna gain should be constant and dispersion low for all frequencies and all incoming angles. This is important, although known gain patterns and dispersion can in principle be reconstructed afterwards: When the initial direction is unknown, the correction is applied iteratively, starting with a rough uncorrected direction estimation. For the convergence of this approach, rather flat antenna patterns are very helpful.
- Besides the wide main lobe, the radiation pattern of the antenna should show **backward attenuation**, otherwise the signal reflected by the surface has a strong impact on the antenna response. This perturbation is hard to correct with the antenna calibration, as the reflectivity of the ground depends on its conductivity and relative permittivity according to the Fresnel equations, and varies from ≈ 0.3 for dry sand up to ≈ 1 for plashy ground [64].
- An obvious requirement for the antenna is a sufficiently **high sensitivity**, as radio pulses may be close to the detection threshold. First of all this requires the internal noise from the preamplifier to be small compared to the external noise. At good, radio-quiet locations, the background then is dominated by the inevitable galactic noise with a noise temperature of 5000 K at 60 MHz [65].
- In addition to the electrical properties the **sturdiness** of the antenna is an important factor, as it may have to withstand high wind speeds up to 200 km/h during storms. The antenna should also be reasonably priced and easy to handle and mount.

During the development of the self-trigger system three different antennas were proposed and tested, which will be characterized in the following. All the antennas are linearly polarized. To record both polarizations, two identical antennas rotated by 90° are arranged together.




Figure 3.2.: *Top:* Pictures of the antenna candidates. *Bottom:* Internal antenna schematic, all antennas feed the coaxial cable via balanced-unbalanced transformers.
a) Wired LPDA b) SALLA c) Butterfly antenna

The antenna type used for LOPES^{STAR} from the beginning is the Logarithmic Periodic Dipole Antenna (**LPDA**), as shown in figure 3.2a. The LPDA is a broadband directional antenna, well known and widely used in the field of telecommunication. In simplified terms, it consists of a stack of $\lambda/2$ -dipoles, each one getting in resonance at a different frequency, while the longer dipole behind works as a kind of reflector, and the shorter dipole in front as a director. For the lower frequency of 30 MHz, this calls for a dipole length of 5 m, resulting in a large total antenna size. The LPDA has several design degrees of freedom, like the number of dipoles and the slope of dipole length and distance, to tune it towards the desired properties. The result is a wide main lobe with an opening angle of 60° and a backward attenuation of 15 dB (s. fig. 3.3). The frequency response is rather flat (s. fig. 3.4), with small bumps at the resonances of each dipole. The impedance of the antenna remains real over the full frequency range, strongly simplifying the subsequent impedance matching. The LPDAs at LOPES^{STAR} were build of aluminum profiles for the carrying structure and aluminum tubes for the dipoles. In strong wind however, the dipoles tend to swing resonantly, and soon suffer fatigue breakage. An improved LPDA with similar electrical properties was built using flexible wires for the dipoles, spanned like a harp between glass fiber spreaders. This so-called wired LPDA proved to resist wind and weather for several years now [66]. The main disadvantage of the LPDA however remains: its sheer size and complexity makes mass production difficult and expensive.

A much simpler approach is the Small Aperiodic Loaded Loop Antenna (**SALLA**), which is basically a dipole operated below its resonance and damped with a resistor to widen the resonance curve (s. fig. 3.2b). By bending the dipole to a circle, with the resistor at bottom and the feed on top, the SALLA achieves a backward attenuation of 15 dB with an upward main lobe wider than the LPDA (s. fig. 3.3). The dominant resistor also keeps the antenna impedance real for all frequencies, which makes matching of the preamplifier easy. Towards lower frequencies, and therefore further away from




Figure 3.3.: Radiation patterns of LPDA, SALLA, and Butterfly antenna depending on the zenith angle at 60 MHz in free-space normalized to their upward gain. Solid lines show polarization in E-Plane, dashed lines in H-plane.




Figure 3.4.: Frequency response of LPDA, SALLA, and Butterfly antenna for upward direction. To fit in the same diagram, the active Butterfly antenna is shifted by -30dB.

its resonance, the antenna gain drops however (s. fig. 3.4). The loop diameter defines frequency characteristics and gain. First SALLA prototypes had a diameter of 75 cm, the final version has a size of 120 cm, increasing the gain and particularly improving the low frequency behavior. Due to its compactness and low number of parts, the SALLA is very robust, easy to manufacture and cheap. A problematic issue with the SALLA is its low sensitivity: On the one hand, the large beam-width also immanently increases the picked up external noise, but mainly the resistive load reduces the gain, so that the internal preamplifier noise may become dominant. Therefore the question whether the SALLA is able to detect the galactic noise modulation during a sidereal day becomes important. According sensitivity studies of the SALLA are described in section 4.2.

The third antenna proposal is the **Butterfly** antenna [67], consisting of a simple dipole with triangular shaped arms and a span width of 2.1 m (s. fig. 3.2c). Being an undamped dipole, the Butterfly shows a distinct resonance, where the complex impedance swaps from capacitive to inductive. To reduce the antenna insertion loss from the variable impedance, a suitable matching network is necessary, which is mounted directly at the feed point together with a 30 dB preamplifier. The Butterfly is therefore an active antenna, rendering antenna calibration measurements with the two-antenna-method impossible, and making antenna simulations more important. According to these simulations, the fat dipole arms slightly improve the broadband characteristics and widen the main lobe, compared to a classical dipole. Similar to the SALLA, the Butterfly's simplicity makes it robust and easy to build. As a non-resistive antenna the Butterfly's sensitivity is high, although the missing backward attenuation introduces also high ground sensitivity. To get defined ground properties, an additional conductive metal mesh should be placed on the surface.

The antenna of choice for the LOPES^{STAR} experiment was the LPDA, the D19 cluster (s. fig. 3.1) was equipped later with SALLAs. The first stations of AERA also use LPDAs, the antenna type for the final extension will be the Butterfly.




Figure 3.5.: Block diagram of the LOPES^{STAR} signal chain. Bold lines correspond each to two polarization channels of the analog radio signal, thin lines are digital connections.

3.1.2. The signal chain

The first device behind the antenna is the preamplifier, also called Low Noise Amplifier (LNA) because its noise figure is very critical due to the very low amplitude of the unamplified antenna signal. To avoid further signal loss, the LNA is placed as close to the antenna as possible. A good linearity is also important for the LNA to avoid intermodulations of frequencies with strong amplitudes outside the used band. For the LPDA, which itself works as a first order band filter, this is less critical than for the SALLA, which must be equipped with a dedicated filter of low order to attenuate FM frequencies. At LOPES^{STAR} the LNA provides a gain of 22 dB, before the signal is fed to the data acquisition through a 100 m long coaxial cable with an attenuation of about -5 dB (s. fig. 3.5).

At the data acquisition (DAQ) the signal is at first filtered with a 32nd order band-pass filter between 40 and 80 MHz and further amplified by 37 dB. The filter removes both shortwave and FM radio interference from the measured signal, but its high steepness and order are mainly required for the function as an anti-aliasing filter for the subsequent digitizing. The digitizing is performed with 12 bit resolution and a sampling rate of 80 MHz, using sub-sampling in the second Nyquist domain, as described in section 3.1.4. To strictly comply with the sampling theorem, the analog filter must suppress the signal at the frequency borders below an amplitude of one LSB (least significant bit). As the filter is designed to only have its -3 dB points at 40 and 80 MHz, minor ambiguities are hazarded. The ADC¹ samples are written into a small ring buffer. If an event is triggered, 2048 samples are read out with a predefined pre- and post trigger time to correct for trigger latency. The event is tagged with a timestamp from a GPS clock² and stored to a database for later analysis. All gain factors in the signal chain as mentioned above are only approximate values, the exact gain values are frequency dependent, during analysis an accurate calibration corrects for this. Filter-amplifiers and ADCs for a full field of 8 channels are housed conveniently in a small VME-crate³, a regular PC handles the DAQ.

¹ Analog Digital Converter.

² Global Positioning System - Satellite system providing precise location and time.

³ Versa Module Eurocard - a standardized commercial bus system.

3.1.3. Analog self-trigger investigations

The first attempt of cosmic ray detection with the LOPES^{STAR} setup was made with an analog self-trigger in 2005. In this trigger approach the first step is to generate the envelope of the high frequency signal and add both polarizations of each antenna. The joint antenna envelope is then fed through a high pass filter, which lets pass pulses, while rejecting continuous background contributions. Then a fixed threshold of around $150\text{ }\mu\text{V/m}$, when translating into local field-strength, is applied, resulting in a trigger rate of some kHz for each antenna. Random triggers of single antennas are suppressed with a coincidence condition between the antennas in a triangle configuration. By demanding a coincidence time lower than the signal propagation time within the triangle, horizontally incoming pulses can also be rejected, when assuming a large source distance and thus a planar radio front. When also ignoring multiple pulses within $1\text{ }\mu\text{s}$ this leads to a coincidence trigger rate between 1 Hz and 100 Hz.

This setup was measuring for a full time of 9 months, but no air shower pulses could be detected in coincidence with KASCADE-Grande. Details about these measurements including a description of the full analog chain and trigger can be found in the thesis of O. Krömer [62].

The major problem of the analog self-trigger was the background situation. Assuming Gaussian distributed noise amplitudes, the threshold crossing rate should drop exponentially when rising the threshold value, meaning that already a small threshold increase can drop the rate significantly. Unfortunately, the noise level varies slightly over time, leading to huge changes in trigger rate, when using a fixed threshold. A threshold regulation is therefore necessary for stable trigger rates. Additionally the exponential drop of trigger rate with threshold works only at rather high rates, when the threshold level still lies within the noise floor. At higher amplitudes, man made pulses dominate, so that a threshold increase only slightly reduces the rate. The fact that background does not only contain uncorrelated white noise seems obvious for urbanized environments as at KIT, but at a different level holds also true at rural areas as seen in section 4.2.

Also the fact that the pulses are not at all detected independently by the different antennas, strongly increases the coincidence rate. Even horizontal events from close-by sources can pass the coincidence constraint due to their curved signal front. The sum of these background misconceptions required a too high threshold value, rendering air shower detection impossible. The measurements revealed that a more sophisticated self-trigger approach was necessary.

3.1.4. Offline self-trigger development

As a consequence, LOPES^{STAR} was equipped with an external trigger from KASCADE-Grande in autumn 2006, each trigger causing the readout of an $25.5\text{ }\mu\text{s}$ (2048 samples) long radio event. Because the energy threshold for radio detectable events is much higher, than the threshold of the external trigger, the trigger rate of about 0.05 Hz is much higher than the expected cosmic radio pulse rate of 1 per week. Thus most triggered events contain only background data, the trigger is just an effective event pre-selection, enriching the data with cosmic events. The recorded data was used in the following to learn more about air shower radio pulses, with the long-term objective to build a digital self-trigger to overcome the restrictions of the analog trigger system.

Upsampling

The first step in the digital analysis chain, is to correct for the effects of the sub-sampled digitizing. According to the generalized Shannon-Nyquist theorem, the band-limited analog signal does not need to be sampled above twice the maximum frequency at 160 MHz, but can also be sub-sampled with 80 MHz, the double band-width. Thereby the frequencies between 40 and 80 MHz are mirrored as an alias down to the baseband. As the baseband between 0 and 40 MHz is empty due to the analog filter, no ambiguities occur, the signal can be reconstructed without losing information, a higher sampling rate would just introduce redundancy. For the analysis however, additional sampling points in between must be reconstructed by upsampling, to fully capture the signal shape (s. fig. 3.6). In particular the average signal amplitude is raised, as the probability of having a sampling point close to an oscillation maximum rises with the point density. This is especially important for short pulses containing only few oscillations. There are two methods for upsampling:

The straightforward method calculates the spectrum of the signal by using a discrete Fourier transform, normally with the Fast Fourier transform (FFT) algorithm. With a sampling rate of 80 MHz the spectrum spans from 0 to 40 MHz. This alias spectrum is mirrored up to its actual range between 40 and 80 MHz, and the baseband below is filled with zeroes, as it is known to be empty due to the filter. Thereby the number of points is doubled. By using the knowledge, that the signal does neither have components above 80 MHz, as many zero points as desired may be filled at higher frequencies. When the number of points in the spectrum is increased by a factor of z , also the number of sampling points in time domain are multiplied by z , after using an inverse FFT. As the covered time range remains the same, the time resolution is increased. Usually powers of two are used as upsampling factor z to use the FFT algorithm most efficiently.

The second method operates in time domain only: Between each two values of the time series, $z-1$ zeroes are inserted. These fast signal jumps introduce high frequency




Figure 3.6.: Example of the raw radio signal sub-sampled with 80 MHz (*black triangles*). The high frequency oscillations of the original signal are not directly visible until upsampling by a factor of 8, which adds intermediate sampling points (*black squares*) [52].

components. As shown in detail in appendix A.1.3, these frequency components are just mirror images of the original spectrum. By applying a digital band-pass filter to the desired frequency range, the upsampling process is completed. Because this algorithm works as a pipeline, and not block-wise as necessary for the FFT, it is perfectly suited for processing continuous data streams. Therefore the first method is usually used in high-level event reconstruction, when the data is already cut into event blocks, and a Fourier transform is necessary anyway for frequency dependent calibrations. The second method is preferably employed at an early level of real time data taking.

Suppression of radio frequency interference

At the site of LOPES^{STAR}, narrow band radio carriers are a major background contribution, as visible in the radio spectrum in figure 3.7. The most dominant signals stem from television broadcast of different channels by the transmitter Raichberg. Besides these strong and continuous lines in the spectrum, lesser contributions are randomly visible close to 40 MHz from the ISM⁴ band, and at 74 MHz from the BOS⁵ band. Despite their small bandwidth, these carriers significantly increase the noise floor.

To suppress this Radio Frequency Interference (**RFI**), a digital signal processing chain was developed: First the radio signal of each recorded event is switched to frequency domain using a FFT. Then the amplitude spectrum is flattened with a moving median filter, the median is chosen, because under all possible averaging methods, it is the most robust towards outliers as expected by carriers. The width of the moving median window must be larger than the double carrier bandwidth, which is about 5.5 MHz for strongly modulated television transmitters. The median filtered amplitudes are used for the whole spectrum, which showed a better background rejection than only modifying frequencies with high excess amplitude. As it does not matter, in which direction the




Figure 3.7.: Dynamic radio spectrum at the site of LOPES^{STAR} antenna field D17 in east-west polarization.

⁴ The Industrial, scientific and medical bands contain 12 frequency ranges from MHz to GHz, including a frequency at 40.7 MHz, used for short range communication and open for everybody.

⁵ The BOS-band (Behörden und Organisationen mit Sicherheitsaufgaben) is used for radio communication by police and emergency services in Germany.

median filter runs over the spectrum, the RFI-suppression can be as well applied before upsampling on the mirrored alias spectrum.

The short radio pulse from an air shower is broad in frequency domain (s. fig. 2.5), and thus hardly affected by the median filter. The pulse information is actually hidden in the phase, causing all the different frequencies to pile up exactly at pulse time. By removing frequencies dominated by external transmitters, the pulse amplitude is not necessarily reduced, as the phase of these contributions is determined by their transmitter, not the pulse. Thereby these frequencies are a kind of uncorrelated noise to the pulse, which may increase or decrease pulse amplitude randomly. For pulse height reconstruction of course, the reduced effective band-width must be taken into account.

The effect of the RFI-suppression obviously depends on the noise situation on site, the stronger the mono-frequent carriers are, the more the signal to noise ratio is improved. Without frequencies to suppress, the noise level may even be slightly increased due to numerical noise. A comparison of different RFI-suppression approaches is given in chapter 6. At the site of LOPEs^{STAR} the signal to noise ratio of radio pulses was improved by more than a factor of 4 in power until November 2007. Since then the television broadcast was switched from analog PAL⁶ modulation to digital DVB-T⁷, the background situation improved abruptly. Although RFI-suppression is not indispensable anymore at LOPEs^{STAR}, it is important at other locations like at the Pierre-Auger Observatory (s.fig. 6.1).

Envelope generation

When the radio pulse from a cosmic ray as shown in fig. 2.5 is detected, it stimulates an oscillation of the analog band filter with a frequency equal to the band center. To remove the high frequency oscillation and access the shape of the filtered pulse, the signal envelope must be calculated. The classical way to create the envelope is to rectify the alternating signal, before low-pass filtering it to get rid of the high frequency ripple. The dimensioning of the low-pass is always a tradeoff between remaining ripple on the one hand and lag smearing at fast signal changes on the other hand. The amplitude of such an envelope can be at most as high as the maximum of the original signal. However the actual envelope of a short pulse can be higher, when its maximum falls just between two maxima of the oscillation.

A more sophisticated way to reproduce the envelope, can be done by means of the Hilbert transform, which in Fourier domain shifts positive frequencies by -90° , and negative frequencies by $+90^\circ$. For a real signal this means that an oscillation $\sin(\omega t)$ becomes $-\cos(\omega t)$ (s. appendix A.1.4), thereby each oscillation zero point becomes an oscillation maximum. This does not correspond to a simple time shift, each frequency contribution of a signal is shifted by a quarter of its period. Thus the frequencies superimpose differently, pulse position and amplitude are slightly modified. In principle, Hilbert transform calculates the missing imaginary part for a real signal to get the complex analytical signal, a common procedure in telecommunications engineering. The signal envelope then is the absolute value of the analytical signal thus equals the quadratic sum of Hilbert transform and original signal.

⁶ Phase Alternating Line is a former analog television modulation standard

⁷ Digital Video Broadcasting Terrestrial is successor of the analog television, using the available band width more efficiently. The VHF I band (47 - 68 MHz) is not used anymore.

Pulse detection

After improving the signal to noise ratio and calculating the signal envelope, the actual pulse detection is the next step. The aim is an algorithm using only the radio data and finding all cosmic rays, which are known from KASCADE-Grande to be high-energetic and close enough to potentially cause a radio pulse. At the same time the number of false detections should be as low as possible. Most externally triggered events do not contain a significant radio pulse, while others have several pulses due to background noise. Therefore a threshold is applied, each time the signal envelope crosses it, one pulse is counted. A fixed threshold value would be highly desirable, as it corresponds to a constant detection limit. However the background conditions vary with antenna position, time of day, weekends, and season so strongly, that a dynamic threshold regulated to the noise level is necessary. Of course, the elevated threshold during noisy periods significantly decreases the detection sensitivity, yet it is the only way to keep the background pulse at a reasonable rate, besides switching off data taking completely. Due to the noisy background this results in a projected real time trigger rate of 24 kHz, when dividing the number of detected pulses by the summed measuring time.

The rate can be reduced based on the pulse characteristics. Cosmic ray induced pulses are for example known to be short and steep, while most background pulses are rather wide. Additionally background pulses often occur in multiple bursts under elevated noise conditions. With a set of cuts on suitable pulse parameters, the trigger rate is lowered to 600 Hz, when projecting real time measurements.

Although the east-west polarization is expected to contain most of the radio signal power, due to the mainly geomagnetic origin of the radio emission, the pulse detection is performed for both antenna polarizations. Both polarizations are processed separately, as the two high frequency signals can not be merged, particularly not by simply adding them, as they represent two components of the vectorial electric field. The Envelope signals could be added quadratically, which would correspond to the total power of the electric field. However, the pulse shape would change, and the signal to noise ratio of predominantly polarized pulses would be reduced. Therefore, both channels are only combined after pulse detection with an or-gate as a last step.

The next step is, to combine the detected pulses from different antennas, and accept only events where the pulses occur in three antennas within the coincidence time t_c . The random coincidence rate of n antennas, triggering independently with a rate of f_a each is negligible for a reasonably chosen t_c :

$$f_{random} = f_a^n \cdot t_c^{n-1} \quad t_c=220 \text{ ns}, f_a=600 \text{ Hz} \quad 10 \mu\text{Hz} \quad (3.1)$$

However, the pulses do not occur independently, but the same pulse from a distinct noise source is detected in one antenna after the other. These horizontally incoming pulses are rejected, when the coincidence time is further decreased, though at the expense of a reduced angle of aperture for cosmic ray detection. The maximal accepted zenith angle θ in a triangle setup for a plane wave front is

$$\theta = \arcsin\left(\frac{c \cdot t_c}{a}\right) \quad (3.2)$$

Where a is the corresponding triangle dimension, which varies azimuth dependently between the side length and the triangle height as shown in fig. 3.8. This zenith discrimination however works only for far away noise sources, where the plane wave front approximation is valid, a curved wave front can strongly reduce the time difference between the antennas. At the KIT site the distance to the next buildings is always lower




Figure 3.8.: *Left:* Sketch of the relevant triangle dimension a depending on azimuth ϕ . *Right:* Accepted zenith angle θ (red area) versus azimuth angle for $t_c = 162.5$ ns.

than 100 m, which makes the horizon suppression less efficient. With a coincidence time $t_c = 162.5$ ns and a triangle side length of 65 m the projected coincidence trigger rate is 50 Hz.

The full data acquisition system and analysis pipeline of LOPESTAR^{STAR} was developed and analyzed in the PhD thesis of T. Asch [52] and confirmed in the diploma thesis of M. Helfrich [68], which was supervised during this thesis. The detection efficiency of cosmic rays above $E = 5 \cdot 10^{17}$ eV was found to be around 16 %. However, most events are just rejected due to high noise conditions. When only silent periods of measurement are taken into account, the detection efficiency rises to 90 %. The percentage of real cosmic ray events in the accepted data, also called purity, is about 6 % including the pre-selection by the external trigger. When projecting self-triggered, live data taking, the purity drops however to a poor value of $1.4 \cdot 10^{-6}$.

The main reason for the moderate trigger performance is obviously the special noise situation at LOPESTAR^{STAR}. Although the offline data analysis gives a first idea about the common features of this background, the full amount of data recorded during two years corresponds to only 60 s of continuous data taking. Rare background events and short-term features are therefore not covered by the recorded data. To further improve the self-trigger algorithm and finally validate it, the trigger must be implemented with a prototype system and confronted with the real challenges of live measurement in this background conditions.

3.2. The prototype for digital self-triggering

The data amount which has to be handled for real-time digital triggering amounts to 1 Gbit/s per channel when digitizing with 80 MHz and 12 bit resolution. Processing one Fourier transform with a window length of $N = 1024$ on this data stream, as needed for the RFI-suppression, already uses the entire computing resources of a 3 GHz Intel Core2TM processor core, even when using the highly optimized FFTW3 algorithm [69]. The realization of the full trigger chain on regular personal computer hardware is therefore impractical, at least for a larger number of antennas. Instead the trigger engine was implemented on Field Programmable Gate Array (FPGA) chips, as described below.

3.2.1. Properties of a Field Programmable Gate Array

A Field Programmable Gate Array is an integrated circuit which can be programmed to incorporate any digital logic from different fields of applications, up to complex mathematical calculations or several full processor cores. Its main component is a grid of ten thousands of logic elements (LE), each with several inputs and one output. The value of the LE output is defined by its inputs according to a programmable truth table. The result can also be stored in a LE-internal flip-flop. The inputs and outputs of different LEs can be connected with each other via programmable switches and can be patched to external pins of the chip. The external pins can be programmed as inputs or outputs, and are compatible to different digital signaling voltages. As a FPGA does not execute a sequential program, it is not configured with a programming language, but with a hardware description language like for example VHDL⁸. The resulting firmware must be loaded to the FPGA during each power-up, for example from a Flash memory device, because FPGAs do not contain nonvolatile memory.

All high speed logics with a certain complexity use a clock signal, to allow for a storage of intermediate results in a pipelined design. As an exact clock synchronicity in all nodes is crucial for high speed, all FPGAs contain a complex clock tree to fan out the clock from one pin to all LEs. The power consumption of the clock tree alone can account for more than half of the total FPGA power budget. The clock tree of a modern FPGA also contains a number of PLLs⁹ to increase the device flexibility, by enabling the use of different clock frequencies for different partitions of the device, a feature also widely used in this work.

For the internal buffering of data, a FPGA also includes some dedicated memory to relieve the LE flip-flops from this task. It is organized in several independent memory banks, each with an own memory controller with variable bit-width to increase the deployment flexibility. As the multiplication is a relatively costly operation, which is very common especially in digital signal processing applications, an FPGA also includes several dedicated multipliers to save LEs and increase the performance.

Although the clock frequency of a FPGA is limited to some 100 MHz and modern CPUs¹⁰ run at some GHz, the computing power of a FPGA can easily exceed the performance of any CPU for a suitable algorithm, just because of its massive parallelism: Within one clock cycle hundreds of multiplications, thousands of additions, and reams of boolean decisions can be executed. For example a real-time FFT, which uses a CPU to full capacity as mentioned above, can be implemented as a fully parallel engine within 5000 LEs, using less than 10 % of the resources of a mid-range FPGA. Because of the lower production quantity, FPGAs usually drag one generation behind the current CPU production technology. Despite its larger semiconductor feature size, an FPGA still has large advantages over a CPU in terms of power efficiency. The FPGA resources are employed bit by bit, each calculation is performed exactly with the desired bit-width, if only to reduce the usage of LEs. As a result, the FPGA's logic circuit is made-to-measure, no unneeded components are driven during operation.

For the self-trigger implementation, FPGAs from the Altera Corporation are used. The firmware is developed with the associated Quartus II design suite. During FPGA programming several constraints must be fulfilled and tested:

⁸ Very high speed integrated circuit Hardware Description Language

⁹ Phase Locked Loop - Device to create a clock signal with a defined frequency and phase relative to a reference clock.

¹⁰ Central Processing Unit

- First, the **functional** simulation proves that the VHDL code behaves as intended. This is absolutely crucial, because VHDL is a low-level programming language, Hardware-specific behavior is used frequently. VHDL is also rather error-prone because unlike usual programming languages, commands are not processed sequentially but simultaneously, as they basically describe connections in a logic circuit. Debugging the firmware on the target hardware is no option, as iteration cycles would be too time-consuming due to the elaborate compilation process. The simulation is performed with Modelsim¹¹, using a dedicated test bench, which provides the design under test with clock signals and input stimuli.
- Second, the correct **timing** of the signals in the FPGA must be simulated, to ensure that its input signal is valid for each node before reading it (sample time) and stays constant long enough (hold time). This is particularly important between different clock domains and for attached peripheral devices, where the required sample and hold times are taken out of their data sheets. The list of timing constraints¹² is fed to the Altera TimeQuest timing analyzer, which performs an automated timing simulation for all nodes. The simulation is based on a supply voltage and temperature dependent FPGA model, and uses a worst case estimation. The speed of each logic operation depends on its complexity, on how many input nodes are compared with each other for the result. Sometimes the timing of a path must be relieved by processing the logic operation step-wise in a pipeline.
- The available **FPGA resources** of logic elements limit the design complexity. To reduce the number of occupied LEs, the source code is first optimized during analysis and synthesis. For example logic paths which are never read are completely removed. The resulting netlist is then fitted to the FPGA architecture, also considering the applied speed constraints, as the detailed placement and signal path routing defines the timing of the circuit. The fitting process is a NP-complex computing problem¹³, the used algorithm starts with various pseudo random placements, and tries optimizations. In particular for complex projects the fitting is very time consuming and may need several iterations. The fitting effort rises dramatically with increased FPGA usage and required speed. A high-speed design may only allow for a low occupancy of less than 70 %, otherwise the fit may fail completely. Finally, the fit result is translated into a programming file, containing the settings of all programmable switches and look-up tables of the FPGA.
- The **power consumption** of an FPGA is, like any other CMOS¹⁴ circuit, dominated by switching losses, and therefore dependent on the number of gates, and their switching frequency. Because the power consumption is important for heat dissipation and supply dimensioning, it can be simulated with the Altera PowerPlay power analyzer. Early power simulation is only a rough estimation using the FPGA occupancy and an average switching rate. Later in the design process

¹¹ Mentor Graphics Modelsim is a simulation suite for analog and digital systems reading VHDL code, any system node can be displayed time dependently like on an oscilloscope screen. For this thesis the Altera Edition 6.4a was used.

¹² The timing constraints list contains about 1000 programming lines.

¹³ In computational theory, a problem is NP-complex (Nondeterministic Polynomial), if there is no algorithm to solve it efficiently (in polynomial time), although it can be verified in polynomial time.

¹⁴ Complementary Metal Oxide Semiconductors is a semiconductor technology. By using complementary transistors for logic gates the quiescent current is negligible.

the power simulation becomes very accurate, when the fitted design and realistic input signals are used.

3.2.2. The prototype hardware

To test the feasibility of the trigger approach in FPGA logic, a target hardware was necessary. Because the development of a custom printed circuit board (PCB) would be too costly during this early project stage, the electronics from a pipeline inspection robot, which was also developed at IPE, was reused. The inspection robot, also called 'pipeline pig', analyzes digitized data from ultrasound sensors with an FPGA to search for metal corrosion and cracks. Although originating from an entirely different application field, the system properties are also well suited for radio detection.

The electronics as shown in the block diagram fig. 3.9 supports four channels, each starts with an analog amplifier with programmable gain. This simplifies the matching to different voltage ranges, when exchanging the signal source from ultrasound sensors to amplified radio antennas. Two dual ADCs digitize the signals with 80 MHz¹⁵ and 10 bit resolution, using sub-sampling in the second Nyquist domain like LOPES^{STAR}.




Figure 3.9.: *Left:* Block diagram of the prototype trigger hardware: Yellow blocks show the two analog channels, interfaces to the exterior are colored blue, the processor devices are red, and auxiliary units green. *Right:* Corresponding picture of the printed circuit board (PCB).

The digitized signals are fed to an Altera Stratix I FPGA containing 40.000 logic elements to implement the actual trigger algorithm. The corresponding firmware is booted from a persistent Flash memory device on power up. During malfunctions a JTAG¹⁶ interface provides access for debugging. The FPGA is connected to a processor module on the PCB backside, which contains a Blackfin 561¹⁷ processor, working memory, a flash drive, and an Ethernet interface. It runs a μ Clinux¹⁸ operating system, providing a local control terminal which is simply accessed via network. The operating system interfaces the FPGA as a memory mapped device, thus selected FPGA registers are addressed through a specified main memory range. Despite its potential high performance, the processor module is only used to control and monitor the trigger setup,

¹⁵ The original ADC, sampling with 40 MHz, had to be exchanged with a pin-compatible type (Analog Devices AD9216), supporting a higher sampling rate.

¹⁶ Joint Test Action Group is the common name for a standardized serial port used to debug electronics.

¹⁷ The BlackfinTM 561 from Analog Devices is a powerful digital signal processor with two 32 bit cores running at 600 MHz.

¹⁸ μ Clinux is a fork of the Linux kernel for microcontrollers without memory management unit.

like for example threshold level or trigger rate. For a fast low-level communication, the FPGA also has access to 10 TTL¹⁹ and 5 RS485²⁰ level converters.

3.2.3. The prototype data acquisition system

The trigger realization sticks closely to the proposed algorithm from the LOPES^{STAR} investigations discussed in section 3.1.4. Three antennas in triangle configuration are each equipped with one FPGA board. Both antenna polarizations are digitized, the remaining two ADC channels are left unused. Inside the FPGA, the signals are subject to an elaborate RFI-suppression chain, which is described in chapter 6.

To reconstruct the original amplitude, the sub-sampled signal must be upsampled (s. sec 3.1.4). Due to performance limits, the signal is only upsampled by two, instead of a usual upsampling factor of eight during offline analysis. As the signal is already available in frequency domain from the RFI-suppression, the method of upsampling in frequency domain suggests itself. This would double the data amount to be processed through the inverse FFT, and would thus require twice the clock frequency. Because the FPGA performance does not suffice for this approach, the upsampling must be performed with the second method by adding zeroes in time domain and subsequent filtering. The filter is a FIR²¹ high-pass of 65th order with rectangular cutoff, providing a trap attenuation of -40 dB. Although the data rate is also doubled to 160 MHz, the employed operations are less complex than an FFT and thus within reach.

After upsampling, the signal is rectified and low-pass filtered to generate the envelope for the final pulse detection. Because the envelope signal does not contain high-frequency components anymore, it is down-sampled again to 80 MHz to reduce the amount of logic running at double frequency and thereby relieve the timing demands. To prevent the loss of a high amplitude value from a pulse, the maximum of two consecutive values is always kept.

A pulse is detected when the envelope signal exceeds a threshold value. The threshold is set dynamically to a multiple of the signal root mean square calculated over the last 0.2 s, which directly implies a cut on the signal to noise ratio of the pulse. The multiplication factor is variable, but was kept at a value of eight during measurements. This is comparable with the LOPES^{STAR} quality cut, requiring a signal to noise ratio of 75 in power or $\sqrt{75}$ in amplitude. After the threshold crossing, the maximum value of the pulse is located, and the time of the maximum is used as pulse time. Additionally, the pulse width at half the pulse height, the integral sum of the pulse, and the integral of the after-pulse region are calculated. Based on these pulse parameters, the offline analysis proposes further cuts to reduce the false trigger rate. However, most of the time the pulse shape cuts were not applied, as their significance was uncertain for the signals recorded under different conditions, and upsampled to a different rate. During the early prototype stage a higher false trigger rate seemed to be rather acceptable, than the danger of missing one real event.

The trigger signals from both separately processed polarization channels are finally combined with an or-gate to a joint antenna trigger, also called 'T2-trigger' as it is the second trigger stage after threshold crossing. The full FPGA design occupies 86 % of the available resources, leaving only small margins for any further extension.

¹⁹ Transistor-Transistor Logic is a standard for digital signaling using a single ended voltage of 0 V for low level and 5 V for high level.

²⁰ RS485 is a differential signaling standard using a voltage difference of ± 0.2 V.

²¹ Finite Impulse Response Filters are discussed in appendix A.2.1




Figure 3.10.: Block diagram of the prototype self-trigger setup. The LOPES^{STAR} setup is extended by three self-trigger units. The signal from the fourth antenna is not used for triggering, but still recorded.

The coincidence trigger of the three antennas is set up in FPGA logic, using a simple master/slave hierarchy: Two FPGAs act as slaves, sending out each of their trigger signals as a RS485 strobe. The rising edge of the strobe defines the trigger time. The master FPGA reads both RS485 edges, corrects for I/O latency, and causes a coincidence trigger if both slaves and the master have a T2-trigger within a configurable coincidence time. The resulting trigger, called 'T3-trigger', is fed to a DAQ system as used at LOPES^{STAR}²², starting the readout of the ADC-board, the generation of a GPS-timestamp, and their storage to the data base, like illustrated in figure 3.10. By extending the existing system, the time-consuming and costly development of a new DAQ can be avoided, the corresponding analysis tools can be reused. With the simple strobe-based communication system, the three trigger FPGAs can run independent from each other with free floating clocks, a complicated synchronization is unnecessary. To keep the housing, power supply, and communication between the FPGA boards simple, the trigger units are placed at a central spot, the antennas are connected with long analog cables. A disadvantage of the modular design is, that the radio signal must be digitized twice, first at the trigger boards and second at the DAQ-system. To avoid signal reflections due to impedance mismatching, the analog cables are connected in a daisy chain, from the filter amplifier, over the trigger to the DAQ, the terminating resistors on the FPGA board in the middle of the chain are removed.

The master trigger board accepts the additional input of an external trigger, which is merged with the coincidence trigger using an or-gate. Additionally, traces marked as background are read out every 20 s to have an unbiased insight into the noise situation. The output of the trigger is equipped with a dead-time of 30 ms to avoid exceeding the maximum readout rate of DAQ system. Because of the latency of the RFI-suppression of $21.9 \mu\text{s}$ and the trigger logic of $4.6 \mu\text{s}$, the radio pulse is expected $26.5 \mu\text{s}$ before the trigger signal. An adequate pre-trigger time set at the ADC-board ensures, that the pulse is recorded at the expected time

The functionality of the setup was tested in laboratory, trigger latencies were measured, system settings were adjusted. Then the trigger was operated for 3 weeks at LOPES^{STAR} antenna field D17 to verify reasonable trigger thresholds and rates, and prove its long term stability. In October 2008 the entire self-trigger setup was trans-

²² The trigger signal is level-converted from TTL to NIM (Nuclear Instrumentation Module) standard, to comply with the ADC system.

ported to the Pierre Auger Observatory and installed at the Balloon Launching station. The ensuing measurements and their results will be described in the following chapter. Although the trigger and data acquisition layout is well qualified as a fast approach for a simplified prototype setup, it is obviously unsuitable for a larger experiment, because it does not scale for a large number of antennas. The development of a system consequently adjusted to the requirements of a large scale radio array will be described in chapter 5.

4. Measurements with the prototype self-trigger

The analysis of LOPES^{STAR} data reveals that the background noise from the industrial surroundings disqualifies the KIT as a location for long term self-trigger measurements. The expected coincidence rate of 50 Hz is hard to handle and is only expected to lead to an improved knowledge about KIT-specific noise details. Obviously, a large-scale radio array must be built at a less radio-polluted, rural site. Additionally, the radio detector must be collocated with an established surface detector (SD), to verify cosmic rays measurements. For the relevant energy range above $10^{17.5}\text{eV}$, the Pierre Auger Observatory, being the leading cosmic ray experiment, is suited best to host the efforts in radio detection.

4.1. Setup at the BLS

To find an appropriate spot for the prototype radio detector on the large area of the Pierre Auger Observatory, a number of criteria are essential. On the one hand the location should be remote enough from the local village of Malargüe and other potential radio noise sources, on the other hand accessibility and electric power are essential to run the measurements.

From the beginning, the Balloon Launching Station (BLS), which is used to start weather balloons for the atmosphere dependent calibration of the Pierre Auger Observatory [70], makes a good compromise between remoteness and infrastructure. The BLS is located in the west of the observatory (s. fig. 2.4), and can be reached by car from Malargüe within 20 minutes. It consists of an insulated 20 ft container, which provides weather protection, connection to the power grid, and access to the computer network of the central campus. As the local land owner had no objections, three antennas in an equilateral triangle configuration with a side length of 100 m were installed in 2006 (s. fig. 4.1). The setup was similar to LOPES^{STAR}, using aluminum LPDAs, a central housing for DAQ, and 160 m long analog cables to connect the antennas. Instead of burying the antenna cables, they were mounted as overhead lines on wooden poles to save installation costs. After construction of the antenna array and a few weeks of unsuccessful operation with the analog self-trigger (s. sec. 3.1.3), a collaborating group from the Netherlands used the setup for radio detection with an scintillator-based external trigger [71].

In October 2008 we started a new measuring campaign at the BLS with the prototype self-trigger implemented on the pipeline-pig hardware and the LOPES^{STAR} based DAQ system as described in the previous section 3.2.3. The installation of the prototype trigger, the maintenance of the measurement, and the data analysis were major parts of this thesis.

The campaign followed three main objectives:

- The proposed **new antenna**, the SALLA with a diameter of 75 cm (s. sec. 3.1.1), was first used in field. It replaced the aluminum LDPAs, which suffered heavy




Figure 4.1.: Map of the prototype self-trigger setup at the BLS (*white container*), underlying satellite image by Google earthTM. The initial triangle of SALLAs mounted on Pole 1, 2, and 3 (*green symbols*) was later extended by the easternmost SALLA (*red*) to a parallelogram. The accessing gravel road and the power line cross in the northwest, the infill tank Olaia (*yellow*) is located to the east.

damage in the previous two years of operation and lost some dipole beams due to high wind speed. The SALLA was supposed to prove its superior durability under the same conditions. At the same time, it had to be shown that the lower sensitivity of the SALLA is sufficient to detect cosmic rays. As the antenna gain is only a problem, if the internal preamplifier noise becomes dominant over external noise sources, the detectability of the inevitable galactic noise was chosen as an antenna sensitivity benchmark.

- The most important aim was of course the test of the **self-trigger** approach and the detection of cosmic rays. Particularly interesting were the effectiveness of the RFI-suppression, the achieved thresholds, and the resulting background trigger rates. This improved the understanding of the specific noise situation at the Pierre Auger Observatory.
- Additionally, the technical feasibility of **remote measuring** over several months without field maintenance was an important issue, particularly for a later array with a large number of stations. Besides trimming the software for stability, this led to the development of fail-safe methods for remote update and reset.

In July 2009, the antenna signals dropped out. It turned out that the preamplifiers and the filter amplifier were damaged, probably by a lightning stroke, while the subsequent hardware in the signal chain was saved by the lightning protection circuits. Fortunately a second journey to Argentina was already booked at that time, so the damage could be repaired one month later. On that occasion the setup was upgraded in two respects to improve the detection threshold, which had been the original purpose of the second trip. First, the antennas were replaced by SALLAs with a diameter of 120 cm, providing a better sensitivity. Second, an additional antenna was installed at Pole 6¹, creating another equilateral triangle. This was necessary because the antenna at Pole 1 turned out to be strongly contaminated by a close-by noise source, while both of the more distant antennas were hardly affected (s. fig 4.5). By replacing the noisy antenna for the trigger generation with Pole 6, the general threshold of the whole array was lowered. When a trigger occurred, the noisy antenna was still read out, and the signal of all 4

¹ The numeration of antenna poles follows a prior plan of a larger radio array at the BLS.

antennas was stored. The new antenna had to be connected with a 250 m coaxial cable, which increased the signal loss and propagation time in the cable, compared to the other antennas. In later analysis this could be corrected easily, the self-trigger setup used a different setting of the programmable gain amplifier on the PCB to adjust the amplitude, and FPGA internal delay lines for the timing.

Measurements with the upgraded system continued until February 2010, when the power supply of the DAQ crate was destroyed. Instead of repairing the damage, the prototype setup was dismantled in favor of the system development for AERA. During the whole campaign, frequent power outages were an issue, each time requiring a manual restart of the system via remote connection. The blackouts caused the major part of the total system downtime of 36 %, besides short scheduled interruptions for updates and special background measurements.

4.2. Background conditions at the BLS

Continuous background

During the first days after installation, the radio background on site was characterized. The dynamic spectrum (s. fig. 4.2) shows that the background noise power at the BLS is lower than at the KIT site (s. fig. 3.7). In the frequency regions around 50 MHz and 60 MHz the noise level is low enough to be compatible with the galactic noise level [72].

Horizontal lines in the dynamic spectrum indicate mono-frequent interferences, at the BLS the power contribution of these carriers is only minor. Therefore the sophisticated RFI-suppression is hardly necessary [73], at least it is not harmful for the measurement. Additional interference is introduced by intermodulations, visible as wider bands of elevated noise at 55 MHz, 67 MHz and with some distinct lines around 77 MHz. The intermodulations are caused by unwanted frequency mixing due to nonlinearities of the LNA, which is placed in front of the band filter and encounters the high signal levels from the FM band. With the SALLA, the system is more sensitive towards intermodulations, as its antenna gain increases towards higher frequencies. The LPDA in contrast acts like an intrinsic first order band-pass and avoids intermodulations with the FM band. For the succeeding radio detector, an LNA design towards low intermodulation is crucial, which can be achieved with a low order band filter in front of the first active element.

The dynamic spectrum also shows a general variation of the noise level with a day – night periodicity which is caused by human activity. Vertical lines indicate a time span with a strongly elevated noise level, during these periods, which account for about 3 % of the measuring time, a reasonable measurement is impossible. The reason for the violently increased background remains unclear. Possible sources are thunderstorms, cars driving by, or airplanes flying over [74].

Transient background

Although the recorded frequency spectrum gives a precise description of the continuous background level, which allows for an estimation of the rate of random pulses as shown in [62], this model does not include sporadic transients, for example from electrostatic discharges. As these transient pulses are short in time and, with a rate in the order of kHz, relatively rare, their average power contribution is negligible. Thus, they are invisible in the frequency spectrum. On the other hand, these transients are common enough to dramatically increase the occurrence of high background amplitudes, adding heavy tails to the expected Gaussian amplitude distribution. Accordingly the




Figure 4.2.: Dynamic power density spectrum of the east-west polarized electric field recorded with the 75 cm SALLA at Pole 2.

background amplitudes, as shown in figure 4.3, only follow a normal distribution for small amplitudes. The contribution from transients can be described with an additional Cauchy distribution. For large statistics, the best modelling of the radio background can be done with the more general Lévy alpha-stable distribution [75], [76].

As a result, the T2-trigger rate does not drop for rising threshold values as fast as predicted in [62], but almost stagnates around 1 kHz, as shown in figure 4.4. For this plot the threshold was systematically varied over a full day, while monitoring the trigger rate. The errorbars display the variance of the rate during the day. The trigger rate is determined with a 14 bit counter and a 1 second period. Accordingly, the value saturates at 16,383 Hz. As expected from the measured frequency spectrum, the RFI-suppression has only negligible effect on the threshold dependency. The antenna position however has a large impact. While Pole 2 and 3 show a comparable behavior, the noise level at Pole 1 is about twice as high.

Threshold level

In principle the prototype trigger setup, which uses low level strobe signals for communication as described in section 3.2.3, could handle trigger rates up to 1 MHz. At the BLS however, the rate was kept between 1 kHz and 2 kHz during long-term operation, because the lower rate is more realistic for the later trigger setup, using high-level communication.

The threshold necessary to achieve this rate must be translated from a voltage threshold to a value of measured field strength, in order to be comparable with a variation of the setup or other experiments. The translation of the threshold level is equivalent to the field strength calibration of the peak amplitude of short pulses, as the threshold is only crossed by these pulses. As the analog system was already used in the LOPESTAR experiment, the known laboratory calibration of the analog chain components [62] can be used. The response of the SALLA is taken from an antenna simulation [77], which was proven to be reliable by a cross calibration in [52]. Because the gain of the analog chain and antenna depend on direction and frequency, the threshold reconstruction requires also assumptions on the measured pulse signals. For a realistic treatment of




Figure 4.3.: Distribution of radio amplitudes at ADC level. The underlying background traces were taken between September 2008 and June 2009 at Pole 2, east-west polarization, and account for 45 s of data.




Figure 4.4.: T2-trigger rate versus threshold value for the east-west polarization of Pole 1 and Pole 2 with and without RFI-suppression

the frequency spectrum of cosmic ray induced pulses, a set of exemplary events from LOPES^{STAR} is used to get an average conversion factor. The direction dependence can be covered with a constant factor as well, at least for moderately inclined cosmic rays, as the directivity pattern of the SALLA (s. fig. 3.3) is almost constant for zenith angles below 60° . With a claimed system calibration uncertainty of 0.8 dB or 10 % [52], and a variance of 8 % within the set of example pulses, the systematic uncertainty for the threshold calibration is estimated to be 13 %. For the setup at the BLS with the 75 cm SALLA, a pulse with a recorded amplitude of 1 mV translates to a local field strength peak of $(13.2 \pm 1.7) \mu\text{V/m}$. With the 120 cm SALLA, the conversion factor decreases to $(6.8 \pm 0.8) \cdot 10^{-3} \text{ m}^{-1}$.

The threshold towards noise transients is different, as the noise mainly arrives from the horizon and may have a different spectrum. Due to the horizontal suppression of the antenna, the threshold for noise is significantly higher. As the quantitative values of the background transient amplitudes are of minor interest, the use of constant calibration factors for the threshold is still a valid approach.

The resulting threshold during the full period of measurements is shown in figure 4.5 as average values for each day, with the error bar indicating the threshold variance during the day. It is clearly visible that the replacement of the noisy Pole 1 in July 2009 with the more distant Pole 6 leads to a large improvement of sensitivity, as the trigger threshold is cut in half. At the same time the exchange of the 75 cm SALLA with the larger 120 cm SALLA has no effect on the electric field threshold, as shown at Pole 2, while of course the voltage threshold is increased due to the higher antenna gain. During the periods of installation however, when there is human activity on site, the noise level seems to be elevated.

Although the average background noise power is lower at the BLS than at the KIT, the heavy tail in the amplitude distribution, which is caused by prevalent transients, results in a comparably high threshold. A further increase of the threshold is not very effective to reduce the trigger rate.




Figure 4.5.: Daily average of the trigger threshold during the full measuring time of the east-west polarization of selected antenna positions at the BLS.

Transient direction

A large drop of the rate by several orders of magnitude is however achieved by requiring coincident T2-triggers from all three antennas. To examine the effect of the coincidence time window on the resulting T3 rate, the coincidence time was varied over one day, while counting the rate over 10 s each. The averaged rate values versus their corresponding time window are shown in figure 4.6, the error bars illustrate the statistic variance of the rate. To avoid exceeding the readout capabilities of the DAQ-system, the trigger rate is limited by a 30 ms dead-time. For low trigger rates, the dead-time has only negligible effect, but at higher rates the underlying rate must be reconstructed from the recorded value.

In contrast to the assumption that the three antennas trigger independently from each other, which leads to a negligible² T3 rate of $f_{T3} = f_{T2}^3 \cdot t_{\text{coincidence}}^2$, a large part of the background pulses is registered together in all antennas. These pulses can be rejected, because they arrive horizontally, and thus trigger the antennas one after the other. Due to the triangular geometry, the transit time of a background pulse depends on the azimuth direction according to figure 3.8. The shortest signal transit time is achieved parallel to the triangle heights. This direction defines the maximal allowed zenith angle for a certain coincidence time and is mentioned as a second x-axis in figure 4.6. While the T3-rate saturates around 30 Hz for large coincidence times, as expected when still allowing horizontal pulses, the rate does not drop significantly directly at the 90° rejection threshold. This is due to sporadic misassignment of the pulse time by the T2-trigger, which results in a fraction of events with a reduced time difference. During the first months of operation, a coincidence time of 262.5 ns was used, later the coincidence time was set to 225 ns, to reduce the necessary harddisk space.

During the initial test operation of the self-trigger setup before shipment to Argentina, the same analysis was performed at the LOPES^{STAR} site. As the antenna triangle at the KIT is smaller, the same zenith angle cut is achieved with a smaller coincidence time.

² With a T2-rate of 1 kHz and a coincidence time of 250 ns, the random coincidence rate is 62.5 μ Hz.




Figure 4.6.: Comparison of the coincidence trigger rate at the BLS and at LOPES^{STAR} versus the corresponding coincidence time and the allowed zenith angle. The red squares illustrate the actual trigger rate, when correcting for dead-time.




Figure 4.7.: Sky-map of triggered events as an azimuthal equidistant projection. The polar angle illustrates azimuth, the radius shows the elevation with zenith in the center and horizon at the rim of the coordinate system.

Therefore the measured values (blue triangles in fig. 4.6), are scaled in the horizontal direction, so that they correspond to the upper x-axis, to be comparable with the BLS points. It is interesting that, although the T2 conditions due to background transients at the KIT are similar to the BLS, the coincidence suppression works much better in the latter place. On the one hand, the background sources at the KIT may be closer to the antenna field, leading to a curved signal front, which can result in a reduced time difference. On the other hand, the high building density may permit multiple signal reflection, so the initial direction is not even recognizable anymore.

For the T3-triggered events, the incoming direction of the recorded radio pulse can be analyzed off-line. A directional sky-map of all reconstructible event directions recorded during the second phase of the campaign with the 120 cm SALLA is shown in figure 4.7. This half year of measurement is suitable for the diagram, as it took place under steady conditions with constant settings. Although the trigger, set to a coincidence time of 225 ns, should only accept events with a valid signal plane and a zenith angle smaller than 50°, only 1.49 million out of 2.41 million self-triggered events are reconstructible.

Of the remaining events, 82 % are not reconstructible, because a significant radio signal from at least one antenna is missing. In most cases the absolute pulse height is enough to trigger, but the signal after the pulse is very noisy and contains several other pulses. Because the FPGA trigger algorithm does not wait for subsequent pulses, these events are recorded although they are obviously caused by unwanted noise. This type of background can be rejected relatively easy in a new trigger system. For the other 18 %, the time difference between the pulses is too large for a successful plane fit. This discrepancy between the real-time self-trigger decision and the subsequent reconstruction can be explained by the independent signal digitizing with ADCs of different resolution for the trigger and the DAQ. In addition, the off-line reconstruction is more elaborate, as the signal is upsampled by a factor of eight instead of two, the envelope is calculated using the Hilbert-transform instead of a low-pass filter, and the pulse finder itself is more advanced than the limited implementation in FPGA firmware.

This also explains why two thirds of the events have a low elevation with a zenith angle above 50° , contrary to the trigger acceptance setting. As the rate of horizontally incoming events is about two orders of magnitude higher than the target T3-rate (s. fig. 4.6), already a small fraction of events with a falsely determined pulse time will lead to many false triggers of horizontal events. This means for future developments, that an improvement of the pulse time determination in terms of robustness and precision is crucial for a reliable coincidence suppression.

The sky-map also shows three distinct hotspots where most of the background events come from. Although the BLS container itself was suspected as the main noise source in the beginning, as it is filled with different electronics, none of the hotspots points towards it. Instead, the two noisy regions close to horizon seem to point parallel to the power line. The averted region towards south-east seems to be rather silent. The third hotspot is located close to zenith, which can not be assigned to a regular noise source. The most probable explanation is the simultaneous injection of noise pulses at the central DAQ system itself, rather than at the antennas.

Periodic transients

While the origin of the background transients can hardly be narrowed down with the geometric reconstruction, the sequential arrangement of the events delivers some insights. For a random process an exponential distribution of the time interval between two self-triggers is expected. In contrast, the values as shown in figure 4.8 primarily group around multiples of 10 ms with an emphasis on every second value. A time difference below 30 ms does not occur, due to the according dead-time.

In frequency domain the time distribution translates to a 50 Hz periodicity with a secondary 100 Hz harmonic, which is a strong indication towards the power grid as the dominant background source. Although the analog band filter blocks the line frequency itself, unwanted electric sparks due to electrostatic discharge cause radio transients. The sparks occur mainly at the voltage maxima of the alternating current with an increased




Figure 4.8.: Distribution of the time difference between two self-triggered events at the BLS. The Plot contains all 135 million events measured during the campaign at the BLS.

breakdown probability in one voltage direction, which explains the proportion of the 50/100 Hz components.

The BLS draws its power from the grid with a private transformer connected to a medium voltage line, which runs in a distance of 200 m from the antenna field along the road to the nearby village La Junta. The power line is realized as a so-called single-wire earth return (SWER) line, which is a common method for low-cost electrification of rural areas. It uses, as the name suggests, only a single wire to transmit electrical power at a voltage of some 10 kV, while the ground serves as return path. As the return current can spread over a huge cross section, the resistance of the ground is small even over a large distance. A proper coupling to the ground at the feed points with massive earth stakes is however crucial. Close to the earth stakes the voltage gradient rises, particularly in dry ground conditions like in the Argentinian pampa, and may exceed the electrical breakdown voltage. Additionally, the return current may leak into the shielding of the antenna cables, which are grounded both at the antenna and at the BLS container for lightning protection. This stray current then causes interferences directly at the signal line.

Unfortunately, a closer inspection of the frequency behavior of these periodic transients is not possible based on the T3 triggers, as the T3-rate is in the order of 1 Hz is too low for short term analysis. The T2 triggers, which have a rate of several 100 Hz, would be more suitable, with the prototype system it is however impossible to record the T2 timestamps. This kind of examination is performed with the subsequent AERA system and will be described in chapter 7.3.

Detection of galactic noise

Besides an examination of the self-trigger characteristics, the sensitivity of the newly proposed SALLA antenna was a key aspect of the measuring campaign. Although the antenna sensitivity does not limit the trigger performance at the BLS, where most of the false triggers are caused by the transient noise polluted environment instead of internal antenna noise, it might become important at a more radio silent location. In the used short wave radio band with an emphasis on the lower frequencies, the dominant natural background source is the galactic noise [65], which provides an inevitable sensitivity limit. The galactic noise amplitude is modulated by the variable fraction of the galaxy visible above horizon, or within the antenna's field of view, during a sidereal day. To distinguish the galactic noise modulation from man-made noise variation during the day – night cycle, the measurement must span over a period long enough to identify the reduced period of a sidereal day of 23 hours 56 minutes. The detectability of this modulation is a good benchmark for a radio system, because if it is visible, the system is as sensitive as possible at all. According to simulations the internal noise of the 120 cm SALLA is about as high as the galactic noise, so it should be close to the detection limit [77], which has to be verified by measurement.

As triggered events would be tainted by the included noise pulses, this background examination is performed on special traces, which are recorded randomly each 20 s. Direct analysis of the dynamic frequency spectra (s. fig. 4.2) shows no modulation parallel to the rising and setting of the galaxy, particularly not in the lower frequency range as expected, but only a variation of the man-made intermodulations at higher frequencies with an 24 h periodicity. The signal power integrated over selected frequency ranges or over the full frequency band does not contain a sidereal modulation either. A closer inspection reveals that over periods of some hours the background level is




Figure 4.9.: Comparison of the hour of the fitted daily radio noise maximum with the corresponding time of galactic culmination.

significantly higher than during silent periods. These noisy periods roughly account for half of the measuring time, and dominate the averaged background signal power.

For the detection of galactic noise the periods with a low noise level must be selected, and these selected periods must cover all 24 hours of a day. For this purpose, the full measuring time is divided into slices of 20 days each, within these 20 days the 5 days with the lowest average signal power are selected³. After plotting the average power of the 5 days versus the time of day, a sine function in the form $A + B \cdot \cos(2\pi \cdot (t - t_0)/24h)$ is fitted, using the fit parameters offset A , amplitude B , and phase t_0 . Assuming a variation with a fixed period of 24 hours is a valid approach both for galactic and man-made noise, as the period of a sidereal day is hardly distinguishable from a solar day during the short time span of only 20 days. For galactic noise it is expected that the fitted phase t_0 , which indicates the time of day with maximum noise amplitude, moves steadily over the year, parallel to the galactic culmination.

The plots of the noise power versus daytime are produced for each antenna channel, except of Pole 1, which is known to be notoriously noisy, and fitted with the sine function. Some of the noise power plots have no sinusoidal shape at all, as they are still heavily contaminated with man-made noise despite previous selection. These cases, when the fit fails and the reduced χ^2 of the fit exceeds 10, are ignored in the further analysis. They account for 60 out of 138 plots. Figure 4.9 shows the remaining fit values t_0 versus the measuring time, using the fitting uncertainty as error bar, and compares them with the known culmination of the galactic center⁴. It is clearly visible that the noise maxima are randomly distributed over the day for the 75 cm SALLA. The poor agreement with the known position of the galaxy is reflected in a χ^2 value over the number of degrees of freedom of only 545/44. The noise modulation measured with the 120 cm SALLA, however, is well aligned with the expected trend, the χ^2 /NDF is 25.3/32.

³ Selecting only a quarter of the time takes into account the uptime of 64 % due to power outages.

⁴ The position of the galactic center, which is co-located the zodiac sign Sagittarius, is taken from the open source software planetarium Stellarium.

This proves that the 120 cm SALLA is sensitive enough to detect the galactic noise, although careful data selection is necessary to identify the small signal within the interfering man-made noise at the BLS. The competing LPDA, however, which was used in the Maxima setup at the BLS [78], delivers a much clearer signal of the galactic noise modulation. It is directly visible by human eye in the frequency spectra [79]. The disadvantage of the SALLA is only partly due to its lower gain and thus increased internal antenna noise. A second factor is the increased exposure to man-made noise from horizontal direction because of its wider main lobe. Otherwise the selection of low noise periods would have no effect, as the internal antenna noise is constant and unavoidable. While the wide acceptance angle of the SALLA has advantages for the future detection of inclined cosmic rays, it also aggravates the background noise issues, we are facing during this early phase of self-trigger development.

4.3. Search for cosmic rays

To verify that a detected radio pulse is caused by an air shower, the reconstructed data from the surface detector is used. The SD-data is processed by the so-called Herald reconstruction, a simplified on-site analysis tool. Although it does not apply all possible corrections like the elaborate Offline reconstruction, it provides the central parameters of the air shower with minor deviations of a few percent [80]. Because an air shower must hit at least three water Cherenkov tanks to be reconstructible, the surface detector gains its full efficiency only at energies above 10^{18} eV. Events with lower energies however, may already cause a radio trigger if they are close-by. Therefore an additional infill tank was installed in the middle of neighboring detector hexagon, to lower the local energy threshold of SD. This infill tank, named Olaia, is located 300 m northeast from the BLS. Because only local air showers are of interest for the further analysis, only SD events which contain a signal of Olaia are used. This amounts to about a dozen events per day. The first step of merging radio events and SD data is to search for events with a similar timestamp from both detectors. If both reconstructed directions are in agreement as well and the radio amplitude distribution is plausible for the core position from SD, the radio event is proven to have a cosmic ray origin.

The time difference between radio and SD is caused to some extent by the propagation time of the shower front from the shower core to the radio antennas. This geometric effect can be corrected with the known position and propagation direction from Herald. A larger contribution to the time difference is given by the added latencies of the data acquisition systems. The Herald reconstruction already corrects for the latencies of SD and provides the estimated time when the shower core touches surface. The radio latency, which is dominated by the RFI-suppression engine, was measured in laboratory to be $26.5 \mu\text{s}$, and is also corrected. The remaining time difference is in the order of 100 ns and is mainly due to clock uncertainties and reconstruction uncertainties.

Externally triggered measurements

To ensure that DAQ system works as intended, and its timing is fully understood, we measured for one week with an external trigger system, which we could borrow from our Dutch collaborators. The external trigger system consists of two scintillator detectors, reused from the HiSPARC⁵ project. If both detectors have a coincident signal, the external trigger is provided as an TTL strobe. The TTL signal is fed to the master

⁵ High School Project on Astrophysics Research with Cosmic Rays [81].




Figure 4.10.: The difference between the timestamps of externally triggered radio events minus corresponding SD events (*X-axis*) versus the propagation time of a plane shower front from its core to the position of the radio setup (*Y-axis*). The black line describes the expected trend, taking into account the known trigger latency.

trigger board, delayed by $19.4 \mu\text{s}$ to be slightly earlier than a corresponding self-trigger, and finally orred with the expected self-trigger signal. Whether an event was caused by the self-trigger, externally, or both, is marked with an auxiliary user bit input to the ADC card. Due to the small distance between the scintillator plates, the energy threshold of the external trigger is low, resulting in an average trigger rate of 0.09 Hz .

During the one week of measurement with the external trigger, 17 events in coincidence with SD were recorded. The time difference towards SD accord well with the values estimated based on the known latencies and shower geometry (s. fig. 4.10). Most of these events show no visible radio signals, as it is expected because of the higher sensitivity of the external trigger system. Three out of the 17 events have significant pulses in all antennas, and 2 other events contain pulses which hardly stand out of the noise floor. Two of the events with strong pulses are also found by the self-trigger setup, which looks very promising at first glance.

However, when using the pulse times of the three antennas to reconstruct the primary arrival direction, none of the 5 events is compatible with a plane wave front. Instead they all have a similar timing, which, when taking into account the curved wave front, points towards the close-by BLS. The source of these pulses is apparently noise, emitted by the external trigger setup, probably due to the high voltage recharge of the photomultipliers, an effect which is very common at the LOPES experiment, and was also observed in earlier measurements at the BLS [71].

Estimation of the expected rate

After verifying that the assigned timestamps of the radio setup are compatible with SD, the borrowed external trigger was returned and data taking continued solely self-triggered for more than a year. During this time no further events together with SD were detected, neither with the original setup nor after the upgrade in August 2009.

The expected rate of events can be calculated on the basis of the known cosmic ray flux, the resulting radio pulses, and the applied threshold. In doing so, the main source of uncertainty is the radio emission model, in particular the lateral distribution of the radio signal strength. The common parameterization of the lateral distribution implies a linear dependence of the radio amplitude on the primary energy E_p and an exponential drop with the distance to the shower axis r . The most advanced examination of the lateral distribution at that time was performed at the LOPES experiment [82, 49], resulting in the following parameterization of the peak amplitude of the east-west component ϵ_{EW} with an additional sensitivity towards the azimuth angle θ and the geomagnetic angle α .

$$\epsilon_{EW} = \epsilon_0 \cdot \left(\frac{E_p}{10^{17} \text{eV}} \right)^{0.95} \cdot e^{-\frac{r}{R_0}} \cdot (1.16 - \cos \alpha) \cos \theta \quad (4.1)$$

The field strength at the shower center ϵ_0 has a fitted value of $(11.0 \pm 1.0) \frac{\mu\text{V}}{\text{m MHz}}$, and is scaled with the effective LOPES bandwidth of 33 MHz, to improve the comparability with other experiments, assuming that the coherent radio pulse height is increased linearly with the bandwidth. The scale parameter R_0 peaks around a value of 125 m, but a minority of events show a rather flat lateral distribution with a large scale parameter, shifting the average value of R_0 to $(236 \pm 81) \text{ m}$. The plausibility of this parameterization was proven by simulation [45], other experiments like CODALEMA [56] showed similar results. The parameterization from LOPES^{STAR} [52] uses a smaller scale parameter of $R_0 = (122 \pm 12) \text{ m}$, the values are however still compatible within their errorbars.

Figure 4.11 compares the pulse height expected for the setup at the BLS⁶ versus shower distance for various primary energies according to the parameterizations from LOPES [49] and LOPES^{STAR} [52]. The shaded bands, which are drawn for each primary energy, show the resulting pulse heights for different incoming directions. The highest signal is achieved with a zenith angle of $\theta = 0^\circ$ and a geomagnetic angle of $\alpha = 90^\circ$, the lower edge of each band describes a zenith of $\theta = 60^\circ$ and $\alpha = 45^\circ$. Both




Figure 4.11.: Comparison of the lateral distribution of the east-west radio component according to two different parameterizations for different primary Energies.

⁶ The BLS setup has an effective bandwidth of 42 MHz.




Figure 4.12.: Expected event rate with its error bars depending on the applied threshold for the antenna triangle at the BLS.

parameterizations give similar results close to the shower center, at larger distances, however, the field strength according to $\text{LOPES}^{\text{STAR}}$ drops much faster. This obviously has a significant impact on the obtainable detection range with a certain threshold. While the LOPES parameterization is based on larger statistics, and is widely approved and accepted within the radio detection community, it puts an emphasis on lower energies and short distances. The $\text{LOPES}^{\text{STAR}}$ parameterization, on the other hand, contains less events, but concentrates on a slightly higher primary energy. Because it was measured with the exact same electronics, the second parameterization is supposed to be more significant for the setup at the BLS.

With the parameterized radio pulse height, the number of self-triggered events for a certain threshold can be estimated with a simple Monte Carlo simulation, implementing the known cosmic ray flux above energies of 10^{17} eV . To detect an air shower, all three antennas in the 100 m triangle must contain a radio signal higher than the used threshold. The coincidence time of 225 ns keeps the detectable zenith angle below 50° , which defines the simulated solid angle of acceptance. Although the pulse description only covers the east-west polarization, it still contains the full detectability, because the major part of the radio pulse power is incorporated in this polarization anyhow, due to the mainly geomagnetic origin of the radio emission. Because of the low sampling rate of 80 MHz combined with the small upsampling factor of only two, the sampled value does typically miss the real pulse maximum in the trigger implementation, which reduces the detected pulse height systematically [62]. This is considered by reducing the measured height of the simulated pulses by 20 %. The dispersion of the analog filter has an additional reducing effect on the pulse height. Because the filter dispersion is not corrected by the prototype trigger, this accounts for another 20 % amplitude reduction.

The resulting simulation of the self-trigger rate is shown in figure 4.12, the shaded bands illustrate the error margins due to uncertainties in the parameterizations. For the LOPES parameterization, hundreds of self-triggerable cosmic rays are detected for any realistic threshold. With this optimistic expectation the prototype measurement was planned and operated. The model from $\text{LOPES}^{\text{STAR}}$, however, predicts only two events per year for the effective threshold of $320 \mu\text{V/m}$ during the first phase of mea-

surement. During the second period with a threshold of $150 \mu\text{V}/\text{m}$, about 12 events per year are predicted, which translates to 3 expected events during these 5 months, when considering the effective uptime of 64 %. The actual observation of no cosmic rays at all lies well within the error margins. Several recent results now also favor a smaller scaling factor: The analysis of MAXIMA data from the BLS, which was taken with an external trigger, indicates that $R_0 = (156 \pm 12) \text{ m}$ [83]. The newest LOPES analysis [84] abstains from the questionable events with a flat lateral distribution, which stem mainly from close-by inclined showers, and results in $R_0 = (157 \pm 54) \text{ m}$. Recent publications from CODALEMA now favor a smaller scalar factor of 160 m as well [85].

The lack of any self-triggered cosmic rays can be understood even better, when taking into account that both parameterizations significantly overestimate the absolute value of the expected field strength. The corresponding calibration factor ϵ_0 was namely measured at the KIT in Germany, where the magnitude of the Earth's magnetic field is twice as high as in Argentina. Because there is neither a properly calibrated measurement, nor a convincing simulation of the absolute value of the radio emission under these conditions to date, the effect can only be estimated. Additionally the models are biased towards higher amplitudes, because the underlying analysis includes only events with well visible radio pulses out of all measured cosmic rays. This holds particularly true for the LOPES parameterization, which originates from measurements at lower energies, close to the detection limit. The impact of this bias is hard to estimate, but it may at least to some extent explain the divergence of both models.

4.4. Survey for future site

The disadvantageous noise situation at the BLS required an increased detection threshold, which finally hindered the self-triggered detection of a cosmic ray event with the prototype setup. For the successor radio experiment the selection of a better site is important. A survey of the background noise was conducted after the first period of the prototype self-trigger campaign, during the update of the BLS setup in August 2009. The radio background is usually measured with a spectrum analyzer, which describes only continuous signal contributions, but ignores background transients, which are very relevant for the later necessary trigger threshold. To properly estimate the local transient pollution, the background trigger rate has to be measured with a similar trigger algorithm.

For this purpose, a mobile radio detection system was assembled with the field-tested 120 cm SALLA, LNA, and active band filter, followed by a Picoscope 3206 USB⁷ oscilloscope, in order to digitize the signal and read it into a laptop computer. The USB scope provides a simple and commercially available method of data taking with an adjustable input voltage range, a sampling rate of 200 MHz, and a resolution of 8 bits. It is accessed with a LabVIEW⁸ program, which emulates the trigger algorithm originally implemented in FPGA hardware. The LabVIEW program consists of the RFI-suppression, which can be switched off, the envelope creation and a final pulse-finder testing a full series of thresholds. While this trigger realization with a computer program has advantages in terms of flexibility, it is by a factor of 450 slower than the FPGA-based real time system. Analyzing 1000 events of 200.000 samples each, which is equivalent to a time slice of one second background data, requires about

⁷ Universal Serial Bus - A standard interface for computer equipment.

⁸ National Instruments LabVIEW is a visual programming language, which is designed for comfortable data acquisition and processing.




Figure 4.13.: False trigger rate at the BLS versus the threshold for different positions. Because no strong monofrequent carriers were present, the RFI-suppression has only little effect and was not employed.

eight minutes, which is fast enough for exemplary background measurements. As all electronics is powered via USB from the Laptop battery, and the antenna is compact and sturdy enough to be mounted to the roll bar of a pickup truck, the system is highly mobile to survey various positions. The disadvantage of a mobile setup is however, that the recorded data at each spot spreads only over a few minutes. Highly variable conditions, for example due to weather conditions, are hardly covered. Instead, the measurement tries to give a characteristic representation of the background noise under regular conditions. Outliers can be identified by comparing neighboring measuring positions for plausibility.

At first the system was employed at the BLS, where the background conditions are already well known, in order to cross check the measurements. The resulting threshold behavior of the east-west polarization at four selected positions is shown in figure 4.13. Two measuring positions are located directly besides the known antenna Poles 1 and 6 respectively. The found trigger rates comply well with the values of the prototype trigger setup at comparable thresholds, the noise level close to the BLS is increased. The next location is about 300m further into the field close to the SD tank Olaia, where the background situation is significantly better. At the last position on the half way between BLS and powerline, the noise level is the most elevated, the transient rate stagnates around 10 kHz and hardly drops towards higher thresholds. The worse noise situation close to the powerline is a final, strong hint towards the powerline as responsible transient noise source rather than the measurement equipment at the BLS.

A second series of measurements was executed at the envisaged site of the successor radio detector in the northwest of the Auger observatory (s. fig. 2.4). From the survey, four representative spots as marked in the map 5.1 are chosen and their east-west polarization is presented in figure 4.14. As the new field is again bordered by a powerline to the north, the worst case location directly under the overhead wires are in the focus of attention. Although realized as symmetric three-phase transmission, which should strongly reduce interferences compared to an asymmetric SWER-line, the high-voltage line is still a dominant source of transient noise, which forbids the placement of antenna




Figure 4.14.: False trigger rate at the AERA site versus threshold value for different positions. Solid lines show the rate without, dashed lines with RFI-suppression.

stations at least directly under the wires. With larger distance to the transmission lines towards south the noise situation becomes better. In the middle of the field, the transient rate is lower than at the most silent Pole6 at the BLS. At the fence which limits the field to the south the rate is even lower than at the remote Olaia position. Because the east-west polarization contains a distinct carrier signal at 67 MHz, the RFI-suppression can effectively improve sensitivity by removing the monofrequent signal. The north-south polarized component of the radio signal shows the same behavior. It is not plotted here, as it is less important for self-triggering of cosmic radio pulses, which are predominantly east-west polarized,

Other measurements were conducted to search for noise from the surface detector, but no increase of the transient rate close to the water tanks was found. Because noise pulses from the photomultipliers or the communication of the tanks may be masked due to their low rate, they can not be finally excluded by this measurement. Consequently the grid of the radio detector stations will still be arranged in a way to keep a maximum possible distance to the tanks.

4.5. Conclusion

The measurement campaign at the BLS successfully proved the SALLA as a suitable antenna candidate for radio detection in the rough environment, as no mechanical failures occurred despite heavy storms with wind speeds up to 200 km/h. The sensitivity of the 120 cm SALLA is good enough to detect the inevitable galactic noise, although the antenna picks up more artificial noise from horizon, due to its wider main lobe. The LNA turned out to be rather intermodulation-prone, particularly in combination with the SALLA, which misses an intrinsic low-pass behavior. Meanwhile both the sensitivity and the intermodulation issues were finally solved by an upgrade of the LNA.

Initial expectations of the campaign assumed at least the detection of one cosmic ray event per week. After a successful proof of principle it was planned to tune trigger parameters to achieve a false trigger rate as low as possible. The first objective, the self-triggered detection of cosmic rays, was however missed. On the one hand, the fail-

ure was caused by the noisy radio background at the BLS, which was underestimated from the beginning. Although the continuous background level was lower, the transient rate at the BLS was at least on par with the industrial environment at the KIT, which led to a higher necessary threshold than expected. On the other hand, the radio emission from the cosmic rays was overestimated. The parameterization obtained from the LOPES-experiment could not be transferred directly to the conditions in Argentina, as it was measured at a different height above sea level in lower geomagnetic field, with an emphasis on lower energies and smaller distances.

When the first self-triggered events failed to appear, the simplified prototype setup did not allow for an alternative strategy. Neither a lowering of the trigger requirements nor a long-term operation with an external trigger was possible, an extensive monitoring of rejected events in so-called minimum bias files was not foreseen. This led to demands for a more flexible data acquisition system for the successor experiment with the possibility to monitor every trigger stage. We also gained valuable experience in the remote operation of an FPGA-based trigger system, leading to additional demands for the successor electronics, like a watchdog circuit and redundant flash memory. The progress of the campaign and its result were reported to the Pierre Auger collaboration in an internal GAP-Note [86].

The planned site of the successor experiment in the northwest offers an improved noise environment at least at some distance of the local power line, according to survey measurements. Because the survey has been conducted within only one day, the knowledge about temporal noise variations is limited though. The development of strategies to effectively suppress periodic noise transients from the power grid, or horizontal events, could therefore be crucial. To keep a high amount flexibility for the future trigger and data acquisition system, the new electronics is based again on widely reconfigurable FPGA logic.

The trouble to self-trigger a cosmic ray also illustrates that even with an improved system, most of the recorded events will be close to the detection threshold, and therefore only trigger few antenna stations. To gain as much information as possible from the radio events, the data of all neighboring antenna positions is however desirable. Even if the recorded signal is too weak to trigger, it can be still reconstructable, and if no signal is visible at all, at least an upper limit information is provided. The requirement to read out data from untriggered stations means that a local ring buffer must bridge the communication latencies until the readout request arrives. The ring buffer also enables the possibility to externally trigger the radio detector, which finally expands the Pierre Auger observatory from hybrid detection, with an implemented mutual trigger of the fluorescence and surface detectors, to a joint triple detector.

While the prototype campaign was planned, built up, operated, and analyzed with the manpower of just two people, the full-scale successor experiment is obviously much more elaborate and requires the effort of a full-grown collaboration. Therefore, this thesis will concentrate on the low-level signal processing and early trigger stages in the following.

5. The AERA experiment

Radio detection of cosmic rays works well only above primary energies of $5 \cdot 10^{17}$ eV. At these energies the cosmic ray flux has dropped to only 3 events per square kilometer per day. Therefore an adequate radio array needs to cover an area of several square kilometers to collect enough statistics. In this early stage of radio measurements, it is important to collocate the antenna array with an established experiment, in order to compare measurements and cross-calibrate the detectors. The Pierre Auger Observatory as the leading cosmic ray experiment provides an optimal site for a large radio detector. This way the hybrid approach of parallel fluorescence and surface detection can be expanded by a third method.

To build such a full scale detector the efforts from three radio groups at the Pierre Auger Observatory were joined in the Auger Engineering Radio Array (AERA) collaboration. Besides the measurements based on LOPES^{STAR} described in the last chapter, the MAXIMA group from the Netherlands also operated a setup of antennas at the BLS using wireless communication [78]. The third partner, the French RAuger group originates from the CODALEMA collaboration, and was measuring close to the CLF¹ [88]. The shared experience led to the design proposal of the AERA detector as described in the following.

5.1. Key features of AERA

The site of AERA was chosen in the north-west of the Pierre Auger Observatory (s. fig. 2.4), where the upgrades HEAT² and AMIGA³ are collocated. The field is limited in the west by the river Rio Atuel and in the south by a fence separating the property of a land owner, who does not allow the installation of radio antenna stations. In the north, a high voltage power line introduces additional background noise, and thereby adds a third border. As a result, the AERA field has a polygonal shape as shown in map 5.1, and covers an area of about 7 km². The antenna stations are arranged in equilateral triangles according to the closest circle packing. Their density varies from a 150 m grid to achieve a high detection efficiency for cosmic rays close to the detection limit of $2 \cdot 10^{17}$ eV, up to a 350 m spacing to cover a larger area for the very rare events well above 10^{18} eV. AERA is constructed in three stages, the dense inner core of 24 stations was build in autumn 2010, 52 additional stations are planned for 2012, the last stage will finally complete the array in 2013 to a full number of 150 stations. To get a perspective of the leap forward resulting from the AERA experiment: Already the first stage of AERA covers 10 times effective area of the LOPES experiemnt.

¹ Central Laser Facility - A calibration setup for the fluorescence detector located at the center of the Pierre Auger Observatory[87].

² The High-Elevation Auger Telescopes expand the energy range of the fluorescence detector towards lower energies [89].

³ Auger Muons and Infill for the Ground Array - a setup of buried muon detectors and additional surface detectors to improve the primary mass sensitivity, and lower the energy threshold [90].




Figure 5.1.: Map of the AERA experiment, with the underlying satellite image from Google earthTM. Stations of the first stage are marked with red/orange antenna symbols, subsequent antennas are colored yellow. The green house symbol shows the position of the central radio station during stage 1, yellow circles indicate surface detector tanks, and red dots mark the survey measurement positions.

Because of the remote installation, each station must operate autonomously at a large extent. Electrical power is only supplied by solar panels and buffered with batteries, the available power budget is strongly limited. To allow for synchronous measurement, the radio stations need a common time base, which is provided by a GPS clock at each station. During the first stage, the antenna stations are connected with a military grade optical fiber, which is buried 10 cm deep, and allows a large communication band width, using commercially available Ethernet hardware. The data acquisition system is located in the close-by Central Radio Station (CRS), using a 40 ft ISO-container for housing and a large array of solar panels for power supply. For the final extension of the AERA field, however, installation and maintenance of many kilometers of fiber network is not economic anymore. Instead, a wireless communication system will be used, which also allows to move the CRS to the Coihueco fluorescence detector building on a hill 4 km west of AERA, where power grid connection is available.

The scheme of data acquisition is designed to comply with the latency and band width limitations of the wireless communication system (s. fig. 5.2): Each antenna station sends only the timestamps of its local T2-triggers to the CRS, where the data is sorted. Only if T2-timestamps from multiple stations fall within a certain coincidence time, a T3-request asks the stations for a chunk of ADC data from that timestamp. The ADC data sent from all requested stations is finally combined and stored as a joint radio event to a database at the CRS.

The tasks of the central radio station are mainly realized in software and implemented on standard computer hardware. For the local station, the system requirements suggest the following six main components:

The **housing** of the local station is made of welded sheet steel and provides weather protection and electromagnetic shielding. Besides holding off rain and sand storms, a proper thermal management is important to prevent overheating from the blazing sun during summer. For this purpose the used casing has double




Figure 5.2.: Sketch of the data acquisition topology of AERA.

walls with air circulating in between, the solar panel offers additional shadow. A plastic fence around each station prevents cattle and goats from damaging the setup, especially the antenna.

The **photovoltaic system** consists of solar panel, battery and charge controller. By using a supply voltage of 12 V, a large variety of commercial systems is available. For a continuous power draw of 10 W, a solar panel with a peak power of 130 W and a 150 Ah lead battery are sufficient [91].

The **GPS-clock** allows for a synchronous measurement of the distributed stations. Commercial GPS modules have a power consumption of only 350 mW and a claimed time resolution of better than 5 ns when applying all corrections. Additionally a reference beacon permanently transmits a set of constant frequencies in the measuring band, which are recorded within the event data. During posterior analysis, the phases of the beacon signal at the different stations can be used to improve the relative timing below 1 ns, as necessary for interferometric beamforming analysis [45].

For the **antenna**, three candidates are available as discussed in sec. 3.1.1. In phase 1, the wired LDPA is used, the antenna for phase 2 and 3 shall be the Butterfly, although the final decision is still to be made. A corresponding LNA amplifies the signal, a band filter removes the short wave band below and FM radio above the measuring band, and serves as an anti-aliasing filter before digitizing.

The **communication** system in the first stage consists of commercial optical fiber converters. For the following phases, a custom-made wireless system is being developed, which uses the known WLAN⁴ band at 2.4 GHz. With regular CSMA⁵ multiplexing of more than 150 subscribers, the network would be completely jammed by collisions. Instead, the precisely known time from the GPS-clock is exploited by using a TDMA⁶ policy with fixed timeslots. With a raw bitrate of 5.5 Mbit/s, each station has a payload of 12 kbit/s at its disposal, distributed over 3 timeslots per second. The development of the communication system is the topic of the PhD thesis of Yue Zhu [92].

⁴ Wireless Local Area Network in this case based on the IEEE 802.11b standard.

⁵ Carrier Sense Multiple Access is a method to share a channel between several nodes. To send, a node waits until the channel is free, indicates a request-to-send, waits if the channel stays free, and starts transmission. It becomes inefficient with a large number of nodes due to collisions.

⁶ Time Division Multiple Access is a method to share a channel between several nodes. Each node has a defined timeslot during which he may send. As no collisions occur, it is efficient for many nodes.

The **digitizer** implements the tasks of sampling, signal processing, triggering, and the local station parts of data acquisition and status monitoring. The digitizer and its firm- and software are the technical keystone of the radio detector, as discussed in the following sections.

5.2. Hardware development for AERA

The roadmap for the construction of AERA was developed during several collaboration meetings in 2009. To work out the technical specification of the digitizer electronics during these meetings, based on the experience from the prototype setup, was an important topic of this thesis. The negotiations finally reached the agreement to build the new digitizer hardware at the KIT⁷ with support in driver and software development from the Bergische Universität Wuppertal. The system is therefore referred as the KIT/BUW digitizer [93]. The actual design, construction, and commissioning of the digitizer, which is described in this section, was performed by the fellow engineers of the IPE. The block diagram of the resulting digitizer (s. fig. 5.3) and the hardware board itself (s. fig 5.4) are discussed briefly to understand the firmware and trigger development in this work.

An important overall requirement is a low power consumption, as the solar power supply is a vital system cost multiplier. The most power efficient approach is an embedded system consisting of an FPGA for the early high speed signal processing and a rather low performance processor for the later data acquisition. A total power budget of 10 W is the development target. The FPGA was chosen from the Altera low cost family Cyclone IIITM, which shows significant advantages over the high-end StratixTM series in terms of power consumption, complexity, and price, while the performance is sufficient (s. app. A.3.1). The used EP3C80F780C6 is the largest available model in the fast speed-grade 6, and contains 80,000 logic elements, which is twice the number of LEs in the prototype setup. The FPGA is the main current consumer, when running the firmware described in section 5.3 it consumes 3 W.




Figure 5.3.: Block diagram of the AERA digitizer hardware. Yellow blocks show the two analog channels, interfaces to the exterior are colored blue, the processor devices are red, and auxiliary units green.

⁷ Involved institutes are the Institute for data processing and electronics (IPE) and the Institute for nuclear physics (IKP).



Figure 5.4.: Photograph of the AERA digitizer.

To increase the initial flexibility, the KIT/BUW digitizer offers the option of either using the processor module from the prototype setup with a 2x600 MHz Blackfin processor, or implementing a 125 MHz Altera NIOS II⁸ soft core inside the FPGA. In the latter case the processor module is exchanged with a module only containing 128 MB working memory and an Ethernet physical driver. By abstaining from dedicated processor hardware the power dissipation is reduced by 1 W. As the NIOS II processor proved to be powerful enough to run the μ Clinux operation system and fulfil the DAQ tasks, the higher performance Blackfin option was not seized. The processor has access to several peripheral devices on the board through a variety of interfaces, which are described in appendix A.3.2.

Because the amplitude of a coherent pulse with a flat spectrum is increased linearly with the recorded band-width, while the uncorrelated noise level grows only with the square-root, the signal-to-noise ratio of cosmic ray induced radio pulses can be improved with a larger band-width. Because the emphasis of the pulses lies particularly at lower frequencies (s. fig. 2.5), it is highly desirable to expand the measuring band down to the full frequency range of 30–80 MHz available between the short wave and FM interferences. This forbids the method of sub-sampling, and requires a sampling rate of at least 160 MHz. By using a slightly higher sampling rate of 180 MHz, the requirements for the anti-aliasing filter are relieved. In order to prevent saturation of high amplitudes from close showers on the one hand, and allow a sensitivity down to the galactic noise level on the other hand, a high dynamic range with an effective resolution of 16 bits is necessary. Unfortunately, no low power, high speed ADCs with this resolution are available, even 14 bit ADCs usually only have an effective resolution below 12 bit⁹. A combination of two 12 bit ADCs¹⁰ per antenna polarization channel turned out to be the best solution: The analog signal is split into a low gain channel and a high gain

⁸ The Altera NIOS IITM is a proprietary processor architecture, optimized for implementation on FPGAs.

⁹ The considered Analog Devices AD9643 14 bit ADC for example has effective resolution of 11.5 bit according to its datasheet.

¹⁰ Using the Analog Devices AD9626 ADC with a power consumption of 0.3 W each.

channel with 20 dB more amplification, each channel is fed to a separate ADC. For regular data taking the high gain channel contains all information, only in the very rare, but important case of saturation the low gain channel data is used. This design approach also increases the application flexibility, for example a modification to record a third antenna polarization would be possible with ease.

Most of the recorded air shower events will only trigger 3 antennas in a triangular arrangement. To increase the obtained event information and learn more about characteristics like lateral distribution or shape of the radio front, it is very important to read out also the surrounding antenna stations. They may have recorded a signal as well, although it is too low to cause a trigger. To make the read out of untriggered data possible, all ADC data must be buffered for the time span needed for a trigger decision. This time span is the sum of the local station trigger delay, the communication latency for the T2 timestamp upload, the coincidence trigger delay, and the communication latency of the event request. With the wireless communication approach of a TDMA period of 1/3 s, the needed buffer time is in the order of seconds. As a result, the data rate of 540 MB/s from two 12 bit channels at 180 MHz, calls for a large buffer memory. Using one regular 2 GB SO-DIMM¹¹ module of DDR2-SDRAM¹² provides a buffer time of almost 4 s for a reasonable price with a power budget of 2 W. The large buffer time also allows for an external trigger from the surface detector, which complies well with the general hybrid approach at the Pierre Auger Observatory, to let one detector improve the other.

Regarding a large array with many local stations low maintenance is very important, which requires superior software stability and the possibility to upgrade firm- and software remotely. For this purpose a Texas Instruments MSP430 ultra low power¹³ microcontroller serves as a watchdog: A program running in the μ Clinux operating system will reset the watchdog frequently. If the system hangs and the watchdog reset is omitted for a certain time, it will power cycle the system. The watchdog can even switch between two flash boot devices, which store different firm- and software images. If one image is faulty and crashes, the system uses the secondary device, an on-site interaction is not necessary. For the variety of components, a number of different voltages must be generated from the 12 V supply with a network of switching regulators. They account for an additional power loss of 1.8 W due to inefficiencies.

To dissipate the accumulated heat the digitizer PCB includes two dedicated copper layers to spread the heat to the sides, and dissipate it to the metal enclosure. The FPGA is fitted with a small additional heat sink. During temperature tests, proper operation in an ambient temperature range between -15°C and 65°C was proven. The electromagnetic compatibility (EMC) of the digitizer is another very important issue in order not to increase the own radio noise level. Switching supplies and clock signals are usually critical noise sources using feed cables as antennas. Measurements in an EMC box showed the compliance of all important noise limits. The digitizer boards were subject to several further laboratory tests. For example, the ADC inputs were checked for random noise level, offset, crosstalk, amplitude linearity, frequency linearity, and ring tones from clocked signals. All measurements showed satisfying results [94].

Besides the new made-to-measure electronics, a second digitizer is provided by our Dutch collaborators. Because this digitizer is only an updated version of the Maxima setup [78], it misses important features like the large ring buffer, or an interface for the

¹¹ Small Outline Dual In-line Memory Module is a form factor used for notebook memory modules.

¹² Double Data Rate Synchronous Dynamic Random-Access Memory is an interface standard used for computer memory modules.

¹³ MSP430 typical power dissipation is 0.3 mW

planned wireless communication system. It was however available almost immediately, so the AERA stations of phase 1 were first equipped with the Dutch electronics, and later replaced with the KIT/BUW digitizer.

5.3. Self trigger firmware

The digitizer hardware, as described in the last section, is a truly universal data taking system with fast ADCs, an all-purpose FPGA, and a variety of communication possibilities. It is the FPGA firmware which specifies the field of application, and turns the digitizer into a piece of radio detection equipment. To handle the multitude of necessary tasks, several modules are implemented in parallel: The central NIOS processor module is superordinate over all peripheral entities, which are connected via the Avalon bus¹⁴. All modules are thereby accessible from the processor through a certain bus address range, and are, with proper drivers, visible from the μ Clinux operating system as Linux devices. The modules are assembled to a so-called system-on-a-chip with the SOPC-builder¹⁵, which delivers VHDL source files to compile the FPGA firmware, and a system description file to compile a made-to-measure μ Clinux kernel.

Most of the modules in the resulting firmware (s. fig. 5.5) are used for communication with hardware outside the FPGA. These standard applications are all available as IP-cores¹⁶, which are ready to use or need only minor platform-dependent modifications. They are described in the appendix A.3.2. In contrast, the two main modules, the data management unit (DMU) and the trigger, are custom-designed, as they are very specific for the radio detection purpose.




Figure 5.5.: Block diagram of the AERA firmware. The central processor (*red*) communicates with all modules (*grey*) via the Avalon Bus. The main user modules are the data management unit (DMU) (*green*) and the trigger unit (*blue*). External hardware (*yellow*) is only shown for illustration.

¹⁴ Altera AvalonTM is a proprietary parallel bus standard, used to connect entities inside the FPGA.

¹⁵ The Alters System-On-a-Programmable-Chip builder merges a NIOS processor with peripheral devices on one FPGA.

¹⁶ An Intellectual Property core is a reusable hardware design, which may be licensed to use.

5.3.1. The Trigger module

The only task of the trigger module is to detect notable pulses in the ADC data, and report them to the subsequent DMU.

Although high gain and low gain ADCs digitize the same signal with merely a different analog amplification, the two 12 bit signals are not merged to one 16 bit signal, but handled separately. On the one hand, merging the signals without introducing discontinuities is not trivial. On the other hand, the low gain channel only contains additional information, if the high gain channel saturates. This happens less than once per hour according to field measurements. Permanently processing a 16 bit wide signal, although the most significant bits (MSB) are almost exclusively zero, is less efficient, than handling the rare cases of saturation separately. A simple threshold applied on the low gain data safely detects high gain saturation, and is, with a programmable delay, used as low gain trigger.

In contrast, the threshold of the high gain trigger is critical, as it defines the cosmic ray detection limit. A sophisticated RFI-suppression module is therefore used, which removes interfering radio transmitters from the measuring band, and lowers the noise level as much as possible (s. ch. 6). The subsequent pulse detection module registers each signal above a threshold as a so-called T1-event. The T1-events are further analyzed. If pulse characteristics like shape or timing bear a suspicious resemblance with typical noise background, they are rejected (s. ch. 7). Radio pulses accepted by the high gain trigger, the T2-events, are provided with basic pulse shape parameters like its width and height, which may be used by the next trigger level.

Besides the radio self-trigger, the use of an external trigger supplied for example by a local scintillator detector is also possible. It had been foreseen to simplify the commissioning of the system with an immediate external trigger signal, although this option was not used so far. Because high gain and low gain trigger are implemented for each antenna polarization separately, a total of 5 trigger sources are merged to a joint antenna trigger. The trigger strobe together with information about the trigger source, and for high gain triggers with pulse shape parameters, are forwarded to the DMU, to finally start the event readout.

The entire trigger is by far the most elaborate module in the FPGA design. It accounts for 34.000 logic elements corresponding to 42 % of the FPGA resources.

5.3.2. The data management unit

If an event is triggered, not the RFI-suppressed data from the trigger module is read, but the raw, uncorrected data from the ADCs. Based on the raw data, any signal processing can still be applied during off-line analysis. The dedicated module, the data management unit (DMU) serves as an interface between the processor and the ADCs, for this purpose it fulfils 4 tasks:

The GPS-clock provides two types of time information, on the one hand a rather slow time telegram, which contains the full UTC-time¹⁷ with a 1 s resolution, and is read by the processor, and on the other hand a simple signal strobe exactly at each start of a second (see also A.3.2). The **time manager** basically consists of a counter running at clock frequency of 180 MHz, which is reset by each pulse-per-second (PPS) signal. Thereby the counter reading assigns a sub-second time value to each clock cycle, which

¹⁷ Coordinated Universal Time (UTC) is the primary global time standard.

has a time resolution in the order of 10 ns ¹⁸. An additional 3-bit counter is increased at each PPS signal. It contains the corresponding second value and can be set by the processor to agree with the 3 LSBs of the UTC-second. Both counters together form an internal timestamp, which is unique within 8 s.

The **event builder** throws a processor interrupt for each event from the trigger and calculates the timestamp of the triggered pulse, corrected by the latency of the pulse detection algorithm. The resulting pulse time is stored with the triggering channel and the pulse parameters from the trigger in a FIFO¹⁹, so that they can be read from the processor. The FIFO can store 50 entries. If it is not read and gets filled up, new triggers are ignored.

The low gain channel is only necessary for very high amplitudes, when the high gain channel saturates, which rarely happens. Therefore it does not make sense to waste precious capacity of the large ring buffer with data which is needless most of the time. Instead the **low gain buffer** stores this data in a tiny 2048-word buffer of FPGA internal memory. If a low gain trigger occurs, the data is retained until read from the processor. The latency of the low gain trigger is adjusted, so that the signal pulse is in the middle of the stored window. In the meantime a second 2048-word buffer fills in. If it is also triggered, before the first one is cleared, new low gain triggers are ignored.

Data from the high gain channels is continuously stored by the **high gain buffer** in a 2 GB memory module, so that even data which was not self-triggered can be read upon external request. Fortunately, most of the complicated interfacing towards the memory module is done by an Altera IP-core²⁰, which provides access through a continuous 27 bit address range and a 128 bit data word. Each data word is filled with 5 ADC values from each polarization channel, which accounts for 120 bits. The remaining 8 bits per word are left empty to avoid a complicated overlapping over word boundaries. The resulting reduction of the fill ratio to 94 % is acceptable, leading to a buffer time of 3.73 s. Writing the buffer basically consists of continuously increasing the write address up to the overflow. At the same time, the write address offset at the start of each second is stored. When the processor requests a read access at a certain timestamp, the stored address offset of the requested second is added with a fifth²¹ of the requested subsecond value to receive the corresponding memory address. Starting from this address, a burst of 2060 samples is read and stored to a FIFO, which is then read out by the processor. To continue seamless data recording also during read accesses and regular DRAM-refresh cycles, the incoming ADC data must be fed through an additional write buffer of FPGA internal memory. This also enables the use of fast burst write accesses, which is indispensable in order to achieve the required data throughput of 580 MB/s. With substantial modifications, the DMU can also handle other memory module sizes than 2 GB. To allow for an external trigger with a latency of up to 7.4 s, a 4 GB version was created as well.

¹⁸ The sub-second time resolution can be improved, by post corrections of the GPS-clock sawtooth, and the crystal clock deviation of 10^{-6} from its denomination frequency.

¹⁹ First In First Out is a memory, from which words are read in the same order as they were written to.

²⁰ The Altera DDR2 SDRAM High Performance Controller handles the double data rate transfers at both clock edges, the refresh of the dynamic memory each 64 ms, the bank, column, and row based addressing, and keeps the preset timing constraints of the memory module. It also supports burst access for consecutive addresses, which can easily increase data throughput by a factor of 10.

²¹ Divided by 5, because 5 ADC-cycles are stored at one memory address, the division is actually calculated by the software driver. The driver also cuts the read out window by the modulo of the division, to get exactly the requested ADC-sample.

The entire DMU occupies 6200 logic elements and 250 kbits of memory, which accounts for 8 % of the FPGA resources. The complex DDR2 SDRAM controller must be driven at a reduced clock frequency of 125 MHz, to fulfil the timing requirements. The resulting data throughput is still sufficient. The reduced memory clock also reduces the power dissipation of the memory module. For proper DMU operation, a data acquisition process must be running on the processor, which reads the event builder FIFO and the low gain buffer in time, if data is present.

6. Suppression of radio frequency interferences (RFI)

The background which is encountered during radio detection of cosmic rays consists of three different components. First, there is the typical random noise, which is widely distributed both in time and frequency domain. Inevitable sources are the galactic noise and preamplifier noise. The second background component are continuous carrier signals from radio stations, which have a narrow frequency pattern. And the third contribution comes from noise transients, which are for example caused by electrostatic discharge and switching power supplies. These noise pulses are short in time, but wide in frequency domain.

The combined background level defines the dependency of the unwanted background trigger rate on the detection threshold according to e.g. fig. 4.14. These diagrams usually show a distinct kink, where the false trigger rate stagnates towards higher thresholds, but steeply rises already for a slight decrease of threshold. The kink is a result of the underlying amplitude distribution (s. fig. 4.3), and is located at the transition from the continuous noise floor to transient noise. It indicates an operating point, where the threshold requirement discriminates signal from background most efficiently. Because the threshold can not be lowered into the continuous noise floor without massively increasing the false trigger rate, the continuous background from the first two components defines a lower limit on the threshold level. Above the limit, the false trigger rate is dominated by the noise transients, a higher threshold hardly decreases the false trigger rate.

While the false trigger rate can still be reduced later on in a second trigger stage (T2), as described in chapter 7, the threshold level can only be improved by lowering the continuous noise floor beforehand. For this purpose the artificial radio frequency interference (RFI) from radio stations is suppressed: After Fourier transforming the original signal into frequency domain, the interfering monofrequent contributions can easily be identified and removed, then the signal is transformed back to the time domain. The signal from a potential radio pulse spreads flat over all frequencies in the measuring band. The actual pulse information is hidden in the corresponding phases, which define where all frequencies superimpose to the full pulse height. The removal of some frequencies hardly affects the pulse height, it only reduces the effective bandwidth slightly. In fact, the bandwidth is already reduced by the interference itself, as the phases of the tainted frequencies are determined by the external signal anyhow, and not by the pulse. Therefore they do not contribute coherently to the pulse height, but only in average power. On an individual basis, they might even reduce the pulse height.

This kind of signal processing has already been used in the posterior analysis of recorded radio pulses at the LOPES and LOPES^{STAR} experiment (s. sec. 3.1.4). Now it is also implemented in the real-time data taking to improve the self-trigger conditions. The feasibility of the RFI-suppression engine in FPGA logic was shown in the diploma thesis of M. Manewald [95]. The implementation was resumed during this thesis, and then field-tested during the prototype measurements at the BLS. With the new AERA




Figure 6.1.: Dynamic Spectrum measured with antenna station 20 in the South of the AERA field. The noise level of the east-west polarization (*left*) is comparable to the north-south polarization (*right*) apart from the different carriers.

electronics a new and improved version of the RFI-suppression was developed, which benefits from both the increased FPGA performance and the experience made before. The final implementation and its validation were published in [96] and will be described in this chapter.

Obviously, the RFI-suppression is only advantageous when the radio background contains monofrequent interferences. At the BLS for example, the RFI-suppression has hardly any effect on the threshold (s. fig. 4.4), as the radio spectrum shows no distinct carriers (s. fig. 4.2). The spectrum at the AERA site, however, is polluted by strong transmitter at 67 MHz in the east-west polarization (s. fig. 6.1). The north-south polarization contains three frequencies from the beacon, which is used to improve the system's time resolution [45]. As a result, an improvement due to the RFI-suppression is already visible in the background survey (s. fig. 4.14).

According to the calibrated spectrum, the single carrier in the east-west polarization concentrates 84 % of the average noise power and spans over 2.8 % of the used bandwidth. This means that the average background amplitude is reduced by 60 %, if the carrier is removed completely. The impact of the carrier on the measured pulse height depends on the individual case, whether the two interfere constructively or destructively. On average, the peak power is increased by the power of the carrier¹. By applying the RFI-suppression, the pulse is returned to its original height, taking into account that the bandwidth is reduced by 2.8 %. In the north-south polarization, the effect is lower, the three beacon frequencies together account for 57 % of the background power and cover 1.0 % of the measuring bandwidth. Although the implementation of the RFI-suppression will also introduce additional numerical noise, the achievable improvements are promising, in particular in the more important east-west polarization (s. fig. 6.2).

6.1. FPGA Implementation of the RFI-suppression

Using the RFI-suppression algorithm, which is well-proven in offline event analysis, generates a couple of problems in real-time operation : First, there is the sheer computing power, which is required to process the input data rate at a sampling frequency of 180 MHz. When taking into account the tight power budget due to the solar supply,

¹ The mean power density of the 67 MHz carrier is about 3.3 pW/m²




Figure 6.2.: Effect of the firmware-implemented version of the RFI-suppression on a real cosmic ray induced pulse. The trace was measured by the Dutch electronics in the east-west polarization.

the only remaining implementation possibility are FPGAs, as mentioned earlier. This means that the whole algorithm must be translated into low-level hardware description language, including the transition from floating point algebra to integer calculations. A special attention is paid to the numerical noise introduced hereby.

Second, the algorithm was only applied on short time slices of event data so far, which is highly advantageous for the already block based Fast Fourier transform. Numeric artefacts at the block edges are for example unimportant, as the interesting signal pulse is usually centered in the event data. For the implementation in front of the trigger, however, a continuous data stream must be handled. Therefore the data is split into blocks, fed through the RFI-suppression, and then again reassembled as a seamless stream.

The sketch of the implementation for both polarization channels is shown in figure 6.3:

The first unit prepares the incoming data for the subsequent algorithm by dividing the continuous stream into blocks. In order to avoid frequency leakage, each block is applied with a **window function**. To prevent the loss of signal pulses which are located directly at the block edge, a certain overlap of consecutive blocks is necessary. The overlapping areas must be processed twice, which means that the full RFI-suppression has to handle even more data, and thus run at a faster clock frequency of 210 MHz.

After windowing, the signal is transformed into frequency domain using a **FFT engine**. As the input signal is purely real, the positive and negative frequencies are symmetric to each other, for an imaginary signal they are antisymmetric (s. appendix A.1.2). Therefore both polarization channels can be handled with one FFT engine, when feeding them into the real and imaginary inputs respectively.

By separating the symmetric from the unsymmetric contributions, both polarization channels are **demerged** again. In doing so, the demerging unit has to await the complete FFT-block, before passing the first value to the further pipeline, which significantly increases the latency of the RFI-suppression engine.

As the signal is already available in frequency domain, the unwanted dispersion of the analog filter can be removed by just a complex multiplication with a set of correction factors. The **deconvolution** of the filter's phase shift compresses




Figure 6.3.: Block diagram of the RFI-suppression. Thin arrows stand for the real inputs and outputs, bold arrows illustrate the internally used complex signals.

the pulse to its original shape, the pulse height is increased by 25-30 % according to [97].

By applying a **median filter** on the amplitudes of the spectrum (s. sec. 6.4), while keeping the phases constant, the monofrequent carries are removed. The pulses which have a smooth frequency spectrum are not affected.

Finally the signal is **transformed back** to time domain. To get the signal envelope, as necessary for the subsequent pulsefinder, the preferable method is to use the analytical signal, which consists of the original real trace plus its Hilbert transform as imaginary part (3.1.4). Coming from frequency domain, this would normally require one FFT-engine for the regular signal, and one for the Hilbert transform. By exploiting again symmetries of the Fourier transform, both transforms can be performed simultaneously with one FFT-engine as explained in appendix A.1.4.

The envelope is then just the **absolute value** of the analytical signal.

To get back to a continuous data stream, the consecutive **windows are merged** again, overlapping areas are thrown away. At the same time, the transition from the higher clock frequency down to the regular 180 MHz is performed. A proper synchronization with the initial window generation module guarantees a seamless concatenation and a constant pipeline latency.

The use of an additional, faster clock domain aggravates latent issues of the RFI-suppression to keep the FPGA timing requirements. In particular at the more complicated entities like the FFT-engines or the median filter, the implementation is often a tightrope walk between the desired numerical accuracy and bit resolution on the one hand, and the attainable performance on the other hand. Some of the applied optimizations including parallelization and pipelining, will be described in the following.

6.2. Choice of the FFT-engine

One central module of the RFI-suppression is the Fast Fourier transform, which is a common application in digital signal processing. Therefore several FFT implementations in HDL already exist, some with an emphasis on high speed, others on low consumption of FPGA resources. The first choice would be an open source engine

provided by OpenCores², which are however not flexible enough for the requirements regarding bit-width and block length. Second, there are professional companies, selling made-to-measure, well-performing solutions. Unfortunately the target market of these IP³-cores is large-scale production, thus the licenses are hardly affordable. The third possibility are IP-cores provided by Altera with licenses already included in the Quartus development environment, which also bear the advantage that they are particularly optimized for the Altera FPGA architecture.

Bit resolution

From the offered IP-cores, only the fully pipelined implementations supporting real time data streaming are qualified. This leaves three different engines, which differ in terms of numerical accuracy, execution speed, and FPGA resource consumption (s. tab. 6.1): The most elaborate one is the floating point implementation, for which a round trip calculation, consisting of a FFT into frequency domain followed by an inverse FFT back into time domain, delivers the exact input signal down to the least significant bit. Because a single floating point FFT already consumes more than a quarter of the FPGA resources, and the operating speed is far from the requirements, it is however out of question.

The variable streaming implementation is less complex. It uses integer algebra for the complex multiplications, and thereby introduces a rounding error of one LSB in the round trip reconstruction. The bit resolution of the signal is increased gradually during the iterative summation in the FFT algorithm, so the resolution of the output signal is by $\log_2(N)$ higher for a FFT with a block length of N . As the increased resolution hinders the further data processing, particularly after the inverse FFT, which again increases the bit-width. The bit-width must be cut down to 12 bits again, which introduces an additional reconstruction residual as shown in figure 6.4. But even with this reduction of complexity, the limited operation speed is problematical (s. tab. 6.1).

The fastest implementation is the fixed streaming method, which does not increase the bit-width during the FFT iterations, but scales down the intermediate results with a right shift when needed. The number of applied scalings is remembered, and serves as block-wise constant exponent in a type of floating point representation. Because all frequency bins of the Fourier transformed signal share the same exponent, the representation does not increase their dynamic range, which stays limited to the input bit-width. This causes numerical errors, when one strong frequency contribution forces an extensive scale down, so that the other frequency bins, which contain the potential signal pulse, are limited to a range of a few bits. This happens particularly when the input signal from the ADCs contains a strong DC offset, which translates to an elevated first frequency bin. Therefore the constant offset is subtracted beforehand with a small FPGA entity.

To evaluate the effect of the monofrequent carriers at the AERA site, the round trip reconstruction error is calculated, using measured background data from both antenna polarizations at AERA. The background noise was measured with the competing Dutch digitizer, because it was already installed in the field. As both systems uses the same analog path, the measurement is transferable, when correcting for the different sampling rate and ADC gain. The background signal is enriched with several artificial pulses. For the pulses the response of the analog filter to the simulated radio signal from figure 2.5 at a distance of 100 m is used. Then the signal is down-sampled from the simulation

² OpenCores.org is a large open source hardware developing community.

³ Intellectual property

FFT-engine	block	bits	LE usage	f_{max} (alone)	f_{max} (combined)
fixed streaming	1024	12	3976	250 MHz	220 MHz
fixed streaming	1024	14	4497	250 MHz	218 MHz
fixed streaming	1024	16	4945	250 MHz	211 MHz
variable streaming	1024	12	5851	219 MHz	199 MHz
floating point	1024	12+4	23000	111 MHz	–
fixed streaming	512	14	4267	250 MHz	218 MHz
fixed streaming	2048	14	5937	250 MHz	216 MHz

Table 6.1.: List of considered FFT implementations, their number of required logic elements, and the obtainable operation frequency for the FFT engine only, and for the combination with other logic, which accounts for an FPGA occupancy of 50 %. 250 MHz is the upper limit of the clock feeding tree. The chosen FFT engine is printed in bold face.

frequency of 1.44 GHz to the sampling rate of 180 MHz. By shifting the simulated signal slightly in time, the degree how good the sampling points meet the actual pulse maximum is varied. The result is scaled to an average pulse height of (300 ± 16) ADC units, when the background noise is added, the variation is doubled to a standard deviation of 30.

Figure 6.4 shows that the reconstruction error can be reduced strongly, when a FFT engine with a higher bit-width is used, introducing a reserve of dynamic range. From the available FFT implementation, a 14 bit fixed streaming engine was chosen, as it offers a good compromise between acceptable numerical noise and sufficient execution speed. For this engine, the reconstruction error accounts for an RMS value of 0.77 bits, which hardly increases the radio noise of 38.32 bits to 38.36 bits. The fact, that both noise contributions are added almost quadratically shows, that the reconstruction residual is mostly uncorrelated with the original signal amplitude. This means that it is randomly distributed in time, and does not predominantly affect high amplitudes, for example from pulse signals. All the values of numerical noise are already evaluated with the resource-saving method of feeding both polarizations through one FFT-engine.

In the implementation, the rescaling by the exponent is not applied directly after the Fourier transform, but kept, added with the block exponent of the subsequent inverse FFT, and finally charged up against the DFT-intrinsic scaling factor of $1/N$. The rescaling is performed in the entity calculating the absolute value of the envelope.

Block length

An other key parameter of the Fourier transform is the block length. A shorter block length N has advantages in resource consumption, as the effort of the Fast Fourier algorithm grows with $N \cdot \log(N)$ (s. app. A.1.1). A minimal block length is however necessary, as it defines the width of the bins in the discrete spectrum, and thus the frequency resolution. If the resolution is much worse than the typical carrier width, the RFI-suppression will reduce the effective detection bandwidth more than necessary, when frequency bins affected by RFI are disabled. Furthermore the block length should be much larger than the typical recorded pulse-width, which is about 100 ns, so that the spectrum is dominated by the background signal, and the additional power from the pulse stays negligible. Otherwise the shape of the spectrum would be tampered by a potential pulse contribution, a consistent carrier suppression would be hindered.




Figure 6.4.: Histogram of the round trip reconstruction residual of different FFT implementations.




Figure 6.5.: Effect of the FFT block length on the signal to noise improvement of the RFI-suppression.

To estimate the interaction of these effects, their combined impact on the performance of the RFI-suppression is evaluated with the test data set used before. To exclude the difference of improvement between both antenna polarizations due to their different carrier content, only background noise from the east-west polarization was used. The improvement of the signal to noise ratio (SNR) of the pulses relative to the original signal is shown in figure 6.5. It is clearly visible that a certain block length is necessary to achieve full efficiency, above 1024 points however, the improvement saturates. Because of disadvantages towards larger block lengths in matters of required resources and offered performance, as shown in table 6.1, a block length of 1024 values was chosen.

This corresponds to the earlier study with the prototype system [95], where an optimal block length of 512 samples was found, which translates to a similar frequency resolution, taking into account the lower sampling rate of 80 MHz.

6.3. Effect of the window function

When the continuous data stream is split into blocks as necessary for the Fast Fourier transform, the block edges need special attention to avoid three harmful effects:

Spectral leakage

When the block creation is thought as a multiplication of the endless input signal with a rectangular window function, the resulting frequency spectrum is the convolution of the original spectrum with the Fourier transform of the rectangular window, the *sinc*-function. The Fourier transform of a monofrequent signal, which is a delta function, is widened by this convolution. The signal power, which was concentrated at one frequency, leaks into the full measuring band, mainly into the neighboring frequencies, hindering the removal of the carrier's entire signal power. This so-called spectral leakage can be reduced by applying a different window function with a Fourier transform, which drops off faster than *sinc*.

An alternative leakage description starts from the Discrete Fourier transform's precondition of a periodic input signal, which implies that the examined data block is repeated periodically. The additional frequencies are introduced by discontinuities at the block edges. A window function, which provides a smoother signal drop towards the block edges, reduces these discontinuities, and thus the leakage. A large variety of




Figure 6.6.: Illustration of selected window functions. We chose a Tukey window with $n = 71$, $N = 1024$ (red line).




Figure 6.7.: Detail of the frequency spectrum of the east-west polarization at AERA. The Tukey window significantly reduces signal leakage.

functions for this purpose exists in literature. To keep the algorithm simple, the Tukey window was chosen (s. eqn. 6.1) [98]. The Tukey window, also known as tapered cosine window, consists of a middle part of unity gain, framed with a cosine-shaped fade-in and fade-out (s. fig. 6.6). The number n of samples, which are affected by the fading is variable, and defines the efficiency of leakage removal. For $n = N/2$ the Tukey window becomes a von-Hann-window:

$$w_i = \begin{cases} \frac{1}{2}(1 - \cos \frac{i \cdot \pi}{n}) & i < n \\ \frac{1}{2}(1 - \cos \frac{(N-i) \cdot \pi}{n}) & i > N - n \\ 1 & \text{else} \end{cases} \quad (6.1)$$

The effect of the spectral leakage and its reduction with the Tukey window is illustrated in figure 6.7 with the 67 MHz carrier measured at the AERA site. Without window function, the detected peak amplitude is decreased by about 20 %, the missing power is distributed over a wide frequency range, and could not be removed by RFI-suppression.

Signal bulging

The remaining power from removed carriers, which leaked into other frequencies is not distributed uniformly in time, but concentrated close to the block edges, where the frequency is not as definite, due to the signal discontinuity. Therefore the background level close to the edges remains increased, even after removal of the elevated frequency bins by the RFI-suppression. The remaining spurious signal bulges (s. fig. 6.8) close to the block edges become smaller, when leakage is reduced with a window function, but they can only be removed completely when the affected areas are thrown away and replaced with unaffected data from an overlapping block.

By using overlapping blocks the implementation of the window function is also simplified: Instead of correcting for the effect of the window function after RFI suppression, which would require a spendthrift division, only the middle part of each block is kept, where the Tukey window is unity. Additional numerical noise due to the window reconstruction is avoided, the windowing only requires a simple multiplication with a set of pre-calculated coefficients, stored in FPGA-internal memory.




Figure 6.8.: Example for the signal bulges with and without overlap application at the block edges (dashed lines).




Figure 6.9.: Reduction of the average noise level with increased windowing, with and without overlapping.

The length of the overlap is a compromise between optimal leakage removal (s. fig. 6.9) and the increased processing effort. With a chosen overlap of $n = 71$, which translates to 14 % redundant data, the effects of spectral leakage are almost completely removed.

Pulse reconstruction

Because of the considerable width of the wanted radio pulses, a pulse can fall directly between two blocks. In this case, half of the pulse would be processed in each block, a proper pulse reconstruction would be questionable, the RFI-suppression would cause effective trigger dead-time. This is particularly important for rather distant cosmic rays with a low radio pulse close to the detection limit, which may be up to 100 ns wide (s. fig. 2.5). Additionally, the radio pulse is flattened by the dispersion of the analog filter, the original pulse height can be reconstructed with a deconvolution in frequency domain as mentioned above. For this purpose, the full filter response, which lasts for about 100 ns [97], must be contained in the FFT block. With an overlap of $n = 71$, which translates to 400 ns, both requirements are safely covered.

The alignment of the overlapping blocks is illustrated in figure 6.10. Because the payload of each N -point FFT block is reduced by $2n$ redundant values, the amount of data to handle is raised to $N/(N - 2 \cdot n) = 1.16$. To process this data, a working frequency of 210 MHz needed, created from the sampling frequency of 180 MHz with an FPGA internal PLL, set to a ratio of 7/6.

The data transition between both clock domains and the overlap is created by means of a two port memory with 147 words (s. fig. 6.11). The write port is continuously fed with the input data stream, while the write address loops over the memory range with the sampling frequency of 180 MHz, just like a ring buffer. The read process in contrast, follows a more complicated flow control. Each sixth time the write address reaches zero, thus after 882 input samples⁴, the read address is also resetted to zero, and the start of a new FFT-block is signaled. Although the first memory position is written at the same time as it is read, the read process delivers the old entry due to the internal latency of the memory controller. The memory readout is continued with a rate of 210 MHz, each time increasing the read address. Because the read process increases its address faster than the write address, it keeps reading the old memory entries. After reading 1024 values, an end-of-block signal is issued and the read process

⁴ 882 is the number of unique samples in a 1024-entry block with a $2 \cdot 71$ overlap.




Figure 6.10.: Alignment of the overlapping blocks. The blue border area of each block is discarded, only the yellow parts are kept and concatenated.




Figure 6.11.: Scheme of the overlap generation and merging engines. The used dual-port memory supports synchronous access with separate clocks.

is stopped. At this time the read address is 141 in its 7th round, almost overtaking the write address, which is close to finishing its 6th round. As soon as this happens, the next block is started, the read process is reset, and the memory entries 0 – 141 are read a second time, thus creating the overlap. Because the rational clock ratio of 7/6 leads to a slightly higher frequency than necessary for a 71-point overlap, the RFI-suppression processes the block of 1024 entries faster than 882 values are digitized by the ADCs. Therefore the RFI-suppression engine is stopped for 5 cycles after each block to keep synchronization, without losing any data.

After overlap generation, the window function is applied, followed by the rest of the RFI-suppression engine. The final transition back to a continuous data stream and into the regular clock domain is conducted with a second dual-port memory. Each time a new block arrives, the write address is reset. As soon as it reaches position 72 for the first time, the read address is synchronized to 71, where the first sample is stored, which is not affected by the windowing. The read pointer runs directly behind the write address, due to the faster write frequency, the gap between read and write address increases. When the 1024-point block is finished, writing is paused, while the read process runs continuously. When the next block is written, and the read synchronization is actually due, the read pointer is already located at position 71, the data from the next block is slid in unnoticed.

The implementation of the overlap is flexible enough to adapt different overlap lengths and clock frequencies, by changing a few parameters.

6.4. Median filter implementation

To flatten the spectrum and thereby remove frequency spikes exceeding the random noise floor, a kind of averaging is necessary. Because the arithmetic average is too much influenced by the frequency outliers caused by the carriers, the more robust median value is used (s. fig. 6.12). The sliding window used to calculate the median value must be more than twice as wide as the carrier. The finally used median over 31 values, which corresponds to 5.4 MHz, is therefore a conservative choice for a carrier spanning over up to 8 frequency bins.

The calculation of the median value is basically the task of sorting an array of values and selecting the middle one. Although there are several well-known sorting algorithms with a complexity of $O(n \cdot \log(n))$, most of them are too complex to be implemented in parallel FPGA logic. A simpler, straightforward approach uses a systolic array of primitive two-port sorters (s. fig. 6.13), which resembles a fully parallel implementation of




Figure 6.12.: Amplitude spectrum as calculated in the RFI-suppression by the 14 bit fixed streaming FFT, together with the 31-value average and median values for the east-west (*left*) and north-south (*right*) polarization.

the Bubble-sort algorithm. It offers a very high processing speed that remains constant for large median widths, which is characteristic for a systolic array of independent cells. The $O(n^2)$ increase of complexity is directly mapped on the FPGA resource consumption, which soon becomes impractical (s. tab. 6.2). Even with topological optimizations slightly reducing occupancy [99]. The main problem of this approach is that for each new point, the updated array is sorted from scratch again, instead of taking advantage of the already sorted data. by just inserting the new value.

The Serial Rank Computer is a more efficient method to sort a consecutive stream [100]. It uses a FIFO with two attributes: the value of each entry and its rank. With each clock cycle, a new value is added to the FIFO and the oldest one drops out. The new and the old value are both compared with each FIFO entry. If they are both higher or both lower, the associated FIFO rank is kept. If the new entry is higher and the old one is lower, the associated rank is decreased by one. Otherwise, the rank is increased by one. The initial rank of the new entry is calculated by a sum over all its comparison results. The median value is selected by a multiplexer choosing the value with a rank of half the median width. The serial rank computer offers a large advantage in matters of FPGA occupancy, as its complexity only grows linearly with the median width. The speed is limited, however. In particular for a large median width, the output multiplexer and the sum of the new entry's rank are critical, as they depend on many input nodes. To reach the required performance as listed in table 6.2, these two operations had to be pipelined over 2 clock cycles.

While the spectrum is delivered as complex numbers by the FFT, and the iFFT requires complex numbers again, the median filter needs the absolute value of the amplitudes. For this purpose, the real and imaginary parts are first squared, which is unproblematic due to the dedicated multipliers in the FPGA, and added. Extracting the root is more complex. The used IP-core must be pipelined over 28 cycles to keep the speed requirements, it occupies 631 logic elements. Although the square-root engine could also be saved, as the median calculation works just as well with squared amplitudes, it is more resource-saving to actually implement it than conducting the median filter with the doubled bit-width.

To reconstruct the complex amplitude from the absolute median value, not the amplitude/phase representation is used, as the detour via the trigonometric functions would be absolutely inefficient. Instead the quotient of the median divided by the original ab-




Figure 6.13.: Implementation of a median filter with a systolic array of primitive sorters, each arrow indicates clock-synchronous transfer.




Figure 6.14.: Scheme of a serial rank computer used as median filter.

solute value is calculated as a scaling factor, and multiplied with the complex original amplitude:

$$A_{median} = \frac{|A_{median}|}{|A|} \cdot A \quad (6.2)$$

The complexity is kept low, by using an integer division. To enable fraction numbers for the scaling factor, the denominator is shifted left by 8 positions, the final result after multiplication is then shifted right again. Therefore the scaling factor has a resolution of 1/256, which introduces additional numeric noise, particularly when a frequency bin is suppressed and the scaling factor is small. Despite this simplification and maximum pipelining, the division, which consumes 612 logic elements, is speed critical, an operation frequency of 210 MHz is not safely reached. To relieve the timing, two division engines must be implemented, which run at 105 MHz and take alternating turns at calculating the division result.

Width	Algorithm	LE usage	f_{max} (combined)	Latency
9	Systolic array	822	250 MHz	5
	Optimized array [99]	689	250 MHz	5
	Serial rank computer	555	250 MHz	10
15	Systolic array	3160	250 MHz	8
	Serial rank computer	957	231 MHz	13
25	Systolic array	6819	250 MHz	13
	Serial rank computer	1627	231 MHz	18
31	Serial rank computer	1976	216 MHz	21

Table 6.2.: List of different median filter implementations, their consumption of logic elements in the FPGA, operating speed, and latency in clock cycles. The chosen median filter is printed in bold face.

6.5. Results of the RFI-suppression

To make an implementation of the RFI-suppression in FPGA logic possible, the algorithm had to be simplified, for example by using integer instead of floating point arithmetic, which introduces additional numerical noise. To assure that this noise contribution is acceptable, the difference between the output signals of the original RFI-suppression and the FPGA-implemented version is examined (s. fig. 6.15). The residual has a bias of about half a bit, which is due to rescaling of the block exponent: Because most blocks are shifted right, the least significant bits are cut off, which effectively rounds down the signal. The bias can be corrected easily by adding '0.5' before rescaling. The variation of the residual with an RMS of 2.6 bits marks the numerical noise. It increases the measured radio noise level only slightly from an RMS of 18.2 bits to 18.6 bits, which demonstrates that the numerical noise is mostly uncorrelated with the signal amplitude.

The RFI-suppression engine delays the recorded signal by a constant latency of 4794 clock cycles or $27 \mu\text{s}$. The main contributors to the delay are the FFT-entities with 2108 cycles each, and the demerging of the antenna polarizations with 1026 cycles, where their latency is reduced by $1/7$ due to their calculation in the 210 MHz domain. The delay can be easily corrected, the event data is buffered for several seconds anyway, to allow for the larger communication latency. The RFI-suppression occupies 32 % of the FPGA resources, making it the largest entity in the design. The power consumption of the entity accounts for 1640 mW which is 63 % of the dynamic power budget, according to simulation with the Altera PowerPlay Estimator. The unproportionally high power requirement can be explained with the high clock frequency on the one hand, and the above average toggle rate in the processed data on the other hand.

As discussed earlier, the effect of the RFI-suppression on the signal to noise ratio depends on the amount of monofrequent noise in the measuring band. Basically, the noise level is expected to drop by the spectral power of the carriers, the height of pulses should be reduced by the percentage of the occupied bandwidth. For the conditions in both antenna polarizations at AERA, the reduction of the background level as shown in table 6.3, slightly exceeds the expectations. At the same time, the average pulse




Figure 6.15.: Comparison of the software-calculation of the RFI-suppression and the simplified FPGA-implementation. *Left:* Exemplary radio signal, the FPGA-calculated envelope closely follows the trace calculated with floating point accuracy. Also the close-by block edge (*dashed line*) does not introduce discontinuities. *Right:* Resulting histogram of the residual.

height is also lowered more than expected by the reduction of the effective detection bandwidth.

Input signal	signal processing	Noise level	Pulse average	SNR
EW-polarization	no RFI-suppression	49.8 bit	328 bit	6.6
	with RFI-suppression	18.2 bit	298 bit	16.4
	Difference:	-63 %	-9.1 %	+149 %
	Expectation:	-60 %	-2.8 %	+143 %
NS-polarization	no RFI-suppression	34.7 bit	323 bit	9.3
	with RFI-suppression	18.4 bit	295 bit	16.0
	Difference:	-47 %	-8.7 %	+72 %
	Expectation:	-35 %	-1.0 %	+50 %
no carriers	no RFI-suppression	21.7 bit	321 bit	14.8
	with RFI-suppression	17.4 bit	299 bit	17.2
	Difference:	-20 %	-6.9 %	+16 %
	Expectation:	0 %	0 %	0 %

Table 6.3.: Results of the RFI-suppression for the test data set consisting of background data from both AERA polarizations and Gaussian noise without carriers, each time enriched with artificial signal pulses. The improvement from the Hilbert envelope is already included, the correction of the filter dispersion with the deconvolution is not applied. All ratios relate to the signal amplitude, not the power.

If the background does not contain any carriers, and the median filter should not have any effect, still both background level and signal pulses are lowered, with a small net increase of the signal to noise ratio. This effect is caused by a systematic signal down-scaling of the RFI-suppression, based on the fact that the amplitude spectrum, which is always positive, is Raleigh-distributed. For the asymmetric Raleigh-distribution the median is a biased estimator, systematically underestimating the spectral power. The estimator bias is smaller for short pulses, because the variance of their amplitude spectrum is smaller than the one of uncorrelated noise. The small remaining improvement of the signal-to-noise-ratio, makes up for the introduced numerical noise. Even without monofrequent interferences, the RFI-suppression has a beneficial effect on the radio detection, which makes it particularly suitable for varying background conditions.

6.5.1. Effect on the random trigger rate

The random trigger rate caused by the remaining noise floor can be estimated: Assuming a normal distribution for the noise signal, which is a valid approach when the carriers are carefully removed, the envelope amplitude is Rayleigh distributed. The probability for the amplitude a is then

$$p(a) = \frac{a}{\sigma^2} \cdot e^{-\frac{a^2}{2\sigma^2}} \quad (6.3)$$

where σ is the RMS of the background noise⁵. With the sampling rate f_S , the rate of random samples above a threshold A is then determined by the cumulative distribution function:

$$f(A) = f_S \cdot \int_A^\infty p(a) da = f_S \cdot e^{-\frac{A^2}{2\sigma^2}} \quad (6.4)$$

⁵ Here the RMS value of the high-frequency signal is used; the RMS of the envelope which is used above is by an exact factor of $\sqrt{2}$ higher.




Figure 6.16.: Comparison of the simulated random trigger rate with the expectation of the east-west polarization. Because the test data set is limited to 39 ms, low trigger rates are not covered by the simulation.

When using the test data set without artificial pulses to simulate the random trigger rate, the result matches the expectations, as long as the RFI-suppression is applied (s. fig. 6.16). Without RFI-suppression, however, the theory significantly overestimates the random trigger rate for the given RMS value. This is due to the fact that the carrier signal voids the initial assumption of Gaussian distributed amplitudes. Instead, the power of the carrier must be added to the random noise power as a constant contribution. The resulting random trigger rate is lower than expected, because the additional carrier signal has no amplitude outliers.

Consequently, the threshold improvement due to the RFI-suppression is lower than suggested by the noise level reduction in table 6.3, but still significant: In the east-west polarization, the threshold necessary for the same random trigger rate is reduced by 46 %, while the noise RMS diminishes by 63 %. In the north-south polarization the threshold improvement is 39 %, when the RMS decreases by 47 %.

For a threshold setting between five and six times the noise RMS, the resulting trigger rate is acceptable. Because the trigger rate rises violently already for a slightly lower threshold, this setting is a soft lower limit for the threshold. At higher thresholds, the trigger rate from the continuous noise floor is negligible, and transient noise becomes dominant (s. ch. 7).

6.5.2. Results of the envelope calculation

The final step of the signal processing is the calculation of the signal envelope. This is necessary to get a unipolar pulse out of the high frequency signal oscillating with the analog filter's center frequency⁶. Although the pulse can also be detected directly based on the high frequency signal, which was proven by the competing MAXIMA setup [78], the timing precision of this approach is limited to half of the oscillation period. For the envelope signal in contrast, the time resolution is determined by the sampling frequency, with special attention to the rising edge, it can be lowered to a fraction of this (s. sec. 7.4). An additional benefit is that the pulse height of the envelope can be larger than the maximum value of the oscillation, if the peak happens to fall close to a zero-crossing.

⁶ The center frequency of a filter is the geometric mean of its cutoff frequencies (here: 49 MHz).




Figure 6.17.: Illustration of different methods to calculate the signal envelope.

The classical approach of an envelope detector is to rectify the high frequency signal and feed it through a low-pass filter. As the specifications of the filter are rather loose, the implementation is simple. A 13th order FIR-filter (s. app. A.2.1) with a cutoff frequency of 30 MHz, von-Hann windowing, and a coefficient width of 8 bits already results in an acceptable envelope (s. fig. A.2.1). Due to the symmetry of the filter coefficients, the envelope detector requires only seven multipliers and about 250 LEs, which is a minimal expense of FPGA resources.

A disadvantage of the low-pass filter is, however, that all fast signal changes are smoothened, which flattens the rising and falling edges of a fast signal pulse. The inherent inertia of the low-pass filter can be avoided with the Hilbert envelope, which is the absolute value of the analytical signal. The analytical signal, which consists of the original signal as real part and its Hilbert transform as imaginary part, can be calculated efficiently from the frequency spectrum with a single FFT-engine (s. app. A.1.4). As the Hilbert transform reconstructs a contribution at the zero-crossings of the high frequency signal, the pulse height is kept relatively constant for all peak positions. This increases the average pulse height, but more importantly reduces the random amplitude variation caused by an unfortunate pulse timing. In the test data set, the average pulse height is increased by 6 % from 280 ± 21 ADC counts for the low-pass to 298 ± 15 for the Hilbert envelope. The remaining standard deviation of the latter method is mainly caused by radio noise which interferes constructively or destructively with the pulse amplitude at random. When the time difference is fixed in the simulated data, the standard deviation of the low-pass amplitude drops to 14 ADC counts, the deviation of the Hilbert envelope stays constant at a value of 15. This means that the fraction of pulse height uncertainty, which is introduced by the random position of the sampling points, can be mostly eliminated with the Hilbert envelope.

If the signal is only available in time domain, the Hilbert transform can alternatively be calculated with an all-pass FIR-filter. For the band limited signal, already a 15th order filter with 8 bit wide coefficients is sufficient to result in a comparable precision, introducing only 0.7 bits numerical noise (s. fig. 6.17).

In this case, the most demanding part of the envelope calculation is the square root needed to get the absolute value. Taking into account its advantages, this kind of envelope generation should always be used, even when abstaining from the elaborate RFI-suppression engine.

6.6. Alternative RFI-suppression methods

Besides the RFI-suppression engine as described above, several alternative methods exist to enhance the signal pulses in a noisy background. On the one hand, there are much simpler approaches, which abstain from the elaborate block-wise Fourier transforms, and use regular digital filters in time domain. On the other hand, the existing RFI-suppression engine could be extended to reshape the signal much more aggressively. This way, the detectability of the pulses can be optimized, when the conservation of their shape is ignored. The following section will give a review of selected approaches, and compare their performance.

6.6.1. Alternatives to the median filter

One problem of the RFI-suppression is that the median filter applied on the frequency spectrum is a non-linear operation. As a result, the superposition of multiple signals is not always trivial, which becomes visible when more than one pulse occurs in a FFT-block: The time difference between two similar pulses translates to a linear phase difference in their individual spectra. In the combined spectrum, this linear phase shift leads to alternating cancellation and superposition of both contributions, resulting in a kind of 'inverse beat frequency' (s. fig. 6.18). Because the double pulse spectrum is not flat anymore, the median filter distorts the wanted signal, leading to spurious pulses, when transforming back to time domain. Because multiple pulses within the short block length of $5.7\text{ }\mu\text{s}$, are a strong indicator for noise transients, these events are rejected (s. sec. 7.2). Their signal distortion by the RFI-suppression is therefore only a minor issue. Nevertheless an alternative to the median filter was examined to get rid of the spurious contributions:

Instead of replacing the full amplitude spectrum with its median, the original spectrum is kept, and only frequencies which are elevated by carriers are modified. The frequencies in question are identified when they exceed the corresponding median value by a factor of three. On the one hand, the amplitude of the identified frequencies should be changed to the median value, which represents the expected amplitude without interference. This way, the affected frequency bins still contribute to the signal. On the other hand, the phase of the corresponding frequency bins is tampered by the dominant




Figure 6.18.: Simulation of the behavior of the RFI-suppression towards multiple pulses without background noise. The frequency spectrum (*left*) of a single pulse is broad and smooth, a second pulse introduces frequency ripple. The median filter removes the ripple, which leads to spurious pulses in time domain (*right*).

carrier, and only the correct phase leads to a constructive pulse contribution. Therefore the elevated bins can be cut off completely.

The comparison of both methods (s. tab. 6.4) shows very similar results, with a tiny advantage for the complete removal of tainted frequencies. When compared with the original median filter, both methods result in a significantly worse signal to noise ratio. One reason for the reduced performance are problems to identify the full width of carriers. More precisely, the threshold factor of three above the median is only a compromise between complete carrier removal and the misidentification of random amplitude variations. Moreover, the biased flattening of the spectrum by the median filter, which leads to a beneficial reduction of the random noise (s. sec. 6.5), is omitted. This becomes obvious for the data set without carrier noise, where the cutoff method has no effect at all.

The cutoff method is very robust towards multiple pulses. Even with three and more pulses per block, the cutoff threshold is hardly reached, because the multiple beat frequencies are uncorrelated as long as the time difference between the pulses is not exactly the same. In the current LOPES analysis pipeline the cutoff-to-median method is implemented.

Input signal	filter method	Noise level	Pulse average	SNR
EW-polarization	no RFI-suppression	49.8 bit	328 bit	6.6
	median filter	18.2 bit	298 bit	16.4
	carrier cutoff to median	22.6 bit	317 bit	14.0
	carrier cutoff to zero	22.5 bit	316 bit	14.1
NS-polarization	no RFI-suppression	34.7 bit	323 bit	9.3
	median filter	18.4 bit	295 bit	16.0
	carrier cutoff to median	22.5 bit	317 bit	14.0
	carrier cutoff to zero	22.1 bit	316 bit	14.3
no carriers	no RFI-suppression	21.7 bit	321 bit	14.8
	median filter	17.4 bit	299 bit	17.2
	carrier cutoff to median	21.7 bit	321 bit	14.8
	carrier cutoff to zero	21.7 bit	321 bit	14.8

Table 6.4.: Comparison of the effect of different filter methods on the noise level, the average pulse amplitude, and the signal to noise ratio. All values relate to the signal amplitude and include the full FFT/iFFT engine with Hilbert envelope.

6.6.2. Matched filter

While the median filter, and even more the cutoff method, are designed to conserve the signal shape, other filter methods change the original signal rather aggressively with the only goal to optimize the pulse detectability. The matched filter, for example, which is commonly used for pulse detection in radar systems [101], is designed to achieve a maximum peak amplitude from a noisy signal: If the spectra of the expected pulse $A_{pulse}(f)$ and the average background $A_{noise}(f)$ are known⁷, the optimal signal to noise ratio is achieved with a filter response of $H(f) = c \cdot \frac{A_{pulse}(f)}{A_{noise}(f)}$, where c is a normalization constant [102]. If the filter is applied to the expected pulse signal, the multiplication of

⁷ The pulse spectrum is complex, because its phase is generally known. The noise spectrum is real, as it derives from an average power spectrum.

Input signal	used $A_{pulse}(f)$	Noise level	Pulse average	SNR
EW-polarization	no filter	49.8 bit	328 bit	6.6
	flat (=1)	18.2 bit	273 bit	15.0
	linear	19.0 bit	306 bit	16.1
	from pulse	27.8 bit	524 bit	18.9
NS-polarization	no filter	34.7 bit	323 bit	9.3
	flat (=1)	18.1 bit	269 bit	14.8
	linear	18.9 bit	301 bit	15.9
	from pulse	27.8 bit	515 bit	18.5
no carriers	no filter	21.7 bit	321 bit	14.8
	flat (=1)	17.5 bit	271 bit	15.5
	linear	18.2 bit	303 bit	16.6
	from pulse	26.6 bit	515 bit	19.4

Table 6.5.: Performance of the matched filter on the test data set, when provided with different expected pulse spectra. All values relate to the signal amplitude and include the full FFT/iFFT engine with Hilbert envelope.

$A_{pulse}(f)$ with its conjugate-complex corresponds (for a real signal) to the autocorrelation of the signal, which, in the ideal case, compresses the energy of a wide pulse to a δ -function.

In contrast to a radar system, the spectrum of a cosmic ray induced pulse is not exactly known. It depends on various shower parameters like lateral distance and incoming direction. As a result, $A_{pulse}(f)$ must be either set to 1, which translates to no ex-ante knowledge at all, or replaced with a general spectral trend according to simulations. In both cases no pulse phase information is used, which reduces the matched filter to weighing every frequency contribution according to its expected signal to noise ratio. Frequency bins with strong interferences are scaled down accordingly. For the noise term in the denominator, the average spectrum from the permanent measurement is used, so it follows background changes dynamically. The matched filter is linear, because the filter coefficients are constant for each FFT-block.

For the simulation of the matched filter with the test data set, three implementation variants are compared (s. tab. 6.5):

In the first case no pulse information is provided, the expected pulse spectrum $A_{pulse}(f)$ is constant. As a result, the input signal is filtered to a completely flat spectrum. This is comparable to a radical variant of the median filter, where the spectrum is scaled to its median, and at least the general shape of the original spectrum is kept. The performance of this approach is comparable to the median cutoff method.

The second implementation estimates the pulse spectrum with a linear drop from full amplitude at 30 MHz to the half value at 80 MHz, the phase is constant. This corresponds roughly to the pulse spectrum at distances between 100 m and 200 m (s. fig. 2.5). The resulting signal to noise ratio is comparable with the median filter.

The third approach profits from the special case of the simulated test data set, where the pulse characteristics are exactly known. It employs the complex pulse spectrum, and sets up a real autocorrelation filter. Obviously this case is not

realistic for the implementation in the field, but it demonstrates the limit of possible signal improvements.

For all three variants, the noise spectrum $A_{noise}(f)$ is taken from the background data before the artificial pulses are added. This is necessary, because the pulses would have a significant impact on the noise spectrum due to their high rate. In contrast, the field implementation would calculate the noise spectrum from the regular measured data, as the pulse rate is low enough to be neglected.

Although the matched filter could potentially outperform the median filter when the applied pulse spectrum is tuned, it is currently deferred because of its signal shape deformation. It might become interesting again when the radio signal is understood better. In any case, the chosen RFI-suppression and the matched filter are very similar, most of the used components like windowing and Fourier transforms are the same.

6.6.3. Notch filter

Instead of removing the radio frequency interference in frequency domain with an elaborate mechanism of Fourier transforms, the application of digital filters in time domain suggests itself. To suppress narrow band contributions, a notch filter is required, which is best implemented as an Infinite Impulse Response (IIR) filter (s. app. A.2.2).

A FIR approach is generally not suitable, because the needed filter order would be reciprocal to its bandwidth. To suppress the triple carrier in the north-south polarization without losing too much bandwidth, a frequency resolution of at least 1 MHz is necessary. This requires a FIR-filter of 180th order, which would occupy 90 multipliers for a symmetric filter design. At this point, a FFT-based engine is more efficient, the proposed RFI-suppression engine needs only 25 multiplications⁸.

A notch filter has unity gain except close to its center frequency, where the theoretical trap attenuation is infinite. The Q-factor defines the steepness of the fall-off (s. fig: 6.19). Consequently, the trap attenuation is not an issue in any realistic case of application, although it is reduced to a finite value by the numerical resolution of the filter coefficients. A second order IIR-filter is sufficient to efficiently suppress even strong carriers. In contrast, a high Q-factor is crucial to confine the effect of the filter to a frequency range as small as possible. This is particularly important if more than one notch filter is applied to suppressed multiple carriers. For high Q-factors the numerical stability of an IIR-Filter becomes critical, with a 14 bit coefficient resolution Q-factors up to 100 are achievable.

To compare the performance of the notch filters with the other methods, the known test data set is used. For the east-west polarization a single IIR-filter is used, the north-south polarization with its three carriers requires the concatenation of three filters. To include the same pulse height increase like the other methods, a Hilbert FIR-filter is applied to calculate the envelope (s. sec. 6.5.2). Figure 6.20 shows that the notch filters of low Q-factors reduce the effective bandwidth too much, resulting in low signal to noise ratios. Around $Q=60$ the signal to noise ratio saturates in both polarizations, the achieved value of about 14 is comparable with the performance of the median-cutoff method.

The biggest advantage of notch filters is their low consumption of FPGA resources. A second order IIR-filter can be implemented with three multiplications, which translates to six FPGA multipliers, taking into account the required bit-width. Unfortunately, the effort rises with every additional carrier to suppress. Additional signal processing

⁸ Both numbers correspond to 14 bit multiplications, which need 2 hardware-multipliers each.




Figure 6.19.: Behavior of IIR-filters in 14 bit integer implementation. The introduced phase is zero except for a pole at the suppressed frequency.




Figure 6.20.: Achieved signal to noise ratio of the RFI-suppression based on IIR-filters with different Q-factors.

like a Hilbert envelope or the deconvolution of the analog filter dispersion must be implemented with separate filter structures. Furthermore, the filter setup must be manually adjusted to the expected background situation, variable background is not handled efficiently.

Notch filters are currently used in the competing Dutch electronics to suppress the narrow band noise at AERA.

6.7. Conclusion

The RFI-suppression method known from offline radio event analysis could be successfully adopted for the real-time self-trigger. In doing so, the algorithm had to be modified to cope with a continuous data stream instead of confined events. Only with several numerical tricks and simplifications the algorithm could be fitted into the chosen FPGA-chip. By comparing the hardware-implemented RFI-suppression engine with the original software, it was proven that the amount of numerical noise introduced by the simplifications is negligible.

Compared with regular digital filters, the described RFI-suppression engine shows a superior performance. However, the performance of the matched filter is on par with the RFI-suppression, it might even be improved in the future with a more detailed knowledge of the signal pulse shape. The matched filter is however closely related to the suggested RFI-suppression engine, as they share most of their components.

In the changing background conditions at the AERA site, the autonomous adaption of the filter to background changes is an important feature of the presented algorithm. Additionally, the established frame of FFT and iFFT provides a universal access to the measured signal in frequency domain. This opens a door to various signal manipulations besides the removal of carriers, like the used correction of the analog filter dispersion.

7. Pulse detection

The first step to detect a pulse is to apply a threshold. If the signal exceeds the threshold, a first level trigger (T1) occurs. The T1-events are dominated by noise transients on the one hand and random triggers from the continuous background on the other hand. In the previous chapter an efficient method to reduce the continuous noise was discussed, as the continuous noise floor puts a lower limit on the threshold level. At this limit, the random trigger rate can be modulated with small variations of the applied threshold.

The transient noise, in contrast, is not affected by the RFI-suppression. Its rate is essentially specified from outside, and is hardly changed by variations of the threshold. To decrease the T1-rate below the external transient rate, the threshold must be increased massively (s. fig. 7.1). The radio detector becomes virtually blind for cosmic rays.

Instead, it is more efficient to run with a variable T1-rate and reject the false triggers with the help of other criteria than the threshold to a second level trigger (T2). The T2-decision, which will be described below, is rendered solely on the basis of local station data, without information from neighboring stations. It is implemented in FPGA-logic to avoid performance bottlenecks in the processor during event bursts. The resulting T2-rate must be low enough to comply with the capacity of the communication and data acquisition system. During the first phase of AERA, the communication system is not a limiting factor, as it is based on a fast glass fiber network, but the proposed wireless system of the subsequent phases will only allow a T2-rate of 200 Hz [93].

After transmitting the timestamps of the T2-events to the central data acquisition system, the coincidence condition between multiple stations forms the third level trigger (T3). The T3-triggered events are then read out and stored to a data base. The future wireless system will provide capacity for an event readout rate of 0.1 Hz per antenna station. Taking into account that the expected cosmic ray rate is below one event per day per station, this still provides reserves for many false triggers.

7.1. First measurements at AERA

When the installation of AERA started in October 2010, the new KIT/BUW digitizer was still in development, so 20 antenna stations were at first equipped with the competing Dutch system. In April 2011, the regular self-triggered data taking started, during the first month 24 identified cosmic ray events were recorded, each event containing data between two and eight stations. Then the detection of cosmic rays suddenly dropped out, in the next four months only 11 additional events could be registered. The reason for this irregular event distribution is not completely understood yet, due to incomplete detector monitoring. Most probably the background noise was suddenly getting worse.

Background measurements

In July 2011, the southernmost three stations of AERA, which are marked as red antenna symbols in figure 5.1, were replaced with the KIT/BUW digitizer electronics¹. Because the central data acquisition system still needed modifications in order to work with the new local stations, the measurement had to be simplified: Instead of forming a coincidence trigger, and reading out the ADC traces, each station only stores its T2-timestamps locally. This way, the detection of a cosmic ray can be registered by comparing the stored timestamps offline with the surface detector data. When all three antennas trigger, even a simple directional reconstruction based on the T2 time differences is possible. For a more elaborate analysis, however, the complete analog signal is missing.

During this period no second level trigger cuts are applied, so T2 equals T1, the trigger rate is kept constant by regulating the threshold level. Consequently, the trigger threshold, which is monitored each 10s, is highly variable. When the transient noise rises, the threshold reaches values where an efficient detection of cosmic rays is not possible anymore (s. fig. 7.1). These intervals of high threshold account for one third of the measuring time. In the remaining time, when the noise level is low, the threshold shows a small daily modulation. A closer inspection of this variation reveals, that the maxima coincide with the daily culmination of the galaxy. The modulation period, observed over three months, corresponds to a sidereal day. This means that the sensitivity of the system is high enough to directly visualize the galactic noise. A distinct analysis as with the prototype system (s. sec. 4.2), is no longer necessary.

This early data taking started in August 2011 and continued for three months. During this time, the measurement was however far from a regular operation, frequent interruptions for test and modification of the data acquisition and trigger system reduce the




Figure 7.1.: Variation of the background conditions at AERA in August 2011: The threshold must be strongly regulated over time to keep a constant T1 trigger rate of 200 Hz. During periods of low noise a modulation from the galactic noise is visible. The electric field threshold on the right vertical axis is an approximate estimation as described in section 4.2, and is based on the known analog signal chain.

¹ The stations are identified as number 20, 21, and 22 in the field.

operating time to 64 days. When only taking into account periods with a reasonable threshold below 150 ADC units, the effective measuring time is decreased to 42 days.

Cosmic ray detection

The data from the surface detector which is used to identify cosmic rays has to be narrowed in time and location to get a reasonable event selection. For this purpose, only air showers are considered which hit the close-by SD-tank 'Kathy Turner' and occurred when the radio measurement was running with a threshold below 150. Because each antenna station is treated separately, one physical air shower translates into three potential station events. If a radio T2-timestamp coincides with one of the selected SD-events within a small time window, the T2 is most probably caused by an air shower. With the known time offsets of earlier measurements (s. fig 4.10), a coincidence window of $\pm 1 \mu\text{s}$ seems suitable. This window results in seven coincident events as listed in appendix A.4. Each of the coincidences corresponds to a different cosmic ray event, which is detected in a single antenna station. Because no event was detected in multiple antennas, a further direction reconstruction is not possible.

To estimate how many random coincidences are expected, 1000 neighboring windows are examined, 500 windows both before and after the actual event time. The corresponding histogram in figure 7.2 shows that the considered event window hardly stands out against the random coincidences. When the number of coincidences is plotted in an appropriate histogram (s. fig. 7.3), the random coincidences should follow a Poisson-distribution, as for any random counting statistics. In this case, however, multiple events are more probable, as the antenna stations do not trigger independently. A single strong transient pulse, for example, can be detected by all of them. Consequently, the single, double, and triple events must be identified, and each treated with its own Poisson statistic. Their accumulated probability describes the observed data much better than a single Poisson distribution. Although the occurrence of one window with seven random entries is not very improbable in a collection of 1000 windows, the fact that the event window is selected beforehand, and not based on the number of entries, makes it significant already. The probability that one window contains seven or more entries by chance, is calculated with the cumulative distribution function to be $3.6 \cdot 10^{-4}$.




Figure 7.2.: Histogram of the coincidences between SD-events and T2-timestamps for 1000 neighboring coincidence windows of $\pm 1 \mu\text{s}$.




Figure 7.3.: Statistics of the number of random coincidences in comparison with the real coincidence rate in the event window.

Event selection

The significance can be increased by reducing the SD-event selection to air showers which are supposed to deliver a large radio pulse at the corresponding antenna station. For this purpose, the expected radio signal S is estimated similar to the parameterized lateral distribution (s. eq. 4.1) as a function of the primary energy E_p , the distance to the shower axis r , the zenith θ , and the geomagnetic angle α :

$$S = \frac{E_p}{10^{17} \text{eV}} \cdot e^{\frac{r}{150 \text{ m}}} \cdot (1 - \cos(\alpha)) \cdot \cos(\theta) \quad (7.1)$$

The absolute scale of the radio pulse estimator S is unknown. Consequently, the required value when an air shower is supposed to be detectable in each antenna remains a free parameter of the wanted selection quality. With a requirement of $S > 0.3$, the SD-event list is reduced from 1040 to 140 candidates, while all seven coincident events are kept. As a result, the triggered cosmic ray events stand out much better against the floor of random coincidences (s. fig. 7.4), the probability that the event window contains only random coincidences drops to $1.3 \cdot 10^{-8}$ (s. fig. 7.5).




Figure 7.4.: Histogram of the coincidence with selected SD-events for 1000 coincidence windows of $\pm 1 \mu\text{s}$.




Figure 7.5.: Statistics of the number of random coincidences, when the SD-events are preselected.

Analysis of triggered events

Besides the T2-timestamps, the low-level measurement also records height and width of the triggered pulses. While the pulse width is fairly uncorrelated with an average of 18 ± 4 clock cycles, the pulse height corresponds nicely to the signal expectation S (s. fig. 7.6). In principle, the pulse height should be proportional to S , but an offset seems to be added by background noise and maybe a trigger bias towards high amplitudes. A further analysis of the pulse amplitudes is very vague: To translate the voltage level into a local field strength, the frequency spectrum must be estimated according to known pulses. With an approximate $1.1 \mu\text{V/m}$ per LSB, a scale factor of $\epsilon_0 \approx 140 \mu\text{V/m}$ is suggested to translate the pulse estimator S into a local field strength. This scale factor is at least in the same order of magnitude as known parameterizations (s. eq. 4.1), which makes the seven detected events further plausible.

Although demonstrating the detectability of air showers with this temporary setup, the detection efficiency of only 7 out of 140 events is not satisfying. When considering events with a higher supposed signal S , the fraction of detectable events rises (s. fig. 7.7), but even with a very strong signal, the efficiency stagnates around 50 %. Interestingly,




Figure 7.6.: Measured pulse height versus the signal expectation from the air shower event.




Figure 7.7.: Efficiency to detect an air shower with an expected radio signal S in a single antenna. Because of the small statistics, a Bayesian calculation is used.

inclined air showers with a large zenith angle are rejected with higher probability, although the zenith angle is already accounted for in the calculation of S . When only regarding vertical showers with a zenith angle below 35 %, the detection efficiency is increased significantly. The reason, why a large fraction of air showers is not detected although all preconditions are fulfilled, remains unclear: A high threshold level can be excluded by the threshold monitoring. Higher order pulse cuts, which could falsely reject an event, are not applied, and dead-time is not an issue with the low-level trigger. A similar measurement with the Dutch electronics, using twelve times the detector area, showed the same saturation around 50 % detection efficiency [103].

Following measurements

In November 2011, three additional stations were equipped with the KIT/BUW digitizer and the low-level measurement was deferred to concentrate the efforts on the development of the actual data acquisition system. The DAQ system fulfills the tasks of collecting the T2-timestamps, creating the coincidence trigger, and reading the event traces. Additionally, the new DAQ provides the possibility to be externally triggered, when a promising air shower is detected by the surface detector. This way, the issue of reduced detection efficiency can be investigated, and the radio measurement becomes less dependent on the troublesome self-trigger in general. The biggest obstacle towards an external trigger mechanism is the latency of the large SD array. Normally the SD-trigger has a delay of 7.3 s, but for more than 99 % of the events it can be accelerated to 4.3 s, by canceling wait periods for communication retransmissions [104]. In order to buffer radio signal during this delay, the memory of the digitizer must be doubled to 4 GB, providing 7.4 s buffer time. Further changes of the local station are not necessary. The external trigger in parallel to the self-triggered operation was started end of January 2012. During 37 days of measurement, 35 events were recorded in the six antennas, which have a reconstructed incoming direction in agreement with the surface detector reconstruction. 14 of these events were also detected by the self-trigger in at least one antenna station. Seven are even triggered in at least three stations, so that they would have been read out by the coincidence trigger system. In March 2012 the number of antenna stations equipped with the KIT/BUW digitizer was increased to 17. Since then the statistics of comic ray data improves every day.

7.2. Pulse parameter cuts

To achieve a proper detection efficiency, the trigger threshold must be as low as possible, values with an electric field strength below $100 \mu\text{V}/\text{m}$ are considered acceptable. At the same time the station trigger rate is limited, mainly by the communication system. The long term target is an average trigger rate of less than 200 Hz [93]. To reach these contradicting aims, a higher order selection of pulse candidates is necessary on station level. In his thesis [52] T. Asch proposed a number of cuts on pulse parameters, based on externally triggered measurements with the LOPES^{STAR} experiment:

Firstly, each recorded event frame must contain a single dominant pulse with a signal to noise ratio of 75 in power. If more than one pulse of that height is found within the $25.6 \mu\text{s}$ trace, the event is rejected. Secondly, the pulse must be short and steep, as expected from cosmic ray induced pulses. Therefore the pulse width, measured as the full width at half maximum height (FWHM), is required to be below 125 ns. To get rid of background pulses with a slow signal fall-off, the ratio is calculated between the pulse integral in the range of the width and the afterpulse integral over twice the width. If this pulse-to-post-pulse ratio is below one, the event is rejected. The combination of these cuts rejects 97.5 % of the background pulses under the local background conditions.

During the first measurements at the Pierre Auger Observatory, in contrast, the background rejection of the cuts was found to be far less efficient than proposed, as the transient noise seemed to be different. Additionally, it was not clear whether the shape of the wanted pulses changed as well due to the modified measuring chain or the different altitude and geomagnetic field, which could reduce the cosmic ray acceptance. Because during the prototype measurements the proof of principle of the self-triggered radio detection was more important than an efficient trigger operation, the cuts were not applied so far. Now that the first cosmic rays are detected, and the further detection is guaranteed by the external trigger, the focus shifts back to a reduction of the false trigger rate.

For this purpose, the calculation of the proposed pulse parameters is implemented in the trigger firmware, the characteristics of background pulses are recorded during random measurements at AERA. To get the same pulse parameters for cosmic ray events, the ADC traces of identified radio events were fed through a simulation of the FPGA trigger. The used data set contains 22 events recorded by the KIT/BUW digitizer with an external trigger, and 122 events from the competing Dutch setup. Because the Dutch data was digitized with a sampling rate of 200 MHz at a different gain, it had to be re-sampled by 9/10 and scaled in amplitude to become comparable. The analog system of both setups is the same fortunately.

Signal quality cuts

The first category of proposed cuts aims to reject noisy pulses, and select signals with a high data quality.

Transient noise often occurs in bursts of many pulses during a short period, so the rejection of multiple pulses suggests itself. The simplest approach for this purpose is to count the **time between consecutive threshold crossings**. If the time is too small, the pulse is rejected. Normally this method would not reject the first transient of a burst, but when the gap between pulses is smaller than the latency of the trigger decision, even the first pulse can be revoked before being submitted to the next trigger stage. The distribution of the time difference between pulses in figure 7.8 contains actually an increased amount of double pulses within less than $6 \mu\text{s}$. The fraction of




Figure 7.8.: Detail of the time difference between two pulse candidates.

double pulse events is however smaller than the histogram suggests, as only a detail of the full distribution is shown. Actually 3.7 % of the background events are rejected, when a minimal gap of $\Delta t > 6 \mu\text{s}$ between pulses is required.

To estimate the rejection of cosmic ray events, the cut on Δt can be interpreted as an applied dead-time window around each T1-trigger. The combined fraction of dead-time is $\Delta t \cdot f_{T1}$, which amounts for 0.6 % for a typical background trigger rate of $f_{T1} = 1 \text{ kHz}$. When taking into account an overlap of adjacent dead-time windows, the fraction of dead-time is slightly reduced.

A second parameter leaping into the mind, when rejecting noisy events, is the **signal to noise ratio** (SNR). Usually the pulse height is used as the signal value, because it corresponds to its peak power. The noise level is averaged over a large period², so that it is not influenced by the surrounding of the pulse. In this case however, a cut on the signal to noise ratio just equals a regular amplitude threshold, which is already applied. An other idea is to use a kind of 'local' SNR, were the noise level is estimated from the direct pulse neighborhood on purpose. This way, the local SNR becomes a discriminator, whether the pulse is accompanied by a noise surge, or if the pulse occurs suddenly and independently from the noise floor.

In the FPGA implementation, the local noise level is calculated by averaging the envelope signal over 106 samples before and 150 samples after a triggered pulse, including the pulse itself. Because the envelope signal is Rayleigh-distributed, the average value is just $\sqrt{\pi/2}$ times the RMS of the underlying RF-signal. The SNR is then calculated in fixed point representation by dividing the detected pulse maximum times 64 by the noise value. Because the SNR division is not required every clock cycle, but only once per trigger, it can be implemented in a very slow, but resource conserving way in plain logic, without pipelining³.

When the measured values of the local SNR are compared for background transients and cosmic ray events (s. fig. 7.9), the background strongly tends towards small SNR values, although no clear separation is achieved. The cumulative distribution function (s. fig. 7.10) illustrates that a significant amount of cosmic ray events would be lost to reject a majority of the background. If for example a local SNR of more than 5 is required, 49 % of the background transients are rejected, but also 11 % of the cosmic ray events are lost. Both cuts on the signal quality together reject only 50 % of the background events, as they are not independent from each other.

² In offline analysis, when the amount of data is limited, a data section far from the pulse is used.

³ The footprint is reduced from 797 to 410 LEs, the division requires 26 clock cycles.




Figure 7.9.: Histogram of the local signal to noise ratio for cosmic ray events and noise pulses.




Figure 7.10.: Rejection probability versus SNR requirement. The error bars illustrate the statistical uncertainty.

Pulse shape cuts

Besides the surroundings of the triggered pulse, also the shape of the pulse itself can be used to discriminate noise transients from cosmic ray induced pulses.

One basic characteristic of a pulse is its width. Using the full pulse width at half maximum height is a common method to achieve a robust value. The width of cosmic ray induced pulses is known to be smaller than 100 ns from simulation (s. fig. 2.5), the response of the analog filter is slightly wider. At the LOPES^{STAR} experiment, most noise pulses were much wider than that, and could be removed properly. The background at AERA, in contrast, has a tendency towards very narrow pulses (s. fig. 7.11), but it does not separate from the cosmic ray events. If only pulses wider than 38 ns are kept, which corresponds to seven clock cycles and more, 56 % of the background is rejected. At the same time, however, also 17 % of the cosmic ray events are lost. The cut on the pulse width and the cut on the local SNR are rather independent from each other, their combination would reject 79 % of the background and 25 % of the cosmic ray events.

A second pulse shape cut, which was proven to be effective at the KIT, rejects transients noise with a slow signal fall-off after the maximum. For this purpose the




Figure 7.11.: Histogram of the occurring pulse widths, measured at half maximum height.




Figure 7.12.: Distribution of the ratio between the pulse integral and the area covered by the afterpulse.

integral of the pulse over its width is divided by the integral of the afterpulse signal over twice the width. With the background situation of AERA however, this pulse-to-post-pulse ratio shows no discrimination between noise transients and cosmic ray events at all (s. fig. 7.12).

In order to find other differences between the typical noise pulses and the wanted cosmic ray signal, a fraction of background pulses was examined manually. Except for a partially slower rise time of noise pulses (s. sec. 7.4), no further distinguishing characteristics were found. The main problem of the second level trigger decision is that the cosmic ray pulse itself does not contain distinctive features to be identified. The remaining difference towards the background is mostly concealed by the analog filter. Additionally, the T2-trigger must have a near perfect acceptance, as the loss of an event is hardly acceptable, given the low event rate. The rather good performance of the T2-trigger at LOPESTAR must be seen in the context that it reduces the false trigger rate from a projected 24 kHz down to 600 Hz. At AERA, in contrast, the T1-rate is usually already in the order of some 100 Hz, a further rejection is much more difficult.

7.3. Periodic background

If the background triggers were independent from each other, the time difference between them would be distributed exponentially. Instead, the underlying exponential trend in figure 7.13, is superimposed by a strong accumulation of events at multiples of 10 ms, similar to the measurement at the BLS (s. fig. 4.8). Based on the high-rate T2-timestamps, a closer examination of the periodicity is possible, in contrast to the earlier prototype measurement, where only the T3-events were accessible. When each T2-trigger is regarded as a unity impulse on the time axis, which is fed through a Fourier transform, the resulting trigger spectrum (s. fig. 7.14) shows a distinct 50 Hz periodicity with a strong 100 Hz harmonic. In contrast to the time difference between neighboring events, the periodic contribution on a large time scale is well-defined against the underlying floor of random triggers. An efficient removal seems possible.

To bring out the characteristics of the periodicity, the positions of the T2-events within a virtual 20 ms period are plotted versus time (s. fig. 7.15). This way, the time axis is folded in a zigzag pattern from the lower left to the upper right. An




Figure 7.13.: Histogram of the time difference between successive T2-triggers. Multiples of 10 ms are more probable.




Figure 7.14.: Fourier transform of the T2-timestamps from one hour with a detail of the 50 Hz contribution.

exact 50 Hz signal would appear as a vertical line in this plot, a sloped line indicates a frequency deviation. Most of the T2-timestamps describe a winding pattern around the 20 ms reference period. The frequency changes within some ten seconds. This behavior is well-known from the power grid, where the frequency variation is used for load regulation and the constance of the line frequency is only guaranteed on a large timescale.

Rejection by individual dead-time

A simple method to reject these periodic triggers is implemented by the Dutch AERA stations. Each T2-timestamp is compared with 50 previous events⁴. When the time-stamp falls close to a previous event within ± 0.5 ms, it is rejected. In principle, each event causes an individual dead-time rectangular in figure 7.16 with a variable rectangular height and a fixed width of 1 ms. With the example data, the method rejects 90 % of the T2-events, while blackening 44 % of the measuring time. The approach is however inefficient, because it also applies dead-time to random triggers, which occur outside the band of periodic transients. Additionally the total amount of applied dead-time can hardly be monitored. When the periodic noise is dominant, most of the individual dead-time windows overlap, the total amount of dead-time is acceptable. However, when the periodic noise is weaker and the triggers are distributed rather uniformly, the dead-time may inhibit almost all of the measuring time.




Figure 7.15.: Developing of the T2-timestamps (*blue dots*) within a 20 ms period. The expected pulse time of the PLL is plotted as a *black line*, when not locked, a *red line* is used.




Figure 7.16.: Periodic noise suppression by the Dutch stations: Each T2-event causes a *black* dead-time window, in which future events are ignored.

⁴ Only events are used which happened less than half a second earlier, so depending on the actual trigger rate, either the last 0.5 seconds or the last 50 events are compared [105].

PLL implementation

Instead of applying a dead-time on each individual event, it is advantageous to identify the periodic transients on the base of several events, and predict future transients. For the purpose of isolating a frequency and its phase from a noisy input signal, a phase-locked loop (PLL) is the method of choice. Usually a PLL is a dedicated circuit, consisting of an oscillator, which is controlled according to the input signal. For the low frequency of 50 Hz, however, a software PLL, which is implemented in FPGA logic, provides a better configurable alternative.

In the software PLL (s. fig. 7.17), the **oscillator** is replaced by a counter, running in an infinite loop from zero to an upper value T_0 . The upper value T_0 corresponds to the oscillation period, the momentary counter reading ϕ to the oscillation phase. The target is to control the oscillation, so that the periodic transients regularly occur at phase zero. The synchronized oscillation is then used to predict the time of future transients.

For this purpose, the **phase comparator** checks if the oscillator is too early or too late compared to each incoming trigger. To avoid disturbances from random triggers, the phase comparator is gated, using only events at a relatively small phase within $\pm \pi/2$.

The phase difference $\Delta\phi$, which is basically the counter reading at the moment of the trigger⁵, affects the oscillation through the **loop filter** in two ways: Firstly, the period is modified proportionally by $k_P \cdot \Delta\phi$, which represents a P-term. Secondly, the phase is modified by $k_D \cdot \Delta\phi$, which corresponds to a derivative D-term in respect to the period. The combined loop filter results in a PD-term, while the closed-loop transfer function of the full PLL is one order higher and describes a PI-controller. With the gain factors k_P




Figure 7.17.: Block diagram of the implemented PLL.




Figure 7.18.: T2-timestamps relative to the PLL prediction. When the PLL is not locked, the vertical axis is red.

⁵ If the oscillator is early, the phase difference is $\Delta\phi = \phi$, else it is $\Delta\phi = T_0 - \phi$

and k_D the properties of the PLL are adjusted. A setting of $k_P = 1/4096$ and $k_D = 1/16$ turns out to be a good compromise between PLL stability on the one hand, and the ability to follow fast frequency changes on the other hand. To prevent a runaway PLL-frequency, when no periodic background is found, the controlled frequency range is also limited between 49.8 Hz and 50.2 Hz.

Relative to the PLL prediction, the distribution of T2-timestamps is rather constant in time (s. fig. 7.18), with a substructure of two peaks in the 20 ms period. This detail is not an effect of the measurement, for example due to a dead-time, but is a characteristic of the background signal. The excess events can be rejected easily by **applying a veto** to a certain period around the PLL prediction. The amount of introduced dead-time is constant, it can be manually adjusted to the background conditions. For the example data, 78 % of the T2-events are rejected, when events within ± 2 ms are vetoed, corresponding to 20 % dead-time.

When the PLL is started from arbitrary initial values, it is free-floating until the gate of the phase comparator randomly crosses the band of periodic pulses (s. fig. 7.15). Then the PLL aligns quickly to periodic signal, and stays locked until the periodic background stops. Even with unfavorable initial values, the PLL is usually locked within less than a minute. To **detect a valid lock** of the PLL, the fraction of events which fall into a ± 3 ms window around the expected pulse time is counted⁶. If the fraction is significantly higher than 0.3, as expected for randomly distributed events, the PLL is certainly locked⁷.

Instead of using a PLL, two alternative methods were examined to isolate the 50 Hz periodicity from the background triggers. The first alternative feeds a unity pulse for each T2-trigger to a narrow-band IIR-filter. When the input signal contains a 50 Hz component, the filter oscillates resonantly with a phase shift of $\pi/2$. The second method divides the 20 ms period in five windows, counting the events falling into each window. Based on the number of events in each window, a linear prediction of the future trigger timestamp is possible. Although both methods deliver valid results, they do not follow frequency variations as precisely as the PLL [106].

Long term performance of the PLL

When the PLL is operated over several days, the variability of the periodic background becomes visible (s. fig. 7.19). Firstly, a strong 100 Hz harmonic appears at times, causing triggers ± 10 ms from the PLL prediction, which can be removed with an additional veto window. In doing so, the amount of introduced dead-time is however doubled. Secondly, the total fraction of triggers with a periodic pattern varies, while also the substructure and the number of peaks in the trigger distribution is changed. During some periods, the contribution of periodic trigger becomes so weak, that the PLL loses the lock, and starts to float. Fortunately the background in these periods is already low enough, that a reasonable threshold alone achieves a sufficiently low false trigger rate (s. fig. 7.20). Further cuts are unnecessary.

Because the periodic transients have a common source, they occur almost simultaneously in all antenna stations. The distance between antennas is negligible compared to the wavelength of the 50 Hz signal, so that the PLLs at all antenna stations are in phase. This way, the veto of each station is applied synchronously, the dead-time of the full

⁶ The width of the lock detection window is freely adjustable.

⁷ The probability of a false lock detection can be estimated using Poisson statistics: When the lock detection counter runs over 1024 events with a ± 3 ms window, the standard deviation is $\sigma = \sqrt{0.4 \cdot 1024} = 18$. A measured fraction of 0.4 would have a significance of 5.9σ .




Figure 7.19.: Long term measurement of the periodic background relative to the PLL prediction. When PLL is not locked, the vertical axis is red.




Figure 7.20.: Threshold applied during the long term measurement to keep a constant trigger rate of 200 Hz. When PLL is not locked, the time axis is red.

array is not increased by a random overlap of the individual station dead-times. A measurement of the PLL phase difference between the three antenna stations (s. fig. 7.21), shows a typical standard deviation of ± 0.7 ms. During a fraction of the time, the time difference is much larger and peaks around ± 10 ms, which happens when both PLLs follow different 100 Hz branches of the periodic signal. To avoid this behavior, the lock detection counter is upgraded with a second counter for the events 10 ms from the PLL prediction. When the latter counter shows significantly more entries, the PLL phase is shifted by π . If both branches are comparably strong, however, this mechanism may fail. In reality, the effect has no consequences, because when a significant 100 Hz components is present, two symmetric veto windows are applied anyhow.

To examine how much of the periodic veto rejection can be translated into a threshold reduction, a run with a fixed threshold and a high false trigger rate is conducted. By storing the height of each triggered pulse, a higher threshold can be simulated posterior. The two hour run is characterized by distinct periodic noise, without the periodic veto, an average threshold of 250 ADC-units is necessary to keep a tolerable false trigger rate of 200 Hz (s. tab. 7.1). With this threshold the periodic noise is well isolated (s. fig. 7.22), and can be rejected almost completely by applying a dead-time window from -2.5 ms to 2.5 ms, and a second one from 7.5 ms to -7.5 ms. With the reduced false trigger rate, threshold can be lowered. At a threshold of 100, the rejection efficiency of the periodic veto is still good. But when the threshold is lowered any further, an increasing number of random triggers obliterates the periodic signal, the rejection efficiency drops. Keeping an acceptable trigger rate, the threshold can be lowered to 85, almost down to the level of periods with a low background.

To avoid that the PLL is disturbed by the less distinct signal, when measuring with a low threshold, the PLL-engine is equipped with its own threshold, which can be set to a higher level. The massive improvement of the trigger level also comes at a price: The




Figure 7.21.: Comparison of the PLL phase of neighboring radio detector stations, when both PLLs are locked.




Figure 7.22.: Effect of the trigger threshold on the clarity of the periodic signal.

effective measuring time of the system is cut in half during the noisy periods. Therefore the periodic veto, which is indeed implemented in the system, is currently not switched on.

Threshold (ADC-units)	Trigger rate raw	Trigger rate after veto	Rejection efficiency
250	204 Hz	3.5 Hz	98 %
100	502 Hz	32 Hz	94 %
85	755 Hz	153 Hz	80 %

Table 7.1.: Performance of the periodic veto for varying threshold levels.

7.4. Constant fraction discriminator

Besides the reduction of the false trigger rate on station level, the precision of the T2-timestamps is important for the subsequent trigger stages. The straightforward approach to determine the pulse time is to use the moment of the maximum. At the maximum, however, the slope of the signal amplitude is zero, the exact position can be ambiguous between neighboring sample values. Already a small amount of noise can additionally shift the maximum. The pulse time is better defined at the rising edge, where the amplitude slope is large. A common method to access the leading edge is a constant fraction discriminator (CFD), which adds the inverted signal, reduced by a constant fraction, to a delayed version of the signal (s. fig: 7.23). When a suitable pulse occurs, the output is first pulled into negative values by the inverted addend, until the delayed signal overrides (s. fig: 7.24).

In principle, this resembles a high-pass FIR filter of low order, sensitive to the fast amplitude increase of a pulse. The actual trigger decision is still made with a threshold, the zero-crossing of the CFD-signal defines only the exact timestamp. The resulting timing resolution can be examined with synthetic pulses, which are superimposed by realistic noise. The time difference between the pulses is predefined, the measured value (s. fig. 7.25) shows a significantly reduced variance of the constant fraction method. In principle, the time resolution can be additionally improved beyond the limit of the sampling rate by means of linear interpolation: The ratio of the CFD-values before and




Figure 7.23.: Block diagram of the constant fraction discriminator




Figure 7.24.: Envelope of an example pulse and its CFD-signal

after the zero-crossing indicates which of the two points is closer to zero. By introducing three additional interim values with this method the initial result is however hardly improved.

The improved time resolution has a large impact on the angular resolution of the early direction reconstruction, which is conducted by the coincidence trigger. The hotspots, where many background triggers originate, are concentrated significantly tighter for the CFD-assigned timestamps in the directional sky-map (s. fig. 7.26). In particular, the accidental elevation of background triggers above horizon is reduced. Most of the noise transients originate from the north-west, the hotspot points directly towards the close-by village of El Sosneado. A specific suppression of this incoming direction seems reasonable.

An ancillary effect of the constant fraction discriminator is that pulses with a slow rise time do not set off the constant fraction discriminator, because the CFD-signal never becomes negative. This affects one third of the false triggers, mainly pulses with a low amplitude, which are embedded in a general increase of the background level. Therefore the pulse rejection of the CFD has a large overlap with the cut on the local




Figure 7.25.: Measured time difference between simulated pulses with different methods, including a simulation of the full RFI-suppression engine. The predefined time difference is 1320 clock cycles.




Figure 7.26.: Sky-map of background events based on the T2-timestamps as an azimuthal equidistant projection. The polar angle illustrates azimuth, the radius shows the elevation with zenith in the center and horizon at the rim of the coordinate system. *Left:* Measured between 6. Sept. and 7. Sept. 2011 with a T2-timestamp defined by the pulse maximum. *Right:* Recorded between 30. Sept. and 1. Oct. 2011 with the constant fraction discriminator.

signal-to-noise ratio value mentioned above (s. sec. 7.2), both cuts combined rejection 60 % background events. At the same time, the CFD-trigger also loses 9 % of the identified cosmic ray pulses, which suffer an unfortunate noise situation. Again, most of the events in question are also affected by the SNR cut, the combined event rejection is 13 %. As the advantages of the CFD-trigger prevail, it is not only implemented as a trigger option, but also used in the current data taking at AERA.

7.5. Coincidence trigger

When the T2-triggers are transmitted to the central radio station, the final third trigger stage combines them to a joint radio event. The performance of the T3-trigger can be simulated accurately with the T2-timestamps measured during the first months of low-level data taking. The results from three antenna stations can be extrapolated to the future self-triggered operation of the full array.

The first step to merge the data stream from each station is to sort the incoming T2-timestamps and select events, which coincide within a certain time window. For the used three antenna stations 20, 21, and 22, which are located inside a maximum distance of 250 m, a coincidence time of $1 \mu\text{s}$ is chosen. Similar to earlier measurements, the stations are found not to trigger independently from each other, the coincidence rate is much higher than expected by pure coincidences (s. eq. 3.1). In fact, about one third of the detected pulses of each station are also found in both neighboring stations. The fraction of coincident events can be reduced by decreasing the allowed coincidence time, which would apply an incidental zenith angle cut (s. fig. 3.8). Because of antenna arrangement in a non-equilateral triangle, the zenith angle cut would however be very irregular. For the large radio array with a variable antenna spacing, the event rejection based on a simple coincidence time is not suitable.




Figure 7.27.: Coincidence trigger rate depending on the allowed zenith angle. The resulting effective aperture angle of the detector refers to the right vertical axis.

Instead, the found triplet of timestamps is used to reconstruct the event direction. Assuming a plane radio signal front, the reconstruction leads to a simple quadratic equation, which can be solved very efficiently. If the equation has no solution due to a negative determinant, the implausible event is rejected. More than two thirds of the coincident events however, have two solutions, one with an upward moving signal, the second, downward one is the valid event direction. The direction reconstruction heavily depends on a correct assignment of the T2-timestamp. The arc-shaped artifacts in the sky-map for example, which are suppressed against the hotspots by about four orders of magnitude, originate from faulty T2-assertions in one station. For most of the background events, the source is located close to the horizon (s. fig. 7.26). In particular with the more precise CFD-assigned T2-timestamp the false trigger rate can be reduced significantly with a cut on the zenith angle (s. fig. 7.27). With an accepted zenith angle of 60° , which corresponds to a sensitive solid angle of π , the T3-rate drops below 0.1 Hz, as required for the future wireless communication system [93].

As the noise sources are not distributed evenly around the horizon, a permanent masking of all inclined events is not very efficient. Also variable noise sources, like for example airplanes flying over [74], are not rejected. Instead, a dynamic, self-learning suppression is examined, which masks regions, from which several pulses occur within a short period of time. This dynamic T3 has essentially three adjustable parameters:

The size of the **cone angle** around each reconstructed event direction, which is regarded conspicuous for subsequent events. The value corresponds to the angular resolution, which can be estimated based on the size of the hotspots in the sky-map.

The **storage time**, how long the directions of earlier events are remembered. A temporary hotspot is masked for at least this time.

The **number of events** from one region, which are required to mask a conspicuous direction. When the second event from one direction is already rejected, it means that each random trigger on the sky is followed immediately by a blinded spot. A better solution is, to wait for at least two events, before a region is declared noisy.




Figure 7.28.: Performance of the dynamic T3 with different settings, determined for the run of 6. – 7. Sept. 2011 using the CFD-timestamp. The upper/lower value for each combination of cone angle and storage time corresponds to waiting for one/two events from one direction before rejecting the following. *Left*: remaining background trigger rate. *Right*: resulting effective aperture angle of the detector.

The setting of the parameters is a compromise between a fast and reliable suppression on the one hand, and a resulting aperture angle as large as possible. For any plausible setting however, the dynamic T3 clearly outperforms a constant cut on the zenith angle (s. fig. 7.28). When masking a cone angle of 20° for 10 minutes after two events from one direction occurred, a rejection mask as shown in figure 7.29 is the result, the remaining fraction of background events are distributed rather evenly over the sky (s. fig. 7.30). The method also provides a consistent performance over the full period of available data, almost regardless of the recording method (s. tab. 7.2).

The dynamic T3 can be implemented very efficiently, using a table of incoming angles for the rejection mask. The table is updated with a circle of entries around the reconstructed angle of each event candidate. The circle diameter corresponds to the applied cone angle, the written entry to the storage time. Each minute, all table entries are decreased by one, until they reach zero again. A second table stores the number of event from each direction, which is reset, as soon as the corresponding entry in the time table reaches zero. If more than a predefined number of events from the same direction as an event candidate are registered, the candidate is rejected. The entries of the tables can be monitored each minute, to register the effective aperture of the detector. For the three antennas, the dynamic T3 algorithm could be executed in a single processor thread by a factor of 1000 faster than the underlying data taking period. This provides sufficient performance reserves for real-time data taking, even with a larger number of antennas. With this efficient algorithm, the most promising approach for the large antenna field is to split the array into triangles again. With a joint table of noise-conspicuous directions, the individual T3-trigger and readout rate of each station should stay almost constant, even for a station inside the hexagonal grid, which is contained in six surrounding triangles.

A similar method to suppress noisy regions was already proposed to select promising cosmic ray candidates from the measured radio data without using external information from the surface detector [107]. That examination was however based on the higher-level reconstruction of the radio traces of the full array, and profits from a much better angular resolution. Also the offline selection is much more restrictive than the suggested settings of the dynamic T3.




Figure 7.29.: Example of the sensitivity mask of the dynamic T3 for the run recorded between 6. and 7. Sept. 2011.




Figure 7.30.: Sky-map of the remaining background triggers after applying the dynamic T3-algorithm.

Besides the T2-timestamps from each detected pulse, it is also foreseen to transmit the pulse height for a higher order T3-cut, which would reject implausible lateral amplitude distributions. So far, the method is not fully implemented, its realization currently does not enjoy a high priority, taking into account the good background rejection based on the timestamp information alone.

	All runs	Runs with CFD T2-timestamp	Runs using Maximum T2
Joint effective uptime	39 days	8 days	31 days
T2-rate	350 Hz	340 Hz	350 Hz
Coincident events within $1\mu s$	107 Hz	64 Hz	118 Hz
Reconstructible events	73 Hz	48 Hz	79 Hz
Events with zenith $\theta < 60^\circ$	6.8 Hz	114 mHz	8.4 Hz
Effective aperture angle	π	π	π
Events passing dynamic T3 ¹	41 mHz	25 mHz	45 mHz
Effective aperture angle ¹	1.78π	1.90π	1.75π
Events passing dynamic T3 ²	6.5 mHz	5.0 mHz	6.9 mHz
Effective aperture angle ²	1.42π	1.63π	1.36π

Table 7.2.: Coincidence trigger performance averaged over all available T2-data.

¹ Mild setting with rejection cone angle 8° , event storage time 5 min.

² Regular setting with rejection cone angle 20° , event storage time 10 min.

7.6. Conclusion

During the first measurements with the KIT/BUW electronics at three antenna stations of the AERA setup, the trigger characteristics were examined in detail. The background noise situation was found to be highly variable: During two thirds of the time, the noise level is low enough to conduct an efficient data taking with a simple threshold trigger. During the remaining periods, however, the rate of noise transients is strongly increased. Second level trigger cuts, which are applied at each antenna station, are rather ineffective, other than proposed in earlier work. The maximum T2-rate allowed by the data acquisition system can only be kept by increasing the detection threshold.

Instead of pulse shape cuts, a novel method to reject periodic transients from the power grid was developed and tested. Although it is capable to reduce the applied threshold almost down to the level of quiet measuring periods, the introduced veto time cuts the effective uptime practically in half.

The third level trigger, in contrast, which combines the pulse candidates from individual stations to a joint event, performs much better than expected. To support a precise reconstruction of the event direction in the early trigger stage, an improved method to determine the pulse timestamp, based on a constant fraction discriminator is implemented. With a dynamic suppression of noisy directions, the coincidence rate can be reduced by four orders of magnitude, while less than a quarter of the detector's effective aperture is lost. This results in an event readout rate in the order of 5 mHz, which is a factor of 20 better than required.

Already during commissioning of the self-trigger, when using only three antennas, placed on an area hardly larger than the earlier prototype setup at the BLS, the first cosmic ray events could be detected. Besides the time difference towards the known surface detector events, which is already highly significant, also the measured pulse amplitudes are compatible with the expectations. At the same time, however, some events which were assumed to be detectable, could not be triggered. The trigger efficiency saturates around 50 %, even under the best possible conditions. This behavior was also found with the competing Dutch stations, which use a slightly different trigger approach. The reason for the reduced detection efficiency remains unclear, systematic effects, like for example disregarded dead-time, were excluded as far as possible.

Based on these results, three different modes of operation come into question for the close future of the AERA experiment:

The first possibility is, to accept a reduction of the effective uptime due to variable artificial noise. Using restrictive cuts on antenna station level, the measurement can be continued during high-noise periods with a reduced detection efficiency. As long as the reason for the variable noise remains unclear, and a worsening of the situation is possible, this approach is unsatisfying.

For a future radio experiment, a completely isolated site is advisory, a long term background survey is necessary to exclude rare background variations.

The second idea is, to exploit the very efficient background rejection of the coincidence trigger with dynamic hotspot suppression. This shifts the emphasis of the trigger decision from the antenna station level to the central data acquisition. For this purpose, the data acquisition and the communication system must be modified to cope with a strongly increased T2-rate, the search for such an alternative wireless system is currently underway [108].

The antenna arrangement of a future radio array can be optimized for this type of measurement, using a structured grid of antenna triplets. This way an additional

trigger stage is introduced between the station-level T2 and the experiment-level T3. It performs a direction reconstruction for each triplet, and benefits from the good angular background suppression. The dimensions of the antenna triangles would be small enough for a fast, probably cable-based communication with a very high T2-rate, but large enough for a sufficient direction resolution. A lateral length in the order of 50 m seems to be a good compromise. The distance between different triplets should be in the order of several 100 m, to efficiently cover a large area. On this scale, a wireless communication is feasible again, as the required bandwidth is small due to the low false trigger rate.

The third possibility, which was chosen for the upcoming operation, combines the regular self-triggered measurement with an external trigger from the surface detector. This way, the recording of every interesting cosmic ray event is guaranteed, even during noisy periods. It is also the most promising approach to examine the issue of the stagnating detection efficiency. Although the use of an external trigger is in principle a step backwards in detector development, it would be negligent to rely solely on the self-trigger as long as its stable performance is not completely proven. Implementing the external trigger is only possible with the KIT/BUW digitizer, as it provides a large buffer memory to bridge the latency of the surface detector communications.

At the Pierre Auger Observatory, the joint measurement complies well with the general philosophy of hybrid detection. For the future, the trend also favors hybrid experiments, which deliver a better data quality. Consequently it seems more probable, that the radio detection method is used as a powerful extension in a combined detector, than replacing the well-proven surface detector technique in a stand-alone implementation. Within a hybrid detector, an external trigger can be provided with ease, the radio self-trigger is then indispensable only for events, which are not detectable by particle detectors, for example very inclined showers.

8. Summary and Outlook

In recent years, the field of radio detection of cosmic rays made huge progress, the lateral distribution, polarization, and shape of the shower front were examined in detail. At the same time, a steady improvement of simulation models led to a better understanding of the underlying radio emission mechanisms. All these results were achieved with data from externally triggered radio detectors in combination with a surface detector array. So far, the radio observation could not emancipate itself as a full-fledged detection alternative.

For this final breakthrough, a reliable self-trigger is essential, but early attempts with an analog threshold detector followed by a coincidence trigger failed mostly due to the variable background conditions. Only at very remote locations, during selected low-noise periods, this simple trigger approach was occasionally successful. Therefore, the systematic research for a robust trigger strategy was started, using data from externally triggered measurements. The investigation identified monofrequent signals from radio transmitters as an important source of interference and proposed a filter method to suppress them. Most of the remaining background pulses could be rejected with a number of cuts on pulse shape parameters. The realization of this strategy as a digital self-trigger system is the main topic of this thesis.

The work started with a fast proof of concept, where the trigger was implemented using already available hardware for a **prototype measuring campaign** at the Pierre Auger Observatory. Based on the knowledge at that time, a rate of one detected cosmic ray event per week was assumed. It was planned to tune the trigger parameters for an optimal background suppression after recording the first events. In reality however, the radio emission had been overestimated, in particular the signal drop towards large lateral distances is much steeper than predicted. At the same time, the local radio noise level had been strongly underestimated, transient noise from a close-by powerline increased the detection threshold. As a result, not a single event could be detected during 15 months of measurement. Unfortunately, the simplified prototype setup did not provide for a fall-back method of operation, for example with an external trigger. Nevertheless, the campaign could prove the suitability of antenna and analog system during the long term of stable operation and improve the understanding of the background situation.

To finally tackle the difficulties of radio detection with a full-scale radio detector, the AERA collaboration was founded. For the new experiment, a made-to-measure **electronics was developed**, introducing the experiences from the prototype campaign into the new system. Its outstanding feature is a large ring buffer, which allows for an externally triggered readout with a latency of several seconds as a back-up solution. Within the AERA collaboration, this thesis covers most of the firmware development for the digitizer, including digital signal processing and early trigger stages.

In this context, a self-adapting method to **suppress radio frequency interferences** was implemented, to improve the background situation as much as possible before triggering. The result is an elaborate combination of a Fourier transform into frequency domain, removal of elevated frequency contributions, followed by a transform back to time domain. The realization of this complex filter processing the real-time data rate is only possible with the limited hardware resources when using numerical tricks and extensive optimization. In comparison with conventional digital filters, the performance of the suppression algorithm is superior, it results in a typical threshold reduction of 40 %.

After the first measurements at the AERA site, the variability of the background turned out to be an additional obstacle for an efficient trigger operation. While two thirds of the measuring time show a low-key noise level, where cosmic ray measurements are possible without further ado, the other third is heavily contaminated with noise transients. A discrimination of these transients from the wanted signal pulses is difficult, because the short pulse, which is emitted by an air shower, does not hold a strong, unique feature for identification. The analog filter additionally conceals potential differences of the pulse shape. Consequently, the proposed **pulse shape cuts** are rather inefficient to reduce the noise transient rate without rejecting cosmic ray induced events.

The only remarkable characteristic of the noise transients is their periodicity with line frequency, which can be exploited for their rejection. For this purpose, a digital phase-locked loop was designed, which identifies the periodicity in real-time and predicts future transients, which are then vetoed. This **periodic veto** method successfully reduces the background situation to the level of the low-noise periods. Due to the temporal blur of the periodicity, however, the veto time has to cover half of the effective uptime. As the allowed station trigger rate is limited by the capacity of the communication system, the loss of uptime is acceptable as a last resort. The only alternative would be a strong increase of the threshold level, which would limit cosmic ray detection to energies above $5 \cdot 10^{18} \text{eV}$.

Although an artificial origin of the noise transients is proven by their correlation to the power grid frequency, their exact source remains unclear, just like the reason for their strong variability. This variability explains, why the transient noise was not discovered in the survey measurements conducted before the construction of AERA. This means that for the selection of a future radio detector site, a long-term survey over several weeks is required.

When combining the trigger candidates from multiple antenna stations to a joint event, the incoming direction of the pulse can be reconstructed as an additional distinctive feature. Because most of the background transients stem from confined directions, a dynamic suppression of these hotspots is extremely efficient: The background is reduced by four to five orders of magnitude, while less than a quarter of the detector's effective aperture angle is lost. The remaining event readout rate is in the order of 5 mHz. The high performance of the **dynamic hotspot suppression**, opens a new perspective for self-triggered operation without uptime losses, despite the suboptimal background conditions at the AERA site. By strongly increasing the allowed background trigger rate of each station, and thus moving the emphasis of transient rejection from station level to coincidence level, an efficient measurement should be possible even during noisy periods. For this purpose, the communication system and the coincidence trigger must be updated in order to be capable of handling the increased load.

Already during commissioning of the trigger system, using only a triangle of antenna stations, the first **cosmic ray events were successfully detected**. In the meantime, a field of 17 stations was equipped with the new electronics, the number of detected events rises every day now.

By using a complementary **external trigger** from the surface detector, the recording of events close to the detection threshold is guaranteed even under poor background conditions. In doing so, also remaining uncertainties about the achievable self-trigger efficiency can be examined in detail. Currently, the next extension stage of the AERA experiment is prepared, increasing the total number of antenna stations by 110. With the increased event statistics, the most advanced investigation of the radio emission at high energies will succeed, finally proving whether the method of radio detection can live up to the high expectations.

A. Appendix

A.1. Characteristics of the Fourier transform

Because the radio signal and also the detection instruments have many frequency-dependent properties, transformations into frequency domain and back to time domain are very common in the business of radio detection of cosmic rays. For the time-discretized sampled signals, the Discrete Fourier transform (DFT) is used. The complex frequency components A_k are calculated from the complex input time series a_n with $n = 0 \dots N-1$ like:

$$A_k = \sum_{n=0}^{N-1} e^{-2\pi i \frac{k \cdot n}{N}} \cdot a_n \quad k = \{0 \dots N-1\} \quad (\text{A.1})$$

The signal in time domain is again reconstructed from the frequency domain with the inverse Fourier transform:

$$a_n = \frac{1}{N} \sum_{k=0}^{N-1} e^{2\pi i \frac{k \cdot n}{N}} \cdot A_k \quad n = \{0 \dots N-1\} \quad (\text{A.2})$$

Although the Fast Fourier transform (FFT) algorithm is normally used, Fourier transforms are still relatively computing-intensive operations. The DFT contains symmetries, which can be exploited to simplify calculations as shown in the following. These optimizations are especially important when using real-time systems and embedded platforms with very limited resources.

A.1.1. Fast Fourier transform

For the calculation of a DFT, the multiplying is numerically most expensive, to Fourier transform vector of N points, N^2 complex multiplications are necessary, resulting in $(4N)^2$ real multiplications. The Fast Fourier algorithm reduces the number of multiplications with a divide and conquer approach. J. Cooley and J. Tukey are generally credited with the algorithm [109], though already C.F. Gauß used a variation for asteroid orbit calculations [110].

The DFT sum over an even number of coefficients a_n can be split into even and uneven components a'_n and a''_n :

$$\begin{aligned} A_k &= \frac{1}{N} \sum_{n=0}^{N/2-1} e^{-2\pi i \frac{k \cdot 2n}{N}} \cdot a_{2n} + \frac{1}{N} \sum_{n=0}^{N/2-1} e^{-2\pi i \frac{k \cdot (2n+1)}{N}} \cdot a_{2n+1} \\ &= \frac{1}{2N'} \sum_{n=0}^{N'-1} e^{-2\pi i \frac{k \cdot n}{N'}} \cdot a'_n + \frac{1}{2N'} e^{-\frac{\pi i}{N'}} \sum_{n=0}^{N'-1} e^{-2\pi i \frac{k \cdot n}{N'}} \cdot a''_n \\ &= \begin{cases} \frac{1}{2} A'_k & + \quad \frac{1}{2} e^{-\frac{\pi i k}{N'}} A''_k & \text{if } k < N' \\ \frac{1}{2} A'_{k-N'} & - \quad \frac{1}{2} e^{-\frac{\pi i (k-N')}{N'}} A''_{k-N'} & \text{if } k \geq N' \end{cases} \quad (\text{A.3}) \end{aligned}$$

The N -point DFT is thereby reduced to two N' -point DFTs ($N' = N/2$), the number of complex multiplications is only $2 \cdot (N/2)^2$. If N is a power of 2, the splitting can be repeated until a trivial sum, resulting in a total complexity of $O(2N \cdot \log_2(N))$.

A.1.2. Fourier transform for real signals

The complex frequency components A_k can be imagined as rotating vectors in the complex plane. To receive a sinusoidal oscillation, two vectors with opposing rotation directions must be superimposed:

$$A'_k \cdot \sin(kt + \phi) = A_k \cdot e^{-2\pi i \cdot kt} + A_{-k} \cdot e^{+2\pi i \cdot kt} \quad (\text{A.4})$$

One rotation direction is assigned to positive frequencies k , the other one to negative frequencies $-k$. According to the Shannon-Nyquist sampling theorem, the frequencies only span unambiguously between $-N/2$ and $N/2$. Outside this range they are repeated periodically, also called aliasing. Utilizing this, the range $N/2 < k < N$ can be conveniently used for the negative frequencies. In most application cases input signals a_n are real, in this case, the negative frequency components are the conjugate-complex of the positive frequencies, and contain only redundant information.

$$A_{N-k} = \sum_{n=0}^{N-1} e^{-2\pi i \frac{(N-k) \cdot n}{N}} \cdot a_n = \sum_{n=0}^{N-1} e^{+2\pi i \frac{k \cdot n}{N}} \underbrace{e^{-2\pi i \cdot n}}_{=1} \cdot a_n = \overline{A_k} \quad (\text{A.5})$$

Similarly for a purely imaginary signal a_n , the negative frequencies are the negative conjugate of the positive frequencies. In these cases, the calculation of the negative frequency components may be economized.

A.1.3. Upsampling in time domain

A signal can be upsampled by a factor of z , by modifying the original time series a_n , ($n = 0 \dots N-1$) to a'_n , ($n = 0 \dots zN-1$) with $z-1$ zeroes filled in between each pair of samples. The spectrum of the zero-padded signal is:

$$\begin{aligned} A'_k &= \sum_{n=0}^{zN-1} e^{-2\pi i \frac{k \cdot n}{zN}} \cdot a'_n & k &= \{0 \dots zN-1\} \\ &= \sum_{m=0}^{N-1} e^{-2\pi i \frac{k \cdot zm}{zN}} \cdot a'_{zm} = \sum_{m=0}^{N-1} e^{-2\pi i \frac{k \cdot m}{N}} \cdot a_m = A_k \end{aligned} \quad (\text{A.6})$$

This means that the DFT of the zero-padded signal a'_n is for $n < N$ equal to the original signal's Fourier transform. The spectrum at the higher indices n contains just copies of the original spectrum:

$$A'_{k+N} = \sum_{m=0}^{N-1} e^{-2\pi i \frac{(k+N) \cdot m}{N}} \cdot a_m = \sum_{m=0}^{N-1} e^{-2\pi i \frac{k \cdot m}{N}} \cdot \underbrace{e^{-2\pi i \cdot m}}_{=1} \cdot a_m = A_k \quad (\text{A.7})$$

In the process, the total signal power stays the same, fulfilling the Parseval theorem: The number of frequency bins with a signal is multiplied certainly, but the bin width is reduced by the same factor. When a signal is sub-sampled in the second Nyquist domain, the actual frequency range of the zero-padded signal contains only the negative frequencies of the original signal, the order is reverse - as required. For real input

signals, negative and positive frequencies are just conjugate. The imaginary parts cancel each other out, so the zero-padded signal directly has the wanted up-sampled frequency contents. Upsampling from odd-numbered Nyquist domains even works for complex signals, because in these cases the positive frequencies are copied to the wanted frequency range. To get rid of additionally introduced frequencies, a band-pass filter must be applied. This step reduces the total signal power by a factor of z , which can be corrected already during the zero-padding process. The filter can be realized as a digital Finite Impulse Response (FIR) filter for example. This way, the upsampling chain is processed entirely in time domain, no Fourier transforms are calculated. As the upsampling algorithm works as a continuous pipeline, it is particularly suited for endless data streams.

A.1.4. Hilbert transform

Recorded signal time series usually contain only real values, leading to redundant information in frequency domain in form of the negative frequencies. Another possibility, called analytical representation, abstains from the negative frequencies, and thereby introduces redundant information in time domain in form of an imaginary signal part. The imaginary part is calculated from the real signal a_n with the Hilbert transform $H(a_n)$, a convolution operation, which in frequency domain translates to a phase shift by $-\pi/2$ for each positive frequency component, negative frequencies are shifted by $+\pi/2$. For example a real oscillation $\sin(\omega \cdot t)$ is transformed to $-\cos(\omega \cdot t)$, the Hilbert transform always has maximum amplitude, where the original signal has a zero-crossing. This is very useful to calculate the signal envelope by quadratically adding Hilbert transform and original signal.

When the signal is already given in frequency domain, both Hilbert transform and original signal can be calculated very efficiently with just one Fourier transform by doubling positive frequency coefficients and setting negative frequencies to zero. The original signal is then found in the real part, the Hilbert transform in the imaginary part of the complex output.

$$\begin{aligned}
a_n + i \cdot H(a_n) &= \frac{1}{N} \sum_{k=0}^{N-1} e^{2\pi i \frac{k \cdot n}{N}} \cdot A_k + \\
&\quad \frac{i}{N} \left(\sum_{k=0}^{N/2-1} e^{2\pi i \frac{k \cdot n}{N}} e^{-\frac{i\pi}{2}} A_k + \frac{1}{N} \sum_{k=N/2}^{N-1} e^{2\pi i \frac{k \cdot n}{N}} e^{\frac{i\pi}{2}} A_k \right) \\
&= \frac{1}{N} \sum_{k=0}^{N/2-1} A_k \left(e^{2\pi i \frac{k \cdot n}{N}} + i e^{2\pi i \frac{k \cdot n}{N} - \frac{i\pi}{2}} \right) + \\
&\quad \frac{1}{N} \sum_{k=N/2}^{N-1} A_k \left(e^{2\pi i \frac{k \cdot n}{N}} + i e^{2\pi i \frac{k \cdot n}{N} + \frac{i\pi}{2}} \right) \\
&= \frac{2}{N} \sum_{k=0}^{N/2-1} e^{2\pi i \frac{k \cdot n}{N}} \cdot A_k
\end{aligned} \tag{A.8}$$

A.2. Digital filters

A.2.1. FIR filter

Instead of filtering signals in frequency domain with the help of Fourier transforms, a filter can also be implemented in time domain. The most simple approach is the Finite Impulse Response Filter (**FIR**) as outlined in figure A.1, where the output signal is determined by the filter coefficients a_i as:

$$y_k = \sum_{n=0}^N a_n \cdot x_{k-n} \quad (\text{A.9})$$

With the help of the Z-transform, this response function can be translated into the Z-transfer function,

$$H(z) = \sum_{n=0}^N a_n \cdot z^{-n} = \frac{\sum a_n \cdot z^{N-n}}{z^N} \quad (\text{A.10})$$

The transfer function has N zero points, that can be placed freely in the Z-plane to define the filter characteristics, the N poles are fixed to the origin. The frequency response is determined by the behavior of the transfer function on the unit circle $z = e^{i\Omega}$, with the normalized frequency $\Omega = 2\pi f/f_s$. As a result the frequency response is the Fourier series of the filter coefficients, the filter coefficients can be calculated with an inverse Fourier transform of the desired filter shape:

$$H(f) = \sum_{n=0}^N a_n \cdot e^{-2\pi i \frac{n \cdot f}{f_s}} \quad (\text{A.11})$$

The Fourier transform suffers from spectral leakage, which reduced the stop-band attenuation. By using window functions as explained in section 6.3, this leakage can be reduced at the cost of a lower filter steepness.

As the name suggests, the response of a FIR filter is limited to its width N . Because the FIR filter does not contain any feedback lines, no oscillations can build up, the filter is always stable. Therefore the required numerical accuracy of the coefficient multiplication is moderate, implementation in FPGA logic is simple. The bit-width of the coefficients is often low enough, so it is more efficient to pre-calculate all possible multiplication results, and store them in memory, to spare valuable multiplications.

With a FIR-chain every linear filter can be realized, both frequency response and phase response can be selected independently. Typical applications are filters with a constant group delay, or phase-correcting all-pass filters. Unfortunately, the order N of the filter must be high in order to achieve a good frequency resolution, as needed for a steep filter slope, for example. This is due to the fact that each filter element is




Figure A.1.: Block diagram of a 3rd order digital FIR filter.

over-damped and suffers from a low Q-factor, as all poles of the filter in Z-plane lie in the origin.

Although the effort can usually be cut in half by choosing symmetric coefficients, the complexity of a FIR filter rises linearly with the required resolution, while the complexity of a FFT-engine would only rise logarithmically¹. Therefore FIR filters are normally used for filter orders below 100 or low-speed applications, e.g. audio signals, where the multiplications are conducted sequentially.

A.2.2. IIR filter




Figure A.2.: Block diagram of a 2nd order digital IIR filter.

The Infinite Impulse Response Filter (**IIR**) is a generalization of the FIR filter, where also the poles of the filter in Z-plane are freely configurable. This is achieved by introducing a feedback of the output signal, as illustrated in figure A.2. The output signal and the Z-transfer function are:

$$y_k = \sum_{n=0}^N a_n \cdot x_{k-n} + \sum_{n=1}^N b_n \cdot y_{k-n}$$

$$H(z) = \frac{\sum a_n \cdot z^{-n}}{\sum b_n \cdot z^{-n}} \quad (\text{A.12})$$

Due to the feedback, the pulse response is in principle endless, although IIR filters are commonly designed with a certain damping, so oscillations wear off eventually. The according Q-factor can be increased arbitrarily by moving the poles of the transfer function towards the unit circle.

An important issue of IIR filters, particularly for high Q-factors, is the required numerical accuracy to ensure filter stability and avoid self-excited oscillations. To improve the tolerance towards quantization errors, high order IIR filters are implemented as cascades of second-order-structures (SOS). Compared with a FIR filter, the IIR approach has large advantages in implementing a low-order filter with a sharp discrimination of neighboring frequencies. A typical application would be a narrow notch filter. Such a low-order filter has only few degrees of freedom, so the phase response can not be adjusted, and a significant dispersion is introduced. For filters of higher order, a high Q-factor is often not important anymore, so that the simpler FIR-filters are preferred.

A.3. Hardware details

A.3.1. The FPGA choice

During hardware development, different FPGA models were considered. The main question was, whether to use high-end Stratix, or low-cost Cyclone FPGAs. While high-

¹ Taking into account that a $O(N \cdot \log(N))$ -complex FFT processes a block of N samples.

end FPGAs allow much higher clock frequencies, their complex design, in particular the elaborate clock tree, leads to a disadvantage in terms of power consumption even though a smaller feature size. To simulate the performance of the different FPGAs, an earlier version of the self-trigger implementation was used. The following table compares the contemplable Stratix IV and Cyclone III FPGAs with the Stratix I from the prototype setup, and motivates the choice of the Cyclone III model.

FPGA Family	Stratix I	Stratix IV	Cyclone III
Model	EP1S40	EP4SGX70	EP3C80
Speed grade	5	4	6
Logic elements	41,250	72,600 ¹	81,264
Feature size	130 nm	40 nm	65 nm
FPGA usage	86 %	53 %	44 %
Maximum speed	161 MHz	299 MHz	218 MHz
Power consumption ²	16.8 W	3.2 W	2.9 W
Cost in US\$ ³	1900 ⁴	800	335

¹ Equivalent LEs, Stratix IV architecture uses Adaptive Logic Modules (ALM).

² As estimated by the Altera PowerPlay Estimator for a 180 MHz clock .

³ Price estimate from buyaltera.com, state: January 2012.

⁴ Stratix I is outdated, devices are only offered for legacy purpose for high prices.

A.3.2. Interfaces of the AERA board

The AERA digitizer consists of several modules, all connected to the central FPGA. The communication with these modules on the digitizer board and to the outside world is provided with several interfaces.

A **JTAG** connector is important during commission to install the first firmware, but also for later debugging. Besides the FPGA, the MSP430 microcontroller and a potential Blackfin module have their own JTAG connector.

The **Ethernet** interface is the communication method during AERA stage 1 via optical media converters. The Media Access Controller (MAC) is implemented in the FPGA and provided by OpenCores². The physical driver chip is located on a small piggy-bag board, which could be exchanged by a full Blackfin processor board.

The **CAN**³ bus will be used to access the wireless communication system in AERA phase 2 and 3. The controller is again provided by OpenCores. To fully comply with the CAN speed specifications, a dedicated PLL clock of 20 MHz was necessary, and thereby a wrapper to ensure a proper clock domain transition of the data.

RS 232⁴ was used during commissioning for a provisional terminal connection, and later to operate the solar charge controller. The employed UART⁵ is an IP-core from Altera, the line driver uses a MAX232 chip.

² OpenCores.org is a large open source hardware developing community.

³ Controller Area Network is a communication bus standard widely used in automotive industry.

⁴ RS 232 is a single ended, UART-based communication standard, also known as serial port.

⁵ A Universal Asynchronous Receiver/Transmitter is a piece of hardware, basically consisting of a shift register to allow the simplest possible form of serial asynchronous data transfer.

A second **UART** is used to set up the GPS clock and read out its time telegram each second, thereby the Processor has access to the full UTC time. The IP is the same as for the RS 232 port.

The **SD-RAM⁶ controller** connects the NIOS processor with its working memory. The controller is an Altera IP-core, the memory chip is located on the extension piggy-bag board.

Through an **SPI⁷** the SD-card⁸ is accessed, where a persistent file system for the operating system is stored.

A **second SPI** is used to configure the ADCs, options like reference voltage or output timing can be selected. The IP-core is, as for the SD-card interface, an Altera IP-core.

I²C⁹ connects the NIOS processor with a small EEPROM¹⁰ containing an unique board id, and with the MSP430 microcontroller. The MSP430 not only runs the Watchdog, but also controls all voltages, and reads out local temperature sensors for monitoring. The I²C IP-core is provided by OpenCores.

With the **EPCS¹¹** interface, the processor can access the boot Flash memory to apply remote updates. The IP-core is, like the related chip, proprietary from Altera.

Besides the interfaces mentioned so far, all connected via Avalon Bus to the NIOS processor, there are low-level ports as well, connected directly to FPGA-logic. This approach is used when the timing of the signals is crucial:

4 **LVDS¹²** ports, selectable as input or output provide low level fast communication of future extensions. For example, an external trigger signal from local scintillators could use this port.

A **Test-pulse** output can be used to apply a stimulus pulse to the input of the analog filter, for remote functional tests, and calibration in the field. The height of the pulse is selected by a Digital-to-Analog-Converter from the MSP430, a fast switch controlled from the FPGA and accessed through the DMU defines the pulse width.

The GPS clock not only provides the slow time telegrams via UART, but also a strobe signal each second. This **PPS-strobe** (pulse per second) defines with its rising edge a precise time signal, which is used in the FPGA by the DMU entity to derive a proper event time. Because the PPS-strobe output is gated with the crystal-clock signal from the GPS-clock, which slightly drifts in time, the PPS-strobe also drifts within one clock cycle. This is the so-called sawtooth-drift. The GPS-clock also monitors its drift and calculates, when the pulse actually should

⁶ Synchronous Dynamic Random-Access Memory.

⁷ Serial Peripheral Interface is a method of simple synchronous serial data transfer.

⁸ Secure Digital Card - a common Flash memory card in consumer electronics.

⁹ I²CTM (Inter-Integrated Circuit) is a low speed communication bus, also known as Two Wire Interface.

¹⁰ Electrically Erasable Programmable Read-Only Memory is a persistent Memory device.

¹¹ Enhanced Programming Configuration Serial Device is a proprietary Altera chip necessary to boot firmware and software during power-up, the data is stored in on-chip Flash memory.

¹² Low Voltage Differential Signal, a digital signaling standard with very good EMC.

have occurred. This sawtooth correction is also available via UART, and used by the NIOS processor to reach the full timing resolution.

A.4. Self-triggered cosmic ray events

During the first months of measuring at AERA only three antenna station were operated. A primitive setup, only recording timestamp and a basic pulse height for each trigger at each local station was used, to get early data, before the data acquisition system was operational. This measurement resulted in nine events coincident with the surface detector (see also section 7):

Auger ID	Antenna station	Polarization	time difference	expected amplitude S	measured amplitude
12449200	20	E/W	99 ns ¹	1.58	216
12449516	20	E/W	-93 ns ¹	0.31	138
12494290	22	E/W	-12 ns ¹	2.53	296
12505312	22	E/W	56 ns ¹	1.07	159
12527244	22	E/W	29 ns ¹	1.63	238
12640182	22	E/W	87 ns	0.67	144
13016465	20	E/W	293 ns	1.54	167

Table A.1.: List of the detected events with the antenna station and polarization they were detected in, their time difference (SD-Radio) corrected for known offsets, the expected unscaled amplitude according to equation. 7.1, and the measured pulse height in ADC units.

¹ In the beginning, the T2-timestamps were mistakenly assigned to the start point of the trace, rather than to the pulse position. The time difference of 5.9 μ s is corrected here.

Auger ID	Antenna	E_p	Azimuth	Zenith	Distance
12449200	20	$3.9 \cdot 10^{17} eV$	327°	35°	68 m
12449516	20	$1.1 \cdot 10^{17} eV$	327°	44°	129 m
12494290	22	$2.6 \cdot 10^{18} eV$	79°	16°	120 m
12505312	22	$3.1 \cdot 10^{17} eV$	177°	24°	26 m
12527244	22	$3.1 \cdot 10^{17} eV$	277°	35°	64 m
12640182	22	$1.3 \cdot 10^{17} eV$	249°	34°	63 m
13016465	20	$3.6 \cdot 10^{17} eV$	205°	18°	34 m

Table A.2.: Air shower parameters of the detected events, taken from the surface detector: Identification number, antenna station where the coincidence was detected in, primary energy E_p , incoming azimuth angle (0° = east), zenith angle, and distance to the antenna station.

Bibliography

- [1] J. Linsley. Five-Year Wilkinson Microwave Anisotropy Probe Observations: Data Processing, Sky Maps, and Basic Results. *The Astrophysical Journal Supplement*, 180:225–245, 2009.
- [2] J. Linsley. Evidence for a Primary Cosmic-Ray Particle with Energy 10^{20} eV. *Physical Review Letters*, 10:146–148, 1963.
- [3] V. Berezhinsky, P. Blasi, and V. Ptuskin. Clusters of Galaxies as Storage Room for Cosmic Rays. *The Astrophysical Journal*, 487:529, 1997.
- [4] J. Blümer, R. Engel, and J. Hörandel. Cosmic rays from the knee to the highest energies. *Progress in Particle and Nuclear Physics*, 63:293–338, 2009.
- [5] E. Fermi. On the Origin of Cosmic Radiation. *Physical Review*, 75:1169–1174, 1949.
- [6] F. Aharonian et al. - HESS Collaboration. Primary particle acceleration above 100 TeV in the shell-type supernova remnant RX J1713.7-3946 with deep H.E.S.S. observations. *Astronomy & Astrophysics*, 464:235–243, 2007.
- [7] W. Apel et al. - KASCADE-Grande Collaboration. Upper Limit of the Spectrum of Cosmic Rays? *Astroparticle Physics*, 31:86–91, 2009.
- [8] J. Abraham et al. - Pierre Auger Collaboration. Correlation of the Highest-Energy Cosmic Rays with Nearby Extragalactic Objects. *Science*, 318:938–943, 2007.
- [9] J. Abraham et al. - Pierre Auger Collaboration. Update on the correlation of the highest energy cosmic rays with nearby extragalactic matter. *Astroparticle Physics*, 34:314–326, 2010.
- [10] K. Greisen. End of the Cosmic-Ray Spectrum? *Physical Review Letters*, 16:748–750, 1966.
- [11] G. Zatsepin, V. Kuz'min. Upper Limit of the Spectrum of Cosmic Rays? *Journal of Experimental and Theoretical Physics Letters*, 4:78–80, 1966.
- [12] D. Allard, E. Parizot, A. Olinto, et al. Measurement of the energy spectrum of cosmic rays above 10^{18} eV using the Pierre Auger Observatory. *Astronomy & Astrophysics*, 443:29–32, 2005.
- [13] J. Abraham et al. - Pierre Auger Collaboration. Measurement of the energy spectrum of cosmic rays above 10^{18} eV using the Pierre Auger Observatory. *Physics Letters B*, 685:239–246, 2010.
- [14] J.W. Cronin. The highest-energy cosmic rays. *Nuclear Physics B - Proceeding Supplements*, 138:465–491, 2005.

- [15] P. Auger et al. Extensive Cosmic-Ray Showers. *Reviews of Modern Physics*, 11:288–291, 1939.
- [16] T. Gaisser and A. Hillas. Reliability of the method of constant intensity cuts for reconstructing the average development of vertical showers (of cosmic rays). In *Proceedings of the 15th. ICRC*, volume 8, pages 353–357, 1977.
- [17] O.C. Alkhofer. *Introduction to cosmic radiation*. Verlag Carl Thiemig München, 1975.
- [18] D. Heck, J. Knapp, J. Capdevielle et al. COSIKA: A Monte Carlo code to simulate extensive air showers. FZKA report 6019, Forschungszentrum Karlsruhe, 1998.
- [19] S. Sciutto. AIRES: A system for air shower simulation (Version 2.2.0). *ArXiv Astrophysics e-prints*, 1999. astro-ph/9911331.
- [20] W. Apel et al. - KASCADE Collaboration. Test of interaction models up to 40 PeV by studying hadronic cores of EAS. *Journal of Physics G: Nuclear and Particle Physics*, 34:2581–2594, 2007.
- [21] T. Bergmann et al. One dimensional hybrid approach to extensive air shower simulation. *Astroparticle Physics*, 26:420–432, 2007.
- [22] V. Hess. über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten. *Physikalische Zeitschrift*, 13:1084–1091, 1912.
- [23] J.-P. Vialle et al. - AMS Collaboration. AMS: A cosmic ray observatory. *Nuclear Physics Proceedings Supplements (B)*, 150:24–29, 2006.
- [24] E. Seo et al. - CREAM Collaboration. Approaching the spectral knee in high energy cosmic rays. *Journal of the Physical Society of Japan*, 78:63–67, 2009.
- [25] P. Auger et al. Grandes gerbes cosmiques atmosphériques contenant des corpuscules ultra-pénétrantes. *Comptes rendus, Académie des Sciences*, 206:1721–1722, 1938.
- [26] J. Abraham et al. - Pierre Auger Collaboration. Properties and performance of the prototype instrument for the Pierre Auger Observatory. *Nuclear Instruments and Methods in Physics Research A*, 523:50–95, 2004.
- [27] N.M. Budnev et al. - Tunka Collaboration. The Tunka-133 EAS Cherenkov array - status, first results and plans. In *Proceedings of the 31st ICRC*, page 1069, 2009. astro-ph/1003.0089.
- [28] Rasha Abbasi and Paolo Desiati - IceCube Collaboration. Large Scale Cosmic Ray Anisotropy With IceCube. *Astrophysical Journal*, 718:194, 2010.
- [29] G.A. Askaryan. Excess negative charge of an electron-photon shower and its coherent radio emission. *Soviet Physics JETP*, 14:441, 1962.
- [30] F. Kahn and I. Lerche. Radiation from cosmic ray air showers. In *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, volume 289, page 206, 1966.
- [31] J. Jelley, J. Fruin, N. Porter, et al. Radio Pulses from Extensive Cosmic-Ray Air Showers. *Nature*, 205:327–328, 1965.

- [32] T.Huege and H. Falcke. Radio emission from cosmic ray air showers. Monte Carlo simulations. *Astronomy & Astrophysics*, 430:779–798, 2005.
- [33] T. Huege and H. Falcke. Radio emission from cosmic ray air showers: Simulation results and parametrization. *Astroparticle Physics*, 24:116, 2005.
- [34] C.W. James, H. Falcke, T. Huege and M. Ludwig. An 'endpoint' formulation for the calculation of electromagnetic radiation from charged particle motion. *submitted to: Physical Review E*, 2010. astro-ph/1007.4146.
- [35] M. Ludwig and T. Huege. REAS3: Monte Carlo simulations of radio emission from cosmic ray air showers using an 'end-point' formalism. *Astroparticle Physics*, 34:438–446, 2011.
- [36] T. Huege, M. Ludwig, O. Scholten, and K.D. de Vries. The convergence of EAS radio emission models and a detailed comparison of REAS3 and MGMR simulations. In *Nuclear Instruments and Methods A; Proceedings of the ARENA 2010 conference*, 2010. DOI:10.1016/j.nima.2010.11.041.
- [37] O. Scholten, K. Werner, and F. Rusydi. A macroscopic description of coherent geosynchrotron radiation from cosmic-ray air showers. *Astroparticle Physics*, 29:94–103, 2008.
- [38] T. Huege and H. Falcke. Radio emission from cosmic ray air showers. Coherent geosynchrotron radiation. *Astronomy & Astrophysics*, 412:19–34, 2003.
- [39] G.A. Askaryan. Coherent radio emission from cosmic showers in air and in dense media. *Soviet Physics JETP*, 21:658, 1965.
- [40] P.W. Gorham et al. - ANITA Collaboration. Observations of the Askaryan Effect in Ice. *Physical Review Letters*, 99:171101, 2006.
- [41] P.W. Gorham et al. Observation of Microwave Continuum Emission from Air Shower Plasmas. *Physical Review D*, 78:032007, 2008.
- [42] M. Ender. Radiodetektion von Luftschauern unter dem Einfluss starker elektrischer Felder in der Atmosphäre. FZKA report 7506, Forschungszentrum Karlsruhe, 2009.
- [43] S. Buitink et al. - LOPES Collaboration. Amplified radio emission from cosmic ray air showers in thunderstorms. *Astronomy & Astrophysics*, 467:385–394, 2007.
- [44] H. Falcke et al. - LOPES Collaboration. Detection and imaging of atmospheric radio flashes from cosmic ray air showers. *Nature*, 435:313–316, 2005.
- [45] F. Schröder. *Instruments and Methods for the Radio Detection of High Energy Cosmic Rays*. PhD thesis, Karlsruhe Institute of Technology, 2010.
- [46] O. Vedenev. Depth of the maximum of extensive air showers and mass composition of primary cosmic radiation at an energy of $4 \cdot 10^{17}$ eV according to data on radioemission from extensive air showers. *Physics of Atomic Nuclei*, 72:250–256, 2009.
- [47] S.P. Knurenko et al. 32 MHz radio measurements at the Yakutsk EAS array. In *Proceedings of the 22nd ECRS, to appear in Astrophysics and Space Science Transactions*, 2010.

- [48] A. Haungs et al. - KASCADE-Grande Collaboration. Latest Results and Perspectives of the KASCADE-Grande EAS Facility. In *Nuclear Instruments and Methods A, Proceedings of the ARENA 2010 conference*, 2010. DOI:10.1016/j.nima.2010.11.126.
- [49] S. Nehls. Calibrated Measurements of the Radio Emission of Cosmic Ray Air Showers. FZKA report 7440, Forschungszentrum Karlsruhe, 2008.
- [50] P.G. Isar et al. - LOPES Collaboration. Polarization studies of the EAS radio emission with the LOPES experiment. In *Proceedings of the 31st ICRC*, page 1128, 2009.
- [51] A. Horneffer et al. - LOPES Collaboration. Primary Particle Energy Calibration of the EAS Radio Pulse Height. In *Proceedings of the 30th ICRC*, volume 4, pages 83–86, 2008.
- [52] T. Asch. Self-Triggering of Radio Signals from Cosmic Ray Air Showers. FZKA report 7459, Forschungszentrum Karlsruhe, 2009.
- [53] H. Röttgering et al. LOFAR: a new radio telescope for low frequency radio observations: science and project status. In *Texas in Tuscany. XXI Symposium on Relativistic Astrophysics*, pages 69–76, 2003.
- [54] O. Scholten et al. First results of the NuMoon experiment. *Nuclear Instruments and Methods in Physics Research A*, 604:102–105, 2009.
- [55] D. Ardouin et al. - CODALEMA Collaboration. Radioelectric field features of extensive air showers observed with CODALEMA. *Astroparticle Physics*, 26:341–350, 2006.
- [56] D. Ardouin et al. - CODALEMA Collaboration. Geomagnetic origin of the radio emission from cosmic ray induced air showers observed by CODALEMA. *Astroparticle Physics*, 31:192–200, 2009.
- [57] J. Auffenberg, D. Besson, T. Gaisser, et al. A radio air shower detector as an extension for IceCube and IceTop. *Nuclear Instruments and Methods in Physics Research A*, 604:53–56, 2009.
- [58] I. Kravchenko et al. - RICE Collaboration. Recent results from the RICE experiment at the South Pole. In *Nuclear Instruments and Methods A, Proceedings of the ARENA 2010 conference*, 2010. DOI:10.1016/j.nima.2010.11.2012.
- [59] A.-Y. Sun, P. Chen and M. Huang. Simulation of the Event Reconstruction of Ultra High Energy Cosmic Neutrinos with Askaryan Radio Array. ArXiv e-prints astro-ph/1002.0023, 2010.
- [60] S. Hoover et al. - ANITA Collaboration. Observation of Ultrahigh-Energy Cosmic Rays with the ANITA Balloon-Borne Radio Interferometer. *Physical Review Letters*, 105:151101, 2010.
- [61] D. Ardouin et al. First detection of Extensive Air Showers by the TREND self-triggering radio experiment. *Astroparticle Physics*, 37:717–731, 2011.
- [62] O. Krömer. *Empfangssystem zur Radioobservation hochenergetischer kosmische Schauer und sein Verhalten bei Selbsttriggerung*. PhD thesis, Karlsruhe Institute of Technology, 2008.

- [63] A. Haungs et al. - KASCADE-Grande Collaboration. The KASCADE-Grande experiment. In *Proceedings of the 28th ICRC*, 2003.
- [64] CCIR. Electrical characteristics of the surface of the earth. Technical Report 527-2, International Telecommunications Union, 1991.
- [65] CCIR. Worldwide minimum external noise levels, 0.1 Hz to 100 GHz. Technical Report 670, International Telecommunications Union, 1990.
- [66] O. Seeger et al. - Pierre Auger Collaboration. Logarithmic periodic dipole antennas for the Auger Engineering Radio Array. *Nuclear Instruments and Methods in Physics Research A*, 2010. doi:10.1016/j.nima.2010.10.140.
- [67] D. Charrier - CODALEMA Collaboration. Antenna development for astroparticle and radioastronomy experiments. *Nuclear Instruments and Methods in Physics Research A*, 2010. doi:10.1016/j.nima.2010.10.141.
- [68] M. Helfrich. Untersuchung der Nachweisschwelle und Effizienz eines Detektorsystems für Radiopulse von kosmischen Luftschauern. Diploma thesis, Karlsruhe Institute of Technology, 2010.
- [69] M. Frigo and S.G. Johnson. The Design and Implementation of FFTW3. In *Proceedings of the IEEE*, volume 93, 2005.
- [70] B. Keilhauer et al. - Pierre Auger Collaboration. The Balloon-the-Shower programme of the Pierre Auger Observatory. *Astrophysics and Space Sciences Transactions*, 6:27–30, 2010.
- [71] A. v. d. Berg et al. First detection of radio signals from cosmic rays at the Pierre Auger Observatory. GAP-Note 65, Pierre Auger Observatory, 2007.
- [72] H. Cane. Spectra of the non-thermal radio radiation from the galactic polar regions. *Monthly Notices of the Royal Astronomical Society*, 189:465–478, 1979.
- [73] S. Grebe, H. Schoorlemmer. The impact of RFI suppression on signal-to-noise calculations of events measured with the radio detection stations near the BLS. GAP-Note 69, Pierre Auger Observatory, 2009.
- [74] B. Revenu. Airplane signals detected by the RAuger radio stations: spectral and time domain characteristics, inter-calibration and polarization studies. GAP-Note 78, Pierre Auger Observatory, 2011.
- [75] R.D. Pierce. Application of the positive alpha-stable distribution. In *Proceedings of the IEEE Signal Processing Workshop on High Order Statistics*, pages 420–424, 1997.
- [76] A. v. d. Berg et al. An Ansatz for the parameterized analysis of transient noise studies. GAP-Note 31, Pierre Auger Observatory, 2011.
- [77] O. Krömer et al. Simulation of antennas proposed for radio detection of UHECRs. GAP-Note 35, Pierre Auger Observatory, 2009.
- [78] A. v. d. Berg et al. - Pierre Auger Collaboration. Radio detection of high-energy cosmic rays at the Pierre Auger Observatory. In *Proceedings of the 30th. ICRC*, volume 5, page 885, 2008.

- [79] J. Coppens et al. Analysis of Galactic Radio Background Noise. GAP-Note 123, Pierre Auger Observatory, 2007.
- [80] A. Gorgi et al. On the SD reconstruction uncertainties and trigger efficiency. GAP-Note 048, Pierre Auger Observatory, 2008.
- [81] C. Timmermans et al. - HiSPARC collaboration. Setup and first results of the HiSPARC experiment. In *Proceedings of 29th ICRC*, volume 6, pages 345–348, 2005.
- [82] A. Horneffer. *Measuring Radio Emission from Cosmic Ray Air Showers with a Digital Radio Telescope*. PhD thesis, Universität Bonn, 2006.
- [83] M. Erdmann et al. Measurement of the Lateral Distribution Function of Radio Emission from Ultra High Energy Cosmic Ray induced Air Showers. GAP-Note 119, Pierre Auger Observatory, 2010.
- [84] W.D. Apel et al. - LOPES collaboration. Lateral distribution of the radio signal in extensive air showers measured with LOPES. *Astroparticle Physics*, 32:294–303, 2010.
- [85] O. Ravel - CODALEMA Collaboration. The CODALEMA experiment. In *4th International workshop on Acoustic and Radio EeV Neutrino detection Activities - ARENA*, 2010. Slides available at <http://arena2010.in2p3.fr>.
- [86] M. Melissas, A. Schmidt, et al. Results from the 2008-2009 SALLA measurement campaigns at the BLS. GAP-Note 011, Pierre Auger Observatory, 2012.
- [87] B. Fick et al. - Pierre Auger Collaboration. The Central Laser Facility at the Pierre Auger Observatory. *Journal of Instrumentation*, 1:11003, 2006.
- [88] B. Revenu et al. - Pierre Auger Collaboration, CODALEMA Collaboration. Radiodetection of cosmic air showers with autonomous radio detectors installed at the Pierre Auger Observatory. *Nuclear Instruments and Methods in Physics Research A*, 604:S37–S40, 2009.
- [89] H.-J. Mathes - Pierre Auger Collaboration. The HEAT Telescopes of the Pierre Auger Observatory Status and First Data. In *Proceedings of the 32th. ICRC*, page 0761, 2011.
- [90] B. Wundheiler - Pierre Auger Collaboration. The AMIGA muon counters of the Pierre Auger Observatory: performance and first data. In *Proceedings of the 32th. ICRC*, page 0341, 2011.
- [91] Michiel A. Hevinga et al. Photo-Voltaic System for the Auger Engineering Radio Array. GAP-Note to be published, Pierre Auger Observatory, 2011.
- [92] Y. Zhu. AERA Wireless Comms - status report. In *Collaboration internal AERA meeting Nijmegen*, 2011.
- [93] 'The AERA group'. AERA: Proposal for the construction of the 20km² Auger Engineering Radio Array at the Southern Auger Observatory. GAP-Note 172, Pierre Auger Observatory, 2009.
- [94] H. Gemmeke et al. . A fast Embedded System for Radio Detection of Cosmic Rays. In *Proceedings of IEEE NSS*, 2011.

- [95] M. Manewald. Entwicklung und Implementierung eines Hardware-Selbsttriggers für die Radiodetektion kosmischer Strahlung. Diploma thesis, Universität Karlsruhe (TH), 2008.
- [96] A. Schmidt et al. FPGA Based Signal-Processing for Radio Detection of Cosmic Rays. *Transactions on Nuclear Science, IEEE*, 58:1621–1627, 2011. doi:10.1109/TNS.2011.2141151.
- [97] C. Rühle et al. Advanced digital self-triggering of radio emission of cosmic rays. *Nuclear Instruments and Methods in Physics Research A*, 662 suppl. 1:146–149, 2012. doi:10.1016/j.nima.2010.11.017.
- [98] F. Harris. On the use of windows for harmonic analysis with the discrete Fourier transform. In *Proceedings of the IEEE*, volume 66, pages 51–83, 1978.
- [99] J. L. Smith. Implementing Median Filters in XC4000E FPGAs. *Journal for Xilinx Programmable Logic Users, Xcell*, 23:16, 1996.
- [100] B. Hirschl, L. Yaroslavsky. FPGA implementations of sorters for non-linear filters. In *Proceedings of the XII European Signal Processing Conference EUSIPCO2004, Vienna*, 2004.
- [101] D. O. North. Analysis of the factors which determine signal/noise discrimination in radar. Report ppr-6c, RCA Laboratories, 1943.
- [102] G. Turin. An introduction to matched filters. *IRE Transactions on Information Theory*, 6:311–329, 1960.
- [103] C. Timmermans. Results from a 12 station setup using the Nikhef Digitizers. In *Collaboration internal AERA meeting Paris*, 2012.
- [104] X. Bertou. Study of time delay to SD T3 decision. GAP-Note 057, Pierre Auger Observatory, 2008.
- [105] C. Timmermans. private communication, 2011.
- [106] A. Schmidt et al. Analysis and rejection of periodic noise in radio measurements of cosmic rays. GAP-Note under review, Pierre Auger Observatory, 2012.
- [107] M. Erdmann et al. A Novel Method of Selecting Cosmic Ray Candidates from Self-Triggered AERA Events. GAP-Note 108, Pierre Auger Observatory, 2011.
- [108] J. L. Kelley et al. Design and Testing of a 5 GHz Commercial Wireless Network for AERA. GAP-Note 054, Pierre Auger Observatory, 2012.
- [109] J.W. Cooley, J.W. Tukey. An algorithm for the machine calculation of complex Fourier series. *Mathematics of Computation*, 19:297–301, 1965.
- [110] M.T. Heideman, D.H. Johnson, and C.S. Burrus. Gauss and the history of the fast Fourier transform. *Archive for History of Exact Sciences*, 34:265–277, 1985.

List of Figures

2.1. Energy spectrum of cosmic rays	4
2.2. Sky-map of ultra high-energetic cosmic rays	5
2.3. Air shower development	6
2.4. Map of the Pierre Auger Observatory	9
2.5. Simulated radio emission of a cosmic ray	10
3.1. Map of the radio experiments at the KIT	18
3.2. Pictures of antenna candidates	20
3.3. Antenna radiation patterns	21
3.4. Antenna frequency response	21
3.5. Block diagram of the LOPEs ^{STAR} measurement	22
3.6. Effect of the signal upsampling	24
3.7. Radio spectrum at LOPEs ^{STAR}	25
3.8. Sketch of the coincidence trigger condition	28
3.9. Block diagram of the prototype electronics	31
3.10. Block diagram of the prototype DAQ	33
4.1. Map of the prototype measurements at the BLS	36
4.2. Radio spectrum at the BLS	38
4.3. Radio amplitude distribution	39
4.4. Trigger rate versus threshold	39
4.5. Threshold development during the prototype measurement	40
4.6. Coincidence trigger conditions	41
4.7. Sky-map of incoming triggers	41
4.8. Periodic transients	42
4.9. Detection of galactic noise	44
4.10. Time distribution of externally triggered events	46
4.11. Expected lateral distribution of the radio amplitude	47
4.12. Expected cosmic ray detection rate	48
4.13. Noise survey at the BLS	50
4.14. Noise survey at the AERA field	51
5.1. Map of AERA	54
5.2. Block diagram of the AERA DAQ	55
5.3. Block diagram of the AERA digitizer	56
5.4. Photograph of the AERA digitizer	57
5.5. Block diagram of the AERA firmware	59
6.1. Radio spectrum at AERA	64
6.2. Effect of the RFI-suppression on an example event	65
6.3. Block diagram of the RFI-suppression algorithm	66
6.4. Comparison of different FFT-engines	69
6.5. Effect of the FFT block length	69

6.6.	Illustration of window functions	70
6.7.	Spectral leakage with window functions	70
6.8.	Example for block edge bulging	71
6.9.	Signal improvement by windowing	71
6.10.	Illustration of the overlapping	72
6.11.	Block diagram of the overlap creation	72
6.12.	Effect of the median filter on the spectrum	73
6.13.	Block diagram of a systolic array median	74
6.14.	Block diagram of a serial rank computer	74
6.15.	Comparison of software and firmware results	75
6.16.	Effect of the RFI-suppression of the random trigger rate	77
6.17.	Example for the envelope calculation	78
6.18.	Effect of multiple pulses on the median filter	79
6.19.	Frequency response of an IIR Notch filter	83
6.20.	Results of the IIR Notch filter on the SNR	83
7.1.	Threshold development at AERA	86
7.2.	T2-timestamp coincidences with the SD-events	87
7.3.	Statistics of the coincidences	87
7.4.	T2-timestamp coincidences with selected SD-events	88
7.5.	Statistics of the selected coincidences	88
7.6.	Measured pulse height versus expectation	89
7.7.	Pulse detection efficiency	89
7.8.	Cut on multiple pulses	91
7.9.	The local signal-to-noise-ratio	92
7.10.	Cut on the local SNR	92
7.11.	Cut on the pulse width	92
7.12.	Cut on the pulse-to-afterpulse-ratio	92
7.13.	Periodic time difference between triggers	93
7.14.	Frequency spectrum of the trigger timestamps	93
7.15.	Winding frequency of T2-triggers	94
7.16.	Suppression of periodic noise based on individual events	94
7.17.	Block diagram of the PLL	95
7.18.	Result of the PLL prediction	95
7.19.	Long term operation of the PLL	97
7.20.	Threshold during the PLL operation	97
7.21.	Phase comparison between different PLLs	98
7.22.	Periodic signal at different thresholds	98
7.23.	Block diagram of the constant fraction discriminator	99
7.24.	Envelope of an example pulse and its CFD-signal	99
7.25.	Improvement of the time resolution with CFD-trigger	99
7.26.	Improved angular resolution with the CFD-trigger	100
7.27.	Coincidence trigger rate versus accepted zenith angle	101
7.28.	False trigger rate and effective aperture of the dynamic T3	102
7.29.	Sky-map of the rejection pattern of the dynamic T3	103
7.30.	Sky-map of the T3-remaining background triggers	103
A.1.	Block diagram FIR filter	114
A.2.	Block diagram IIR filter	115

List of Tables

6.1.	Table of FFT implementations	68
6.2.	Table of median filter implementations	74
6.3.	Results of the RFI-suppression	76
6.4.	Comparison of different Filter methods	80
6.5.	Table of the matched filter performance	81
7.1.	Performance of the periodic veto for varying threshold levels.	98
7.2.	Average coincidence trigger performance	103
A.1.	Table 1 of self-triggered cosmic ray events	118
A.2.	Table 2 of self-triggered cosmic ray events	118