

Extending the Physics Reach of the Pierre Auger Observatory using Low-Level Trigger Data

Zur Erlangung des akademischen Grades eines
Doktors der Naturwissenschaften (Dr. rer. nat.)

von der KIT-Fakultät für Physik des
Karlsruher Instituts für Technologie (KIT)

und der

Universidad Nacional de San Martín (UNSAM)

genehmigte

Dissertation

von

M. Sc. Martin Leo Schimassek

aus Reutlingen

Tag der mündlichen Prüfung: 17.12.2021

Referent: Prof. Dr. Ralph Engel

Korreferent: Prof. Dr. Xavier Bertou

Betreuer: Dr. Darko Veberič



This document is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0): <https://creativecommons.org/licenses/by/4.0/deed.en>

Extending the Physics Reach of the Pierre Auger Observatory using Low-Level Trigger Data

For the attainment of the academic degree of

Doctorate in Science

from the

Karlsruher Institut für Technologie (KIT)

and the

Universidad Nacional de San Martín (UNSAM)

accepted

Dissertation

of

M. Sc. Martin Leo Schimassek

of Reutlingen

Day of the oral examination: 17.12.2021

Referee: Prof. Dr. Ralph Engel

Co-referee: Prof. Dr. Xavier Bertou

Supervisor: Dr. Darko Veberič

Extending the Physics Reach of the Pierre Auger Observatory using Low-Level Trigger Data

Tesis aceptada para optar por el título de

Doctor en Astrofísica

del Instituto de Tecnología "Prof. Jorge A. Sabato" de la
Universidad Nacional de San Martín (UNSAM)

y del

Karlsruher Institut für Technologie (KIT)

por

M. Sc. Martin Leo Schimassek

de Reutlingen

Fecha de la defensa oral: 17.12.2021

Director: Prof. Dr. Ralph Engel

Co-director: Prof. Dr. Xavier Bertou

Colaborador: Dr. Darko Veberič

Abstract

The Pierre Auger Observatory, the world's largest cosmic-ray observatory, has significantly changed the understanding of ultra-high energy cosmic rays during its almost 20 years of operation. However, the sources of the highest energy particles are still unknown and the findings have raised (even) more questions including the surprising observation of heavy nuclei at high energies and deficits in describing the muonic data in simulations. In this thesis, I reach beyond the usually-used event data to explore more possibilities and to shed light onto the remaining questions in cosmic-ray science. I present the first-ever large-scale analysis of data collected on a trigger level of the Auger surface detector.

In searching for signals from neutrons, appearing as delayed particles, I open a new observation channel for hadronic interactions that can help identify the sources of discrepancy between data and currently available models. While delayed particles, presumably neutrons, have been observed in scintillators before, this thesis is the first to analyse and explicitly search for their signal in the standard surface detectors of Auger. The found signals, reported in this thesis, might make a reevaluation of the detector simulations necessary for any future analysis and, as such, highlight the importance of the provided clarification of the existence of neutron signals in the surface detector data.

Another possibility to provide insight on the nuclear mass of cosmic rays at high energies is to use the Gerasimova-Zatsepin (GZ) effect. It predicts that a small fraction of cosmic-ray nuclei at high energies are disintegrated by solar photons creating coincident cosmic-ray events if the geometry is favourable. However, the predicted rates are small and no compatible events have been observed. With the unique opportunity created by the available non-event data set used in this thesis, I extend the searched phase space to regions with small – but observable – predicted GZ-event rates. With a significant deviation from background, I find indications that these events are present in the recorded Auger data. Future analyses triggered by this result have to evaluate whether the observed events are compatible with standard GZ events and what they imply for the mass composition.

Additionally, we can use the same data to try to find the origin of lightning by understanding the possible role of cosmic rays in lightning initiation. I show that the Auger surface detector trigger data is sensitive to showers in the relevant energy range and above. After considerations of the associated background, I show the results of the first large scale analysis correlating cosmic-ray event candidates with detected lightning strikes.

During thunderstorms unusual events have been observed in the detectors, hinting at high-energy gamma-ray emission during thunderstorms. I provide a first explanation of the unexpected ring-shaped topology using the trigger data and develop a new algorithm to increase the detection efficiency of these events so that future work can rely on much larger statistics than the currently available pool of only about 20 events.

Zusammenfassung

Das Pierre-Auger-Observatorium, das weltgrößte Observatorium für ultrahoch-energetische kosmische Strahlung, hat das Verständnis derselben erheblich verändert. Trotz beachtlichen Fortschritts haben die Ergebnisse viele weitere Fragen aufgeworfen, und die Quellen der energiereichsten Teilchen sind nach wie vor unbekannt. Zu den neuen Rätseln gehören unter anderem die unerwartete Beobachtung schwerer Kernen bei hohen Energien, sowie Defizite der Simulationen in der Beschreibung der Muon-Daten. In dieser Arbeit nutze ich Daten, die über die üblichen Ereignisdaten hinaus gehen, um weitere Möglichkeiten zu erforschen, die offenen Fragen der kosmischen Strahlung zu beantworten.

Mit der Suche nach zeitverzögerten Signalen von Neutronen, eröffne ich einen neuen Beobachtungskanal für hadronische Wechselwirkungen, der dazu beitragen kann, die Ursachen der Diskrepanz zwischen Daten und den derzeit verfügbaren Modellen zu identifizieren. Während verzögerte Teilchen schon früher in Szintillatoren beobachtet wurden, analysiere ich solche Signale in Wasser-Tscherenkov-Detektoren. Die hier gefundenen Signale machen eine Überprüfung der Detektorsimulationen für künftige Analysen erforderlich und unterstreichen somit die Bedeutung dieser Arbeit.

Um einen Einblick in die Massenverteilung der kosmischer Strahlung bei hohen Energien zu erhalten, ist die Nutzung des Gerasimova-Zatsepin-(GZ-)Effekts möglich, der vorhersagt, dass ein Teil der Kerne der kosmischen Strahlung bei hohen Energien mit Photonen der Sonne wechselwirken und bei günstiger Geometrie zu messbaren zeitgleichen Ereignissen führt. Jedoch sind die vorhergesagten Raten klein und GZ-Ereignisse wurden noch nie beobachtet. Mit der Gelegenheit, die sich durch den in dieser Arbeit verwendeten neuen Datensatz ergibt, erweitere ich den erkundeten Phasenraum. Ich finde signifikante Hinweise darauf, dass GZ-Ereignisse in den aufgezeichneten Auger-Daten vorhanden sein könnten. Zukünftige Analysen müssen bewerten, ob die hier beobachteten Ereignisse mit der GZ-Hypothese kompatibel sind und was ihre Bedeutung im Detail bedeutet.

Darüber hinaus können dieselben Triggerdaten verwendet werden, um die Entstehung von Blitzen zu verstehen, indem die mögliche Rolle kosmischer Strahlung bei der Blitzentstehung untersucht wird. In dieser Arbeit zeige ich, dass die Triggerdaten des Auger-Oberflächendetektors kosmische Strahlung im relevanten Energiebereich messen können, und präsentiere die erste Korrelationsanalyse auf diesen Triggerdaten.

Schon zuvor wurden in Gewittern ungewöhnliche Ereignisse beobachtet, die auf hoch-energetische Gammastrahlung während Gewittern hinweisen. Ich liefere eine Erklärung für die unerwartet-ringförmige Topologie der Ereignisse und entwickle einen neuen Algorithmus, um die Effizienz der Datennahme für solche Ereignisse zu erhöhen. So können künftige Arbeiten auf eine größere Ereignisanzahl, derzeit etwa 20, zurückgreifen um den Ursprung dieser Strahlung zu ergründen.

Resumen

El Observatorio Pierre Auger, el observatorio de rayos cósmicos más grande del mundo, ha profundizado significativamente nuestro conocimiento en el área de los rayos cósmicos de ultra alta energía. Sin embargo, aún se desconocen las fuentes de las partículas de mayor energía, y de los descubrimientos han surgido aún más preguntas. Éstas incluyen la sorprendente presencia de núcleos a altas energías, así como déficits en las simulaciones para describir los datos. En esta tesis, exploro posibilidades por fuera de los datos habituales para arrojar luz sobre las incógnitas pendientes. Aquí presento el primer análisis de los datos recogidos a nivel de trigger del detector de superficie del Observatorio Pierre Auger.

Al buscar señales de neutrones, que aparecen como partículas tardías, abro una nueva ventana de observación de las interacciones hadrónicas, que puede ayudar a identificar las fuentes de discrepancia entre los datos y los modelos actuales. Aunque las partículas tardías, supuestamente neutrones, se han observado antes en centelleadores, este análisis es el primero que analiza y busca explícitamente su señal en los detectores Cherenkov de superficie de Auger. Las señales aquí encontradas prueban necesaria una revisión de la simulación de los detectores para cualquier análisis futuro, evidenciando la relevancia de este análisis.

Otra posibilidad de aportar información sobre la composición de los rayos cósmicos a altas energías es utilizando el efecto Gerasimova-Zatsepin. Éste predice que algunos núcleos a altas energías son desintegrados por fotones solares, creando eventos de rayos cósmicos coincidentes si la geometría es favorable. Sin embargo, las tasas predichas son muy pequeñas y nunca se ha observado ningún suceso de este tipo. Con la oportunidad única dada por el conjunto de datos subumbrales utilizado en esta tesis, extendiendo el espacio de fase de búsqueda a regiones donde las tasas de eventos GZ predichas son pequeñas pero observables. Con una desviación significativa respecto del fondo, encuentro indicios de que estos eventos están presentes en los datos de Auger. A raíz de este resultado, futuros análisis deberán evaluar si los eventos observados son compatibles con eventos GZ estándar y lo que implican para la composición de los rayos cósmicos.

Además, utilizo los mismos datos subumbrales para tratar de entender el posible papel de los rayos cósmicos en la iniciación de relámpagos. Muestro que los datos de trigger son sensibles a las lluvias atmosféricas a las energías relevantes y superiores, y presento el primer análisis de la correlación entre los eventos de rayos cósmicos y los relámpagos detectados.

Asimismo, se han observado eventos inusuales que sugieren posibles emisiones de rayos gamma de alta energía durante tormentas eléctricas. Utilizando los datos de trigger, proporciono una primera explicación de la inesperada topología de los eventos con huellas en forma de anillo, y desarrollo un nuevo algoritmo para aumentar la eficiencia de detección de estos eventos. Consecuentemente, amplíé significativamente la estadística disponible (actualmente de unos 20 eventos), que servirá de base para los trabajos futuros.

Contents

1	Introduction	1
2	Physical Background	3
2.1	Cosmic Rays	3
2.1.1	Extensive Air Showers	4
2.1.2	Detection of Cosmic Rays	6
2.1.3	Acceleration and Possible Sources	6
2.1.4	Propagation of Cosmic Rays and Multi-Core Events	7
2.1.5	Solar Modulation	9
2.2	Physics of Thunderstorms	10
2.2.1	Lightning Strikes	11
2.2.2	Detection of lightning strikes	12
2.2.3	Terrestrial Gamma-Ray Flashes and Other Phenomena	13
3	The Pierre Auger Observatory	15
3.1	The Fluorescence Detector	15
3.2	The Surface Detector	17
3.2.1	The SD Station Hardware	17
3.2.2	The SD Trigger System	18
3.2.3	Calibration of the Surface Detector Stations	21
3.2.4	Scalers	23
3.2.5	Reconstruction of Ultra-High Energy Cosmic Rays	23
3.3	Upgrades	24
4	Analysis of Scalers	27
4.1	Improvements of the Data Processing	27
4.1.1	Identification of Bad Stations	28
4.1.2	Station Outliers	29
4.1.3	Lightning and Localised Excesses	31
4.1.4	Averaging and Scaling	33
4.1.5	Applied Corrections for Long-term Trends	35
4.1.6	Baseline Correction	37
4.2	Results	39
4.2.1	Forbush Decreases	40
4.2.2	Frequency Spectrum	41
4.2.3	The Diurnal Signal	43
4.2.4	Observation of the Solar Cycle	44

4.3	Summary and Outlook	45
5	Surface Detector Events associated with Lightning	47
5.1	Observations of SD-Rings and SD-Discs	47
5.2	Identification of Trigger Effects on Lightning Events	49
5.2.1	Available Data	49
5.2.2	Station Dead Time	50
5.2.3	The used Events	52
5.2.4	Analysis Strategy	53
5.2.5	Results	59
5.3	Information from Delayed Events	64
5.3.1	Data Treatment	64
5.3.2	Observations	66
5.3.3	Recovering the Signals	69
5.3.4	Concluding Remarks	73
5.4	Modification to CDAS	74
5.4.1	A Concept for Prioritized Read-Out	74
5.4.2	Events to Detect	75
5.4.3	Evaluation of Algorithms	76
5.4.4	Investigated Algorithms	78
5.4.5	Implementation in the UB	83
5.5	Possibilities with the Upgraded Electronics	89
5.5.1	Input Data	89
5.5.2	Dedicated Triggers	91
5.5.3	Extended Traces	93
5.5.4	Read-Out Adjustments	97
5.6	Search for Correlation of Cosmic Ray Showers and Lightning	98
5.6.1	Sensitivity of SD-T2s to Low Energy Events	98
5.6.2	Lightning Events	103
5.6.3	Characterisation of the Data Set	104
5.6.4	Correlation with Station Triggers	107
5.6.5	Correlation with Doublets	109
5.6.6	Correlation with Trigger Pairs	110
5.7	Summary	112
6	Correlations of SD Triggers	115
6.1	Characterisation of the Data Set	115
6.1.1	Time-Difference Spectra	116
6.1.2	Identification of Unusual Stations	116
6.2	Temporal Correlations	119
6.2.1	Time Difference Spectra	119
6.2.2	Expectations for Delayed Triggers	120
6.2.3	Statistical Search for Delayed Signals	121
6.2.4	Event Correlation Search for Delayed Signals	124
6.2.5	Effects of the Station Electronics	127
6.2.6	Prospects of Detecting Very Late Particles	129
6.2.7	Detection of Pulses in Event Traces	130
6.2.8	Late Particles in SD	131
6.2.9	Late Particles in SSD	135
6.2.10	Discussion	137
6.3	Searching for Multi-Core Events	139

6.3.1	Spatial Clustering Methods	141
6.3.2	Elementary Reconstruction Using Triggers	141
6.3.3	Checking CDAS: Comparison with Recorded Events	143
6.3.4	Multi T3s	144
6.3.5	Correlation of Events with Sub-Threshold Candidate Events	145
6.3.6	Search for Sub-Threshold GZ Events	154
6.3.7	XAS: eXtended Air Showers	161
6.4	Summary	165
7	Summary and Conclusions	171
A	Additional Information on SD-Rings	177

Introduction

Cosmic-rays at the highest energies have always been mysterious because of their incredible energies and the lack of knowledge about the sources. The facilities necessary to measure these particles at the highest energy are also impressive in size and sensitivity. For instance, the Pierre Auger Observatory [1], whose data we use in this work, covers an area of 3000 km² in the Argentinian pampa with more than 1600 detector stations of which each is able to detect more than 2000 of atmospheric muons per second. Despite its impressive size, the data collected so far have not been sufficient to answer multiple questions that arose ultra-high energy cosmic-ray research. Partially, this is due to the fact that the obtained data raised more questions than it answered, for example an increasingly heavy mass composition at high energy was found [2] in tension with the previous expectation of protons. This also means that the predicted flux suppression from photo-pion production of protons on the cosmic microwave background might not be the cause of the strong cutoff found in the energy spectrum [3]. However, composition measurements at the highest energies are still lacking statistics which motivates a currently on-going upgrade of the Auger Observatory [4].

A historic idea how the mass-composition of cosmic-rays can be measured without the need for simulations and hadronic-interaction modelling is the Gerasimova-Zatsepin (GZ) effect [5]. Nuclear disintegration in the photon field of the Sun can create coincident pairs of cosmic-ray events whose energies are linked to the primary mass. However, the expected rate and separation for these events is smaller than commonly thought accessible because of the large area necessary at low energy thresholds. Thus, no measurement of this effect has been reported so far.

In this work, we will explore the possibilities of using the data collected at the Pierre Auger Observatory differently than done previously. Using the low-level trigger data, we can lower the energy threshold substantially and thus try to search for these GZ events effectively. If we were to find a signal, or even discover that the background for such searches is low and on the order of the predicted signals, a new window into measuring the mass composition of cosmic rays might be opened.

Closely linked to the problem of mass composition at the highest energies are the hadronic-interaction models that are used to describe the particle production at high energies in simulations. An observed discrepancy in the muon number with data is not resolved yet [6]. In a complementary way of investigating the relevant hadronic interactions, also neutrons, usually not recorded because of large time delays, could provide important information. Using the trigger data itself instead of the event data, we can explore the possibility of studying such late particles in this analysis and open another window into hadronic interactions.

Additionally, in other fields long-standing questions are still open. One of the fields, lightning physics, is linked to cosmic rays with the important question on how a lightning

strike initiates. Free electrons from extensive air showers were proposed as suitable candidates multiple times [7]. Observations during thunderstorms are naturally difficult, including the observation of lightning itself as it is often – if not always – obscured by the cloud and so clear observations that could confirm or reject such hypotheses are sparse. Observationally, it is also not clear if the (small but dense) cosmic-ray facilities used at the energies (likely) relevant to the lightning problem have sufficient aperture to observe the correlation due to geometric effects. So, in a similar way as for the GZ events, in this work we want to make use of the large exposure of the Pierre Auger Observatory to contribute to the lightning initiation problem. While clearly the usual energy threshold of Auger is too high, we can use the same low-level trigger data as for the GZ search and make use of the fact that single detector stations of the Pierre Auger Observatory are also sensitive to low-energy showers.

That cosmic-ray observatories have the power to discover and observe effects related to thunderstorms has been shown in the recent past by the observation of ELVES, light emission in the ionosphere caused by the electromagnetic pulse from lightning, with the Auger Fluorescence detectors [8]. But also the Telescope Array surface detector reported the observation of a special, lightning-related phenomenon: high-energy gamma-ray emission preceding the lightning strike, a Terrestrial Gamma-Ray Flash (TGF), was observed [9]. In the Auger surface detector (SD) data peculiar events during thunderstorms were also observed [10], but of unknown origin and with topologies hinting at instrumental effects. Thus, we can also use the additional data from the triggers to investigate such periods and arrive at a conclusive statement whether such types of events, possibly TGFs, are observed in the Auger SD.

The analysis of all these effects with the trigger data can also help to develop new strategies to increase the efficiency of data taking. This applies to special thunderstorm events, as well as potentially sub-threshold GZ events. We explore such possibilities in this work, especially in light of new capabilities in the Auger SD station electronics available after the AugerPrime upgrade.

In the following Chapters, we introduce the physics relevant for extensive air showers, the phenomenon used to detect cosmic rays at the Pierre Auger Observatory, introduce development of lightning as well as the details of the detectors used. In the presented analyses, we focus on data not used in most standard Auger analyses and move up the threshold levels in successive Chapters. We start with the lowest-threshold counters, and show how sensitive they are to the rate of cosmic rays and then move to the trigger data to extend the knowledge about peculiar lightning events. After the correlation analysis with lightning strikes, we finally explore the possibilities with the lowered-energy threshold using trigger data in addition to the event data with respect to cosmic-rays. We search for delayed particles to constrain hadronic interactions, and perform the first dedicated search for GZ events in Auger.

Physical Background

In this Chapter, we briefly summarise the basic principles behind the measurements done in cosmic-ray physics and this work in particular. We include background information about cosmic rays and their propagation and sketch the detection techniques used. Additionally, we add a brief introduction to lightning strikes as introduction to the topic of lightning initiation and its link to cosmic rays that is investigated in this work. We start with a general explanation of cosmic rays and then develop further detailed aspects relevant to this work in separate Sections.

2.1 Cosmic Rays

Cosmic rays, particles of high energy impinging on Earth's atmosphere, were discovered more than a century ago as ionising radiation [11]. Today, we know that also electrons and photons are present in the flux of astrophysical particles arriving at Earth, but the majority are ionised nuclei and mostly protons. Fig. 2.1 shows the energy spectrum of particles above 10^{13} eV which can be mostly characterised as a power law $J \propto E^{-3}$ with inflection points marked in the Figure. The lowest-energy cosmic rays measured directly above the atmosphere have an energy of about 1 GeV and the measured spectrum spans to 10^{20} eV covering 11 orders of magnitude [12].

The highest energy part of the spectrum above about 10^{18} eV is known as ultra-high energy cosmic-rays (UHECR) and is primarily targeted by the Pierre Auger Observatory whose data we use in this work. In this energy range, the spectrum shows three significant features deviating from the pure power-law spectrum [3]. The most notable is the so-called "ankle", a steepening at about 6×10^{18} eV in the energy spectrum, and a flux suppression at the highest energies. In terms of arrival directions, the sky at ultra-high energies is isotropic to a percent level, with the only deviation beyond 5σ reported so far by the Pierre Auger Observatory being a dipolar structure at the amplitude level of 6.5% for events with energy above 8 EeV [13]. While at low cosmic-ray energies (up to ~ 100 TeV) direct detection with balloons and satellites (e.g. AMS-02 [14]) makes the determination of the mass composition of the cosmic rays possible, at ultra-high energies the measurement of the shower maximum (cf. next section) can only provide indirect information. The current understanding of the data suggests mostly protons at about 10^{18} eV and an increasingly heavy composition at higher energies [15, 16]. We shortly summarize proposed models and source candidates for these UHECRs in Section 2.1.3. For a more in-depth review of the current state of UHECRs see Refs. [16] and [17].




Figure 2.1: Energy spectrum of cosmic rays measured with air-shower detectors. Taken from Ref. [18]

At the other end of the spectrum, the Sun is modifying the observable cosmic-ray flux at energies below about 100 GeV. We discuss effects associated with transient events of the Sun and the modifications of the spectrum briefly in Section 2.1.5.

2.1.1 Extensive Air Showers

Because of the low flux, the measurement of cosmic rays at high energies requires the indirect detection via extensive air showers (EAS), which are large particle cascades in the atmosphere. Therefore, we briefly introduce concepts and naming used for describing EAS.

The concept of an air shower, or particle cascade in general, can be most easily understood in electromagnetic cascades initiated by gamma rays. Ignoring minor contributions of e.g. photo-nuclear interactions, the cascade is governed by the processes of bremsstrahlung and pair creation in the field of the air nuclei. In a simple model by Heitler [19], illustrated in the left sketch of Fig. 2.2, interactions occur after traveling an interaction length $\ell_{\text{em}} = X_0 \ln 2$, where X_0 is the radiation length, after which the particle number n is doubled and the energy is split in half. Expressing the length in terms of traversed column density $X = \rho d$, typical values are $\ell_{\text{em}} \approx 25 \text{ g cm}^{-2}$. In this way, the number of particles grows exponentially until the energy of the particles in generation n reaches the critical energy $E_c \approx 84 \text{ MeV}$ when energy loss due to ionisation and interactions with atomic electrons become relevant. The depth corresponding to this maximum of particles $N_{\text{max}} = E_0/E_c$, is $X_{\text{max}} = n_{\text{max}} \ell_{\text{em}} = X_0 \ln(E_0/E_c)$ where the initial energy of the primary particle is E_0 [12].

A simple extension of this model to hadronic showers is given with the Heitler-Matthews model [20]. The right sketch in Fig. 2.2 illustrates the idea. Analogous to the electromagnetic case, particles interact after one interaction length λ_π and create n_{tot} (multiplicity) new particles with equal energy splitting. In hadronic interactions, mostly pions are created with an ratio of 2:1 of charged to neutral pions. The neutral pions decay promptly and fuel the electromagnetic cascade as shown in Fig. 2.2. Typical values of n_{tot} are tens to hundreds at




Figure 2.2: Illustration of particle cascades as they happen in air showers. On the *left* an electromagnetic cascade characterised by the radiation (doubling) length l_{em} . On the *right*, a hadronic cascade with hadronic, muonic, and electromagnetic components. Taken from Ref. [12]

the very high energies and the typical interaction length of pions $\lambda_{\pi} \approx 120 \text{ g cm}^{-2}$ is larger than for the electromagnetic case. Also in this case, the cascade stops when reaching a critical energy, which is however governed by the decay of pions rather than ionisation losses and leads to $E_d \approx 30 \text{ GeV}$ [12].

The decaying pions form the muonic and neutrino component that in reality are also fed by other hadronic decays which are neglected in this simple model. Using the fact that about one third of the energy is transferred to the electromagnetic component in each generation, we can get a simple estimation of the depth of shower maximum as

$$X_{\text{max}} \approx X_1 + X_0 \ln \left(\frac{E_0}{2 n_{\text{tot}} E_c} \right) \quad (2.1)$$

with the depth of the first interaction X_1 and the average energy of neutral pions in the first generation E_0/n_{tot} . If we further assume that for a nucleus the binding energy of nucleons is much smaller than the scale of the interaction, we can use the *superposition model* to infer properties of a shower initiated by a nucleus of mass number A . The average shower maximum shifts to shallower values using $E_0 \rightarrow E_0/A$ in Eq. (2.1). For the number of muons, in this simplified model taken as number of pions at critical energy E_d after n_d generations, we get

$$N_{\mu}^A = A^{1-\alpha} \left(\frac{E_0}{E_d} \right)^{\alpha} \quad (2.2)$$

with $\alpha = \ln(n_{\text{ch}})/\ln(n_{\text{tot}})$ and the number of charged pions per interaction n_{ch} . In the central part of the shower, there is also a remaining hadronic component that is not discussed here as it is not relevant for most observations [12].

Modern simulations are carried out with dedicated simulation codes like CORSIKA [21] or AIRES [22]. These make use of models of the hadronic interactions at the highest energies that are tuned to describe the available accelerator data and extrapolate to the inaccessible phase-space and energies necessary for these simulations. The most prominent of these models are EPOS-LHC [23], QGSJetII-04 [24], and Sybill 2.3d [25].

2.1.2 Detection of Cosmic Rays

The way cosmic rays are detected depends crucially on the energy. At low energies (up to several TeV) the fluxes are large enough so that direct detection with particle physics like detectors in balloons or in satellites is possible. Examples of such experiments on balloons are CREAM [26], PAMELA [27], and DAMPE [28] as free flying satellites, and AMS-02 onboard the ISS [14].

For higher energies, the atmosphere is used as a calorimeter using extensive air showers for the detection. Multiple ways of detecting the cascades have been developed ranging from particle detection on the ground to the emission of radio pulses from the deflection of electrons in the geomagnetic field. For particle detection, the most frequently used detector types are scintillators and water-Cherenkov detectors. We describe in Section 3.2.5 the techniques used to evaluate signals from such detectors for the reconstruction of the primary energy on an example of the surface detector of the Pierre Auger Observatory.

The particles in the cascade also emit Cherenkov light directly in air which can also be used for detection. Imaging telescope like HESS [29] use this technique while mostly targeting gamma rays rather than hadronic showers. Also non-imaging techniques are used for example in TUNKA [30].

The energy deposition of charged particles from the cascade in air also leads to emission of UV fluorescence light from excited nitrogen molecules. Several observatories use this emission for the detection of air showers, among them the Telescope Array [31], the largest observatory in the northern hemisphere, and the Pierre Auger Observatory [1], the largest observatory in the world. We describe in detail the Pierre Auger Observatory in Chapter 3 as it is extensively used in this work.

2.1.3 Acceleration and Possible Sources

In this Section, we briefly summarise acceleration mechanism of charged particles in astrophysical sources. To accelerate particles, energy needs to be transferred to charged particles. Fermi [32] proposed a mechanism to use macroscopic energy from moving magnetised plasma for this conversion. A charged particle colliding with a moving magnetised “cloud”, scattering elastically on magnetic irregularities, on average gains energy when averaging over all possible incoming and outgoing directions. After n such collisions, starting with an initial energy of E_0 , the average resulting energy is

$$E_n \simeq E_0 \left(1 + \frac{4}{3}\beta^2\right)^n \quad (2.3)$$

with the velocity of the cloud $v = \beta c$, i.e. scaling quadratically in β . To get a more efficient acceleration mechanism, the randomly moving clouds can be replaced with a shock front which the particle can cross multiple times. There the energy scales as

$$E_n \simeq E_0 \left(1 + \frac{4}{3}\beta\right)^n. \quad (2.4)$$

To get the resulting energy spectrum of particles, we can assume a constant escape probability per encounter p and obtain a power law

$$N(>E) \propto \frac{1}{p} \left(\frac{E}{E_0}\right)^{-\gamma} \quad (2.5)$$

with the spectral index

$$\gamma \simeq -\ln(1-p) / \ln\left(1 + \frac{4}{3}\beta\right) \quad (2.6)$$




Figure 2.3: Hillas plot showing the sizes and average magnetic field strengths of classes of objects. The Hillas condition [34] (cf. Eq. (2.7)) is a necessary condition for acceleration of cosmic rays and is visualised as lines of maximal energy in this diagram. Taken from Ref. [16]

in the right ballpark necessary to describe the cosmic-ray spectrum if transport effects are taken into account [33].

The maximal energy of acceleration in such magnetised regions is governed by the size of the source and the magnetic field providing the scattering centres. A necessary condition for the maximal energy of a certain type of source is given by the Hillas condition [34]

$$E_{\max} \propto \beta Z e B R, \quad (2.7)$$

where Z is the charge number of the cosmic rays, B the magnetic field strength, and R the size of the source [33].

Fig. 2.3 shows the sizes and typical magnetic-field strengths for a number of source candidates for ultra-high energies. Keeping in mind that Eq. (2.7) is only a necessary condition, it is clear that it is not straightforward to identify one out of a few classes of objects as prime candidates for UHECR acceleration. The published indications of correlations of UHECRs with gamma-ray sources, e.g. in Ref. [35] with starburst galaxies, are hard to explain in theoretical models of the corresponding acceleration sites (see e.g. Ref. [36, 37]). More work on both experimental and theoretical side is necessary to finally identify the sources of UHECRs.

2.1.4 Propagation of Cosmic Rays and Multi-Core Events

To arrive from an (unknown) sources to the observer at Earth, UHECRs have to pass through the intergalactic space and the galaxies. On this path, deflections in magnetic fields and interactions can modify the energy and direction. Here, we discuss mainly two effects of relevance for UHECRs: the cut-off at the highest energy and the effects of the solar radiation field.

Greisen [38], and Zatsepin and Kuzmin [39] predicted a suppression (GZK suppression) of the energy spectrum for protons at the highest energies due to photo-pion production,



with the photons of the cosmic microwave background providing the target photon population. The threshold for this process is $E_p^{(\text{th})} \sim 7 \times 10^{19}$ eV but due to the black-body spectrum of the photons it is a smooth transition. Similarly, electron-positron pairs can be created in a similar reaction, referred to as Bethe-Heitler pairs, that however lead to a lesser loss of energy per reaction. For nuclei the principle is the same, but the main process is photo-disintegration of the nucleus rather than photo-pion production, which is however still possible. For these nuclear reactions also other photon fields are of relevance, e.g. the infrared background. Numerically, the suppression for heavy nuclei happens to be very similar in threshold energy and resulting range of about 100 Mpc at 10^{20} eV as that of protons [33].

A similar effect, happening however close to the Sun, using solar photons as target rather than the cosmic microwave background, was proposed by Gerasimova and Zatsepin (GZ) [5]. These GZ events are characterised by two simultaneous events observable at Earth – if the primary cosmic ray is a nucleus – because the nucleus can emit a nucleon via photo-nuclear disintegration close to Earth. Due to the strong boosting the initial direction of both fragments is identical. Because the ratio of charge to mass will in general be different for the two fragments, the interplanetary magnetic field causes separation of the two events. If the two fragments are measured with energies E_1 and E_2 then the mass of the primary particle can be estimated as $A \approx (E_1 + E_2)/E_1$, where E_1 is the smaller fragment [40].

The expected rate of these events depends on the primary composition, the spectrum, and the (known) cross section for photo disintegration and was previously calculated, cf. Refs. [40–43]. In Fig. 2.4 we show the result of one of these calculations as integral spectrum, highlighting that the probability of such an interaction is of the order of 10^{-5} at the highest energies.

This probability can be estimated using the mean free path of a cosmic ray arriving from direction $\hat{\xi}$ and distance $|\xi|$,

$$\frac{1}{\lambda} = \int_0^\infty d\epsilon \frac{dn}{d\epsilon}(\xi) \sigma(\epsilon) (1 + \beta \cos \alpha) \quad (2.9)$$

where α is the angle between photon (emitted radially from the Sun) and the cosmic ray. The integrated splitting probability is then

$$\eta_{\text{GZ}} = 1 - \exp\left(-\int_0^\infty d\xi \frac{1}{\lambda(\xi)}\right) \quad (2.10)$$

resulting numerically in the values $\ll 10^{-3}$ [40].

For experiments, the distribution of separations is important because it is mainly responsible for what could be detected and what not. So far, no experimental evidence for this effect exists and searches were mostly limited to groups of small detector arrays [44, 45] searching for correlations over very large distances. Fig. 2.5 shows the distribution of separations δ at Earth for iron primaries and three different directions. It is clear that the highest rate is obtained close to the Sun with the largest photon density. However, the separation is unfortunately large so that it is experimentally next to impossible to measure it. Even if multiple small clusters of detectors, size scale $\ell \sim 100$ m separated by large distances, e.g. $r \sim 150 \text{ km}^1$, are used, the geometrical detection probability $p(r) \propto \ell/2\pi r \approx 10^{-4}$ is very small.

¹Values inspired by the long baseline used in the Large Area Air Shower experiments [44].




Figure 2.4: Spectrum of GZ events as function of primary particle energy. For the dashed line a separation less than the Earth’s diameter is required. From the ratio to the total cosmic-ray flux (lower panel), the efficiency for the splitting of nuclei in this energy range can be read off as $10^{-8} < \eta < 10^{-4}$. Taken from Ref. [42].

Thus, counter-intuitively it might be better for experiments to use the night side, i.e. a direction away from the Sun to limit the separation. At higher energies, the separation scales like δ/E such that the day side might be of interest again [42]. To give a scale of the expected rates and the necessary apertures, we give a number for the required aperture of a detector to detect one GZ event per year if both fragments are to be detected above the energy threshold. At $E_0 = 10^{17}$ eV this results for a detector located at 45° latitude and a minimal separation $\delta_{\min} = 1$ km in a required aperture of $A_{\min} \approx 2.5 \times 10^5$ km² sr [42]. For comparison, the Pierre Auger Observatory (cf. Chapter 3) the world’s largest cosmic-ray observatory has an aperture of about 9×10^3 km² sr and a full efficiency threshold of $10^{18.3}$ eV [46] showing how difficult it is to detect these events.

2.1.5 Solar Modulation

Moving from special propagation effects at the highest energies to the lowest energies, we discuss here the influence of the Sun on the spectrum of cosmic rays at low energies. In transient events, the Sun injects particles with energies up to $\sim 10^{10}$ eV [47] but is also responsible for a continuous stream of charged particles called “solar wind”, carrying with them magnetic field lines that create the heliosphere [48]. But beyond being the source of low-energy particles, the magnetic field of the heliosphere has influence on the “usual” low-energy cosmic rays. In 1937, Forbush [49] discovered that transient magnetic storms decrease the rate of cosmic rays. Fig. 2.6 shows a typical example of such a “Forbush-decrease” that is characterised by a fast decrease (several hours) and a gradual recovery (several days) of the rate of low-energy cosmic rays. The drop in rate is caused by the magnetic field encapsulated by the plasma produced during coronal mass ejections (CMEs) at the Sun. The first drop




Figure 2.5: Distribution functions of separations between fragments of GZ-events for different directions. The coordinate system is centred at Earth and the sun is at (0,0) in latitude and longitude ϕ . ϕ is positive to dusk. Three different directions are shown with their relative normalisation to scale, all calculated with iron primaries at $E_0 = 6.3 \times 10^{17}$ eV. Taken from Ref. [41].

in Fig. 2.6 is caused by the turbulent magnetic field behind the shock front of the ejected material. The second one corresponds to the ejecta hitting the Earth [50].

Another class of transient events related to the Sun are related to the solar cosmic rays mentioned in the beginning. During “ground-level enhancements” (GLE) the measured rate increases due to energetic particles emitted by the sun. Thus, in contrast to Forbush decreases, these GLEs are not linked to effects on the galactic cosmic rays [47] but rather are a separated class of particles. These particles are accelerated by the transfer of magnetic energy to kinetic energy in magnetic reconnection events [51]. These magnetic reconnection events can cause both the aforementioned CMEs and the acceleration of particles that can lead to bursts of electromagnetic radiation, solar flares, visible from radio (synchrotron) to gamma (bremsstrahlung) energies. As such, these flares and the CMEs often happen in coincidence but are not causing each other [47, 51].

The frequency of these transient events, but also the overall heliosphere, changes with a 11-year cycle linked to the solar dynamo. Especially the change in the solar magnetic field implies also changes in the rates of low-energy cosmic rays arriving at the Earth. Thus, we expect to see an 11-year modulation of the rate of low-energy particles at the Earth as has been observed with neutron monitors before [52].

2.2 Physics of Thunderstorms

Following the short introduction to cosmic rays, we want to briefly introduce the physics of lightning strikes because in this work, we investigate the lightning initiation problem and the




Figure 2.6: A classical Forbush decrease as measured by three neutron monitors (Deep River, Mt. Wellington, Kerguelen). The black line indicates the average rate. Two steps in the decrease and the recovery typical for Forbush decreases are visible. Taken from Ref. [50].

connection to cosmic rays. Beyond the very basics of lightning strikes that we can introduce here, we refer the reader to recent reviews, e.g. Ref. [7].

2.2.1 Lightning Strikes

While the phenomenon of lightning strikes is known to everyone, the details of the processes involved are not clearly understood. In parts, this is due to the thundercloud obscuring the observation and the complexity of the processes involved. The general situation in thunderclouds with charge separation is sketched in Fig. 2.7. Due to friction of soft hail (graupel) with small ice crystals charge is separated. The lighter particles are carried upwards in the cloud creating the typical charge separation while the heavier ones remain at the bottom of the cloud. However, the exact distribution of charges varies between storms and can even be inverted for some clouds [7].

The discharges are classified into several categories with the two overarching ones being cloud-to-ground and cloud discharges with the majority (about 90%) of lightning discharges falling into the second category. Within the cloud discharge category, there are several possibilities including inter- and intra-cloud lightning connecting charge layers of two clouds or within a single cloud. These classes, together with cloud-to-air discharges, are shown and labelled in Fig. 2.7. A few rare cases, e.g. cloud-to-ground lightning from the positive layer, are not shown in the sketch [7].

The general development of lightning is shown in Fig. 2.8. After a – not yet clearly identified initiation [7, 53] – a plasma channel, called “leader” is formed that propagates downwards towards Earth. This leader moves in discreet steps of tens of meters with some 10-50 μ s between the steps, hence it is referred to as “stepped leader”. Moving downwards (with branching) as sketched in the first panels of Fig. 2.8, the (also not well understood [7]) attachment process with a positive discharge induced by the strong negative charge of the leader from the Earth creates a conductive channel. In this channel, the main discharge visible to the human eye happens with a peak current of typically 30 kA. About 80% of the negative discharges are followed by three to five subsequent strokes that happen within 30-50 ms. Because of the existing channel, the leader for this second stroke can propagate




Figure 2.7: Schematic drawing of thunderclouds with typical charge layers formed by friction of hail and graupel. Common names for the different kinds of lightning are given. Taken from Ref. [7].

more smoothly (“dart leader”) as shown in Fig. 2.8. However, also forms of leaders in between the dart-leader and the stepped-leader for subsequent strokes are observed [7].

In the whole process of lightning strikes several topics are not yet well understood. In Ref. [7], even a list of the ten most urgent unanswered questions is given. The first one, the initiation problem, raises the question about how the breakdown in air occurs leading to the formation of the leader. Typical measured values of the electric field are one order of magnitude below the conventional breakdown field of air. So, local enhancements of the field with hydrometeors, or the influence of free charged particles from cosmic-ray EAS are discussed to create the breakdown [53]. Other work [54] suggests that exact treatments of the breakdown field in air, depending on the concentration of different molecules, can lead to significantly lowered breakdown fields. With the Pierre Auger Observatory, we can investigate the hypothesis of the cosmic-ray induced breakdown by searching for small EAS correlated to lightning strikes which we present in Chapter 5.

2.2.2 Detection of lightning strikes

We give a brief introduction to the systematic detection of lightning strikes because for correlation analyses of lightning with cosmic-ray information about lightning strikes are necessary.

Beyond the use of photography for lightning, the main tool to detect processes in lightning strikes is radio measurements at various frequencies. For instance, the streamer discharge at the beginning of the strike can be measured in the frequency range ~ 0.4 to 1 GHz, and emission from the propagating leader in the range of about 60 to 200 MHz [55].

The strongest emission from the lightning process is coming from the return stroke. It produces a broadband radio emission peaking at frequencies between 5 and 15 kHz. Using these pulses, detector networks can use similar patterns in the waveforms measured in different places to locate the lightning strike by finding intersection points of contours of constant propagation time to all the measurements [55].

Using the same principle but on scales of several kilometres on the ground, the interferometric approach can be used to follow the development of the lightning within the cloud. With this technique a three-dimensional and time-resolved tracking of the radio sources is possible. See for instance Ref. [56] for an example of such measurements.




Figure 2.8: Sketch of the development of a lightning strike from the stepped leader to second return strokes. The time scale of the processes is given below the sketches. Taken from Ref. [7].

2.2.3 Terrestrial Gamma-Ray Flashes and Other Phenomena

A special type of emission from lightning strikes is the emission of gamma rays of multiple MeV energy. These gamma rays are emitted in so called Terrestrial Gamma-Ray Flashes (TGFs) and were discovered by BATSE, an instrument for the detection of gamma-ray bursts (GRBs) from space on board of the Compton Gamma-ray Observatory (CGRO) satellite [57]. Fig. 2.9 shows light curves measured by BATSE for some example TGFs. They typically last for a millisecond and show a wide range of shapes, some even repeated signals. In addition to these flashes of gamma rays, also X-ray emission has been observed, both presumably linked to relativistic runaway electron avalanches (RREAs). In these avalanches the electrical field strength is large enough that electrons are accelerated during their free path to energies enough to release new electrons upon collision. The exact origin of the emission and mechanisms producing the RREAs and gamma-rays are not identified yet [58]. Recently, these gamma rays have also been observed at ground in rocket-triggered lightning [59] and in cosmic-ray observatories during thunderstorms [9, 60]. For an in depth review of the current understanding of TGFs see Ref. [58].

Beyond the emission of high-energy radiation there are also other luminous events related to thunderstorms that go beyond the visible flash of the return stroke. In Fig. 2.10 these Transient Luminous Events (TLEs) are shown and named. Correlations of such TLEs with TGFs are also subject of research and modelling [58]. The Pierre Auger Observatory, is also




Figure 2.9: Light curves of different Terrestrial Gamma Ray Flashes (TGFs) recorded by BATSE. Taken from Ref. [57].




Figure 2.10: Illustration of different Transient Luminous Events (TLE). These are electric events caused by thunderstorms that are visible in different wavelengths. Taken from Ref. [61].

capable of measuring TLEs, especially the ELVES² [8]. However, as we do not focus on the analysis of these events in this work, we refer the reader for more details on this class of events to Ref. [8].

²ELVES is an acronym for Emission of Light and Very Low frequency perturbations due to Electromagnetic pulse Sources.

The Pierre Auger Observatory

After the general introduction to the relevant physics for this work in the previous Chapter, we focus now on the Pierre Auger Observatory [1] and its components of relevance for this work. We give a brief overview of the individual sub-detector systems of the Observatory and explain the necessary details of the trigger data used throughout this work.

The Pierre Auger Observatory is built as a hybrid detector, combining fluorescence detection (FD) with particle detection at the ground to provide high statistics with low systematic uncertainties. A calorimetric measurement, as provided by the fluorescence technique [1], sets the energy scale [62] of the Observatory and helps reduce the impact of systematic uncertainties from simulation-based energy determination. In total, the Observatory covers an area of 3000 km² with the surface detector (SD) overlooked by four FD sites, as illustrated in Fig. 3.1, and is thus the largest cosmic-ray observatory in the world.

In the following, we briefly discuss the FD and present a more in depth introduction to the surface detector (SD). Finally, we shortly introduce the main components of the on-going upgrade of the Observatory [4].

3.1 The Fluorescence Detector

The fluorescence detector of the Pierre Auger Observatory is designed to measure the ultraviolet light emitted by excited nitrogen when an extensive air shower passes through the atmosphere above the surface detector. The design requirement was to detect every shower with energies above 10¹⁹ eV hitting the array during FD on-time [1]. Geometrically, the FDs are positioned on the edge of the SD array overlooking it from four positions, each of these “eyes” consisting of six Schmidt telescopes as can be seen in Fig. 3.1. In addition, there is the low-energy extension with three high-elevation Auger telescopes (HEAT) at one of the sites [1].

Schematically, Fig. 3.2 shows the construction of these telescopes consisting of a 13 m² segmented mirror, a PMT camera with 440 pixels, a UV-filter, and an aperture with a shutter. Each telescope has a field of view of 30° × 30° in azimuth and elevation with a minimum elevation of 1.5° [1]. The filter is more than 80% transmissive at the UV wavelength of interest (330 to 380 nm). To protect the telescope during the day and from strong winds, the shutters close automatically. Measurements are possible given clear moonless nights to reduce the background with an up-time of about 15%.

The basic principle for the reconstruction of air showers with the FD is the isotropic emission of light from the excited nitrogen molecule. Because of the optical setup of the FD, the shower will be seen as a source of light viewed at different angles. The shower axis forms

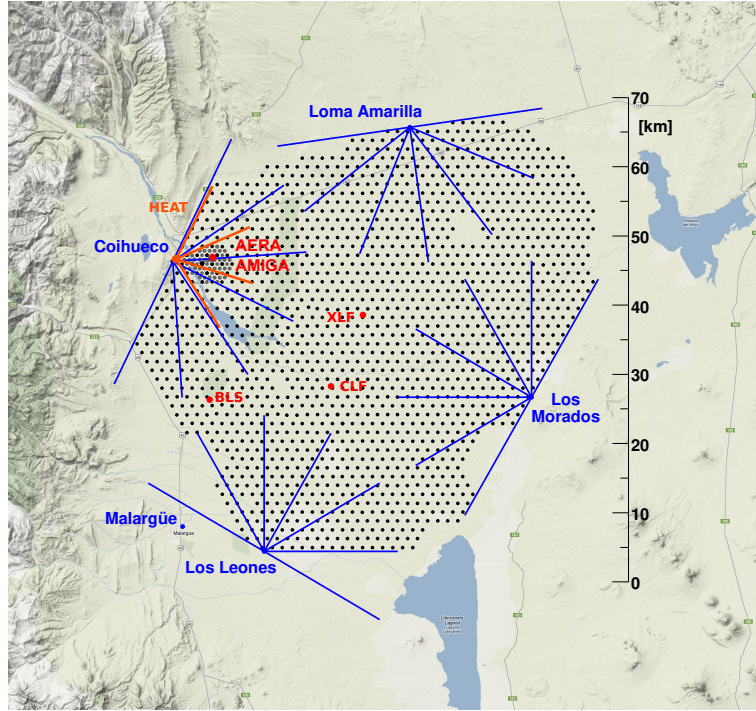


Figure 3.1: Map of the Pierre Auger Observatory and its detector systems. Black dots represent surface detector stations and blue lines indicate the field of view of fluorescence telescopes. Additional systems are marked with their respective positions in the field. Taken from Ref. [63].

a plane with the telescope, and the position of the propagating shower core along the axis defines the viewing angles. From the arrival time of the fluorescence photons in the different pixels associated to these different angles, the geometry of the shower can be determined. In the hybrid approach used in Auger, a triggered surface detector station helps significantly in this geometry fit as it is far away in phase space compared to the points of FD and thus provides a strong leverage on the axis. To be able to use the SDs in this analysis with high efficiency, a hybrid trigger is implemented that allows the read-out of stations without an SD-array trigger.

With the known shower geometry, the light at the aperture of the telescope can be converted into light emitted from the shower core taking into account the attenuation. The emitted light is proportional to the deposited energy. We can convert the emitted light to energy deposit with the fluorescence yield, as given e.g. in Ref. [64]. The resulting profile, shown for an example event as function of the slant depth X in the right panel of Fig. 3.2, is then fitted with a Gaisser-Hillas [65] function

$$f_{\text{GH}}(X) = \left(\frac{dE}{dX} \right)_{\text{max}} \left(\frac{X - X_0}{X_{\text{max}} - X_0} \right)^{(X_{\text{max}} - X_0)/\lambda} e^{-(X_{\text{max}} - X)/\lambda} \quad (3.1)$$

to determine the depth of shower maximum X_{max} [1].

An estimate of the shower energy is obtained by integrating the profile and correcting for the energy carried away by particles not visible in the UV wavelength range, e.g. neutrinos [66] and muons.

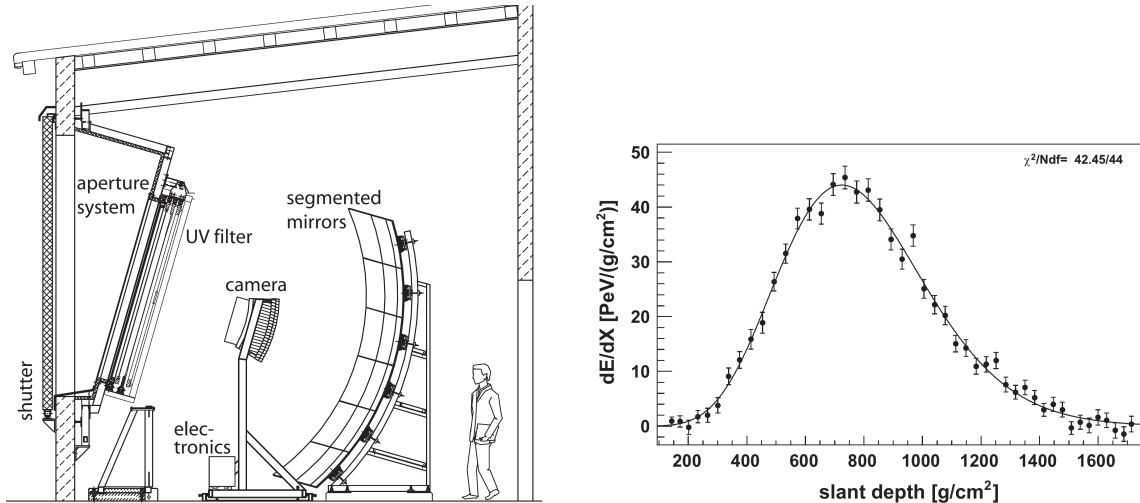


Figure 3.2: Schematics of the fluorescence detector and the associated reconstruction. Figures are taken from Ref. [1]. *Left:* Schematic drawing of the telescopes. *Right:* Example reconstructed shower profile in terms of energy deposit.

3.2 The Surface Detector

In this Section, we describe the SD system [67] of the Pierre Auger Observatory in more detail. We start with a brief description of the hardware followed by a more detailed discussion of the triggers and the calibration. Similar to the FD, we also sketch the reconstruction methods used in the analysis of air showers measured with the SD.

3.2.1 The SD Station Hardware

The design goal of the SD of the Pierre Auger Observatory was to be able to measure with nearly 100% efficiency all showers above $10^{18.5}$ eV with an almost 100% uptime and with a well defined aperture [1]. This is achieved with a design of the SD consisting of 1660 autonomous stations covering $3\,000\text{ km}^2$ in a triangular grid with an 1.5 km spacing (SD-1500). To extend the capabilities of the Observatory to lower energies, also an infill region of SD stations with a 750 m spacing (SD-750) is deployed [68].

Due to the large area covered the stations have to run autonomously with a requirement of no or little maintenance. The design, illustrated in the left panel of Fig. 3.3, is a 3.6 m diameter water tank with a sealed diffusive-reflective liner that is 1.2 m high containing 12 metric tons of purified water [67]. The Cherenkov light emitted by relativistic charged particles in this tank is recorded with three Photonis XP1805 PMTs [1] symmetrically placed on the top of the tank.

To achieve the autonomous operation, a solar cell system with batteries capable of providing on average 10 W of electric power [1] is in use and visible in Fig. 3.3. The station electronics performing calibration, data-taking, communicating, and providing the high voltage for the PMTs is powered with this system and its enclosure is visible on the top of the tank in Fig. 3.3. For data-taking, six channels with 10 bit flash analogue-to-digital converters (FADC) sampled at 40 MHz [1] are available. For each PMT the output of the last dynode – amplified to 32 times the level of the anode – and of the anode are fed to these FADCs to sample the signal and to increase the dynamic range. Saturation occurs at about 650 times the signal of a single vertical through going muon [1].

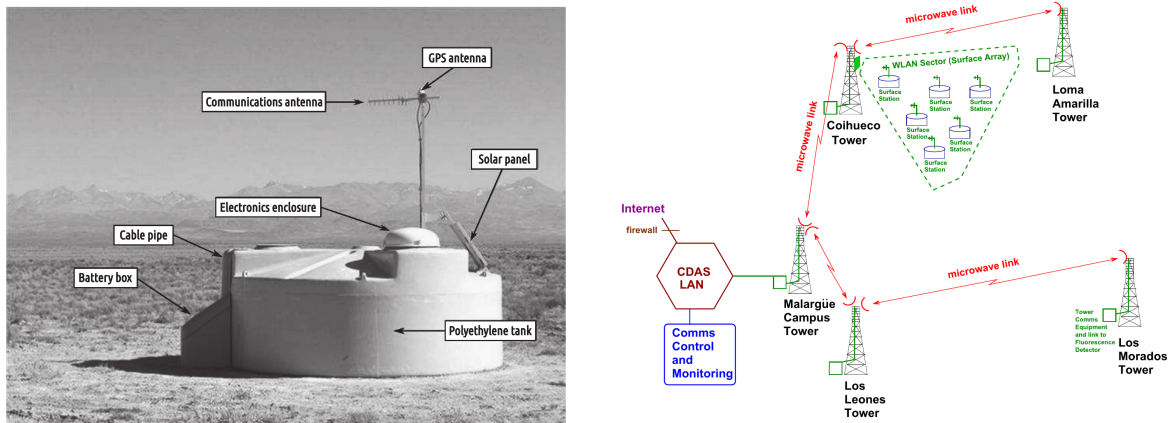


Figure 3.3: Systems of relevance for the surface detector. Taken from Ref. [1]. *Left:* Photograph of a surface-detector station with labels for its hardware components. *Right:* Scheme of the communication system for the data transfer.

For timing, a GPS receiver is used that provides, in combination with a custom ASIC¹, absolute timing to about 10 ns RMS [1]. Triggering and data-taking is performed with a programmable logic device (PLD) implementing in firmware the different trigger modes, described in the next section.

An IBM PowerPC 403 GCX CPU with clock rate of 80 MHz acts as station controller and is placed together with a slow-control microprocessor on a unified board (UB) that interfaces the front-end, the timing, and the communication transceiver [1]. On the station controller the data acquisition system (DAQ) is implemented, configuring trigger threshold, taking calibration data, and sending collected data via the communication system.

The right panel of Fig. 3.3 shows a scheme of the communication system of the Observatory with the backbone microwave links between the FD sites connecting it to the central data acquisition system (CDAS) [1]. A WLAN between the station in one of 28 sectors and a communication tower at a FD site is the providing the link to the individual stations. The custom protocol using a GPS synchronised time division multiple access (TDMA) scheme reaches an effective bandwidth of 1200 bits per second uplink from the station with 6 downlink slots per second. The TDMA scheme assigns each station a certain time slot in the network to achieve contention free communication. Due to the bandwidth, the sustainable trigger rate per station is about 20 Hz (sent to CDAS) and the event rate to less than 1 per hour [46]. We discuss the triggers achieving this in the next section.

3.2.2 The SD Trigger System

Here, we will describe in more detail the trigger system used for the SD. We go into more depth because the data from these triggers is used in this work.

The main challenge for the data taking with the SD is to reduce the rate of single particles crossing the detector stations at 3 kHz to a rate of about 20 Hz that can be sent to CDAS to form event triggers without losing high-energy events. A hierarchical scheme is employed to achieve this, with two levels of triggers on the local station and one to form event-level triggers. Fig. 3.4 depicts this hierarchy and shows the approximate rates associated with each trigger [46]. In 2013 new triggers [1, 68] were added to extend the sensitivity to lower energies that are incorporated into the ToT-branch in the scheme of Fig. 3.4.

¹Application Specific Integrated Circuit

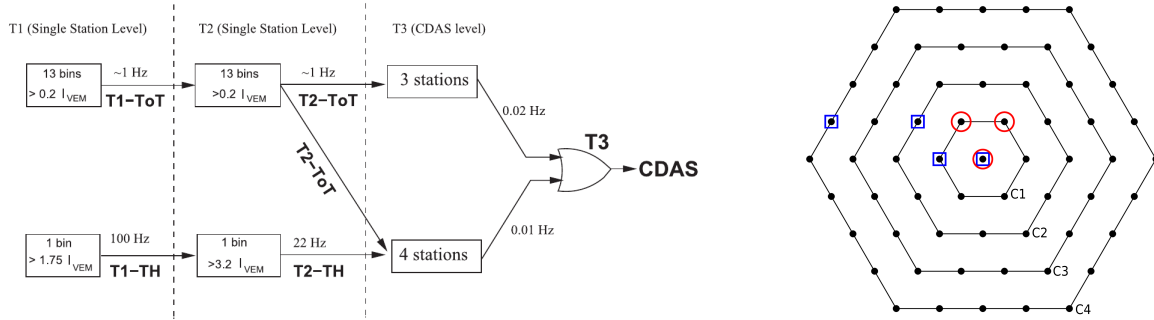


Figure 3.4: Illustrations of the trigger system for the Auger SD. *Left:* Trigger chain from the station hardware level (T1), station software level (T2), to the CDAS forming of coincidences on array level (T3). Taken from Ref. [1]. *Right:* Illustration of two configurations in the hexagonal grid with the crowns C_i that fulfil the geometric requirements imposed by the T3 trigger. Modified from Ref. [1].

On station level, two trigger levels are used because it is possible to buffer more traces to send to the CDAS upon request than it is possible to send to CDAS to form the trigger. Thus, the first level (T1) is performed by the station PLD and the resulting traces are processed in the station controller and reduced there to the bandwidth fitting 20 Hz (T2) [46]. We will describe these station-level triggers in the following.

The first trigger mode, the threshold trigger (“Th-T1 or T2”) is simply requesting the maximal amplitude of a signal to be above a value high enough to reduce random background to about 100 or 20 Hz respectively for T1 / T2. In practise, this threshold is set to $1.75 I_{VEM}$ in coincidence of all three PMTs for the T1 where I_{VEM} is the average maximal amplitude of a vertical through-going muon [46]. The calibration for I_{VEM} is available also online in the station through an online calibration method that we describe in the next section. For the reduction to 20 Hz with the T2 trigger the threshold is set to $3.2 I_{VEM}$. In case not all three PMTs are working, the thresholds are adjusted accordingly to achieve the same trigger rate [46].

These threshold triggers are good at recording inclined showers that are dominated by muons and create large signals not spread in time. The right panel of Fig. 3.5 shows an example trace triggered by only this type of trigger from an inclined event. The elongated footprint of the event is shown as inset and the characteristic short pulse of large amplitude is clearly visible.

The second mode of triggers are targeting signals spread out in time, but also possibly with lower amplitudes. These time-over-threshold triggers (ToT) target signals far from the core of the showers in high-energy events or signals close to the core in lower-energy events. The spread of the signals arises from scattering of the electro magnetic component present in the signals and from geometric delays for the muons [46]. The left trace in Fig. 3.5 shows as a typical example a station 2 km from the core that was triggered by such a ToT. In comparison with the signal from the inclined event on the right, the time spread as well as the small amplitude are clearly visible.

In practise, the trigger condition for such ToT-T1 is more than 12 time bins (i.e. > 300 ns) in a $3 \mu\text{s}$ sliding window with a signal above a threshold of $0.2 I_{VEM}$ in coincidence of at least two of the three PMTs. The rate of these triggers is about 1 Hz and no further condition is applied to the T1s and all events recorded with a ToT are promoted to a (ToT)-T2 [46]. In the case of an event passing both the T1-threshold and the ToT, the ToT is given higher priority such that the event will appear as ToT in the formation of array-level triggers.

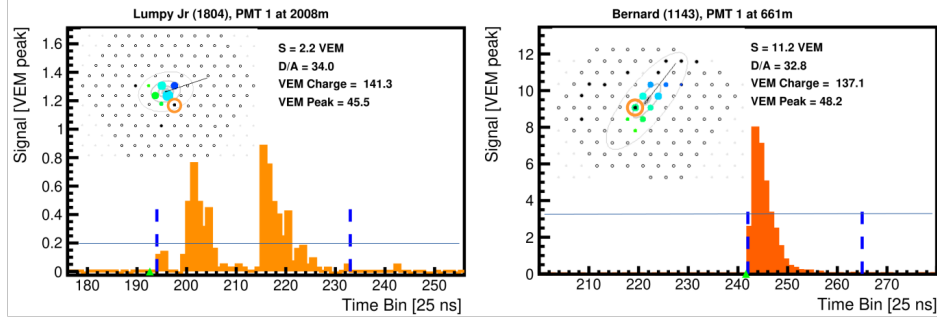


Figure 3.5: Illustration of the two main trigger modes of the SD. *Left:* ToT-triggered trace of station 1804 (PMT 1) from the event shown in the inset. The threshold of $0.2 \text{ VEM}_{\text{peak}}$ is illustrated as a line. *Right:* Threshold triggered station in an inclined event with larger total signal but smaller time dispersion.




Figure 3.6: Illustration of the deconvolved ToT (ToTd) trigger. Taken from Ref. [68]. *Left:* Trace triggered by a ToTd. *Right:* Deconvolved trace of the same event used to form the ToTd trigger.

Since June 2013 two new triggers are used in addition. They work similarly to the ToT trigger and target electromagnetic parts of the shower to reduce the energy threshold and to increase sensitivity to photon and neutrino-induced showers [1].

The first of the new triggers, the deconvolved time-over-threshold (ToTd), is deconvolving the exponential decay of the typical signal of muons to reduce the impact of coincident muons on the trigger. In Fig. 3.6 an example trace that was triggered by the ToTd is shown, both as actual trace and as deconvolved trace used for the trigger. The deconvolved traces with bins D_i is obtained as

$$D_i = \frac{S_i - S_{i-1} e^{\Delta t/\tau}}{1 - e^{\Delta t/\tau}} \quad (3.2)$$

where S_i are the bins of the measured trace. Here, $\tau = 67 \text{ ns}$, the typical decay time of light in the water-Cherenkov detector, and the time bin width is $\Delta t = 25 \text{ ns}$. On the deconvolved trace, the same condition as for the ToT is applied to obtain a trigger [68] adding only a condition on the integral of the signal of about 0.5 VEM^2 in a window of 125 bins.

The second new trigger, “multiplicity of positive steps” (MoPS), makes use of the fact that muons are the main random background and lead to characteristic strong pulses of amplitude $I_{\text{VEM}} \approx 50 \text{ FADC}$. To count particles, monotonically increasing parts of the trace (called “positive steps”) are counted in a $3 \mu\text{s}$ window and the trigger condition is fulfilled if more than four such steps are found. A condition on the step size j , defined as the maximal difference of amplitude in the monotonically increasing interval, is imposed to reduce noise by requiring $j > 3 \text{ FADC}$. And to exclude muons, $j < 32 \text{ FADC}$ is used. Finally, the same condition on the integral signal is applied as in the ToTd trigger to reduce the influence of instabilities [1, 68]. Both new triggers are used as ToT-like triggers in the trigger hierarchy.

²The Vertical equivalent muon is the calibration unit used and is explained in Section 3.2.3

To form an event trigger T3, the location and the timing of triggers is used to form coincidences compatible with a shower front moving at the speed of light. Making use of the triangular grid of the SD, the geometric conditions can be expressed in terms of the neighbourhood of a given station. Each station is surrounded by rings of stations forming a hexagon, called crown C_i as visualised in the right panel of Fig. 3.4, which can be used to impose numbers of triggered stations in the neighbourhood of a chosen station. There are two modes implemented which make use of the different background rates of the threshold and ToT triggers. The first one “ToT-2 C_1 & 3 C_2 ” requires at least one nearest ToT triggered neighbour to a ToT triggered station (i.e. two ToTs within C_1) and one ToT triggered station next-to-nearest neighbour of a station (i.e. three within C_2). A compact formation of three triggers, like the red circles in Fig. 3.4 is fulfilling this condition. For more inclined events, a second mode “2 C_1 & 3 C_2 & 4 C_4 ”, is used that takes all T2s into account [46].

For both T3 conditions, the timing of the triggers has to match the assumption of a plane front moving at the speed of light. As such, the necessary condition of time differences to the first trigger $\Delta t_{\text{first}} < (6 + 5C_n)\mu\text{s}$ is imposed. The resulting rate of these triggers is about 0.03 Hz [46] or 1.8 events per minute.

When a T3 is formed the data from the stations are requested, also for those within six crowns of a station contributing to the T3 that was not part of the event trigger. At this point also the traces not passing the T2 conditions (i.e. T1 traces) are returned if their timing matches within 30 μs [46].

The trigger information is continuously being sent to CDAS in packages containing all T2s of one GPS second, encoding the trigger microsecond and the trigger type for each of the T2s. For the type of the trigger four bits are available which can encode also additional information beyond the ToT/Th-information necessary for the T3. In October 2016 an additional algorithm was implemented [69] to select small showers characterised by a charge-over-amplitude ratio that exceed that of instantaneous signals (like muons). Fig. 3.7 shows the distribution of this ratio in background triggers (collected with a random T3) in red with a clear break in the spectrum at about 1.3. In the implementation of the station software this selection is performed using the online calibration values and placing a cut of charge-amplitude ratio > 1.2 for all “latch-threshold” triggered events³. The distinction of a latch from a post-latch trigger refers to the case when more than one trigger is activated for the same event. Because the same event should not be recorded more than once, it is done so only after the first trigger mode is activated – the latch trigger – and additional triggers for the determination of the trigger type are computed on the existing trace (“post-latch”). The signal close to the core of a sub-threshold shower is expected to be sharply rising thus triggering as threshold trigger motivating this choice while also reducing the necessary computation in the station controller [70]. Because after 2016 the trigger information received at CDAS is stored in addition to the data, this information propagated through the trigger type bits is available for offline analysis.

Another special algorithm is flagging low-gain saturated traces in the same way. It searches for more than six bins above 1020 FADC-channels (coincident in all three PMTs) in latch-triggered traces and uses a dedicated trigger flag (later called “low-gain saturation”) available in the stored trigger data.

3.2.3 Calibration of the Surface Detector Stations

In this Section, we briefly outline the procedure used to calibrate the response of the SD stations. The physical unit that is used to have a common reference for the stations is the

³As extracted from the source code available for the Auger collaboration from <https://gitlab.com/auger-observatory/sde/ub/-/blob/master/Src/Trigger2/xbtrig.c>




Figure 3.7: Normalised distribution of Charge (“Area”) divided by maximal amplitude (“Peak”) calibrated with the muon response of random signals in data (red) and for data read by a special lightning trigger (blue). The special lightning trigger is using commercial lightning detection devices [71] to detect lightning strikes and requests data from stations close in time and space. Taken from Ref. [69].

most probable signal of a vertical, centred, through-going muon called “VEM”. The goal of the calibration procedure is then to obtain a conversion of electronics ADC channels to VEM [72].

There are two parts that require the calibration: the analysis of the shower signals and the triggering during data taking. We will cover the employed algorithms for both briefly.

Fig. 3.8 shows histograms of the integrated signal (charge) and amplitude (peak) of signals recorded in a SD station. In both histograms, clearly a peak corresponding to atmospheric muons is visible that can be used to obtain a value for the VEM. Because not all atmospheric muons are centred and vertical, the actual value of the VEM is not the peak value visible in these histograms but shifted by a constant factor, visible from the shifted red histogram. These VEM-histograms were obtained using an additional detector to select vertical, central muons only [72]. For the analysis of shower data, these histograms – for each PMT separately – are fitted and the peak value is converted to the VEM [73].

In the station software such algorithms are not possible to run due to the instabilities and the limitations of the station controller. Instead, the background rate is used to define the thresholds and to continuously monitor the mean peak value I_{peak} . After initially setting the high voltage of the PMTs such that the rate of each PMT for a threshold of 150 ADC above baseline is 100 Hz (implying $I_{\text{peak}} \approx 50$ ADC), the high voltage is not changed during data-taking and only the conversion of 1 VEM in terms of ADC is tracked [72].

For this purpose a “calibration trigger” is used that is counting the events satisfying a T1-threshold and has a peak above $2.2 I_{\text{peak}}$ in the PMT to calibrate. The rate of these events should be 70 Hz and the estimate of I_{peak} is adjusted using a σ - δ -algorithm to achieve this [72]. A similar approach is used to track the charge Q_{VEM} of the VEM using signals with amplitude of $1.75 I_{\text{peak}}$. For further details on this calibration procedure see Ref. [72] and Ref. [73] for the offline calibration.




Figure 3.8: Histograms of charge and amplitude (summed over the PMTs) of a station taken with a 3-fold coincidence trigger at 5 ADC threshold. The first peak comes from low-energy particles convolved with the trigger response, while the second one comes from through-going atmospheric muons. With the help of a muon telescope centred, vertical through-going muons are selected and shown in red. Taken from Ref. [72].

3.2.4 Scalers

The Auger Scalers [74] (Scalers) are particle counters that provide information about the background rate in the SD stations. Conceived to be able to search for gamma-ray bursts [75] by observing a simultaneous increase of the rate in many or all SD stations, they can also give insight to the rate of low-energy cosmic-rays influenced by solar modulation [74, 76].

The Scalers are implemented as a trigger in the station firmware counting all pulses with an amplitude of more than 3 ADC above baseline in coincidence of all three PMTs. After an initial testing period, an upper cut of 20 ADC above baseline was implemented to reduce the background from muons ($I_{\text{peak}} \sim 50$ ADC) in gamma-ray bursts searches [75]. Thus, a signal is counted if the maximal amplitude I fulfils $3 < I/\text{ADC} < 21$.

The Scalers are sent every second together with the T2-triggers to CDAS and stored. On average the count rate is about 2 kHz per detector and the deposited energy is about 15 to 100 MeV [74]. Using simulations, the primary particles creating the secondary particles primarily counted by the Scaler are determined to be between 10 GeV and about 10 TeV with a median of 90 GeV [76].

We show details of the necessary corrections before the Scalers can be used in physics analysis in Chapter 4 and use them in example analyses to highlight their potential.

3.2.5 Reconstruction of Ultra-High Energy Cosmic Rays

After taking the data with the trigger system described in the previous sections, the events can be reconstructed to obtain information about the primary particle: most importantly arrival direction and energy. Here, we sketch the basic principles of the reconstruction. For a more detailed description, cf. Ref. [73] that is used as basis for this short summary.

The arrival direction can be reconstructed using the relative timing of the stations. With a basic model of a plane front shower moving at the speed of light, three stations are sufficient to constrain the shower axis and the direction in the zenith angle θ and the azimuth angle ϕ . A sketch of the relative timing in a high-energy event can be seen in the left panel of Fig. 3.9. The footprint of stations is shown with the timing as colour code and the time evolution across the footprint is clearly visible.




Figure 3.9: Illustrations of the event reconstruction with the Auger SD. Figures are taken from Ref. [73]. *Left:* Footprint of stations of an event coloured by arrival time (blue early, green late). The projection of the reconstructed shower axis ending in the shower impact point (red square) is shown as the black line. *Right:* Lateral distribution of signals in a shower. The fits correspond to two different approaches for the reconstruction and the functional form. For details see Ref. [73].

To reconstruct the energy and the impact point of the shower core, an axisymmetric model of the lateral distribution of signals (LDF), e.g.

$$S(r) = S_{1000} \left(\frac{r}{r_{\text{opt}}} \right)^{\beta} \left(\frac{r + r_s}{r_{\text{opt}} + r_s} \right)^{\beta + \gamma} \quad (3.3)$$

with radial distance r is used. β , γ , and r_{opt} are parameters describing the shape of the lateral distribution and are parametrised on subsets of the data. The symmetry constrains the impact point, while the scaling of the function S_{1000} is providing an estimator for the primary energy [73].

To convert S_{1000} to an energy, the attenuation in the atmosphere has to be accounted for obtaining an attenuation-free shower-size estimator which is calibrated to the FD energy scale using the hybrid subset of the data set. For the functional form of the attenuation the isotropy of the arrival directions for high energies can be used to constrain the zenith dependence of the signal with a constant-intensity-cut method [3].

3.3 Upgrades

The Pierre Auger Observatory is currently undergoing an upgrade [4] to further enhance its capabilities to answer the open questions in ultra-high energy cosmic-ray physics. The goal of the upgrade dubbed AugerPrime is to enable mass sensitivity on an event-by-event basis by enhancing the surface detectors with scintillators [77, 78], new electronics [79], an additional small PMT [80], and radio antennas [81], as well as extensions of the measuring time of the FD to nights with more background light [4, 80], and an underground muon detector [4].

In this work, we focus on the data taken already with the current hardware and only outline possibilities for the future with AugerPrime. For more detailed information see e.g. Ref. [4].

The Surface Scintillator Detector (SSD) provides a handle on the mass of the primary particle because of its different response to the muonic and electromagnetic components of showers when compared to the water-Cherenkov detector. With a parametrisation of the

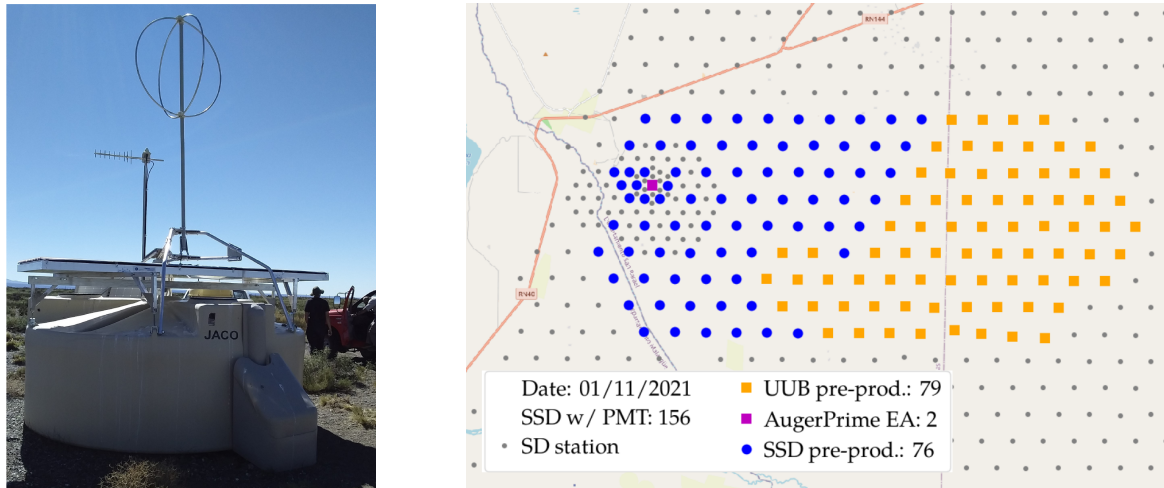


Figure 3.10: Components of the upgrade of the Pierre Auger Observatory (AugerPrime). *Left:* Photograph of an upgraded station with Scintillator and radio antenna. Taken from Ref. [78]. *Right:* Deployment map of the upgraded components in the field. Kindly provided by A. Streich.

response for the two components, the two measured signals can be decomposed into muonic and electromagnetic parts to obtain the muon content that is mass sensitive. For details of such an approach see e.g. Ref. [82].

To accommodate the new detectors and to provide more powerful processing capabilities, also the electronics of the stations are upgraded. The unified board (UB) will be replaced with the upgraded unified board (UUB) and provide faster sampled FADCs (120 MHz instead of 40 MHz) with a larger dynamic range of 12 instead of 10 bits. A tenfold increase in computation power and memory is provided with the new hardware. The first preproduction UUBs have been deployed in the field in December 2020 [79].

In Fig. 3.10 a photograph of an upgraded SD station is shown with the radio antenna, targeting the electromagnetic component for inclined showers, and the SSD mostly efficient for vertical showers. The status of deployment at the time of writing of this thesis is shown in the right panel in Fig. 3.10.

Analysis of Scalers

In this Chapter, we present a first analysis of the Auger Scalers on more than 15 years of data, covering more than a full solar cycle. We extend previous and existing analyses [74, 83–86] with new selections and introduce new corrections to increase the data quality. We then focus on showing that this data set is suitable for studying solar physics by searching for known and expected signals. In the conclusion of this Chapter, we briefly discuss as well the possibilities for improving the Scaler in the future with the upgraded AugerPrime electronics.

Results presented in this Chapter were previously used in Ref. [87] and are extended here with more data. The selections were first presented internally within the Auger collaboration in Ref. [88] and are summarised again in this chapter.

4.1 Improvements of the Data Processing

Already in previous work with the Auger Scaler data (e.g. Refs. [74, 83–86]) it proved critical to select the appropriate data from the Auger Scalers to obtain reliable results. This is due to the large influence minor instabilities of the hardware can have at very low thresholds. In this Section, we show how we extend and refine these selections and cuts for the analysis of long-term trends of the data. We focus on clear concepts to facilitate easy future modifications and adjustments to different analyses.

In this work, the main implementation challenge is to gather all necessary auxiliary data for the analysis of the Scalers. The aforementioned cuts need the detector monitoring information which is not available in the Scaler files. Detailed information about the PMT status is available in the event data, and need to be accessed for this analysis as well to ensure the selection of properly working PMTs only. For the known dependencies on atmospheric conditions also weather information is needed.

This situation is summarised in Fig. 4.1, where also the different time intervals for data collection are indicated. The idea behind the analysis (software¹) ansatz used here is thus that collecting the data and merging them in a meaningful way is a task that can be separated from the actual analysis. Because of this separation, the analysis is defined in “modules” that can be analysis-specific while the details of handling the input are hidden. In Fig. 4.1 the places where modifications are meant to happen are indicated with red-shaded boxes.

The data that are not available on the scale of single seconds like the Scalers are interpolated. This is the case for the monitoring and the weather data. For the PMT quality, we use a single value per day, defining any PMT as bad for the whole day if it appeared as “bad” in at least one SD event in this period. In the end, a cleaned data sample is constructed

¹The software written for this analysis is made available to the Auger collaboration to facilitate future analyses.

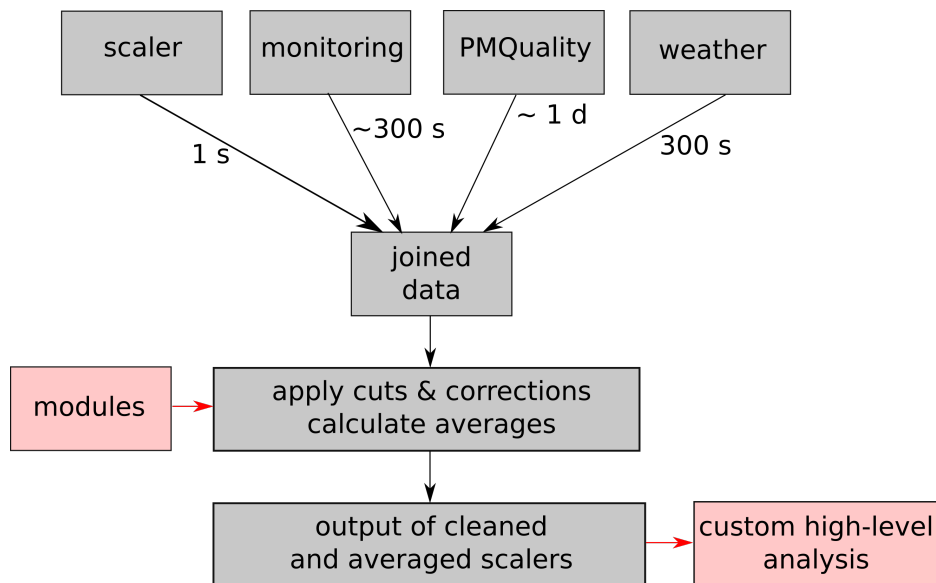


Figure 4.1: Sketch of data flow used in this Scaler analysis framework. The used time scale of the input data is indicated as number at the arrows. Parts meant for modification in future analyses are indicated with red boxes.

containing averages of the Scaler rate in a custom time interval. For our purpose a length of five minutes is adequate and adopted here. In future extensions, this time scale can be changed if necessary.

In the following, the details of this work with the associated cuts and corrections are described leading to the corrected data set used in the analysis. We will apply the analysis on the data from January 2006 to end of June 2021. Details presented in this sections were already published internally within the Auger collaboration as Ref. [88].

4.1.1 Identification of Bad Stations

After merging all the necessary data, the first data-selection step consists of identifying those stations that are in general ready to take valid data and are sending valid data for each second. All the following data selections are applied on single seconds on data only.

To exclude corrupted or unknown data we select from the input only known stations by comparing the station Ids. In normal conditions, this selection is not removing any data and is there only to avoid contamination from unexpected data. Similarly, we exclude stations that do not send data. In our case, this includes the stations with available Scaler data but no monitoring that is needed to select the working stations. This happens occasionally because the Scaler data is sent as part of the trigger data with the highest-priority in the network while the monitoring data has the lowest priority.

At this stage, we also add a “sanity cut” on the rate itself. We want to avoid imposing a strict condition on the data but we remove those data points that are above 65 535 counts to avoid overflow problems with the unsigned short int variable type used in the software. Given the expected rate of about 2 kHz this cut does not have any influence on normal data.

The removal of stations without three properly working PMTs also happens at this stage. We reject the data of stations that have at least one PMT flagged by the official PMT cuts [89, 90] which is merged into the Scaler data from the official SD productions of event data.

The two final cuts use two parameters from the monitoring data of the station: the signal shape parameter “Area-over-Peak” (AoP) and the online calibration of the most frequent




Figure 4.2: Statistics on the number of stations left in the analysis after applying quality cuts. *Left:* Number of stations left in the analysis with relative rates. *Right:* Removed stations because they do not fulfil the $35 < I_{\text{peak}} < 65$ condition. A slight recovery at the very end of the data set can be seen and it originates from the restart of the stations after the shut down of the array during the austral winter of June 2021 [91].

amplitude of the muon signals I_{peak} . We request $2.5 < \text{AoP} < 4.5$ similarly to previous studies [84] to exclude stations well outside the normal working range. The online estimate of the muon signal I_{peak} is of importance because the threshold of the Scaler is fixed in FADC rather than kept constant relative to I_{peak} . Thus, if the calibration – tracked by I_{peak} – changes over time from the expected 50 ADC the implicitly selected energy range changes. To avoid biases due to such drifts in calibration, we select only stations with all PMTs fulfilling $35 < I_{\text{peak}}/\text{ADC} < 65$. The range corresponds to quite a large $\pm 30\%$ deviation from the design value since we try to keep as many stations as possible in the analysis. In future analyses, these choices might need to be revisited. Should it be possible to correct the Scaler rate for the drift in calibration, this cut could be loosened.

Applying these cuts to the data, we observe that the number of stations not rejected shows significant long-term trends as visible from the left panel in Fig. 4.2. The right panel in Fig. 4.2 shows the number of stations rejected by the condition on I_{peak} . It is clear that the ageing of the detectors has influence on this analysis. To identify the other contributions to the rejection of stations, we show the fraction of passing and rejected stations in Fig. 4.3. Clearly the PMT cut has the largest influence apart from the online calibration. The influence of this low statistics has to be studied in more detail in future work and especially the decrease of valid stations in recent years has to be closely monitored. A special case is the deficit of stations in the years 2013 and 2014, visible in Fig. 4.3. The monitoring data used to place the cuts is not available as in the rest of the data leading to a loss of nearly all stations due to the conservative approach in not using stations without all the requested data. This also implies that with further manual checks, these data can be recovered for future analysis.

4.1.2 Station Outliers

After removing those stations identified by “auxiliary” data as unreliable, we proceed to using the data itself to remove outliers. In this section, we describe the removal of stations because they report a very high count rate in a single second without other stations being affected as well. These single-station outliers in the Scaler rate were reported previously e.g. in Ref. [86]. We want to stress that these are not to be confused with “unstable” stations because we observe outliers in all stations and their occurrence is stable in time. Fig. 4.4 shows this with the station ids and the time of outliers that are rejected with the method explained in the following. To illustrate how these outliers look in a time series, we show




Figure 4.3: Fractions of stations valid or rejected in the analysis as a function of time covering the full data set. For clarity, only two five-minute intervals per day are shown. Note that the fractions do not have to add up to unity because a station can be rejected due to multiple conditions and we do not show rejection due to missing data. Missing monitoring data is responsible for the drops of number of active stations in 2013 and 2014.

every Scaler rate value of station 452 in the left panel of Fig. 4.5 over a period of about 5 days with some outliers clearly visible.

To be able to detect outliers, we need an estimate of the mean and variance at the time of interest. We make use of the structure given by the desired result of the analysis – a five minute interval with average quantities – and use the median and the median absolute deviation (MAD) to estimate mean and variance. With both the median and the MAD available, the outlier detection is simple: for every second and station in the input data, the pull quantity

$$z = \frac{\Gamma - \tilde{\Gamma}}{\tilde{\sigma}} \quad (4.1)$$

is calculated, with $\tilde{\Gamma}$ denoting the median and $\tilde{\sigma}$ the MAD scaled to match a Gaussian sigma. If this pull is larger than $z_{\max} = 3$ this station and the second are flagged as outlier. In the case that there are more than two stations in the same second with such outliers, this rejection is not performed because we do not want to remove “obvious” signals by accident.

During investigations of the origin of these excesses, we find a link to the topic of electronic response after large signals that is discussed in depth in Section 6.2.5. Due to the construction of the PMT base, there is a time window after a very large signal when the signal overshoots the true baseline before relaxing to the original value. Thus, we start from the hypothesis that given these overshoots, these outliers might occur due to the changed baseline after a small shower arrives very close to a station.

To confirm this hypothesis, we extract all GPS seconds and station IDs that show such a Scaler excess and correlate them with low-gain saturated triggers (cf. Section 3.2.2) from 15 June 2019 (see Section 6.2.3 for similar searches with SD triggers alone). The right panel of Fig. 4.5 shows the resulting distribution of z for all stations in the data. It is clearly visible that there is a tail of random fluctuations reaching up to about $z = 10 \approx 10\sigma$ from the very large number of trials $N_{\text{tries}} \approx 86\,400 \text{ s} \times 1600 \text{ stations}/(\text{s station}) = 1.38 \times 10^8$.

Ignoring effects of the second boundary, we match the GPS-seconds of low-gain saturated T2 triggers with the Scaler excesses and mark those coming from the same station as




Figure 4.4: Properties of the “outliers” in single stations observed in the Scaler rate. Using a cut at $z = 10$ (Eq. (4.1)) we find about 329 000 excesses in June 2019 data. The majority of these correspond to absolute Scaler rates of 3-5 kHz. *Left:* Distributions of station Ids for the excesses. With the exception of very few stations that show a higher rate, all stations have such outliers. *Right:* Time of the selected excesses. Only minor increases in rates are visible excluding thunderstorms as origin of these outliers.




Figure 4.5: *Left:* Example of Scaler data from a single station showing singular outliers (taken from Ref. [86]). Such outliers are present across all stations and all times. *Right:* Distribution of excess z -scores of the Scaler rate during 15 June 2019 and the seconds and stations associated with a high-gain saturated T2.

“matching”. Those matching excesses are shown in orange in the right panel of Fig. 4.5. In total 1461 out of 1935 selected T2s have a corresponding Scaler excess. To estimate the contribution of random coincidences, we add a random integer value from $[-10, 10]$ to the GPS second of the T2s and repeat the analysis. In this case, only 94 T2s have a corresponding Scaler excess, making the association highly significant.

However, it should be noted that the number of Scaler excesses is much larger than the number of low-gain saturated T2s. 35 654 excesses above 5σ and 9950 above 15σ show that also large signals which do not saturate the low-gain channel have to contribute if no other effect is present. Given the very low threshold of the Scaler, it is expected that also much smaller deviations of the baseline have a significant effect on the measured rate. It thus is plausible for all those excesses to be in fact associated with large signals of small showers hitting a single station.

4.1.3 Lightning and Localised Excesses

Since the threshold of the Scaler is low, the Scaler rate is very sensitive to the electromagnetic pulses associated with lightning strikes. Identifying seconds with lightning strikes is thus

important for background rejection in analyses of the Scaler rate. In this Section, we describe how lightning effects can be detected using the Scaler data itself and how to remove part of the data to obtain a sample unaffected by such effects. The content of this section were published previously by the author in Ref. [87].

The general concept for identifying seconds with lightning strikes is similar to the search of rate increases correlated with gamma-ray bursts [75]. By comparing the Scaler rate measured in the test second $\Gamma(t_i)$ with the mean rate recorded in a five minute interval $\langle \Gamma \rangle$, significant excesses are selected. Here, as for the selection of station outliers, we take the median of five-minute intervals as a robust estimate of the mean rate $\langle \Gamma \rangle$, and estimate the current standard deviation $\hat{\sigma}$ with the MAD under a Gaussian assumption for the underlying distribution. These quantities are defined for each available SD station separately.

We use a Gaussian background model with a common additional signal s to estimate the significance of excesses. With the rates of all stations $\vec{\Gamma}$, their respective median rates $[\vec{\Gamma}]$, and the estimated standard deviations $\vec{\sigma}$, the likelihood of an additional excess signal is

$$\mathcal{L}(\vec{\Gamma} | [\vec{\Gamma}], \vec{\sigma}, s) = \prod_i \frac{1}{\sqrt{2\pi}\sigma_i} \exp \left[- \left(\frac{\Gamma_i - [\vec{\Gamma}]_i - s}{\sqrt{2}\sigma_i} \right)^2 \right], \quad (4.2)$$

where i is covering all stations.

The maximum of this likelihood gives

$$\hat{s} = \sigma^2 \sum_i \frac{\Gamma_i - [\vec{\Gamma}]_i}{\sigma_i^2} \quad (4.3)$$

with

$$\frac{1}{\sigma^2} := \sum_i \frac{1}{\sigma_i^2}. \quad (4.4)$$

After maximising with respect to s , Wilks' theorem [92] is used to estimate its significance

$$p = 1 - \text{erf} \left(\sqrt{\frac{\hat{s}}{2}} \sum_i \frac{\Gamma_i - [\vec{\Gamma}]_i}{\sigma_i^2} \right). \quad (4.5)$$

We discriminate lightning seconds from other increases using the fit quality of a true common signal of selected significant excess seconds. When the excess is originating from a lightning strike, the SD stations with a high signal are few and localised so that the quantity

$$\chi^2 = \sum_i \frac{(\Gamma_i - [\vec{\Gamma}]_i - \hat{s})^2}{\sigma_i^2 + \hat{s}} \quad (4.6)$$

where \hat{s} is the estimated common signal, is very large due to the rest of the SD array showing no excess. Seconds with lightning strikes can be identified by using the Gaussian transformation of the goodness-of-fit χ^2 with the criterion

$$d = (\chi^2 - N_{\text{DoF}}) / \sqrt{2N_{\text{DoF}}} \gtrsim 100 \quad (4.7)$$

where N_{DoF} denotes the number of degrees of freedom. For a lightning-strike event the pull $\Delta = (\Gamma - [\vec{\Gamma}]) / \hat{\sigma}$ of each station is visualised in the left panel of Fig. 4.6. A localised excess is clearly visible in the eastern part of the SD.

By stacking all identified seconds with lightning strikes during the single day of 24 January 2016, as well as the ‘‘quiet’’ seconds, the difference in the distribution of the Scaler counts




Figure 4.6: *Left:* Visualization of a lightning strike as seen in one second of the Scaler data. The color and sizes of the markers represent the deviation of the measured Scaler rate in the second from the median rate obtained on a 5-min interval. See text for more details. *Right:* Distribution of measured Scaler rates during seconds including lightning strikes (red) and during normal conditions (black).

during lightning strikes is made visible. The right panel of Fig. 4.6 shows the histograms for this example day. A clear deviation from the Gaussian distribution during lightning periods is visible. The seconds identified as affected by lightning strikes are not used in the analyses presented here.

Similarly to lightning strikes, coincidences of single-station outliers described in the previous section should be removed as a form of “localised” excesses. To include such coincident outliers in this rejection, we use a cut $d > 3$ from Eq. (4.7) as rejection criteria for single seconds with $p < 10^{-5}$ (Eq. (4.5)). If in an interval of averaged data more than two seconds are removed due to this condition, the whole set is removed to be conservative with respect to residual lightning contamination.

4.1.4 Averaging and Scaling

Before we explain the corrections of the Scaler rate for long-term use, we briefly discuss what averaging the Scaler rate can imply and how scaling it to known values can help in the interpretation of data.

In general, the computation of the average Scaler rate per data interval is straight forward. For every station, we use the arithmetic mean over all seconds in the interval that are not removed by one of the criteria discussed before. Then, the average over all such stations can be computed. We want to emphasize here, that we chose to keep the analysis on the level of single stations as long as possible and thus use this order of averaging. If no weighting with the number of active seconds is used, this is numerically different to using the average over the stations per second and taking the mean in the time interval. Because we will apply corrections for long-term effects on station basis, the aforementioned way of calculating the mean per station is conceptually cleaner. For these corrections, we will also use the averaged quantity, e.g. the average atmospheric pressure, or the average Area-over-Peak that we calculate in the same way.

We foresee the possibility to reject stations in the average if the data sample available after the application of all selections is too small. Here, we use the loose cut that only if a station has less than two seconds of valid data the station is removed. In future analyses, this number can be revisited but we think the influence of this cut is marginal.

A cut that is more important is the number of stations available to form an array average. Because the stations are deployed on different altitudes, have different ages, and different PMT-gains among other differences, the Scaler rate of two stations is not identical. Thus, only if many stations are available for analysis these differences average out. We define the threshold for this to be 150, about 10% of the array.

However, we want to emphasize that this cut is not strict at all and allows for systematic effects to show if the pure average of the Scaler rate is used for analysis. In the left panel of Fig. 4.7 we show the Scaler rate of 2013 – a year with problems in the monitoring data leading to low numbers of available stations – with outliers clearly visible.

We can compensate these differences between stations by scaling the corrected Scaler rates $\Gamma_i^{(c)}(t)$ of a station with a known reference value

$$r_i(t) = \frac{\Gamma_i^{(c)}(t)}{\langle \Gamma_i \rangle}. \quad (4.8)$$

The idea behind this scaling is that the relative response of the stations to changes of the physical rate is more uniform than the absolute count rate. To prove that this approach works in practise, we show this relative rate in the right panel of Fig. 4.7 with the same data as in the left one. Clearly, the outliers are removed. In the following we use the whole year of 2013 – that is roughly in the middle of our data set – as a reference period.

A second possibility to achieve this more universal response of the full SD to changes is to use a pull-like quantity

$$z_i(t) = \frac{\Gamma_i^{(c)}(t) - \langle \Gamma_i \rangle}{\langle \sigma_i \rangle} \quad (4.9)$$

where we use the average rate $\langle \Gamma_i \rangle$ of the same reference period and in addition scale it with the average variance in the five minute intervals $\langle \sigma_i \rangle$. This estimator has very similar properties to Eq. (4.8) but because it relies on a correct estimate of the variance $\langle \sigma_i \rangle$ is more prone to outliers and systematics. Therefore, we will use Eq. (4.8) in the following.

If we want to estimate the variance and thus the statistical error of the Scaler rate in the given time interval, we can use two ways. The first way of calculating the width of the Scaler distribution is

$$\hat{\sigma} = \text{StdDev}(\{\langle \Gamma_i^{(c)} \rangle_t\}) \quad (4.10)$$

with the corrected and time-averaged Scaler rate $\langle \Gamma_i^{(c)} \rangle_t$ of station i , and i running over all available stations. This quantity is used as error bar in the left panel of Fig. 4.8. It is obvious that this estimate does not reflect the statistical fluctuations of the mean. However, as a measure of systematic uncertainty due to the presence or absence of given stations, it is a useful first estimate.

The second way of calculating the variance is

$$\hat{\sigma}_2 = \text{StdDev}(\{\langle \Gamma_i^{(c)}(t_j) \rangle | t_j \in [t_0, t_0 + T]\}), \quad (4.11)$$

with the time-interval start-time t_0 and its length T . The estimated errors are largely reduced and shown in the right panel of Fig. 4.8 as error bars. Due to the scale, these error bars are hardly visible reflecting the actual statistical power of the large number of SD detectors in Auger.

With this take on the statistical-error estimate it might be beneficial to reduce the time intervals for averages from five to maybe one minute. It is visible from the right panel of Fig. 4.8 that there is almost no overlap within uncertainty between adjacent measurements which hints at resolvable substructure. However, any reduction of the interval length should make sure that all values for quality cuts are adjusted accordingly.




Figure 4.7: Comparison of the mean rates and rescaled quantities to create a more uniform response of the whole SD-array to effects in the Scaler rate. *Left:* Mean corrected Scaler rate in 2013. *Right:* Mean r_i from Eq. (4.8).




Figure 4.8: Comparison of the mean corrected Scaler rate for different ways of calculating the variance. Note that the offset between different days arises from excluding different stations due to PMT-quality cuts. We do not use this quantity for any analyses but refer to the rescaled quantities that do not show such behaviour. *Left:* Using the spread of mean rates over stations as estimate of the variance (Eq. (4.10)), it contains mostly systematic effects. *Right:* The variance estimated as the variance of the mean (over stations) in a given time interval (Eq. (4.11)). Here this interval is five minutes. Note the change in scale on the y-axis.

4.1.5 Applied Corrections for Long-term Trends

Before we can use the Scaler data on time scales beyond minutes, we have to account for the known long-term trends caused by detector ageing and changes in the atmosphere. First corrections of these effects have been shown previously [74, 84]. We revisit these corrections to improve where possible and to extend the analysis to low station statistics by defining the corrections on station level.

The first important correction is accounting for the change of rate of measured particles with changes of the atmospheric pressure. It can be easily understood that an increased atmospheric pressure corresponds to higher mass overburden above the detector and thus an anti-correlation with the Scaler rate is expected. This anti-correlation can be fitted by using e.g. a simple linear model $\Gamma = a_0 + a_1(p(t) - \langle p \rangle)$ and then corrected with

$$\Gamma_{\text{corr}}(t) = \Gamma(t) - a_1(p(t) - \langle p \rangle). \quad (4.12)$$

Extending this model to single stations, we observe a correlation of the coefficients a_1 of the stations with their average rate (a_0). The left panel of Fig. 4.9 shows this correlation.

Additionally, such an additive correction is changing the relation of the variance σ_{Scaler} with the mean $\langle \Gamma \rangle$ that is to be expected to follow roughly $\sigma \approx \sqrt{\Gamma}$ under a counting assumption. While, if properly taken into account, this is of no further concern, we can also use a multiplicative correction to avoid these additional tasks.




Figure 4.9: Comparison of different possibilities to correct the Scaler rate for changes in atmospheric pressure. *Left:* Correlation of the two coefficients of a fit of the form $\Gamma = a_0 + a_1(p - p_0)$. Stations with a larger average Scaler count have on average a more negative pressure coefficient. *Right:* Distribution of the scaling factors for a pressure correction of the form $\Gamma_c = a_0(1 - \gamma_p(p - p_0))$ as function of the respective average rate of a station a_0 .

Thus, we fit $\Gamma(t)/\langle\Gamma\rangle = a + \gamma_p(p - p_0)$ for each station and then get the corrected rate as

$$\Gamma_{\text{corrected}} = \Gamma(1 - \gamma_p(p - p_0)) \quad (4.13)$$

and the fitted pressure coefficient γ_p . From the right panel of Fig. 4.9 it is clear that the correlation of the correction with the average rate per station is removed with this ansatz. For data, we therefore use this correction with data fitted to the full year of 2013 in the centre of the data set. Based on the little spread in the γ_p between stations, it might be beneficial for future analyses to test using a single fit for all stations making sure no further systematic effects are introduced with such an approach.

Next, we perform a correction with respect to the Area-over-Peak (AoP) of a station that captures changes in the pulse shape in the SD. The AoP-correction is necessary for long-term studies because it changes on a scale of years. Although such a correction has been reported previously [84], we change it and switch from a linear-fit-based to a model-based scaling approach for a variety of reasons.

Firstly, the fit of an AoP correction is biased due to the coincidental correlation of the main decrease of AoP, starting exponentially in about 2008, with the decline of the galactic cosmic-ray rate due to the solar cycle. Secondly, we observe two populations in the AoP-data, which cannot be fitted with a single linear function. Additionally, with the use of a (microscopic) model for the correction the need for any fit is removed and with this, the danger of removing accidentally coincident signals is drastically reduced.

Our model is based on two assumptions: Firstly, the total available (time-integrated) signal, i.e. number of photons, in the tank scales with the optical quality of the watery liner. Since AoP scales with the reflectivity of the liner, we expect the number of photon-electrons

$$n_{\text{ph}} \approx \int dt s_{\text{ADC}}(t|n_0) \propto \text{AoP} n_0 \quad (4.14)$$

to scale with the integrated signal s_{ADC} and as a function of the number of created Cherenkov photons n_0 . This means that there is a higher probability of crossing the threshold for lower energy particles in a detector with a higher AoP. Secondly, the probability of triggering all three PMTs with a low energy signal also scales with the reflectivity of the liner, again likely proportional to AoP. In total, we therefore expect a scaling behaviour $\Gamma_1/\Gamma_2 = \text{AoP}_1/\text{AoP}_2$.

Based on this, we define a correction factor as the ratio of the current AoP value to a reference value, i.e. $\gamma_{\text{AoP}}(t) = 3/\text{AoP}(t)$. The corrected rate is then given by $\Gamma_c(t) = \Gamma(t)\gamma_{\text{AoP}}(t)$.

In Fig. 4.10 the profiles of the Scaler rate vs. AoP histograms with and without the correction are shown. It is clearly visible that the correction flattens the Scaler rates in this




Figure 4.10: *Left:* Correlation of the Scaler rate with Area-over-peak. The red points represent the profile of the distribution. *Right:* Residual correlation of the AoP-corrected Scaler rate with AoP. Red points indicate the profile distribution.




Figure 4.11: *Upper panel:* Scaler rate of station 751 for a day in April 2012 averaged in 61 s intervals. The intervals correspond to the station calibration intervals. *Lower panel:* Baseline of PMT-1 in the same period showing the characteristic drift with temperature over the day.

phase-space. The resulting line is not perfectly flat, which – in parts – is expected and desired because we do not want to account for coincidental declines of the Scaler rate as mentioned earlier.

4.1.6 Baseline Correction

In this section, we discuss effects related to the drift of the baseline in the stations. The baseline of the PMTs in the field is constantly changing with temperature. Because of the integer nature of FADCs, and thus the thresholds of the Scaler, this can lead to residual changes in the measured rate. Fig. 4.11 shows this effect for a single station by plotting the Scaler rate as a function of time compared with the baseline of PMT-1. It is clear that the shape of the measured Scaler resembles the baseline drift, however, broken into a piece-wise function. These discontinuities come from the online-calibration, which is performed in the station during data taking and which corrects the baseline drift. It can do so only when the baseline changes by more than 1 ADC. In Fig. 4.11 there are more changes of the threshold than visible from the baseline of PMT-1 in the lower panel, because this happens simultaneously for all three PMTs. We outline the correction performed for this analysis in the following.

The first step to correct for these threshold-changes – or even to get a picture as clearly as in Fig. 4.11 – is to find the correct intervals. Because the station performs the online-calibration and resetting of the thresholds every 61 seconds during data taking, it is important to find the correct intervals of constant settings first. The idea we use is to first identify time intervals without a possible station reset and then determine the corresponding calibration phase.

The first part is solved by searching for missing data in the time series. However, we keep in mind that missing data does not correlate 100% to a station reset. In practise, we use a maximal length of missing data of exactly one calibration phase, i.e. 61 s, interpreting zeros in the data stream as missing data.

The second stage is implemented as a minimisation of the mean variance of the Scaler rate in the set of 61 distinct intervals defined by the integer of the GPS second modulus 61. With mean variance defined as

$$\langle \sigma \rangle = \frac{1}{n} \sum_i \hat{\sigma}_i, \quad (4.15)$$

the arithmetic mean of the sample variances $\hat{\sigma}$ in the 61 s intervals.

Fig. 4.12 shows the resulting minimum for an example station. The error bars indicate the variance in the sample of sample-variances and not the error of the mean itself. It can be used to determine instabilities that may affect this phase finding. We use the obtained minimum and in a post-processing step, join continuous intervals based on the determined calibration phase. If two adjacent intervals have the same calibration phase, they are joined to form a single interval with the common calibration phase. In the case where exactly one out of two adjacent intervals has been created with less than 15 sample variances, they are joined with the value of the larger set of intervals. These operations are repeated until no further steps are possible.

With these calibration intervals known, we can try to correct for the residual drift. In the following, we use the notation

$$b_i = \text{baseline}_i - \lfloor \text{baseline}_i \rfloor \quad (4.16)$$

for the non-integer part in individual PMTs and

$$\Delta_b = 1/3 \sum_i (\text{baseline}_i - \lfloor \text{baseline}_i \rfloor) \quad (4.17)$$

for the average non-integer part of the baseline.

The general idea of our correction is that after accounting for all station-specific (calibration) effects, all SD stations should behave universally. Motivated by this ansatz, a function of the form

$$\Gamma = \Gamma_{\text{ref}} k_{\text{nius}} \prod_i (1 + k_i \Delta b_i) \quad (4.18)$$

is fitted with a χ^2 -minimisation. In this ansatz the baseline coefficients k_i for the i^{th} PMT, and the overall scaling, k_{nius} , are free parameters. k_{nius} includes also known dependencies not included in this model for now, e.g. the altitude of a station, or the VEM-peak position.

The resulting parameters k_i fitted in April 2012 over all stations are shown in the left panel of Fig. 4.13. It is clearly visible that the distributions peak at a universal value.

The (large) spread of the distribution can be associated with the large correlation of the baseline drifts and the resulting instability of the fit. The right panel of Fig. 4.13 is indicative of these correlations, because the phase space of the b_i is not completely covered but rather populated along lines of correlated increase of the b_i . The baselines are correlated because of the common origin of the drift: the ambient temperature.

Motivated by these peaks at $k_i \approx 0.05$, we develop a new model that has just a single free parameter: k_b the scaling constant for all PMTs and stations. It is determined by minimising the variance of all Scaler rates as function of k_b , where the distribution of the mean Scaler rate

$$\hat{\Gamma} = \Gamma_i(t) \frac{1 + 1.5k_b}{1 + b^{(i)}k_b} \gamma_{\text{AoP}} \quad (4.19)$$

is used to calculate the variance. The AoP correction factor γ_{AoP} is optional and follows the method described in Section 4.1.5. $b^{(i)}$ refers to the sum of the b_i of a station. Scaler rates from a station that has a reported baseline value closer than 0.05 to an integer value at that time are not used in the fit to avoid influence of the outliers visible in the right panel of Fig. 4.13. These arise from the different time intervals of transmitted monitoring data: the baseline-values are taken from monitoring with a time-interval of about five minutes while the calibration can reset the threshold up to 5 times in this interval. The resulting value is very close to $k_i \approx 0.05$ and slightly dependent on time and available stations. To avoid overfitting, we simply chose $k_i = 0.05$ for the correction and proceed with corrections on individual station level.

We add an algorithm to search for these outliers, by using a low-pass filter

$$\langle \Gamma(i) \rangle = \alpha \langle \Gamma(i-1) \rangle + (1 - \alpha) \Gamma(i) \quad (4.20)$$

with $\alpha = 0.1$ to get a prediction in the next 61 s-interval $\langle \Gamma(i) \rangle$ based on the previous ones that is then compared to the one obtained from data. In case of an absolute deviation greater than 10σ , where σ is the estimated error of the mean, the 61 s-interval is removed from future analyses. The threshold is chosen fairly high to avoid too many false positives. One should keep in mind that the number of comparisons over 1600 stations and 13 years of data is very large.

Fig. 4.14 shows the performance of this method by showing the Scaler rate as a function of Δ_b (Eq. (4.17)) before and after applying Eq. (4.19). To remove the residual dependence in some stations, like station 750 in this example, we use a simple linear correlation analysis. We create a single variable from the three PMT baselines b_i using Eq. (4.17) and fit

$$\Gamma = \Gamma_0 + \alpha_b \Delta_b \quad (4.21)$$

only in regions where the distance of Δ_b to an integer value is large. These regions are already preselected in the right panel of Fig. 4.14 by requesting all three PMT-baselines to have a distance to the closest integer of more than 0.05. Using $\Gamma_c = \Gamma - \alpha_b \Delta_b$ results then in a final corrected value that we use in the analysis. As for the pressure correction we fit the pressure correction with data from 2013. We observe more than 90% of stations with a residual dependence of less than 100 Hz / FADC, compared to the e.g. 300 Hz visible in the left panel of Fig. 4.14 this shows the success of the correction.

4.2 Results

In this Section, we present first results obtained with the data set prepared by the procedure described above. We highlight that after the performed selections the Auger Scalers have potential to be used to study solar effects on the cosmic-ray rate on scales from minutes to years. For future analysis, this can be seen as a starting point to increase the systematic understanding of the effects visible in the Scalers but also to understand remaining uncorrected effects and differences to established measuring methods like neutron monitors. The detailed evaluation of all the systematics associated with the Scalers is then the next step in providing a statistically powerful data set which can add complementary information




Figure 4.12: Distribution of the average variance in 61 s intervals for station 73 as function of the phase. The error bars indicate the standard deviation of the used sample. A clear minimum at GPS second modulo 61 = 32 is visible.




Figure 4.13: *Left:* Fitted coefficients for baseline non-integer parts for all stations separated by PMT-number in April 2012. *Right:* Visualisation of the dependence of the Scaler rate on the baseline non-integer parts. The colour and the size of the markers indicates the Scaler rate measured in a 61 s interval while the b_i are defined as $b_i = \text{baseline}_i - \text{round}(\text{baseline}_i)$.

compared to neutron monitors by opening a different energy window into the low-energy cosmic ray physics.

4.2.1 Forbush Decreases

In this first Section of results, we focus on short time scales and transient events. The classical case of transient events related to solar activity are Forbush decreases. It has been observed and analysed before [76] that after applying a pressure correction the Auger Scaler rates are sensitive to Forbush decreases.

With the updated analysis presented in this work, these observations are confirmed and profit from the updated corrections. In the right panel of Fig. 4.15 an example of a Forbush decrease, as observed with this new Auger Scaler data set, is shown. For comparison, data from three different neutron monitors [93], Athens [94] in the northern hemisphere, McMurdo [95] in the Antarctic, and Tsumeb in Namibia [96], are shown for the same time period. The Forbush decrease during mid of June 2015 is clearly visible in all stations, with




Figure 4.14: Comparison of the dependence of the Scaler rate on the non-integer part of the baselines. We use data from station 750 in June 2019 for this illustration. *Left:* The uncorrected Scaler rate as function of Δ_b from Eq. (4.17). Outliers are caused by wrongly “assigned baselines” that come from the unknown values of Δ_b . *Right:* Scaler rate corrected with Eq. (4.19) as function of Δ_b . A small residual dependence is visible and corrected for on station level. Data with a distance to the closest integer of the baselines of less than 0.05 are not used for the fit and are not shown here.

the Auger Scaler rate showing the smallest relative amplitude due to the high magnetic cut-off.

Besides the Forbush decrease itself, there are more features in the neutron-monitor time-series which can also be observed in the Auger Scaler rate. One example is the enhancement in the rates during the end of July 2015. This correlation with the observations of neutron monitors shows the potential of the Auger Scaler rate to complement neutron-monitor data at high-rigidity cut-off in the southern hemisphere.

We want to stress that while the amplitude of Forbush decreases in the Auger Scalers is smaller than in neutron monitors, the statistical power is high enough to resolve them. Due to the relatively strong diurnal variation seen in the Scalers (c.f. Section 4.2.3) the automated detection of Forbush decreases is difficult and subject of future work. In the example in Fig. 4.15 also another possibility for future improvement is visible: during a period just before the decrease on 14 June 2015 data is missing. While these periods with missing data are not significantly interfering with long-term analyses, they should be investigated further for analyses on short scale transients to potentially recover these data.

4.2.2 Frequency Spectrum

As pointed out in the Introduction, the intensity of cosmic-rays with energies lower than 100 GeV measured at Earth is modulated by solar effects. In this part, we want to briefly discuss the periodic parts of this modulation as observed in the Auger Scaler rate.

The left panel of Fig. 4.16 shows the discrete Fourier-series expansion of the observed Scaler rates in logarithmic scale. The sharpest peaks in the spectral power are visible for a signal with a frequency of one solar day and the corresponding overtones. This diurnal signal is due to the Earth’s rotation in the magnetized solar wind and has been observed in the Auger Scalers before [83]. A more detailed analysis of this particular signal is presented in Section 4.2.3.

Another expected periodicity of the cosmic-ray intensity is associated with the rotation of the Sun and corresponding magnetic features. Therefore, an enhancement in the spectrum at frequencies corresponding to the Carrington period (about 28 days) is expected. On the logarithmic scale of Fig. 4.16 this is not clearly visible. However, small enhancements at the indicated range are visible and the frequency spectrum is not expected to show a clear peak

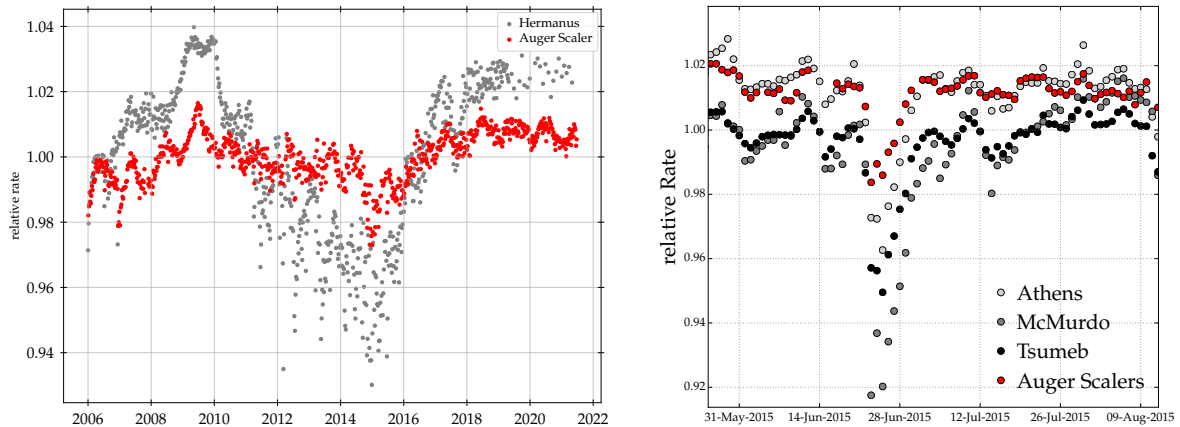


Figure 4.15: *Left:* Long-term comparison of the measured rates of the Hermanus [97] neutron monitor with the corrected Auger Scaler rate. Data from 2006 until the end of June 2021 is shown, averaged in 7 day intervals. *Right:* Comparison of the measured Auger Scaler rates during a Forbush decrease with data from the Athens [94], McMurdo [95], and Tsumeb [96] neutron monitors.

at the Carrington frequency, because of the contributions from the different frequencies of individual transient events.

Because the rate of magnetic features of the Sun correlates with the solar cycle, a time dependence of the spectral power of these frequencies is expected. The right panel of Fig. 4.16 shows the integrated spectral-power for Fourier modes with periods between 20 and 70 days. For reference, the number of sunspots is shown during the same period as well. The two minima and the solar maximum in sunspot number also appear as two minima and a maximum of the spectral power. But, as expected from the presence of non-periodic solar events like Forbush decreases, there is no one-to-one correlation of the spectral power with the sunspot number. However, the general trend proves that the data set is consistently observing the magnetic features over the whole data range and thus can be used in future work to extract more detailed information.

In the long-period domain an enhancement of the Fourier modes with a period of 365 days is visible, corresponding to the movement of the Earth around the Sun but potentially also to instrumental effects. A change in the average temperature in the atmosphere could lead to a seasonal modulation similar to the one observed due to the changes in the Molière radius, however, we do not see any correlation of the Scaler rate with temperature. This does not exclude a correlation with temperature at higher altitudes, that can influence the rate of muons due to the low energy of the primaries measured by the Scalers. Additionally, we can think of systematic effects due to the influence of stations with unstable electronics or PMTs in summer. From Fig. 4.3 we see that there are many instrumental effects with a strong seasonal modulation. These effects might propagate into the averaged rate by contributing only in certain seasons. For future analyses, this means that we have to ensure all stations contribute similarly by imposing corrections that remove the differences between them as much as possible.

On scales beyond a year, we cannot resolve the 11-year solar cycle with this Fourier analysis with the length of the data set available. A time correlation based approach for this observation is presented in the Section 4.2.4.

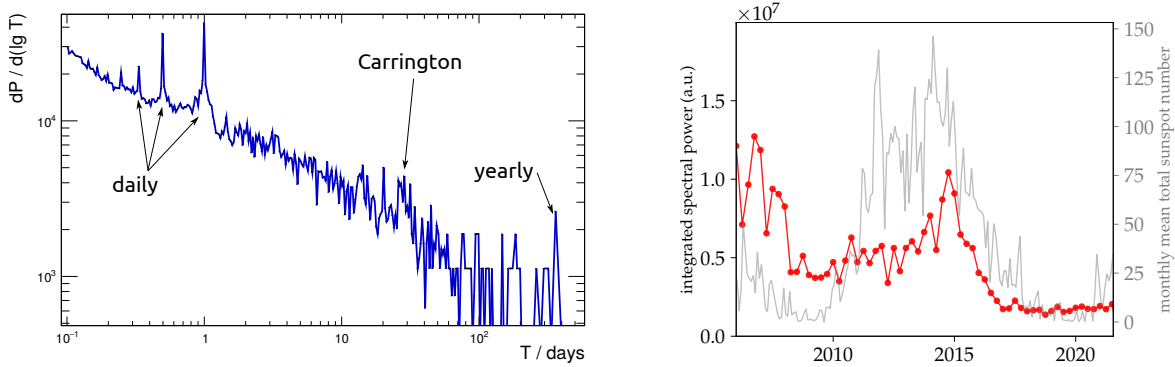


Figure 4.16: *Left:* Observed frequency spectrum of the data with annotations of frequencies of special interest. *Right:* Time dependence of the integrated spectral power around the Carrington period $T \in [20, 70]$ days including overtones. For reference, the solar sunspot-number data from Ref. [98] is shown in gray.

4.2.3 The Diurnal Signal

Following the clear indication that there is a diurnal signal in the Fourier analysis, we present a more detailed analysis of this signal in this Section. The rotation of the Earth in the magnetized solar wind gives rise to a diurnal change in the measured rate of particles at the ground. Due to the tilt of the Earth's rotation axis and the location of the Observatory at latitude 35° south, as well as the minor changes in distance to the Sun, yearly modulations of this signal are also expected.

The data set of the corrected Auger Scalers show these modulations that are not correlated with pressure or temperature very clearly. The left panel of Fig. 4.17 shows these daily modulations of the measured Scaler rates. For each month, we average the Auger Scaler rates in five-minute bins in time of day (UTC) and rescale with the average Scaler rate in the respective month. This rescaling emphasizes the shape of the diurnal signal, with an amplitude of about 1%, instead of highlighting the long-term variations caused by the 11-year solar-cycle which is measured with an amplitude of about 4%.

From Fig. 4.17 the seasonal variation of the diurnal signal is clearly visible. The peak position of the signal shifts by approximately three hours over the course of a year. In January the daily signal peaks at about 14:00 UTC, whereas in austral winter the peak moves to 17:00 UTC.

During the solar maximum in the years from 2012 to 2014 the daily signals are less similar. In this period the regular structure is washed out by a higher rate of transient events and higher variability. It might be possible to recover the signal if a correct filtering of periods with known transients is applied.

The long-term stability of the daily profiles can be derived from comparing the daily profiles across the 15 years shown in Fig. 4.17. Apart from resolution effects from the SD still being deployed during 2006, the profiles of 2006 match those measured in 2021.

With the high statistics available at Auger, the diurnal modulation of about 1.2% peak-to-valley is resolved every day. This is also the reason for the strong enhancement of the spectral power for frequencies of one per day and its overtones in the Fourier spectrum, shown in the left panel of Fig. 4.16.

In the right panel of Fig. 4.17 this ability to resolve the signal is visualized by comparing the average profile of June 2018, drawn in black, with the profiles of two individual days of this month, in red and blue.

The data of the individual days are rescaled with their average rate, to enable the comparison of the shapes. For both chosen days the diurnal variation is clearly visible. However, the

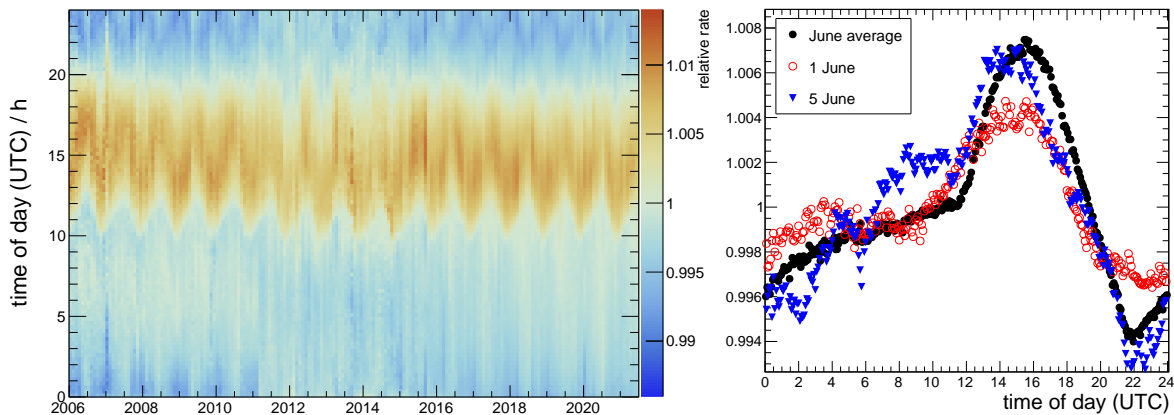


Figure 4.17: Observed daily pattern of the Scaler rate as a function of the time of day (UTC). *Left:* Monthly average values: Each month is normalized separately to emphasize the shape rather than the long-term trend observed in Fig. 4.15. *Right:* For June 2018, the average daily shape of the observed diurnal pattern of the Scaler rate together with two single days as examples of the variations in this diurnal pattern.

amplitudes and details of the signals are distinctively different: the peak-to-valley amplitude for June 1st with 0.7% is smaller than in the average profile, which shows an amplitude of 1.4%. In contrast to June 1st, the peak-to-valley amplitude of June 5th is very similar to the average profile, but this day also shows that with the high statistics of the Auger Scaler rate, other features, only present in singular days, can be resolved. In the example of June 5th there is an additional rate enhancement at about 8:00 UTC, as well as two additional depressions of the rate at 2:00 and 6:00 UTC.

This comparison highlights the abilities of the Auger Scaler rates in resolving structures on hourly scales with high statistics. This is the basis for future analysis of the rates on these scales also during transient events like Forbush decreases.

4.2.4 Observation of the Solar Cycle

In this Section, we want to show that the Auger Scaler rate can be used for long-term measurements of solar modulations of the cosmic-ray intensity as it was shown in previous work [84] on scales containing more than a full solar cycle (11-year). To achieve the necessary long-term stability of the measurements, the corrections for the detector response as presented above are essential.

The left panel of Fig. 4.15 displays the measured Scaler rate from January 2006 until the end of June 2021. For comparison, the rate as measured by the Hermanus neutron monitor [97] is shown during the same time period.

Over the whole period we observe a variation in the average Scaler rate of about 4%. With a maximum of +3% (relative to January 2006) in 2009 and a minimum with -1% in 2015. In about 2018 the amplitude reaches again about +3% (relative to Jan 2006) during the current solar minimum. The correlation with the neutron monitor - as well as the match with the expected behaviour from solar sunspot data [98] - showing a recovery of the Scaler rate towards the second solar minimum during the operation of Auger, highlights the long-term stability and reliability of the corrected Scaler rate. Even the flatter maximum in the rate, as seen by neutron monitors is visible in the data too. However, the difference of fairly strong annual modulations discussed in previous sections remains. Future work is necessary to fully understand these modulations and determine whether they are of instrumental or physical origin.

4.3 Summary and Outlook

In this chapter, we presented analyses using the present data taken with the Scalers. With new corrections applied to the data, we see clearly the potential for contributing to measuring the solar influence on low-energy cosmic-rays with the Scalers.

The key for being able to obtain stable measurements of the flux of low-energy cosmic rays with the Scalers on scales of years is a careful selection of data and stations as well as the appropriate corrections. We show that we are able to perform such corrections and to obtain a data set suited for physics analysis. Bearing in mind that this is a first application of such an analysis on the full 15 years we observe solid performance on the full time scale. However, it is clear that the number of stations left in the analysis is low, especially in the last years. Thus, the details of the cuts could be revisited in future analyses to recover as much statistics as possible without introducing systematic biases. New corrections, for example making use of the online-calibration information, can help in this task. Such future corrections can also help detecting the origin of the seasonal modulation visible in the frequency and long-term analysis. We expect that these remaining issues can be resolved in the near future.

Despite these remaining limitations, we can present a long-term measurement that shows which we do observe the 11-year solar cycle and see a recovery of the signal in the last solar minimum proving that the corrections perform well. The observations of Forbush decreases show that the obtained data are stable on scales of several days, proving the same for the pressure correction, and proving that the corrected data are sensitive to such physical variations.

For the very short time scales, we introduce a new correction to reduce the impact of the drifting baseline which cannot be corrected online due to the integer nature of the threshold. This correction reduces the spread across different stations and is crucial e.g. for the search of excesses on time scales of minutes as it was for instance done in Ref. [86]. We believe this novel correction is a very good starting point to facilitate new analyses focusing on short time scales by reducing the systematic variance considerably.

For the future with the upgraded electronics (UUB), we can learn from this analysis and especially from the discussion around the baseline correction on how the Scaler could be improved. The first question for the future of the Scalers with the UUB is whether or not the focus is on gamma-ray search or solar physics. It has been shown already previously that the upper limit for the amplitude introduced for the search for gamma-rays [99] is decreasing the sensitivity of the Scalers for e.g. Forbush decreases [74]. If the main-target of the data set is solar physics, a removal of the upper threshold seems favourable because a discontinuous data taking is given due to the change of the electronics in any case. At the same time, there might be a possibility to use the increased computing capabilities of the UUB to directly correct the baseline online in a way as presented in Section 4.1.6. The details will depend on the exact implementation of the Scaler in the UUB and the implementation should be targeted in the near future during the deployment phase of the UUBs.

In summary, we show that the Scalers have the potential to be a high-statistics data sample at low energies using a different approach than neutron monitors if all the known effects are corrected for. We think this analysis can be seen as a starting point for the future, formalising all corrections and unifying the approaches. In the near future the upgraded electronics will open new possibilities to use the gained knowledge to further improve the measurement of low-energy cosmic rays at Auger.

Surface Detector Events associated with Lightning

In this Chapter, we will explore the possibilities that the Auger SD offers for measurements in thunderstorms. We start by analysing exotic events observed accidentally during thunderstorms that show characteristics of relativistic particles, arriving at the SD stations, of not yet known origin [10]. We will discuss how peculiarities in the footprint of these events are expressions of the trigger system of the Observatory which is not optimised to work during these conditions or for these events. Furthermore, we will present possible extensions and modifications of the data acquisition system (DAQ) that can help to record more information more frequently about these interesting exotic events. In addition, we discuss how with the upgrade of the SD electronics for AugerPrime [4] also new possibilities for the detection of such events arise. By highlighting and conceptualising possibilities of efficient triggering on the signals of these events we argue on how such new detection channels can be designed.

In the second part of this Chapter, we search for the possible causal relation between cosmic-rays and the initiation of lightning [53]. By extending the available data down to trigger data of single stations, we can significantly lower the energy threshold of the observations with the Auger SD while still covering a large area in a region with strong lightning activity during the austral summer. The lower threshold makes energies relevant for lightning initiation accessible, increasing the chances of answering the question whether or not cosmic rays contribute or are crucial for lightning initiation.

The contents of this Chapter were by the author in parts previously shown in several Auger collaboration internal notes [100–103]. In addition, the results presented in this thesis contributed significantly to the work published as Ref. [104].

5.1 Observations of SD-Rings and SD-Discs

We start this Section with a brief introduction to the terminology and phenomenology of peculiar events observed during lightning conditions in the Auger SD. Fig. 5.1 [104] shows typical SD signals for, both normal cosmic-ray data taking and so-called lightning data. While the cosmic-ray signals are characterised by short pulses of various amplitudes with all PMTs measuring roughly the same signals, the typical case for events triggered by lightning activity is different. Panel *d* of Fig. 5.1 shows an example of the characteristic high-frequency noise measured in all PMTs. The signals for the three PMTs are clearly not the same. These signals are caused by the electric fields picked up in the PMT cables, and are different due to the different orientations of the cables which here act as antennas.

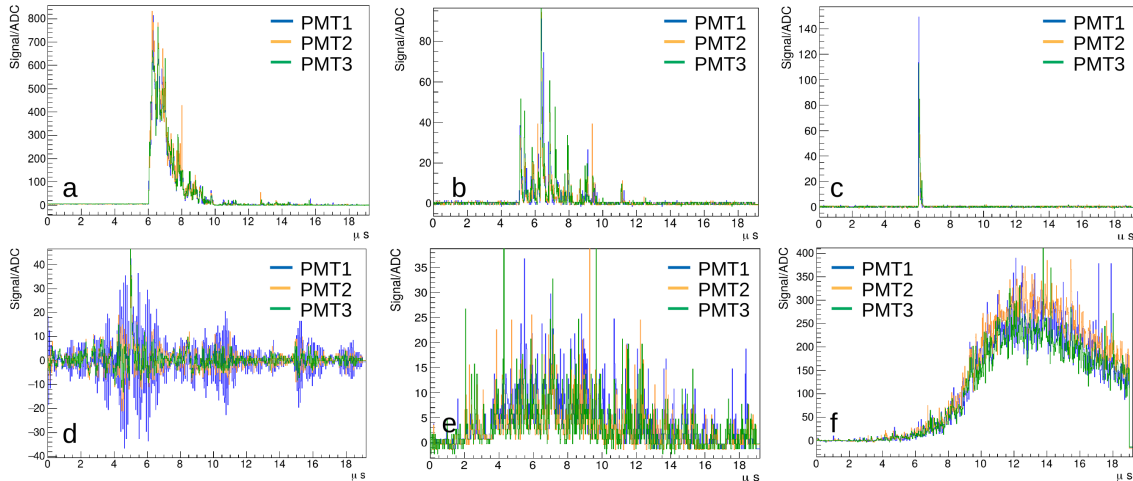


Figure 5.1: Illustration of typical signals in the surface detectors: *a to c* are signals from cosmic-ray showers at different distances to the core. *d* is a typical example trace caused by RF noise picked up in the PMT cables during lightning activity. *e and f* show examples of the so called “long-signals” rarely observed during thunderstorms, which are due to the signals in detectors rather than the pick-up in the cables. The two examples differ in amplitude but are thought to be originating from the same type of events. Taken from Ref. [104].

In contrast to such spurious signals, the traces of the peculiar events we discuss in this Chapter, are caused by signals in the SD stations [10]. The conclusion of real signals in the tank is based on the observation that all the three PMTs measure the same signals (within fluctuations), as evident from the examples in Fig. 5.1 *e*, and *f*, and also because the independent low-gain channels agree perfectly with the signals measured in the high-gain channels [10].

On the ground, these events are characterised by the large and roughly circular shape of the triggered stations. Fig. 5.2 shows three examples out of the total 23 identified events [104]. The varying sizes also reflect the amplitude variation shown in Fig. 5.1. However, the remarkable feature of most large events is the apparent lack of stations in the central part of the event footprint. This peculiar shape leads to the name “SD-rings” for these types of events, and their long-lasting signals are referred to as “long-signals”. Especially in the comparison with the signal produced by single muons (panel *c* in Fig. 5.1) it is clear why these signals covering roughly $10 \mu\text{s}$ are referred to as “long”. To distinguish the peculiar ring shape from those events with long-signals that have a filled circular footprint, the term “SD-disk” is used. One example of such an event is given as the left most panel of Fig. 5.2.

The physical origin of these signals is unknown. Similarities with events recorded at the Telescope Array (TA) detector [9] that are identified as gamma radiation suggest that they are originating from TGFs too. However, the signal amplitudes and the extent of the events are different. The energy deposit is much larger than in the TA events but matches that measured with NaI scintillators in rocket-triggered lightning related TGFs [59, 104]. Both detectors, water-Cherenkov and NaI-crystals, have a high detection efficiency for gamma radiation making a difference in signal amplitude relative to the plastic scintillators plausible.

However, for the central region of TGFs no lack of signal is known in the form observed in the SD-rings. In the following, we will thus use the trigger information to try to answer the question whether or not these holes in the footprint are of physical origin or rather due to the DAQ of the Auger-SD.




Figure 5.2: Footprints of example peculiar events observed in the SD. Circles mark stations with “long-signals” (c.f. Fig. 5.1 e, f) with the size corresponding to the signal size and the colour representing time (blue early, to late in red). So called lightning stations are marked with stars and muon/cosmic-ray stations with crosses. A ring-like shape for these events is evident, leading to the naming “SD-ring”. For smaller events, the ring is filled creating “SD-disks”. Figure is taken from Ref. [10].

5.2 Identification of Trigger Effects on Lightning Events

We start the discussion on the SD-rings with trying to understand how the SD-trigger performs in lightning and high-rate conditions. Thanks to a new data set available and presented in Section 5.2.1, in conjunction with an SD-ring event recorded in May 2017, we are able to discuss in detail the effect of the trigger on what is recorded.

5.2.1 Available Data

For the data used in this work, we use the so called “T2Dump” files. These files store all the information about the SD T2 triggers available for the building of the T3s. For each GPS second, these consist of a list of triggers, defined as a trigger time in μs , the station ID, and the trigger type. The trigger time in these data is the end-of-trace time rounded down to μs . The trigger type discriminates threshold from ToT triggers, or better, threshold from all other (new) types. In the most recent version of the station software also the types of ToT, ToTd, and MoPS can be distinguished. These data are available since January 2016.

In the following, we use different samples from this data set for different parts of this work adjusted to the requirements. For instance, for the analysis of the minimal time difference between triggers in the next section, we chose 18 May 2017 01:00-02:00 (UTC) because of strong lightning activity leading to high trigger rates. The high trigger rate is an ideal testing ground to check for dead times of the station. In addition, there is an SD-ring [105] which makes the study of short-term trigger activity especially interesting. In general, using all available trigger data is difficult because of the large amount of data. For every hour of normal SD-operation, the T2s form about 500 MB of compressed binary data, exceeding the total amount of regular SD-data by a factor of 20.

In all analyses of this work, we apply sanity cuts on the input data. If the GPS microsecond of a trigger exceeds the valid range of $[0, 1\,000\,000]$ or the station sends undefined trigger flags, we reject all triggers from the station for the whole GPS second. The reason for rejecting the whole second of triggers is the communication system that sends all triggers of a GPS second in a single package. Thus, the conservative approach to keep the sample clean is to reject the whole package, i.e. second, if indications of corrupt data are present. The overall rate of such invalid triggers is very low (a fraction of about 10^{-6}). In Fig. 5.3 we show the microseconds and flags of removed triggers during a full hour of data (01 Jan 2019 23:00




Figure 5.3: Illustration of rejected T2s from the T2-dump files during 01 Jan 2019 23:00 (UTC). The trigger microseconds are shown on the y-axis, with some exceeding the 1 000 000 value leading to rejection. Valid trigger flags are $\{1, 2, 8, 9, 10, 11, 12, 13, 14, 15\}$ and values other than these are rejected. Trigger flag 7 indicates scaler data and is removed already during the reading of the data. In total 124 T2s are rejected out of 1.29×10^8 in one hour of data.

UTC). Only 124 triggers, out of 1.29×10^8 are removed, indicating that the station triggers are, as expected, already a clean data sample.

5.2.2 Station Dead Time

In this Section, we present results related to the minimal possible amount of time difference between consecutive triggers in a station. This includes effects of dead times and can help understand why, for example, some stations in the centre of the so-called SD-rings are missing in the event data. In relation to this topic, we use here a data set that includes an SD-ring and with that also a lightning period.

To determine the minimal time difference we collect for all stations the time difference between triggers. We also include higher orders, i.e. the time between a trigger and the next-but-one-th trigger, here denoted as Δt_{13} . In general the quantity Δt_{1j} denotes the time difference of a given trigger of a station and its next j^{th} trigger.

Fig. 5.4 displays the histograms obtained by summing these quantities over all stations in this one hour of data. On long time scales (left panel) the expected behaviour is clearly visible and can be described as

$$p(\Delta t; n = j - i) = \frac{\lambda^n (\Delta t)^{n-1}}{(n-1)!} \exp(-\lambda \Delta t) \quad (5.1)$$

with the trigger rate λ . The fit of the rate and the overall normalisation is performed only on the Δt_{12} data. For the higher orders, the values are fixed to those values and only n is adjusted.

In the right panel of Fig. 5.4, the behaviour at small times is shown by zooming into the first 25 ms of the ratio of observed over expected counts. Due to the presence of lightning in




Figure 5.4: *Left:* Time differences between a trigger and the next i^{th} trigger per station summed over all stations, shown for $i = 2, \dots, 6$. The expectations (solid lines) are computed using Eq. (5.1). *Right:* Zoom into the first 25 ms of the same data. The logarithm of the ratio of the observed data to the expected is used to display all orders on a similar scale.




Figure 5.5: Illustration of the amount of dead time as a function of the number of consecutive triggers. The numbers are estimated from the histograms of Fig. 5.4. A linear fit, excluding the lowest order, is performed and the parameters are given in the parameter panel.

these data, there are large excesses of triggers at short times. At times < 2.5 ms sharp cut-offs are visible. Those indicate dead-times of the station triggers.

To estimate the earliest time when new triggers are possible we determine the centre of the bins of the histograms with more than five entries. We do not want to take single entry bins, as the results might be influenced by a single corrupted package of triggers that might pass the filters.

We display the result of this determination in Fig. 5.5 together with a linear fit to the higher orders. It is clearly visible that after the initial $20 \mu\text{s}$ veto period¹ a constant scaling of about $300 \mu\text{s}$ per additional trigger is necessary to trigger again. The interpretation of this is that this dead-time is related to the copying of the ring buffers into the event buffer. With the upgraded electronics (UUB) this should be less of a problem due to the vastly increased computational power of the board.

¹ We want to stress that technically this is not a dead time, because in principle all parts from this time are stored in the previous event trace.

5.2.3 The used Events

Previously, the study of SD-rings with T2-dump files proved to be difficult, as there were no known events with such long-signal traces in the period with the T2-dump data available which could be used to identify potentially new events.

The presented analysis is made possible by the identification [105, 106] of two events with long-signal stations in the period starting in January 2016 which have T2-dump data available.

In the following, we describe briefly the chosen events and their apparent topology.

The most important event for this study is a lightning event occurring on 18 May 2017 at the GPS second 1179105785. It consists of several raw SD events, with the majority of the long-signal stations being recorded as SD event 42852075. For this event, Fig. 5.6 shows the distribution of triggered stations on the ground. Additionally, an example of a long-signal trace is given from station 1521 with its position in the footprint marked with a crossed circle in Fig. 5.6. The characteristic feature of this event topology – which also has been seen in previous studies and different events [105] – is the “missing” centre. This lack of triggered stations in the middle of such events persists even when merging all raw events and is also not an effect of the merger². To illustrate the enormous scale of the event, we give an overview of the activity of the triggers at CDAS in a few numbers:

- 591 requested traces,
- 11 T3s,
- 278 traces read out
- 994 T2 triggers in 10 ms.

The amount of triggers has to be compared with the background expectation of

$$N_{\text{bg}} \approx 20 \text{ Hz} \times 1600 \times 0.01 \text{ s} = 330, \quad (5.2)$$

and the number of T3s with the normal T3 rate of about 0.1 Hz. Further details on the event structure are given in the Section 5.2.4, together with the explanation of how to classify the observations.

In the data of January 2017 another event with long-signal stations was identified [106]. On 22 January 2017 at GPS second 1169159329, SD event 41257078 has several long-signal stations. Fig. 5.7 shows the distribution of triggers of this event on the ground. For station 1779, a trace of PMT 1 in the high-gain channel is also shown. For reference, we give the number of requested and recorded traces also for this event:

- 735 requested traces,
- 2 T3s,
- 720 traces read out,
- 1082 T2 triggers in 10 ms.

There are several differences compared to the SD ring event shown in Fig. 5.6. Firstly, there is no hole in the centre of this event. Also, most of the stations marked as “T3 lost” in Fig. 5.7 have data recorded in a successive event. The lack of the hole in the centre makes it also a candidate for the class of “SD-disks” [107]. In addition to the missing centre, the amplitude

²We want to emphasize that the event merger is not expected to work perfectly for exotic signatures with signals that are longer than the trace length.

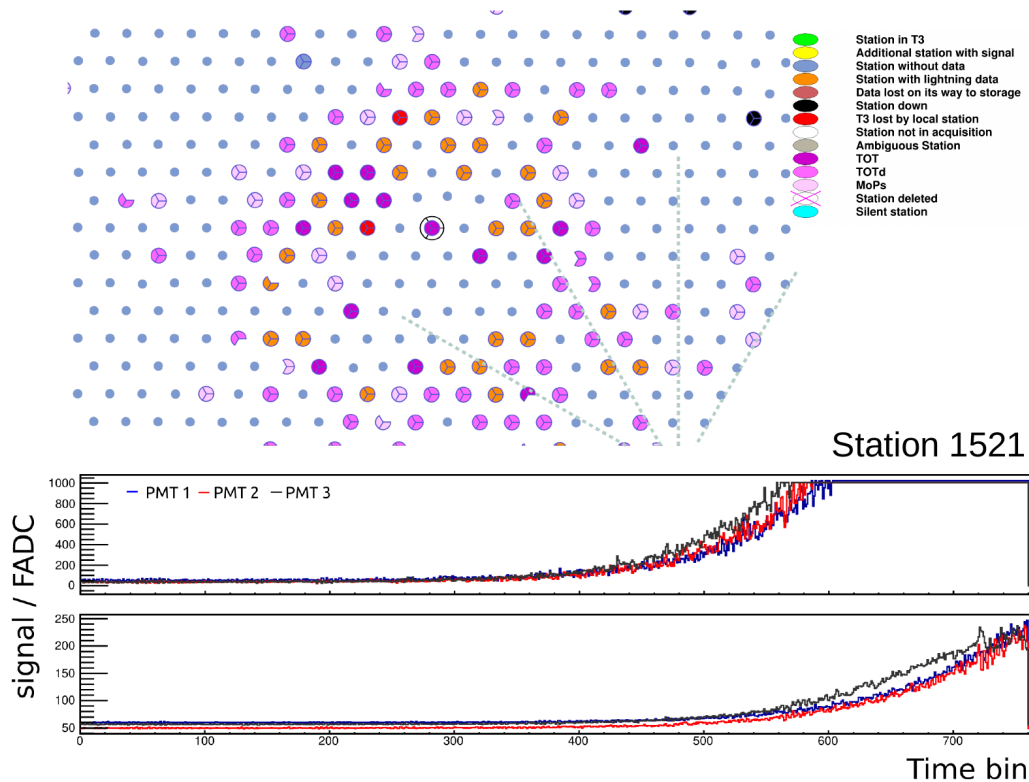


Figure 5.6: Display of the triggered stations of event 42852075. A hole in the middle of the event indicating the “ring” structure is visible. Also in subsequent events within the same GPS second no data from those stations are recorded. Example traces of the three PMTs of station 1521 are given in the lower panel. Both high-gain (upper panel) and low-gain (lower panel) channels show the same behaviour in all three PMTs of a very large and long-lasting signal.

of the long signals is different to the ones typically observed in the SD rings. The signals in the SD disks are generally of smaller amplitude, as it is visible when comparing the example traces in Fig. 5.7 and Fig. 5.6.

Furthermore, it should be noted that for the SD disk the overall level of trigger activity in the array is not smaller than for the SD ring. In the 10 ms interval around the event centre, there are some 100 T2 triggers more across the array when comparing to the SD ring event used in this note. However, this higher number of triggers arises from more coherent events, such that they are read out (more successfully) in only two T3s, compared to 11 for the SD ring.

5.2.4 Analysis Strategy

In this section, we briefly describe the methods used to gather, match, and merge the available data to allow drawing conclusions from these singular observed events.

We try to use all available data to understand what might cause (un)physical holes in the SD Ring events. Most of this information is part of two sources of data: the SD data (`sd*.root-files`), and the T2-dumps.

The SD data contains all read out traces and therefore is the way to identify SD-rings with the long traces themselves. Additionally, the lightning traces can help identify the relative timing of long signals and one or multiple RF pulses from lightning strikes. The last important part of these data is the error flag. It indicates lost data, so it is possible to

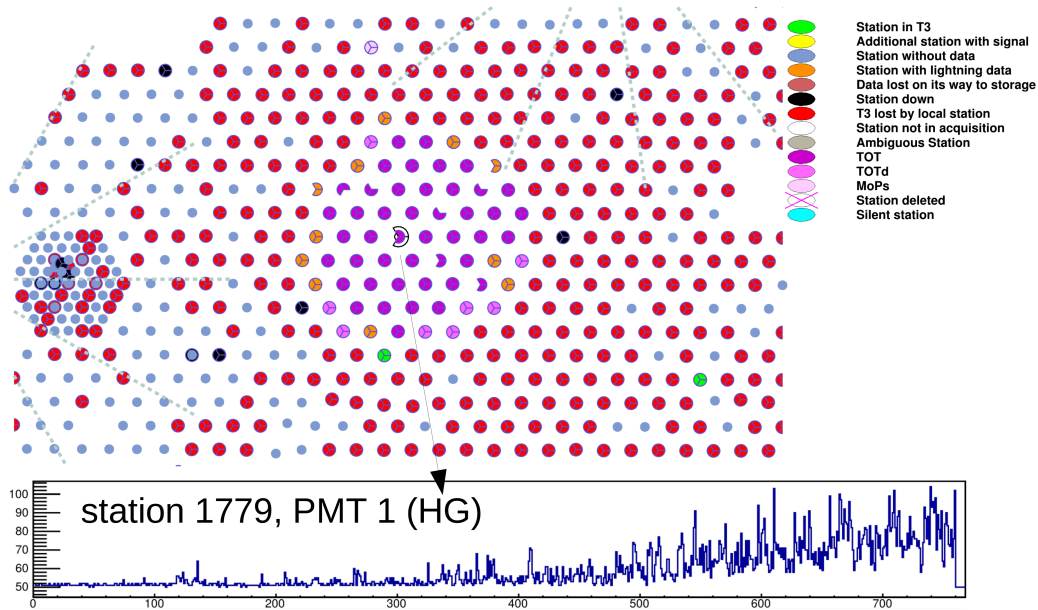


Figure 5.7: Display of the triggered stations of event 41257078. Stations in the middle of the event have long signals, as shown for the example with station 1779 in PMT 1. Note that the overall scale of the signals is smaller than for the event shown in Fig. 5.6. The stations marked as “T3 lost” in this display are in fact read out and the data are available in the previous event. We opt to show the raw data, as there are stations appearing twice in this event and therefore the merging of these special events is more delicate than in the usual conditions.

determine whether data are lost due to bandwidth or timing issues, or, if a corresponding trace was not even requested.

The information about triggers that were not requested can be reconstructed with the help of the other data set. The T2 dumps, first used in [108] and presented briefly in Section 5.2.1, store all T2 messages received at CDAS with information on the station, trigger type, and the GPS microsecond. In addition, these data can be used to derive dead-times of stations based on previous triggers as we have shown in Section 5.2.2.

For cross-checks of the time ordering of T3 requests we also use the Ik logs. To identify potentially interesting clusters of triggers in the T2s that are not a T3, we use an analysis based on spatial and directional clustering, further detailed in Section 6.3 with minor modifications. We allow propagation speeds slower than the speed of light for nearest-neighbour stations in the clustering, to account for possible slower effects from lightning.

The goal of this work, is to combine all these data into a single structure to facilitate an easier overview of what happens during these events. As basis of this data structure we chose the T2s extracted from the T2 dumps because these represent all available (T2) triggers and the higher level data affects only a (strict) subset of these triggers. T1-triggered stations are technically not a part of this, however, they are not decisive for the triggering of SD rings (or T3s in general) and are thus not a priority here. We still keep the trace data obtained via T1-triggered stations by inserting an invalid T2 with trigger flag “0” into the data.

In summary, the methods of this work aim at answering the question whether a certain T2 was requested by an event (and if so by which one), has returned data, has a certain type of trace, or is part of a T2-cluster. In the following, we detail the steps in the process of answering these questions.

Trigger Extraction To limit the amount of data, we chose a data set of interest consisting of all triggers ± 5 ms around an (arbitrarily) defined event time. For the SD-ring of 18 May 2017,




Figure 5.8: Representation of all SD triggers around the SD-ring and SD-disk events studied in this note. Different kinds of trace types and lost data are indicated by different markers. For details see text. *Left:* Data for the period around the SD-ring event 42852075. The event time is defined as microsecond 94 000 of the GPS second 1179105785. *Right:* Data around the SD disk event 41257077/41257078. The event time is defined as microsecond 573 600 of the GPS second 1169159329.

we chose the microsecond 94 000 of GPS second 1179105785 as event time to roughly match the centre of the main long-signal event and keep all activity around it contained within 5 ms. The left panel of Fig. 5.8 shows the distribution of the triggers in this period. For better visibility, the time scale is limited since the triggers outside of this interval do not take part in any SD event and are considered noise here. The classification of the different triggers is explained in the next paragraphs.

The right panel of Fig. 5.8 shows the data for the period of the SD-disk in the same representation. Here we pick the microsecond 573 600 of GPS second 1169159329 as event time. Due to the more “coherent” nature of this event the time scale in the figure is changed to only 300 μs . Comparing the two panels of Fig. 5.8 the difference in the “noise” level of the lightning activity is obvious. We will discuss those differences further in Section 5.2.5.

Merging The basis of the analysis presented here is a successful merging of the trigger information from the T2s with the available data from SD-events. This means merging the trace information from the raw SD-events, as well as the information from the cluster analysis.

Because the pair of trigger time and station Id form a unique identifier for every trigger, matching the triggers from events to ones in the T2 dumps is straightforward. The trigger GPS microsecond can be obtained from the member offset of the `IoSdT2Trigger` class in combination with the T3 time. Those times match exactly the microseconds written to the T2 dumps. The GPS nanosecond used in the SD reconstruction is not used for this because by being more precise it can differ from the raw value, sent as T2 message (c.f. Fig. 5.46). In addition, the GPS nanosecond is only available for returned data such that this method of calculating the T2-time is necessary in any case, so to avoid complications, only the raw microsecond is used for merging.

For both SD-ring and SD-disk events, we use the raw `sd_*.root` files and the “raw” events for merging. This implies that multiple instances of the same trigger can occur in the aggregated data. We associate a trigger with the event that has the correct data: If the GPS block exists, the nanosecond has to match the trigger time to within 3 μs . Otherwise, the trigger is considered as “T3 lost”. This is necessary because in the case of requesting two different traces from a single station within a single T3, the second trace is never returned

because the station software is not processing the second request. In this rare case, there are still valid data for this station in the raw event data. However, these data are an exact copy of the first trace and are therefore ignored.

In the case of requesting a trace in multiple T3s without having the trace data returned to any, the trigger is marked as “T3 lost” in association with the first event that requested it.

Merging the information from the T2 cluster analysis is more straight-forward: We are only interested in a yes/no information and a trigger is either associated with a single event or none.

In summary, every trigger is classified with the following flags:

- is part of a T2 cluster (c.f. Section 6.3.1),
- was requested by at least one T3,
- has trace data.

A trigger that has none of these properties is classified as “noise” or background trigger. In Fig. 5.8, these background triggers are visible mostly “outside” of the main event activity. Due to the dead time arising after many triggers, those triggers still can be relevant (c.f. Section 5.2.2). After checking that the time sequence of T3 requests coincides with the raw event Id with the help of the Ik logs, the request number for the associated event ids is also added as property of the trigger. It will be used to identify issues with sending too many requests within a second.

For those triggers that have a trace, the type of the trace can be further classified, especially keeping the peculiar signature of the “long-signals” in mind. The applied classification is described in the following.

Trace Classification Within this work, the goal of classifying traces is not to develop a perfect algorithm to differentiate lightning noise from long-signals and muons or shower signals. We rather aim at a simple classification that does not need excessive tuning. The reasoning is that to understand potential trigger issues with such events some misclassified traces are of minor importance.

Fig. 5.9 shows the comparison of a long signal with a usual CR signal in the left panel. The distinction between these two, for such large signals, is easy. However, already distinguishing between a low-amplitude long signal, e.g. like the one shown in Fig. 5.7, and the high-frequency noise from lightning strikes is more difficult.

Corner cases, as for example shown in the right panel of Fig. 5.9, are also beyond the scope of this analysis. This station records something that resembles very much a long signal for two of the three PMTs. However, PMT 3 shows clear indications of RF noise, as well as a sudden drop of the signal at the end of the trace. We think that such cases should be addressed in future work analysing more than two events with such stations and with rigorous treatment of corner cases. We show the example traces solely to illustrate the inherent difficulty of designing a good classification algorithm when not even the target categories are entirely known.

In order to get the most important long signals correct, a part of the traces in the centre of the SD-ring event is visually scanned. The resulting classification of traces is presented in Table 5.1. Those flags overwrite any classification done by the algorithm presented here. A trace that is solely classified as long-signal trace by the algorithm is shown differently in the figures.

For the classification, the following classes of traces are used:

- undefined, for stations without unmasked PMTs,




Figure 5.9: *Left:* Comparison of a “long signal” from event 4067441 with a normal CR signal. Taken from Ref. [105]. *Right:* Comparison of the high- and low-gain channels of station 1237 in the event 42852075. This station shows parts of the typical long-signal-like traces for all PMTs. PMT 3 has additional strong lightning-trace-like noise on top of a potential long-signal.

- cosmic-ray, intended for small showers or muons,
- lightning, for high-frequency noise from lightning strikes,
- long-signal, for the starting long-signals, i.e. the rising edge of the long signal,
- long-signal (end), for long signals that have only the falling edge of the signal,
- ambiguous, for cases that are not easily classified as one of the others.

The automated classification of the traces is based on two properties: the number of bins above the threshold of 50 FADCs above the baseline, and the maximum bin of the (magnitude) of the Fourier transform. Both properties are averaged over all valid PMTs³. In case no PMT is valid, the trace is classified as “undefined”. The FFT of the (binned) HG trace returns spectra of 768 bins length. We use the maximum bin as indication of the main frequency component in the signal.

With these two quantities, long signals are signals of stations that either

- have more than 50 bins above the 50 FADC threshold in all valid PMTs,
- or have a maximum bin of the FFT magnitude (averaged over all valid PMTs) of less than 10.

The high-frequency RF noise is also detected using the FFT by defining all stations with an average maximum FFT bin above 100 as “lightning” traces. The frequency content of a typical cosmic-ray trace is in between the one from a long signal and the one of RF noise. Therefore, we define all stations with an average maximum FFT bin below 50 and above 10 as a cosmic-ray trace. The range of maximum bins not covered by these definitions is classified as “ambiguous”, as it might be a combination of other cases such that a clear classification is more complex and beyond the scope of this work.

In Fig. 5.8, those stations classified as “long signal (algorithm)” are stations that were not identified as long-signal stations by visual inspection – not necessarily because it is not a long signal – but only with the algorithm. It is clearly visible that the algorithm is misclassifying some stations far away from the actual ring or disk event. As stated earlier, this is not a problem in the scope of this study.

³For defining a “valid PMT” we use the TubeMask of the IoSdCalib-class.

Table 5.1: List of the traces for the SD-ring events that are visually classified. The trigger time is the time reported in the T2 messages. All triggers are from GPS second 1179105785. An ending long signal refers to a signal that resembles long signals without the rising edge of the trace. This distinction can be important for timing analyses.

Station Id	trigger time / μs	type
1415	94633	long-signal
1288	94634	long-signal
1534	94639	long-signal
1236	94632	long-signal
1714	94628	long-signal
1521	94648	long-signal
1345	94639	long-signal
1404	94648	long-signal
1720	94655	long-signal
1416	94665	long-signal
1237	94215	other
1521	94207	long-signal
691	94220	long-signal
697	94224	long-signal
1006	94222	long-signal
1473	94214	long-signal
1491	94224	long-signal
1459	94213	long-signal
1412	94226	long-signal
1003	94229	long-signal
1460	94231	long-signal
1715	94231	long-signal
1290	94247	long-signal (end)
1492	94248	long-signal (end)

Dead Times Due to the configuration of having two buffers in the PLD for triggering, there can be substantial dead-time in the case of multiple triggers as shown in Section 5.2.2. Because such dead-times could explain “missing” stations in the ring events, we add estimations of the dead-time for each station based on the model presented in Section 5.2.2. Concerning this dead time, we have to emphasize two points. Firstly, it is a strict lower limit, because a potential (random) T1 between the triggers considered here would not appear in the T2 dumps and therefore the dead-time would be underestimated. Secondly, there is a signal dependent dead-time associated with the definition of the triggers as crossing from below to above threshold [109]. This means that if the signal does not fall below the threshold within a single trace, the following signal in the next (virtual) trace is not triggering. It is not possible to reconstruct the dead-time associated with this without the signals themselves that are not recorded due to this very same caveat.

Finding the Centre We frequently display and discuss distributions with respect to their dependence on the distance to the event centre. Here, we describe briefly how this centre is chosen. The emphasis of this work is not on the analysis of the events themselves, therefore

we opt for simpler methods for defining and classifying the event data. We define the location on the ground as *centre* where an assumed timing model would create the earliest signal.

We approximate this point using simply the centre of gravity of stations in the event. Depending on the topology, we modify the set of stations used in this calculation. For the SD-disk, the centre contains triggered (long-signal) stations and therefore the estimation using only the visually scanned (i.e. central) stations with a long signal gives a very good estimate. The quality of this estimate can be assessed from the dispersion in the right panel of Fig. 5.8. Due to the very large extent of the event with 700 stations, the false positives of the algorithmically classified stations would bias the core estimate if they were to be included.

For the SD-ring event, the situation is different because the centre has missing stations. However, the event seems to be contained in the SD array, such that using the bary-centre of all stations that returned trace data gives a reasonably good estimate. This is not entirely clear from the left panel of Fig. 5.8 because of the multiple lightning events surrounding the SD-ring event itself. With the timing analysis presented in the following section, it will become clearer that this choice is sufficient.

5.2.5 Results

In the following, we discuss the results we obtain from the merged available information collected with the methods described in the previous section. We separate different sub-topics into paragraphs that cover the signal timing, read-out ordering, double requests, and finally dead-times, and the discussion of the missing central signals. The signal timing is important to make sure that the timing conditions imposed by the T3 are met by the signals of interest. We investigate the read-out ordering because not all events return data and the time of request can be crucial in determining what data is read-out and what is lost. Keeping our goal at answering whether or not the central stations have a signal or not, we finally check double-requests and dead-times as rare cases that can lead to data-loss and thus are of special interest to us.

Signal Timing Firstly, we discuss the relative timing of the long signals with respect to each other and also the surrounding lightning pulse(s). This can be of relevance because the T3 algorithms assume propagation of the trigger times with speed of light. In Fig. 5.10 we show the trigger times as a function of the distance to the centre of the events. The left panel contains the data for the SD-ring event, while the SD-disk data is shown on the right. As in this case the main focus lies on explaining missing stations in the centre of the event, we fit the times of the firmly identified long-signal stations only, with

$$t_{\text{long}}(r) = ar + t_0 = \frac{r}{v} + t_0 \quad (5.3)$$

where r is the radial distance from the centre and t_0 a arbitrary time offset. We include stations that only have the end of long signals. We shift these trigger times (ad-hoc) by $30 \mu\text{s}$ backwards in time to try to account for the missing start. The fit yields $a = 4.36 \pm 0.81 \mu\text{s}/\text{km}$ and $t_0 = (200 \pm 4) \mu\text{s}$. This corresponds to a speed of propagation of $v \approx 230 \text{ m}/\mu\text{s}$. We want to stress that this speed is a speed of propagation of the *trigger* times and is not a statement regarding the actual signal seen in the traces. It is however interesting that the assumption of having signals propagating with the speed of light is not observed with the trigger times for such signals. Given the observed “speed” of the trigger front the time differences between crowns are expected to be above the allowed limit⁴ for three crowns, and just on the limit for two crowns. As it is within the allowed interval for one crown and – we think – typically

⁴cf. XbAlgo.cc of the CDAS.

almost all neighbouring stations trigger in such events, this is probably not a large problem. However, it is good to keep in mind, that the assumptions made in the T3s might not be suited for all cases of exotic events like the SD-rings.

In addition to this result that the apparent speed of propagation is slower than the speed of light, we use this fit to extrapolate to smaller distances to the centre to have benchmarks when there should have been triggers in case the hole is a central-trigger (i.e. T3) effect.

The long signals of a second smaller event, visible in the top left corner of the left panel of Fig. 5.10, are not fitted and not considered here. We assume that those signals arise from a second (but probably related) event that will be investigated further in Section 5.3. It also seems that the centre of this second event is slightly shifted. However, due to the low statistics and the nature of this study, we do not want to give firm statements on whether the centre is indeed shifted or not.

Interestingly, there is also a third cluster of triggers with a similar separation in time as it is visible in the left panel of Fig. 5.8. However, those triggers do not form a T3 and we therefore do not have any further information on the types of signals causing it. This small cluster of triggers is marked differently than just pure noise, because the alternative cluster finding analysis, presented in Section 6.3, identified those triggers as a significant cluster. In conjunction with the observation of multiple long-signal events in short periods [106] and the multiple-peaked elves [8], it seems possible that such clusters do contain similar signals, possibly at lower amplitude, making it harder to trigger on it with the non-optimised T3.

The right panel of Fig. 5.10 indicates that for the SD-disk the situation is different. Firstly, there is no hole, such that we do not need to extrapolate the timing to smaller distances. Secondly, there is no clear separation of the lightning pulse and the long signals. Thus, we try to guide the eye to the small relative difference of the long signals and the lightning pulse by adding a line that corresponds to a point source emitting a signal moving at speed of light at the relative time offset $t_0 = 48 \mu\text{s}$. With this line it becomes apparent that in contrast to the SD-ring event the long signals appear before the lightning strike. Additionally, we do not observe a large difference in the propagation speed to the speed of light. However, in the cross-over region of long signals and lightning pulses, the trigger times do not follow the same “propagation speed”. This might lead to some loss of stations in similar events at the border of the long-signal region. In the case of this event, this is however not the case and all traces are successfully recorded.

Read out Ordering After making sure that the time distribution of the relevant triggers is not influencing the read out, we discuss now the effect of the read out ordering (in time) of events. In Fig. 5.11 the requested traces are marked in different colours for the different raw events. The sorting of the events in time follows the raw event id that was checked with the I_k logs.

Interestingly, the general approach is to send the T3 found last as first request. This ordering is however influenced by the “merging” of raw T3s (c.f. discussion in Section 6.3.4) such that the ordering is not strictly reversed in time any more. A general tendency to request later events earlier is still visible in the left panel for the SD-ring event.

That the ordering is of importance can be inferred from Fig. 5.12. We show the read-out efficiency, i.e. the ratio of the received traces over the requested traces, as red points (right y-axis) with a decreasing trend with the request number clearly visible. The absolute numbers are displayed with black markers to help put the ratios in perspective. Fig. 5.12 shows that the probability of getting the requested traces drops significantly after the fifth request. For these SD-ring data, it was a lucky coincidence that both parts with long signals are the request number 3 and 4 such that most data were returned. We want to emphasize that the drop in

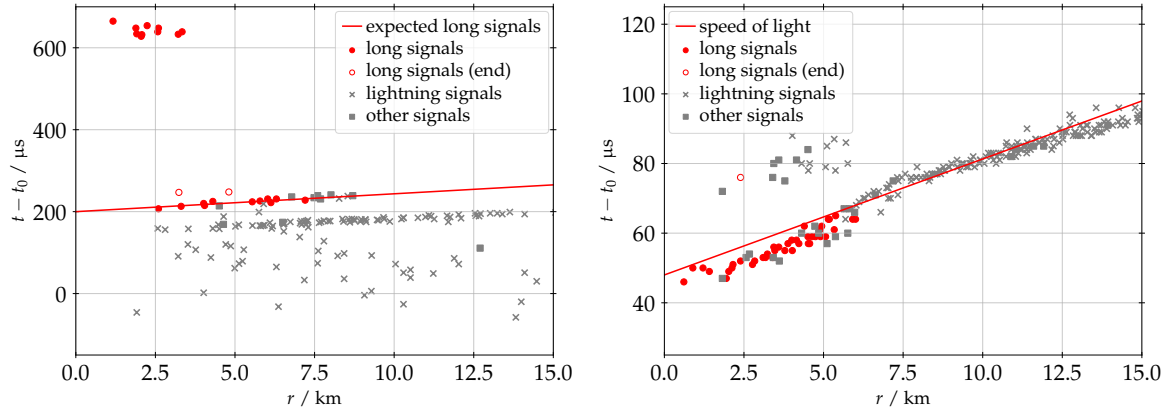


Figure 5.10: Relative timing of the long signals together with a model. *Left:* Triggers with data for the SD-ring event. The linear fit is performed on the long signals only, shifting the ending long signals by $30 \mu\text{s}$. Long signals as only identified by the algorithm are included as “other”. *Right:* Data for the SD-disk event. The model is not a fit, but added manually because fits are difficult due to the overlap region of long signals and lightning signals. The choice is to present a signal propagating with the speed of light that emphasises the difference in propagation of the lightning pulse and long signals. The small amplitudes of all signals in this event can also explain deviations from a simple propagation at the speed of light. Different altitudes of stations can also contribute.

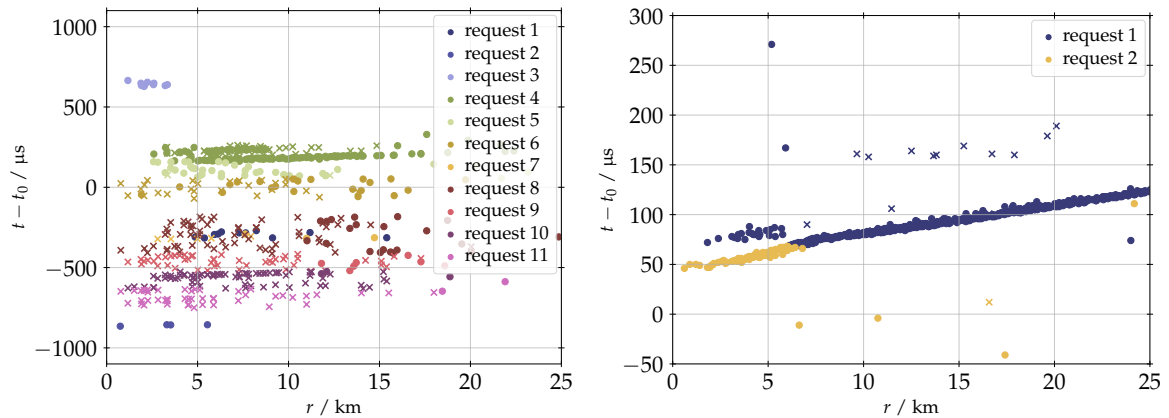


Figure 5.11: Comparison of the read out of the stations for the SD ring and the SD disk. *Left:* SD-ring event(s) of 18 May 2017, showing the association of time and space of triggers requested by events with the time ordered request number of the GPS second. Circles indicate read out traces and crosses stations with errors. *Right:* SD-disk event of 22 Jan 2017. Illustration of the request order for the two raw events containing the long signals.

efficiency for the fourth request (i.e. the big SD ring) is due to a different technicality related to double requests for stations, detailed in the next paragraph.

The conclusion from these observations is that there is a potential to read out more long-signal traces if such events are requested first. This implies of course that these events can be identified on the T2 level. The advantage would be that such a change is not requiring any change of normal data-taking because in normal conditions no change would be applied and even in lightning conditions no change in what is requested (in total) is made. We discuss the implementation of such an algorithm in Section 5.4.

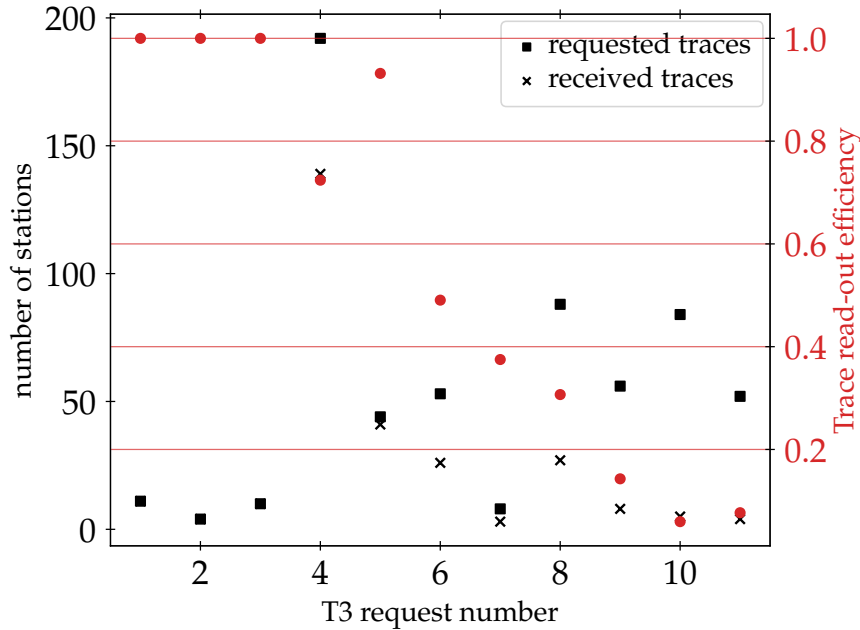


Figure 5.12: Number of requested and received traces per T3 request for the SD-ring event. The efficiency is plotted on the right y-axis. The long-signal traces are recorded as events 3 and 4. The dip at event 4 is caused by requesting two traces in one T3 for multiple stations.

Double Requests To further investigate the read out efficiency presented in Fig. 5.12, especially the drop for the SD-ring event, we analyse what is requested for this particular raw event 42852075. To facilitate this, Fig. 5.13 shows all the requested stations of this event categorised into returned and lost data. We add the time fit of Fig. 5.10 to give a time scale as to where the long signals should (roughly) be. From the right panel of Fig. 5.13, that zooms into a small subset of the stations involved, it becomes clear that it is possible that a single station is requested more than once in a single T3. However, in the current implementation of the station code, these second requests are never answered. It is obvious that due to this a few stations with potential long signals did not return the (potentially) most interesting trace. We want to stress that not only the stations present in the right panel of Fig. 5.13 are affected. We chose this range solely for illustration.

Note however, that in the raw events the stations with lost data do not appear with an error. Rather they do have the data members filled with valid data. That data is however an exact copy of the first trace and therefore should not be discarded.

In summary, this issue limiting the read out performance is significantly affecting the number of long signal stations in at least the event under study. However, it is possible to gain (station) statistics of long signals with fairly small effort by either splitting all doubled stations of a T3 into multiple T3s, or by modifying the station software accordingly.

Dead Times and Missing Centre Lastly, we discuss the effect of station dead times on the events and try to address the question if such dead-times could explain the missing centre. The left panel of Fig. 5.14 shows the (requested) traces of the SD-ring event(s) as function of the squared radial distance in a narrow radial interval around the centre. We use the squared radial distance because in this variable the density of stations is uniform, i.e. the true “missing” stations are easier to recognise. With this argument it becomes apparent that at least 7 stations within 2.2 km ($= 5 \text{ km}^2$) of the centre are active in the second (i.e. they had at least one trigger in the considered period) and have not triggered for the long-signal




Figure 5.13: Display of the read out status of the traces for the raw SD event 42852075. For each station that was requested, also the expected timing of the long-signal event from Fig. 5.10 is shown. On the left, all requested stations are shown. On the right, we show a zoom to the Id range of 1000 to 1020.

part of the event. Using the dead-time model from Section 5.2.2, indicated as grey boxes in the figure, we can conclude that only two of these stations are (for sure) affected by these dead-times. This means that at least 5 stations could potentially have triggered but did not. In parts those stations might be affected by the same kind of dead-time but caused by a (non-lightning) T1. However the rate makes this explanation for all of these stations very unlikely because the number of expected triggers for these 5 stations would be approx. $80 \text{ Hz} \times 100 \mu\text{s} \times 5 \text{ stations} = 0.04$. Thus, we conclude that the buffer read-out dead-time is unlikely as explanation of the “missing” centre.

A different form of “dead-time” arises from the definition of the triggers: to trigger again, the threshold needs to be crossed from below. This means that if the signal does not fall below the threshold for an extended period of time, no further trigger is possible. To see if indications of this situation exists in the SD-ring event, we display the types of the triggers around the long-signal part in the right panel of Fig. 5.14. It becomes evident that the ending parts of the long signals are triggered as threshold triggers, even though the “long-signals” are a prime example of a signal that should satisfy the ToT. This indicates that at this point the ToT(s) were already triggered before and only by a lucky coincidence these traces were triggered as threshold T2s. The second observation for this period is that the MoPS is (sadly) very effective at picking up lightning noise and thus creates a lot of additional T3s leading to loss of data. It should be noted that not all stations send extended trigger flags, so that it is not possible for all stations to tell if a ToT was in fact a new trigger or not. For simplicity, we do not include this information here.

The conclusion on these two types of dead-times for this event is so far inconclusive such that we cannot identify with certainty a trigger reason why the central stations did not trigger.

In summary of all the effects listed here, we have to conclude that while many effects reduce the efficiency of reading out the ring-events themselves, no firm reason for a missing centre could be identified. This leaves us in the unsatisfactory position that, while improvements of how to measure more events are clearly pointed out, no answer to the initial question “is the hole a trigger effect?” can be given. In order to try to add more insights to this, we will go back to the delayed events mentioned earlier and try to infer an answer to this question from there in the next section.

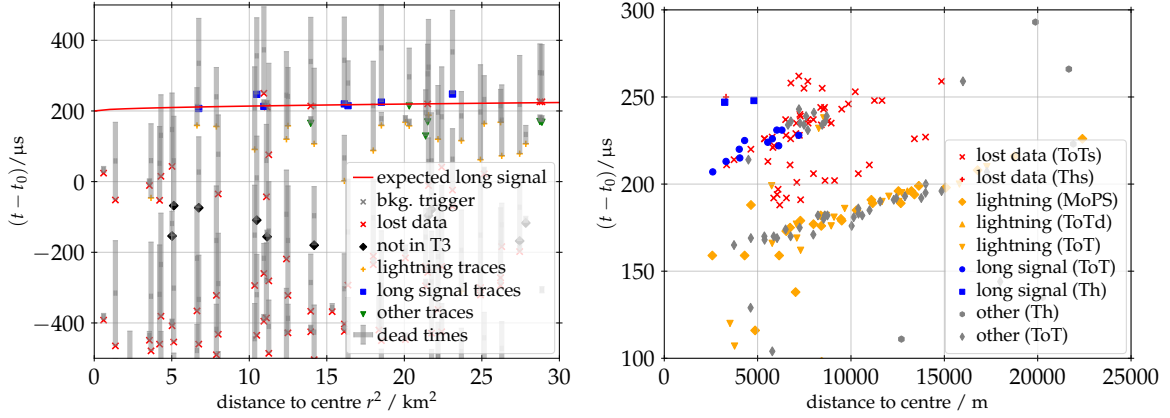


Figure 5.14: Analysis of potential station dead-time effects for the SD ring event. *Left:* Presentation of the triggers around the long-signal part of the event as function of r^2 . Guaranteed dead-times from the limited amounts of buffers (cf. Section 5.2.2) are marked in gray. The expected signal timing is taken from Fig. 5.10. *Right:* Illustration of the different trigger types contributing to the events. Especially interesting is that the ending parts of the long signals are triggered as threshold triggers.

5.3 Information from Delayed Events

Following the general investigations of trigger effects on the data taking of SD-rings of the previous section, we focus now on additional information that can be derived from events that happen after the events of interest. The contents of this section were already published collaboration internally as Ref. [103] and have been used in the collaboration proceeding Ref. [104]. Here, the contents are summarised to come to the conclusion that the most likely explanation of the ring shape in these special events is indeed a trigger effect.

5.3.1 Data Treatment

In contrast to the discussion in the previous section, we now aim at using all available information from the SD data, rather than only the trigger data, that we have for the events that are usually referred to as SD-rings [105, 110]. This enables us to also use the events that happened before the trigger data was stored. In the following, we explain which events are used and how the contained data are merged and cleaned.

We start by selecting all (raw) SD-events within 2.5 ms of a given SD-ring event of interest, similar to studies done previously [111]. The event times are determined manually for the small number of known interesting events from the traces that are clearly identified as *long-signal-like*. Table A.1 has a summary of the times of the events under study. Not all events with these kind of signals are in this table as in this note we are solely interested in finding information about the missing centre rather than the events themselves. The previously found SD-disk events are thus not included.

To highlight that the 2.5 ms time window is large enough to cover all events, the left panel of Fig. 5.15 shows the time distribution of received and lost traces around the event 4067414 recorded in October 2007. In the right panel a different, even larger event, from May 2017 is shown with the full time window. It is clear from the time distribution that there are several events participating and the 2.5 ms windows are wide enough to select all relevant events. For the overall time structure of lightning activity in this region larger time scales might be necessary, as used in [111]. As we will show in the following, all these events have a sub-structure of point-source-like events and thus the 2.5 ms window is more than wide

enough to gather all information for answering the question whether or not the stations in the centre have signals.

The second step is to extract and merge all available data for every station. We omit the event information at this point because the structure of the events – designed for cosmic rays – does not suit the processes for the SD-ring events. Unlike for events past January 2016 when T2Dumps can be used to identify all (T2) triggers in the period of the event, as shown in the previous sections, for the historic data only the information stored in the events is available. However, even if the event data were lost, the information of the trigger associated with the requested data is still available. For every requested T2, the time of the trigger can be reconstructed from the `IoSdT2Trigger` data member. In addition, the information from the answer of stations to the T3 request that did not have a T2 can be used to infer whether or not the station had some trigger in that period.

When using these data to reconstruct which triggers occurred in the lightning periods several effects have to be considered. Firstly, with many stations triggering, it can happen that the same station appears in a single T3 request more than once with different T2 offsets. This leads to lost data because the station (software) never replies to the second request, as it was pointed out previously. It is however still possible to know when that trigger occurred by using the trigger time of the corresponding `IoSdStation` instance.

Secondly, requests for the same trigger from multiple T3s are fairly common in lightning conditions. All available data thus have to be carefully matched and merged to avoid interpreting available data as lost. Lastly, lost data from overwritten buffers in the station happen frequently due to the high number of T3 requests happening in these periods. Therefore, the error flag 2, indicating “T3 not found”, has to be carefully interpreted because in the case of a T1 we might not know whether there was no T1, or if it was simply overwritten and thus “not found” when the request finally arrived at the station.

To merge multiple instances of lost data, we arbitrarily define lower error codes as of higher priority. The reasoning is that the smallest possible error code “0” indicates available data. In this way we can determine an error code for every T2 we find in the data.

A special case is the aforementioned double requests of a station. In case two T2s of the same station are requested in a single T3 and data is returned, the data of the second requested trace is an identical copy of the first one, a mistake most probably done in the CDAS-DAQ event builder. Thus, we identify these cases only when data was returned in which case we set the error code of the second trace to “10” because the trace in fact was lost due to an interface bug in the DAQ⁵.

After the merging of the T2 data, the T1s are added. We make sure, that also data from T1s is merged if appropriate with the corresponding T2s from other T3s.

To illustrate the scale of the events and the data available, we show the position of the stations participating in the events around the aforementioned event in the left panel of Fig. 5.16. “Black” stations are not drawn, neither are stations that are not deployed yet – the event 4067414 was recorded before the completion of the SD array. Stations that have no T2 requested are drawn in black, while the ones with at least one T2 are shown with blue markers.

Due to the position of the event at the border of the array and the presence of random-coincidence muon triggers, the estimation of the true event centre requires a careful selection of the correct stations. Because the goal of this analysis is to find hints of lost data rather than a description of the events themselves, we select a station in the centre manually as proxy for this centre. For the event shown in Fig. 5.16 we take station 1346 as highlighted with a red

⁵Most likely the best solution is to change the local-station code to reply correctly to such T3 requests. But it is also possible to require that a valid T3 contains a station only once.




Figure 5.15: Visualisation of the events and triggers around the SD events 4067414 on 14 October 2007 and 42852075 on 18 May 2017. *Left:* Distribution of the trigger times of all stations in all events within 2.5 ms of the long-signal event. We define the event time t_0 as GPS time 876427835.195820 s. A zoom into the range $[-500, 700] \mu\text{s}$ is shown. *Right:* Time distribution of the triggers around the time 1179105785.94200 and the May 2017 SD-ring event. The full ± 2.5 ms are shown. It is clearly visible that the event is shorter than this search window.




Figure 5.16: Visualisation of the events and triggers around the SD event 4067414 on 14 October 2007. *Left:* Map of the stations involved in the event. We define station 1346 as centre of the event. The red circle indicates a 7.5 km radius around this chosen centre that will be used in the visualisations of the recorded traces. We highlight the stations with a long-signal with red-circles. Stations having a linear signal (overshoot) after the main event are marked with green circles. In case a trigger with lost data is reported at the time of the other overshoot we mark those stations with grey diamonds. *Right:* Zoom into the region with the long signal stations.

cross. For the other events the choices of centres are given in Table A.1. When the barycentre of all T2 stations works well enough, no such choice has to be made.

The red circle around the chosen centre in Fig. 5.16 indicates a 7.5 km radius around the centre that contains the relevant stations. In further figures only the stations within this area are shown to increase the visibility of the features.

5.3.2 Observations

To facilitate a better understanding of the relative timing of different events that happen around the SD-rings we try to visualise all available data in a single figure. Fig. 5.17 is using the events around SD event 4067414, the same event we used as an example also in the previous subsection. As mentioned in Section 5.3.1, we solely use the stations within the red

circle of Fig. 5.16 indicating a 7.5 km radius. In Fig. 5.17, the stations are sorted with respect to their radial distance to the centre from bottom to top of the figure, i.e. the first station is station 1346 by definition. For each of these stations, we add the trace data from all events of the period. For the relative timing a time offset of $t_0 = 195820 \mu\text{s}$ is subtracted, and the propagation time $t_{\text{prop}} = r/c$ of the signal from an assumed point source at the ground is taken out. The point source assumption at the ground is a good (enough) approximation for the purpose of showing which traces belong to a single luminal event, but is not meant as statement that there is a source at the ground. Such studies have to be carried out more carefully including a precise estimate of the signal time. Nonetheless, it is visible from the good (vertical) alignment of the long-signal traces that the estimated centre is not far off.

To help recognising the different types of signals, we use the online-calibration parameters to have a roughly calibrated trace that is scaled such that the maximal value is always at the same amplitude to increase visibility. To keep the information of the signal size, we add a colour code that is based on the logarithm of the integrated signal. Should the online estimate of the baseline be overestimating the real baseline such that the integrated signal appears negative, we shift the trace until we obtain a positive signal.

Lost data is indicated with red brackets where the error code is given as number next to the “trace”. The length of the brackets is the same as the one of traces and the timing corresponds to the microsecond of the T2 message sent to CDAS. Thus a jitter of some $3 \mu\text{s}$ is possible.

T1s are left out in Fig. 5.17 to increase visibility of the main features of the events. In Fig. 5.18, T1s are added to complete the picture with all information available. For stations returning “T3 not found”, meaning that no trace is available in the requested time window, we show the requested window as grey exclusion region. If the station had a trace but no T2, we draw a red region corresponding to the time window requested by the T3.

There are several important observations to be made from Fig. 5.17 and Fig. 5.18. Firstly, stations that do not have a long signal in the centre are alive and have data in previous events, as was pointed out before e.g. in Ref. [111]. However, these signals do not line up with the same (point) source, even though it is likely that the same general physical phenomenon leads to both the early lightning traces and the later long signals. Although, in general also the previous events seem to originate from a point source in a similar position, just at a different time, as indicated by their almost perfect alignment with respect to each other.

Secondly, apart from few exceptions, the central stations did not have a T2 at the time of the long-signal event. One of these exceptions is e.g. station 1347 that lost a trace at the correct time due to requesting two traces in one T3 as indicated by error code 10. Thus, we know that the loss of signals in the centre is not a CDAS/bandwidth issue: if no trace was lost, there simply was no trigger (trace) to readout in the first place. This fact is highlighted even more in Fig. 5.18 with the stations 1118 and 1369 that responded with “T3 not found” to the T3 request because there was no trace available.

Lastly, and most importantly for this work, there are delayed traces at about $500 \mu\text{s}$ in relative time that are neither lightning traces, nor cosmic rays, nor long signals. We assume these signals are the result of the overshoot of the baseline of the PMTs after very large signals. We will discuss in Section 6.2.5 how this effect of overshoots known from laboratory measurements [112, 113] is visible in the trigger data. From the electronics we expect that the undershoot after a large signal recovers and turns into an overshoot after about $400\text{--}500 \mu\text{s}$. This time scale is exactly matched in the event here. In Fig. 5.20 we marked this time window on the recorded signal in an SD test-tank after a very large pulse. Usually the triggers associated with these overshoots do not lead to T3s because in cosmic-ray events very few stations, usually just one, have a signal large enough to cause delayed triggers and for a T3 at least three stations are required.

To further highlight the fact that these traces are indeed overshoots and not long signals with small amplitude, we show an example of these two classes in Fig. 5.19. The left panel shows an example of an overshoot with a very linear rise of the ADC values with little noise. On the other hand, the right panel shows a long signal of small amplitude taken from a different event in January 2017 that is very clearly different. The overshoot recovery is not a linear process, as visible in the overall signal in Fig. 5.20. However, the traces are so short in comparison to the time scales of the exponential that it will appear linear. This feature is highlighted with the tiny red box in Fig. 5.20 that is drawn to scale with respect to the trace length.

The key point of observing overshoot signals is that they require a large signal in the station, independently of whether it triggered or was readout or not. Therefore, we can infer from the observation of an overshoot signal that a station had a large signal about $450 \mu\text{s}$ prior. For example, in Fig. 5.18 we can see that station 1369 has an overshoot signal even though no trigger was reported during the time of the long-signal event. Another example, where the long signal was recorded as well, is station 1110 with a strongly saturated long signal and an overshoot afterwards. Following this chain of thoughts, we can see that the central hole of the event is most likely an effect of the local station trigger rather than physics. We highlight the presence of the overshoots in the central parts of the event in the right panel of Fig. 5.16 by highlighting the stations that either have an overshoot signal, or lost a trace in the time window of that T3: almost the entire central part has overshoots while none of them had a long-signal trace.

To highlight this point even further, Fig. 5.21 shows the data used already in the previous section again in a radial projection. The absence of triggers in the central part at the time of the event of interest – indicated by the red line – is clearly visible. However, a re-evaluation of the traces of the late triggers revealed that they are overshoot signals rather than long signals with small amplitude. In the appendix, in Fig. A.1, the traces are visible in the same way as in Fig. 5.17. From Fig. 5.21 it is evident that the central stations all have delayed triggers at the time of the overshoot. Some of these triggers were requested with a T3 and can be looked at to confirm that it indeed is the overshoot. We mark these station in Fig. A.3 to highlight that these stations are indeed in the very centre of the event.

The reason that not all of these overshoot traces that lead to a T2 are requested is the timing and scatter of the overshoots. The T3 algorithm allows up to $8 \mu\text{s}$ time difference between neighbouring stations to be a valid first crown coincidence. Without a single first crown pair, no T3 can be formed. The probability of some stations not having the T2 from the overshoot in the correct time window for the T3 is fairly high given that the time scale of the overshoot is much larger than the tolerance of the T3 algorithm.

This last point is also the reason that in many SD-ring event we do not observe the overshoot traces. If the events are not large enough and the traces barely saturate the stations, the chance of having three stations – as required for the T3 – with a T2 from the overshoot is very small. An example of an SD-ring event that has only a single saturated station is shown in Fig. A.2. We can assume that also here the central stations have a signal. Because of the smaller amplitudes overall, it just does not lead to overshoots large enough to trigger (enough) stations. Thus, the information about the central signal is lost completely here, but it is compatible with the same explanation of a local station issue.

The time scatter also explains why some stations, like for example station 1108 in Fig. 5.17, have no overshoot signal even though the long-signal trace was clearly saturated. If the trigger is just outside the window it will not appear in the data. For the majority of the stations in events like 4067414 the absence of an overshoot signal is caused by the small amplitude of the signals. The decrease of the amplitude of the long signals can be seen in




Figure 5.17: Data from T2s in event requests from events around the SD event 4067414. Traces from all stations that were requested are presented. The signals are calibrated using the online calibration values and scaled to the same maximal value. The colour code indicates the integrated signal size without taking saturation into account. Red brackets are drawn for data that is lost at the time of the trigger with the corresponding error code shown. Multiple separated luminal (compatible with the speed of light) pulses are visible. The stations are sorted with respect to their distance to the centre of the event, here chosen as station 1346, ascending from the bottom. The distances are visualised on the left with the grid structure clearly visible, because a station was used as estimate of the centre.

the colour code in Fig. 5.17 even though the smallest values (and thus lightest colours) are present in the lightning traces before the actual long signals.

5.3.3 Recovering the Signals

Inspired by the way the overshoot depends on the signal size [112, 113], we can try to estimate the lost signal size with the help of the overshoot signal. In Fig. 5.22, we show the correlation of the integrated signal with the difference of the baseline for the few stations that have both an overshoot trace and the long signal. We use the integrated ADC values of the low-gain trace as signal estimator. The baseline difference is calculated as the difference of the average values of the first and last five bins of the trace.

Because the recovery is a slow process compared to the trace length, the relative timing of the recorded part of the overshoot might be important. In Fig. 5.22 we add this information with a colour code. It is clear from the lack of statistics and the large spread that a precise recovery of the signal is very difficult.

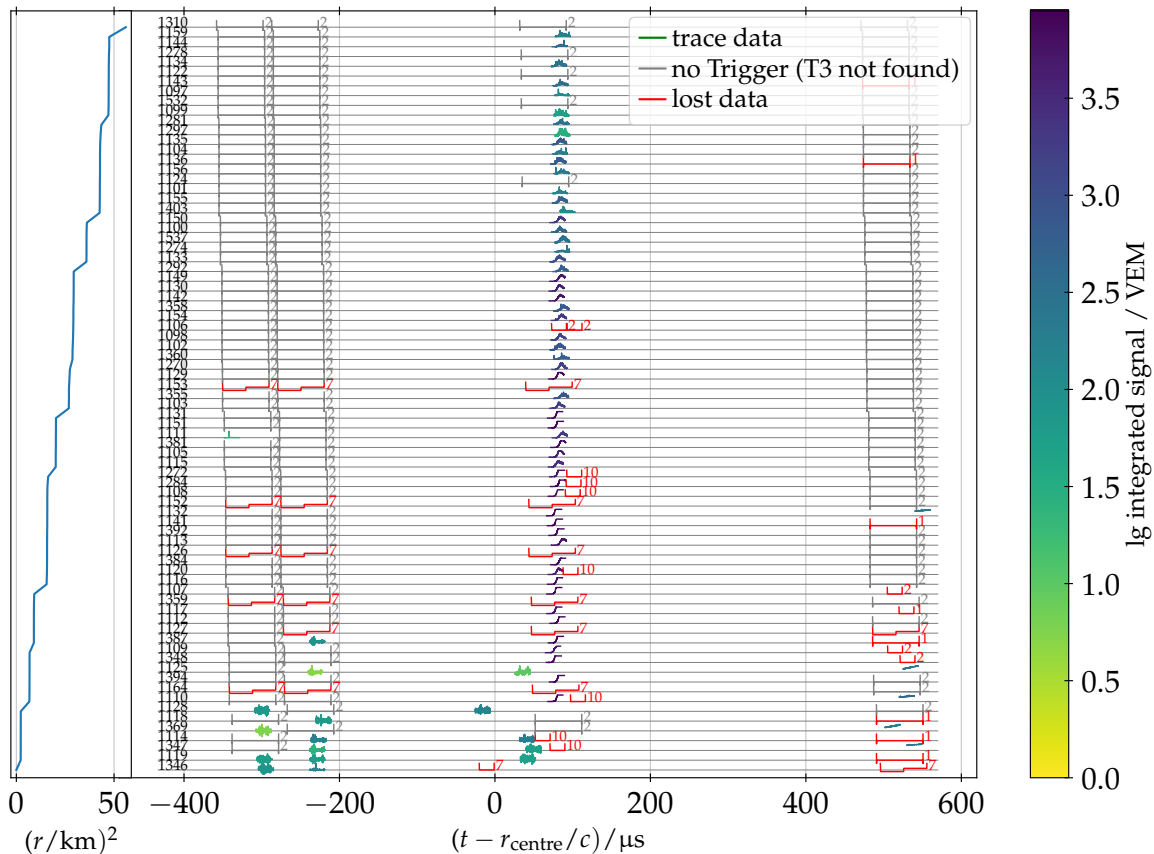


Figure 5.18: Data from T2s and T1s in event requests from events around the SD event 4067414, where the GPS time $t_0 = 876427835.195820$ s is chosen to approximately centre the event at 0. The presentation of the traces is the same as in Fig. 5.17 with the stations that sent data, or lost it, from T1s included. Gray brackets indicate the time windows of T3 request that did not have data in the station, i.e. an exclusion window for that station.

To highlight why some of the scatter originates from the incomplete recording of the long signal, we pick the two data points at about 5.7 ADC difference. In Fig. 5.23 the corresponding long signals are shown with both low- and high-gain signals. The most important difference is that one signal contains the peak of the long signal, while the other one does not, making it clear why the signal amplitude is seemingly so different.

Also the level of lightning noise at the start of the traces is different as is visible from the zoom in Fig. 5.24. The latch bin of 250 (corresponding to $t = 6.14$ μ s) can give a hint to why one trace is seemingly starting later. In the left panel of Fig. 5.24 we have a pulse of lightning-like noise at the latch bin. For the other signal, in the right panel, it is the long signal itself that starts to rise significantly there. This shift of about 1 μ s paired with differences in the signal width and rise time lead to the significant difference in the integrated low-gain signal.

Further studies might improve the precision of these estimates by using fits of the long signals that make use of the similar signal shape. In light of missing peaks of the signal, e.g. in the left panel of Fig. 5.23, this approach is limited too. With more signals recorded and by potentially using models of the electronics, further improvements might be possible in the future.




Figure 5.19: Comparison of long-signal traces with those of overshoots. The rise in amplitude of the long signal and the linear rise of the baseline are different, even if in some displays they may seem similar due to their slow component of the signal. *Left:* Overshoot trace for station 1125 in event 4067413. All PMTs show an linear increase of the baseline with very little noise. *Right:* Example of an low-amplitude long signal from station 1438 of event 41257078.




Figure 5.20: Measurement of the response of the SD-electronics in a test tank from Ref [112]. The measurement is inverted with respect to the signal, i.e. negative values are positive signals in the stations. After a very large pulse (saturated) there is an undershoot which saturates the measurement device. As indicated with the first red line, the signal recovers above the true baseline after about $300 \mu\text{s}$. The second red line is estimating the time when triggers appear in data after low-gain saturated signals (c.f. Section 6.2.5). To highlight which part of this recovery process the overshoot traces are capturing, we marked it with a red box. A small box within this marked time window is to scale of the trace length showing why the overshoot will look linear in the trace.




Figure 5.21: Radial projection of the data of the events around SD event 42852077 from 18 May 2017. The triggers are taken from the T2Dumps and matched to the data in the SD data. All the traces are classified as in Section 5.2 with the exception that the late traces are re-classified due to a visual re-analysis of the data.




Figure 5.22: Correlation of the slope of the overshoot traces with the size of the preceding long-signal. The baseline drift is calculated from the difference of the first and last five bins of the trace. As estimate of the signal size we use the average over working PMTs of the low-gain signals in ADCs. The time difference between the signal trace and the overshoot is displayed with the colour-code.




Figure 5.23: Two examples of stations that have a long signal and a following overshoot with a drift of about 5.7 ADC in the overshoot signal. The difference, that in one case the peak of the long signal is contained while it is not in the other one, is apparent. The online estimate of the baselines are subtracted in this visualisation. *Left:* The trace of station 1521 of event 42852075 from May 2017. Δ Baseline = 5.6 ADC. *Right:* Station 1110 in event 4067414 from October 2007. Δ Baseline = 5.8 ADC.




Figure 5.24: Zoom into the start of the two signals from Fig. 5.23. Differences in the rise time and the level of lightning noise are clearly visible. *Left:* Trace from station 1521 of event 42852075, Δ Baseline = 5.6 ADC. *Right:* Trace from station 1110 of event 4067414, Δ Baseline = 5.8 ADC.

5.3.4 Concluding Remarks

In this section, we show that we have information about the signals in the centre of the SD-rings that tell us whether or not there was a signal. By observing an overshoot signal from stations that before were just known to be not-triggered, we can infer that these stations indeed have a large signal at the right time. We thus conclude that it is very likely the rings are in fact disks.

For the cases that do not have stations with signals in the centre and no overshoot signal, we have to resort to reasoning and the principles of similarity of these events. When the signals are not very large for at least three stations, the overshoot signals will never be recorded due to the requirements of the T3 algorithm. This constraint paired with the expected scatter of timing in the overshoot signals make it very likely that for smaller events no overshoots are recorded even if they are in fact there. From the study of the last known SD-ring from May 2017 with the T2Dumps, described in Section 5.2, we can see that this effect also occurs with many stations. For that particular event, we have overshoot signals

from a number of stations, but also a fairly large number of T2s in the time interval of the overshoots that are not requested by any T3 (c.f. “not T3”-triggers in Fig. 5.21), highlighting that the T3 is inefficient at picking up these (unwanted) signals.

Combining the likelihood that small events will not have enough overshoot triggers to cause a T3 with the assumption that the SD-rings are all similar, we conclude that it is very likely that all SD-rings are filled and should thus be referred to as “SD-disks”. In future studies the exact problem with the station triggers in the centre should be identified. Such knowledge can help support this conclusion further and help develop adequate changes that in the future these events are recorded entirely.

5.4 Modification to CDAS

In the previous sections, we discussed why the centre of the SD-ring events is not present in the data and how the ordering of the T3-requests governs which traces are successfully readout and which ones are not. Here, we will outline why this very same effect leads to the reduction of the observed rate of SD-disk events after the introduction of the “new triggers” ToTd and MoPS [68, 114, 115] and how we can improve the read-out efficiency without changing the trigger conditions themselves.

The new triggers, aimed at measuring the electromagnetic component, are by construction more sensitive to smaller signals in the traces. While in normal data-taking conditions this leads to an increase of the trigger efficiency without negative effects, during thunderstorms it creates many more T3s leading to an overload of the communication system and thus creating more lost events. For the cosmic-ray analyses, this effect is minimal because these periods are removed from the analysis by identification as “bad period”. Thus, no event influenced by this is used in the analyses and the exposure is accounted for correctly.

For recording SD-disks, these lost events are creating the major loss in efficiency observable in the rate: Taking the subset of events used in the previous analyses (c.f. Table A.1), we observe 15 events before 2014 and only two afterwards. However, at the same time, the new triggers add long signals of smaller amplitude in the events if they are readout, with the SD-disk event from Jan 2017 used before in this work being a prime example.

Thus, the question to be answered in this section is whether it is possible to keep all advantages of the “new triggers” but also recover the SD-disks events. We will start discussing the concept and the necessary ingredients in Section 5.4.1, followed by the descriptions of the details of the implementation and testing.

The contents of this section have in parts been previously published collaboration internally as Ref. [102] and contributed to Ref. [104].

5.4.1 A Concept for Prioritized Read-Out

In order to understand how the efficiency of recording SD-disk events can be increased without changing the normal data taking, we first need to understand the information flow in the trigger decision and read-out. Fig. 5.25 illustrates this system: While the station has all the information about the traces, it can only send 4 bits with the trigger message that is used in the formation of the central triggers in the “Xb”-process. Currently, only the information on the trigger type – the black label in Fig. 5.25 – are sent. After forming the T3s, the T3-requests are sent to the stations and the data, if still available, is returned.

Because we do not want to change the data taking in normal conditions, a change of the T3 formation is out of question. However, the ordering of the T3-requests that are sent to the station(s) can be changed without altering the normal data taking. In simple terms, the rate of T3s in normal conditions is low enough (about 0.1 Hz) that almost no GPS second contains

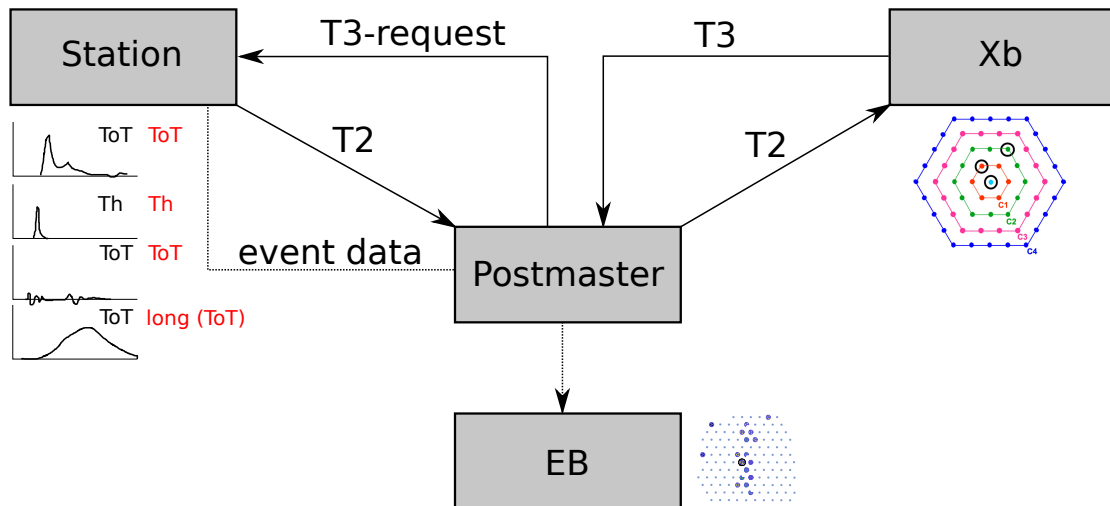


Figure 5.25: Illustration of the information flow from the station to the event level. The station sends the T2-triggers to the central trigger process “Xb” via the Postmaster that manages the network traffic. For these triggers only the type (black labels) is sent. After forming the coincidences, the T3s are processed into T3-requests to the stations that then (ideally) return the data that end up in the event structure via the Postmaster and the Eventbuilder “Eb”. If new information in the T3-request or formation is required, e.g. the “long” flag, it has to be created at the station and sent with the T2s.

two or more T3s. Since a list of length one has no ordering, any ordering condition imposed has no effect. Even in the case of multiple T3s in a second, a reordering of the requests does not change anything fundamentally: If all event data is returned, nothing changes. If events have lost data, the event affected might change but for exposure estimation the same conditions can be applied.

Thus, we aim at changing this ordering to prioritise interesting events in conditions with multiple events. However, to have the necessary information to facilitate such a reordering, the information from the station about which trace is “interesting” has to be transferred to the relevant processes. This means that within the 4 bits available for the trigger type the information *IsLong* has to be encoded too. In the sketch of Fig. 5.25 this is done with the tags written in red. To keep the information necessary for the T3, i.e. ToT or not, two separate flags are necessary. After considerations of the available bits that allow 16 combinations, while only 3 – ToT, not-ToT, and Scaler – are strictly needed for normal data taking, we pursue this solution.

However, the main requirement for this solution is an algorithm that can efficiently identify long signals with as little computation power as possible. In the following sections, we will describe how we find such an algorithm and finally present the implementation in Section 5.4.5.

5.4.2 Events to Detect

With our goal of identifying and tagging long signals with the station software in mind, we need to define more closely what “long signal” includes. Besides, to evaluate the performance of different algorithms a test data set is necessary. Due to the lack of knowledge about the physical origin of the long signals, no simulations are available. Therefore, we use recorded SD data which will contain both signals and various backgrounds. The challenge when using recorded data for both signal and background is to correctly identify the interesting traces that are considered “signal”.

When referring to “long signals” usually examples like those in the left panel of Fig. 5.26 are shown. To highlight that these are not the only signals seen in lightning periods that are unusual, we picked another example in the right panel. Also for this example event, as well as for those shown in Fig. 5.27, the condition that all PMTs and high- and low-gain channels see the same signal is fulfilled, thus excluding RF-induced signals. We aim at correctly identifying all these different signals because the proposed algorithm is part of the DAQ system and thus should be as inclusive as possible.

As a starting point we define the data set by using previous high-level analyses of SD-rings and other exotic event structures [105, 106]. However, these analyses were performed on merged SD-data. But because the proposed algorithms are to be used in the DAQ system before the merging, the first step in obtaining a list of signal events is to check which raw SD-events contain the interesting long signals within the merged events. We perform a visual scan of the raw events with the event display ED to convert the merged event IDs to raw events.

In addition, inspired by the recent analysis of stations appearing with multiple (different) signals within a merged SD-ring event [107], we inspected also events around the original selected ones. In this process we added smaller events that show an interesting time structure of repeated long signals on time scales of several 100 μs and were not part of the original analyses. Fig. 5.27 shows such an example of four raw events with four or five pulses within 400 μs . Thus for our analysis, this example contains four signal (raw-)events due to the number of raw SD events that have at least one of the interesting signals. In total, we obtain a list of 90 signal events⁶ in 38 different GPS seconds. As background we use all the other events in the raw `sd_*.root` files containing the selected signal events.

The data set requires some preprocessing to mimic the situation in the DAQ system. We apply a loose input selection on the data and remove stations with errors and all non-T3 stations⁷ as those would not be available to send a tag indicating long signals. Furthermore, we use only the raw (high-gain) FADC-traces of the SD data and, if necessary, the online estimates of the baseline and the VEM as well as the PMT-mask. Other information is not available on the station that has to run this algorithm and thus is not used in this work.

5.4.3 Evaluation of Algorithms

Before we describe the proposed algorithms and their performance in detail, we have to be clear about the “loss-function”, i.e. what has to be optimised by any algorithm. In general, the goal for any trigger algorithm is a high efficiency and a low false-positive rate. However, in our case the situation is more complex. First of all, the algorithms here are not trigger algorithms but “tagger” because they use the data recorded with the already existing triggers and they will just add a new flag to the existing T2s. On top of that, the signals to be detected are defined on (raw-)event level only and not on a station level. Lastly, in normal data taking, as well as with pure lightning noise, a false positive result of the tag has no detrimental effect in contrast to the usual situation of triggering.

Keeping these basic ideas in mind, we still base our evaluation of the algorithms on detection efficiency and false positive rate. However, the decision on right or wrong is moved to the event level, by defining a raw event as tagged if at least one station returned a long-signal flag, defined by the respective algorithm. The efficiency implied by this definition includes cases that the algorithm actually flagged a lightning station as signal and signal

⁶This event list is available in the git-repository in the cosmo-geo group at <https://gitlab.iap.kit.edu/auger/cosmogeosd-ring-trigger>.

⁷This is mainly T1 stations, but it can also affect stations with a T2 which were not part of the T3.

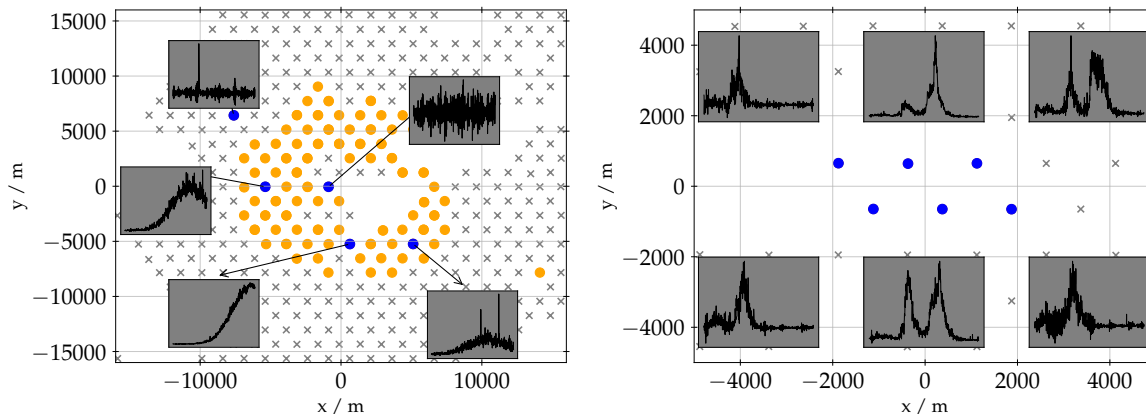


Figure 5.26: Example events selected as “true” long-signal events. *Left:* (Raw-)Event with SD Id 4364128. The averaged high-gain signal of the stations 836, 1012, 704, 1018, and 1441 is shown in the insets (clockwise, starting from the top left). Shown are signals from random muons, RF-noise, and long-signals of various sizes. *Right:* (Raw-)Event with SD Id 4364157. The averaged high-gain traces of all stations are shown next to the station in the footprint. The traces show a double-peak structure which is not in line with the “typical” long signals seen in SD-rings.




Figure 5.27: Activity across multiple raw events. Data from the SD events 4364783, 4364784, 4364785, and 4364786 are shown. *Left:* The trace data of all stations having at least one trace in the four events with the station Ids given on the left. The data is averaged over active PMTs and the online estimates of baseline and VEM are used to get an “online” calibrated trace. For better visibility the traces are scaled with a factor $\log(S)/S$ with the integrated signal S where the colour represents the integral S on a logarithmic scale. The delay from the propagation of the signals, starting at an assumed source in the barycentre of all stations, is accounted for. The stations are sorted by decreasing distance to the estimated centre from top to bottom. *Right:* Position of the stations having at least one trace on the ground in relative positions. The used barycentre is marked with a red cross.

stations were not flagged. In this sense, our definition of efficiency and false positives is station-agnostic.

Furthermore, the total (event) false-positive rate includes many cases that have no influence on the data taking. As outlined in the introduction, a change of the ordering of events in a list of length one does not change anything. Since significant parts of the input data set have no lightning activity but rather normal cosmic-ray events only, these false positives fall into this category. Therefore, we define the false-positive rate of signal GPS-seconds as the amount of GPS seconds that have both (at least) one signal event and (at least) one

false-positive event. With this redefinition we include only the periods when false-positive tags of an algorithm actually are detrimental to the goal. The use of GPS seconds as basis is due to the set-up of the CDAS that processes the triggers in blocks of GPS seconds.

For our goal of recording more SD-disks, or similar events, the most important quantity is the efficiency. Since there always is a – to some extent – arbitrary choice in the trade-off of false positives and efficiency, we define a common working point for the comparison. We define the working point of the algorithms, such that the parameter settings are chosen to have the highest (event-)efficiency with at most one signal GPS second having a false-positive event. We use the absolute number of one GPS second rather than a relative measure because of the very limited statistics available. Future studies, especially if dedicated to new possible triggers in the UUB, should revisit these choices.

5.4.4 Investigated Algorithms

In this section, we explain the details of the algorithms we propose and test them on our test data sample from the previous section. In total, three algorithms are tested that are motivated by their simplicity and the likelihood of being fast enough to work on the station. With larger data sample and more processing power in the upgraded electronics, this list can be extended in the future to more and more complex algorithms. However, this is beyond the scope of this work while first ideas for the UUB are discussed in Section 5.5.

N_{bins} over threshold A simple approach, which was already used to identify SD-disk events [105] in offline analyses, is to count the number of bins in the trace above a sufficiently-high threshold. The motivation behind this algorithm is the simple observation that the long-signals are long and large. Not many other traces have similar amounts of bins above a large threshold. However, we should keep in mind that signals observed in some of the more unusual events like Fig. 5.27 or those in SD-disks are generally of much smaller amplitude, making this algorithm possibly not suited for such signals.

Formally, we define the algorithm as

$$|\{t(i) > x_{\text{thr}}; i \in I \subseteq [0, 768)\}| > n_{\text{thr}} \quad (5.4)$$

on a subset of intervals (or the complete trace) I of the trace bins $t(i)$ where $|\cdot|$ denotes the cardinality of the set. We chose to use an FADC threshold x_{thr} rather than an (online) VEM one to simplify calculation. In addition, the coincidence conditions have to be defined. There are two main possibilities to impose coincidence: The first one is applying the algorithm on each PMT separately and then requiring the number of PMTs passing the condition Eq. (5.4) $n_{\text{PMT}}(n_{\text{mask}})$ to be above a required number depending on the number of working PMTs n_{mask} . The second option is to apply the coincidence directly on a bin level, i.e. for bin i

$$n(i) = |\{t_j(i) > x_{\text{thr}}; \text{for } j \in \text{PMTs}\}|, \quad (5.5)$$

and then

$$|\{n(i) > n_{\text{PMT}}(n_{\text{mask}}); i \in I \subseteq [0, 768)\}| > n_{\text{thr}}. \quad (5.6)$$

Such a bin-wise approach, although more complicated, is potentially more sensitive because one of the main differences between long signals and lightning noise is that the long signals are coincident in all PMTs.

Following first tests, we decide to request full coincidence of all working PMTs and do not use stations with only one working PMT. This choice should be revisited, especially should the amount of stations with only one working PMT increase in the future. For a rigorous test




Figure 5.28: Efficiency of flagging signal events correctly in the test data set for the N_{bins} over threshold algorithm. *Left:* The algorithm is applied on each PMT separately and a coincidence of passing PMTs is required. *Right:* The algorithm is applied on a “coincident” trace, i.e. the condition on n_{PMTs} over threshold is applied per bin.

of these conditions together with the threshold the statistics of the available test data is not large enough.

For the remaining two parameters x_{thr} and n_{thr} we determine the optimal settings by scanning the parameters. In all cases we use the full trace as interval I . Also for this possibility of meta-parameter optimisation, the statistics of the data set is not sufficient to avoid overfitting.

In Fig. 5.28, we show the obtained efficiencies of tagging events correctly for different thresholds and both coincidence conditions. The more restrictive bin-wise coincidence condition leads to lower efficiencies for the same parameter settings. However, also the false-positive rate is lower, as indicated by Fig. 5.29. Using the previously described choice of the working point we obtain the parameters $x_{\text{thr}} = 25$ FADC with $n_{\text{thr}} = 45$ for applying the coincidence at the end and $x_{\text{thr}} = 20$ FADC with $n_{\text{thr}} = 45$ for applying the coincidence bin-wise. The achieved efficiencies with $\varepsilon_{\text{PMT}} = 0.61$ and $\varepsilon_{\text{bin}} = 0.57$ are not significantly different.

Consecutive bins above threshold A straight-forward modification of the first algorithm is to use the distinguishing fact of long signals over RF-noise that long signals are above threshold mostly in consecutive bins. Thus, we define the second algorithm by comparing the largest number of consecutive bins n_{above} above a (FADC-)threshold x_{thr} per trace with a threshold number n_{thr} . Also here the coincidence of the PMTs can be treated in the same two ways as described in the previous section. For consistency, the coincidence conditions are kept identical.

Similarly, we scan over the two remaining parameters to determine the best working-point. Fig. 5.30 shows the obtained efficiencies and Fig. 5.31 the corresponding number of wrongly tagged signal GPS-seconds.

The resulting efficiency at the working point is then $\varepsilon_{\text{PMT}} = 0.61$ and $\varepsilon_{\text{bin}} = 0.49$ for the two possibilities of forming the coincidence.




Figure 5.29: Number of signal GPS seconds that contain in addition a false positive tag. Different settings of the N_{bins} over threshold algorithm are used. *Left:* PMT wise application of the algorithm. *Right:* The algorithm is applied to a coincident trace.




Figure 5.30: Efficiency of flagging signal events correctly in the test data set for the consecutive bins above threshold algorithm. *Left:* The algorithm is applied on each PMT separately and a coincidence of passing PMTs is required. *Right:* The condition on n_{PMTs} over threshold is applied per bin.

Integral differences Many of the SD-disk signals are fairly low in amplitude and only picked up due to the new triggers. Because the concept of using an FADC threshold is not very efficient for very small signals, we try also an algorithm that is considering more than one bin at a time. The algorithm is based on two main ideas: The first one is that the oscillating structures of RF-noise present in lightning events will average out if not the peak but the integral is used. The second one is that a difference of parts of the trace is independent of the baseline. This makes any such algorithm robust against undershoot or overshoot after large signals. In the case of long signals, this is important when there are repeated bursts since a




Figure 5.31: Number of signal GPS-seconds that contain in addition a false-positive tag for the consecutive bins above threshold algorithm. The two panels show the different ways of forming a PMT coincidence. *Left:* PMT-wise application of the algorithm. *Right:* The algorithm is applied to a coincident trace.

first large long signal changes the baseline for the following events. That such signals exist is highlighted with the example in Fig. 5.27. Even though this event is not used in high-level analyses so far, because the overall number of stations with a long signal is small, we can assume that similar time scales and effects might occur with larger events too.

In practise, we define the algorithm as differences of integrals taken in predefined parts of the trace. This requires to choose a set of intervals $\{I_i\}$ that can be used for these differences.

Formally expressed, the algorithm has one threshold parameter x_{thr} and can be written as

$$\sum_{i=2}^{n_{\text{inter}}} |S_i - S_1| > x_{\text{thr}}, \quad \text{where} \quad S_i = \sum_{j=a_i}^{b_i} t(j) \quad (5.7)$$

is the integral of the trace in the i th interval $[a_i, b_i]$ and where n_{inter} is the number of such intervals.

The performance of the algorithm is critically dependent on the choice of the intervals. Based on the idea that in the data we see starting-signals, ending long signals, and contained (smaller) signals, we use three intervals, even though the obvious choice would be two. The first interval covers the first 50 bins, $I_1 = [0, 50)$, and is intended to either capture baseline or the ending long signals that decay in the first part of the trace. The second one, $I_2 = [275, 325)$, is intended to find the long signals that are positioned correctly around the latch bin leaving out the first few signal bins to avoid false positives with shorter signals. The last interval, $I_3 = [664, 714)$, should contain baseline in case of an ending long signal or the largest amplitudes of the starting long signals.

We do not scan these intervals because it is challenging to systematically cover the available phase space, especially if the number of intervals is not fixed. In addition, with the limited amount of events available in the test sample, “overfitting” might occur. A handful of different values for the integration intervals were tested by hand but no improvement with respect to the given set of intervals was found, solely small decreases in efficiency.

As for the other algorithms, two ways of handling the multiple PMTs are possible. The traces of the available PMTs can be averaged before applying the algorithm, i.e. $t(i) =$




Figure 5.32: Efficiency and number of false-positive signal GPS-seconds for the integral differences algorithm as function of the chosen threshold x_{thr} . *Left:* Efficiency for both options of combining the PMTs. The lines mark the respective chosen working point for the two options of combining the PMTs. *Right:* Number of signal GPS-seconds having at least one false-positive flag.

$1/n_{\text{PMTs}} \sum_{\text{PMTs}} t(i)$, or the algorithm is applied to each PMT and then the number of passing PMTs is checked. In our test sample, we check both options and apply the same coincidence conditions as for the other algorithms: at least two working PMTs and full coincidence. We have also tested other options for the coincidence conditions and found only lower efficiencies at the working point.

Fig. 5.32 shows the efficiency and the false positive rate in signal GPS-seconds for different threshold settings of the algorithm. At the working point we obtain the efficiencies $\epsilon_{\text{avg}} = 0.7$ for a threshold of 865 FADC and $\epsilon_{\text{PMT}} = 0.74$ for a threshold of 655 FADC. Both values are larger than the efficiencies obtained with the other algorithms.

In conclusion, we have three algorithms that perform similarly on our test data, while the algorithm “integral differences” slightly out-performs the other two. The fact that all three algorithms provide a reasonably good efficiency is important because it gives options in the implementation in the field should one or the other algorithm not be fast enough. We discuss the actual implementation and its test in the field in the next section. To highlight the performance advantage of the “integral differences” algorithm, we show the range of false positive rate and efficiency for all three algorithms together in Fig. 5.33. We limited the scanned parameters to the range of one false-positive signal GPS-second to make the choice of the working point, i.e. the highest efficiency in this plot, clearly visible.




Figure 5.33: Comparison of the efficiency and false positive rate of the different algorithms for the subset of parameters that allow one false positive tag during a GPS second with a signal event. It is evident that the “integral differences” algorithm is performing better than the other two. However, also the very limited statistics of the sample is visible as steps in the curves.

5.4.5 Implementation in the UB

In the previous sections, we discussed possible algorithms to tag long signals in the station and which of the proposed ones is the optimal choice. Here, we describe the challenges and constraints in the actual implementation in the station software and how we successfully implemented the algorithm in the field.

The delicacy in the implementation is the choice of the interface, i.e. the trigger flags that are sent to CDAS. In the 4 bits available already other auxiliary analyses are performed, e.g. the search for very large signals in single stations as pointed out in Section 3.2.2. From Table 5.2 it is clear that there are no “free” trigger flags that are interpreted as ToTs by the central trigger process. To avoid complications and due to the risk of data loss, we prefer to not modify the central trigger code. Therefore, we have to change the meaning of the different trigger flags. Even in the current software the different trigger flags are determined by prioritising for instance the MoPS over ToTs. This hierarchy is illustrated in the flowchart in Fig. 5.34.

The intricate structure is an expression of the complications of being economic with computation time and using shortcuts when possible because the online calibration is running in the same program as the determination of the trigger flags. For example, the algorithm for finding “wide” signals, i.e. those with an area over peak in units of Q_{VEM}/I_{VEM} larger than 1.2, is running only on the subset of latch-threshold events. In this context, the “latch” trigger refers to the type of trigger that first triggered in the FPGA and lead to the readout of the trace.

The selection of latch thresholds is motivated by the observation that (small) showers close to the core trigger mostly as threshold-triggers and thus skips computation for all ToT-only triggered traces. The overall sensitivity of station triggers of stations supporting this selection is discussed in detail in Section 5.6.1.

The limitations on the choice of the trigger flags from the central trigger algorithms are that all trigger flags smaller than “7” are interpreted as threshold and those above 7 as ToTs. From Table 5.2 it is obvious that for threshold triggers, there are still unused trigger flags.

Thus, we chose “3” as new special value to indicate a threshold T2, i.e. not any ToT-like trigger active, flagged as long-signal by the algorithm to be implemented. For the ToT-side, we change the interpretation of all flags out of the necessity to free a flag and our goal to keep all additional algorithms running as much as possible. The result, shown as flow-chart in Fig. 5.35 gives the “long-signal” flag “15” the highest priority followed by the search for low-gain saturation that is kept at trigger flag 8. For the other flags, the usual hierarchy ToT > ToTd > MoPS is followed and the even flags stay associated with “wide signals”.

We attribute the long-signal information the highest priority because it is the only flag that is intended to be used directly during the data taking. Due to this requirement of running online, we also have to remove some shortcuts in the program, such that requirements on the latch trigger are removed. Programmatically, this leads to a minor change of what is passing through the calibration trigger. Because the ToT-triggers make up only about 1% of the total shower traces, this change has only very small effects on the actual calibration itself. We will make sure that there are no detrimental effects of this choice in real data when testing the software in the field.

The final implementation of this algorithm is available as branch *long_signal_tag* in the software repository of the UB⁸ and uses the exact values determined in the previous section. Should a re-implementation in the upgraded UUB be necessary, or should a significantly larger data set be available, these values should be re-visited and the use of the online calibration should be considered.

After successful initial testing of the new software in station “Didi” on the Observatory campus by R. Sato, the software in the stations of the test-hexagon⁹ is updated to the new version called V0R9B6P12 on 23 March 2021. With the data from these stations, we can see if the software works as expected.

Crucially, the stations keep sending triggers and participate in events after the update showing no obvious signs of malfunction. With the collected trigger data, we can then test whether or not the new flags are working as expected. We expect false-positive tags regularly even in normal data taking conditions, due to events as the one shown in Fig. 5.36. The tails of saturated signals have enough signal in the integration windows, highlighted as red boxes in the right panel of Fig. 5.36, that the flag is triggered.

Taking a data sample of one hour well before the update (15.03.2021) and one after it (30.03.2021), we check the distribution of trigger flags observed in the data in the left panel of Fig. 5.37. The change in the distribution of the flags is very apparent, with the main points being that the newly introduced flag “3” appears with a low rate, the rate of flag “15” drops to the expected background levels of the long-signal tag, and the ToT-flags are shifted. For the ToT flags, the change is mainly due to two effects: Switching the determination of the flag to the “usual” priority reduces the amount of MoPS triggers and most ToTs are actually “wide” as expected from their construction.

To make sure that the increased need of processing with the new software does not lead to significant lag in the processing, we check the time distribution of the triggers within the second. The idea is that if the software is not able to process the data fast enough predominantly the last part of the second might be missing. From the right panel of Fig. 5.37 it is clear that this is not the case and that the time distribution is perfectly uniform.

Using the monitoring data, a minor change in the overall T1-trigger rate is observed for the new software that can be explained by the necessary modifications to the program structure. In Fig. 5.38 we show the T1-rate of two stations over the time scale of several days

⁸Accessible collaboration-internally at <https://gitlab.com/auger-observatory/sde/ub/> as member of the Auger-collaboration.

⁹This refers to the stations with Ids (943, 944, 945, 948, 949, 954, 972). The data from these stations is by default excluded from high-level analyses to facilitate software tests without the danger of compromising used data.

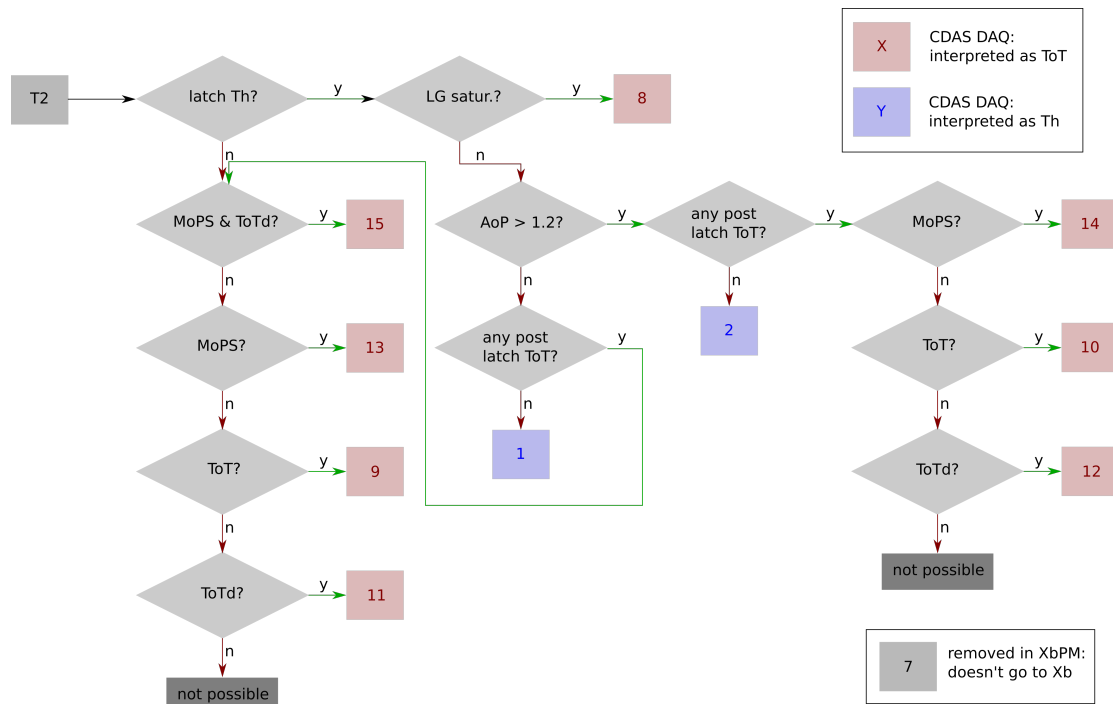


Figure 5.34: Schematic flow chart of the determination of the trigger flags as sent by software version V0R9B6P11 of the local station software. Red boxes indicate trigger flags that are interpreted as ToT-like by the T3 algorithms in Xb. Blue boxes are used for those representing threshold triggers. The removed trigger flag 7 is used for sending the scaler counts. The term “latch” trigger refers to the trigger in the FPGA that triggered first, while post-latch ones are those that fulfil their respective condition after the trace was already created.

in March 2021. The update during the 22.03.2021 is visible as minor change in the rate. The change of about 2 Hz for Station 943 (top in Fig. 5.38) is the maximal change observed and well within the tolerance of the system.

After continuous monitoring of the test hexagon for several weeks, without observing any problem with the software, the software of 21 more stations¹⁰ was updated on the 21st of April 2021. Monitoring these in the period afterwards showed also no problems.

At the time of writing this thesis, the update of the remaining stations was just performed (09.11.2021) but no data was available to be included in this work. The implementation of the priority read-out in the CDAS readout is still pending because of concentration of man-power on the on-going upgrade of the Observatory and the necessity to have all stations running the new software.

¹⁰Covering the corner of the SD-grid with the test hexagon. The stations 600, 601, 617, 618, 620, 621, 627, 630, 633, 648, 940, 941, 1053, 1057, 1058, 1059, 1060, 1061, 1063, 1067, 1314 were updated.

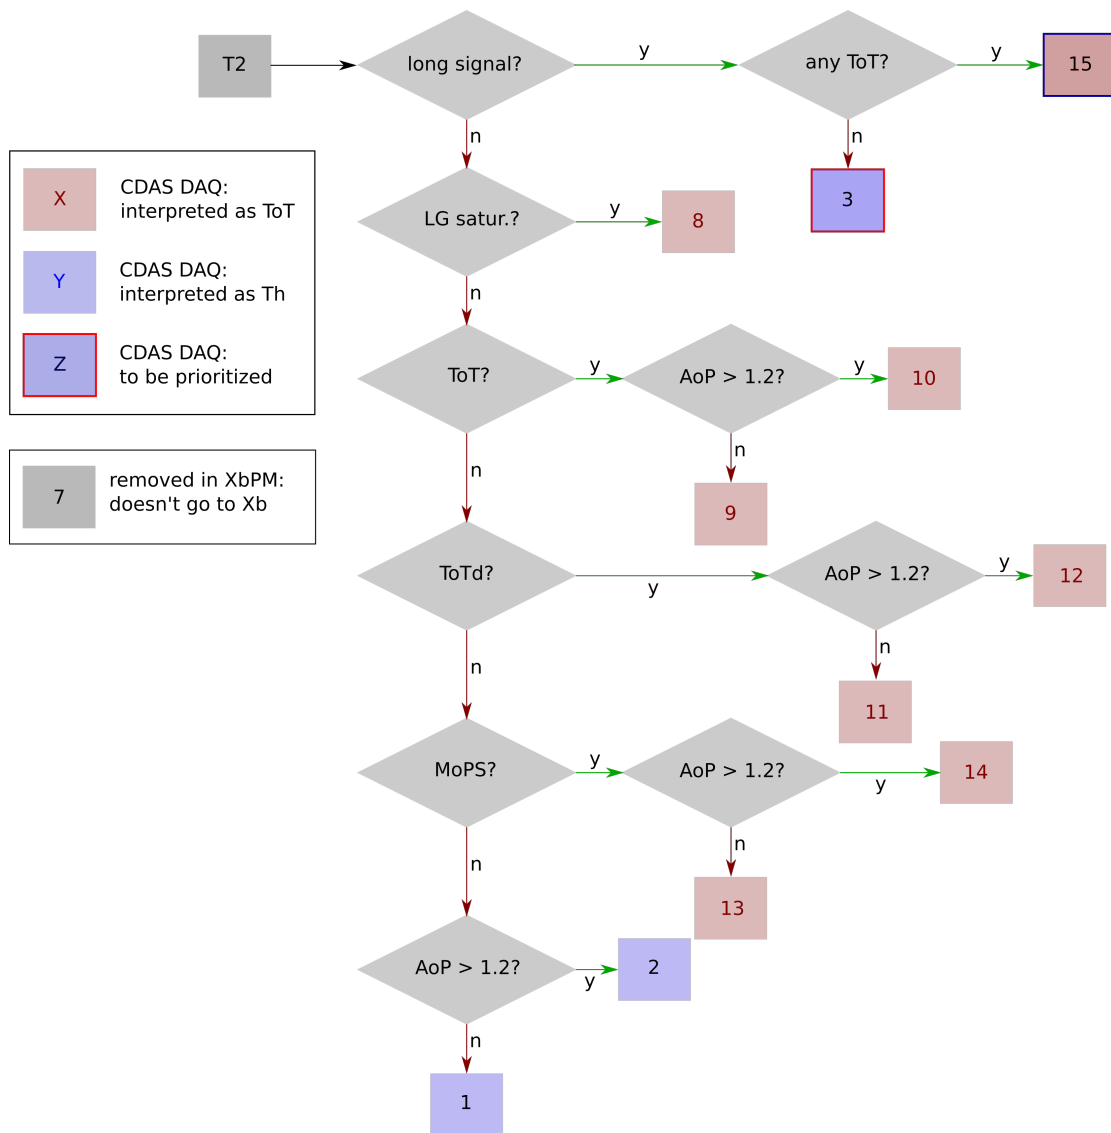


Figure 5.35: Schematic flow chart of the determination of the trigger flags as sent by the new software version V0R9B6P12 of the local station software. Blue boxes indicate flags of threshold triggers and red boxes ToT-like triggers. If a box has a border of the other colour, the respective flag is to be used in prioritizing the read out.




Figure 5.36: Example event with a false-positive flag of long signals. *Left:* Footprint of SD event 62641929 with station 1694 marked. The trace of this station as shown in the *right* panel would have been tagged as long-signal candidate by the implemented algorithm. The windows the algorithm is using are marked with red boxes.




Figure 5.37: Comparison of the trigger data after the software update to version V0R9B6P12 from V0R9B6P11 on the 23rd of March 2021. All updated stations from the test-hexagon (943, 944, 945, 948, 949, 954, 972) are stacked in the figures. The sample for the old software is taken from 15th of March 2021 00:00 UTC while the sample for the new software from the 30th of March 2021 00:00 UTC. *Left:* Frequency distribution of the trigger flags. The new flag “3” is appearing (rarely) indicating that the algorithm is working and compatible with expectation of a low false-positive rate in normal data. *Right:* Distribution of the fractional-second part of the trigger times. No indications of missing triggers towards the end of the second are visible, thus no obvious problems with computational load are detected.




Figure 5.38: Distribution of the T1-rate in the monitoring stream for two of the stations (943, 944) with the updated software. The time of the software update is visible as a minor shift in the T1 rate that is of different amplitude for the two stations.

Table 5.2: List of possible T2 flags in a T2 message sent by the station and their respective interpretation in the current (V0R9B6P11) and new software (V0R9B6P12). The column “interpretation in Xb” gives information on how the central trigger process “Xb” is interpreting a given flag. Flags marked as “invalid” are not used and can thus only appear due to corrupted data.

flag	interpretation in Xb	P11	P12
0	Th	invalid	invalid
1	Th	threshold-T2	threshold-T2
2	Th	threshold (wide)	threshold (wide)
3	Th	invalid	threshold (long)
4	Th	invalid	invalid
5	Th	invalid	invalid
6	Th	invalid	invalid
7	scaler	scaler	scaler
8	ToT	LG-saturation	LG-saturation
9	ToT	ToT	ToT
10	ToT	ToT (wide)	ToT (wide)
11	ToT	ToTd	ToTd
12	ToT	ToTd (wide)	ToTd (wide)
13	ToT	MoPS	MoPS
14	ToT	MoPS (& MoPS + ToTd) (wide)	MoPS (wide)
15	ToT	MoPS & ToTd	any ToT (long)

5.5 Possibilities with the Upgraded Electronics

The upgrade of the Pierre Auger Observatory [4] includes the upgrade of the station electronics. With the upgraded unified board (UUB) several possibilities of enhancing the observatory's capabilities for detecting these exotic events arise. In this section, we give an overview of what components of the DAQ are affected and where improvements are possible.

There are three main parts that determine the efficiency of detecting SD-disks (or cosmic-rays) apart from the hardware: the station firmware with the (T1) triggers, the station software with calibration and threshold settings for the triggers, and the central part of the DAQ CDAS with the T3-algorithms.

While the old electronics is limited in both CPU-time and available resources for additional triggers, the new electronics are capable of having new firmware triggers. Thus, we can try to design a firmware trigger that is very efficient in detecting long signals and at the same time good at rejecting lightning traces. Such a trigger would however be only a real improvement if the threshold can be lowered in this way, i.e. long signals of very small amplitude can be detected. The other possibility of improving the detection capabilities of the SD with a new trigger is related to the fact that the missing central stations are likely to originate from "saturated" triggers. In case too much RF-noise is present in the station, the triggers stop working properly because they are active all the time and no new trace is created. A new trigger that is not sensitive to the lightning noise could thus increase the efficiency in this crucial part.

A second option to increase the information recorded for the long signals is to try to extend the length of the traces to catch the complete signal also for the large amplitude cases. To not overload the communication system the information has to be compressed in any such approach and the triggering has to be efficient and selective enough to keep the amount of data to transmit small. Rebinning the trace can be such an approach that we discuss in Section 5.5.3.

With the increased computational power of the UUB, there also are possibilities to do better processing of the traces in the station software. A tagging algorithm similar to the one implemented in the old electronics could lead to a dedicated T3-level algorithm for these events. One caveat in this is the false-positive rate that might need to be adjusted to the limitations of the communication systems and the handling of interferences with normal data taking.

In the following sections we investigate the aforementioned possibilities in more detail focussing on sketching what can be possible. We do not want to give definite recommendations on what to implement in the field due to a lack of detailed knowledge about the signals in the UUB. We hope that with the implementation of the prioritized read-out in the UB we can gather more data and at the same time obtain simulations that can tell us what these signals are in detail. With such simulations the development of a dedicated firmware trigger can also be done with much more rigour.

5.5.1 Input Data

For the development of any new trigger or software tag the input sample to tune and test the proposals is crucial. In our case, the situation is complicated because we neither have long signals with the UUB nor a simulation of the signal. Therefore, all the work done here is to be seen as preliminary study on what is potentially possible but should be followed up by extensive tests once data and simulations are available.

Nevertheless, we can work with the traces we do have, even if they are recorded with the UB, keeping in mind that in the UUB there is also a down-sampled signal that is running the compatibility UB triggers. Using the UB traces as first and main data set, we visually scan the

events around the known SD-rings and SD-disks and classify the traces in “long”, “cosmic-ray”, and “lightning”. The selection leads to a total of 6175 traces, with 1185 long-signal traces and 3442 lightning traces.¹¹

A clear caveat of this sample is that the traces are only available because they were already triggered by an existing trigger, i.e. there is little to gain from trying to trigger on the very same signals. However, adding a new T1 that is only close to fully efficient for the signal traces – that are triggered by the other triggers – still has advantages. We learned that a major problem in the central part of the events is that the stations do not trigger, likely due to all triggers being active because of RF-noise. An additional trigger that does not trigger on most of the lightning noise but is efficient for the large amplitude long signals could potentially fill these holes. This requires however careful tuning and testing. The question for this part of this work is therefore: How much background suppression can we achieve with reasonable signal efficiency?

To make the conditions for any tested algorithms similar to the situation in the field, we use the online estimates of the baselines, the online information about masked PMTs, as well as potentially the online VEM-peak to set thresholds. We will however focus on using simple metrics, like FADC thresholds, since the purpose of this study is to highlight possibilities and not to give definite answers that would not be reliable in the first place due to the limitations of the input sample. Also to avoid extensive scans of coincidence conditions, we simply use the average PMT signal as the only trace for the tested algorithms.

To sketch the potential of detecting smaller signals than the observed ones, we will try to use the recorded signals themselves rather than parametrisations. We scale the traces with a known factor < 1 to obtain a signal of similar shape but smaller amplitude. In detail, we apply the scaling only if the signal is more than one FADC count above the baseline estimate and round the result to obtain a new integer trace. In the right panel of Fig. 5.39 an example of such an rescaled trace together with the original is shown. Once available, these trace should be replaced with properly simulated ones of small amplitude.

The last data set we can use in these tests is a sample of one million UUB traces that were triggered with a random trigger [116]. This corresponds to about 17 s of data that can be used to determine how many background triggers a given algorithm would produce. Due to the lack of the decisive lightning noise in this sample, it is of limited value for the purpose of developing a long-signal trigger even if just targeting the rejection of lightning noise.

Nonetheless, valuable information can be obtained from this sample with respect to the question of how such a trigger would behave during normal data taking. In addition, we can do implicit cross-checks of the traces and behaviour of the UUB in the field because any trigger sensitive to exotic signals should not find any. If it does, the most likely origin would be station specific problems.

The sample does contain signals from air showers, in some cases even of more than one in a single trace, like Fig. 5.40 shows. We can also see that the SSD has a much smaller signal and the traces are in principle available for use in a trigger. However, due to our lack of knowledge of how a long signal would appear in the SSD, we do not use the SSD-traces in our tests here. This is of course another point that should be addressed by future studies when finalising a real trigger.

¹¹A complete list of the event and station Ids that make up this list can be found in the repository that can be used for reproduction of the results https://gitlab.iap.kit.edu/schimassek-m/long_signal_t1/-/blob/master/configs/all.cfg.




Figure 5.39: Performance of the ToT-like algorithms on rescaled traces. *Left:* Efficiency for detecting long-signals once they are rescaled to smaller amplitudes. The ToT-algorithm is used with different settings: The “high-purity” parameters lead to a lightning background rejection of 99%, while a less strict set at 90% rejection is used for the “normal purity”. Parameters similar to the standard ToT (13 bins out of 120 at 0.2×50 FADC) as it is used in the SD give a reference. *Right:* Example trace from station 1164 in the SD-event 4067441. Both the original and the rescaled version (factor 0.5) are shown with the respective bin that triggered the “high-purity” configuration of the ToT-algorithm.




Figure 5.40: Example of a trace from the sample of 1 000 000 random UUB traces. It is trace number 446595 and it was selected using a simplified ToT-trigger with modified parameters. *Left:* Complete traces of the WCD PMTs. *Right:* The average of the large PMTs of the WCD- and the SSD-trace zoomed into the signal part of the trace.

5.5.2 Dedicated Triggers

In this section, we present the algorithms we test to enhance the SD’s capabilities for detecting long signals with their respective performance in our test sample. In line with the scope of this work, we do not perform an exhaustive search for the best algorithm because both a lack of signals with the UUB and the lack of detailed simulations make any prediction beyond order of magnitude estimates questionable from the beginning. To stress the difference to Section 5.4.4, we are interested in a real trigger here rather than a software tag added after the trace was already triggered. Thus, we will also change the processing such that the trigger algorithm has to work iteratively and never gets the trace as a whole.

The algorithms presented here are based on the very same ones that were already tested for tagging long signals in the UB: comparisons of integrals of the trace and Time-over-Threshold (ToT) triggers. In Fig. 5.41, the concepts of the two types of triggers are visualised

with real traces that were selected with the respective triggers. In the left panel, we show a ToT-triggered trace that requests at least 40 bins out of a window of 100 to be above a threshold of 20 FADC above the baseline, while in the right one an integral trigger with two windows (grey bands) is used. In the figure, the first window of 100 bins that fulfils this conditions is marked as grey band while the baseline and the corresponding threshold are marked by coloured lines. It is clearly visible that the baseline estimate from the online data is not describing the true baseline in this example. Due to the high number of bins above threshold requested, a counting and marking of those bins in the figure is not attempted. In our test, we scan the window length from 100 to 375 bins in steps of 25 bins, and the threshold from 5 to 45 FADC above threshold in steps of 5 FADC. The occupancy we scan relative to the window length from 5% to 45% of the window length.

In addition to this usual parameters of the algorithm, we try adding a condition on the integral, i.e. requesting $\int s(t)dt > x_{\text{th}}$ with a threshold x_{th} . Such an integral should in theory be able to select real signal instead of oscillating noise if the conditions and the length is appropriate. We try values from 25 FADC·time-bin to 200 but also without this condition.

For the integral difference trigger, the simplest configuration is the one presented in the right panel of Fig. 5.41: two integration windows of the same lengths (grey bands), with a delay d and a threshold t . In the figure, we plot the average value in the integration regions as visualisation of the value of the integrals. From the comparison of this dashed red lines with the black threshold it becomes apparent that this trace is (just) fulfilling the trigger condition. To test the parameters of such an algorithm, we test all combinations of the window length $l \in \{25, 50, 75, 100\}$ bins, $d \in \{50, 100, 200, 400\}$, and $t \in [100, 525]$ FADC·time-bins in steps of 25 FADC·time-bins for two integration windows. We generalise the trigger to many integration windows by introducing a threshold $t_n \leq \sum_{I_n} t - \sum_{I_0} t$ for each window with respect to the first one. In addition, we add the option of using a threshold on the sum of the differences (as done in the previous section):

$$t_{\text{sum}} \leq \sum_i \left| \sum_{I_i} t - \sum_{I_0} t \right|. \quad (5.8)$$

In addition to the configuration with two windows, we test sets of parameters for three windows, motivated by the success of the tagging algorithm using three such windows implying there are two delays d_1 and d_2 . We use all combinations of the parameters

$$\begin{aligned} t_1 &\in \{100, 150, 200, 250, 300\} \text{ FADC·time-bins,} \\ t_2 &\in \{100, 150, 200, 250, 300\} \text{ FADC·time-bins,} \\ t_{\text{sum}} &\in \{0, 50, 100, 150, 200, 250, 300, 350, 400, 450\} \text{ FADC·time-bins,} \\ l &\in \{50, 100\} \text{ bins,} \\ d_1 &\in \{0, 50, 100, 200\} \text{ bins,} \\ d_2 &\in \{100, 300\} \text{ bins.} \end{aligned}$$

The resulting detection efficiencies for all algorithms, paired with the rejection of lightning traces are shown in Fig. 5.42. The left panel shows the important combination of detection efficiency of long signals as function of detection efficiency of (triggered) lightning noise while the effect on the selection efficiency of cosmic-ray signals is shown in the right panel. There is not one single ROC-curve per algorithm in this figure because the different parameter settings allow for very different tunes of the trigger. Fig. 5.43 shows the efficiencies for the ToT-trigger, split into the three different types of signals, as a function of the occupancy parameter. From the many parallel curves it becomes evident that the spread visible in Fig. 5.42 is expected and that the optimisation problem is highly dimensional. It is however

clear, that the general feature that many bins above threshold correspond to being more “long-signal”-like.

To make the comparison of the different algorithms easier, we will quote here the best efficiency for long signals at a given minimal lightning-rejection power that we chose as 95%. Detailed studies of the necessary rejection power in the field with UUBs are required before any decision for a real trigger can be taken. For the ToT-trigger, we get an efficiency of 73% at 95.6% lightning-rejection efficiency for 67 out of 150 bins above a threshold of 10 FADCs without a condition on the integral signal.

For the integrator trigger, we get 60% efficiency at 95.4% lightning rejection for two integration windows of length 25 bins delayed by 400 bins. The necessary threshold is 525 FADC bins. For three integration windows, no set of tested parameters reaches 95% background rejection. Judging from Fig. 5.42, the performance with three integration windows is similar to that with two, except that the parameter space of high background rejection is not occupied. There might be parameter settings that achieve the required 95% rejection, but given that the overall performance is similar to that with two integration windows, we do not try to find these parameters.

While for the signal tagging (c.f. Section 5.4.4) the integrator trigger outperforms ToT-like algorithms, the case is different here. The fundamental difference here is that for tagging, the latch bin is already fixed due to the mere existence of the trace, while when designing a trigger no such fixed reference time exists. In this set-up, the ToT performs better because the region of the signal that is “skipped” in tagging can not be simply skipped when computing triggers.

To estimate how the algorithms here were to perform on signals of smaller amplitude, we use the scaled data set presented in the previous section with the example of the ToT-like algorithm. In Fig. 5.39, we show the performance of three settings of the algorithms as a function of the scaling parameter. It is clear that there is a set of large signals that can be selected even if they are scaled down significantly. However, from the comparison to the standard ToT-like setting¹² we also see that there is little room for actually improving on the small signals. With the new triggers, ToTd and MoPS, that likely make up the 15% missing to full efficiency even at unity scaling, the SD is already very efficient for these small signals. Thus, we conclude that the only real improvement to the overall sensitivity can come from triggering without lightning-noise with a high efficiency on large signals.

Using the real random UUB traces and translating the 95%-lightning rejection parameters roughly to scale for the UUB (time lengths $\times 3$ and peak thresholds $\times 4$), we get no triggers. This only shows that the expected rate of such triggers would be less than 0.06 Hz. While the distortions on data taking by such a rate would be minimal, the remaining question on how it might alter CR-event triggers has to be answered before any such implementation should be attempted. With the deployment of the UUB in progress during the writing of this thesis, the necessary data for these tests will be available in the near future.

5.5.3 Extended Traces

As mentioned in the introduction of this section, a second option to increase the information available for SD-disks is to extend the time scale covered by the data. Because the amount of data to be transmitted cannot be extended arbitrarily, we have to find a way to compress the information and be efficient in selecting the interesting parts only. A simple possibility of such an approach is to re-bin the trace such that with the same amount of bins the time covered can be increased n -fold. One possibility to implement such a rebinning in addition

¹²This is only an approximation as this comparison is performed on the average trace rather than taking the detailed coincidence conditions into account.




Figure 5.41: Illustrations of how the tested trigger algorithms work on real traces. *Left:* Data from event 4067441 of station 1164 showing a long signal. We use the signal to illustrate how the time-over-threshold trigger with modified parameters could work. The black line indicating the baseline is taken from the values of the online calibration of the station, such that the offset to the actual baseline would also appear in the field. The used threshold of 20 FADC above baseline is drawn with a red line. In this example we use a window of 100 bins and request 40 bins above the threshold. This window is highlighted as grey band. *Right:* Trace of station 326 from SD event 1332954 with a normal cosmic-ray signal. The grey bands show the integration windows used in a trigger comparing integrals of the trace.

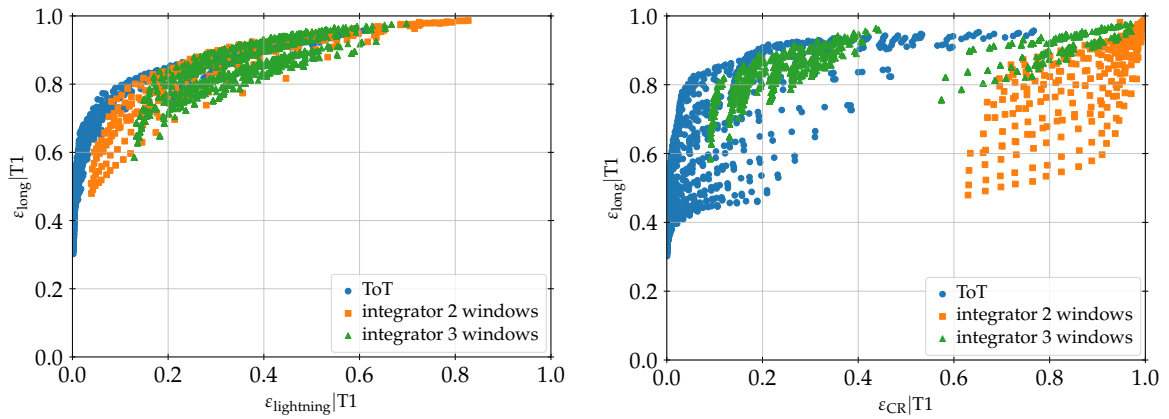


Figure 5.42: Efficiencies of different algorithmic approaches for a potential T1 and different parameter settings. The test-sample consists of UB-signals because of the lack of UUB signals. The choice of UB signals implies that all efficiencies are conditional on having triggered with one of the existing triggers. *Left:* Detection efficiency of (already triggered) long signals as a function of lightning background rejection, in this case the efficiency of tagging lightning signals. *Right:* To highlight the properties of the algorithms and settings, the efficiency of the algorithm as a function of the efficiency of finding cosmic-ray signals.

to the currently available data would be a dedicated buffer in the FPGA that is constantly filled and triggered together with the normal traces but only sent if a dedicated trigger on this buffer is active. Dedicated triggering independent of the normal traces might lead to interferences with normal CR-data and thus should be evaluated carefully only. Here, we will highlight that an average of the active PMTs averaged in time into several bins still carries significant information on the long signals and that triggering long signals on such a trace is possible.

We start from the same sample of UB traces used in the previous section and described in Section 5.5.1 and form “ n -rebinned” traces by taking the average value of n time-bins of the




Figure 5.43: Dependence of the detection efficiency of the different types of signals on the occupancy parameter of the ToT trigger. Different apparent curves correspond to different values of the other parameters. It is clear that the occupancy has little influence on the efficiency for long signals, while it has a much stronger effect on cosmic-ray or lightning signals. For the scanned range and parameters see text.

original trace, averaged over the active PMTs. To keep the length (in bins) comparable, we add baseline samples of 200 bins before and after this rebinned data. The baseline samples are generated as MC-values with a mean value corresponding to the mean baselines of the active PMTs and a spread of $\sigma = 1/\sqrt{n_{\text{PMT}}}$ FADC. In Fig. 5.44, we show an example trace of a long signal station in the original form on the left and averaged and rebinned with $n = 4$ on the right. It is clear that the rebinned version keeps the main features of the signal and has the possibility of extending the time range four-fold.

To obtain first estimates of how such rebinned traces can be triggered, we take inspiration from the studies on the T1-proposals from the previous section. Since our goal here is to only showcase that a rebinned trace concept is possible and can be triggered on, we only use the ToT-algorithm that worked better as T1. As set of parameters to test, we use

- the window length l from 5 to 120 bins in steps of 5 bins,
- the minimal occupancy o , i.e. the number of bins above threshold, from 5 to 100 bins in steps of 5,
- the threshold t as 5 to 50 FADC above baseline in steps of 5 FADC,
- the integral signal I , in addition to $o \times t$, larger than $I \in \{-\infty, 0, 40, 80, 160\}$.

While some of these parameter settings will most likely not work well in practise, we only want to span enough phase-space to show that such an algorithm can in principle be optimised enough to work sufficiently well. We show the results of the scan with these parameters in Fig. 5.45. It is evident from the comparison of the $n = 1$ points that the loss of performance with respect to the non-rebinned traces is minimal. Lightning rejection of more than 90% is certainly possible even with a rebinned trace while the efficiency for long signals remains acceptable at about 70%. We also have to keep in mind that for the long signals of small amplitude the extended read out is of less importance such that the loss of efficiency associated with them is acceptable.

At the same time, the efficiency of (mistakenly) triggering on cosmic-rays is reduced significantly for $n = 4$ because the signals are short such that after rebinning the number of




Figure 5.44: Example trace to illustrate the concept of rebinning traces. Data from station 878 of event 1332968 is shown as case of a long-signal trace that extends beyond the time range covered by the trace. *Left:* Signals of the individual PMTs in raw FADC. *Right:* Average signal in blue and the rebinned trace averaging 4 bins into one.




Figure 5.45: Efficiency of detecting long signals with rebinned traces as function of the rejection power against lightning traces and cosmic-ray traces with a ToT-algorithm. We show two possibilities of rebinning the traces $n = 2, 4$ and keep the original trace as reference with $n = 1$. *Left:* Efficiency of triggering on long-signals as a function of the efficiency on lightning traces for different parameters of the ToT-algorithm. *Right:* Using the efficiency on cosmic-ray traces as background instead of lightning traces. This background is important to estimate the influence of possible rebinned traces on normal CR-data taking.

bins above baseline is simply too small to be detected. This feature is actually good because this means that such a rebinning approach is intrinsically bad at adding unwanted data to the cosmic-ray data stream.

All in all, we can see from these efficiencies that such an approach is promising and should be pursued in detail. However, the correct timing, trigger logic, behaviour in non-triggered background, and all possible interferences with data taking have to be considered before a prototype can be added to the DAQ of the UUB. The exact implementation will depend on the available resources in the electronics and should thus be designed once the DAQ of the UUB is working stably in the full array in the near future.

5.5.4 Read-Out Adjustments

After discussing possible extensions of the trigger capabilities in the UUB in the previous sections, we focus now on inefficiencies generated by the software and further ideas that can be tested in future work.

As we point out in Section 5.3.2, we do not record some long-signal traces because they are requested together with a lightning trigger that happened immediately before the long signal. This means that we can increase the efficiency of the DAQ-chain by modifying the station software in a way to properly handle these sort of double requests. Because it touches the core of data taking, such updates should be implemented and tested carefully, however, we consider this modification a “low-hanging fruit” because there is no conceptual change to the DAQ.

Another idea that can help enhance the data already taken without the need for additional triggers would be another trigger flag. Because the UUB should be able to process all T2s with more computationally expensive algorithms than the UB, we could add a T2-flag that indicates overshoots. This is desirable because it allows inference on potentially missing signals, as in Section 5.3, even when no T3 read the overshoot traces out because of the T2Dumps. Algorithmically, it is surely possible to detect a linear increasing trace as signature of overshoots. We leave the discussion of details on algorithms for future work, noting the issue of the again limited data sample because the number of overshoot traces known is even smaller than those of long signals.

Going away from easy and little invasive changes to the DAQ, we could imagine that a dedicated T3 that tries to capture repeated events as seen in some of the smaller SD-disk-like events might be able to lower the threshold on the size of the events because fewer stations might be necessary. However, because it is likely that also from RF-noise there will be repeated signatures any such T3 probably would need a working long-signal tag at the T2-level. As the implications of adding a new T3 mode are large and the potential gain is hard to estimate with the data we currently have, we consider a detailed proposal beyond the scope of this work.

As final remarks on the DAQ software related to SD-disks and lightning events, we want to add that it might turn out to be necessary to handle lightning data differently with the UUB than before with the UB. Because of the larger buffers of the UUB data loss is much more unlikely, reducing the need for a prioritised read-out, but also increasing the probability that excessive amounts of (lightning) event data “block” the communication system during thunderstorms. In the pre-production array of UUB this effect likely lead to a loss of monitoring data during 31 January 2021 [117]. Should there be the possibility that large thunderstorms lead to a loss of monitoring and calibration data that have a lower priority than event data, the DAQ should be adapted. In this case, we assume that a lightning detection in the station software with a corresponding T2-flag can be a minimally invasive solution. It can enable the central trigger algorithms to scale down the rate of T3s with only such lightning triggers randomly to keep possible biases low while keeping the event rate at acceptable levels. However, as with any other change of the DAQ, such an algorithm should be extensively tested and only used if required to keep the interferences with continuous data taking low.

As a summary of these thoughts and the concepts shown in the previous sections, we can say that the UUB offers unique chances to increase the observatory’s capabilities for recording exotic events. However, there is still considerable amount of work to be done to arrive at detailed proposals. Given that at the time of writing of this thesis the first UUBs are being deployed, the near future is the right time to follow up on these ideas.

5.6 Search for Correlation of Cosmic Ray Showers and Lightning

Moving away from the detection of exotic events that we observe in the surface detector, we explore the possibilities of contributing to the lightning initiation problem with the capabilities of the SD. The long-standing issue of how the lightning is initiated in nature is still not (entirely) understood (for a recent summary and proposed model see e.g. Ref. [53] and references therein). Among the proposed models are initiation by charged hydrometeors (e.g. [118]) that can cause local field strengths high enough to create electron avalanches [119], and electron-avalanches seeded by cosmic-rays [120–122].

For the latter model, primary energies above about 10^{15} eV are necessary to reach the required charged particle densities within the cloud [53, 123]. In the proposed model in Ref. [53], the main mechanism of initiating the lightning is a network of electron avalanches (following Ref. [122]) initiated by a cosmic-ray. While the energy of these showers are well below the threshold usually quoted for the Pierre Auger Observatory, we have already discussed in Chapter 4 that this threshold can be substantially lowered if the usual ways of reconstructing events is omitted. Thus, we will discuss here how sensitive the Auger-SD is to such low-energy showers and perform a search for such initiating cosmic-ray events to try to confirm or potentially rule out the model of Ref. [53].

5.6.1 Sensitivity of SD-T2s to Low Energy Events

There are several possibilities to test the sensitivity of the Auger-SD to low-energy events. The usual way to obtain trigger efficiencies is to use Monte-Carlo simulations of the air-showers together with detector simulations. However, because the validation of the trigger simulations code was not finished until the time of writing of this thesis [124] and a good energy resolution is not required for our purpose here, we want to use the data itself to estimate the sensitivity. In future work, the results of this work should thus be confirmed with detailed simulations.

In the following, we present a set of analyses that continuously try to reduce the threshold of detection. Also in this section, we make use of the T2 triggers from the T2Dumps to carry out the low-energy searches. We start by using coincidences of nearest neighbours in the SD-grid, effectively removing the necessity of a third trigger compared to the event trigger. To reduce the threshold even further down, we test the same construction on the infilled SD-750 to finally discuss the sensitivity of a single station to small showers.

Nearest Neighbour Coincidences The first possibility to significantly lowering the threshold of the Auger-SD is to remove the requirement for three triggers. Thus, we want to estimate here the effective energy range selected by requesting two nearest neighbour stations to trigger in coincidence. To select such pairs of triggers, we scan through the T2-data and form clusters of *connected* triggers where we use the connectivity condition:

$$\text{con}(M_1, M_2) = \begin{cases} \text{true} & \text{for } |d(M_1, M_2) - 1500 \text{ m}| < 250 \text{ m} \ \& \ \Delta t \leq 5 \ \mu\text{s}, \\ \text{false} & \text{else,} \end{cases} \quad (5.9)$$

for two triggers M_i with the distance of the corresponding stations d and the time difference of the reported trigger time Δt . The chosen maximal time difference is motivated by the time that a signal needs to propagate 1500 m at speed of light. In practise, this algorithm will select also events with more than two triggers, as long as they are connected through nearest-neighbours. A simple cut on the amount of triggers per such cluster can then be used to sub-select true sub-threshold events formed by two triggers only.

Due to jitter in the raw timing provided in the trigger-time tag, the GPS timing of the trigger can be different from the time tag by more than the rounding¹³ error. An estimate of this time jitter is necessary to estimate whether or not the tolerance of 5 μs is inclusive enough to detect all relevant coincidences. Fig. 5.46 shows the distribution of the difference in time from T2-time tag to GPS nanosecond as read from event data, i.e. $\Delta t = t_{\text{T2}} - t_{\text{GPS}}$ where both times denote the end-of-trace time. The two panels, left and right, show different months of data, January and May of 2021 respectively, highlighting the influence of thunderstorm periods in the tails of the distribution for the austral summer when thunderstorms occur more frequently. Overall, the fitted width of the time jitter is 470 ns in winter and 486 ns in summer. Compared to the variance of a uniform distribution originating in the rounding of 288 ns, the time jitter is increased roughly by a factor two, however still small enough that the time tolerance of 5 μs is good enough. This is mostly due to the lowered sensitivity to highly-inclined showers that transverse more atmosphere than vertical ones. For example, for a shower at 45° zenith angle, the expected maximal time difference is solely 3 μs leaving 2 μs tolerance.

For the upgraded stations (UUB) in Fig. 5.46, the sole contribution to the time jitter is the rounding, improving the precision of the trigger-time tag significantly. This means that in the future better selections can be possible. Due to the constraint that the trigger time is not necessarily solely a shift of the signal time, especially if random background events can trigger a station too, a change in the time tolerances for any algorithm has to be evaluated very carefully. For the study presented here, this result is of no relevance as no UUB data can be used in the correlation analysis.

For estimating the efficiency of triggering two stations, the information about the triggered surface detector stations should be available in the event data, due to the hybrid set-up of the DAQ chain [1]. We opt for not using the event data though, so that the efficiency of selections in the T2-data is directly included in the efficiency estimate. Using reconstructed hybrid events [125, 126], we apply a reduced set of quality cuts, starting with the standard X_{max} quality cuts [15] to get a reliable sample. We remove the fiducial field of view cut [2] because for our purpose event statistics is more important than an unbiased X_{max} -distribution. With the same argument we remove the cut on the observed X_{max} . Finally, we reduce the required minimal track length from 200 to 150 g/cm^2 and add a cut on the energy uncertainty to be less than 50%. The cut on minimal reconstructed energy is removed to go as low in energy as possible. These last modifications are all motivated by the fact that statistics at low energy is more important to roughly estimate the sensitivity of the SD than the precision of the FD-reconstruction. Future work should be dedicated to quantifying the systematics of these estimates.

With this input sample, we select the GPS seconds of selected events and collect all nearest-neighbour clusters of triggers using Eq. (5.9). We then search for a cluster corresponding to the event by calculating the time difference to the expected arrival time of the shower plane, using

$$t_{\text{shower}} = \vec{a} \cdot \vec{r} / c, \quad (5.10)$$

where $\vec{r} = \vec{x}_{\text{station}} - \vec{x}_{\text{core}}$ is the distance of the station to the estimated core. For simplicity, we use the first station of the tuple and request that the estimated arrival time of the shower is not more than five microseconds different from the trigger time subtracting 13 μs to account for the latch bin position.

Fig. 5.47 shows the resulting efficiency estimates in the data sample from January to end of March 2017. As expected the efficiency saturates for $\lg(E/\text{eV}) > 18.5$ save for fluctuations that are caused by non-active hexagons and border effects. Because the purpose of this

¹³The integer nature of the provided time tag leads to a flooring of the GPS nanosecond rather than a classical rounding.

study is not the confirmation of the full-efficiency threshold, and the selection conditions for sub-threshold events can be different and more relaxed, we do not try to include these small effects in this study. The important observation in Fig. 5.47 is that the SD array is more than 90% efficient with pairs of triggers down to $\lg(E/\text{eV}) = 17.7$, lowering the threshold almost a full decade. To illustrate the contribution of pairs, i.e. events that have exactly two nearest neighbours triggered within $5 \mu\text{s}$, we show the efficiency for this selection in red and blue. It peaks at $\lg(E/\text{eV}) = 17.3$ and extends with more than 10% efficiency down below $\lg(E/\text{eV}) = 17.0$.

Interestingly, the contribution of threshold triggers to this efficiency is almost negligible, as highlighted by the vanishing difference between all triggers (red) and ToT-only triggers (blue) in Fig. 5.47. If an efficient selection is possible without threshold triggers, the contribution of random coincidences is strongly reduced due to the approximate quadratic dependence of the coincidence rate on the trigger rate. We can calculate the random coincidence rate of a pair of stations with trigger rate Γ with

$$N_{\text{coinc}} = \Gamma \times p(1 \text{ trigger} | \Delta t < 5 \mu\text{s}) \quad (5.11)$$

$$= \Gamma \times (\Gamma \Delta t \exp(-\Gamma \Delta t)) \quad (5.12)$$

$$\approx \Gamma(\Gamma \Delta t(1 - \Gamma \Delta t + \mathcal{O}((\Gamma \Delta t)^2))) \quad (5.13)$$

$$\approx \Gamma^2 \Delta t \quad (5.14)$$

where we assume $\Gamma \Delta t$ to be small. In the case of T2s and a coincidence within $5 \mu\text{s}$, this quantity evaluates to roughly 4×10^{-5} . Thus, a reduction of the rate by a factor 20, from 20 Hz (Th-T2) to about 1 Hz (ToT), equates to a reduction in random coincidences of 400.

In the right panel of Fig. 5.47, we use further sub-selections based on the reported trigger flags in the pairs to evaluate possibilities to make efficient shower selections possible in difficult conditions, e.g. thunderstorms. Due to the presence of high-frequency noise in the traces during thunderstorms, the selection of real events that include the new triggers is very challenging. However, due to the presence of the “wide” flag in a significant fraction of the true events, it seems still possible to use those triggers to reject lightning events at the expense of losing 80% efficiency.

To estimate the efficiency of trigger pairs in the SD-750, we repeat the same analysis with a few modifications. We select only events that are closer than 1 km to any SD-750 station and reduce Δt_{max} in Eq. (5.9) to $4 \mu\text{s}$ to account for the different spacing. The resulting efficiency is shown in Fig. 5.48. In this data sample the statistics is an even more limiting factor due to the smaller size of the array and the constraint that events of lower energy have to be close to the FDs to be detected. Nonetheless, we see that the threshold energy is further reduced by about a decade in energy, however at the expense of exposure and covered area.

Single Station Efficiency To reach even further down in energy, we focus now on showers that trigger only a single station. Because the FD-hybrid reconstruction requires a triggered station, it cannot serve as independent data set for estimating the efficiency of single stations. Even with simulations, the task of correctly estimating the efficiency is not straightforward, because the uncertainty on the spectra and composition lead to potentially significant systematics. However, because with only the trigger information available there it is not possible to reconstruct the energy, we can resort to rate arguments and measurements with doublets to estimate what energies contribute to what extent to the 20 Hz of T2s.

Already in the beginning of the Auger observatory such estimates were performed [127] to estimate the expected rate of coincidences in so-called “doublets”. Doublets are pairs of stations that are about 10 m apart and as such will measure the same shower in coincidence, if that shower has an extent beyond 10 m. To get a rate of (true) coincidences, the radial




Figure 5.46: Determination of the time jitter of the T2-message. Using event data the time of the T2 can be compared to those of the event GPS tag as $\Delta t = t_{T2} - t_{GPS}$. *Left:* Distribution of Δt in the SD-events of January 2021 for both UB and UUB stations. The σ of the gaussian fit for the time jitter of the UB stations is (470.6 ± 0.3) ns. It is likely that the tails of very large Δt s are related to effects of lightning strikes that happen in austral summer with high frequency. *Right:* Using all data from May 2021, the time jitter changes to (485.5 ± 0.7) ns for the UB stations, while the (lightning related) outliers are not present any more.

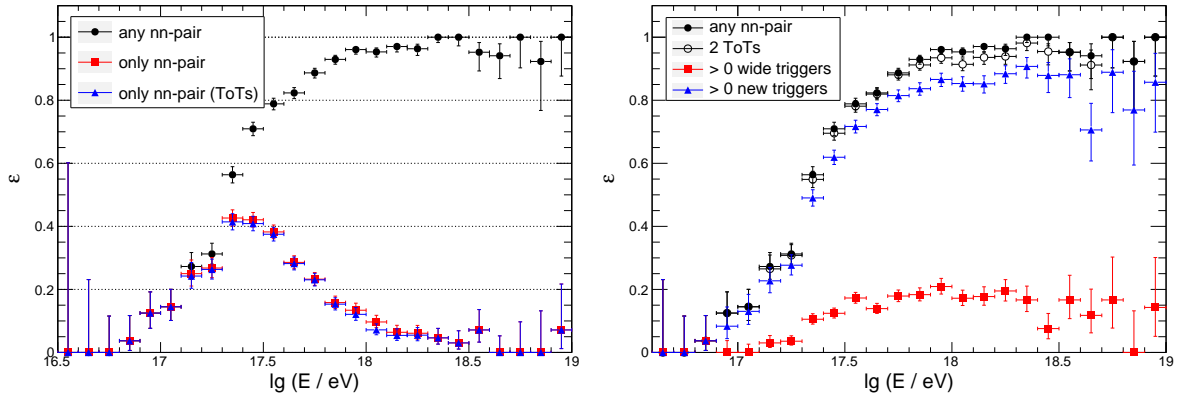


Figure 5.47: Sensitivity of pairs of T2-triggers in nearest neighbour stations. *Left:* Total efficiency (black) as determined with hybrid events [125, 126] with reduced quality cuts [15] to move the energy threshold of the hybrid events down to lower energies. The red dots indicate the efficiency if in addition to the condition of two neighbouring triggers within $5 \mu s$ no additional triggers, i.e. no triplet of triggers, are allowed. The blue triangles make the additional selection on the nearest-neighbour-only events that both triggers are a ToT type. Deviations from full efficiency at highest energies are due to reduced quality cuts and missing cuts on non-working stations. *Right:* Details of the trigger types contributing to the detected events.

contour of simulated showers at trigger threshold is calculated, as shown in the left panel of Fig. 5.49, and convolved with the spectrum of cosmic rays taking the radii that lead to a coincidence into account. The resulting spectrum is shown in the right panel of Fig. 5.49, and integrated it predicts a coincidence rate of 0.9 Hz.

To confirm these results, we can use the T2s from the T2Dumps of existing doublets with station Ids (72, 688), (79, 819), (80, 669), (89, 734), (90, 651), (91, 643), and (93, 710). Taking the time jitter shown in Fig. 5.46 into account, we allow for up to $3 \mu s$ of time difference between the stations. We can be conservative in this choice because the number of random coincidences for a pair of stations is, following Eq. (5.14), with about $\Gamma_c \approx 1.2$ mHz much smaller than the expected signal. In the left panel of Fig. 5.50 we show the results of this coincidence search separated by trigger type available in the station software version V0R9B6P11 (c.f. Fig. 5.34) for data from the 1st of June 2020. Overall, we can confirm the expected rate at a measured

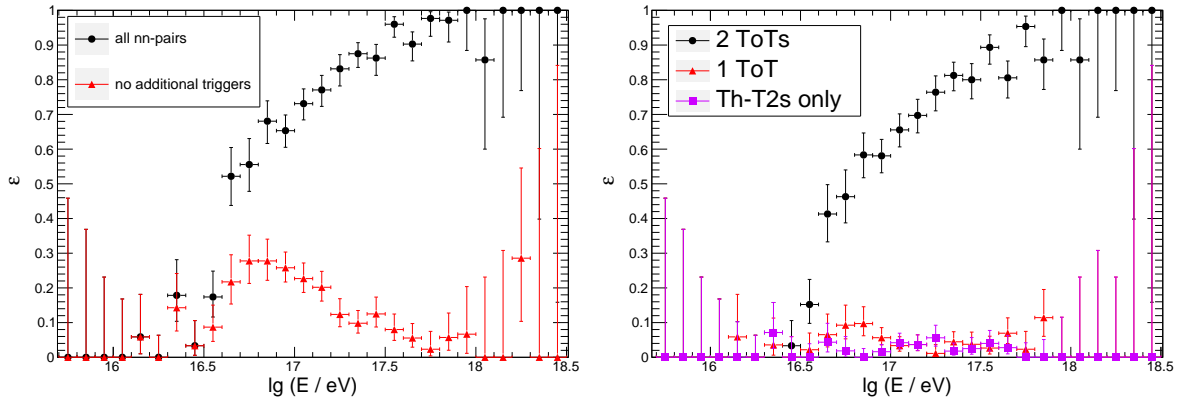


Figure 5.48: Sensitivity of pairs of T2-triggers in nearest neighbour stations in the infilled SD-750 array. *Left:* Total efficiency (black) as determined with hybrid events [125, 126] with reduced quality cuts to move the energy threshold down. The red dots indicate the efficiency if in addition to the condition of two neighbouring triggers within $4\ \mu\text{s}$ no additional triggers, i.e. no triplet of triggers, are allowed. Deviations from full efficiency at highest energies are due to reduced quality cuts and missing cuts on non-working stations. *Right:* Details of the trigger types contributing to the detected events. Out of the two triggers forming the nearest-neighbour pair, none, one, and two are requested to be a ToT for the purple, red, and black points. It is very clear that the vast majority of events is detected with two ToT triggers.

0.7 Hz. Interestingly, about 80% of the coincidences appear as simple threshold triggers, as the red markers in the right panel of Fig. 5.50 show. The black points in this figure show the signal fraction per trigger flag, i.e. how many of the triggers are coming from showers that also triggered the doublet partner. The wide trigger flag (flag 2) for the threshold trigger doubles the fraction of “true” shower events, however, it still only contains 3% of the showers and has a signal fraction of less than 10%. For the ToT-types the signal fraction is significantly larger, reaching a maximum of 20% for ToTs. However, in terms of overall fraction of small showers, these ToT-coincidences contribute only 10%. Thus, we have to conclude that while the single SD-stations are sensitive to low-energy showers ($E \approx 10^{15}$ eV) in a radius of about 100 m around the station, we cannot efficiently select them from the T2s. Therefore, a detailed simulation study aiming at identifying small showers on single station basis and a successive implementation in the field might increase the sensitivity of correlation studies as performed here significantly. However, as such an algorithm requires extensive simulation and testing, it is considered beyond the scope of this work.

A remaining possibility is the trigger flag “8” indicating low-gain saturation in the station. We expect that these triggers are to 100% due to showers, but we cannot show this in Fig. 5.50 because this flag is only active for stations with three working PMTs and the doublet stations in our data set are working with two PMTs only¹⁴.

¹⁴This is due to the installation of the so called SSD-PPA [77]. For tests of the SSD without the UUB a PMT of the SD was unplugged and the input used for the SSD.

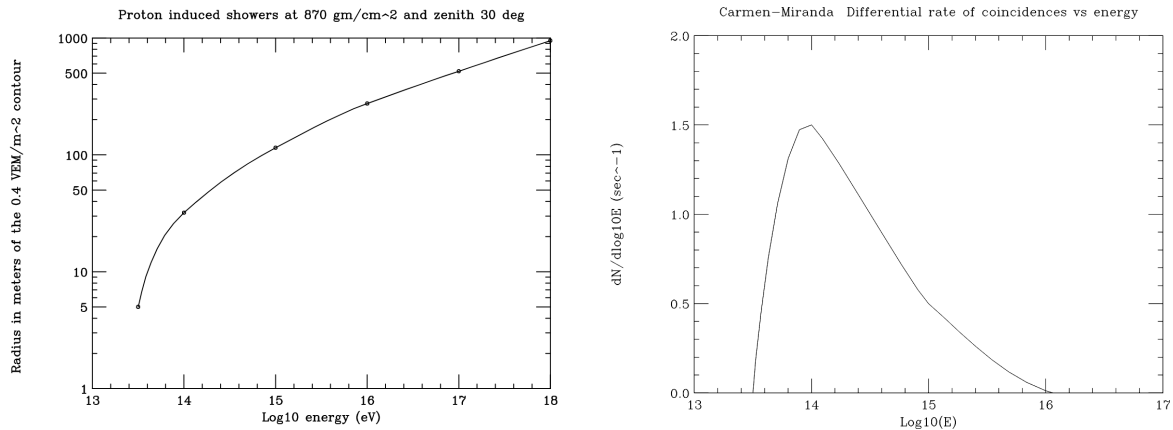


Figure 5.49: Estimation of the rate and spectrum of cosmic rays causing coincidences in doublet stations (two stations about 10 m apart). Figures taken from Ref. [127]. This study used preliminary simulations and simpler trigger logic than what is currently implemented in the field. Given these limitations the agreement with the observed rate, shown in Fig. 5.50 is remarkable. *Left:* Size of the footprint able to trigger the station if a threshold of about 4 VEM is assumed. *Right:* Convolution of the trigger threshold from the footprint size with the flux of cosmic rays in order to obtain the coincidence rate in doublets. The integral of the spectrum gives a rate of 0.9 Hz with an estimated uncertainty of about 50% [127].




Figure 5.50: Dependence of the rate of coincidences between doublets on the trigger type. The trigger flag is generated by the station software version V0R9B6P11 and its logic is illustrated in Fig. 5.34. ToT-types are encoded as flag > 7 and the “wide” signals are the even flags. The data shown here is taken from the 1 June 2020 and averaged over the doublets (72, 688), (79, 819), (80, 669), (89, 734), (90, 651), (91, 643), (93, 710). The expected number of random coincidences is negligible with about 1900 on a signal of 425 000. *Left:* Absolute trigger rate of the different flags in the *master* station (black) and those from doublet coincidences (red). Summed over all triggers the average T2 rate is 21.28 Hz and the coincidence rate of the doublets is 0.70 Hz. *Right:* Number of triggers that happen in coincidence with the *slave* station relative to the total number of triggers of that type (black). In addition, the number of coincidences per trigger flag to the total amount of coincident triggers is shown in red.

5.6.2 Lightning Events

The second important ingredient to search for correlations of cosmic-rays and lightning strikes is the detection of the lightning itself. While the SD itself is very efficient at triggering on lightning activity, as shown in Section 5.1, no analysis of these events beyond the identification of lightning activity is available. For a correlation analysis a good time and positional reconstruction is important to increase the sensitivity. Thus, we use the external lightning

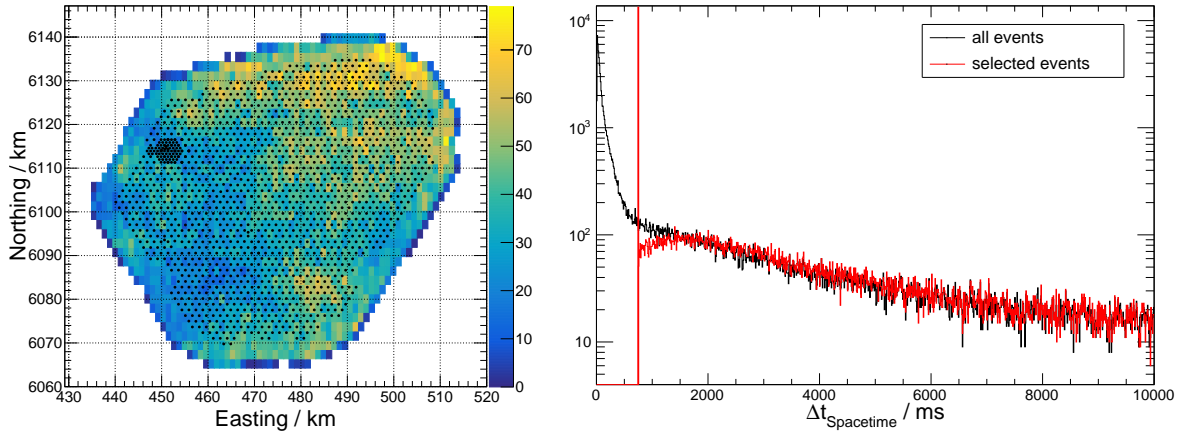


Figure 5.51: Characteristics of the used lightning strikes from the GLD360 [128] network. Selections on the proximity to the SD-array as well as on the time distributions in tuples of strikes are used. *Left:* Positions of the selected lightning strikes within the array. The detectors are shown as black dots. *Right:* Distances of consecutive lightning strikes in space and time with $\Delta t_{\text{Spacetime}}^2 = (\Delta t)^2 + (\Delta s/c)^2$. The selection of $\Delta t_{\text{Spacetime}} < 750$ ms is highlighted in red and serves to select the first lightning strike in lightning tuples.

network data GLD360 [128] that reconstructs lightning strikes using the radio signals that has been used in Auger analyses before [8, 129]. It is available for this analysis as a dataset of lightning locations and times around the Auger Observatory for the time period of 2016 until April 2020. For the analysis presented here, we select lightning strikes within 5 km of the Auger-SD. The resulting distribution of the reconstructed lightning strikes on ground is shown in Fig. 5.51.

The data sample contains many consecutive lightning strikes in the same location. Since we are interested in finding a correlation of the initial strike with cosmic-rays, we have to select the first strikes in tuples. From the black histogram in the right panel of Fig. 5.51, that shows the distribution of $\Delta t_{\text{Spacetime}}^2 = (\Delta t)^2 + (\Delta s/c)^2$ for all data, where Δt is the time difference and Δs is the spatial distance. It is clear that the multiple strikes dominate the sample if they are not removed.

The idea behind the quantity $\Delta t_{\text{Spacetime}}^2 = (\Delta t)^2 + (\Delta s/c)^2$ is that because of the large extent of the Auger-SD, the position of lightning strikes matters in the formation of tuples too. For example, two lightning strikes happening within 10 ms but 60 km apart are unlikely caused by the same initial event. Thus, we include the distance weighted with the speed of light and introduce a cut at $\Delta t_{\text{Spacetime}} < 750$ ms to remove tuples. The value of this cut is motivated by the cross-over point of the two exponential slopes seen in the total spectrum in Fig. 5.51 and leaves a data set of about 64000 lightning strikes.

5.6.3 Characterisation of the Data Set

Before we perform the correlation analysis, we want to make sure that the relative timing of the Auger-SD with respect to the external lightning data, as well as the positioning is correct. Thus, we discuss here the correlation of the SD-ToTs that are sensitive to lightning activity (c.f. Section 5.1) with the lightning data.

The first test we perform is concerning the relative timing of the lightning strikes and the SD-triggers. To be more sensitive to the lightning-induced triggers, we take only ToT-like triggers, i.e. T2-flags > 7 , in the following. We proceed by using the data from November

and December 2016 as “burn”-sample¹⁵ and extract all ToTs in the GPS seconds of selected lightning strikes. When computing the time difference $\Delta t = t_{\text{trigger}} - t_{\text{GLD}}$ for stations close to the reconstructed position ($r < 5$ km), we observe a clear enhancement at about $\Delta t = 0$ in the left panel of Fig. 5.52. The central peak at $(12.4 \pm 0.4) \mu\text{s}$ matches perfectly with the expectation of the latch-bin delay of the triggers. This delay is the difference between the end-of-trace time, reported in the trigger time, and the typical signal time at time bin 250 of 768. The total distribution is likely widened by triggers that trigger before the return stroke on other RF-sources associated with the lightning strike. The second peak around $30 \mu\text{s}$ is caused by the veto-period of the SD-electronics and corresponds to second triggers that can happen after the a initial trigger at $\Delta t \approx 10 \mu\text{s}$ was formed.

To highlight the sensitivity of the SD-ToTs to lightning activity before the final return stroke, we show the same data on a scale of ± 1 s in the right panel of Fig. 5.52, scaled to the rate observed in a randomly selected sample of the same size. The trigger rates are enhanced for several hundred milliseconds around the reconstructed lightning strike. Crucially, this also means that we cannot use these triggers to search for correlations because the influence of RF-noise triggers is too large. However, the “wide”-flag reduces this influence such that the distribution of wide triggers in Fig. 5.52 is flat in relative time to the event.

After ensuring correct timing, we also test the positional accuracy. We extract all SD-ToTs (i.e. trigger flag > 7) in a narrow time window $-500 < \Delta t / \mu\text{s} < 250$ and with $r < 25$ km around the lightning that should only contain the main event, and show the two dimensional difference $\vec{p}_{\text{T2}} - \vec{p}_{\text{GLD}}$ in position in Fig. 5.53. The enhancement of triggers above background in the very centre of the histogram is clearly visible indicating accurate position reconstruction. A minor enhancement to the north-west of the central point is visible too that might also be an effect of event selection slightly outside of the SD-grid in the particular sample chose here.

In the right panel of Fig. 5.53 we show the distance distribution fitted with a Rayleigh distribution to obtain a measure of the positional resolution. The estimate from the fit of $\sigma = 5.1$ km includes the resolution of the lightning network, quoted for lightning strikes in North America as 5-10 km [128], and the extend of the event in the SD. If more than a single station was triggered by the lightning strike, and our experience from Section 5.1 shows that happens regularly, the width of the footprint is influencing σ too. Therefore, we assume in the following when discussing selections for the final correlation analysis that the reconstruction precision is better than 5 km.

Knowing that no shift neither in time nor position of the lightning data is necessary, the final step before attempting the correlation analysis is to show that the SD-T2s are stable in the necessary period. This check should include both the stability in time over the period of analysis from 2017 to 2020 and the stability across different stations. Of special interest for this analysis is the behaviour of the “wide” trigger flags that should be present since end of October 2016. The left panel of Fig. 5.54 shows the rate of stations with the trigger flag 2 using a sample of triggers randomly extracted from 2020 data following the spatial distribution of the input lightning strikes. It is evident that some stations are triggering more often with these special flags than others. To reduce the impact of such outliers, we reject all stations more than 2.5σ away from the mean of the Gaussian fit in further analyses. This chosen threshold roughly coincidences with the cross-over of the exponential tail with the Gaussian peak and removes 353 stations for this period.

In the right panel of Fig. 5.54, we show the relative number of stations sending the flag “2” – taken as proxy for the new trigger flags – sampled on the 15th of each month for the period of 2016-2020. The introduction of the new flags in November 2016 is clearly visible, but also a decay of the number of stations sending these extended flags with time afterwards. We

¹⁵The T2 data for the second half of November 2016 are missing and thus not included here.

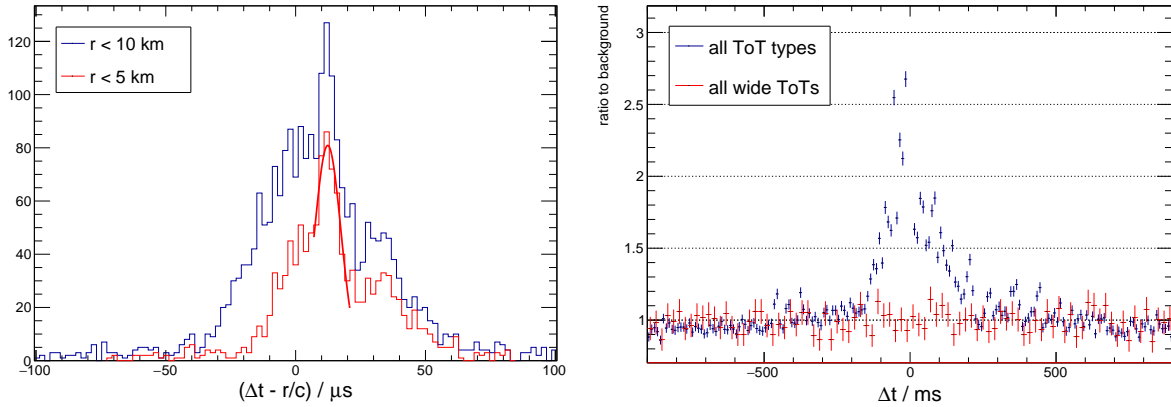


Figure 5.52: Time distributions of ToT triggers (flag > 8) with respect to the reconstructed time of the lightning strike. Data from lightning strikes are November and December 2016 is used in this distribution. *Left:* Distribution of triggers close to the strike itself taking out the expected time of propagation of the signal (at c). A Gaussian fit of the most central peak yields a mean of $(12.4 \pm 0.4) \mu\text{s}$. Given the latch bin of about 250 and that the timing is taken as end of trace, the expected time shift is $12.9 \mu\text{s}$ in perfect agreement with this observation. A second hump is visible about $20 \mu\text{s}$ after the main peak corresponding to second triggers of stations that can only occur after the veto period of the trace from the first trigger. *Right:* Distribution of the ToT triggers on time scales of ± 1 second. The observed counts are divided by the expectation from a uniform background estimated in scrambled data. Without the “wide” flags the influence of the lightning strike and the RF-noise is clearly visible and extends from -500 ms to about 500 ms. In case of the wide triggers, the distribution is compatible with background only.

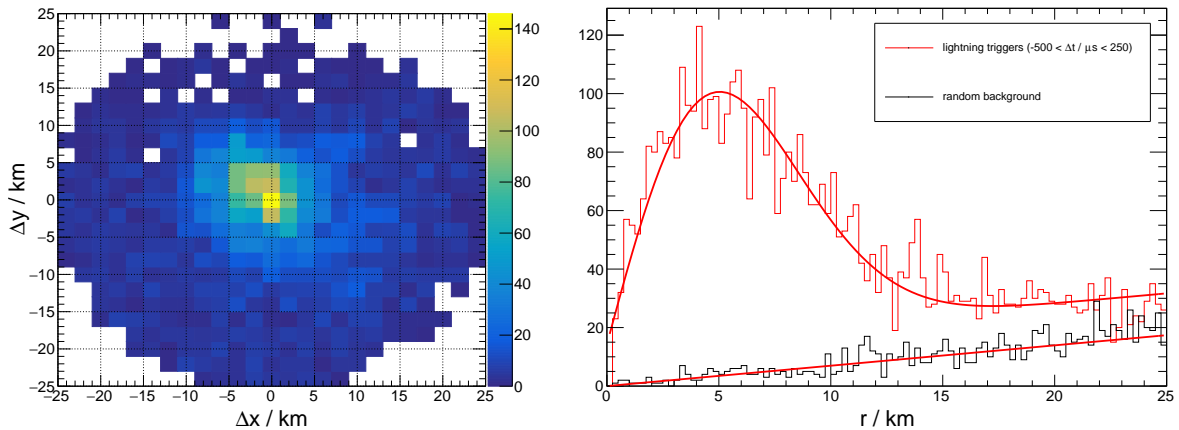


Figure 5.53: Using the response of the SD ToT-triggers to lightning induced RF-noise to test the location accuracy of the lightning strike reconstruction. *Left:* Distribution of the 2D distances of ToT-triggered stations to the reported location of the lightning strike. A time window of $-500 < \Delta t / \mu\text{s} < 250$ is used to select only triggers caused by the return stroke itself. *Right:* Radial distribution of these triggers fitted with a Rayleigh distribution plus background. The obtained width of the distribution is 5.05 km, that is a combination of the resolution of the lightning location and the size of the footprint caused by the lightning. The background is taken from the same data set at times 500 ms away from the lightning strike.

assume it is due to a misconfiguration such that the stations started using the old software again after reboots or similar events. With the introduction of new software to handle SSDs with the old electronics (SSD-PPA [77]) in March 2019 the new trigger flags are running stably in all stations. For the analyses this means that selecting the new trigger flags only implies a reduction of the effective area of the observatory that has to be taken into account in the exposure estimates.




Figure 5.54: Performance of the new trigger flags indicating wide signals in data. *Left:* Identifications of stations with problems in the tagging. The number of times the *wide trigger flag* “2” appears in background data extracted by randomly shuffling lightning-strike data for 2020 data. A clear Gaussian distribution with mean about 600 is visible together with a tail of stations with high counts of wide triggers. The red perpendicular line at 2084 entries is the applied cut for “good” stations at 2.5σ of the Gaussian distribution. *Right:* The fraction of stations sending the new trigger flags as a function of time. For simplicity the 15th of every month is used to sample the data and the most frequent new flag “2” is used to identify a station sending the new flags. The new software was introduced at the end of October 2016, clearly visible in the timeline. The following decline is probably due to stations falling back to an old software version after reboots or similar events. With the introduction of new software for the SSD-PPA [77] the flags are working stably.

5.6.4 Correlation with Station Triggers

Starting with the sample that is in principle the most sensitive to low energies, we present in this section the correlation of single station triggers with lightning strikes. For each of the selected lightning strikes, we collect all station triggers within 15 km and with $-950 \text{ ms} < \Delta t < 25 \text{ ms}$ that are flagged as “wide”. The reason for this selection is that for simple threshold triggers, the background is too large, while for normal ToT types the uncontrollable background from RF-noise is rendering them useless for our purposes.

From other observations, e.g. Ref. [56], it is clear that the time scale from lightning initiation to return stroke δt_{init} can be several hundred milliseconds. This poses a problem because the number of random coincidences with SD-triggers in windows of such sizes is large. However, if there were a characteristic time scale of this process such that $\delta t_{\text{init}} \approx t_{\text{mean}} \pm \sigma$ with σ of the order 50 ms, we could potentially detect this as an excess. Because of the unknown value t_{mean} and σ we will incorporate a scan into the analysis. This scan has to be accounted for when talking about significance.

In more detail, for every trigger type (2, 8, 10, 12, 14) we compare a signal window with the expectation for background in that window. The signal window is defined by the maximal radial distance $r_{\text{max}} \in \{1, 3.54, 5, 7.07, 10, 14.1\}$ km of the trigger to the reconstructed lightning position sampling the parameter space roughly uniformly in r^2 . Overall, the scale σ of the reconstructed position governs the step-size chosen here. For the time delay Δt to the event, we use five different interval lengths ($\{50, 100, 150, 200, 250\}$ ms) and scan them through the data in discrete steps of 10% of the interval length starting at $\Delta t > -500$ ms. We do not want to use larger time differences than 500 ms to avoid the influence from previous lightning strikes that are excluded only to 750 ms, as described previously, and can lead to significant effects on the triggers for several 100 ms afterwards. To exclude residual effects from the return stroke, all triggers in the interval $-150 \mu\text{s} < \Delta t < 150 \mu\text{s}$ are excluded.

During the scan, the interval with the highest signal significance is kept as final result. We use the Poisson-CDF as estimator for the (local) significance, with the background rate estimated from data itself. To reduce the influence of low statistics, we randomly shift the selected lightning strikes by a time taken from a uniform distribution of $t \in [0, 20]$ s with a time scale exceeding that of visible lightning activity in the triggers by more than a factor 20. We repeat this procedure 50 times and use the assumption that a random background sample leads to a uniform distribution of Δt to reduce the uncertainty of the background estimate N_{bg} for arbitrary interval choices by calculating the density $dN/dt(r_{\text{max}})$ over the full Δt sample, i.e.

$$\frac{dN_{\text{bg}}}{d\Delta t}(r_{\text{max}}) = \frac{N(r_{\text{max}})}{\Delta t_{\text{max}} - \Delta t_{\text{min}}}. \quad (5.15)$$

Thus, for a given interval of length T we have

$$p = 1 - \text{CDF}(n_{\text{obs}}, \mu = T \frac{dN_{\text{bg}}}{d\Delta t}) \quad (5.16)$$

as significance that is minimised by the scan.

Applying the scan to the data, we get two highly significant excesses for the trigger types 12 and 14, i.e. the new triggers ToTd and MoPS. From the left panel of Fig. 5.55, it is however clear that these excesses at $\Delta t \approx 0$ are an effect of the RF-induced background that is not filtered out by the “wide” signal selection. In the figure, we show the combined selection 12 or 14 for $r < 5$ km to highlight the similarity of the excess with the distribution of the unfiltered triggers in Fig. 5.52.

For the triggers not affected by this background, i.e. threshold triggers (flag 2) and low-gain saturated signals (flag 8), we do not find significant signals. We get the smallest local p-value of 0.002 for flag 8 and $r < 7.07$ km and $-220 \text{ ms} < \Delta t < -120 \text{ ms}$ with $n_{\text{obs}} = 4$ at expected 0.9. Given the amount of scanning involved, this is clearly no significant excess and we do not try to estimate the significance rigorously. The order of magnitude of scanned intervals can be found as $n_r n_t n_{\text{scan}} \approx 6 \times 5 \times 40 = 1200$, when we take the number of scanned intervals n_{scan} for the intervals of length 100 ms. These intervals are not entirely statistically independent because the data selected by the different intervals have common subsets. However, given that the amount of scans is about 10^3 , p-values of order 10^{-3} are clearly not significant.

These results are not surprising given the low efficiency achieved with the wide flags. To illustrate the lack of expected sensitivity, we use the remaining flag with the highest efficiency: trigger flag 2 with about 3% efficiency in doublets. Integrating over $-500 \text{ ms} < \Delta t < -50 \text{ ms}$ at $r < 7.07$ km – the value found for the highest local significance in the scan – we obtain a 90% C.L. Feldman-Cousin limit [130] of 607 signal counts. Using

$$n(\text{CR}|\text{lightning}) \approx \frac{\mu_{\text{max}}}{\varepsilon_{\text{flag}} \varepsilon_A} \approx \frac{\mu_{\text{max}}}{0.03 (r/r_{\text{hex}})^2} = 5 \times 10^5, \quad (5.17)$$

with the radius $r = 150$ m for detecting showers of energy 10^{15} eV to 10^{16} eV (from Fig. 5.49), and the distance between stations in the hexagon $r_{\text{hex}} = 750$ m, to estimate the geometric efficiency ε_A . This number is one order of magnitude larger than the size of the event sample, thus no constraint can be derived from this data.

For the low-gain saturated triggers the argumentation is very similar because the efficiency is very small (but unknown). However, due to the low background (0.9 per 100 ms at $r < 7.07$ km) it might have been possible to observe a signal if the signal came from higher energies than 10^{15} eV which is why we performed this search. Due to the lack of knowledge of the efficiency, we do not perform the same calculation of the limit here and leave it for future work together with optimised trigger flags that take the problems found here into account.




Figure 5.55: Visualisation of the distributions used to search for excesses. The time difference of all triggers with the selected type in a radial window are selected and counted. To prevent large biases from the lightning strike itself, a window of $150 \mu\text{s}$ width around the strike itself is excluded from this analysis and thus not included in the last bin. *Left:* Comparison of the observed distribution of Δt in the signal sample with the expected background distribution for trigger flags 12 and 14 (i.e. wide ToTd and MoPS) for $r < 5 \text{ km}$. A clear excess at $\Delta t \approx 0$ is visible caused by residual influence of the RF-noise on the triggers. *Right:* Distribution of the wide threshold triggers with $r < 7.07 \text{ km}$ with the maximal excess window found by the scan marked with red lines. At a local p-value of 0.004 ($n = 38351$, $\mu = 37841.1$) it is clearly not significant, especially accounting for the scan.

5.6.5 Correlation with Doublets

Because the analysis with single stations is not constraining due to a low efficiency in selecting low-energy showers with the trigger flags alone, we discuss in this section the extension of the analysis to doublets. The idea is that with doublet stations we can efficiently select the small showers with threshold triggers that are not affected by RF-noise. From Fig. 5.50 we know that 85% of the low-energy showers can be selected this way. Most of the remaining 15% are not accessible even with doublets, because also doublets are affected by RF-noise since it will likewise appear in both stations of the doublet. To highlight this, we show the Δt distribution of ToT-like triggers in the right panel of Fig. 5.56 that shows the same excess at the time of the lightning strike as the single stations.

Obviously, when we restrict the analysis to doublets we lose the aperture of the full SD. The number of lightning strikes in the “sensitive area” of 7 km around the doublets of 852 compared to the full 64000 for the full array illustrates this loss in area well. However, due to the increased sensitivity in the measurement with the selection efficiency increasing by almost two orders of magnitude, it is still more likely to observe a correlation in this measurement.

We perform the same scan as for single stations to find the best possible window for correlation restricting the types to threshold triggers. We obtain the smallest local p-value of 1.2×10^{-3} for $r < 1 \text{ km}$ and $-365 \text{ ms} < \Delta t < -215 \text{ ms}$ at $n_{\text{obs}} = 10$ and $\mu = 3.58$. For the same reasons as in the previous section, this is not a significant excess.

To test whether or not we can constrain the fraction of lightning strikes initiated within 500 ms by a cosmic ray, we get the Feldman-Cousin limit for $r < 7.07 \text{ km}$ and $-500 \text{ ms} < \Delta t < -50 \text{ ms}$: $\mu_{\text{max}} = 85.6$ at 90% C.L. Using

$$n(\text{CR}|\text{lightning}) \leq \frac{\mu_{\text{max}}}{\epsilon_1 \epsilon_A} = 2526 \gg n_{\text{lightning}} = 852 \quad (5.18)$$

we see that also in this sample the background is too large such that we cannot put any constraint on the fraction of cosmic-rays initiating lightning. However, in this sample the

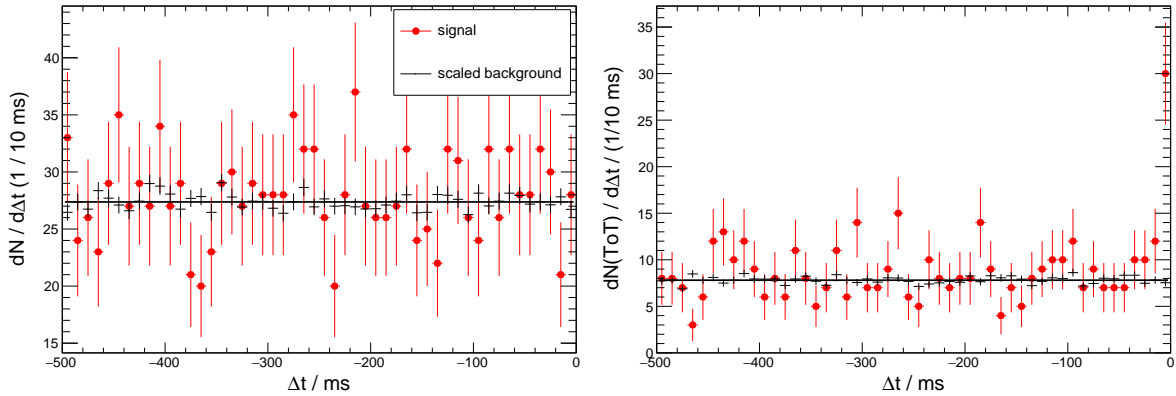


Figure 5.56: Correlation of triggers of doublets with lightning strikes. Only T2s that were coincident in both stations of the doublet are used here to select small showers with high efficiency. *Left:* Distribution of coincidences where the *master* station ($ID > 100$) had a threshold trigger and the lightning strike was within 7.07 km. No significant enhance of the rate is detected above the random background. *Right:* Using all ToT-flags (i.e. trigger flag > 7) for coincidences between doublet stations the influence of lightning noise is visible. Similarly to the analysis with single stations, we do not use these data to be conservative.

limit is only a factor three larger than the number of events. Thus, it seems possible that with accurate lightning detection, including the time scale before the actual return stroke like done by LOFAR [56], meaningful limits could be derived in the future.

5.6.6 Correlation with Trigger Pairs

The final part of the search for a correlation of cosmic rays is aiming at higher energies and tries to use coincident pairs of nearest neighbours (nn-pairs) in the SD-grid to increase the sensitivity by suppressing the background. However, also for pairs of triggers the RF-signals in the stations are the most important background. Thus, we first have to show that we can select subsamples of data that have an understood background.

The left panel of Fig. 5.57 shows the influence of RF-noise on pairs of triggers, by selecting pairs of ToTs (including MoPS and ToTd). Clearly, these triggers are not usable for the correlation analysis because of this uncontrolled background even though they constitute the majority of true cosmic-ray events (c.f. Fig. 5.47). Already in Section 5.6.1, we discussed that it might be necessary to select samples with “wide” triggers to reduce the influence of RF-noise. With such selections, the efficiency drops by about a factor three and as it is visible from the peak at $\Delta t \approx 0$ in the right panel of Fig. 5.57 there is still RF-noise left in the sample.

Using threshold triggers, as for the single station analysis, is not promising because the efficiency is much less than 1% for pairs of triggers. Therefore, the only option left is to ask for two “wide” triggers. The efficiency for such a selection is about 1% centred at energies of $\lg(E/\text{eV}) = 17.5$, following the same procedure as in Fig. 5.47. Thus, we have to resort to a similar argument as for the use of the low-gain saturated triggers in the single station search. The efficiency with such a selection is too low to constrain the fraction $n(\text{CR}|\text{lightning})$, however, if the background is low enough a significant signal can still be detected. For the infilled SD-750 the efficiency is even smaller than 1% and so we do not use this data set in this analysis.

We approach this analysis analogously to the previous analyses by selecting all nearest-neighbour clusters following Eq. (5.9) in a time window of $-500 \text{ ms} < \Delta t < -150 \mu\text{s}$ and within 15 km around the lightning strikes. Here, we define $\Delta t = t_{\text{first}} - t_{\text{lightning}}$ as the time




Figure 5.57: Correlation of pairs of nearest neighbours with lightning strikes using coincident pairs of nearest neighbour stations within 7.07 km of the reconstructed position of the lightning strike. *Left:* Selecting pairs of triggers with two ToT-like triggers. A clear and expected excess of triggers due to lightning-induced RF-noise before and especially at the lightning strike is visible. *Right:* Selection of pairs with at least one “wide” trigger flag. Contamination from RF-noise, visible as an excess at small times, is present in this selection too.

difference of the closest trigger to the lightning strike. Similarly, this closest station is used to determine a distance of the cluster to the event.

From this set of nn-clusters, we then sub-select only pairs, i.e. clusters with exactly two triggers, that both have a “wide” trigger flag that results in a sample of 30 nn-pairs only. We estimate the random background by using the same lightning events shifted randomly in time and repeat the analysis, analogous to the treatment in the previous sections. To not be guided by eye in low-statistics conditions, as evident from the distribution of distances r and time differences in the left panel of Fig. 5.58, we apply the scanning scheme presented in the previous sections also on these data. In contrast to the other analysis, the minimum of $p = 9 \times 10^{-5}$ for $r < 7.07$ km and $-290 \text{ ms} < \Delta t < -240 \text{ ms}$ for $n_{\text{obs}} = 5$ at $\mu = 0.7$ is not easily interpretable as insignificant. We visualise this excess in the right panel of Fig. 5.58 but do not change any prescription of the scanning.

However, to get a proper significance, we use Monte-Carlo method that takes the scans into account. We generate pseudo-data sets where r^2 and Δt are uniformly distributed and the number of events is generated by drawing a Poissonian number from the mean number observed in the background estimates. On each of these pseudo-data sets we perform the same scanning procedure and compare the minimal obtained p-value with the minimum seen in the signal sample. In 1000 MC-data sets, we observe 50 local p-values smaller than in data, translating to a p-value of 0.05 or 1.6σ revealing it as insignificant.

Nonetheless, the enhancement of such pairs about 250 ms before a lightning strike is interesting. Because of the selection of only two nearest neighbours without further triggered stations, no SD-events are available for most of the “candidate” events. Keeping the signatures of SD-disks in mind, the interpretation of these events as long-signal candidates rather than cosmic rays initiating the lightning is tempting. With the updated software, it might be possible in the future to distinguish such signatures from each other if enough data in this configuration can be collected. Another possibility is to study the feasibility of promoting pairs of “wide” triggers to a T3 such that the traces are available for detailed analysis. While with the number of candidates found in this analysis, this seems certainly possible, we did not estimate the background associated with such a selection in normal data taking conditions. Furthermore, we think that with the UUB more processing power is available so that any such selections can be done with much more rigour.

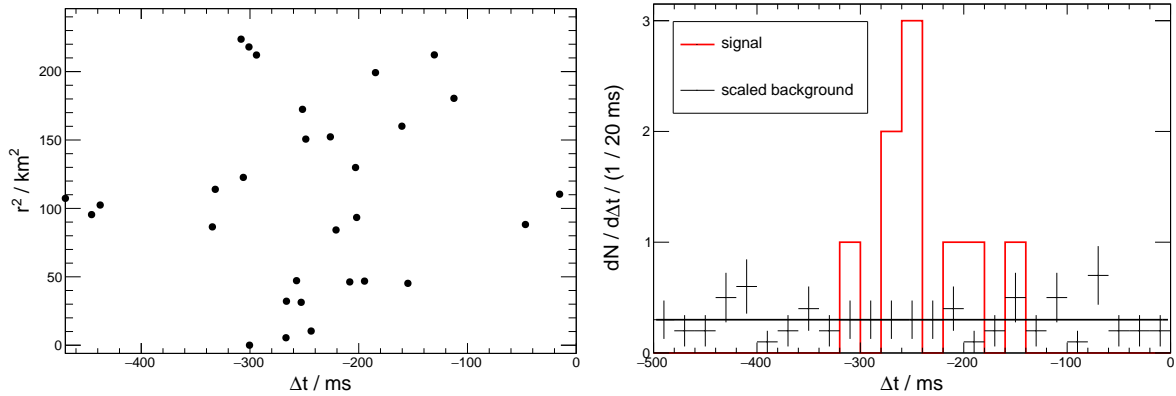


Figure 5.58: Correlation of nearest-neighbour pairs of “wide” triggers with lightning strikes. *Left:* The complete selected sample of 30 pairs. *Right:* Selected signal region with $r < 7.07$ km including the background estimated from randomly shifted events. Note that the binned values of the background are not used in the analysis but rather the fit that is less affected by low statistics.

5.7 Summary

In this chapter, we discussed different effects and events visible in the data of the surface detector of the Pierre Auger Observatory that are related to lightning activity.

While the physical origin of the peculiar SD-rings is still not known for sure, we show that the ring structure is very likely an artefact of the SD-trigger. This means, that rather than having two classes of peculiar events that share the characteristic long-signals but are different in footprint, we only have SD-disks that are sometimes not recorded as such. Using similarity arguments of such completely filled events with other observations [9, 59] it seems plausible that we observe gamma radiation during thunderstorms with the SD in these events. A definite proof, however, is still missing but we are sure that with the possibilities of the upgraded DAQ and electronics, and a larger set of events recorded in the future this question can be answered. In this context, we have shown and successfully tested an updated software for the current electronics of the Auger-SD increasing the sensitivity of the read-out that has just been deployed to the full array at the time of finalising this thesis.

For the new electronics that are currently being deployed in the field at the time of writing of this thesis, we show that with dedicated modifications in the soft- and firmware an increased sensitivity can be easily reached. The details of how exactly these improvements are to be made depend on the final version of the DAQ software and constraints imposed by the bandwidth and other physics goals, e.g. new triggers for improved neutrino detection. We thus only give first ideas to help facilitate an informed discussion within the Collaboration and refer for the details to future, more rigorous work.

Using the SD-detectors to observe cosmic rays correlated with lightning strikes proves to be difficult. Due to the strong influence of the lightning activity as high-frequency signals in the PMT-cables of the SD-stations, strict cuts are necessary that reduce the efficiency significantly. With this reduced sensitivity, the task of correlating single cosmic rays of energies of about $10^{15} - 10^{17}$ eV with lightning strikes becomes very difficult. Especially the lack of knowledge about the time between the supposedly initiating cosmic ray and return stroke increases the random background to unsatisfactory levels for this analysis. Thus, with the current data it is not possible to constrain the influence of cosmic-rays on the lightning problem.

Nevertheless, with the work presented here, it is clear that for future efforts with the upgraded electronics improvements on several sides can make such measurements sensitive enough. On the cosmic-ray detection side, it is imperative to be able to efficiently select the

signals from small showers but also to reduce the influence of any lightning RF-induced signals to obtain a clean data sample. On the lightning detection side, detailed instrumentation adding information about the time scale involved can narrow down the search window far enough to make detection or rejection of the hypothesis much more likely.

While we do not observe any significant signal due to the necessary scanning in time scales and radii, we think the observation using pairs of triggers in nearest-neighbour stations should be revisited with more data in the near future to confirm or reject the hint of a signal in the SD. Note however, that even in the case of an observed correlation, it has to be proved that cosmic rays rather than exotic events are observed. Also here, the tagging algorithm developed to optimise the readout of SD disks will facilitate future searches. This will contribute to solving the question of whether cosmic-rays are involved in the initiation of lightning or not.

Spatial and Temporal Correlation of Surface Detector Triggers

In this Chapter, we want to explore the possibilities of using the T2 data of the Auger SD to go beyond the usually recorded events in space and time. For this we explore the correlations of triggers among themselves or with external events, e.g. reconstructed UHECR showers.

Among the peculiarities we want to search for in this way are the Gerasimova-Zatsepin (GZ) events [5] that would appear as a pair of separated, simultaneous showers. Because one or both showers of such a GZ event could be below the reconstruction threshold, we can extend the phase space available for these searches with this analysis significantly. We know that such extensions of the energy range are clearly possible since we already discussed the sensitivity of the T2s in Section 5.6.1.

We can also use the triggers to search for delayed signals, e.g. of slow, long-lived particles, if we look at times that extend beyond the time scale of the recorded SD traces. While all relevant particles for the air-shower reconstruction are recorded in these traces, it is definitely possible that heavy, (meta-)stable particles, like for example neutrons, are creating signals well beyond the trace time window. In the first part of this Chapter we will focus on such searches using time correlations.

To set the scene, we show that time correlations of the triggers themselves are a useful tool to characterise the hardware. Using the newly available microsecond data we analyse whether the triggers are stable also on a microsecond scale. Exploring such possibilities also helps setting the path for a successful commissioning of the new electronics for the AugerPrime upgrade.

6.1 Characterisation of the Data Set

In this Chapter, we use the same “T2Dumps” data as presented previously in Section 5.2.1. Instead of using the trigger data as a tool to understand effects of the DAQ as in Section 5.1, we use the trigger data here as independent data set that contains information about the background flux of low-energy cosmic rays, sub-threshold events, and the ultra-high energy events themselves, as well as their possible correlations. We start with a brief discussion of the properties of the data focusing on time distributions, keeping in mind that correlated triggers – which we are searching for – should be visible in a time-difference spectra if the random background is taken into account. Some results of this section were previously collaboration-internally published as Ref. [100] and are summarised and extended here.

6.1.1 Time-Difference Spectra

If we want to analyse the trigger times and their statistical behaviour, the first quantity beyond the average rate that can give insights into the stability of the triggers is the time difference between consecutive triggers Δt . The spectra of Δt in pure random background are expected to be exponentially distributed with a scale λ corresponding to the average trigger rate. Already in Section 5.2.2 we have seen that in these spectra the overall-small effect (200 μs while $\langle \Delta t \rangle \approx 50 \text{ ms}$) of dead time was imprinted at small times. Here, we will discuss the shape of the distribution for Δt beyond the dead-time effect and will use henceforth the data of 14 to 15 June 2019, a quiet period not affected by thunderstorms, unlike the sample used in Section 5.2.2. To separate effects of the different triggers, we analyse the two classes of triggers, threshold and ToT, separately. The main reason being that any effect of the ToTs would otherwise be hidden by the vastly greater amount of threshold T2s.

This means that we have two distinct sets of time differences Δt : those between two threshold or ToT triggers of a station, respectively. In Fig. 6.1 the spectrum for the threshold triggers in the June 2019 data set is shown, calculating Δt for each station and stacking the distributions of all stations. It is clearly visible that the expected exponential behaviour is present. The inset shows a zoom into the first 3 ms of the very same histogram. On the time scale of about 300 μs deviations are visible which we will analyse in more detail.

The first feature is the 20 μs veto period visible as absence of any $\Delta t < 20 \mu\text{s}$ which is implemented to prevent triggering more than once on the same trace data [131]. The second feature, a deficit of triggers until $\Delta t \approx 250 \mu\text{s}$, can be explained as the effect of the station dead-time after more than two triggers, discussed previously in Section 5.2.2. The dead-time creates only a small suppression rather than a complete absence of triggers because of its higher-order nature, i.e. $\Delta t = t_i - t_{i-1} = 200 \mu\text{s}$ is possible for two triggers. Only if $\Delta t = t_i - t_{i-2} < 250 \mu\text{s}$ the dead-time affects the spectrum.

Lastly, there is a peak of triggers at about 235 μs . This peak arises from a handful of stations, as can be derived visually from Fig. 6.2. It shows the time spectra Δt for each station separately to highlight such outliers. We discuss an algorithmic approach to identify these stations in the next Section.

Moving to the ToT-triggers, Fig. 6.3 shows the raw spectrum of time differences for these triggers only. Also here the veto period is visible, as well as the overall good agreement with the exponential distribution. Similar to the threshold trigger case, there are deviations from the expected behaviour on small time scales. These last to about 600 μs , and are in part due to singular stations as in the case of the threshold triggers. However, as we will see after the removal of “unusual” stations, some of these unexpected effects persist. The possible origins of these correlations of triggers will be discussed after the removal of the unusual stations.

6.1.2 Identification of Unusual Stations

Following the results from the “raw” Δt -spectra, it is necessary to remove outlier stations to obtain a cleaned sample for further analyses. Thus, we discuss the identification of stations that have unusual trigger-time distributions, as those seen in Fig. 6.2 here.

To not impose a certain type of “abnormal” behaviour, we use a comparison with a default model of what the Δt -distribution should be. This baseline model consists of an exponential distribution which is cut-off at the veto time of rounded 20 μs (trace length), and additionally, has a suppressed contribution until 250 μs , based on the expected influence of the dead-time.

The time scale of the exponential distribution of the model is determined using a fit of an exponential function to the range beyond 300 μs . We chose this range to avoid a bias caused by the dead-time suppressed region for the parameter estimation.




Figure 6.1: The raw time-difference spectrum between consecutive threshold triggers summed over all stations. Overall, the distribution shows an excellent agreement with the expected exponential distribution. *Inset:* Zoom into the first 3 ms of the distribution in a linear scale.




Figure 6.2: Distribution of the time differences between threshold T2s separated by station ID for time differences < 1 ms. For a few stations, strong deviation at about $235 \mu\text{s}$ are visible.

With this model, we perform a Kolmogorov-Smirnov test¹ with the observed trigger time distribution. We reject any station with $p_{\text{KS}} < 0.01$ for threshold T2s, while for ToT triggers $p_{\text{KS}} < 0.05$ is used. The different choice is based on the average difference in statistics available to account for the changed statistical power of the respective data set.

Examples of removed stations are given in Fig. 6.4. In the left panel, station 170 shows the “typical” peak at $235 \mu\text{s}$ which was mentioned already in the previous section. In the right panel an example of a station which is rejected due to unusual ToT distributions is shown. In this case several wide peaks are visible. The method presented here removes 40 out of 1613 stations that sent T2s in the considered period. The station Ids and the origin of the removal are given in Table 6.4.

A disadvantage of the KS-test is its lack of sensitivity if the statistics is low. To alleviate this we add, if necessary, a simple peak-rejection algorithm based on the exponential fits and

¹ For practical reasons the actual test is performed with binned data, via the ROOT [132, 133] TH1::KolmogorovTest method.




Figure 6.3: The raw time-difference spectrum between consecutive ToT-triggers summed over all stations. Overall, the distribution shows an excellent agreement with the expected exponential distribution. *Inset:* Zoom into the first 1.5 ms of the distribution in a linear scale.




Figure 6.4: *Left:* Trigger-time differences of threshold triggers for station 170 in the June 2019 data set. The station is flagged as showing unusual trigger-time distributions. The red line represents the exponential fit to the tail of the distribution to build the expected model. *Right:* Same distribution but of ToT triggers for station 1453.

the number of bins showing deviations of more than 3 and 5σ . We calculate the deviation using the p -value of a Poisson distribution, i.e.

$$p = 1 - \text{CDF}(n|\mu) \quad (6.1)$$

for an expected number of counts μ . This allows us to process the data in blocks of one day rather than multiple days. As thresholds we use $n_{3\sigma} \leq 10$ and $n_{5\sigma} \leq 1$.

In some parts of this Chapter, we make use of stations sending the extended trigger flags which allow for the identification of new triggers [114, 115] or traces that were low-gain saturated. In analyses using these data we use only stations which send the extended trigger flags. The right panel of Fig. 5.54 shows the number of stations that fulfil this condition as a function of time, highlighting that in some periods of data about half of the statistics is lost. We want to stress that the stations not sending the extended flags are not broken, but rather seem to run an older version of the local-station software. We reject them in parts of the analyses for reasons of clarity.

6.2 Temporal Correlations

In the following, we want to explore the possibilities of investigating the “clean” data with respect to possible correlations of triggers in time. With this term we refer to effects where, unlike in pure Poissonian processes, the observation of a *next* trigger may depend on previous events. For the cosmic-ray events that the Auger-SD is designed to measure, no such correlations are expected because the processes of acceleration and the propagation of cosmic rays randomise the arrival times leading to no correlations between events in time. However, local effects, for example slow and long-lived particles like neutrons produced in the hadronic cascade, could lead to temporal correlations of triggers solely because the event is longer than the trace. We can also use the search for temporal correlations as a general test of the assumption that cosmic rays are arriving randomly and independently.

In the following, we first discuss the general observations of trigger-time distributions followed by more specific searches. In these searches we use information from the reconstructed cosmic-ray events combined with trigger data to search for delayed effects occurring after events.

6.2.1 Time Difference Spectra

In this Section, we repeat the analysis of the time difference spectra presented in Section 6.1.1 after removing the stations identified in Table 6.4. Fig. 6.5 shows the result for the threshold and ToT triggers separately. For the threshold triggers this rejection of a few stations completely removes the unidentified peak in the stacked histogram, described in the previous section, leaving only features of the spectrum that are understood. Whether the rejected stations also show undesired behaviour in trace analysis is subject to future investigations and beyond the scope of this work.

For the ToT-triggers the situation is more complex. As visible from the right panel of Fig. 6.5 the cleaned spectrum still shows a fairly irregular behaviour at small times.

To further investigate this behaviour, we show in Fig. 6.6 the distribution of the time differences for the different types of ToTs separately. Those histograms only focus on the observed deviations from the expected behaviour on short time scales and therefore only show the first 1.5 ms of the distributions. It is visible from the three panels that the biggest irregularities come from the ToTd triggers. The peak in the overall distribution right after the veto period is almost exclusively due to these ToTd triggers. Similarly, the recovery feature is almost exclusively a MoPS feature.

For the MoPS triggers the feature is likely an effect of the trigger implementation of the integral condition. With the use of a decaying sum to compute the baseline comes a time scale of 50 μs [134] that influences later triggers. This time scale is matched by the recovery observed in the data.

The origin of the structure for the ToTd spectra remains unknown. While also in the ToTd a condition on the integral signal is implemented that uses the same decaying sum as the MoPS [134], no recovery is evident in the spectrum. Due to the priority in the trigger flags (c.f. Fig. 5.34) these features appear in the small subset of data when two subsequent traces are exclusively ToTd-triggered. In other combinations of triggers, e.g. a ToTd after ToT-triggered traces, these features are not visible. While there is the possibility that these triggers imply long-lasting physical signals in the SDs, due to the fact that these features are visible only exclusively in ToTd, leads us to think that peculiarities of the implementation might cause these features. Note that the double-peaked nature of this distribution is likely linked to the dead-time after two triggers.

We want to stress, that these irregularities only cover a very small fraction of the time and the overall number of triggers. The vast majority of triggers follow the expected distributions.




Figure 6.5: *Left:* Distribution of the time differences between threshold triggers after removing the stations from Table 6.4 and summing over the remaining stations. We show only the range until 2.5 ms. *Right:* Distribution of the time differences between ToT triggers after removing the stations from Table 6.4 and summing over the remaining stations. We show only the range until 3 ms.




Figure 6.6: *Left:* Time differences between “old” ToTs only: the spectrum is very regular. *Middle:* Time differences using only ToT triggers. The combination of ToT and MoPS is not included. A clear peak after the veto period of 20 μs is visible. *Right:* The same distribution for MoPS triggers only. Also here triggers that have both MoPS and ToT are not used.

Other combinations of trigger types, for example a ToT following a threshold trigger, are not investigated here in detail as those would lead to a large number of possible combinations without adding more information. We believe that the most significant features are already explained with the simple combinations presented here.

6.2.2 Expectations for Delayed Triggers

From using the station triggers alone to search for correlations, we here briefly summarise why we may expect delayed particles and thus triggers after large events based on previous observations and simulations.

Recently [135], delayed pulses in the traces from the SSD pre-production array have been observed. Such delayed pulses are similar to observations of subluminal pulses in scintillation detectors by Linsley [136]. Such delayed particles were observed also with other detector systems, mostly with neutron counters [137]. This so-called “neutron thunders” can arise from neutrons created in the soil and in detectors nearby the hadronic core of an air shower. It leads to signals that are up to 1 ms delayed with respect to the initial air shower. Since the size of the hadronic core, and thus the associated neutron signal, are correlated, the measurement of such delayed particles might be interesting to resolve open questions in composition and hadronic interactions. In the following Sections, we want to discuss the possibilities of observing such delayed neutrons with the Auger SD and also present a search for delayed signals in the current data.

From measurements of the neutron-capture efficiency by water in SNO+ [138] we know that the WCD of the SD can measure neutrons when their number is high enough. To estimate in a back-of-the-envelope manner the necessary number of neutrons to trigger, we use the fact that the neutron capture on the hydrogen of water will release a 2.2 MeV gamma ray. This gamma ray can, after subsequent interactions, in the water lead to Cherenkov light with about 2 MeV of deposited energy. We can estimate the number of neutrons required for a ToT trigger by comparing this deposit to the value of the VEM, which corresponds to roughly 250 MeV of deposited energy. The ToT trigger requires at least 13 bins above a threshold of 0.2 VEM. If we assume a “perfect” time distribution of the captures for the trigger we get

$$N_{\text{capt}} \gtrsim 13 \times \frac{50 \text{ MeV}}{(2 \text{ MeV/capture})} = 325 \text{ captures.} \quad (6.2)$$

We can compare this number with the number of neutrons created at the core of a shower with 40 TeV of hadronic energy that is estimated to be about 7000 [139]. Thus, it seems possible to measure these neutrons with the SD if the overall detection efficiency for capture and photon conversion is better than 1%. Given the estimated efficiency of up to 50% for SNO+, this efficiency seems well within reach.

More recently, it was shown [140] that the distribution of neutrons created in an air shower can be predicted using FLUKA [141]. From Fig. 6.7 it is clear that contrary to muons there are significant parts of the neutrons arriving later than 10 μs . While the simulations are performed at energies lower than those relevant for Auger, from the scaling with shower depth in Fig. 6.8 and the change of X_{max} for higher energies it is clear that high numbers of neutrons are justifiably expected.

Motivated by the possibility to see such neutrons we present a search for delayed pulses in the WCDs that exceed the time scale of the trace length (i.e. $>19.2 \mu\text{s}$) by using triggers. In a separate search, we will target the short time scale $\lesssim 10 \mu\text{s}$ with the traces.

For the delayed triggers we use the available trigger data to search for short time scale correlation of station triggers. A similar correlation analysis can be performed by using the reconstructed SD events and searching for additional triggers of stations in the event. The details of these analyses, their results as well as the systematic effects hindering the searches are discussed in the following Sections.

6.2.3 Statistical Search for Delayed Signals

The first approach in search for delayed signals on the time scales that exceed single-trace length is to use the collected T2 triggers of all stations. Correlations of triggers with delayed particles should appear as tuples of triggers of a single station, correlated in time. A sub-selection of triggers following the large-signal triggers can be used to target the analysis more towards particles hitting a station near the cores of small showers. In the following, we describe how we search for such excesses and give our observations resulting from the data.

To search for time correlations of triggers we want to define and count tuples of triggers seen in each station. In this analysis, we use the extended trigger flags and subsequently reject stations which do not send them because we want to be able to select small showers with the newly introduced flags.

For each station, we define a “trigger tuple” as all the triggers appearing within 2 ms of the first one. Fig. 6.9 illustrates this approach for a single station. To avoid extensive computational costs, we adopt an iterative approach that adds a trigger to a tuple if it occurs within 2 ms of the first one. If the time limit is exceeded with this added trigger, the existing tuple is completed and stored if containing at least two triggers, while the current trigger (t_1) is placed in a new tuple.




Figure 6.7: Energy spectra of muons (*left*) and neutrons (*right*) in FLUKA [141] simulations for proton-initiated air showers. The different colours represent different arrival times at the ground. Neutrons arriving with more than 10 μ s delay are clearly present. Taken from Ref. [140].

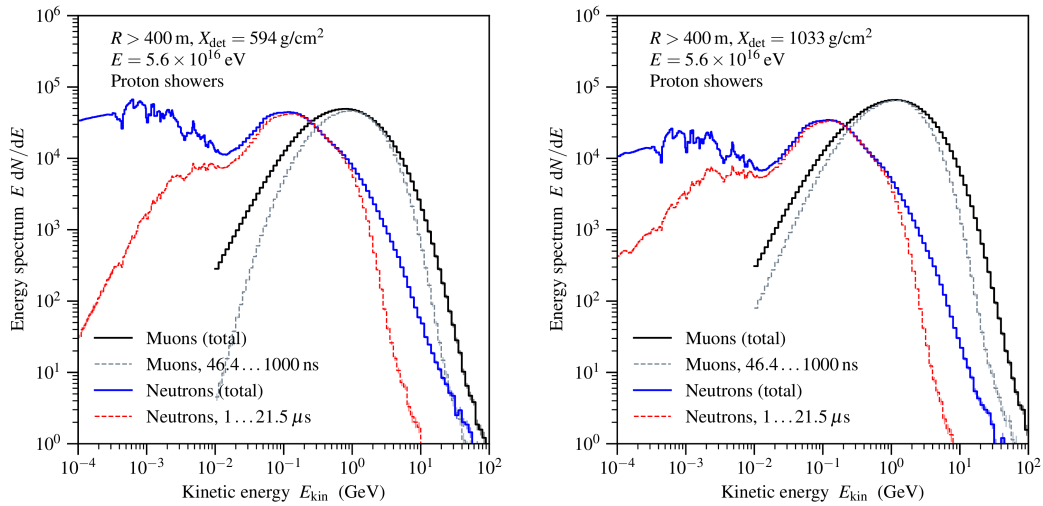


Figure 6.8: Neutron and muon spectra simulated with FLUKA for two observation depths X_{det} at large distances R from the core. The main time scales contributing to the spectra are highlighted in colour. Taken from Ref. [140].

For simplicity and to save computation time, we use only the data of 15 June 2019. It provides sufficient statistics with e.g. 1734 extended triggers marked as “low-gain saturated”.

The goal is then to find excess of occurrences (over Poissonian background) of triggers after certain initial trigger. Selections on the type of the initial trigger or on multiplicity of triggers in a tuple can enhance potential excesses by significantly suppressing the background.

We can accurately describe the background as the expected distributions of trigger-tuple multiplicities assuming a pure Poissonian background, as it was previously used for the detection of dead-time in Eq. (5.1). To recap, for a Poissonian process the probability of a waiting time t between a given event and the next n th event follows

$$p(t; n) = \frac{\lambda^n t^{n-1}}{(n-1)!} \exp(-\lambda t), \quad (6.3)$$




Figure 6.9: Illustration of the definition of the trigger tuple. The arrow shows the time axis, where the t_i mark the times when a trigger in a station occurs. Time differences Δt_{j1} to the first (after a previous tuple) are highlighted with additional labels. For computational simplicity we do not optimise the time window but adopt an iterative approach instead: we add triggers into the tuple until Δt_{j1} exceeds 2 ms. After that time a new tuple is started.

where λ is the mean rate of the events. We assume that for every trigger in consideration the next n th trigger exists, which is a very good approximation given that the whole data set contains billions of triggers. The relative amount of tuples with a given multiplicity within a maximal time difference t_{mas} can be thus calculated with

$$r(n) = \int_0^{t_{\text{max}}} p(t; n) dt = \int_0^{t_{\text{max}}} \frac{\lambda^n (\Delta t)^{n-1}}{(n-1)!} \exp(-\lambda \Delta t) dt. \quad (6.4)$$

The resulting expression can be transformed into

$$r(n) = 1 - Q(n-1, \lambda t_{\text{max}}) \quad (6.5)$$

with the regularized gamma function Q .

If multiple stations with slightly different mean trigger rates are stacked, λ becomes an effective constant. It can deviate from the average of the mean rates if the sample is dominated by stations with higher mean trigger rate. Similarly, if sub-samples of the tuples are considered, all formulae become approximations since no separate analysis on the corresponding subset is performed. For example, ToT triggers could be treated as a separate data stream where all threshold T2s are ignored in the tuple formation.

Following the algorithm described above, from the 24 hours of data during 15 June 2019 we obtain 1.04×10^8 trigger tuples. Out of those tuples, 1.8×10^7 are doublets and another 2.1×10^6 are triplets. Fig. 6.10 summarizes the characteristics of this data set. The left panel shows the different trigger types contributing to the tuples and their overall abundance in the data set. It is clear that for most trigger types the ratio of appearing in data and in tuples is roughly constant. The only real exception is the triggers with low-gain saturation: for the 1734 triggers of this type, 897 tuples are found, about 50% compared to about 10% for the other triggers.

In the right panel of Fig. 6.10 the multiplicity of the tuples is shown for different selections of the types of the contained triggers. The distributions are fitted with the expectations from Eq. (6.5) leaving out the trigger doublets as they deviate clearly from the expected behaviour with respect to the higher multiplicities. A possible explanation for this can be the presence of dead-time after multiple triggers [100].

The obtained λ_{eff} values are (19.2 ± 0.1) Hz for all tuples, in reasonable agreement with the mean trigger rate of about 22 Hz, (27.5 ± 0.2) Hz for tuples containing at least one ToT, and (70 ± 2) Hz for pure ToT tuples. These values are in clear disagreement with the trigger rate given that the average ToT rate is at about 2 Hz. This clearly indicates an excess of such triggers on short time scales.

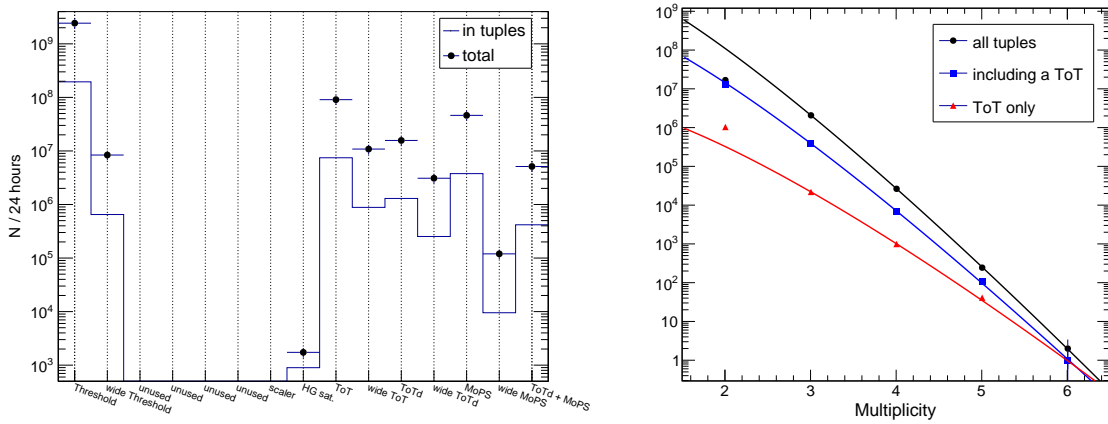


Figure 6.10: *Left:* Distribution of the triggers of different types in the data set. The points show the count of triggers in the data while the histogram shows their abundance in tuples. *Right:* Trigger multiplicities of the tuples observed in data for different selections: all (black), including at least one ToT (blue), ToTs only (red). The fit follows Eq. (6.5), fitting the mean rate λ and the overall scaling n_0 for multiplicities greater than two. The obtained values are (19.2 ± 0.1) Hz, (27.5 ± 0.2) Hz, and (70 ± 2) Hz respectively, in clear disagreement with the overall rate for ToT triggers of about 2 Hz.

As mentioned, the excess of the flag for low-gain saturated triggers is already visible in the frequency of tuples with this flags. This excess can be made even more visible by selecting tuples that start with such a trigger. Fig. 6.11 shows the multiplicity spectrum of this selection in comparison with the multiplicity spectrum of threshold triggers. The excess is not only due to a few singular stations or a single event in time (like in a thunderstorm), as can be seen in Fig. 6.12 showing the distribution of the stations that have such triggers as well as the time distribution of such tuples.

The right panel of Fig. 6.11 shows the corresponding time differences Δt_{21} of the second trigger of a tuple relative to the first. A lack of the time differences for about 250 to 300 μs is clearly visible. The majority of triggers appears about 500 μs after the initial trigger with the large-signal flag. From the distribution of the time differences Δt_{i1} , of trigger i in the tuple taken relative to the initial trigger, in the same Figure also the extent of the trigger tuple can be inferred with some uncertainty. The largest uncertainty here arises from the dead-time associated with multiple triggers in quick succession (c.f. Section 5.2.2). However, overall a time scale of 800 μs of trigger activity after large signals is observed.

This is in agreement with expectations from delayed neutrons if an explanation for the lack of triggers in the beginning can be found. As we cannot exclude instrumental effects, especially given that the two new trigger types are causing separate early peaks, we discuss this result further in conjunction with results obtained from analysing event data where more information about the initial event can help distinguish instrumental from physics effects.

6.2.4 Event Correlation Search for Delayed Signals

Following the discussion of delayed triggers in stations, we now focus on searching for the same effect following reconstructed event data. When searching for correlations with events, we use the vertical SD dataset of the ICRC-2019 [142] release of the Observer [126]. Because a reliable core reconstruction is of relevance for part of the selections used later, we remove




Figure 6.11: *Left:* Multiplicity of the constructed tuples for two different selections on the type of the initial trigger. Most tuples are initiated by a threshold T2 (blue) making up 93% of the data. The signal of delayed triggers is still clearly visible by selecting triggers with the low-gain-saturated flag as the initial trigger of a tuple (red). Such tuples contribute only 8 ppm of the overall data set, but contain 1.1% of quadruplets. *Right:* Distribution of the time difference of the first two triggers of tuples with a high-gain-saturated flag. The distribution for the tuples with a new trigger in the second position in the tuple are given separately. In gray, the distribution of the time differences to the initial trigger are given.




Figure 6.12: Cross checks for tuples initiated with triggers with the low-gain-saturated flag. *Left:* Station IDs of the trigger tuples. *Right:* Time distribution of these tuples during the day analysed.

events with $\lg(E/\text{eV}) < 18.0$, use only 6T5 events, and exclude events with zenith angle $\theta > 60^\circ$.

The goal of this analysis is to find triggers that appear delayed with respect to the known (reconstructed) air-shower events. Given that the focus lies on distinguishing a local (station) effect from air-shower effects, we design this method as “station-based”. This means that rather than taking the event time and searching for triggers delayed with respect to this time, we use the trigger time t_{trig} of each station in a event as the basis for building the time difference. With this time difference, the relative time series for other triggers of that station can be extracted from the trigger files. For every station (with signal) in a reconstructed event we thus have the shower-level variables: reconstructed energy E_{rec} , zenith angle θ , and shower-plane distance r_{sp} , as well as – on station-level – the total signal S_{tot} . For each of these stations, the time difference $\Delta t_i = t_i - t_{\text{trig}}$ and types and flags of all the T2 triggers in a time window of $-1000 < \Delta t / \mu\text{s} < 1000$ are extracted and stored.

To search for signals similar to those observed in the trigger-tuple analysis, we select here only the ToT triggers² and perform selections on the saturation status of the traces of stations. To enable the detection of such an effect in the correlation with air showers we tag stations from events that have at least one trigger in the interval $100 < \Delta t/\mu\text{s} < 1000$ as having *trailing triggers*. A similar definition is applied for early triggers to control random background, such that stations with any trigger within $-1000 < \Delta t/\mu\text{s} < -100$ are referred to as having *leading triggers*.

Then, the properties of the showers with stations that have trailing triggers can be compared to the random background expectation, inferred from the leading triggers. The question whether or not such an effect is of instrumental nature can then be answered based on the degree of correlation of the trailing-trigger tag with local and shower-level observables. For example, we expect that for the neutron hypothesis the number of stations with trailing triggers should increase with reduced distance to the core as well as with increased energy. On the contrary, an instrumental effect of the station would result mainly in a correlation with the station signal, independently of the shower that caused it.

Moving to the results of applying this method to data, Fig. 6.13 shows the distribution of the logarithm of the total signal of all stations selected in the input data sample. Of these 1.1 million stations in events, roughly 300 000 do not appear with the trigger in the T2Dump files. This can be due to two different reasons: traces obtained by requesting T1 triggers with the T3 do not appear in the T2Dump files, and we reject some stations and some triggers for being unreliable (c.f. Fig. 5.3). Because the amplitude of the signals recorded with a T1 is small, the loss is not relevant for this study. On top of that, the number of rejected stations varies with time, but is about 150 for the given period with more than 1400 stations remaining, leaving sufficient statistics.

When searching for a signal of trailing triggers also the results of the tuple analysis should be reproducible: the signal should be enhanced with large signals. As expected, the distribution of the signals that have trailing triggers has a very clear excess for large values, as shown in the left panel of Fig. 6.13. The right panel of Fig. 6.13 highlights this fact even more clearly by selecting the stations that are low-gain saturated. Even here, only those with the larger signals have a trailing trigger on top of an almost negligible background. The total excess is the same as for all stations, showing that only low-gain saturated stations contribute to the excess. It should be noted that in the distributions shown in Fig. 6.13 no recovery of saturated stations is attempted to avoid loss of statistics and because, in general, the integral of the trace is still larger for a “more saturated” station than for the one with a lower signal. The distortions that are relevant for reconstruction are not important for a pure correlation analysis as used here. As low-gain saturation indicates very large signals that can cause short term issues in the stations, this result points towards a possible instrumental effect.

To check the other hypothesis of delayed particles from air showers Fig. 6.14 shows the correlation of low-gain-saturated stations with and without trailing triggers with the reconstructed energy and the distance to the core. As expected for a physical signal, the fraction with trailing triggers increases with higher energy and smaller distance to the core. However, it should be noted that this also correlates directly with the measured signal amplitude.

To solve the ambiguity in this correlation, we look at the correlation of a pure station-level quantity, the total signal S_{tot} and the distance to the core in Fig. 6.15. The neutron or late-particle hypothesis would also for lower energies predict a higher fraction very close to the core. A gradient of the trailing trigger events in the r_{sp} -direction is thus expected. However, we observe what is expected for a local-station effect: (almost) no gradient for r_{sp} and a strong gradient in S_{tot} .

²This includes all the sub-types of ToT, technically: trigger flag >7.

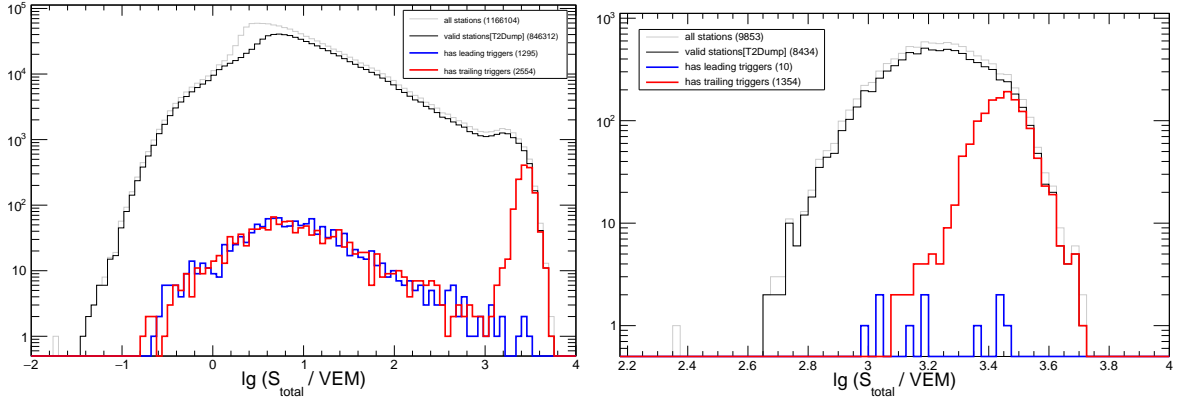


Figure 6.13: *Left:* Total signal of all stations in the data set. T1 stations and stations that were rejected by the T2Dump station selection are present in the gray histograms. The black histograms contain all *valid* stations. Here, the total signal is used and no saturation recovery is attempted. Numbers in parentheses give the total amount of statistics shown for each sub-selection. *Right:* Distribution for stations that are low-gain saturated.




Figure 6.14: Distribution of the distance to the core and shower energy of stations with low-gain saturation. Stations that show delayed triggers are marked in red. They cluster at high energies and closer to the shower core. Blue marked stations with a (random) leading trigger are very rare.

In summary, we observe a clear signal of trailing triggers also for high-energy events which are very similar to the one found in the trigger-tuple analysis. Note that in the data analysed for the tuples (one day) no event with a trailing trigger was found. Given that overall we observe a frequency of 1.2 events with trailing triggers per day, a zero-observation is contained within the likely fluctuations. The vast difference in statistic also means that most of these signals are caused by sub-threshold showers and not high-energy events. The analysis of the properties of stations showing the trailing triggers strongly indicates that it is an instrumental effect of the local station.

6.2.5 Effects of the Station Electronics

In this section, we present a possible explanation for the observed trailing triggers. The question to be answered is whether there is a known effect that can create an almost complete lack of triggers after a large signal lasting for about $200 \mu\text{s}$ and then inducing one or more triggers after that on a time scale of up to 1 ms.




Figure 6.15: Distribution of the total signal (non-recovered) for low-gain-saturated stations and their shower-plane distance to the shower core. Since a “more” saturated station will also show a larger total signal (even if this is not representing the true signal) any effect caused by the station (local feature) will appear as a feature in the S_{tot} -distribution and not as a radial gradient. Stations that show trailing triggers are marked in red. No clear gradient with the shower-plane distance is observed hinting at a local feature rather than an effect of the air showers.

It is a known property of the PMT and its base [143] that after large signals the baseline has an undershoot that can in certain conditions be used to recover the original signal itself [113]. The relaxation of this undershoot is a possible explanation of the observed trailing triggers. As it was already measured in the prototypes of the SD, the baseline relaxes on a time scale of about 1 to 1.5 ms [112]. We used the measurement with an oscilloscope attached to the PMT base already in the previous Chapter in Fig. 5.20. It indicates that not only the undershoot is long enough to explain the lack of triggers, but also that there is a substantial overshoot for an even longer time that perfectly matches the time scales of the trailing triggers we observed in the data.

Furthermore, it is possible to show the correlation of the rate of such trailing triggers with the amplitude of the expected overshoot as this can be calculated from the undershoot and the properties of the PMT base. We estimate the overshoot by calculating the size of the undershoot first. We use the difference of the first and the last 15 bins (averaged) of the high-gain traces and apply a conversion factor. From the circuit schematics of the PMT base we know that this factor is about $1/7 = 0.143$ [112, 113].

To link the trailing triggers more clearly with this expected overshoot, we scale the calculated expected overshoot with the online estimates of the VEM as this is the relevant quantity for the definition of the triggers. For simplicity we take the average (online) VEM trace of all non-masked PMTs. Then, the number of stations having trailing triggers should be close to 100% once this overshoot becomes comparable to the trigger threshold (0.2 VEM for ToT and ToTd [144]).

Fig. 6.16 shows the expected overshoot for all stations and highlights those with trailing and leading triggers. We use the leading triggers again to estimate the contribution from the random background. A clear excess of the trailing triggers is visible for expected overshoots of more than 0.1 VEM.

We show the corresponding ratio of events with trailing triggers as function of the expected overshoot in the right panel of Fig. 6.16. The expected threshold behaviour around 0.1 VEM is clearly observed. Deviations from this expectation – especially at large values – arise most likely from wrongly estimated baseline undershoots as the used algorithm is rather simplistic.

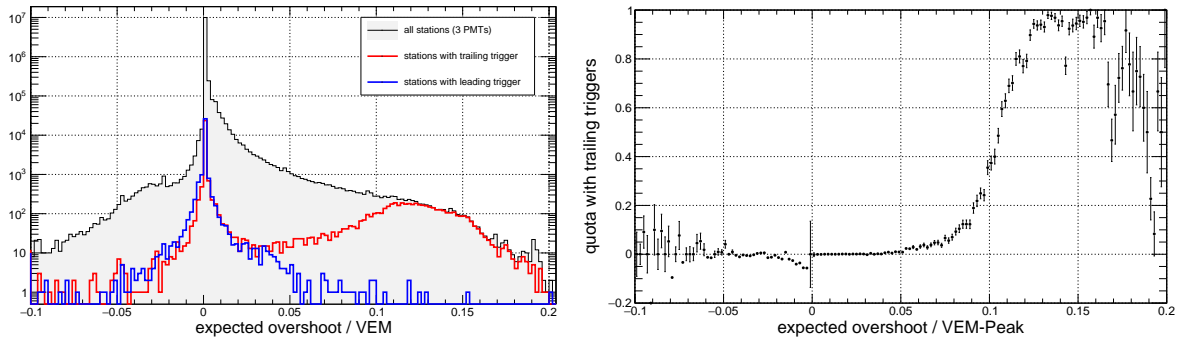


Figure 6.16: Distributions of the expected overshoot for all stations with 3 functioning PMTs. *Left:* All stations (black), those with at least one trailing trigger (red) and those with at least one leading trigger (blue). *Right:* Relative number of stations showing trailing triggers, calculated per bin as $(N_{\text{trail}} - N_{\text{lead}})/N_{\text{tot}}$. Errors are approximated as arising from a binomial process of N_{tot} tries with $N_{\text{trail}} - N_{\text{lead}}$, the estimated number of true excesses, successes. If a deficit of trailing triggers is observed this estimate breaks down and no error is shown. The tail at large expected overshoot can contain singular events which due to the simplicity of the used algorithms have a wrongly assigned undershoot.

This match of the expected value of the baseline overshoot to the trigger threshold and the matching approximate time scale of the undershoot relaxation are proof that most – if not all – trailing triggers originate from the “wrong” baseline estimate used for the trigger in the moments after large signals.

6.2.6 Prospects of Detecting Very Late Particles

In the previous sections, we used two different methods to search for delayed signals beyond the scale of the trace of about $20 \mu\text{s}$. Both methods show a clear signal which nevertheless matches the expectation from the relaxation of the baseline and thus is solely an instrumental effect. We do not think that this effect has a sizeable negative impact on the data taking itself. However, in the case that a further search for the effects of delayed neutrons or other particles, as observed in Ref. [137], is attempted, this instrumental background has to be properly taken into account.

A possibility that still remains in search for delayed neutrons (or, in general, enhanced contributions close to the core) is to identify oddities in the fraction of events showing trailing triggers.

In such an approach, the hypothesis is that the fraction of showers creating a station signal large enough to cause trailing triggers is rising with energy in a smooth way that in principle could be predicted with simulations. Any deviation from a smooth or expected fraction could hint at a change of the signal very close to the core. Such a change could potentially arise from changing composition or from changed hadronic interactions.

To highlight the available statistics and the concept, Fig. 6.17 shows the energies of the observed events and the corresponding fraction of events with trailing triggers. We observe, as expected, a smooth rise of the fraction of events with trailing triggers with energy that can be well described with an exponential function $f(E) = \exp(-a(\lg(E/\text{eV}) - 18) + b)$ with fitted parameters $a = 0.96$ and $b = -2.29$. Solely for illustration, using the small (and insignificant) fluctuation around $\lg(E/\text{eV}) = 19$ as an example, we show how a possible excess could be identified and fitted. The fit of the excess with a Gaussian peak on top of the same exponential background yields an amplitude of the peak of 0.11 ± 0.05 at $\lg(E/\text{eV}) = 19.16 \pm 0.04$ showing that it is only a 2σ fluctuation.




Figure 6.17: *Left:* Distribution of the energies of the events used for the correlation analysis. The distribution of energies of events with trailing triggers is shown in red. The almost negligible random background is estimated from the events with leading triggers and is shown in blue. *Right:* Fraction of events having trailing triggers as a function of the reconstructed energy. The expected rise of the fraction is fitted with an exponential function (black). Unexpected excesses in this evolution could manifest themselves as peaks on top of this background. For example, a Gaussian peak is fitted together with the exponential background to the data (red). The shown excess at $\lg(E/eV) = 19.16 \pm 0.04$ with an amplitude of 0.11 ± 0.05 is not significant.

A complete analysis of such excesses of course has to take the look-elsewhere effect into account and has to be performed independently of the binning. However, as this analysis of trailing triggers is dominated so far by an instrumental effect, its exact understanding also in the new electronics is crucial for any detailed study using this data. In addition, it is probably more precise to use the trace data itself to search for changes in the saturation levels and fractions with energy. Such analysis is beyond the scope of this thesis and is left for future work.

6.2.7 Detection of Pulses in Event Traces

To extend the search for delayed particles to the time scales smaller than several 10 microseconds, as what can be done using trigger-data only, we focus on the trace data recorded with the events. For these time scales, simulations suggest that many neutrons are expected to arrive in the stations, see Fig. 6.7.

In the following, we use a simple algorithm to search for pulses in the recorded traces of events. For every trace for simplicity, we calculate the average VEM-signal of the active PMTs, shown as black solid trace in the example in Fig. 6.18. We use the calibration of the peak-height (VEM_{peak} , c.f. Section 3.2.3) to avoid differences of PMTs to govern the choice of thresholds and the sensitivity.

On this average trace, we define a peak as any compact set of time bins that has at least one bin above a threshold $I_{\text{th}} = 0.03 VEM_{\text{peak}}$. To determine the length of the peak, we use a lower threshold $I_{\text{low}} = 0.02 VEM_{\text{peak}}$ to define the end-points of the compact peak interval A . Formally,

$$A = \{ \text{consecutive } i \mid \forall i \in A : t_i > I_{\text{low}} \ \& \ \exists j \in A : t_j \geq I_{\text{th}} \} \quad (6.6)$$

with the trace bins t_i . The two choices of thresholds are marked in the example Fig. 6.18 as red lines and highlight that the values are tuned to be just above the noise level, taking the quantisation of the values into account. From the right panel of Fig. 6.18 it can be also visually inferred that the difference of I_{th} and I_{low} is relevant to exclude fluctuations but also to properly include the full extent of the single pulses. While it might be desirable to include also pulses at even smaller amplitudes, e.g. the one at $16.5 \mu\text{s}$ in the Figure, the influence of noise becomes too large if a lower threshold is used.

To characterise each pulse p_A , we define the charge Q_{pulse} by computing the integral of the VEM_{charge} calibrated PMT signals in the interval A and define the time of the first bin in A as t_{peak} . This time can be used to relate it to the signal time of the main cosmic-ray signal as used in the reconstruction $\Delta t_{\text{start}} = t_{\text{peak}} - t_{\text{start}}$. We added the signal window for illustration in Fig. 6.18 as vertical grey dashed lines. The end-time is relevant if the effect of any late pulses on the determined total station signal is of interest.

To characterise the distribution of the signals of involved PMTs, we define the signal symmetry of a detected peak as

$$S = \frac{\langle Q \rangle_{\text{other}}}{Q_{\text{max}}}. \quad (6.7)$$

Here Q_{max} is the maximal charge of the active PMTs and $\langle Q \rangle_{\text{other}}$ the average charge in the other two³ PMTs. Thus, $S \rightarrow 0$ for a maximally asymmetric signal (only in a single PMT) and $S \rightarrow 1$ for a perfectly symmetric signal (equal signals in all PMTs).

We apply the algorithm described above to all *candidate* stations of reconstructed events passing the selection for the energy spectrum [3] with energies above $10^{18.5}$ eV, i.e. the full efficiency threshold. To avoid effects of the moving baselines after saturated signals, we also exclude all stations marked as high-gain saturated⁴. This selection leaves us with 841 564 traces from 160 080 events. The resulting distribution of the distances to the shower-core in the shower-plane r_{sp} and of the total signals are shown in Fig. 6.19. The selection of non-saturated stations leads to a “maximal” signal of about 200 VEM and the removal of stations closer than about 500 m to the core. Future work will aim at including also stations closer to the core and with larger signals to complete the available phase-space of the measurement if adequate performance of the peak-finder in such traces can be ensured.

6.2.8 Late Particles in SD

Applying the peak finder presented in the previous section to data, we get a total of 9.94×10^6 pulses. The shower signals themselves are included in this number and are expected to be of large Q_{pulse} and within $\mathcal{O}(1 \mu\text{s})$ of the signal start time. Different contributions to these pulses can be distinguished from the time and charge distributions. Fig. 6.20 shows these distributions in a logarithmic colour-scale. The shower signals are visible as narrow band of high-charge pulses at $\Delta t = 0$ separated by an “empty” area at small negative Δt . This suppression is due to the construction of the signal start time from signal pieces and their separation [73].

As expected, at $\lg(Q/\text{VEM}) \approx 0$ a band of accidental atmospheric muons is visible extending in the trace to the usual range. From the desired position of the signal at the “latch” bin ~ 250 , we expect it to lay in the 768-bin trace in a time range between $-6.25 \mu\text{s}$ and $12.95 \mu\text{s}$. A handful of signals are visible outside of this time range. They are expected due to traces triggered by accidental muons around the shower signal. In these cases t_{start} will deviate significantly from bin 250 and will thus open the time interval of earlier times. Similarly, the

³If only two PMTs are functioning, no averaging is used. For the case of only one PMT in use, the symmetry is ill-defined.

⁴implicitly including low-gain saturated stations




Figure 6.18: Visualisation of the peak-finder algorithm with data from station 157 in event 620441. The start of a peak is defined by crossing a first threshold (solid red) of $0.03 \text{ VEM}_{\text{peak}}$ from below. The width is set by searching for the first bin that crosses a second threshold (dashed red) of $0.02 \text{ VEM}_{\text{peak}}$. *Left:* The full trace with late signals visible as small peaks at $t \approx 14 \mu\text{s}$. *Right:* Zoom into the late part of the trace with the delayed pulses. It is apparent that some of these pulses are caused by single PMTs, requiring a low threshold in the average signal used for the peak-finder.




Figure 6.19: Statistics of the selected data sample of stations in reconstructed SD events. *Left:* Distribution of the distance in the shower-plane r_{sp} of the selected stations. The cutoff at $\sim 500 \text{ m}$ is due to the requirement of no saturation. *Right:* Distribution of the total signal sizes in the selected stations.

signal of the shower can be shifted to earlier times, e.g. due to the veto period of an accidental earlier trigger. Additionally, we observe a range of lower-amplitude pulses placed within $3 \mu\text{s}$ of the signal start time, as expected from the usual width of the shower signals, especially far away from the core.

When considering smaller-amplitude pulses, the change of rate is more visible in the projections of the respective charge ranges as we show in the right panel of Fig. 6.20. A clear increase is visible which upon closer investigation also shows more features we can associate with known processes. For instance, the weak band at $\lg(Q/\text{VEM}) \approx -0.7$, best visible at $\Delta t > 5 \mu\text{s}$, corresponds to electrons from stopped and decayed muons. This correspondence is proved by comparing the deposited energy, proportional to the signal, which should be within experimental uncertainty equal to the maximal energy of the electron in the muon decay. In vacuum, this limit is given by half the muon mass $E_{\text{max}} \approx 53 \text{ MeV}$, however, in the presence of atoms it can be modified due to the momenta transfer to the nucleus. Using E_{max} as an order-of-magnitude estimator because the full decay spectrum is not resolved in

the presented analysis, we get $\lg(Q_\mu/\text{VEM}) \approx \lg\left(\frac{E_{\text{max}}}{E_{\text{VEM}}}\right) = -0.7$ when taking the deposited energy by 1 VEM as $E_{\text{VEM}} = 250$ MeV. In addition, the time scale of the exponential decay of these pulses, visible in the red projection histogram in the right panel of Fig. 6.20, matches within uncertainties the time-scale of the muon decay of about $2 \mu\text{s}$. It seems possible to use these decays to alternatively check the VEM-calibration for electromagnetic particles since the usual VEM-calibration is using through-going muons. The necessary selection, background estimation, and simulations are beyond the scope of this work and should be explored in the future.

The other projections in Fig. 6.20 reveal that for the smallest resolvable charges there is a ‘‘bump’’ in the time profile which is best visible in a charge range between $-1.5 < \lg(Q/\text{VEM}) < -1$. While it is also present at smaller charges, the background from ‘‘noise’’ is much larger at small charges. The fact that these small pulses are not only noise effects of the baseline is clear from the comparison of the levels before and after the signal, with at least an increase by a factor 10 comparing small pulses at $\Delta t \approx -5 \mu\text{s}$ to those at $\Delta t \approx 12 \mu\text{s}$. The main question associated with observing such a ‘‘bump’’ in the time spectrum is whether or not this is related to a physical effect or is it rather caused by the PMT.

PMTs in general can show afterpulses [145] following to the true pulse. These occur mainly through two processes that are linked to generation of electrons in the photo cathode. First, rest gas in the tube can be ionised and the drift of these ions in the field causes a measurable delay until the interaction of the ion with the photo cathode creates free electrons. Second, the electrodes in the PMT can emit photons when bombarded by electrons of the true-signal avalanche which can in rare cases make it back to the photo-cathode. In common is to both cases that the probability should scale with increased signal and that they can happen in individual PMTs more often.

To test, whether or not these delayed pulses are afterpulses we thus use the station signal as parameter to correlate the frequency with the signal size. As a decisive quantity for the opposite hypothesis, that the pulses are from delayed shower-particles, we can use the radial distance of the station. Both quantities are correlated because, in general, stations closer to the core have larger signals. However, because in the data sample we mix different energies and zenith angles there is complementary information included.

Fig. 6.21 shows the distributions of pulse times Δt_{start} and signal size S (left) or radial distance r with a logarithmic colour scale. Overall, the colour scale is dominated by the radial distribution and the excess of delayed pulses is not directly visible. To increase the visibility of the bump, we show the projections on the Δt_{start} axis in bands of S or r . These distributions are shown in Fig. 6.22 with the bump of delayed pulses clearly visible. Comparing the two panels, it seems that the distance to the core r has more influence on the relative amount of these pulses than the signal size itself.

To further highlight this point, we show different radial selections using (almost) constant S in Fig. 6.23 for fairly low and larger signals. It is evident that especially for the lower total signal selection ($|S - 25 \text{ VEM}| < 5 \text{ VEM}$) the relative amplitude of the late pulses changes significantly with radial distance. For the larger signals, hints of the same effect are visible. However, due to the reduced statistics it is less evident.

To exclude that these effects come from known interactions of particles, we can use the normal air-shower simulations and repeat the analysis. Fig. 6.24 shows the distribution of pulses with their delay with respect to the shower start time from such simulations [146] using proton simulations with the EPOS-LHC interaction model [147] at similar energies of $18.5 < \lg(E/\text{eV}) < 19.0$ with a flat spectrum in $\lg E$. However, the detector simulations include neither neutrons, nor radioactive decays in the water, nor PMT afterpulses. Clearly, the standard simulations do not reproduce the feature of delayed triggers. The quantisation of the charges in the simulation is likely due to the exclusion of different calibration values




Figure 6.20: Distribution of the charge and relative time of pulses in event traces. In general, due to the DAQ event builder, the shower signal is expected to start at the bin 250, leading in conjunction with the trace length of $19.2 \mu\text{s}$ to the structures visible between $t_{\text{start}} = -6 \mu\text{s}$ and $t_{\text{start}} = 13 \mu\text{s}$. *Left:* Two dimensional distributions of the relative time of a detected pulse and its charge. We use the start time of the shower signal as determined by the reconstruction [149] as a reference time. Visible structures include the shower signal, a band associated with random muons at $\lg(Q/VEM) = 0$, a band of decaying muons at roughly $\lg(Q/VEM) = -0.7$, as well as a decaying distribution of late peaks with an additional structure at $t_{\text{start}} \approx 8 \mu\text{s}$. *Right:* Several projections of the time of pulses depending on the charge of the pulse. Charge bands are chosen as the two lowest charge ranges and one containing the expected pulses from decaying muons.




Figure 6.21: Correlation of the time distributions of late pulses with quantities from the station (total signal, left) and the reconstructed shower (radial distance in the shower plane, right). A strong correlation of the local quantity “station signal” with late arriving pulses could be a hint that local effects like PMT-afterpulsing play a role in the creation of these signals. Likewise, a strong correlation with the shower quantity makes a physical signal more likely. *Left:* Distribution of the pulse times relative to the start time of the signal as a function of the shower signal on the y-axis. The effect of resolved single particles for low-signal stations is visible as only for these stations peaks close in time to the initial shower signal are found. We do not observe a clear correlation of the late particles with the signal size, in parts due to the steeply falling spectrum of signals. *Right:* Using the radial distance of the station to the shower core as a variable on the y-axis, the distribution of late pulses shows a slight increase with decreasing distance to the core.

for the VEM in FADC. In more recent simulations these steps disappear, but the general observation of a missing tail of small pulses and especially the absence of the “bump” remains.

We discuss dedicated simulations including neutrons of all energies in Section 6.2.10 after discussing first results for the scintillator data from the SSD-PPA [77, 148] in the following Section.




Figure 6.22: Projections of the time distributions in bands of total signal and radial distance. *Left:* Dependence of the pulse times for different shower signals. We observe mostly parallel, i.e. scaled, distributions indicating that the bump caused by late pulses is not (directly) correlated to a local effect. *Right:* Using radial bins, the distributions of pulse times change significantly hinting at an origin in the shower. Note that due to the requirement of no saturation the distributions are cut off at about 500 m from the core.

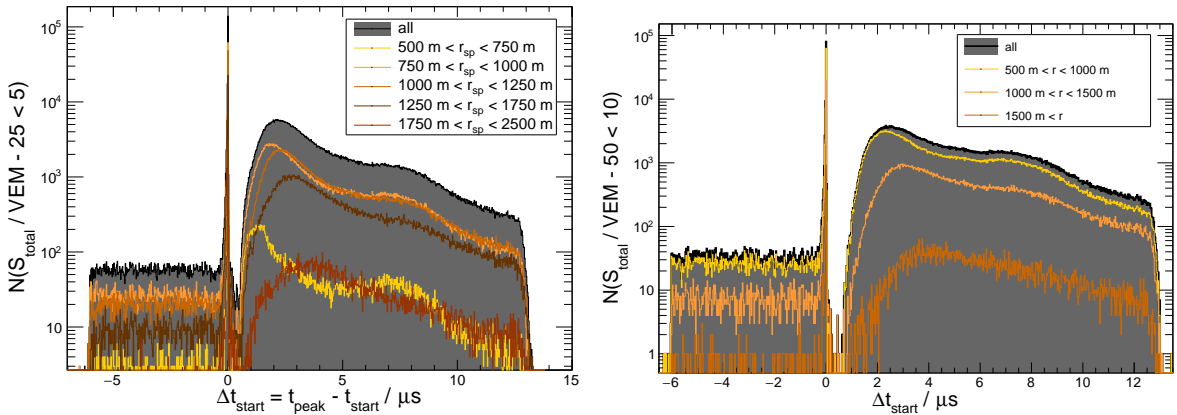


Figure 6.23: Disentangling the influence of the distribution of radial distances and the signal size on the delayed pulses. *Left:* Selecting signals of about 25 VEM and showing the distribution of Δt_{start} for different radial bands. A clear dependence of the relative amplitude of the bump at $\Delta t_{\text{start}} \approx 8 \mu\text{s}$ on the radial range is visible. *Right:* Selecting larger amplitude signals of about 50 VEM. The range is larger than for the smaller signals to account for the lowered statistics.

6.2.9 Late Particles in SSD

Following recent results from the Auger collaboration [135], discussing delayed pulses on the order of $O(1 \mu\text{s})$ with respect to the signal window in the SD, we tentatively apply the same peak-finder to the data of the SSD Pre-Production Array [148]. We use the same thresholds, but now converted to MIP, of $I_{\text{th}} = 0.03 \text{ MIP}_{\text{peak}}$ and $I_{\text{low}} = 0.02 \text{ MIP}$. Due to the smaller available data sample of only about 42 000 events with $\sim 101\,000$ traces the influence of eventual noise contaminating the pulse sample is not as large.

Fig. 6.25 shows the resulting distributions of logarithmic pulse charge and delay, similarly to Fig. 6.20. From the quantised bands for low Q , it is evident that the peak-finder thresholds should be revisited in the future. However, for the purpose of observing delayed pulses of significant charge, we can simply ignore the bands at small Q .

Due to the low statistics, most delayed pulses appear as bins with count “1” in the colour plot. To make the effect of an increase in the rate of late pulses more visible, we show projections onto the Δt_{start} axis in the right panel of Fig. 6.25. We pick two bands that are




Figure 6.24: Distribution of pulse charges and times as in Fig. 6.20 from SD-simulations. We use proton simulations in an energy range of $18.5 < \lg(E/\text{eV}) < 19.0$ simulated with EPOS-LHC model [147] from the Naples library [146]. These simulations do not include neutrons, radioactive decays, and PMT-afterpulsing. *Left:* Two dimensional distribution of logarithmic pulse charge $\lg(Q/\text{VEM})$ and the delay Δt_{start} . Discrete levels probably arise from using the exact same value for the VEM in FADCs. *Right:* Projections of bands in $\lg(Q/\text{VEM})$ onto the Δt_{start} axis.




Figure 6.25: Distribution of pulses detected in SSD data. We use the SSD-PreProduction array (SSD-PPA [77]) data [148] for these measurements. *Left:* Pulse charges as function of the relative time of the detected pulse with respect to the start of the shower signal. The influence of noise is clearly visible. *Right:* Projections in different charge intervals to highlight the frequency of the late pulses.

not affected by the low-amplitude noise pulses, i.e. $-1 < \lg(Q/\text{MIP}) < 0$ and $\lg(Q/\text{MIP}) > 0$ separating small and large pulses. Clearly, the rate of large pulses is above background levels even for $\Delta t_{\text{start}} \approx 12 \mu\text{s}$, while the majority is present at about $3 \mu\text{s}$ after Δt_{start} matching the previous observations [135].

At lower amplitudes, the shape of the distribution is similar. Particularly, we do not observe an increase in the range of 7 to $9 \mu\text{s}$ as in the SD. However, also the background is higher in the SSD. This can be inferred from the comparison of the number of pulses before the main shower signals in Fig. 6.25 and Fig. 6.20. While for the SD, with half the interval-size ($-1.5 < \lg(Q/\text{VEM}) < -1$) about 30 pulses per 25 ns are detected in the full data sample of 160 000 events, in the SSD-case the background is at about 100 pulses per 25 ns for only 42 000 events. In addition, the capabilities of the two detectors to detect small pulses are different so that it is unclear to what extent the same effect should be visible. We discuss hints to such differences in the following Section.

6.2.10 Discussion

With the detection of delayed pulses in both SD and SSD, the decisive question of whether they correspond to physical signals or not is hard to answer by looking at data only. While the different dependence of the signals on r and S point towards physical origins, without a more detailed understanding of the potential signal it is difficult to draw conclusions. For instance, while it can be easily understood that neutrons can be delayed to several μs , as we have seen in Fig. 6.7 already because they are stable (on the scale of air showers) and neutral, it is not clear why this should create a bump rather than a continuous distribution.

To help answer this question, we refer to simulations of air showers done in FLUKA [141] keeping track of neutrons and protons of all energies. Using simplified detector simulations, e.g. simulating the SD as a water volume and counting the energy deposit in Cherenkov light only, we obtain predictions of the time distribution of signals created by neutrons (and protons) [150]. We cite and show the results in Fig. 6.27 for the SD and in Fig. 6.28 for the SSD.

Interestingly, these simulations predict a peaked component of low-amplitude signals for the SD at the time when we observe them in data. These signals come from nuclear excitations of oxygen (mostly at about 6 MeV, c.f. [151]) which decay electromagnetically. Due to the low energy when compared to VEM (about 250 MeV) it is expected that such signals are only visible as “localised” pulses, i.e. if the reaction happens close to one of the PMTs or in a favourable geometry of the Cherenkov photons. Thus, we expect a high-asymmetry in the pulses coming from these signals.

To confirm this, we compute the average signal symmetry S following Eq. (6.7) per bin of $\lg Q$ and Δt_{start} excluding signals from stations with only one PMT working. Fig. 6.26 shows this symmetry colour coded with $S = 1$ being maximally symmetric and $S \approx 0$ fully asymmetric. Negative values of S are possible if the signals in the two smaller-signal PMTs are solely negative fluctuations of the baseline. It is clearly visible that the shower signals and accidental muons create highly symmetric signals, as expected. Concerning the delayed pulses, the signals indeed are very asymmetric, as expected. Interestingly, also the increased level of very-low amplitude background pulses are linked to very asymmetric pulses pointing to similar effects.

From the radial distributions in the simulations, it is also expected that the bump of neutrons shifts to later times with increased distance to the core [150]. We try to confirm this prediction in data by selecting radial bands $r < 1600$ m and $r > 1600$ m for a given energy range corresponding to the energy used in the simulations. For the nearby stations, we use $18.5 < \lg(E/\text{eV}) < 18.6$ and for those far from the core $18.5 < \lg(E/\text{eV}) < 18.75$. We increase the energy interval for the stations far from the core to increase the statistics but we have to keep that in mind in the interpretation of the result. The right panel of Fig. 6.26 shows the resulting time distributions for a narrow band of $-1.8 < \lg(Q/\text{VEM}) < -1.5$ containing the neutron “bump”. No clear shift is visible due to the change of statistics and background, but the turn-over of the bump happens roughly $1 \mu\text{s}$ later for $r > 1600$ m. Given the change in background and the resulting systematic differences in the turn-over, we do not want to over-interpret this result. It is clear however that in future work there is room for improvement in such radial profiles and their dependence on the energy.

For the SSD, the simulations reveal that the delayed pulses are indeed coming from protons and neutrons as Fig. 6.27 clearly shows. Also here, a component at about $7 \mu\text{s}$ is predicted as for the water-Cherenkov detector. However, due to the differences in noise and background, we cannot resolve it in the data. Interestingly, the overall number of signals is larger in the SD than in the SSD by about a factor 4, although most pulses in the SD are of much smaller amplitude.




Figure 6.26: Properties of the observed pulses for comparison with detailed the simulations of delayed particles (c.f. Fig. 6.27). *Left:* Average (PMT) symmetry S from Eq. (6.7) for all pulses per bin of $\lg Q$ and Δt . Traces from stations with only one PMT are excluded. *Right:* Comparison of the time distributions of delayed pulses with $-1.8 < \lg(Q/VEM) < -1.5$ for two radial bands. For $r > 1600 \text{ m}$ a wider energy range is used to increase the statistics. Simulations predict a shift of the low charge pulses at $\Delta t_{\text{start}} \approx 8 \mu\text{s}$ (“bump”) to later times farther from the core [150]. Due to the large background, no direct confirmation is possible.

From the component of the neutron (and proton) signal at small times up to about 1 VEM also in the SD it is clear that in future simulations the contributions of these particles to the total signal should be estimated. While these signals can probably account for a part of missing signals in simulations [6], we do not expect that they can explain the “missing” muon signal completely. For example, in highly inclined events, we expect the neutrons to be attenuated while the deficit of the simulations is still visible there [152]. In conclusion, it should be very clear that a more detailed look into the signals measured at late time in the detector and the corresponding simulations including neutrons of all energies is promising more insights into the details of hadronic interactions.




Figure 6.27: Simulations of energy deposition in Cherenkov light by neutrons and protons in a water tank [150] with the energy distribution and timing following air-shower simulations using FLUKA [141] of protons with $E = 5.6 \times 10^{18}$ eV. Reflections and details in the water tank are not simulated. The signals are scaled to those of vertical through-going muons simulated in the same way to resemble the energy scale of Fig. 6.20. t_c corresponds to the time that a particle produced in the first interaction needs to reach the observation point travelling in a straight line at c , and is used as a proxy for the signal start time used in data. The colour scale is normalised to the number of muons, i.e. per muon of the air shower the given number of pulses from neutrons and protons are expected. The two-dimensional integral evaluates to 2.7, referring to on average 2.7 observable pulses per muon. (E.g. for a station observing 10 muons, 27 pulses from neutrons are expected.) Kindly provided by A. Ferrari [150].




Figure 6.28: Prediction of pulses from neutrons and protons in the same simulations as in Fig. 6.27. Signals are scored in a thin scintillator and are scaled to match MIP. Per incident muon, in total 0.73 pulses are registered. Kindly provided by A. Ferrari [150].

6.3 Searching for Multi-Core Events

In this section, we explore the possibilities of observing multi-core events with the Auger SD. In general, we refer to any set of correlated air showers measurable in the SD as distinct events as multi-core events. More specifically, such events are predicted in the context of photo-nuclear splitting of cosmic-ray nuclei in the solar radiation field, known as GZ effect [5]. While this effect is known to exist, the predictions of the observable rates of GZ-events range from low to basically impossible to detect [40–43] in real detectors. Overall, there is agreement

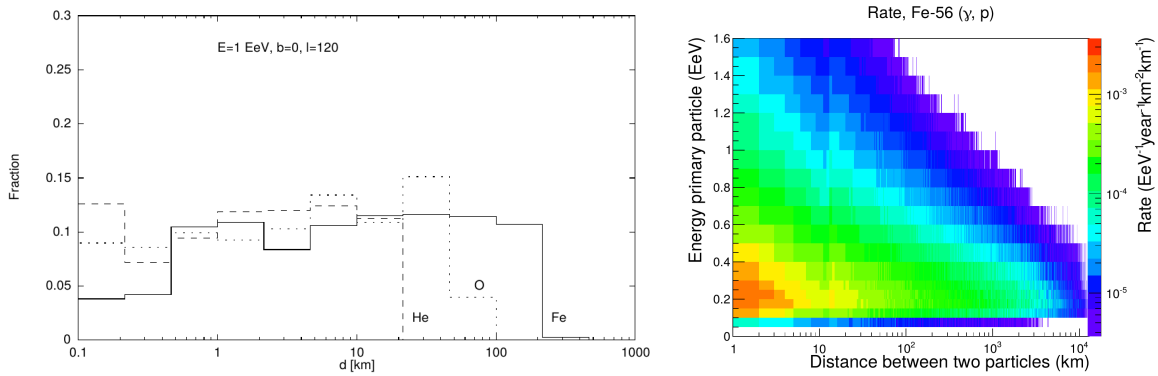


Figure 6.29: Prediction on the separation between the two GZ fragments. *Left:* Distribution of the separation for different primaries at $E = 1$ EeV and a night-side direction. Taken from Ref. [40]. *Right:* Differential flux of GZ events as a function of primary energy and separation. For this prediction, it is assumed that the full cosmic-ray spectra consists of iron nuclei. Taken from Ref. [43].

that the rate of GZ events is about 1 per year for an Auger-sized detector at energies well below the threshold energy of Auger SD. Here, we want to make use of the additional information we have from the triggers to lower the energy threshold of Auger significantly and provide a first measurement or limit of the rate of GZ-events in Auger for these low energies. In contrast to other searches, e.g. [44], which used a detector distributed over multiple sites, we can make use of a (nearly) uniformly monitored area within the Auger SD.

This continuous area can be advantageous if the energy threshold is low enough, as can be seen from the right panel of Fig. 6.29. In this prediction [43], the overall rate of GZ events is maximal for small separations, increasing the importance of complete coverage on the ground. Furthermore, in previous predictions, e.g. in Ref. [40], a significant fraction of events with small separations is predicted. In the left panel of Fig. 6.29 this is shown for a direction away from the Sun. However, due to the apparently long tails of the distance distributions, visible in Fig. 6.29, the average separation might be very large, as reported in other works, e.g. Ref. [42]. Due to the limited size of the detector, in this work we focus on “small” separations on the order of several kilometres.

In the following sections, we first outline general methods to search for clusters, i.e. coincidences of stations matching geometric or temporal constraints. The clearest examples of such constraints and clusters are the normal events which form geometrically compact configurations (almost) coincident in time, where the time can also be expressed as matching roughly a plane shower-front moving at the speed of light. We generalise these both concepts of using the geometry or the plane-front approximation to reach beyond what is usually selected by the T3 trigger. We then apply these analyses descending in energy from reconstructed showers to purely sub-threshold (T2-only) searches. Because we have the (first) opportunity, we will also use this analysis to reproduce the shower triggers of the recorded shower data to provide an independent cross-check of the DAQ-system.

Finally, we extend the analysis from expected GZ events to theoretically not well motivated exotic signatures, sometimes referred to as “super pre-showers” [153], consisting of many coincident low-energy showers. We derive first limits on such signatures using the large aperture of the Auger SD and the opportunity to perform it independently from the recorded shower data because we now have the T2-level data available.

6.3.1 Spatial Clustering Methods

In this Section, we will discuss strategies we can use to search for events in the background of random triggers of all surface detectors focusing on the geometry of the stations involved. The strategy presented here is identical to Eq. (5.9) from Section 5.6.1 which is limited to search for nearest neighbour coincidences. To generalise the same search to less compact patterns, we relax the *connectivity* condition from Eq. (5.9). This relaxation of the conditions enables us as well to detect trigger patterns by using spatial coincidences, like in the implementation of the T3 conditions [46]. In detail, we define a binary connectivity measure c between T2 messages M_i based on their space-time distance,

$$\text{con}(M_1, M_2) = \begin{cases} \text{true} & \text{;for crown}(M_1, M_2) \leq 4 \ \& \ \Delta t \leq (5 \text{ crown}(M_1, M_2) + 3) \ \mu\text{s}, \\ \text{false} & \text{;otherwise,} \end{cases} \quad (6.8)$$

where Δt is the absolute time difference between the messages M_1 and M_2 and⁵

$$\text{crown}(M_1, M_2) = \lfloor \max\{\|\Delta\vec{x}\|_{\vec{a}_1}, \|\Delta\vec{x}\|_{\vec{a}_2}\} / 1500 \text{ m} + 0.8 \rfloor \quad (6.9)$$

is the crown distance with $\|\vec{x}\|_{\vec{a}} = \sum_i |x_i| a_i$, $\vec{a}_1 = (1, 1/\sqrt{3})$ and $\vec{a}_2 = (0, 2/\sqrt{3})$, and $\lfloor x \rfloor = \text{floor}(x)$. With this condition, a spatial cluster can be defined as any set of connected triggers. In reference to Graph theory, this definition is equivalent to taking all non-trivially connected components of the graph formed by all triggers. Figure 6.30 illustrates this scheme in an one dimensional example case.

Due to the transitivity of the connectivity condition, it includes trigger patterns extending over the whole SD array and is not limited only to crowns ≤ 4 . The choice of crown four is motivated by the requirement of being fully ‘‘T3 inclusive’’, i.e. detecting all possible T3 configurations. We will briefly compare our result with the official T3 implementation as a cross check in Section 6.3.3.

By construction, these clusters will contain patterns that are very likely originating purely from random coincidences, e.g. two triggers within crown four.

For each cluster, in the following called ‘‘spatial cluster’’, an event time t_{event} is defined as the average of the first and the last trigger time of the cluster. To account for imperfections of the grid of SD stations, i.e. (T2-)dead stations and missing stations, fake triggers are added for each dead or missing nearest neighbour of a triggered station. These fake triggers can be used in the selection of specific signatures. In addition, the average position $\langle \vec{x} \rangle$ of a cluster can be used to relate clusters from different parts of the SD array to each other.

Besides the cross-check of the T3 (c.f. Section 6.3.3), we use the clusters here as a veto for searching exotic signatures in Section 6.3.7. While it is possible to use this construction as well to search for coincidences of (exotic) semi-compact formations, the background is much larger than for the other analyses presented in this Section increasing substantially the difficulty to extract meaningful and statistically significant limits. We leave such searches for future work and focus on clear cases of rare but predicted signatures, like GZ events, and ‘‘clean’’ exotic signatures, e.g. super-preshowers.

6.3.2 Elementary Reconstruction Using Triggers

After discussing the construction of geometrically compact formations of clusters, we now focus on the timing irrespective of the geometry. The idea is to be able to reconstruct a plane shower-front moving with the speed of light solely based on triggers such that time coincidences of triggers over the whole SD array can be detected.

⁵as directly taken from the CDAS software, `XbArray.cc`




Figure 6.30: Schematic illustration of the spatial clustering analysis: the geometric distance, simplified in a one dimensional example, on the x -axis and the time difference on the y -axis. Red circles represent triggers for different stations, while black lines indicate the connectivity from Eq. (6.8).

In this Section, we present how we can calculate a plane shower-front analytically for a set of three stations, the minimal possible case for a fully determined geometry. In the following, we call such a (successfully) reconstructed plane from three triggers *triplet*. Later in Section 6.3.7 we will use these reconstructed planes in a Hough transform [154] to find shower-plane coincidences. The parameters of the triplet form the three-dimensional Hough space, which can be described with the variables (u, v, t_0) , where $(u, v, \sqrt{1 - u^2 - v^2})$ parametrises the (shower) axis \hat{a} and t_0 is taken as the (projected) arrival time at the centre of the SD array. The reconstruction of such a set of three stations is performed in the same way as the Offline seed reconstruction [155] and is summarised in the following.

With the first station (\bar{x}_1, t_1) taken as local reference point for position and time, the remaining stations define two projections with respect to the yet-unknown shower axis \hat{a} ,

$$\hat{a}(\bar{x}_i - \bar{x}_1) = c(t_i - t_1). \quad (6.10)$$

With the normalised vectors

$$\hat{i} = \frac{\bar{x}_{21}}{x_{21}} = \frac{\bar{x}_2 - \bar{x}_1}{|\bar{x}_2 - \bar{x}_1|}, \quad \hat{j} = \frac{\bar{x}_{31}}{x_{31}} = \frac{\bar{x}_3 - \bar{x}_1}{|\bar{x}_3 - \bar{x}_1|} \quad (6.11)$$

and times

$$\tau_{\{i,j\}} = \frac{c(t_{\{2,3\}} - t_1)}{x_{\{2,3\}1}} \quad (6.12)$$

the Eq. (6.10) can be written as a matrix equation

$$\begin{pmatrix} 1 & \hat{i} \cdot \hat{j} \\ \hat{i} \cdot \hat{j} & 1 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} \tau_i \\ \tau_j \end{pmatrix} \quad (6.13)$$

for the axis ansatz

$$\hat{a} = \alpha \hat{i} + \beta \hat{j} + \gamma (\hat{i} \times \hat{j}). \quad (6.14)$$

This system has a solution

$$D = 1 - (\hat{i} \cdot \hat{j})^2, \quad \alpha = \frac{\tau_i - \tau_j \hat{i} \cdot \hat{j}}{D}, \quad \beta = \frac{\tau_j - \tau_i \hat{i} \cdot \hat{j}}{D}. \quad (6.15)$$

The normalisation of the axis vector defines the absolute value of the parameter γ as

$$\gamma^2 = \frac{1 - |\alpha\hat{i} + \beta\hat{j}|^2}{1 - (\hat{i})^2}. \quad (6.16)$$

The sign has to be chosen such that the shower is coming from above, i.e. the axis vector is pointing upwards. If both solutions are from one hemisphere, e.g. $\hat{a}_z > 0$ for both $\pm\sqrt{\gamma^2}$, the solutions are not physical and the triplet is discarded. In other cases of incompatible timing with a plane front, Eq. (6.16) will yield a negative value for γ^2 . Such sets of stations are not considered in this analysis.

An origin for numerical instabilities in this reconstruction is the denominator of Eqs. (6.15). In the case of an aligned configuration of stations the value of D approaches zero. We only consider configurations of three stations with $D > 10^{-4}$ in the following.

The selection of appropriate triggers to attempt the reconstruction can significantly speed up the processing of the trigger data. In normal data conditions, i.e. no thunderstorms or communication problems, the mean T2-trigger rate of the whole SD array is about 34 kHz. If all possible combinations of three triggers from within a second were to be tested, $n = N(N-1)(N-2) \approx N^3$, the same combinatoric enhancement that makes the Hough transform work can lead to very high computational costs. The obvious choice of only considering a time window $t_c \approx 250 \mu\text{s}$ compatible with the speed-of-light transition time of the maximal distance in the SD array reduces the average number of T2s in the combinations to nine. This still leaves about 500 possible combinations of triplets within t_c per T2, while after the reconstruction only 2.9 triplets / T2 are observed. Thus, we employ a preselection scheme that requires $\Delta t_{ij} \leq r/c$ for all combinations of i, j when choosing three triggers to attempt reconstruction.

In the case of using this reconstruction to aid selection of “coherent” events from already geometrically preselected triggers, as done e.g. in Section 6.3.6, the computational load is greatly reduced and such preselections are of much less importance.

A caveat to this approach is the inefficiency at the edges of the phase space: there is no tolerance with respect to the timing, i.e. if $\Delta t_{ij} = r/c + \delta$ where δ is a small time jitter (cf. Fig. 5.46). Therefore, we expect some inefficiencies when reconstructing these planes which could be taken care of in numerical fitting. However, the computational cost of numerical minimisation is too large when using the full data set of triggers. Thus, we do not attempt such recovery of corner cases and have to account for this in the efficiency estimates.

6.3.3 Checking CDAS: Comparison with Recorded Events

With the availability of the complete T2 information in the T2Dump files it is now possible to perform a cross-check of the implementation of the T3 logic by using independent software. In this section, we discuss such a check on a small subset of the available data: June 2016. We choose a month in the austral winter to avoid complications caused by lightning triggers leading to lost data and making the identification of the corresponding triggers in the T2s more difficult.

Furthermore, in this cross-check we focus only on the SD-1500 and SD-750 arrays and in this comparison do not include special triggers, e.g. from doublet stations or FD ones.

We get the CDAS data sample of triggers by extracting the event times and station triggers from the IoSd-root files selecting all events within the GPS seconds 1148774400 and 1151366400 and having trigger names “TOT2C1&3C2”, “2C1&3C2&4C4”, or “BIG”. This leaves us with a total of 337 042 events.

On the T2 side of the analysis, we use the spatial clustering method, detailed in Section 6.3.1, to obtain clusters of triggers from the T2Dump files. Because the computation of

the crowns in the grid needs positions, and doublets and other special stations have to be excluded, we use the Offline station list⁶ and take all stations of grid 1 and 2 (SD-1500 and SD-750 arrays). To select the T3s from the clusters found by Eq. (6.8), we compute the crowns C_i for the cluster centred around each trigger. If any of these configurations fulfils either of the two T3 conditions, we promote the whole cluster to a T3-like event resulting in 330 733 clusters.

Investigating whether the extent of the discrepancy of about 2% is coming from actually missing events or rather from different splits in larger events, we match the two data sets. For all T3s we search for a cluster contained within 1 ms and proceed by matching the triggers included individually. In this way it is possible that two T3s are matched with the same cluster, which by construction is more inclusive.

However, we match only 330 753 T3s to the clusters, leaving 6289 T3s unmatched. Interestingly, many of the multiple T3s that are matched with a single cluster are repeated requests of the same event across the second boundary. As an example, the SD events 38071815 and 38071816 request the same stations (977, 976, 988, 1254) with an average microsecond of 999980.

Upon further investigation of the 6289 missing clusters, three contributions explain the discrepancy completely: First, the triggers of station 1764 do not appear in the T2Dump files, while the station was active. Nevertheless, in the later datasets, the triggers appear normally also for this station. 1092 events including this stations are therefore not found in the T2Dump files. Second, station 1847 is marked as “off-grid” for the analyses and was therefore not included in the clustering. However, it was still used for the formation of T3s in CDAS, leading to 4273 excess-events in the T3 data. Third, the so-called historic doublet stations are treated differently in the T3 algorithm than the treatment of newer doublets. While for the newer doublets one station is never considered in the T3 algorithms, for the historic doublets both can form a T3 together with other stations.

In our approach, all doublets are treated the same leading to 411 events missing because station *Rouge* (ID 185) of the *Moulin-Rouge* doublet triggered, but not the “master” station *Moulin*. Similarly, for the *Dia-Noche* doublet 510 additional events happen with *Noche* (ID 186) – but not *Dia* – triggering.

Thus, the summary of this comparison is that the official T3 implementation is never missing any event and if in doubt, e.g. due to stations deployed out-of grid, takes the conservative approach of including them. With this comparison we firmly establish the completeness of the current implementation with respect to its designed purpose without testing the physics or selection quality of the trigger. However, it is clear that this approach shows how new ideas or implementations for improving the triggers in AugerPrime can and should be tested.

6.3.4 Multi T3s

A simple yet powerful way of searching for multiple events arriving at roughly the same time is the time difference spectrum of consecutive events. With the assumption of a constant rate of events, the spectrum of waiting times $dN/d\Delta t$ follows an exponential distribution as known for a Poissonian process. This also means, that the smallest time differences are the most frequent, hence creating significant background. One possible strategy to counteract this is to use higher orders, i.e. asking for n additional events within a time window Δt between two events. Then the small times are suppressed with a factor λ^n as already pointed out in Eq. (5.1).

⁶SStationList.xml.in

In practise, the time-difference spectrum will be distorted due to “bad periods”, e.g. due to lightning when the data of the array is not usable. Another effect is the change of the event rate which occurs below full-efficiency threshold and is mostly related to the atmospheric pressure driving the overburden and attenuating the showers. For a non-constant λ the spectra will also look different than Eq. (5.1) predicts.

To take all these non-ideal conditions into account, and obtain a pure background expectation spectrum that removes correlations between events, we randomly shift all events in time by adding a random second $s \in [-5000, 5000]$ to the GPS second and a random value $u \in (-10^6, 10^6)$ to the GPS microsecond. By repeating this procedure multiple times and averaging the results we can obtain the average background expectation. The exact values for s and u are chosen arbitrarily, the important fact is only that they are larger than the correlation length of interest.

For our analysis we use the available reconstructed SD data set [126] from January 2004 to the end of January 2021. We apply the quality cuts for the ICRC 2019 data set [142] but reduce the requirement on core containment to a 5T5, as done for previous anisotropy studies [13, 156, 157]. These reduced requirements increases statistics at the expense of a reduced reconstruction accuracy [73]. Given that our aim here is to use the inclusive measurement and use timing rather than energy information, the use of these looser quality cut seems appropriate. We also use all events implied by this selection, starting at $\lg(E/\text{eV}) = 17.5$ which is clearly sub-threshold.

The left panel of Fig. 6.31 shows the distributions of Δt_{ij} for three example cases. In red, the usual case of consecutive events ($ij = 12$) is shown to highlight the deviations from the pure exponential behaviour due to the real data taking conditions. This behaviour can be accurately described with the scrambled data (black histogram).

As examples for higher order, we show the time difference for $n = 3$ (purple), i.e. $ij = 14$, and $n = 5$ ($ij = 16$, blue) together with the expectation from the pure random distribution (black). The aforementioned suppression of events with small Δt is clearly visible and no deviations from the expected random behaviour is found on time scales of seconds.

Furthermore, for very short time scales, which are in fact expected for GZ-events, we observe no events. The right panel of Fig. 6.31 highlights this by zooming into the first 500 μs of the Δt_{12} spectrum. There is a clear deficit of events with very small time differences, corresponding to the actual expected signal window. We can explain this deficit with the used trigger algorithm and the reconstruction chain.

The T3-trigger algorithm is using a 60 μs window to build the event request but also merges all clusters found within 100 μs . For large events, the trigger algorithm is directly using a 100 μs window. During the standard event reconstruction [155] the available data is not checked for the presence of a second event. Therefore, even in the case of two events recorded as one, only one, or none, are reconstructed with the current algorithms.

From Fig. 6.31 we can read-off the amount of events “lost” in this way: about 8 in 16 years of operation. It is important to note that these events are present in the recorded data. However, the selection and reconstruction strategy needs to be revised in order to restore and reconstruct such events. It is however clear from the low expected background of only 0.5 events per year that this search is promising. The optimisation of the event selection and the evaluation of the background associated with the identification of such events is beyond the scope of this work and a subject for future investigations.

6.3.5 Correlation of Events with Sub-Threshold Candidate Events

To circumvent the limitation of the DAQ and reconstruction chain to record and reconstruct events with very small time differences, we try to use the T2 triggers themselves to find




Figure 6.31: Time difference spectra of reconstructed SD-events. *Left:* The distribution of time differences for consecutive events Δt_{12} (red), third Δt_{14} (purple), and fifth Δt_{16} (blue) triggers respectively. The expectation from randomly shuffled events is shown in black for each of the spectra. *Right:* Distribution of times between consecutive events on a microsecond scale. The real data distribution is shown in red, while the black line and histogram are derived from scrambling the data to obtain an accurate expectation of random background.

the potential counterpart of a reconstructed event. From Section 5.6.1 we know that the SD array is sensitive to small showers mostly by two neighbouring stations triggering with a ToT-type trigger. Thus, we opt to use such pairs of triggers in this search for GZ (-like) events presented here.

We start from the same selection of reconstructed events as used in the previous Section, limited to the time range of years 2017 to 2020 due to the limitations of the availability of the T2Dump data. The T2 data from year 2016 have different trigger flags and will be used in the future for additional cross-checks. From each reconstructed event we use the GPS timing of the core and the reconstructed arrival direction. With this information it is possible to extract all clusters of neighbouring T2s following Eq. (5.9) in the GPS seconds of events. From these clusters, we select pairs of ToT-type triggers only to reduce background. Higher-multiplicity clusters satisfy the T3 condition and thus are recorded in the normal SD-data stream. An attempt to use these data was shown in the previous Section.

For each event, the pairs within ± 10 ms and from the whole SD-array are stored together with the events to allow for further selections and the estimation of random coincidences. We also compute the radial distance of the pair in the reconstructed shower plane, by taking the closest station (in shower plane) as a proxy, and using

$$r_{\text{sp}} = |\vec{r} \times \hat{a}| = |\vec{r}| \sqrt{1 - (\hat{r} \cdot \hat{a})^2} \quad (6.17)$$

to compute the shower plane distance r_{sp} from the distance $\vec{r} = \vec{r}_{\text{station}} - \vec{r}_{\text{core}}$ of the station from the core and the axis \hat{a} . We also compute the time difference to the expected arrival of the shower plane using

$$\Delta t_{\text{plane}} = t_{\text{trigger}} - t_{\text{core}} + \frac{1}{c} \hat{a} \cdot \vec{r}. \quad (6.18)$$

An additional requirement, motivated by the hypothesis of GZ (or similar) events originating far away and thus appearing parallel at Earth, is that the observed time difference of the candidate pair matches the one expected by the arrival direction of the reconstructed shower. Due to the limitations of the trigger time as estimator for the signal timing we require this to be strictly within $3 \mu\text{s}$,

$$\left| \Delta t_{\text{trigger}} - \frac{1}{c} \hat{a} \cdot \vec{r} \right| < 3 \mu\text{s}. \quad (6.19)$$

Furthermore, we need to remove the event data itself as it is contained in the constructed nearest-neighbour pairs due to the redundancy of the T2 data with respect to the recorded event data. For all pairs with

$$|\Delta t_{\text{sp}} - 13 \mu\text{s}| < 25 \mu\text{s}, \quad (6.20)$$

we check the *candidate* stations of the event and remove the pair if one of its stations is candidate station in the event.

An additional complication arises from the fact that the reconstruction [126] is not using the ToTd and MoPS triggers and thus is not marking stations with only those triggers as candidate stations even though they belong to the event. As an conservative approach to reduce the influence of such pairs, we additionally remove all pairs that are nearest neighbour to any candidate station. In this removal step, we do not use any additional time selection than Eq. (6.20) to account for timing tolerances introduced e.g. by shifting latch bins in the new triggers. Such a selection will as well remove some potential GZ events with small distance between the events, but as there is so far no observation of such events at all, our goal is to get a sample as pure as possible to first determine if such events are seen by the Auger SD.

We select a set of *candidate* events by using a time window around the expected arrival time and use only pairs with a distance of less than 5 km to keep the background low. The time selection can be formulated as

$$|\Delta t_{\text{sp}} - t_c| < t_w + \frac{r_{\text{sp}}}{c}, \quad (6.21)$$

with the central time $t_c = 13 \mu\text{s}$. The choice of the width of the interval $t_w = 12 \mu\text{s}$ is motivated by the length of the trace after the latch bin but the results here do not change significantly if a slightly different window is used. To avoid penalisation from scanning, we use only one value of t_w .

The left panel of Fig. 6.32 shows this selection together with the scatter plot of the preselected pairs. An excess at small distances is clearly visible, while in total in the shown window ten events are observed – nine of which in the observation window of $r_{\text{sp}} < 5 \text{ km}$. Using randomly selected $|t_c| > 30 \mu\text{s}$, we can determine the chance probability of observing this amount of pairs in pure background. The normalised distribution dN/dn_{observed} in background is shown in the right panel of Fig. 6.32 together with a Poisson distribution with the mean $\lambda = 0.66$ corresponding to the mean number of events in such random search windows. It is very clear that we observe more than background with this search.

Due to the low number of candidates and to exclude effects we did not yet account for, e.g. dead-stations or imperfections of the SD-grid, we visually scan the candidates. Fig. 6.33 shows three of the nine candidates together with information on which stations were not sending T2s at the time of the event (dead stations). The information for the other candidates are summarised in Table 6.1. It is clearly visible in the central panel of Fig. 6.33 that one event is very likely split due to missing stations in the grid. For three other events there is a single dead station in between the selected pair and the reconstructed event. The right panel of Fig. 6.33 shows one example of such events. As representative of the remaining candidates, we show event 170922876600 in the left panel. For all candidates, the separation is fairly small and there are minor hints that the separation decreases with increasing energy, as visible in Fig. 6.34.

For the interpretation of the five observed candidates, it is also necessary to take non-random backgrounds into account that are not removed by the visual selection. The only known effect contributing is the splitting of real events by stations affected by dead time (cf. Fig. 5.5) due to random triggers before the event. In the following, we try to estimate the contribution of this effect on the observed number of candidates.

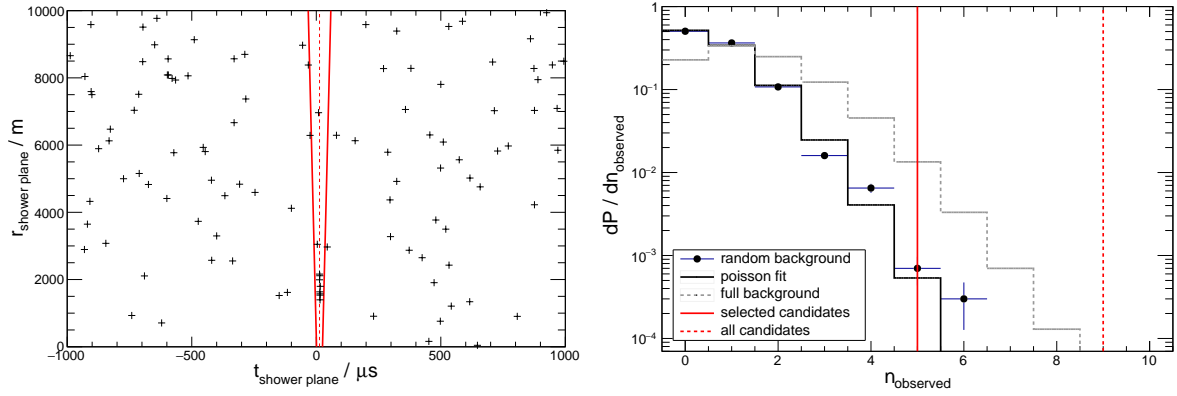


Figure 6.32: Correlation of nearest-neighbour trigger pairs with reconstructed events. *Left:* Distance of the pair to the shower core in the shower plane as a function of the relative timing to the shower plane. The red dashed line indicates the expected arrival at 13 μs due to the nominal position of the signal in the trace at bin 250 out of 768. The search window for the correlation is indicated by the solid red lines corresponding to a 12 μs -wide window adding tolerance based on the distance as r/c . *Right:* Distribution of the number of pairs in the search window if the search window is centred randomly within $[-9, 9]$ ms and excluding the signal region (black points) and pairs more than 5 km away from the core. A Poisson distribution with the mean $\lambda = 0.66$ is shown as black line indicating the expectation for a random uniform background that is sampled with the different search windows. The observed number of events in the signal window of 9 is highlighted as dashed red line, while the remaining 5 candidate events after visual scanning are shown as solid red line.




Figure 6.33: Visualisations of the split event candidates from Table 6.1. *Left:* True candidate event from Auger-Event 170922876600. *Middle:* Removed candidate because the influence of the missing stations cannot be easily estimated. It is Auger-Event 191447249800 reconstructed with $\lg(E/e\text{V}) = 18.1$. *Right:* Removed candidate due to a dead station connecting the reconstructed event with the pair (Auger-Event 203582210000).

In general, the probability of a station to be “in dead time” for any given event can be estimated using the known rate of T1s of 100 Hz and the necessity that two T1s need to happen within 220 μs of the event. Thus, we can use the Poisson probability with the mean

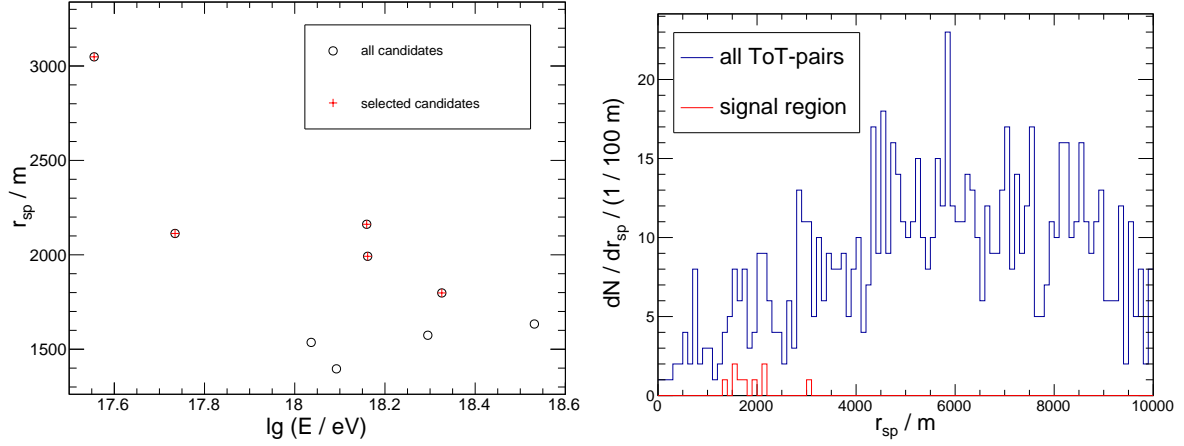


Figure 6.34: Characteristics of the observed event candidates. *Left:* Correlation of the distance in shower plane r_{sp} with the energy of the reconstructed event. The candidates left after visual selection are marked with red crosses. *Right:* Distribution of r_{sp} in the full extracted set of nearest-neighbour pairs of ToTs for all time differences within 10 ms. The signal region is concentrated on short scales, independent of the used cut $r_{sp} < 5000 \text{ m}$.

rate $\lambda = \Gamma_{T1}\Delta t = 2.2 \times 10^{-2}$ to get the probability of seeing two or more events before the events as

$$p(\text{dead}) = p(n \geq 2|\lambda) = 1 - e^{-\lambda} (1 + \lambda) \quad (6.22)$$

$$= 1 - (1 - \lambda + \frac{\lambda^2}{2})(1 + \lambda) + \mathcal{O}(\lambda^3) \quad (6.23)$$

$$= \frac{\lambda^2}{2} + \mathcal{O}(\lambda^3) \quad (6.24)$$

$$\approx 2.4 \times 10^{-4}. \quad (6.25)$$

To split an event into a reconstructable part and a pair of ToTs, we need at least six triggers and the station that is not triggering due to dead time needs to be geometrically favourable. We therefore estimate the resulting rate as

$$\Gamma_{\text{dead}} \approx \Gamma(n \geq 6) \gamma_{\text{geo}} p(\text{dead}) \quad (6.26)$$

where $\Gamma(n \geq 6)$ is the rate of events with more than five stations and γ_{geo} is a factor describing geometry. We estimate γ_{geo} by selecting all events with more than five stations in a representative sample of data and try all combinations of removing a station. By keeping track of how many of the configurations fulfil the conditions that two clusters of nearest neighbours exist with one of them having at least three stations and the other one two ToT triggers, we obtain $\gamma_{\text{geo}} = \frac{1020}{50282} = 2.03 \times 10^{-2}$.

Fig. 6.35 illustrates this search by showing an example footprint of a six station event that can be split if one station is removed. The additional examples show the same principle for five-station events to be split in two pairs, as it is necessary for the background estimation in Section 6.3.6. The right panel of Fig. 6.35 shows the differential rate of events with their nearest-neighbour multiplicity. We used the year 2019 to create the data sample by taking five days from each month to keep the computational costs low. The event rate in the same sample evaluates to $\Gamma(n \geq 6) = 1.40 \times 10^{-3} \text{ Hz}$, as obtained by integrating the histogram in Fig. 6.35 for $n \geq 6$.

Thus, we obtain an overall estimate $\Gamma_{\text{dead}} = 6.8 \times 10^{-9} \text{ Hz}$. Taking the bad periods of data taking into account, we get an expected number of candidates from this source of




Figure 6.35: Illustration of the effect of stations splitting events because of dead time. *Left:* A possible footprint of a six station event (blue dotted circles) that can be split in a three station event plus a pair of nearest neighbours by removing a station. Additionally, two possible footprints of a five station event are shown for the discussion in Section 6.3.6. The solid red lines mark a configuration that cannot be split into two nearest-neighbour pairs by removing one station. For the dashed black configuration such a split is possible by removing the central station. *Right:* Differential rate of nearest-neighbour multiplicity n_{T2} in the representative data sample (cf. text). The integrated rate for $n \geq 6$ is $\Gamma(\geq 6) = 1.4 \times 10^{-3}$ Hz.

$N(\text{dead}) = 0.84$ for the four-year data sample. Adding this background to the random coincidence one, we have a total background level of $N(\text{total}) = 1.50$ for our sample. The corresponding Poisson distribution is shown as grey dashed line in the left panel of Fig. 6.32. We want to stress that our construction of γ_{geo} is on purpose conservative by not selecting only those high-multiplicity clusters that fulfil all requirements for the official reconstruction.

Evaluating the probability of observing five events with this background gives a 2.2σ deviation from the expectation. Using the Feldman-Cousins construction [130] of confidence intervals or upper limits, we get a signal interval of $\mu = [0.9, 8.5]$ at 90% C.L. or an upper limit of 12.2 at 99% C.L.

To compare this number of events with models, we need to estimate the exposure for these events. Because with these events we operate below full efficiency of the array, it is difficult to obtain a good estimate of the exposure as it will depend critically on the zenith-angle distribution, primary composition, efficiency modelling, and the energy distribution in split events. In this work we aim at providing a first estimate of the integral exposure and leave the detailed calculation for future work with more theoretical input. Therefore, we will use the following to estimate the rate of GZ events, interpreting all candidates as possible GZ-signal contributions

$$J_{\text{GZ}} \approx \frac{n_{\text{sig}}}{T \hat{\varepsilon} A_{\text{hex}} \Omega}, \quad (6.27)$$

where $\hat{\varepsilon}$ is an effective efficiency, T the time interval, and A_{hex} the area covered with working hexagons as used for the energy-spectrum estimation [3]. The idea here is to take out the energy integral by defining

$$\hat{\varepsilon} := \int_{E_0}^{\infty} dE \hat{\phi}(E) \varepsilon(\alpha E) \varepsilon_{\text{pair}}((1 - \alpha)E) \quad (6.28)$$

with the spectral shape $\hat{\phi}$ used as weighting function, the SD efficiency ε , the efficiency for measuring a shower as a pair of ToT-triggers $\varepsilon_{\text{pair}}$, and the energy fraction carried by main remnant αE . In most cases, due to the strong boost of the reaction, we expect $\alpha = (A - 1)/A$ for the emission of a single nucleon. Because we do not want to exclude the possibility of




Figure 6.36: Estimation of the effective efficiency for the detection of a split event using Eq. (6.28). *Left:* Varying the minimal energy for three different values of energy fraction transferred to the emitted nucleus α . Vertical lines indicate the minimal energy of the primary particle such that the leading remnant has more than $\lg(E/\text{eV}) = 17.5$ to satisfy the required minimal reconstructed energy condition. *Right:* Dependence of the effective efficiency on the energy distribution in the splitting α . We use two minimal energies close to the actual threshold.

non-dominant processes, e.g. two-nucleon emission, and to account for different isotopes we treat α in the following as a “free” parameter.

To model the efficiency, we use a parametrisation of the efficiency for oxygen primaries [158] modelled as

$$\varepsilon(\lg E) = \left(1 + \exp \left(-4 \frac{\lg(E/E_0)}{w} \right) \right)^{-1} \quad (6.29)$$

with parameters $\lg(E_0/\text{eV}) = 17.87$ and $w = 0.45$ derived from EPOS-LHC simulations with 38° zenith angle. We parametrise the efficiency for measuring a pair of ToTs with a Gaussian fit to Fig. 5.47 and obtain

$$\varepsilon_{\text{pair}} = (0.37 \pm 0.01) \exp \left(-\frac{1}{2} \left(\frac{\lg(E/\text{eV}) - (17.43 \pm 0.02)}{0.37 \pm 0.02} \right)^2 \right). \quad (6.30)$$

As simple model for the flux of primary cosmic rays we use $\phi(E) \propto E^{-3}$ leading to

$$\hat{\phi}(E) = 2E_{\min}^2 E^{-3}. \quad (6.31)$$

The minimal energy E_{\min} is given by the choice of the integration interval for the limit. However, since we neglected the effect of the required minimal SD energy in the reconstructed input sample defined in Section 6.3.4, a natural choice is the minimal reconstructed energy of $\lg(E/\text{eV}) = 17.5$. Due to the resolution of the reconstruction as well as the energy bias below full efficiency threshold the actual minimal energy can differ slightly from this threshold.

Fig. 6.36 shows the numerical result for $\hat{\varepsilon}$ for a set of parameters. We show results for equal energy splitting $\alpha = 0.5$, a splitting following a superposition model with $\alpha = 15/16$ and $\alpha = 0.9$. The resulting efficiencies range from about 0.005 to 0.03 with a mean value of about 0.02.

To highlight the high uncertainties of this estimation, we additionally try to estimate $\hat{\varepsilon}$ based on actual predictions of the energies of GZ events. We take the predictions from Ref. [43] – shown in the left panel of Fig. 6.37 – and, due to the lack of a low-energy cut-off, change the efficiency of the main array to

$$\varepsilon(\lg E) = \left(1 + \exp \left(-4 \frac{\lg(E/E_0)}{w} \right) \right)^{-1} \left(1 + \exp \left(-4 \frac{\lg(E/\text{eV}) - 17.5}{0.25} \right) \right)^{-1} \quad (6.32)$$




Figure 6.37: Using the predicted flux of GZ events for weighting the calculation of the effective efficiency. *Left:* Predicted total flux for all primaries to be considered either iron or oxygen (taken from Ref. [43]). *Right:* Result of Eq. (6.28) if the fluxes from the model are used to get $\hat{\phi}$ (cf. Eq. (6.31)). For the iron primaries, we change the parametrisation of the efficiency to the one of iron [158]. In addition, we add a model of the cutoff due to the requirement of $\lg(E_{\text{rec}}/\text{eV}) > 17.5$.

to include the efficiency of the energy requirement. The resulting efficiency as a function of the energy fraction carried by the main remnant is shown in the right panel of Fig. 6.37. Because the efficiency of detecting iron showers is higher and the distribution of GZ events for iron includes higher energies, the resulting effective efficiency is different by about a factor three.

The other factors in Eq. (6.27) are easier to calculate, as in our approximation the measurement time is simply the amount of time of data taking save for “Bad Periods”. In our data sample this evaluates to $T = 1.237 \times 10^8 \text{ s} = 3.92 \text{ years}$. Averaging the number of active hexagons that can be used to detect the main shower of the array over the measurement time, we obtain $A_{\text{hex}} \Omega = 5756 \text{ km}^2 \text{ sr}$ as instantaneous exposure.

Using Eq. (6.27) we thus obtain the flux of GZ events with separation of roughly more than 2 km

$$\frac{dJ}{d\Omega dA dT}(r \gtrsim 2 \text{ km}) = (3.5_{-3.0}^{+6.3}) / (\hat{\epsilon} 3.92 \times 5756) (\text{km}^2 \text{ sr y})^{-1} \quad (6.33)$$

$$\approx (1.6_{-1.3}^{+2.8}) \times 10^{-1} (\text{km}^2 \text{ sr y})^{-1} \quad (6.34)$$

using $\hat{\epsilon} = 10^{-3}$ to be on the conservative side. Using the 99% C.L. upper limit for the number of signal events, we get $5.4 \times 10^{-1} (\text{km}^2 \text{ sr y})^{-1}$ as the upper limit. For comparisons with predictions, the underlying distribution of α has to be known, as Fig. 6.37 makes very clear. Nonetheless, we can compare our observation with predictions from Ref. [43] by using J/J_{pred} as measure, scaling the rate shown in Fig. 6.37 with $T/\text{years} A_{\text{hex}}/(2\pi 3000) \text{ km}^2 \text{ sr}$ and accounting for efficiency based on Fig. 6.37. We use both the 90% C.L. intervals and the 99% C.L. upper limit and divide the predicted flux by a factor two to roughly account for the part of the events that are not separated enough to be picked up by this analysis. The resulting ratios are shown in Fig. 6.38. More work on the details of the predictions and the understanding of the efficiencies is necessary to draw any firm conclusions from this work.

Furthermore, systematic effects, e.g. the dependence of the true exposure on the separation of the fragments, are not included in the estimates presented here. We assume that these effects are small compared to the detection efficiency, as shown in Fig. 6.37, but nonetheless should be revisited moving forward after this first estimate presented here. Other effects that could explain split events like rare first interactions that emit a proton with high- p_T are




Figure 6.38: Comparison of the obtained rate with the predictions from Ref. [43]. The different markers indicate different assumptions for the efficiency depending on the energy fraction carried by the main remnant. The left-most points per primary correspond to using the maximal efficiency from Fig. 6.37, with the predicted flux divided by two to account for the necessary separation. Next to it, we use the full predicted flux, while for the right-most points we use the efficiency corresponding to $\alpha = (A - 1)/A$ and keep the factor two for the separation.

Table 6.1: List of candidates of ToT pairs adjacent to reconstructed events. Every pair-event combination that fulfils Eq. (6.21) is listed here. After the visual scan, four of the nine candidates are removed due to either dead-stations or imperfections of the SD grid. ψ_{sun} is the angular distance between the reconstructed arrival direction and the Sun. We calculate it only for the candidate events.

GPS second	Auger-event Id	$\lg(E/\text{eV})$	$\theta/^\circ$	$\psi_{\text{sun}} / ^\circ$	comment
1175198381	170922876600	18.2	54.5	111	candidate
1214909940	181868272500	18.3	58.5	131	candidate
1221681158	182652834300	17.6	33.3	63	candidate
1242806913	191447249800	18.1	55.4	-	hole in SD array
1251403185	192442877000	17.7	42.8	39	candidate
1261432745	193603593000	18.2	43.8	112	candidate
1271370603	201103778800	18.5	54.0	-	dead-station
1292782115	203582210000	18.0	55.7	-	dead-station
1293327490	203644907500	18.3	55.2	-	dead-station

beyond the scope of this work but have to be considered if the excess persists with more data available.

6.3.6 Search for Sub-Threshold GZ Events

After discussing the search for split events with one part below the SD threshold, we try to extend this search now to the regime where both showers are below the threshold. We start from a similar starting point as in the previous section: we construct clusters of nearest-neighbour triggers with Eq. (5.9). We change the range of data to all GPS seconds, from hours without detected lightning strikes in the data set in Section 5.6.3. The requirement of having the lightning veto restricts the data range from January 2017 to the end of April 2020 due to the availability of the lightning data in question for this analysis. In total, 1452 hours of data are excluded, leaving 27732 hours of data to be analysed. For future extensions, the use of the SD data as self-veto should be considered.

From the selected clusters, we additionally select only those with two ToTs to obtain a clean data sample without overlap with Section 6.3.5. From the resulting sample of 4.2 million pairs, equalling a rate of 4.2×10^{-2} Hz, we select pairs that are coincident by requiring

$$\frac{t_1^{(j)} + t_2^{(j)}}{2} - \frac{t_1^{(1)} + t_2^{(1)}}{2} < \frac{r}{c} + 5 \mu\text{s} \quad (6.35)$$

with the trigger times of the i -th pair $t_j^{(i)}$ (with $j \in 1, 2$) and the distance r between the centre of mass of the pairs as criterion of coincidence. Applying Eq. (6.35) on the data leaves us with 2639 pairs, corresponding to a suppression factor of 6×10^{-5} .

To further narrow down the sample to only those pairs of pairs that are compatible with a single shower front, we use the elementary reconstruction from Section 6.3.2 on the triggers contained in the selected coincident pairs. Requesting that all combinations of triplets formed by the four triggers are reconstructable, without posing any condition on the reconstructed directions themselves, leaves us with 245 *candidates*. To highlight why we do not consider the direction reconstructed in this way a good estimator, we show in Fig. 6.39 two traces from one of the remaining candidates which is also a T3. It is clearly visible that for the station triggered by the new trigger ToTd, the signal starts μs early with respect to the expectation from the latch bin. This means that the reconstructed direction will be biased if a better timing is not available from the traces, and in extreme cases, the event might not be selected. An accurate estimation of how much inefficiency is introduced with this cut is difficult to obtain because it depends on the detailed modelling of the events and the triggers. For example, in the candidate used in Fig. 6.39, the event passes because for both pairs the trigger types within each pair are similar leading to an offset in trigger timing between the pairs but no rejection with the reconstruction. We leave the detailed discussion of this systematic uncertainty of the efficiency for future work as it is beyond the scope of this work.

Among the selected candidates, a large fraction are normal showers split by a non-working stations or border effects due to e.g. missing stations in the grid. To filter out events affected by non-working stations we generate lists of non-working nearest-neighbour stations \mathcal{N} for both pairs and require the intersection $|\mathcal{N}_1 \cap \mathcal{N}_2| = \emptyset$ to be empty. With 185 candidates removed (75.5%), with one example given in the left panel of Fig. 6.40, this condition removes most of the candidates. However, some influence of higher orders remain, as the right panel in Fig. 6.40 shows. To remove also these effects, we in addition request that the set of nearest-neighbours of the dead stations does not have an overlap with the dead nearest neighbours of the other pair. Only three more events are rejected with this condition. One of rejected events, shown in the left panel of Fig. 6.41, is likely a false positive rejection and should be considered candidate instead. To be conservative, we do not use it as a candidate.

The next class of events, very likely arising from normal showers that get split by detector effects, are those hitting missing stations in the SD grid. One example of such a candidate is shown in the right panel of Fig. 6.41. To filter out these events, we request that at least




Figure 6.39: Traces from the split-event candidate (two pairs) at GPS second 1241378968 and the corresponding SD event 53342934 to highlight the uncertainty associated with the time tag given of the T2. *Left:* Average trace of station 241 using the online calibration values. The red line marks bin 250, approximately where the latch bin is expected. The trace is triggered by a ToT trigger. *Right:* Trace of station 219 with the trigger time displayed in the same way. The shift of the signal with respect to the expected start time is clearly visible. This trace was triggered by a ToTd trigger.




Figure 6.40: Illustration of events split by dead-stations which are removed from the final candidate set. *Left:* The stations in between the two pairs are dead and thus create the signature of split events instead of a single connected footprint. In this example the stations (410, 405) and (416, 414) form the signal pairs in GPS second 1167855849. A total of 185 events like this are removed from the data set due to the presence of dead-stations. *Right:* Example for the rejection of a second order dead-station effect. The pairs (1613, 1610) and (1644, 1639) in GPS second 1168142320 are connected with two crowns of dead-stations making a large event a likely explanation for the observed footprint.

three of the four stations have a full hexagon surrounding them. Another 37 candidates are removed by this condition.

The remaining 20 candidates are visually scanned to take care of split events caused by a station triggering on random muons (or small showers) just before the event of interest. These splits will mostly occur with five station events with the central station triggering early.




Figure 6.41: Further examples of rejected events. *Left:* The dead nearest neighbours of the two pairs are nearest neighbours to each other and therefore lead to the rejection of the candidate. While we believe it is unlikely that these stations influence the real footprint, we remove this candidate anyway to be conservative. *Right:* Example of an event removed because all stations border a hole in the SD-grid. In total, 37 such events are removed.

We can estimate the expected rate of these splits similarly to Eq. (6.26) from the previous section. Adjusting it yields

$$\Gamma_{\text{early}} \approx \Gamma(n_{T2} = 5) \gamma_{\text{geo}} p(\text{early}) \quad (6.36)$$

where $p(\text{early})$ describes the probability of a station to be triggered early by a random small shower. We follow the same procedure as introduced in the previous section to get $\Gamma(5)$ and γ_{geo} . From the right panel of Fig. 6.35, we can read off $\Gamma(5) = 2 \times 10^{-3}$ Hz. The geometric factors, as illustrated by the two footprints with five stations in the left panel of Fig. 6.35, yield 0.01. The fact that this value is lower than for the higher-multiplicity events discussed in the previous Section, can be explained by the existence of many compact footprints which cannot be separated into two pairs. The footprint marked with red solid lines in Fig. 6.35 is an example for such a configuration.

We can obtain an estimate for $p(\text{early})$ by calculating the probability of getting one T1 in a time window that still allows the signal to be present in the trace. Given that there are about 500 bins of trace after a typical trigger at bin 250, we use $\tau \approx 25 \times \text{ns} \times 500 = 12.5 \mu\text{s}$ as time window. Thus, we get

$$p(\text{early}) = p(N = 1 | \lambda = \Gamma_{T1} \tau) \quad (6.37)$$

$$= \lambda e^{-\lambda} \approx \lambda \quad (6.38)$$

$$= 1.25 \times 10^{-3} \quad (6.39)$$

using the rate of T1s $\Gamma_{T1} = 100$ Hz. In total, we get $\Gamma_{\text{early}} = 2.5 \times 10^{-8}$ Hz, resulting in an expected number of candidates of 2.5 in the total data set. The exact number is not important because we are able to filter out these effects by looking at recorded data since all such events fulfil the T3 condition. Therefore, we also do not need to account for uncertainties in this estimate in the calculation of limits or signals.

Performing the visual scan, we remove five such events, in agreement with Poisson statistics and our estimate. Table 6.2 lists these events, together with their SD event ID as “early random trigger”. Fig. 6.42 shows one example of these removed events with the footprint on ground in the left panel and the trace data in absolute time scale in the right




Figure 6.42: Visualisation of SD event 45731478 which is selected as candidate for a split event into two ToT pairs because the central station (1410) was triggered early by an accidental small shower. *Left:* Positions of the stations of the recorded SD-event. In the analysis presented here, stations 1407, 1409, 1401, and 1524 form the two pairs of ToTs that were selected. The red cross marks the station with the earliest (true) signal. *Right:* Trace data from the stations scaled such that the maximal peak is at the same amplitude. The integrated signal is colour coded. The y-position of the traces is sorted by the distance of the stations to the first signal station (1407). Three stations are triggered by accidental muons and do not match the geometry and timing of the shower (1804, 1514, 1502).

panel. Clearly, the central station has a signal belonging to the shower while the trace is shifted because of a (barely visible) early random muon signal.

In addition, we remove four other candidates because they happened in periods with some lightning events. These events make it through the lightning rejection cut because the used lightning detection method is not 100% efficient. In the case of three of these four events we have to rely on events occurring close in time because the candidates themselves do not fulfil a T3 condition and thus are not recorded.

Finally, the eight remaining candidates we list in Table 6.3, selecting only those with separation less than 5.5 km. For these candidates we have to statistically show whether this number is above the expectation for random background or not. The value 5.5 km is motivated by the cut used in the previous section, the predictions of mostly small separations [43], and the grid structure. 5.5 km corresponds to twice the smallest possible distance of pairs that are not nearest neighbours plus some additional tolerance.

To estimate the random coincidence rate, we create random sets from the actual data. Before the coincidences are evaluated, we add a random shuffle time $t \in [-5000, 5000] \mu\text{s}$ to each pair to destroy coherence between the pairs. The rate obtained in this way will be slightly overestimated since the pairs from true events, split by e.g. non-working stations, increase the overall rate of pairs. However, the influence is very small because the rate is steeply falling with multiplicity (cf. Fig. 6.35 (right)) and we chose to include it here to be conservative in the final estimate of the background. For each random data set the same analysis and selection is performed, except for the visual scanning because the effects filtered there are not of a random nature and thus do not appear in the shuffled data sets.

Using 50 runs of shuffled data we get an average number of candidates in the selected radial range of $N_{\text{rand}} = 1.28 \pm 0.16$.

The other contribution to the number of candidates that are not random and not filtered out in visual scanning are those events split by stations affected by dead-time. The discussion




Figure 6.43: Illustrations of the observed excess of coincident pairs of neighbouring ToT-pairs. *Left:* Distribution of the distances r between coincident pairs and the time difference Δt . Orange circles are events from the signal sample while blue squares are from randomized samples that have 10 times the exposure. Lines indicate limits in the selection: compatibility with speed of light propagation (inclined solid line), minimal distance of non-nearest neighbours (dashed line), and maximal allowed distance for the signal window (vertical solid line). *Right:* Effective efficiency for detection of GZ events using the predicted spectra from Ref. [43] and Eq. (6.42). The expression is symmetric around $\alpha = 0.5$ because both parts are detected in the same way so only the upper part is drawn.

is completely analogous to the one in Section 6.3.5, and yields, with the values of $\Gamma(5)$ and γ_{geo} introduced earlier,

$$\Gamma_{\text{dead}} \approx \Gamma(5) \gamma_{\text{geo}} p(\text{dead}) = 4.8 \times 10^{-9} \text{ Hz.} \quad (6.40)$$

With the exposure time $T = T_{\text{tot}} - T_{\text{lightning}} = 1.02 \times 10^8 \text{ s}$ this gives us $N_{\text{dead}} = 0.49$.

With a Poissonian model this corresponds to 3.34σ without accounting for systematic uncertainties. Following Ref. [159], we can extract a 90% C.L. interval for N_{sig} as $N_{\text{sig}} \in [2.91, 12.32]$. The conversion of this rate to the rate of such events follows closely Eq. (6.27), as

$$J = \frac{N_{\text{sig}}}{\epsilon T A \Omega}. \quad (6.41)$$

For a first estimate within the scope of this work we use the area covered by 5T5 and 6T5 hexagons, averaged over the time period of the dataset, resulting in $A = 1875 \text{ km}^2$. As the solid angle Ω , we use the half sphere 2π because no explicit cut on the directions is introduced and all zenith angles up to $\theta = 90^\circ$ are allowed in the elementary reconstruction. With these numbers we arrive at $5_{-2}^{+5} \times 10^{-12} / \epsilon \text{ (s sr km}^2\text{)}^{-1}$ as flux of such events, still depending on the efficiency ϵ for detection. Adjusting Eq. (6.28) from the previous section, as

$$\hat{\epsilon} := \int_{E_0}^{\infty} dE \hat{\phi}(E) \epsilon_{\text{pair}}(\alpha E) \epsilon_{\text{pair}}((1 - \alpha)E) \quad (6.42)$$

we obtain the efficiencies shown in the right panel of Fig. 6.43 for the predicted spectra of GZ-events from [43]. In these efficiencies the effects of the required geometric splitting are not included. We expect about a factor of two since the minimal separation in the ground plane is about the same as in the previous Section and larger than the majority of the predicted flux in Ref. [43] at about 1 km.

While interpretations of this excess as GZ events are possible, further investigations of systematic effects which are beyond the scope of this work are necessary before drawing any firm conclusions. Among the possibilities which are not included are effects of stations that

have GPS problems. Such stations would not appear as “non-working” because they still send triggers, however, they do not participate in events because the timing with the real events is not matching. The frequency of such GPS problems is unknown. Additional (yet unknown) dead-times could contribute as well, but we consider unknown dead-times to be an unlikely explanation of such excesses.

Rare effects of stations stopping to work in the time between sending the triggers and having an early T1 should appear as error stations in the events. We do not observe any such events but did not explicitly search for them either. Nevertheless given that for most candidates we have an event, we expect such cases would be found if they were to contribute significantly.

On the analysis side, there are uncertainties associated with the background estimates. However, the necessary factor four to explain the complete signal as background is unlikely to be an effect of mismodelling especially since most factors are estimated directly from the data, including γ_{geo} , the rate of observed events $\Gamma(5)$, and the random coincidence rate.

To complete the analysis, we use the so-far not used radial range of $r > 5.5$ km to put a limit on split events with such separation. Using the same data samples as used for the nearby part, we get a random background expectation of $N_{\text{bg}} = 4.1 \pm 0.3$ events and observe $N_{\text{obs}} = 3$. Thus, we can put a limit using the Feldman-Cousins construction [130] at 90% C.L. of $N_{\text{sig}} < 3.47$. Converting this limit to a flux with Eq. (6.41), using $\varepsilon = 1$ so that it can be dealt with in the context of concrete models, we obtain

$$J_{\text{split}}(r > 5.5 \text{ km}) < 2.9 \times 10^{-12} (\text{s sr km}^2)^{-1} \text{ (at 90\% C.L.)}. \quad (6.43)$$

The data sample used in Eq. (6.43) does not contain any data beyond 10 km due to the requirement of four elementary reconstructions, as described before. To generalise the analysis to cases when minor time delays or deflections cause the two pairs to not reconstruct to a common plane, we remove the aforementioned condition and analyse the set of data with $r > 5.5$ km again. Therefore, the only condition applied here for the coincidence is Eq. (6.35) and we select only those clusters with exactly two pairs fulfilling the condition.

Fig. 6.44 shows the resulting distribution of observed distances in randomised samples as well as the signal sample. We do not perform additional selections based on dead or missing stations because at the distances considered here these effects are negligible. However, to exclude residual effects from lightning, as seen in the previous part, we remove all candidate pairs that occur in hours when at least one SD event had a lightning station. 62 out of 2617 pairs are removed by this condition. To convert the numbers of events observed in signal and in background to limits on non-random contributions to such coincident pairs, we use again the Feldman-Cousins construction of confidence intervals at 90% C.L.. Due to the large separation, we neglect effects from dead-time affected stations, leaving only random contributions as background.

We show the constructed limits in the right panel of Fig. 6.44, with one bin $15 \text{ km} < r < 20 \text{ km}$ having a confidence interval rather than a limit due to a 2σ upwards fluctuation. As reference, we construct the expected limits and their standard deviations from Monte-Carlo simulations and show them together with the observed limits. All limits agree with the expectation at 1σ level, while the 2σ outlier is expected too.

To convert these limits on the observable numbers of events n_{obs} to fluxes, the exposure $\langle A \Omega \rangle = 2\pi \times 1875 \text{ km}^2$ and the observation time $T = 1.02 \times 10^8 \text{ s}$ can be used. The exact efficiency for separations on the scale of the full SD grid depends on the exact configuration and can be calculated if necessary from the information on the grid structure. Since the detection efficiency is also important for the determination of the flux, we leave this estimation for future work.

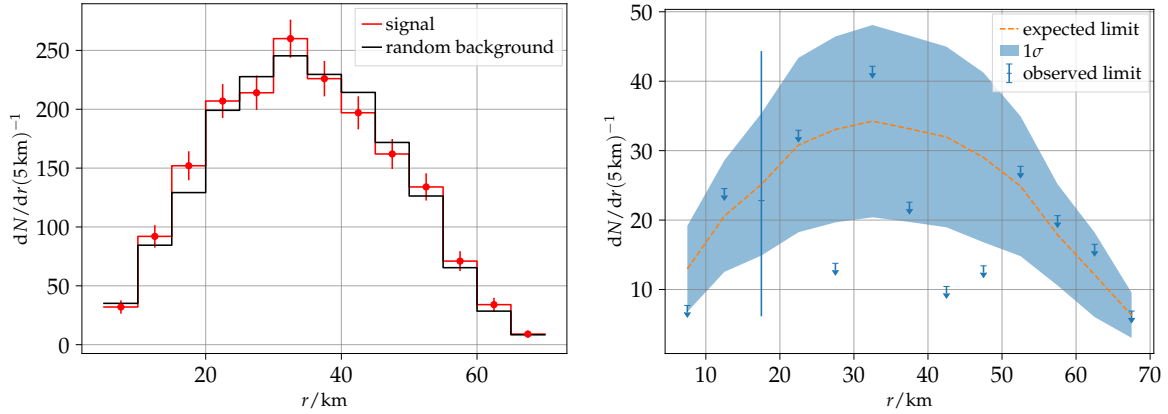


Figure 6.44: Extension of the search for split events to arbitrary distances between the events. The requirement on the reconstructability of plane fronts with the coincident pairs is removed to increase statistics. *Left:* Number of expected coincidences from random background (black) and observed counts for distance bins of 5 km. In the first bin, only $r > 5.5$ km are included not to generate overlap with the analysis of short distances. *Right:* Conversion of the observed number of events to 90% C.L. upper limits with the Feldman-Cousins construction. The expected limit with an 1σ band is drawn for reference. For $15 \text{ km} < r < 20 \text{ km}$ we get a confidence interval rather than an upper limit. Given the number of bins, a single 2σ fluctuation like this is not considered significant.

Table 6.2: List of candidates of two ToT pairs removed after the visual scan (cf. text for details).

GPS second	SD-event Id	reason
1168713211	41184617	early random trigger
1170735853	-	lightning
1197027653	45731478	early random trigger
1235759312	52434532	early random trigger
1238909725	52942518	early random trigger
1239819906	-	lightning
1253054092	55322854	early random trigger
1262972987	56792418	lightning
1265060361	-	lightning

Table 6.3: List of candidates of two ToT pairs left after all cuts are applied.

GPS second	SD-event Id	comment
1168070206	-	-
1191310404	44794740	-
1198409205	45970791	-
1210727643	48123844	two n.n. T1s (no closed footprint)
1241378968	53342934	-
1244019345	-	both at edge of holes in grid
1247546778	54393796	T1 too early (does not cover signal time)
1271648917	58127461	hybrid event (201135690200)

6.3.7 XAS: eXtended Air Showers

In the previous sections, we discussed the search for correlations of small showers as motivated by the predicted GZ events. The geometric compactness is a key condition to enable the identification of small showers in the background of both the reconstructed events and the pairs. However, as we have shown in Section 5.6.1 the sensitivity for showers below 10^{17} eV is very small when pairs of triggers are required. If we want to use single stations to lower the energy threshold, Fig. 5.50 shows that the inclusion of all trigger types is necessary, even if this means that the background from random muons is increased substantially. But this increased background implies an increased threshold such that only coincidences of many stations can be picked out of the background.

In light of these experimental boundary conditions, we aim in this part to find exotic signatures consisting of many small showers triggering only single stations in a coherent way. Thus, we can extend the phase-space of our searches to the extreme edges of multiplicity and separation between fragments. While such signatures are not theoretically well motivated, hence even exotic, there were attempts at searches for generic exotic signatures within the cosmic-ray community before [153, 160]. This work is the first serious attempt to use triggers, i.e. sub-threshold data, to set limits on the rate of such signatures.

The analysis in this work is similar to the classical Hough transform [154] designed to detect lines in images. It relies on the idea that any air-shower-like event will lead to trigger times related to its initial arrival direction. Even more so, every highly relativistic event, producing causally related triggers, will fulfil this requirement. Therefore, the use of an “elementary object” that reconstructs a direction, corresponding to the line through two points in the original Hough transform, will lead to an enhancement in this direction. How such an elementary object for a plane shower front moving with the speed of light can be constructed was already described in Section 6.3.2.

The remaining task in finding a signal with the application of this elementary reconstruction to data is the detection of the combinatoric enhancement of the Hough transform. Figure 6.45 shows such an enhancement originating from a (toy) Monte-Carlo (MC) signal with 25 stations timed according to a plane front indicating that this approach is indeed working for high-multiplicity events.

The identification of such over dense regions in the Hough space is equivalent to our original goal of finding plane-front-like trigger patterns. We use the *DBScan*-algorithm [161] based on the pair wise distance measure d between different reconstructed triplets to detect these clusters. In this analysis, we opt for a χ^2 -like compatibility distance d between two triplets A and B. This distance is based on the expression

$$\chi^2(A, B) = \sum_i \frac{(\tilde{t}_i^A - t_i^B)^2}{\sigma_{\tilde{t}}^2 + \sigma_{t_i}^2} + \frac{(\tilde{t}_i^B - t_i^A)^2}{\sigma_{\tilde{t}}^2 + \sigma_{t_i}^2} \quad (6.44)$$

where \tilde{t}_i^A is the expected trigger time at station i of the triplet B given the plane reconstructed with trigger times t_i^A of the triplet A. The expected time is calculated as

$$\tilde{t}_i^A = \hat{a}^A \cdot (\bar{x}_1^A - \bar{x}_i^B) + t_1^A. \quad (6.45)$$

We estimate the uncertainty of this arrival time as

$$\sigma_{\tilde{t}_i^A}^2 = \sum_{j=1}^3 \sigma_{t_j}^2 \left[\left(\frac{\partial \hat{a}^A}{\partial t_j} \right) \cdot (\bar{x}_1^A - \bar{x}_i^B) + \delta_{j1} \right] \quad (6.46)$$

where the uncertainty of the trigger time $\sigma_{t_j} = 1/\sqrt{12} \mu\text{s}$ is taken from a uniform pdf of width $1 \mu\text{s}$ and where δ_{ij} is the Kronecker delta. This model assumes that the rounding is




Figure 6.45: Illustration of the enhancement in a direction (u_0, v_0) arising from a Monte-Carlo signal with 25 stations. *Left panel:* The reconstructed axis vectors $(u, v, \sqrt{1 - u^2 - v^2})$ as points in the (u, v) -plane. *Right panel:* The positions of the stations with a (simulated) trigger contributing to the reconstructed planes in the left panel. Black markers indicate contribution to the detected cluster, while grey ones are random background triggers.

the dominant contribution to the uncertainty. We know from Fig. 5.46 that the jitter is about 70% larger than the expectation from rounding only. However, as evident from Eqs. (6.44) and (6.46), the overall uncertainty is just a scaling of the distance. Thus, any change of σ_{t_j} will be absorbed in the appropriate hyper-parameter optimisation of the cluster algorithm.

The final distance measure d between two triplets A and B is given by the reduced χ^2 ,

$$d^2(A, B) = \frac{\chi^2(A, B)}{n_{\text{dof}}(A, B)} \quad (6.47)$$

where the number of degrees of freedom $n_{\text{dof}}(A, B)$ is given by

$$n_{\text{dof}}(A, B) = 2 [n_{\text{T2}}(A, B) - 3] \quad (6.48)$$

with the number of different triggers in A and B denoted with $n_{\text{T2}}(A, B)$.

The cluster algorithm *DBScan* has two free parameters: the search-radius ρ around each point and the minimal required number of points within this radius m_{pts} for a cluster to be considered as significant. In this analysis our choice is $m_{\text{pts}} = 25$ and $d^2 < \rho^2 = 1.5$ with d^2 from Eq. (6.47). The choice of m_{pts} defines the minimal number of triggered stations necessary to create a detectable cluster, since the number of triplets from n stations is given by

$$N_{\text{max}}(n) = \binom{n}{3}. \quad (6.49)$$

Thus, we select events with more than six triggered stations since $N_{\text{max}}(6) = 20$.

The value of ρ^2 is motivated by the expected value of the reduced χ^2 , but takes into account the spread for low number of degrees of freedom and is optimised (cf. Fig. 5.46 and the previous discussion) to obtain a low false-positive rate of about 5 mHz, while having a detection efficiency greater than 80% for trigger patterns arising from a plane front with 8 triggered stations, almost independently of their geometry.

Fig. 6.46 shows this detection efficiency as computed with the (toy) MC-signals for plane-front signals as a function of the number of triggered stations for two extreme geometric cases: (a) randomly distributed triggers and (b) neighbouring/compact trigger patterns. We generate the signals with random geometry by drawing random station IDs until the required number of stations is reached and inserting the generated triggers into real data files to have




Figure 6.46: Detection efficiency of plane front signals with the cluster algorithm for two extreme geometric cases as function of the signal size, i.e. the number of triggered stations. Label “random” refers to choosing random stations as part of the plane front and “compact” to neighbouring stations. For comparison, the detection efficiency using a simpler, euclidean-like metric from Eq. (6.50) (instead of Eq. (6.47)) is also shown.

a realistic background. For “compact” geometries, we add the requirement that any new random station is a nearest-neighbour of a previous one. The detection of these MC signals follows the exact same analysis chain as real data, enabling us to calculate the efficiency by matching the found events with the inserted ones.

The aforementioned 80% detection efficiency for 8 triggered stations is clearly reached in both cases. For compact geometries, the detection efficiency is even closer to 100%. This difference arises from the relative timing uncertainty, which is large for compact signatures and increases the probability to include noise triggers into a cluster.

The difference between the two extreme geometries (a) and (b) is however drastically reduced, comparing the distance measure Eq. (6.47) to simple, euclidean like metrics, e.g.

$$d_{\text{eucl}}^2 = \Delta u^2 + \Delta v^2 + k^2 \Delta t_0^2, \quad (6.50)$$

with $k = 1/300 \mu\text{s}^{-1}$. We performed a scan with such a configuration as well and added it in Fig. 6.46 to highlight this fact.

Clusters found by *DBScan* are characterised by the event time, the direction and the stations involved. We define the estimate of the direction and event time as the arithmetic average of the (u, v, t_0) of the triplets in the cluster. To exclude contributions of fully compact events, that are already covered by the more sensitive searches presented in the previous sections, we use the spatial clustering analysis from Section 6.3.1 to veto such configurations. We consider a spatial event with event time t_{spat} to be compatible with a cluster if the time difference between expected arrival-time $\vec{a} \cdot \langle \vec{x} \rangle / c$ and event time is smaller than $50 \mu\text{s}$,

$$|t_{\text{spat}} - \vec{a} \cdot \langle \vec{x} \rangle / c| < 50 \mu\text{s}, \quad (6.51)$$

and we merge clusters if their respective event times are within $50 \mu\text{s}$ of each other. With these condition, we iteratively add all compatible clusters and select only those which are not merged with any other cluster.

The set of remaining clusters is sorted by number of included triggers n_{T2} and compared to the amount of such clusters expected from the random background. To obtain a good estimate of the random background, we scramble the real input data: for station i a random time shift Δt_i is drawn from a uniform distribution with $0 \mu\text{s} < \Delta t_i < 5000 \mu\text{s}$ and added to all trigger times of this station. The maximal shift is about ten times larger than the coherence time, i.e. the propagation time across the SD array with speed of light. To increase statistics in the background estimate the process can be repeated with different random time shifts. Here, we chose 15 *runs* per input file, a good compromise between statistics and computational costs as the analysis needs about 40 mins of CPU time per iteration through one real hour of data.

In summary, the result of this analysis is a spectrum dN_{obs}/dn_{T2} of observed events N_{obs} for the signal run without shuffling and dN_{bg}/dn_{T2} for the random background contribution with an exposure factor of 15. In normal conditions no contributions to such events beyond random coincidences are expected because of the vetoing of spatial clusters. This means that also events affected by dead stations are almost certainly excluded because an event needs to be split in a way where more than four crowns are in between the stations and still about 6 triggers remain. Just the size of the original event necessary to achieve this makes this scenario virtually impossible.

However, in lightning conditions very variable footprints, which can extent over large portions of the SD array, have already been observed. Fig. 6.47 shows an example of a lightning event that did not fulfil the T3 condition and is characterised by fairly large distances between individual stations. The actual lightning strike most likely happened significantly far south of the SD, such that the cone of arrival times from the point-like source is almost compatible with a plane-front pointing to the horizon. While this event is not selected by our event selection, it is easier to identify it as lightning induced and highlights the need to avoid thunderstorm periods for the search for exotic signatures. Thus, we perform the search only in hours of data when not even a single (T3) SD event had a lightning trace. To simplify this selection and avoid limitations of computation time, we choose March to July 2020 as observation period, a period with few lightning events in the SD. The extension to more data is CPU-time consuming but straightforward and can be carried out in the near future.

Finally, we convert the numbers N_{obs} and N_{bg} , shown as histograms in the left panel of Fig. 6.48, into limits or observed rates following the profile likelihood approach [159] for estimating a possible signal contribution with the nuisance parameter N_{bg} using the implementation TRolke within ROOT [132, 133]. In contrast to the previous sections, the uncertainty on N_{bg} is significant because the background exposure is computationally costly and thus lower, making the inclusion of the background uncertainty necessary. With the limit on the number of signal events μ , we get the limit on the flux of such exotic events with

$$\frac{d^4\Gamma}{dn dt dA d\Omega}(n_{T2}) = \frac{\mu}{\varepsilon(n_{T2} T A \Omega)}. \quad (6.52)$$

As efficiency estimates we use the “random” geometry case of Fig. 6.46 and use the covered time period $T = 1.291 \times 10^7$ s and the instant average aperture $A = 2\pi \times \langle n_{T2} \rangle A_{\text{stat}}$ sr km² taken from the area of the Wigner-Seitz cell of the lattice A_{stat} and the average number of active stations $\langle n_{T2} \rangle = 1519.15$. The right panel of Fig. 6.48 shows the resulting limits together with the expectations. To estimate the expected limits, we perform pseudo-experiments with $\mu = 0$ and the given background and get the mean and variance σ of the resulting limits shown as band in the Figure. Clearly, no excess of events is observed and we can provide constraining limits.




Figure 6.47: Example of an “unusual” signature observed during a thunderstorm on 24 January 2016 (GPS second 1137 637 449.299078). This event does not pass a T3-condition. *Left Panel:* Positions of the triggered stations. The type of a trigger is indicated by different markers, while the colour is indicating the cluster it belongs to. Triggers marked as “noise” are triggers which happen in a time window of $[-250, 250]$ μs around the event time and are not part of any cluster. *Right Panel:* Reconstructed directions of triplets in the (u, v) space. The points from the found cluster(s) are marked in black, while the red circles indicate triplets formed by triggers from spatial clusters.

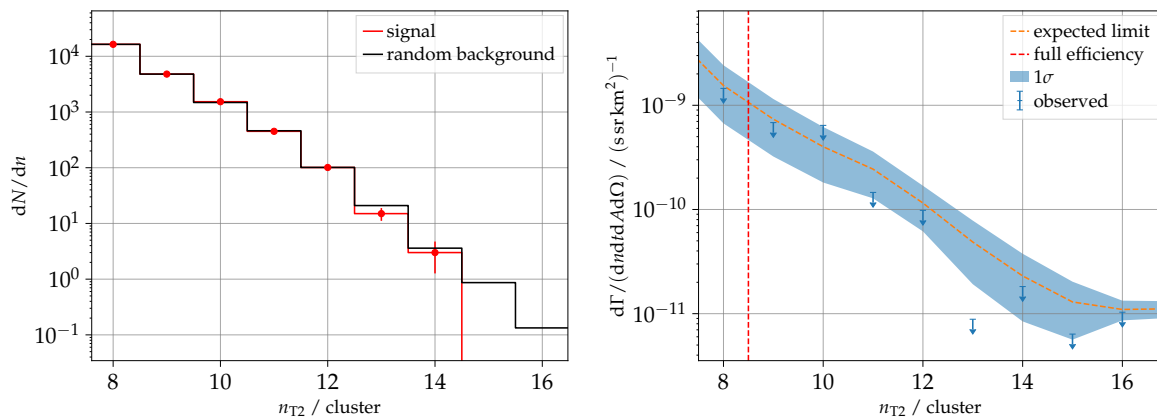


Figure 6.48: Observed and expected number N of clusters found by *DBScan* as a function of their respective trigger multiplicity n_{T2} . Data from the end of austral summer (March 2021) until winter (July 2021) are used here. *Left:* Number of events in the signal sample and the expectation from randomized background. *Right:* Using a profile-likelihood approach [159], the observed and expected counts are converted into limits at 90% C.L. for signals creating exotic signatures. We convert the limits on counts to a rate limit using the observation time $T = 1.291 \times 10^7$ s and the aperture $A = 13946$ sr km² including the efficiency from Fig. 6.46.

6.4 Summary

In this Chapter, we have investigated correlations of the SD triggers by analysing their times and geometry. Simple quantities like the time difference between triggers, timing of pulses and triggers, or geometric coincidences show potential to search for phenomena in so-far unexplored parts of the possible phase-space.

First, we show that we can use the microsecond trigger times to investigate the behaviour of the different triggers on short time scales. We can confirm that the triggers work as expected and do not show instabilities on short time scales beyond those expected from station dead-time or the implementation of baseline tracking in the FPGA. The only (small)

exception to this is the ToTd trigger which reveals interesting autocorrelations. We do not understand the origin of this behaviour yet but observing such deviations shows that such analyses can add valuable information to monitoring of the stability of data taking, or to the commissioning of new triggers or hardware. Thus, we think this work can inspire early tests already during the deployment of the upgraded electronics at the Auger SD in the near future.

Furthermore, when analysing correlations of triggers with reconstructed UHECR showers, we find delayed triggers that are not recorded with the normal data taking procedure. We can show that these delayed triggers are artefacts of the swinging baseline of the PMT signal following large signals. We can observe them without array-level (T3) triggers using tuples of T2 triggers. Inspired by the possibility that such delayed signals could also come from slow particles like neutrons, we investigated the event data of recorded events finding so far unexplained pulses of small amplitude. Comparing the frequency of these pulses with station signals and with shower quantities like the radial distance to the shower axis, we see indications that these pulses are not due to station-specific effects like PMT-afterpulsing. With the help of dedicated simulations⁷, we can find strong indications that the signals are actually from interactions of delayed neutrons. The time distribution, the signal sizes, and also the PMT-signal symmetry follow the expectation leading us to this conclusion. Finding these neutron signals opens a new channel of information on the hadronic interactions in air-showers that was previously ignored. While the signals themselves do not seem to carry a lot of composition information due to the compensation of a shift in the shower maximum with the number of created neutrons, it is important to evaluate the influence of the neutrons on the measured signals. Having shown the presence of neutrons in the showers, mostly as small pulses, future work has to focus on understanding the overall contribution of neutrons to the signal and its effect on the reconstruction.

Besides the time correlations that lead to the discovery of signals caused by neutrons in the trace we also evaluated spatial correlations of coincident triggers. The search for correlated events is also contained in this category. We find that with the current event reconstruction an efficient search for coincident events is not possible due to the details of the DAQ process. This can however, in future analyses, serve as an advantage because all the signals are already available in the recorded SD data, should the existence of such events be confirmed.

Making use of the trigger data, we can reach beyond the DAQ and threshold limitations. We present two separate and independent analyses with the triggers. The first one is using reconstructed events as input and searches for additional separated sub-threshold events detected as nearest-neighbour pairs of ToT-triggers. For simplicity we will refer to this search as “event-plus-pair sample”. The second analysis is removing the necessity for a reconstructed events and uses only coincident pairs of ToT triggers. We dub this analysis and the respective sample of event candidates “pair-of-pairs sample”.

In the event-plus-pair sample we find a 2.2σ excess of separated ToT pairs correlated with reconstructed events over purely random background. These (combined) events are characterised by small separations on ground of less than 5 km between the event and the nearest-neighbour pair of ToT-type triggers which we use to identify the sub-threshold events. Using purely sub-threshold information, in the pair-of-pairs analysis, we find an excess as well. At 3.3σ over background it is more significant than the event-plus-pair sample, associated with higher energies, with separations between the pairs at similar distances.

To add to the hypothesis that these deviations away from the pure background come from true GZ events, we scrutinise the distribution of the detection times and the directions with respect to the Sun. From Fig. 6.49, showing calculated differential rates of GZ events

⁷Performed, analysed, and kindly provided by Alfredo Ferrari.




Figure 6.49: Expected rate of GZ events for a night (*left*) and day (*right*) direction for different primaries as a function of energy. The three lines represent different maximal separations for the two fragments of $d < (1000, 100, 10)$ km respectively. It is clear that for directions close to the Sun, the deflections are generally larger. Both plots taken from Ref. [40].

for a day and night direction for different separations on ground [40]. It is clear that for the same separation on ground at higher energies a trajectory closer to the Sun is required. In simple terms, events close to the Sun on their way to Earth cross large volumes with a strong magnetic field and thus should at the same energy exhibit a larger separation than those coming from the night side of Earth, i.e. from the side opposite the Sun. In Fig. 6.49 this can for instance be seen by comparing the peak of the distribution for the 10 km separation (lowest line) for the day and night Sun direction. A shift by one order of magnitude is clearly visible.

The first visual clue that our data is not uniformly distributed in time can be gained from Fig. 6.50 where we show the time of day of the candidate events as a function of season and time. While the distributions do not seem to be very uniform, an Anderson-Darling test [162] cannot find significant deviations from uniform distributions for the two samples ($p_{\min} = 0.16$). The reason is simply the low statistics of the sample with only 5 and 8 events, respectively.

A more simple ansatz to test whether or not the candidates might point to being GZ events, is the relative direction of the Sun. For the event-plus-pair sample we can calculate the angular distance between CR direction and the Sun. We list these values in Table 6.1. There is an enhancement at about 120° including 3 out of total 5 events. We expect 1.5 events from background in the sample of 5, making an interpretation of the 3 events of similar angular separation from the Sun as GZ signal tempting. However, with such small statistics it is hardly significant.

A simpler approach which we can test statistically is the hypothesis that – based on the similar separations between the parts in the candidates of both samples – the higher energy events from the event-plus-pair sample should come from closer to the Sun than those from the pairs-of-pairs sample. Because we do not have a reconstructed direction for the latter, we can use the Sun’s position on the sky as zero order approximation. Since the shower creating the pairs has to come from above the horizon this introduces a (rough) scale of distance to the Sun’s direction.

We show the position of the Sun on the (local) sky at the time of the events and – if available – the reconstructed direction of the event in Fig. 6.51. Interestingly, for four out of five of the events from the event-plus-pair sample the Sun is above the horizon while for the majority (six out of eight) of events from the pair-of-pairs sample it is below. Motivated by Fig. 6.49 we perform a statistical test that such a separation between day and night is random.

We create Monte-Carlo (MC) data by drawing five and eight GPS seconds from a uniform distribution covering the used data sample for each MC data sample. We then evaluate if

more than three of the GPS seconds in the smaller sample correspond to day time and at least six of the eight to night times. If this is the case, we count this as “similar” to our data. From 10 000 MC samples, 264 pass this condition leading to 1.9σ deviation from a pure background expectation.

We can combine these three significances because all use independent information using Fisher’s method

$$\chi = -2 \sum_i \ln p_i \quad (6.53)$$

of combining the p -values. Converting the χ^2 -distributed sum with six degrees of freedom to a corresponding significance we get 4.1σ .

We have to stress that these 4.1σ are for the rejection of a pure background hypothesis. Thus, contributions from for example rare interactions in air showers could be significant on these scales and might be the missing part in the background modelling for GZ events. To confirm that these candidates are indeed GZ events, detailed modelling of the expected events with the given efficiency of this analysis is required that is beyond the scope of this work. It is clear however, that the background for such searches is low enough to justify further pursuit of these analyses and also to extend them to smaller distances between the events, if possible. Potentially also using the knowledge that coincidences of pairs of ToT triggers are rare, it seems possible to add a special T3 mode to read out these events to facilitate further investigations also for distances above 2 km. Additionally, possibilities of using the new upgraded electronics in the Auger SD should be explored to potentially identify small showers on single-station basis. Further efforts should also be invested in the estimated significance of 4.1σ with respect to systematic uncertainties. We assume they are fairly small because of the large data samples used to estimate for instance the random background and dead-time contributions in the background estimates. However, for example the MC for the direction of the Sun should include the exact uptime of the detectors rather than a pure uniform distribution and model the influence of e.g. atmospheric pressure on the background rate. While this requires a concentrated effort, we think it is worthwhile because we have shown that there are clear indications of unknown and unexpected events in the data.




Figure 6.50: Illustrations of the time distribution of candidates for split events combining data from Section 6.3.5 (blue) and Section 6.3.6 (black). *Left:* To highlight a possible dependence on the relative position of the Sun, we show here the time of day and the season (as fractional year). While the distributions of the time of day may optically seem not very even, an Anderson-Darling test [162] shows that both samples are not significantly ($p_{\min} = 0.16$) non-uniform. However, the small sample size makes any statistical statement about non-uniformity difficult. *Right:* The same data but using detection time instead of the season to obtain a glance of the uniformity across the measurement period.




Figure 6.51: Local skymap of the reconstructed directions (crosses) of the event candidates from Section 6.3.5 together with the corresponding position of the Sun (circles) in the matching colours. The coordinates are local altitude and azimuth at the time of the event so that the relative positions of event and sun on the sky can be directly read off. A direction with azimuth of 0° is pointing north and it is counted positive to the East. For the candidates from Section 6.3.6, we show the position of the Sun only because no reconstructed direction is available.

Table 6.4: A list of rejected stations for the June 2019 data set. Only stations with data are listed here. Additional information is given in the “comments” column. Low statistics refers to a station that had less than 10^5 trigger pairs in the data sample.

station ID	rejected(Th)	rejected(ToT)	comments
116	yes	yes	low statistics (ToT)
132	no	yes	-
140	no	yes	-
149	yes	no	-
151	no	yes	-
170	yes	no	-
172	yes	no	-
174	yes	no	-
181	yes	no	-
190	yes	no	-
203	yes	no	low statistics (ToT)
225	no	yes	-
232	no	yes	-
601	no	yes	low statistics (ToT)
602	yes	no	-
681	yes	no	-
697	yes	no	-
707	yes	no	-
712	yes	no	-
742	no	yes	low statistics (ToT)
792	no	yes	low statistics (ToT)
933	yes	no	-
964	no	yes	-
992	yes	no	-
1079	yes	no	-
1081	yes	no	-
1085	no	yes	-
1135	yes	no	-
1319	no	yes	-
1320	no	yes	-
1453	no	yes	-
1637	yes	no	-
1737	yes	no	low statistics (ToT)
1744	yes	yes	low statistics (ToT)
1800	yes	no	-
1825	yes	no	-
1830	yes	no	-
1831	yes	no	-
1843	no	yes	-
1844	yes	no	-

Summary and Conclusions

In this work we demonstrate that it is worthwhile to reach beyond the usual data analyses used by cosmic-ray observatories, in this case the Pierre Auger Observatory. We show that in low-level data there are interesting effects and physics hiding, waiting for us to explore it.

The first type of data we use are the Auger Scalers, counters operating at a very low threshold. It was previously known that these counters track the rate of low-energy cosmic rays [74]. We extend the range of the analysis of these data by incorporating knowledge of how the baseline is computed in the local station software. With the associated correction, the variance in the data is significantly reduced. On top of revisiting the data-cleaning methods and making use of the standard Auger PMT quality cuts, we also improve the corrections relevant for the long-term evolution of the data. Due to ageing of the detectors, the optical properties are changing and introducing a drift in the scalers, making the observation of the 11-year solar cycle in the raw scaler data impossible. We show that with a simple model of the proportionality of the signal to the optical properties of water-Cherenkov detectors we can correct for this drift. By not having to rely on an (ad-hoc) correlation of the ageing parameter (Area-over-Peak) with the measured rate, we also eliminate the systematic uncertainty introduced by accidental correlations. In detail, if the time scale of the decrease in Area-over-Peak is similar to the change of intensity of low-energy cosmic rays caused by the solar cycle, a correction based on a fitted correlation might over-correct the ageing. Thus, we present a more reliable correction than previously used and show that it is possible to obtain stable corrections to observe the 11-year solar cycle with this auxiliary data from the Auger Observatory. We clearly observe Forbush decreases, find that analyses at the Carrington frequency are correlated with the Sun-spot number, as well as clearly observe the 11-year solar cycle. This shows that it is possible to use these data for studies on solar physics on many time-scales. A point to be addressed in the future is the remaining spurious yearly modulation, likely introduced by stations being rejected with a seasonal dependence. Trying to create a more uniform response of the stations in the Scalers with additional corrections and calibration can help to fulfil this task.

Additionally, the conclusions from these corrections should be used to improve the measurement of such rates with the on-going AugerPrime upgrade. We do not present detailed proposals for such an improved Scaler because the physics of interest has to be decided within the collaboration first. While the current set-up is useful for solar physics, it was already shown previously that the upper limit on the amplitude used – motivated by the search for gamma emission from GRBs – lowers the sensitivity to transient solar events. If limits on ultra-high energy ($\gtrsim 10 - 100$ TeV) gamma ray emission are the main goal, then the three-fold coincidence settings should not be used. However, the influence of such a change of the trigger condition on the stability of the scaler is unknown and has to be tested first.

Thus, the presented effects can be a good foundation for the future development of a new Scaler trigger.

Moving from the lowest-energy data available to the low-level trigger data, used to take the event data for standard analyses, we explore the possibilities that these data offer. Using station doublets and requiring coincident triggers between them, we confirm that these trigger data contain about 4% (0.7 Hz in 20 Hz level-2 trigger (T2) rate) showers of energies above 10^{14} eV. Similarly, we find that additional trigger flags introduced to find such small showers in the trigger data increase this fraction, but do not provide more than 6% purity. Targeting somewhat higher energies, but still staying well below the full efficiency threshold of the event-based SD reconstruction, we find that with pairs of nearest-neighbour triggers we can very efficiently select showers of energies about one order of magnitude below the normal threshold, i.e. at about 10^{17} eV. We obtain this efficiency by comparing reconstructed hybrid events¹ with the pairs found in the trigger data. Due to the low statistics at low-energy, these efficiency estimates should be extended once more data is available, and should also be confirmed with simulations in future work. For the same statistics reason, we cannot determine the efficiency around 10^{17} eV very accurately but it is clear that the window for rate-based searches that do not need the details of recorded events can be extended to much lower energies.

With this trigger data set, we try to test the hypothesis that (some) lightning strikes might be initiated by cosmic rays. The extension to lower energies paired with the large exposure and area of the Auger Observatory create a unique opportunity for such analyses. We find that with single-station triggers it is not possible to search for this correlation because for the ToT-like triggers the background from RFI-pulses is too large and for the threshold triggers the signal-to-background ratio is too small. Extending the analysis to nearest-neighbour pairs (n.n.-pairs) of triggers increases the energy threshold but also reduces the background. Yet, we find similar problems: large RF-influence in ToT-like n.n.-pairs and a very small efficiency to trigger on low-energy showers ($O(1\%)$ or less) for threshold triggers. However, making use of the special trigger flag meant to select small showers, we can suppress the RF background if we request that both of the triggers of a pair are tagged in this way. To an extent, the reduction in background compensates the extreme loss of trigger efficiency to only about 1% at $10^{17.5}$ eV. We find a local excess on the 3.7σ level about 260 ms before the return stroke. However, after accounting for the scanning in time delay and radial distance, it is with 1.6σ not significant any more.

While it is clear that with more data this search window can be used to test the cosmic-ray-initiation hypothesis again without the penalty of the scan, there is a potential problem in the direct interpretation of excesses in this observational channel. We know of a special type of (rare) events occurring during thunderstorms that was observed previously with the surface detector. These special events are characterised by comparatively long signals that are probably picked up by the small shower flag used here and are clearly not normal cosmic-ray events. Thus, it is important for the cosmic-ray-initiation hypothesis-test to be able to distinguish a cosmic-ray air shower from these rare events. Therefore, we propose that adding an array-level trigger of such pairs should be considered because it could provide decisive information if an excess – like the local one reported here – was observed in future analyses.

We also use the trigger data to investigate the aforementioned special events characterised by a circular footprint on the ground with a missing centre, hence dubbed “SD rings”. The main question is whether or not the missing central part is a physical property of the events or rather an artefact of the DAQ. We show that although trigger effects and station dead-times, found in the T2 data, influence what is recorded, no clear problem with the DAQ can be

¹A Fluorescence Detector measurement with at least one triggered SD station.

identified as cause for the central hole. Nonetheless, we find clear evidence that the missing stations in the centre are indeed a (station) trigger effect, because we observe overshoots of the baseline that follow very large signals in the event data and the timing of the overshoots matches with that expected from a large signal at the time of the SD-ring event. Thus, we know that the footprint of these events is actually filled. The similarity of these events with other observations and their correlation with thunderstorms suggests that we observe Terrestrial Gamma-Ray Flashes (TGFs). To confirm this TGF hypothesis, more future measurements are necessary. One possibility is to use additional instrumentation to provide information about the relative timing of these events within the lightning development.

With the addition of new and more sensitive station triggers, the rate of the SD-ring events dropped significantly. The new triggers are very sensitive to RF signals from lightning processes and thus create many (array-level) triggers overwhelming the DAQ. We can show that the communication system of the Auger SD limits the number of events that can be read out successfully per (GPS) second and this significantly affects the periods with SD-ring events. Therefore, we develop a modification to the DAQ that can identify signals from SD-rings with about 75% efficiency to allow for their prioritised read out without affecting normal data taking. With the implementation of this priority read out we will recover the drop in the rates of the SD rings, and also be able to provide a trigger to external instruments to select the interesting lightning events in data. We think that in the near future these events can provide many insights into high-energy emission processes during lightning strikes and that further development of the techniques and triggers is worthwhile.

We also sketch possibilities how the station trigger could be improved for SD rings with the upgraded electronics of the surface detectors. Our results make it clear that it is difficult to improve the sensitivity of the existing station triggers to low-amplitude signals because the cosmic-ray signals, to which existing triggers are optimised, are of similar amplitude but shorter in duration and thus harder to capture than the long-signals of SD rings. However, adding a new trigger that is very efficient at rejecting lightning-induced RF signals can improve the sensitivity in the central part of the events since the most likely explanation for the missing signals is saturation of the triggers, i.e. an activation of all triggers long before the signal of interest occurs. A new less sensitive trigger could alleviate this problem. Additionally, we show that using a coarser time binning by averaging every four consecutive time bins into a dedicated trace could be treated easily with algorithms similar to the standard triggers. This way they can provide additional information on the time evolution of the signals beyond what is possible to record with the standard (CR) setup due to data-rate limitations. The proposals of this work have to be tested before deployment in the field is possible. Yet it is evident that the possibilities for future improvements are very promising.

When searching for delayed signals after (normal) cosmic-ray events we also find the overshoot of the baseline in stations with large signals. Thus, these spurious triggers prevent any identification of small contributions of, for instance, slow neutrons in the central part of air showers in the trigger data. Nonetheless, we extend the search to the event data covering 20 μ s instead of about 1 ms as with the trigger data. Interestingly, we do observe delayed pulses at low amplitudes that cannot be explained with simple afterpulses of the photo multipliers. A comparison with simulations including neutrons show that these signals, about 7 μ s delayed with respect to the first particles of the shower, are indeed caused by neutrons interacting in the detectors. The main signal is caused by neutron interactions with nuclei ($n + N \rightarrow n + N^*$) creating nuclear excitations of the oxygen nuclei in water that emit mostly a 6 MeV gamma ray upon de-excitation. Due to the low energy of the de-excitation gamma rays, a high asymmetry of the signal in the three PMTs of each water-Cherekov detector is expected, and it is indeed observed in data. These signals raise interesting questions for the treatment of neutrons in simulations. So far, they were not explicitly included in the simulations due to

their signals being fairly small, but from our results it is clear that they are very numerous and thus do contribute. Their exact share of the total signal has to be investigated, especially in light of in literature reported differences between data and simulations. We are sure that this result calls for further work on the understanding of the hadronic component of air showers at the highest energies.

In the last part of this work, we analyse correlations between air showers to search for both exotic and rare but predicted effects. With exotic effects, we mean the correlation of many small air showers on distance scales between individual showers of several kilometres that are not predicted and expected. By using the trigger data, we can lower the energy threshold for the air showers within such correlated events to about 10^{15} eV and use the full area of the Pierre Auger Observatory for this correlation search. Using a Hough transform of plane shower fronts, and a selection of geometrically separated triggers we are able to set a strong limit on such correlated showers. We find that in Auger for 10 triggered stations the rate of such events from random background is about 10^{-4} Hz, leading to a flux limit of $6.4 \times 10^{-10} \text{ (s sr km}^2\text{)}^{-1}$.

While we are not aware of any theory predicting such exotic signatures of many coincident showers, there are predictions for the rate of Gerasimova-Zatsepin events. These are two² coincident cosmic rays that originate from photo-disintegration of a primary cosmic-ray nucleus in the photon field of the Sun. However, the typically predicted rates of these events are less than one per year for detectors of the size of Auger if the full-efficiency energy threshold is applied. For these rare events we extend the search to lower energies with the trigger data. Our search strategy includes the use of normal reconstructed cosmic-ray events and covers several implicitly selected energy ranges down to the trigger-only search.

In the search for coincident reconstructed events, we find that for our selection of events a background of about 0.5 events per year is expected for coincidences within 100 μs . A direct observation of such events is however not possible because of automatic merging of such events by the DAQ. Given this low background expectation, it seems promising to extend the search in future work by revisiting the event reconstruction and trying to extract from event data the coincident events merged into a single event. The sensitivity of such a search will depend on the efficiency and specificity of the event-separation algorithm but we think it should be pursued to obtain an actual limit on the rate of GZ events with the Auger SD.

Moving to lower energies, we use the combination of a n.n.-pair of triggers correlated with a reconstructed event as next search channel. We account for several possible background contributions: Random coincidences are estimated using randomly shuffled data. Contributions from normal (higher-energy) events split by stations in dead-time are calculated using event data. And we remove effects from non-functioning stations. We find that there is a minor excess of such correlated events above the combined background at the 2.2σ level. In detail, we observe five pairs correlated with events with a background expectation of 1.5.

While this result alone is not significant we also find an excess of spatially separated and coincident n.n.-pairs (pairs of pairs) in the lowest energy channel of the analysis. For these pairs we observe eight events in the data with an expected background of 1.8 corresponding to 3.3σ excess at the same distance scale (about 2 km separation). To further test if the sample we obtain is randomly distributed or rather linked to the Sun and hence GZ-like, we test the event times with respect to the position of the Sun at the times of the events. We find that for the higher energy sample (event plus n.n.-pair) four out of five events occur during the day, while six of eight in the lower energy sample (pairs of pairs) occur during the night. This is in line with the simplified expectation that at the same separation scale of the events, lower energy events have to come from directions further away from the Sun, because the magnetic

²While in principle multiple splits of a nucleus in photo-nuclear interactions are possible, contributions of multiple interactions are negligible.

field close to the Sun is larger and would split them to larger distances. We use a Monte-Carlo estimate to test how significant such a separation of the two event sets in day and night is and find that it is only a 1.9σ effect due to low statistics. However, combining the three independent observations of event rate in the two channels and arrival times with Fisher's method, we arrive at a combined significance of 4.1σ against a background hypothesis.

At such a significance level we thus have indications that there are effects beyond the background we account for. A link to GZ events seems likely given the predicted rates of a few events per year below threshold for Auger, and the arrival times match this hypothesis as well. However, the detection efficiency we estimate for the most likely process of nuclear disintegration, emission of a single nucleon, seems too low to be in agreement with the predictions. Thus, we do not exclude other causes, for example rare effects in hadronic interactions like an emergence of a separated sub-shower from a high- p_T particle created in the first interaction. To confirm such hypotheses, careful simulations with sufficient statistics have to be performed. Also revisiting the exact predictions for GZ events in the sub-threshold regime explored here seems a promising idea. Furthermore, the data that will be collected with the Auger SD in the next years can help increase the statistics of this analysis significantly as only about three years of data were used in this work (compared to the 15 years of standard Auger data).

Even if this excess should turn out to be of other origin than the GZ mechanism, it still shows the potential of searches as performed here. We have to keep in mind that the obtained event rates show that the background in such analyses is of the order of one per year, similar to the initial predictions that deemed such rates unobservable. Detailed modelling of efficiencies, thresholds, and apertures is necessary to move forward from this point. Given the evidence we find, it is of great interest to further pursue and extend these searches in the future.

In summary, we have shown that the use of trigger data to reach beyond the usual analysis paths at the Pierre Auger Observatory reveals many interesting results that will hopefully trigger further investigations.

From observing delayed triggers after events, we find that the neutrons from air showers can create measurable signals in the surface detectors. Thus, we highlight that the choice to neglect them in simulations should be revisited. Additionally, we show how TGF candidates, observed in the Auger SD, can be better identified and read out based on the trigger signatures we have seen in event data. Most importantly, we find evidence at 4σ -level of GZ events when using the trigger data, thanks to the significantly lower threshold energy they provide. This analysis has shed light onto very interesting physics within the reach of the Pierre Auger Observatory.

Additional Information on SD-Rings

We add here further information on the data used in Section 5.3. Table A.1 has information on the SD-ring events that are used in this analysis together with their GPS-timing information and information on whether or not they contain overshoot signals.

Figs. A.1 and A.2 highlight the relative timing of these effects together with the actual ring event. Fig. A.3 shows the position of the stations for the SD-ring event which is also used for the analysis in Section 5.2.

Table A.1: Parameters of long-signal events used in this study. The station Id of the *central station* gives an estimation of the centre of the event. If no such id is given the centre can be approximated by taking the barycentre of all participating stations. This event list is not complete in terms of long-signal events as *disk*-events are not investigated here. The event GPS microsecond is estimated from the end-of-trace time of some of the clear long-signal stations and is as such only a crude first estimate.

event Id	GPS second	microsecond	central station	overshoot
1332969	799216018	818600	895	no
1668134	812404399	191160	-	no
4067414	876427835	195820	1346	yes
4067430	876428201	343430	1114	yes
4067441	876428578	942950	1114	no
4364128	882598057	28500	-	no
4365403	882622554	3137	-	no
4365444	882623320	497390	-	no
4572085	887560302	559585	926	no
4633093	888868850	274570	321	no
7277834	919258771	587470	-	no
8811846	944230218	122300	-	no
12871544	1002052510	683581	321	yes
21041532	1044210927	132023	-	no
21430560	1048984840	303475	-	no
34812479	1127156524	105627	1676	no
42852075	1179105785	94200	-	yes




Figure A.1: Data from T2s and T3-requested T1s from events around the SD event 42852075. Lost and not-found T1 requests are not shown for better visibility. Cf. Fig. 5.17 for detailed explanation of the plotting.




Figure A.2: Data from T2s and T3-requested T1s from events around the SD event 4067441. Cf. Fig. 5.18 for detailed explanation of the plotting.




Figure A.3: Position of the stations with long-signals and overshoots in the May 2017 SD-ring event. All stations appearing in any event within ± 2.5 ms are marked with black markers, even if they report only lost or not-found data. Stations that had at least one T2 are marked with blue circles. Long and overshoot signals are marked separately based on a visual classification. The red circle around the estimated centre highlights the radius of 7.5 km used for displaying stations in the central parts e.g. in Fig. A.1.

Acknowledgments

I thank the reviewers of this thesis Prof. Dr. Ralph Engel and Dr. Xavier Bertou for making this thesis possible, and for their guidance, support, and ideas, as well as the chance to participate in the joint double doctoral degree, a unique and lasting experience that I am very grateful for. My special thanks go to Dr. Darko Verberič for countless discussions, help, and inspiration, on and off work. For invaluable help with logistics and organisation, I thank Sabine Bucher, Marie-Christine Mundt, and Anna Friedrich.

To Flavia Gesualdi for support in good and bad moments, endless help, and for being the best friend. My thanks go to my family: my parents Walter and Irene; and my siblings Karin and Rudolf for their help and silent support. I am very grateful to all who welcomed me in Bariloche. To Isabel Goos for being a good friend, moral support, and invaluable in unexpected situations. To Miguel Sofo Haro, Evelyn Coronel, Geraldina Golup, Monica Barrera, and María Belén Lovino for welcoming me at CAB and helping me around. To Alexander Streich and Ana Laura Müller for calmly enduring my lack of knowledge about organisational procedures, for help, and relaxed and ironic discussions about current situations. I am grateful to Steffen Hahn for his support with python projects, data-handling, and for enduring my requests for more storage space.

To Daniela Mockler and Joachim Debatin for making the start into my PhD-project easier and the coffee enjoyable. To all I had the pleasure of sharing an office at IAP with: To Alvaro Taboada, Johan Hulsmann, Maximilian Reininghaus, Thomas Moch, Nicholas Gonzales, Fabio Convenga, Tobias Schulz, and Allan Payeras for the good company and enjoyable atmosphere. To Christoph Schäfer, Maximilian Stadelmaier, Emily Martins for enduring my excessive need for coffee and for the nice atmosphere created. To Felix Schlüter, Kathrin Bismark, Olena Tkachenko for the interesting discussions and good company. To David Schmidt for being always helpful and never rejecting a question, and for the good time in Bariloche. To Alfredo Ferrari for support in late analysis and precise explanations.

To Roberto Mussa and the members of the Auger Cosmo-Geo task for fruitful discussions and help. To Roberta Colalillo for always being helpful and open. Special thanks to Ricardo Sato, without whom no such analyses would be possible.

To all the helpful and inspiring people at the institutes and within the Auger collaboration not mentioned here.

I am truly lucky to have met all of you.

Bibliography

- [1] A. Aab, et al., Pierre Auger, The Pierre Auger Cosmic Ray Observatory, Nucl. Instrum. Meth. A 798 (2015) 172–213. arXiv:1502.01323, doi:10.1016/j.nima.2015.06.058.
- [2] A. Aab, et al., Pierre Auger, Depth of Maximum of Air-Shower Profiles at the Pierre Auger Observatory: Measurements at Energies above $10^{17.8}$ eV, Phys. Rev. D 90 (12) (2014) 122005. arXiv:1409.4809, doi:10.1103/PhysRevD.90.122005.
- [3] A. Aab, et al., The Pierre Auger Collaboration, Measurement of the cosmic-ray energy spectrum above 2.5×10^{18} eV using the pierre auger observatory, Phys. Rev. D 102 (2020) 062005. doi:10.1103/PhysRevD.102.062005.
- [4] A. Aab, et al., Pierre Auger, The Pierre Auger Observatory Upgrade - Preliminary Design Report (2016)arXiv:1604.03637.
- [5] N. Gerasimova, G. Zatsepin, Disintegration of cosmic ray nuclei by solar photons, JETP 11 (899) (1960).
- [6] L. Cazon, EAS-MSU, IceCube, KASCADE Grande, NEVOD-DECOR, Pierre Auger, SUGAR, Telescope Array, Yakutsk EAS Array, Working Group Report on the Combined Analysis of Muon Density Measurements from Eight Air Shower Experiments, PoS ICRC2019 (2020) 214. arXiv:2001.07508, doi:10.22323/1.358.0214.
- [7] J. R. Dwyer, M. A. Uman, The physics of lightning, Physics Reports 534 (4) (2014) 147–241, the Physics of Lightning. doi:https://doi.org/10.1016/j.physrep.2013.09.004.
- [8] A. Aab, et al., A 3-year Sample of Almost 1600 Elves Recorded Above South America by the Pierre Auger Cosmic-Ray Observatory, Earth and Space Science 7 (4) (2020) e2019EA000582, e2019EA000582 10.1029/2019EA000582. doi:https://doi.org/10.1029/2019EA000582.
- [9] R. U. Abbasi, et al., Gamma-ray Showers Observed at Ground Level in Coincidence With Downward Lightning Leaders, J. Geophys. Res. Atmos. 123 (2018) 6864. arXiv:1705.06258, doi:10.1029/2017JD027931.
- [10] R. Colalillo, Pierre Auger, Peculiar lightning-related events observed by the surface detector of the Pierre Auger Observatory (2017) 138–145doi:10.22323/1.301.0314.
- [11] V. F. Hess, Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten, Phys. Z. 13 (1912) 1084.

- [12] S. Mollerach, E. Roulet, Progress in high-energy cosmic ray physics, *Prog. Part. Nucl. Phys.* 98 (2018) 85–118. arXiv:1710.11155, doi:10.1016/j.pnpnp.2017.10.002.
- [13] A. Aab, et al., Pierre Auger, Observation of a Large-scale Anisotropy in the Arrival Directions of Cosmic Rays above 8×10^{18} eV, *Science* 357 (6537) (2017) 1266–1270. arXiv:1709.07321, doi:10.1126/science.aan4338.
- [14] R. Battiston, AMS 02, The antimatter spectrometer (AMS-02): A particle physics detector in space, *Nucl. Instrum. Meth. A* 588 (2008) 227–234. doi:10.1016/j.nima.2008.01.044.
- [15] A. Yushkov et al. (Pierre Auger Collab.), Mass composition of cosmic rays with energies above 10^{17} eV from the hybrid data of the pierre auger observatory, in: Proc. 36th Int. Cosmic Ray Conf., Madison WI, USA, 2019.
- [16] R. Alves Batista, et al., Open Questions in Cosmic-Ray Research at Ultrahigh Energies, *Front. Astron. Space Sci.* 6 (2019) 23. arXiv:1903.06714, doi:10.3389/fspas.2019.00023.
- [17] L. A. Anchordoqui, Ultra-High-Energy Cosmic Rays, *Phys. Rept.* 801 (2019) 1–93. arXiv:1807.09645, doi:10.1016/j.physrep.2019.01.002.
- [18] P. A. Zyla, et al., Particle Data Group, Review of Particle Physics, *PTEP* 2020 (8) (2020) 083C01. doi:10.1093/ptep/ptaa104.
- [19] W. Heitler, The quantum theory of radiation, Oxford University Press (1954).
- [20] J. Matthews, A Heitler model of extensive air showers, *Astropart. Phys.* 22 (2005) 387–397. doi:10.1016/j.astropartphys.2004.09.003.
- [21] D. Heck, J. Knapp, J. N. Capdevielle, G. Schatz, T. Thouw, CORSIKA: A Monte Carlo code to simulate extensive air showers, FZKA-6019 (1998).
- [22] S. J. Sciutto, AIREX: A system for air shower simulations (1999) arXiv:astro-ph/9911331, doi:10.13140/RG.2.2.12566.40002.
- [23] T. Pierog, I. Karpenko, J. M. Katzy, E. Yatsenko, K. Werner, EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider, *Phys. Rev. C* 92 (3) (2015) 034906. arXiv:1306.0121, doi:10.1103/PhysRevC.92.034906.
- [24] S. Ostapchenko, QGSJET-II: physics, recent improvements, and results for air showers, *EPJ Web Conf.* 52 (2013) 02001. doi:10.1051/epjconf/20125202001.
- [25] F. Riehn, R. Engel, A. Fedynitch, T. K. Gaisser, T. Stanev, Hadronic interaction model Sibyll 2.3d and extensive air showers, *Phys. Rev. D* 102 (6) (2020) 063002. arXiv:1912.03300, doi:10.1103/PhysRevD.102.063002.
- [26] E. S. Seo, et al., CREAM: 70 days of flight from 2 launches in Antarctica, *Adv. Space Res.* 42 (10) (2008) 1656–1663. doi:10.1016/j.asr.2007.03.056.
- [27] O. Adriani, et al., The PAMELA experiment on satellite and its capability in cosmic rays measurements, *Nucl. Instrum. Meth. A* 478 (2002) 114–118. doi:10.1016/S0168-9002(01)01726-0.
- [28] J. Chang, et al., DAMPE, The DARK Matter Particle Explorer mission, *Astropart. Phys.* 95 (2017) 6–24. arXiv:1706.08453, doi:10.1016/j.astropartphys.2017.08.005.

- [29] J. A. Hinton, H.E.S.S., The Status of the H.E.S.S. project, *New Astron. Rev.* 48 (2004) 331–337. arXiv:astro-ph/0403052, doi:10.1016/j.newar.2003.12.004.
- [30] S. F. Berezhnev, et al., The Tunka-133 EAS Cherenkov light array: status of 2011, *Nucl. Instrum. Meth. A* 692 (2012) 98–105. arXiv:1201.2122, doi:10.1016/j.nima.2011.12.091.
- [31] H. Kawai, et al., Telescope Array, Telescope array experiment, *Nucl. Phys. B Proc. Suppl.* 175–176 (2008) 221–226. doi:10.1016/j.nuclphysbps.2007.11.002.
- [32] E. Fermi, On the Origin of the Cosmic Radiation, *Phys. Rev.* 75 (1949) 1169–1174. doi:10.1103/PhysRev.75.1169.
- [33] T. K. Gaisser, R. Engel, E. Resconi, *Cosmic Rays and Particle Physics: 2nd Edition*, Cambridge University Press, 2016.
- [34] A. M. Hillas, The Origin of Ultrahigh-Energy Cosmic Rays, *Ann. Rev. Astron. Astrophys.* 22 (1984) 425–444. doi:10.1146/annurev.aa.22.090184.002233.
- [35] A. Aab, et al., Pierre Auger, An Indication of anisotropy in arrival directions of ultrahigh-energy cosmic rays through comparison to the flux pattern of extragalactic gamma-ray sources, *Astrophys. J. Lett.* 853 (2) (2018) L29. arXiv:1801.06160, doi:10.3847/2041-8213/aaa66d.
- [36] G. E. Romero, A. L. Müller, M. Roth, Particle acceleration in the superwinds of starburst galaxies, *Astron. Astrophys.* 616 (2018) A57. arXiv:1801.06483, doi:10.1051/0004-6361/201832666.
- [37] A. R. Bell, J. H. Matthews, Echoes of the past: ultra-high energy cosmic rays accelerated by radio galaxies, scattered by starburst galaxies (2021) arXiv:2108.08879.
- [38] K. Greisen, End to the cosmic-ray spectrum?, *Phys. Rev. Lett.* 16 (17) (1966) 748–750. doi:10.1103/PhysRevLett.16.748.
- [39] G. T. Zatsepin, V. A. Kuz'min, Upper Limit of the Spectrum of Cosmic Rays, *JETP Lett.* 4 (1966) 78 – 80.
- [40] L. N. Epele, S. Mollerach, E. Roulet, On the disintegration of cosmic ray nuclei by solar photons, *JHEP* 03 (1999) 017. arXiv:astro-ph/9812130, doi:10.1088/1126-6708/1999/03/017.
- [41] G. A. Medina-Tanco, A. A. Watson, The Photodisintegration of cosmic ray nuclei by solar photons: The Gerasimova-Zatsepin effect revisited, *Astropart. Phys.* 10 (1999) 157–164. arXiv:astro-ph/9808033, doi:10.1016/S0927-6505(98)00056-5.
- [42] S. Lafebre, H. Falcke, J. Horandel, J. Kuijpers, Prospects for direct cosmic ray mass measurements through the Gerasimova-Zatsepin effect, *Astron. Astrophys.* 485 (2008) 1. arXiv:0804.2633, doi:10.1051/0004-6361:200809468.
- [43] J. van Eijden, S. de Jong, C. Timmermans, Cosmic ray interactions in the solar system: The gerasimova-zatsepin effect, arXiv:1606.07693v2 (2016).
- [44] A. Iyono, et al., Cosmic ray composition studies through the Gerasimova-Zatsepin effects of heavy nuclei at LAAS, *Astrophys. Space Sci. Trans.* 7 (2011) 327–333. doi:10.5194/astra-7-327-2011.

- [45] P. La Rocca, et al., Search for long distance time correlations between cosmic air showers with the MRPC telescopes of the EEE network, PoS EPS-HEP2019 (2020) 051. doi:10.22323/1.364.0051.
- [46] J. Abraham, et al., Pierre Auger, Trigger and Aperture of the Surface Detector Array of the Pierre Auger Observatory, Nucl. Instrum. Methods A 613 (2010) 29–39. arXiv:1111.6764, doi:https://doi.org/10.1016/j.nima.2009.11.018.
- [47] L. Miroshnichenko, Solar Cosmic Rays, Springer Heidelberg, 2015.
- [48] R. Bruno, V. Carbone, Turbulence in the Solar Wind, Springer Heidelberg, 2016.
- [49] S. E. Forbush, On the effects in cosmic-ray intensity observed during the recent magnetic storm, Phys. Rev. 51 (1937) 1108–1109. doi:10.1103/PhysRev.51.1108.3.
- [50] H. Cane, Coronal mass ejections and forrush decreases, Space Science Reviews 93 (2000) 55–77. doi:10.1023/A:1026532125747.
- [51] K. L. Klein, A. L. M. (Eds.), The High Energy Solar Corona: Waves, Eruptions, Particles, Springer, Berlin, Heidelberg, 2007. doi:10.1007/978-3-540-71570-2.
- [52] G. Bazilevskaya, et al., Solar cycle in the heliosphere and cosmic rays, In: Balogh A., Hudson H., Petrovay K., von Steiger R. (eds) The Solar Activity Cycle. Space Sciences Series of ISSI, vol 53. Springer, New York, NY. doi:10.1007/978-1-4939-2584-1_14.
- [53] A. Y. Kostinskiy, T. C. Marshall, M. Stolzenburg, The mechanism of the origin and development of lightning from initiating event to initial breakdown pulses (v.2), J. Geophys. Research Atm 125 (22) (2020) e2020JD033191, e2020JD033191 2020JD033191. doi:https://doi.org/10.1029/2020JD033191.
- [54] A. Syssoev, D. Iudin, F. Iudin, V. Klimashov, A. Emelyanov, On the problem of critical electric field of atmospheric air, Atmosphere 12 (8) (2021). doi:10.3390/atmos12081046.
- [55] M. Füllekrug, Introduction to lightning detection, Weather 72 (2) (2017) 32–35. arXiv:https://rmets.onlinelibrary.wiley.com/doi/pdf/10.1002/wea.2810, doi:https://doi.org/10.1002/wea.2810.
- [56] O. Scholten, et al., The initial stage of cloud lightning imaged in high-resolution, J. Geophys. Research Atm. 126 (4) (2021) e2020JD033126, e2020JD033126 2020JD033126. doi:https://doi.org/10.1029/2020JD033126.
- [57] G. J. Fishman, et al., Discovery of intense gamma-ray flashes of atmospheric origin, Science 264 (5163) (1994) 1313–1316. doi:10.1126/science.264.5163.1313.
- [58] J. Dwyer, D. Smith, S. Cummer, High-energy atmospheric physics: Terrestrial gamma-ray flashes and related phenomena, Space Sci Rev 173 (2012) 133–196. doi:10.1007/s11214-012-9894-0.
- [59] J. R. Dwyer, et al., A ground level gamma-ray burst observed in association with rocket-triggered lightning, Geophys. Research Lett. 31 (5) (2004). doi:https://doi.org/10.1029/2003GL018771.
- [60] T. Abu-Zayyad, et al., Telescope Array Scientific, J. W. Belz, P. R. Krehbiel, J. Remington, M. A. Stanley, R. U. Abbasi, R. LeVon, W. Rison, D. Rodeheffer, Observations of the Origin of Downward Terrestrial Gamma-Ray Flashes, J. Geophys. Research Atm. 125 (23) (2020) e2019JD031940. arXiv:2009.14327, doi:10.1029/2019JD031940.

- [61] (retrieved 15.11.2021). [link].
URL https://upload.wikimedia.org/wikipedia/commons/9/93/Lightning_sprites.jpg
- [62] B. Dawson, Pierre Auger, The Energy Scale of the Pierre Auger Observatory, PoS ICRC2019 (2020) 231. doi:10.22323/1.358.0231.
- [63] D. Veberič (2021). [link].
URL https://web.ikp.kit.edu/darko/auger/auger-array/auger_array-pdf/auger_array-ad.pdf
- [64] M. Ave et al. (AIRFLY Collab.), Measurement of the pressure dependence of air fluorescence emission induced by electrons, *Astropart. Phys.* 28 (2007) 41, arXiv astro-ph/0703132. doi:10.1016/j.astropartphys.2007.04.006.
- [65] T. K. Gaisser, A. M. Hillas, Reliability of the Method of Constant Intensity Cuts for Reconstructing the Average Development of Vertical Showers, in: Proc. 15th Int. Cosm. Ray Conf., Vol. 8, Plovdiv, Bulgaria, 1977, p. 353.
- [66] A. Aab, et al., Pierre Auger, Data-driven estimation of the invisible energy of cosmic ray showers with the Pierre Auger Observatory, *Phys. Rev. D* 100 (8) (2019) 082003. arXiv:1901.08040, doi:10.1103/PhysRevD.100.082003.
- [67] I. Allekotte, et al., Pierre Auger, The Surface Detector System of the Pierre Auger Observatory, *Nucl. Instrum. Meth. A* 586 (2008) 409–420. arXiv:0712.2832, doi:10.1016/j.nima.2007.12.016.
- [68] P. Abreu, et al., Pierre Auger, The energy spectrum of cosmic rays beyond the turn-down around 10^{17} eV as measured with the surface detector of the Pierre Auger Observatory, *Eur. Phys. J. C* 81 (2021) 966. arXiv:2109.13400, doi:10.1140/epjc/s10052-021-09700-w.
- [69] X. Bertou, Do cosmic rays and lightning correlate? a first look at a new dataset, Auger internal note GAP-2016-008 (2016).
- [70] X. Bertou, private communication (2020).
- [71] J. Rautenberg, Pierre Auger, Lightning Detection at the Pierre Auger Observatory, PoS ICRC2015 (2016) 678. doi:10.22323/1.236.0678.
- [72] X. Bertou, et al., Pierre Auger, Calibration of the surface array of the Pierre Auger Observatory, *Nucl. Instrum. Meth. A* 568 (2006) 839–846. arXiv:2102.01656, doi:10.1016/j.nima.2006.07.066.
- [73] A. Aab, et al., Pierre Auger, Reconstruction of events recorded with the surface detector of the Pierre Auger Observatory, *JINST* 15 (10) (2020) P10021. arXiv:2007.09035, doi:10.1088/1748-0221/15/10/P10021.
- [74] P. Abreu, et al., Pierre Auger Collaboration, The Pierre Auger Observatory scaler mode for the study of solar activity modulation of galactic cosmic rays, *JINST* 6 (2011) P01003. doi:10.1088/1748-0221/6/01/P01003.
- [75] X. Bertou, Search for gamma ray bursts using the single particle technique at the pierre auger observatory, in: Proc. 30th Int. Cosmic Ray Conf., Pierre Auger Collaboration, Mérida, Mexico, 2007.

- [76] S. Dasso, H. Asorey, Pierre Auger, The scaler mode in the Pierre Auger Observatory to study heliospheric modulation of cosmic rays, *Adv. Space Res.* 49 (2012) 1563–1569. arXiv:1204.6196, doi:10.1016/j.asr.2011.12.028.
- [77] A. Taboada, Pierre Auger, Analysis of Data from Surface Detector Stations of the AugerPrime Upgrade, *PoS ICRC2019* (2020) 434. doi:10.22323/1.358.0434.
- [78] G. Cataldi, et al., Pierre Auger, The upgrade of the Pierre Auger Observatory with the Scintillator Surface Detector, *PoS ICRC2021* (2021) 251. doi:10.22323/1.395.0251.
- [79] P. Abreu, et al., Pierre Auger, AugerPrime Upgraded Electronics, *PoS ICRC2021* (2021) 230. doi:10.22323/1.395.0230.
- [80] A. Castellina, Pierre Auger, AugerPrime: the Pierre Auger Observatory Upgrade, *EPJ Web Conf.* 210 (2019) 06002. arXiv:1905.04472, doi:10.1051/epjconf/201921006002.
- [81] P. Abreu, et al., Pierre Auger, First results from the AugerPrime Radio Detector, *PoS ICRC2021* (2021) 270. doi:10.22323/1.395.0270.
- [82] David Schmidt, Sensitivity of augerprime to the masses of ultra-high-energy cosmic rays, phdthesis (2019).
- [83] H. Asorey *et al.* (Pierre Auger Collaboration), Cosmic Ray Solar Modulation Studies in the Pierre Auger Observatory, in: *Proc. 31st Int. Cosmic Ray Conf.*, Łódź, Poland, 2009.
- [84] J.J. Masias-Meza *et al.* (Pierre Auger Collaboration), Solar Cycle Modulation of Cosmic Rays Observed with the Low Energy Modes of the Pierre Auger Observatory, in: *Proc. 34th Int. Cosmic Ray Conf.*, The Hague, The Netherlands, 2015.
- [85] Ahmed Mohamed Saleh Hassanin Khalil, Studies of astrophysical very-high energy gamma-ray emission with the pierre auger observatory, phdthesis (2017).
- [86] M. Schimassek, Search for transient sources of low-energy cosmic rays with the Pierre Auger Observatory, Master's thesis, Karlsruhe Institute of Technology, Karlsruhe (11 2017).
- [87] M. Schimassek *et al.* (Pierre Auger Collab.), Analysis of Data from the Low-energy Modes of the Surface Detector of the Pierre Auger Observatory, in: *Proc. 36st Int. Cosmic Ray Conf.*, Madison, Wisconsin, USA, 2019.
- [88] M. Schimassek, D. Veberič, R. Engel, A New Scaler Analysis Framework, Auger internal note GAP-2019-060 (2019).
- [89] I. Lhenry-Yvon, About bad PMTs : inventory of spurious data remaining after the two first levels of PMT quality cuts, Auger internal note GAP-2018-022 (2018).
- [90] P. P. I. Lhenry-Yvon, Bad PMTs : the third levels of quality cuts, tuned for PMT traces analysis, Auger internal note GAP-2018-029 (2018).
- [91] R. Sato, Sd array off - cloudy days, private communication. (2021).
- [92] S. S. Wilks, The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses, *Ann. Math. Stat.* 9 (1) (1938) 60 – 62. doi:10.1214/aoms/1177732360.
- [93] The NMDB database (www.nmdb.eu), founded under the European Union's FP7 programme provided the data.

- [94] H. Mavromichalaki, et al., *Ann. Geophys.* 23 (2005) 1–8.
- [95] Data kindly provided by the Univ. of Delaware Dep. of Physics and Astronomy and the Bartol Research Institute.
- [96] TSUMEB Neutron Monitor, North-West University of South Africa. Data accessed via NMDB.eu.
- [97] P. Centre for Space Research, North-West University, Hermanus neutron monitor.
- [98] SILSO World Data Center, International Sunspot Number Monthly Bulletin and online catalogue, 2006–2021.
- [99] X. Bertou, Pierre Auger, Search for Gamma Ray Bursts using the single particle technique at the Pierre Auger Observatory, in: 30th International Cosmic Ray Conference, Vol. 4, 2008, pp. 441–444. [arXiv:0706.1256](https://arxiv.org/abs/0706.1256).
- [100] M. Schimassek, X. Bertou, D. Veberič, R. Engel, SD: Station Trigger Time Distributions, Auger internal note GAP-2020-042 (2020).
- [101] M. Schimassek, X. Bertou, D. Veberič, R. Engel, Trigger effects in SD Rings, Auger internal note GAP-2020-043 (2020).
- [102] M. Schimassek, X. Bertou, D. Veberič, R. Engel, Algorithm for Tagging of Long Signals, Auger internal note GAP-2021-003 (2021).
- [103] M. Schimassek, D. Veberič, X. Bertou, R. Engel, Evidence that SD-Rings are Filled, Auger internal note GAP-2021-037 (2021).
- [104] P. Abreu, et al., Pierre Auger, Downward Terrestrial Gamma-ray Flashes at the Pierre Auger Observatory?, *PoS ICRC2021* (2021) 395. doi : 10.22323/1.395.0395.
- [105] V. Purrello, X. Bertou, Study of strange SD events detected during thunderstorms, Auger internal note GAP-2015-012 (2015).
- [106] R. Colalillo, private communication (2020).
- [107] R. Colalillo, SD discs analysis: split traces, Auger Online Collaboration Meeting, April 2020 (2020).
- [108] X. Bertou, Looking for a correlation between cosmic rays and lightning, Auger internal note GAP-2015-068 (2015).
- [109] D. Nitz, private communication (2020).
- [110] I. Mariş, M. Roth, 65 stations in one shot. Is this lightning?, Auger internal note GAP-2006-079 (2006).
- [111] Roberta Colalillo, Sd-disks footprint checks, Cosmo-Geo Meeting, 15 April 2021.
- [112] B. Genolini, T. N. Trung, J. Pouthas, Baseline stability of the surface detector PMT base, Auger internal note GAP-2003-051 (2003).
- [113] S. Maldera, G. Navarra, Restoring the SD PMT signals through the Laplace transforms, Auger internal note GAP-2005-006 (2005).
- [114] P. Billoir, Proposition to improve the local trigger of Surface Detector for low energy showers, Auger internal note GAP-2009-179 (2009).

- [115] P. Billoir, New proposal to improve the local trigger of the Surface Detector, Auger internal note GAP-2011-089 (2011).
- [116] R. Sato, D. Nitz, Uub-random traces, private communication. (2021).
- [117] R. Sato, Sd operation, Online Collaboration Meeting, March 2021. (2021).
- [118] A. Attanasio, P. R. Krehbiel, C. L. da Silva, Griffiths and Phelps lightning initiation model, revisited, *J. Geophys. Research Atm* 124 (14) (2019) 8076–8094. doi:<https://doi.org/10.1029/2019JD030399>.
- [119] Y. Raizer, *Gas discharge physics*, Springer-Verlag, 1991.
- [120] A. Gurevich, G. Milikh, R. Roussel-Dupre, Runaway electron mechanism of air breakdown and preconditioning during a thunderstorm, *Phys. Lett. A* 165 (5) (1992) 463–468. doi:[https://doi.org/10.1016/0375-9601\(92\)90348-P](https://doi.org/10.1016/0375-9601(92)90348-P).
- [121] A. Gurevich, K. Zybin, R. Roussel-Dupre, Lightning initiation by simultaneous effect of runaway breakdown and cosmic ray showers, *Phys. Lett. A* 254 (1) (1999) 79–87. doi:[https://doi.org/10.1016/S0375-9601\(99\)00091-2](https://doi.org/10.1016/S0375-9601(99)00091-2).
- [122] J. R. Dwyer, A fundamental limit on electric fields in air, *Geophys. Research Lett.* 30 (20) (2003). doi:<https://doi.org/10.1029/2003GL017781>.
- [123] C. Rutjes, U. Ebert, S. Buitink, O. Scholten, T. N. Trinh, Generation of seed electrons by extensive air showers, and the lightning inception problem including narrow bipolar events, *J. Geophys. Research Atm* 124 (13) (2019) 7255–7269. doi:<https://doi.org/10.1029/2018JD029040>.
- [124] D. Nitz, Sd trigger validation roundup, Auger Analysis Foundation Online Meeting, 21 October 2021.
- [125] P. Abreu, et al., Pierre Auger, The Exposure of the Hybrid Detector of the Pierre Auger Observatory, *Astropart. Phys.* 34 (2011) 368–381. arXiv:1010.6162, doi:10.1016/j.astropartphys.2010.10.001.
- [126] M. Stadelmaier, S. Hahn, et al., The observer data v17r0, private communication.
- [127] J. W. Cronin, Carmen Miranda’s Beautiful Song revised, Auger internal note GAP-2001-040 (2001).
- [128] N. W. S. Demetriades, M. J. Murphy, J. A. Cramer, Validation of Vaisala’s global lightning dataset (gld360) over the continental United States, 21st International Lightning Detection Conference, Orlando, Florida, USA. (2010).
URL <https://www.vaisala.com/sites/default/files/documents/6.Demetriades%2C%20Murphy%2C%20Cramer.pdf>
- [129] K.-D. Merenda, *Studies of elves and their connection to lightning with the Pierre Auger Observatory*, phdthesis (2020).
- [130] G. J. Feldman, R. D. Cousins, Unified approach to the classical statistical analysis of small signals, *Phys. Rev. D* 57 (1998) 3873 – 3889.
- [131] Z. Szadkowski, A PLD200 First Level Fast Trigger in the Pierre Auger Observatory Surface Detector, Auger internal note GAP-2002-050 (2002).

- [132] R. Brun, F. Rademakers, ROOT: An object oriented data analysis framework, *Nucl. Instrum. Meth. A* 389 (1997) 81–86. doi : [10.1016/S0168-9002\(97\)00048-X](https://doi.org/10.1016/S0168-9002(97)00048-X).
- [133] F. Rademakers, P. Canal, A. Naumann, O. Couet, L. Moneta, V. Vassilev, S. Linev, D. Piparo, G. GANIS, B. Bellenot, E. Guiraud, G. Amadio, wverkerke, P. Mato, TimurP, M. Tadel, wlv, E. Tejedor, J. Blomer, A. Gheata, S. Hageboeck, S. Roiser, marsupial, S. Wunsch, O. Shadura, A. Bose, CristinaCristescu, X. Valls, R. Isemann, K. Albertsson, root-project/root: v6.20/04 (Apr. 2020). doi : [10.5281/zenodo.3895855](https://doi.org/10.5281/zenodo.3895855).
- [134] D. Veberič, Changes to the sd trigger simulation in offline, Auger Analysis Foundation online meeting 8 Feb 2021 (2021).
- [135] D. Schmidt, SSD Integration Window, Auger Analysis Meeting Nijmegen (2019).
- [136] J. Linsley, Subluminal pulses from cosmic ray air showers, *J. Phys. G* 10 (1984) 191–195. doi : <https://doi.org/10.1088/0305-4616/10/8/005>.
- [137] A. Erlykin, Neutron ‘Thunder’ Accompanying an Extensive Air Shower, *Bull. Russ. Acad. Sci. Phys.* 71 (4) (2007) 191–195. doi : <https://doi.org/10.3103/S1062873807040302>.
- [138] M. Anderson, et al., Measurement of neutron-proton capture in the SNO+ water phase, *Phys. Rev. C* 102 (041002) (2020).
- [139] V. Aushev, et al., *Izv. Akad. Nauk. Ser. Fiz.* 661 (486) (1997).
- [140] R. Engel, A. Ferrari, M. Roth, M. Schimassek, D. Schmidt, D. Veberic, Neutron production in extensive air showers, *PoS ICRC2021* (2021) 492. doi : [10.22323/1.395.0492](https://doi.org/10.22323/1.395.0492).
- [141] T. T. Böhlen, F. Cerutti, M. P. W. Chin, A. Fassò, A. Ferrari, P. G. Ortega, A. Mairani, P. R. Sala, G. Smirnov, V. Vlachoudis, The FLUKA Code: Developments and Challenges for High Energy and Medical Applications, *Nucl. Data Sheets* 120 (2014) 211–214. doi : [10.1016/j.nds.2014.07.049](https://doi.org/10.1016/j.nds.2014.07.049).
- [142] The Pierre Auger Observatory: Contributions to the 36th International Cosmic Ray Conference (ICRC 2019): Madison, Wisconsin, USA, July 24- August 1, 2019. arXiv: [1909.09073](https://arxiv.org/abs/1909.09073).
- [143] D. Veberic, Estimation of the Total Signal in Saturated Stations of Pierre Auger Surface Detector, in: 33rd International Cosmic Ray Conference, 2013, p. 0633.
- [144] A. Coleman, The New Triggers Settings, Auger internal note GAP-2018-001 (2018).
- [145] P. Photonics, photomultiplier tubes – principles & applications, 1994.
- [146] R. Colalillo, et al., The naples shower library, private communication. URL <https://www.auger.unam.mx/AugerWiki/NapoliLibrary>
- [147] T. Pierog, I. Karpenko, J. M. Katzy, E. Yatsenko, K. Werner, EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider, *Phys. Rev. C* 92 (3) (2015) 034906. arXiv: [1306.0121](https://arxiv.org/abs/1306.0121), doi : [10.1103/PhysRevC.92.034906](https://doi.org/10.1103/PhysRevC.92.034906).
- [148] A. Taboada, B. Manning, D. Schmidt, Ssd preproduction array data, private communication (2021).

- [149] A. Aab, et al., Pierre Auger, Reconstruction of events recorded with the surface detector of the Pierre Auger Observatory, *JINST* 15 (10) (2020) P10021. arXiv:2007.09035, doi:10.1088/1748-0221/15/10/P10021.
- [150] A. Ferrari, private communication (2021).
- [151] D. Tilley, H. Weller, C. Cheves, Energy levels of light nuclei $a = 16-17$, *Nuclear Physics A* 564 (1) (1993) 1–183. doi:https://doi.org/10.1016/0375-9474(93)90073-7.
- [152] A. Aab, et al., Pierre Auger, Muons in Air Showers at the Pierre Auger Observatory: Mean Number in Highly Inclined Events, *Phys. Rev. D* 91 (3) (2015) 032003, [Erratum: *Phys.Rev.D* 91, 059901 (2015)]. arXiv:1408.1421, doi:10.1103/PhysRevD.91.032003.
- [153] N. Dhital, et al., We are all the Cosmic-Ray Extremely Distributed Observatory, in: *Proc. 35th Int. Cosmic Ray Conf., Busan, Korea, 2017*.
- [154] R. Duda, P. Hart, Use of the Hough transformation to detect lines and curves in pictures, *Communications of the ACM* 15 (11) (1972). doi:https://doi.org/10.1145/361237.361242.
- [155] D. Veberič, M. Roth, SD Reconstruction Manual, Auger internal note GAP-2005-035 (2005).
- [156] A. Aab, et al., Pierre Auger, Large-scale cosmic-ray anisotropies above 4 EeV measured by the Pierre Auger Observatory, *Astrophys. J.* 868 (1) (2018) 4. arXiv:1808.03579, doi:10.3847/1538-4357/aae689.
- [157] A. Aab, et al., Pierre Auger, Cosmic-ray anisotropies in right ascension measured by the Pierre Auger Observatory, *Astrophys. J.* 891 (2020) 142. arXiv:2002.06172, doi:10.3847/1538-4357/ab7236.
- [158] M. Schimassek, D. Veberič, R. Engel, Useful Parameterisations of SD Simulations, Auger internal note GAP-2019-075 (2019).
- [159] W. A. Rolke, A. M. Lopez, J. Conrad, Limits and confidence intervals in the presence of nuisance parameters, *Nucl. Instrum. Meth. A* 551 (2005) 493–503. arXiv:physics/0403059, doi:10.1016/j.nima.2005.05.068.
- [160] R. Clay, J. Singh, CREDO, A search for bursts at 0.1 PeV with a small air shower array, *PoS ICRC2021* (2021) 298. doi:10.22323/1.395.0298.
- [161] M. Ester, et al., A density-based algorithm for discovering clusters in large spatial databases with noise, *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD-96)* (1996) 226–231.
- [162] T. W. Anderson, D. A. Darling, Asymptotic Theory of Certain "Goodness of Fit" Criteria Based on Stochastic Processes, *Ann. Math. Stat.* 23 (2) (1952) 193 – 212. doi:10.1214/aoms/1177729437.