

RADBOD UNIVERSITEIT NIJMEGEN



FACULTEIT DER NATUURWETENSCHAPPEN, WISKUNDE EN INFORMATICA

Synchronisation Methods for Air Shower Radio Detectors

MASTER'S THESIS PHYSICS AND ASTRONOMY

Author:


Eric Teunis DE BOONE

Supervisor:

dr. Harm SCHOORLEMMER

Second Reader:

dr. Katharine MULREY



November 2023

A digital copy of this document can be found at
<https://etdeboone.nl/masters-thesis/thesis.pdf>

The L^AT_EX source, including figures, can be found at
<https://gitlab.science.ru.nl/mthesis-edebone/m.internship-documentation>
— or —
<https://etdeboone.nl/masters-thesis/documentation>

Code for generating figures, as well as the beacon synchronising pipeline can be found at
<https://gitlab.science.ru.nl/mthesis-edebone/m-thesis-introduction/>
— or —
https://etdeboone.nl/masters-thesis/airshower_beacon_simulation.

Contents

1	An Introduction to Cosmic Rays and Extensive Air Showers	1
2	Air Shower Radio Interferometry	7
2.1	Radio Interferometry	8
3	Waveform Analysis Techniques	11
3.1	Fourier Transforms	12
3.2	Cross-Correlation	14
4	Synchronising Detectors with a Beacon Signal	15
4.1	The Synchronisation Problem	16
4.1.1	Sine Synchronisation	17
4.1.2	Array synchronisation	17
4.2	Pulse Beacon	18
4.2.1	Timing accuracy	21
4.3	Sine Beacon	23
4.3.1	Timing accuracy	25
5	Single Sine Synchronisation	29
5.1	Air Shower simulation	31
5.2	k -finding	33
5.3	Strategy / Result	36
6	GRAND signal chain characterisation	41
7	Conclusion	47
A	Supplementary Figures	49
B	Random Phasor Sum Distribution	55
	Bibliography	57

Chapter 1.

An Introduction to Cosmic Rays and Extensive Air Showers

The Earth is bombarded with a variety of energetic, extra-terrestrial particles. The energies of these particles extend over many orders of magnitude (see Figure 1.1). The flux of these particles decreases exponentially with increasing energy. For very high energies, above 10^6 GeV, the flux approaches one particle per square meter per year, further decreasing to a single particle per square kilometer per year for Ultra High Energies (UHE) at 10^{10} GeV.

At these high energies, the incoming particles are primarily cosmic rays¹, atomic nuclei typically ranging from protons ($Z = 1$) up to iron ($Z = 26$). Because these are charged, the various magnetic fields they pass through will deflect and randomise their trajectories. Of course, this effect is dependent on the strength and size of the magnetic field and the speed of the particle. It is therefore only at the very highest energies that the direction of an initial particle might be used to (conservatively) infer the direction of its origin.

The same argument (but in reverse) can be used to explain the steeper slope from the “knee” (10^6 GeV) onwards in Figure 1.1. The acceleration of cosmic rays equally requires strong and sizeable magnetic fields. Size constraints on the Milky Way lead to a maximum energy for which a cosmic ray can still be contained in our galaxy. It is thus at these energies that we can distinguish between galactic and extra-galactic origins.

Other particles at these energies include photons and neutrinos, which are not charged. Therefore, these particle types do not suffer from magnetic deflections and have the potential to reveal their source regions. Unfortunately, aside from both being much less frequent, photons can be absorbed and created by multiple mechanisms, while neutrinos are notoriously hard to detect due to their weak interaction.

When a cosmic ray with an energy above 10^3 GeV comes into contact with the atmosphere, secondary particles are generated, forming an Extensive Air Shower (EAS). This air shower consists of a cascade of interactions producing more particles that subsequently undergo further interactions. Thus, the number of particles rapidly increases further down the air shower. This happens until the mean energy per particle is sufficiently lowered from whereon these particles are absorbed in the atmosphere.

Figure 1.2 shows the number of particles as a function of atmospheric depth where 0 g/cm² corresponds with the top of the atmosphere. The atmospheric depth at which this number of particles reaches its maximum is called X_{\max} .

¹These are therefore known as Ultra High Energy Cosmic Rays (UHECRs).

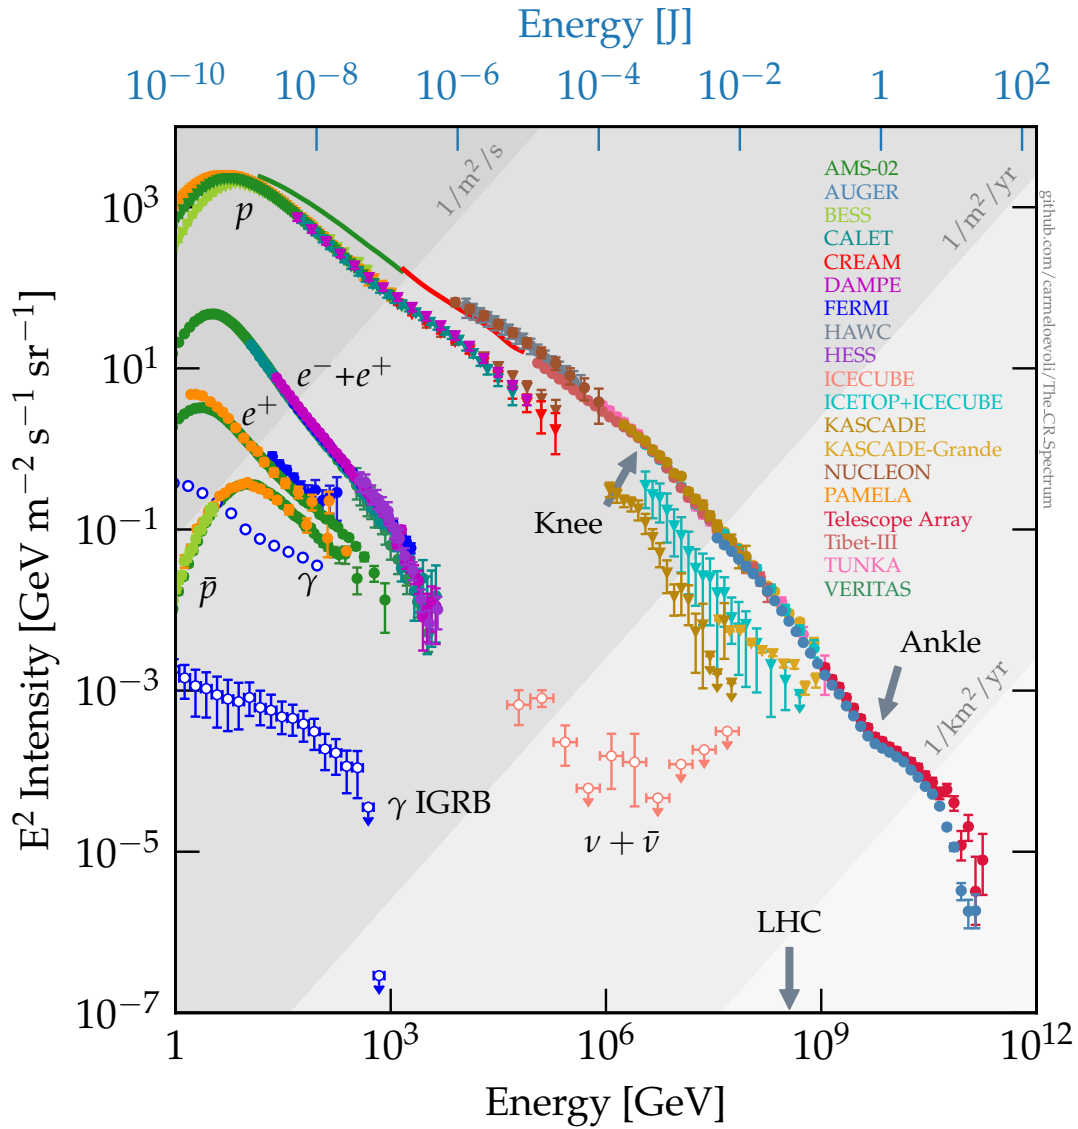


Figure 1.1: From [1]. The diffuse cosmic ray spectrum (upper line) as measured by various experiments. The intensity and fluxes can generally be described by rapidly decreasing power laws. The grey shading indicates the order of magnitude of the particle flux, such that from the ankle onwards ($E > 10^9$ GeV) the flux reaches 1 particle per square kilometer per year.




Figure 1.2: From H. Schoorlemmer. Shower development as a function of atmospheric depth for an energy of 10^{19} eV. Typically, iron- and proton-induced air showers have a difference in $\langle X_{\max} \rangle$ of 100 g/cm^2 [2]. For air showers from photons this is even further down the atmosphere. They are, however, much more rare than cosmic rays.

In Figure 1.2, X_{\max} is different for the air showers generated by a photon, a proton or an iron nucleus. Typically, heavy nuclei have their first interaction higher up in the atmosphere than protons, with photons penetrating the atmosphere even further. Therefore, accurate measurements of X_{\max} allow to statistically discriminate between photons, protons and iron nuclei.

The initial particle type also influences the particle content of an air shower. Depending on the available interaction channels, we distinguish three components in air showers: the hadronic, electromagnetic and muonic components. Each component shows particular development and can be related to different observables of the air shower.

For example, detecting a large hadronic component means the initial particle has access to hadronic interactions (creating hadrons such as pions, kaons, etc.) which is a typical sign of a cosmic ray. In contrast, for an initial photon, which cannot interact hadronically, the energy will be dumped into the electromagnetic part of the air shower, mainly producing electrons, positrons and photons.

Finally, any charged pions created in the air shower will decay into muons while still in the atmosphere, thus comprising the muonic component. The lifetime, and ease of penetration of relativistic muons allow them to propagate to the Earth's surface, even if other particles have decayed or have been absorbed in the atmosphere. These are therefore prime candidates for air shower detectors on the Earth's surface.

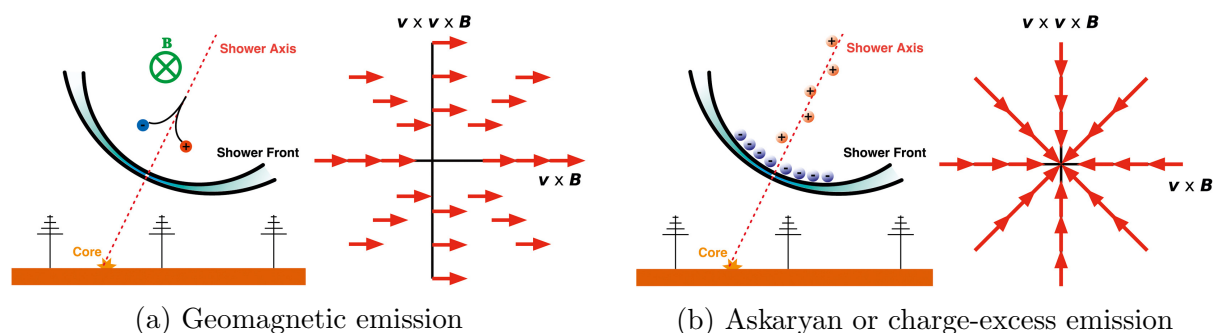


Figure 1.3: From [4, 5] The Radio Emission mechanisms and the resulting polarisations of the radio signal: (a) geomagnetic and (b) charge-excess. See text for explanation.

Processes in an air showers also generate radiation that can be picked up as coherent radio signals. Due to the magnetic field of the Earth, the electrons in the air shower generate radiation. Termed geomagnetic emission in Figure 1.3, this has a polarisation that is dependent on the magnetic field vector (\vec{B}) and the air shower velocity (\vec{v}). An additional mechanism emitting radiation was theorised by Askaryan[3]. Due to the large inertia of the positively charged ions with respect to their light, negatively charged electrons, a negative charge excess is created. In turn, this generates radiation that is polarised radially towards the shower axis (see Figure 1.3).

Due to charged particles moving relativistically through the refractive atmosphere, the produced radiation is concentrated on a cone-like structure. On the surface, this creates a ring called the Cherenkov-ring. On this ring, a peculiar inversion happens in the time-domain of the air shower signals. Outside the ring, radiation from the top of the air shower arrives earlier than radiation from the end of the air shower, whereas this is reversed inside the ring. Consequently, the radiation received at the Cherenkov-ring is maximally coherent, being concentrated in a small time-window. It is therefore crucial for radio detection to obtain measurements in this region.

As mentioned, the flux at the very highest energy is in the order of one particle per square kilometer per century (see Figure 1.1). Observatories therefore have to span huge areas to gather decent statistics at these highest energies on a practical timescale. In recent and upcoming experiments, such as the Pierre Auger Observatory (Auger)[2] and the Giant Radio Array for Neutrino Detection (GRAND)[6], the approach is typically to instrument a large area with a (sparse) grid of detectors to detect the generated air shower. With distances up to 1.5 km (Auger), the detectors therefore have to operate in a self-sufficient manner with only wireless communication channels and timing provided by Global Navigation Satellite System (GNSS).

In the last two decades, with the advent of advanced electronics, the detection using radio antennas has received significant attention. Analysing air showers using radio interferometry requires a time synchronisation of the detectors to an accuracy in the order of 1 ns[7] (see Chapter 2 for further details). Unfortunately, this timing accuracy is not continuously achieved by GNSSs, if at all. For example, in the Auger Engineering Radio Array (AERA), this was found to range up to multiple tens of nanoseconds over the course of a single day[8].

This thesis investigates a relatively straightforward method (and its limits) to improve the timing accuracy of air shower radio detectors by using an additional radio signal called a beacon. It is organised as follows.

First, an introduction to radio interferometry is given in Chapter 2. This will be used later on and gives an insight into the timing accuracy requirements.

Chapter 3 reviews some typical techniques to analyse waveforms and to obtain timing information from them.

In Chapter 4, the concept of a beacon transmitter is introduced to synchronise an array of radio antennas. It demonstrates the achievable timing accuracy for a sine and pulse beacon using the techniques described in the preceding chapter.

A degeneracy in the synchronisation is encountered when the timing accuracy of the GNSS is in the order of the periodicity of a continuous beacon. Chapter 5 establishes a method using a single sine wave beacon while using the radio interferometric approach to observe an air shower and correct for this effect.

Finally, Chapter 6 investigates some possible limitations of the current hardware of GRAND and its ability to record and reconstruct a beacon signal.

Chapter 2.

Air Shower Radio Interferometry

The radio signals emitted by an EAS (see Chapter 1) can be recorded by radio antennas. For suitable frequencies, an array of radio antennas can be used as an interferometer. Therefore, air showers can be analysed using radio interferometry. Note that since the radio waves are mainly caused by processes involving electrons, any derived properties are tied to the electromagnetic component of the air shower.

In Reference [7], a technique was developed to obtain properties of an air shower using radio interferometry.¹ It exploits the coherent emissions in the air shower by mapping the power. Such a power mapping (of a simulated air shower) is shown in Figure 2.1. It reveals the air shower in one vertical and three horizontal slices. Analysing the power mapping, we can then infer properties of the air shower such as the shower axis and X_{\max} .

The accuracy of the technique is primarily dependent on the timing accuracy of the detectors. In Figure 2.2, the estimated atmospheric depth resolution as a function of detector synchronisation is shown as simulated for different inclinations of the air shower. For detector synchronisations under 2 ns, the atmospheric depth resolution is competitive with techniques from fluorescence detectors ($\sigma(X_{\max})$ 25 g/cm² at Auger [2]). With a difference in $\langle X_{\max} \rangle$ of ~ 100 g/cm² between iron and proton initiated air showers, this depth of shower maximum resolution allows to study the mass composition of cosmic rays. However, for worse synchronisations, the X_{\max} resolution for interferometry degrades linearly.

An advantage of radio antennas with respect to fluorescence detectors is the increased duty-cycle. Fluorescence detectors require clear, moonless nights, resulting in a duty-cycle of about 10% whereas radio detectors have a near permanent duty-cycle.

¹Available as a python package at <https://gitlab.com/harmscho/asira>.




Figure 2.1: From [7]. Radio interferometric power analysis of a simulated air shower. *a)* shows the normalised power of $S(\vec{x})$ mapped onto a vertical planer, while *b)*, *c)* and *d)* show the horizontal slices on different heights. On *b)*, *c)* and *d)*, the orange and blue dot indicate the true shower axis and the maximum power respectively.




Figure 2.2: From [7]. X_{\max} resolution as a function of detector-to-detector synchronisation. Note that this figure shows a first-order effect with values particular to the antenna density of the simulated array.

2.1 Radio Interferometry

Radio interferometry exploits the coherence of wave phenomena.

In a radio array, each radio antenna records its ambient electric field. A simple interferometer can be achieved by summing the recorded waveforms S_i with appropriate time delays $\Delta_i(\vec{x})$ to compute the coherency of a waveform at \vec{x} ,

$$S(\vec{x}, t) = \sum_i S_i(t + \Delta_i(\vec{x})). \quad (2.1)$$

The time delays $\Delta_i(\vec{x})$ are dependent on the finite speed of the radio waves. Being an electromagnetic wave, the instantaneous velocity v depends solely on the refractive index n of the medium as $v = \frac{c}{n}$. In general, the refractive index of air is dependent on factors such as the pressure and temperature of the air the signal is passing through, and the frequencies of the signal.

The time delay due to propagation can be written as

$$\Delta_i(\vec{x}) = \frac{|\vec{x} - \vec{a}_i|}{c} n_{\text{eff}}, \quad (2.2)$$

where n_{eff} is the effective refractive index over the trajectory of the signal.

Note that unlike in astronomical interferometry, the source cannot be assumed at infinity, instead it is close-by (see Figure 2.3a). Therefore the time delays for each test location \vec{x} have to be computed separately.




Figure 2.3: Schematic of radio interferometry (a) and the overlap between the recorded waveforms at the source location S_0 (b) and a far away location (c). Δ_i corresponds to the time delay per antenna from (2.2).

Features in the summed waveform $S(\vec{x})$ are enhanced according to the coherence of that feature in the recorded waveforms with respect to the time delays.

Figures 2.3b and 2.3c show examples of this effect for the same recorded waveforms. At the true source location, the recorded waveforms align and sum coherently to result in a summed waveform with enhanced features and amplitudes. Meanwhile, at a far away location, the waveforms sum incoherently resulting in a summed waveform with low amplitudes and without clear features.

An additional effect of interferometry is the suppression of noise particular to individual antennas as this adds up incoherently. The signal in the summed waveform grows linearly with the number of detectors, while the incoherent noise in that same waveform scales with the square root of the number of detectors.

In the technique from [7], the summed waveform $S(\vec{x})$ is computed for multiple locations. For each location, the power in $S(\vec{x})$ is determined to create a power distribution. An example of this power distribution of $S(\vec{x})$ is shown in Figure 2.1.

The region of high power identifies strong coherent signals related to the air shower. By mapping this region, the shower axis and shower core can be resolved. Later, with the shower axis identified, the power along the axis is used to compute X_{\max} .

Chapter 3.

Waveform Analysis Techniques

Radio antennas are sensitive to changes in their surrounding electric fields. The polarisation of the electric field that a single antenna can record is dependent on the geometry of this antenna. Therefore, in experiments such as Auger or GRAND, multiple antennas are incorporated into a single unit to obtain complementary polarisation recordings. Additionally, the shape and size of antennas affect how well the antenna responds to certain frequency ranges, resulting in different designs meeting different criteria.

In each radio detector, the antenna presents its signals to an Analog-to-Digital Converter (ADC) as fluctuating voltages. In turn, the ADC records the analog signals with a specified samplerate f_s resulting in a sequence of digitised voltages or waveform. The n -th sample in this waveform is then associated with a time $t[n] = t[0] + n/f_s = t[0] + n \cdot \Delta t$ after the initial sample at $t[0]$.

The sampling is limited by the ADC's Nyquist frequency at half its sampling rate. In addition, various frequency-dependent backgrounds can be reduced by applying a band-pass filter before digitisation. For example, in AERA and in AugerPrime's radio detector [9], the filter attenuates all of the signal except for the frequency interval between 30–80 MHz. In addition to a band-pass filter, more complex filter setups are used to remove unwanted components or introduce attenuation at specific frequencies. For example, in GRAND [6], the total frequency band ranges from 20 MHz to 200 MHz. such that the FM broadcasting band (87.5 MHz–108 MHz) falls within this range. Therefore, notch filters have been introduced to suppress signals in this band.

From the above it is clear that the digitised waveform is a measurement of the electric field that is sequentially convolved with the antenna's and filter's response. Thus to reconstruct properties of the electric field signal from the waveform, both responses must be known.

Different methods are available for the analysis of the waveform, and the antenna and filter responses. A key aspect is determining the frequency-dependent amplitudes (and phases) in the measurements to characterise the responses and, more importantly, select signals from background. With Fourier Transforms (FTs), these frequency spectra can be produced. This technique is especially important for the sine wave beacon of Section 4.3, as it forms the basis of the phase measurement.

The detection and identification of more complex time-domain signals can be achieved using the cross correlation, which is the basis for the pulsed beacon method of Section 4.2.

3.1 Fourier Transforms

FTs allow for a frequency-domain representation of a time-domain signal. In the case of radio antennas, it converts a time-ordered sequence of voltages into a set of complex amplitudes that depend on frequency. By evaluating the FT at appropriate frequencies, the frequency spectrum of a waveform is calculated. This method then allows to modify a signal by operating on its frequency components, i.e. removing a narrow frequency band contamination within the signal.

The continuous Fourier Transform takes the form

$$X(f) = \int_{-\infty}^{\infty} dt x(t) e^{-i2\pi ft}. \quad (3.1)$$

It decomposes the signal $x(t) \in \mathcal{R}$ into plane waves with complex-valued amplitude $X(f)$ at frequency f .

From the complex amplitude $X(f)$, the phase $\phi(f)$ and amplitude $A(f)$ are calculated as

$$\phi(f) = \arg(X(f)), \quad \text{and} \quad A(f) = 2|X(f)|.$$

Note the factor 2 in this definition of the amplitude. It is introduced to compensate for expecting a real valued input signal $x(t) \in \mathcal{R}$ and mapping negative frequencies to their positive equivalents.

When $x(t)$ is sampled at discrete times, the integral of (3.1) is discretized in time to result in the Discrete Time Fourier Transform (DTFT):

$$X(f) = \sum_{n=0}^{N-1} x(t[n]) e^{-i2\pi ft[n]} \quad (3.2)$$

where $x(t)$ is sampled a finite number of times N at times $t[n]$. Note that the amplitude $A(f)$ will now scale with the number of samples N , and thus should be normalised to $A(f) = 2|X(f)|/N$.

Considering a finite sampling size N and periodicity of the signal, the bounds of the integral in (3.1) have collapsed to $t[0]$ up to t_{N-1} . It follows that the lowest resolvable frequency is $f_{\text{lower}} = 1/T = 1/(t_{N-1} - t[0])$.

Additionally, when the sampling of $x(t)$ is equally spaced, the $t[n]$ terms can be written as a sequence, $t[n] - t[0] = n\Delta t = n/f_s$, with f_s the sampling frequency. Here the highest resolvable frequency is limited by the Nyquist frequency.

Implementing the above decomposition of $t[n]$, (3.2) can be rewritten in terms of multiples of the sampling frequency $f = kf_s/N$, becoming the Discrete Fourier Transform (DFT)

$$X(k) = e^{-i2\pi ft[0]} \sum_{n=0}^{N-1} x(t[n]) \cdot e^{-i2\pi \frac{kn}{N}}.$$

The direct computation of this transform takes $2N$ complex multiplications and $2(N-1)$ complex additions for a single frequency k . When computing this transform for all integer $0 \leq k < N$, this amounts to $\mathcal{O}(N^2)$ complex computations. Fast Fourier Transforms (FFTs) are efficient algorithms that derive all $X(0 \leq k < N)$ in $\mathcal{O}(N \log N)$ calculations.




Figure 3.1: *Left:* A waveform sampling a sine wave with white noise. *Right:* The frequency spectrum of the waveform. Comparison of the DTFT and DFT of the same waveform. The DFT can be interpreted as sampling the DTFT at integer multiple of the waveform's sampling rate f_s .

In the previous equations, the resultant quantity $X(f)$ is a complex amplitude. Since a complex plane wave can be linearly decomposed as

$$e^{-ix} = \cos(x) + i \sin(-x) = \operatorname{Re}(e^{-ix}) + i \operatorname{Im}(e^{-ix}),$$

the above transforms can be decomposed into explicit real and imaginary parts as well, i.e., (3.2) becomes

$$\begin{aligned} X(f) &= X_R(f) + iX_I(f) \equiv \operatorname{Re}(X(f)) + i \operatorname{Im}(X(f)) \\ &= \sum_{n=0}^{N-1} x(t[n]) \cos(2\pi ft[n]) - i \sum_{n=0}^{N-1} x(t[n]) \sin(2\pi ft[n]). \end{aligned} \quad (3.3)$$

The normalised amplitude at a given frequency $A(f)$ is calculated from (3.2) as

$$A(f) \equiv \frac{2|X(f)|}{N} = \frac{2\sqrt{X_R(f)^2 + X_I(f)^2}}{N}. \quad (3.4)$$

Likewise, the complex phase at a given frequency $\phi(f)$ is obtained by

$$\phi(f) \equiv \arg(X(f)) = \arctan2(X_I(f), X_R(f)). \quad (3.5)$$

Applying (3.3) to a signal $x(t) = A \cos(2\pi ft + \phi)$ with the above definitions obtains an amplitude A and phase ϕ at frequency f . When the minus sign in the exponent of (3.1) is not taken into account, the calculated phase in (3.5) will have an extra minus sign.

Figure 3.1 shows the frequency spectrum of a simulated waveform that is obtained using either a DFT or a DTFT. It shows that the DFT evaluates the DTFT only at certain frequencies. By missing the correct frequency bin for the sine wave, it estimates both a too low amplitude and the wrong phase for the input function.

When calculating the DTFT for multiple inputs which share both an equal number of samples N and equal sampling frequencies f_s , the sin and cos terms in (3.3) are the same for a single frequency f up to an overall phase which is dependent on $t[0]$. Therefore, at the cost of an increased memory allocation, these terms can be precomputed, reducing the number of real multiplications to $2N + 1$.




Figure 3.2: *Left:* Two waveforms to be correlated with the second waveform delayed by 5. *Right:* The correlation of both waveforms as a function of the time delay τ . Here the best time delay (red dashed line) is found at 5, which would align the maximum amplitudes of both waveforms in the left pane.

Thus, for static frequencies in a continuous beacon, the coefficients for evaluating the DTFT can be put into the hardware of the detectors, opening the way to efficiently measuring the amplitude and phase in realtime.

3.2 Cross-Correlation

The cross-correlation is a measure of how similar two waveforms $u(t)$ and $v(t)$ are. By introducing a time delay τ in one of the waveforms it turns into a function of this time delay,

$$\text{Corr}(\tau; u, v) = \int_{-\infty}^{\infty} dt u(t) v^*(t - \tau), \quad (3.6)$$

where the integral reduces to a sum for a finite amount of samples in either $u(t)$ or $v(t)$. Still, τ remains a continuous variable. Figure 3.2 illustrates how the best time delay τ between two waveforms can thus be found by finding the maximum cross-correlation.

In reality, both waveforms have a finite size, also reducing the time delay τ resolution to the highest sampling rate of the two waveforms. When the sampling rates are equal, the time delay variable is effectively shifting one waveform by a number of samples.

Techniques such as upsampling or interpolation can be used to effectively change the sampling rate of a waveform up to a certain degree.

Since zero-valued samples do not contribute to the integral of (3.6), they can be freely added (or ignored) to a waveform when performing the calculations. This means two waveforms of different sampling rates can be correlated when the sampling rates are integer multiples of each other, simply by zero-stuffing the slowly sampled waveform. This allows to approximate an analog time delay between two waveforms when one waveform is sampled at a very high rate as compared to the other.

Chapter 4.

Synchronising Detectors with a Beacon Signal

The detection of extensive air showers uses detectors distributed over large areas. Solutions for precise timing (< 0.1 ns) over large distances exist. Initially developed for fibre-optic setups, White Rabbit [10] is also being investigated to be used as a direct wireless time dissemination system [11].

However, the combination of large distances and the number of detectors make it prohibitively expensive to realise such a setup for UHECR detection. For this reason, the time synchronisation of these autonomous stations is typically performed with a GNSS clock in each station.

To obtain a competitive resolution of the atmospheric shower depth X_{\max} with radio interferometry requires an inter-detector synchronisation of better than a few nanoseconds (see Figure 2.2). The synchronisation defect in AERA using a GNSS was found to range between a few nanoseconds up to multiple tens of nanoseconds over the course of a single day (see [8, Figure 3]). Therefore, an extra timing mechanism must be provided to enable interferometric reconstruction of EASs.

For radio antennas, an in-band solution can be created using the antennas themselves by emitting a radio signal from a transmitter. With the position of the transmitter known, the time delays can be inferred and thus the arrival times at each station individually. This has been successfully employed in AERA reaching an accuracy better than 2 ns [8]. For this section, it is assumed that the transmitter is actively introduced to the array and therefore controlled in terms of produced signals and transmitting power. It is foreseeable that “parasitic” setups, where sources that are not under control of the experiment introduce signals, can be analysed in a similar manner. However, for such signals to work, they must have a well-determined and stable origin. See the next Chapter for one such possible setup in Auger.

The nature of the transmitted radio signal, hereafter beacon signal, affects both the mechanism of reconstructing the timing information and the measurement of the radio signal for which the antennas have been designed. Depending on the stability of the station clock, one can choose for employing a continuous beacon (e.g. a sine wave) or one that is emitted at some interval (e.g. a pulse).

Nonetheless, various sources emit radiation that is also picked up by the antenna on top of the wanted signals. An important characteristic is the ability to separate a beacon signal from noise. Therefore, these analysis methods must be performed in the presence of noise.

A simple noise model is given by gaussian noise in the time-domain which is associated to

many independent random noise sources. Especially important is that this noise model will affect any phase measurement depending on the strength of the beacon with respect to the noise level, without introducing a frequency dependence, i.e. white noise.

In the following, the synchronisation scheme for both the continuous and the recurrent beacon are elaborated upon. Before going in-depth on the synchronisation using either of such beacons, the synchronisation problem is worked out.

4.1 The Synchronisation Problem

An in-band solution for synchronising the detectors is effectively a reversal of the method of interferometry in Section 2. The distance between the transmitter T and the antenna A_i incurs a time delay caused by the finite propagation speed of the radio signal (see the Δ_i term in (2.2)). In this chapter it will be denoted as $(t_d)_i$ for clarity.

If the time of emitting the signal at the transmitter t_0 is known, this allows to directly synchronise the transmitter and an antenna since

$$(t'_0)_i = t_0 + (t_d)_i = (\tau_0)_i - (t_c)_i, \quad (4.1)$$

where $(t'_0)_i$ and $(\tau_0)_i$ are respectively the true and measured arrival time of the signal at antenna A_i . The difference between these two terms gives the clock deviation term $(t_c)_i$.

As (4.1) applies for each antenna, two antennas recording the same signal from a transmitter will share the t_0 term. In that case, the differences between the true arrival times $(t'_0)_i$ and propagation delays $(t_d)_i$ of the antennas can be related as

$$\begin{aligned} (\Delta t'_0)_{ij} &\equiv (t'_0)_i - (t'_0)_j \\ &= [t_0 + (t_d)_i] - [t_0 + (t_d)_j]. \\ &= (t_d)_i - (t_d)_j \equiv (\Delta t_d)_{ij} \end{aligned} \quad (4.2)$$

Combining (4.2) and (4.1) then gives the relative clock mismatch $(\Delta t_c)_{ij}$ as

$$\begin{aligned} (\Delta t_c)_{ij} &\equiv (t_c)_i - (t_c)_j \\ &= [(\tau_0)_i - (t'_0)_i] - [(\tau_0)_j - (t'_0)_j] \\ &= [(\tau_0)_i - (\tau_0)_j] - [(t'_0)_i - (t'_0)_j]. \\ &= (\Delta \tau_0)_{ij} - (\Delta t'_0)_{ij} \\ &= (\Delta \tau_0)_{ij} - (\Delta t_d)_{ij} \end{aligned} \quad (4.3)$$

Thus, measuring $(\tau_0)_i$ and determining $(t_d)_i$ for two antennas provides the synchronisation mismatch between them.

Note that t_0 is not required in (4.3) to be able to synchronise two antennas. However, without knowledge on the t_0 of the transmitter, the synchronisation mismatch $(\Delta t_c)_{ij}$ cannot be uniquely attributed to either of the antennas; this scheme only provides relative synchronisation.

4.1.1 Sine Synchronisation

In the case of a sine beacon, its periodicity prevents to differentiate between consecutive periods using the beacon alone. The measured arrival term τ_0 in (4.1) is no longer uniquely defined, since

$$f(\tau_0) = f\left(\frac{\varphi_0}{2\pi}T\right) = f\left(\frac{\varphi_0}{2\pi}T + kT\right), \quad (4.4)$$

where $-\pi < \varphi_0 < \pi$ is the phase of the beacon $f(t)$ at time τ_0 , T the period of the beacon and $k \in \mathbb{Z}$ is an unknown period counter. Of course, this means that the clock defects t_c can only be resolved up to the beacon's period, changing (4.3) to

$$\begin{aligned} (\Delta t_c)_{ij} &\equiv (t_c)_i - (t_c)_j \\ &= (\Delta \tau_0)_{ij} - (\Delta t'_0)_{ij} \\ &= (\Delta \tau_0)_{ij} - (\Delta t_d)_{ij} \\ &= \left[\frac{(\Delta \varphi_0)_{ij}}{2\pi} - \Delta k'_{ij} \right] T - (\Delta t_d)_{ij} \\ &= \left[\frac{(\Delta \varphi_0)_{ij}}{2\pi} - \frac{(\Delta \phi_d)_{ij}}{2\pi} \right] - \Delta k_{ij} T \\ &\equiv \left[\frac{(\Delta \phi_c)_{ij}}{2\pi} \right] T - k_i T. \end{aligned} \quad (4.5)$$

Relative synchronisation of two antennas is thus possible with the caveat of being off by an unknown amount of periods $k_i \in \mathbb{Z}$. Note that in the last step, $k_i = \Delta k_{ij}$ is redefined taking station j as the reference station such that $k_j = 0$.

The correct period k alignment might be found in at least two ways. First, if the timescale of the beacon is much longer than the estimated accuracy of another timing mechanism (such as a GNSS), one can be confident to have the correct period. In AERA for example, multiple sine waves were used amounting to a total beacon period of $\sim 1 \mu\text{s}$ [8, Figure 2]. With an estimated timing accuracy of the GNSS under 50 ns the correct beacon period can be determined, resulting in a unique measured arrival time τ_0 .

A second method consists of using an additional (discrete) signal to declare a unique τ_0 . This relies on the ability of counting how many beacon periods have passed since this extra signal has been recorded. Chapter 5 shows a special case of this last scenario where the period counters are approximated from an extensive air shower.

4.1.2 Array synchronisation

The idea of a beacon is to synchronise an array of antennas. As (4.3) applies for each pair of antennas in the array, all the antennas that record the beacon signal can determine the synchronisation mismatches simultaneously.¹ Taking one antenna as the reference

¹The mismatch terms for any two pairs of antennas sharing one antenna $\{(i, j), (j, k)\}$ allows to find the closing mismatch term for (i, k) since

$$(\Delta t_c)_{ij} + (\Delta t_c)_{jk} + (\Delta t_c)_{ki} = 0$$




Figure 4.1: Reconstruction of a transmitter’s location (tx) or direction using three antennas (a , b , c). For each location, the colour indicates the total deviation from the measured time or phase differences in the array, such that 0 (blue) is considered a valid location of tx . The different baselines allow to reconstruct the direction of an impulsive signal (*left pane*) while a periodic signal (*right pane*) gives rise to a complex pattern (see Appendix A.1 for enhanced size).

antenna with $(t_c)_r = 0$, the mismatches across the array can be determined by applying (4.3) over consecutive pairs of antennas and thus all relative clock deviations $(\Delta t_c)_{ir}$.

As discussed previously, the synchronisation problem is different for a continuous and an impulsive beacon due to the non-uniqueness (in the sine wave case) of the measured arrival time τ_0 . This is illustrated in Figure 4.1 where a three-element array constrains the location of the transmitter using the true timing information of the antennas. It works by finding the minimum deviation between the putative and measured time differences $(\Delta t_{ij}(x), \Delta t_{ij}$ respectively) per baseline (i, j) for each location on a grid.

For a sine signal, comparing the baseline phase differences instead, this results in a highly complex pattern constraining the transmitter’s location.

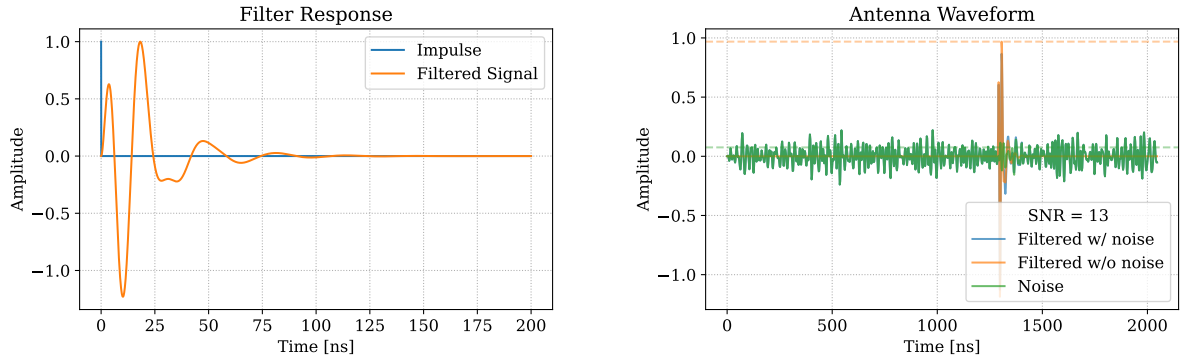
In the former, the mechanism of measuring $(\tau_0)_i$ from the signal has been deliberately left out. The nature of the beacon, being impulsive or continuous, requires different methods to determine this quantity. In the following sections, two separate approaches for measuring the arrival time $(\tau_0)_i$ are examined.

4.2 Pulse Beacon

To synchronise on an impulsive signal, it must be recorded at the relevant detectors. However, it must be distinguished from air shower signals. It is therefore important to choose an appropriate length and interval of the synchronisation signal to minimise dead-time of the detector.

With air shower signals typically lasting in the order of 10 ns, transmitting a pulse of 1 μ s once every second already achieves a simple distinction between the synchronisation and air shower signals and a dead-time below 0.001%.

Schemes using such a “ping” might also be employed between the antennas themselves. Appointing the transmitter role to differing antennas additionally opens the way to



(a) The impulse response of the used filter. Amplitudes are not to scale.

(b) Simulated waveform with noise. Horizontal dashed lines indicate signal and noise level.

Figure 4.2: (a) A single impulse and a simulated filtered signal, using a Butterworth filter, available to the digitiser in a detector. (b) A noisy sampling of the filtered signal. It is derived from the filtered signal by adding filtered gaussian noise.

(self-)calibrating the antennas in the array.

In this section, the idea of using a single pulse as beacon signal is explored.

The detection of a (strong) pulse in a waveform is conceptually simple, and can be accomplished while working fully in the time-domain. Before recording the signal at a detector, the signal at the antenna is typically put through a filter-chain which acts as a band-pass filter. This causes the sampled pulse to be stretched in time (see Figure 4.2a).

We can characterise the response of a filter as the response to an impulse. This impulse response can then be used as a template to match against measured waveforms. In Figure 4.2a, the impulse and the filter's response are shown, where the Butterworth filter band-passes the signal between 30 MHz and 80 MHz.

A measured waveform will consist of the filtered signal in combination with noise. Due to the linearity of filters, a noisy waveform can be simulated by summing the components after separately filtering them. Figure 4.2b shows an example of the waveform obtained when summing these components with a considerable noise component.

Detecting the modelled signal from Figure 4.2a in a waveform can be achieved by finding the correlation (see Section 3.2) between the two signals (see Figure 4.3). The correlation is a measure of how similar two signals $u(t)$ and $v(t)$ are as a function of the time delay τ . The maximum is attained when $u(t)$ and $v(t)$ are most similar to each other. Therefore, this gives a measure of the best time delay τ between the two signals.

When the digitiser samples the filtered signal, time offsets τ smaller than the sampling period $\Delta t = 1/f_s$ cannot be resolved. Still, for many measurements under ideal conditions, one can show that the resolution of the timing asymptotically approaches $\Delta t/\sqrt{12}$. This is an effect of the quantisation of the sampling period, where the time offsets τ are modelled as a uniform distribution in time bins the size of Δt . In that case, the variance of a uniform distribution applies, obtaining this limit.




Figure 4.3: *Top*: The measured waveform and templated filter response from Figure 4.2a. *Bottom*: The (normalised) correlation between the waveform and template as a function of time delay τ . The template is shifted by the time delay found at the maximum correlation (green dashed line), aligning the template and waveform in the top figure.

As can be seen in Figure 4.2a, the impulse response spreads the power of the signal over time. The peak amplitude gives a measure of this power without needing to integrate the signal.

Expecting the noise to be gaussian distributed in the time domain, it is natural to use the root-mean-square (RMS) of its amplitude as a quantity representing the strength of the noise.

Therefore, the signal-to-noise ratio (SNR) will be defined as the maximum amplitude of the filtered signal versus the RMS of the noise amplitudes.

4.2.1 Timing accuracy

From the above, it is clear that both the SNR as well as the sampling rate of the template have an effect on the ability to resolve small time offsets. To further investigate this, we set up a simulation² where templates with different sampling rates are matched to simulated waveforms for multiple SNRs.

First, an “analog” template is rendered at $\Delta t = 10$ fs to be able to simulate small time-offsets. Each simulated waveform samples this “analog” template with $\Delta t = 2$ ns and a randomised time-offset t_{true} .

Second, the matching template is created by sampling the “analog” template at the specified sampling rate (here considered are 0.5 ns, 0.1 ns and 0.01 ns).

Afterwards, simulated waveforms are correlated (see (3.6) in Chapter 3.2) against the matching template, this obtains a best time delay τ per waveform by finding the maximum correlation (see Figure 4.3).

Comparing the best time delay τ with the randomised time-offset t_{true} , we get a time residual $t_{\text{res}} = t_{\text{true}} - \tau$ per waveform.

For weak signals (SNR $\lesssim 2$), the correlation method will often select wrong peaks. Therefore a selection criterion is applied on $t_{\text{res}} < 2\Delta t$ to filter such waveforms and low SNRs are not considered here.

Figure 4.4 shows two histograms ($N = 500$) of the time residuals for two SNRs. Expecting the time residual to be affected by both the quantisation and the noise, we fit a gaussian to the histograms. The width of each such gaussian gives an accuracy on the time offset σ_t that is recovered using the correlation method.

By evaluating the timing accuracies σ_t for some combinations of SNRs and template sampling rates, Figure 4.5 is produced. It shows that, as long as the pulse is (much) stronger than the noise (SNR $\gtrsim 5$), template matching could achieve a sub-nanosecond timing accuracy even if the measured waveform is sampled at a lower rate (here $\Delta t = 2$ ns).

²<https://gitlab.science.ru.nl/mthesis-edboone/m-thesis-introduction/-/tree/main/simulations>

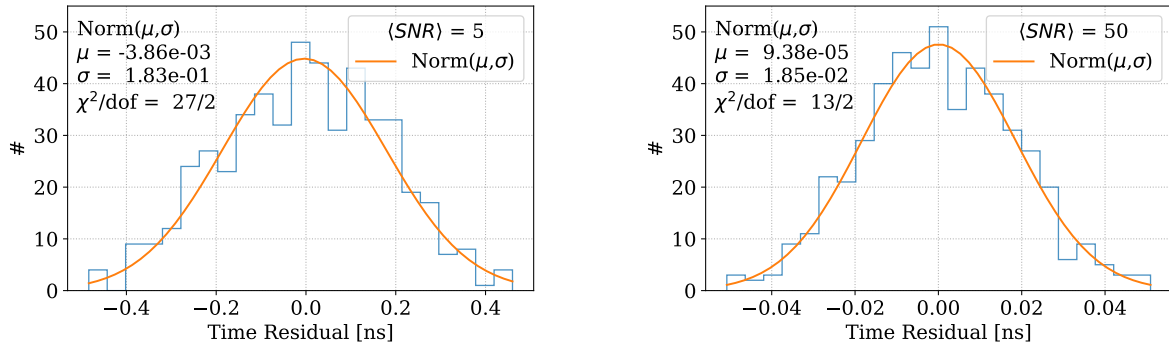


Figure 4.4: Time residuals histograms ($N = 500$) for $\text{SNR} = (5, 50)$ at a template sampling rate of 10 ps.




Figure 4.5: Pulse timing accuracy obtained by matching $N = 500$ waveforms, sampled at 2 ns, to a templated pulse, sampled at $\Delta t = 0.5$ ns (blue), 0.1 ns (orange) and 0.01 ns (green). Dashed lines indicate the asymptotic best time accuracy ($\Delta t/\sqrt{12}$) per template sampling rate.

4.3 Sine Beacon

In the case the stations need continuous synchronisation, a different approach can be taken. Still, the following method can be applied as a non-continuous beacon if required. A continuously emitted beacon will be recorded simultaneously with the signals from air showers. It is therefore important that the beacon does not fully perturb the recording of the air shower signals, but still be prominent enough for synchronising the antennas. By implementing the beacon signal as one or more sine waves, the beacon can be recovered from the waveform using Fourier Transforms (see Section 3.1). It is then straightforward to discriminate a strong beacon from the air shower signals, resulting in a relatively unperturbed air shower recording for analysis after synchronisation.

Note that for simplicity, the beacon in this section will consist of a single sine wave at $f_{\text{beacon}} = 51.53 \text{ MHz}$ corresponding to a period of roughly 20 ns.

The typical Fourier Transform implementation, the Fast Fourier Transform (FFT), finds the amplitudes and phases at frequencies $f_m = m\Delta f$ determined solely by properties of the waveform, i.e. the sampling frequency f_s and the number of samples N in the waveform ($0 \leq m < N$ such that $\Delta f = f_s/(2N)$).

Depending on the frequency content of the beacon, the sampling frequency and the number of samples, one can resort to use such a DFT (3.1). However, if the frequency of interest is not covered in the specific frequencies f_m , the approach must be modified (e.g. by zero-padding or interpolation). Especially when only a single frequency is of interest, a simpler and shorter route can be taken by evaluating the DTFT (3.2) for this frequency directly.

The effect of using a DTFT instead of a FFT for the detection of a sine wave is illustrated in Figure 4.6, where the DTFT displays a higher amplitude than the FFT.

Of course, like the pulse method, the ability to measure the beacon's sine waves is dependent on the amplitude of the beacon in comparison to noise. To quantify this comparison in terms of SNR, we define the signal level to be the amplitude of the frequency spectrum at the beacon's frequency determined by DTFT (the orange line in Figure 4.6), and the noise level as the scaled RMS of all amplitudes in the noise band determined by FFT (blue line in Figure 4.6). Since gaussian noise has Rayleigh distributed amplitudes (see Figure B.1 in Appendix B), this RMS is scaled by $1/\sqrt{2\pi}$.

However, for sine waves, an additional method to increase the SNR is available. In the frequency spectrum, the amplitude with respect to gaussian noise also increases with more samples N in a waveform. Thus, by recording more samples in a waveform, the sine wave is recovered better. This effect can be seen in Figure 4.7 where the signal to noise ratio increases as \sqrt{N} .

Note that the DTFT, as a finite FT, suffers from spectral leakage, where signals at adjacent frequencies influence the ability to resolve the signals separately. Depending on the signal to be recovered, different windowing functions (e.g. Hann, Hamming, etc.) can be applied to a waveform. For simplicity, in this document, no special windowing functions are applied to waveforms.




Figure 4.6: Signal to Noise definition in the frequency domain. Solid lines are the noise (blue) and beacon’s (orange) frequency spectra obtained with a FFT. The noise level (blue dashed line) is the $RMS/\sqrt{2\pi}$ over all frequencies (blue-shaded area). The signal level (orange dashed line) is the amplitude calculated from the DTFT at 51.53 MHz (orange star).

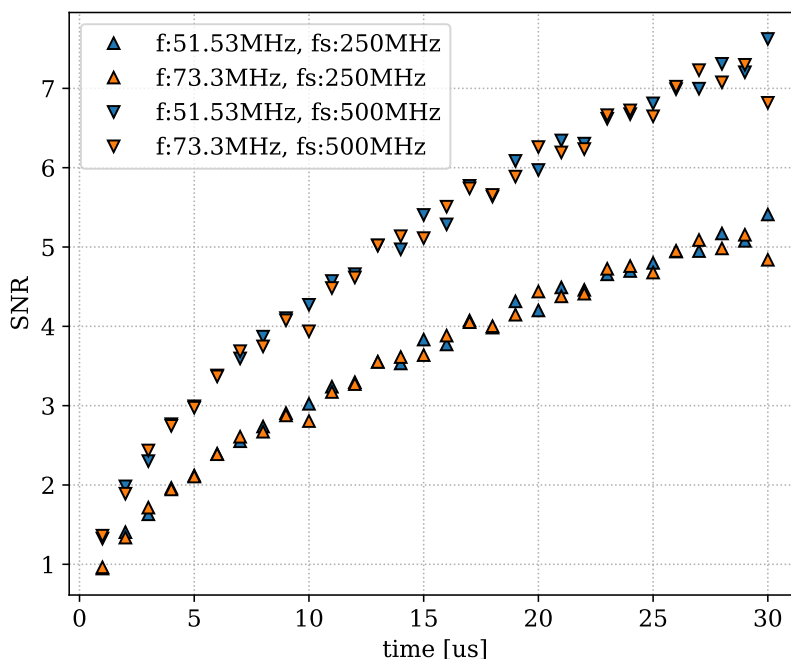


Figure 4.7: Signal to Noise ratio (SNR) as a function of time for waveforms containing only a sine wave and gaussian noise. Note that there is little dependence on the sine wave frequency. The two branches (up and down triangles) differ by a factor of $\sqrt{2}$ in SNR due to their sampling rate.




Figure 4.8: Band-passed waveform containing a sine wave and gaussian time domain noise and the recovered sine wave at 51.53 MHz. Part of the waveform is removed to verify the implementation of the DTFT allowing cut-out samples.

4.3.1 Timing accuracy

The phase measurement of a sine beacon is influenced by other signals in the recorded waveforms. They can come from various sources, both internal (e.g. LNA noise) and external (e.g. galactic background) to the detector.

To investigate the resolution of the phase measurement, we generate waveforms of a sine wave with known, but differing, phases ϕ_{true} . Gaussian noise is added to the waveform in the time-domain, after which the waveform is band-pass filtered between 30 MHz and 80 MHz. The phase measurement of the band-passed waveform is then performed by employing a DTFT. We can compare this measured phase φ with the initial known phase ϕ_{true} to obtain a phase residual $\phi_{\text{res}} = \phi_{\text{true}} - \varphi$.

In Figure 4.8, the band-passed waveform and the measured sine wave are shown. Note that the DTFT allows for an implementation where samples are missing by explicitly using the samples' timestamps. This is illustrated in Figure 4.8 by the cut-out of the waveform.

Figure 4.9 shows two histograms ($N = 100$) of the phase residuals for a medium and a high SNR, respectively. It can be shown that for medium and strong signals, the phase residual will be gaussian distributed (see below). The width of each fitted gaussian in Figure 4.9 gives an accuracy on the phase offset that is recovered using the DTFT.

Note that these distributions have non-zero means, this systematic offset has not been investigated further in this work.

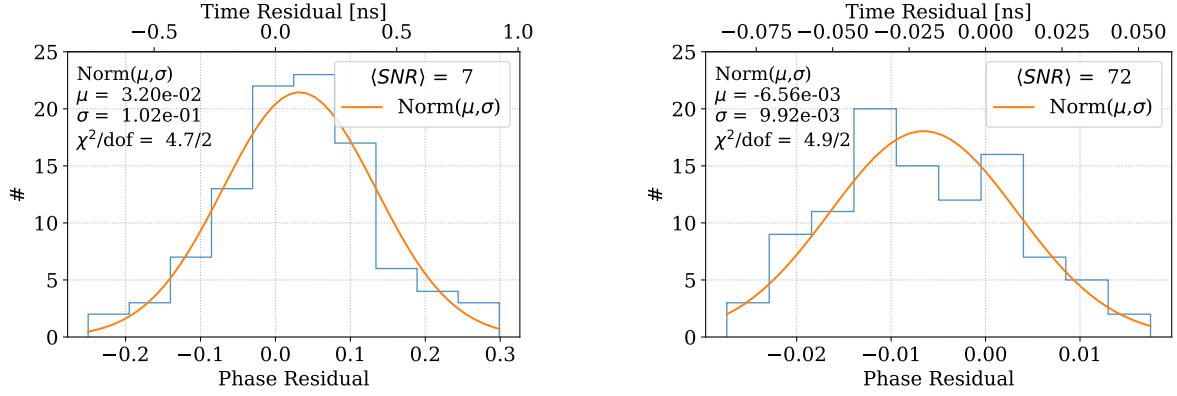


Figure 4.9: Phase residuals histograms ($N = 100$) for $\text{SNR} \sim (7, 70)$. For medium to strong signals the phase residuals sample a gaussian distribution.

For gaussian noise, the measurement of the beacon phase ϕ can be shown to be distributed by the following equation (see Appendix B or [12, Chapter 2.9] for derivation),

$$p_{\Phi}(\phi; s, \sigma) = \frac{e^{-\left(\frac{s^2}{2\sigma^2}\right)}}{2\pi} + \sqrt{\frac{1}{2\pi}} \frac{s}{\sigma} e^{-\left(\frac{s^2}{2\sigma^2} \sin^2 \phi\right)} \frac{\left(1 + \operatorname{erf} \frac{s \cos \phi}{\sqrt{2}\sigma}\right)}{2} \cos \phi, \quad (4.6)$$

where s is the amplitude of the beacon, σ the noise amplitude and $\operatorname{erf} z$ the error function. [12] names this equation ‘‘Constant Phasor plus a Random Phasor Sum’’.

This distribution approaches a gaussian distribution when the beacon amplitude is (much) larger than the noise amplitude. This can be seen in Figure 4.10 where both distributions are shown for a range of SNRs. There, the phase residuals of the simulated waveforms closely follow the distribution.

From Figure 4.10 we can conclude that depending on the SNR, the timing accuracy of the beacon is below 1 ns for our beacon at 51.53 MHz. Since the time accuracy is derived from the phase accuracy with

$$\sigma_t = \frac{\sigma_{\phi}}{2\pi f_{\text{beacon}}}, \quad (4.7)$$

slightly lower frequencies could be used instead, but they would require a comparatively stronger signal to resolve to the same degree. Likewise, higher frequencies are an available method of linearly improving the time accuracy.

However, as mentioned before, the period duplicity restricts an arbitrary high frequency to be used for the beacon. For the 51.53 MHz beacon, the next Chapter 5 shows a method of using an additional signal to counter the period degeneracy of a single sine wave.




Figure 4.10: Phase accuracy (right y-axis) for a sine beacon as a function of signal to noise ratio for waveforms of 10240 samples containing a sine wave at 51.53 MHz and white noise. It can be shown that the phase accuracies (right y-axis) follow a special distribution (4.6) that is well approximated by a gaussian distribution for $\text{SNR} \gtrsim 3$. The time accuracy is converted from the phase accuracy using (4.7). The green dashed line indicates the 1 ns level. Thus, for a beacon at 51.53 MHz and a $\text{SNR} \gtrsim 3$, the time accuracy is better than 1 ns.

Chapter 5.

Single Sine Beacon Synchronisation and Radio Interferometry

As shown in Chapter 4, both impulsive and sine beacon signals can synchronise air shower radio detectors to enable the interferometric reconstruction of extensive air showers. This chapter will focus on using a single sine beacon to synchronise an array due to the simple setup and analysis required for such a beacon. Additionally, at Auger, a public TV-transmitter is broadcasting at 67.25 MHz. This poses an opportunity to use a “free” beacon to synchronise the radio antennas of AERA and AugerPrime (AugerPrime).

Due to the periodicity of sine beacons, the ability to synchronise an array is limited up to the beacon period T . As previously mentioned, the correct periods can be ascertained by choosing a beacon period much longer than the estimated accuracy of another timing mechanism.¹ Likewise, this can be achieved using the beating of multiple frequencies such as the four frequency setup in AERA, amounting to a total period of $> 1 \mu\text{s}$.

In this chapter, a different method of resolving these period mismatches is investigated by recording an impulsive signal in combination with the sine beacon. Figure 5.1 shows the steps of synchronisation using this combination. The extra signal declares a shared time t_0 that is common to the stations, after which the periods can be counted. Note that the period mismatch term Δk_{ij} in (4.5) will be referenced throughout this Chapter as k since we can take station i as reference ($k_i = 0$).

When the beacon transmitter is also used to emit the signal defining t_0 , the number of periods k can be obtained directly from the signal. However, if this calibration signal is sent from a different location, its time delays differ from the beacon’s time delays. For static setups, these time delays can be resolved by measuring the involved distances or by taking measurements of the time delays over time. In dynamic setups, such as for transient signals, the time delays change per event and the distances are not known a priori. The time delays must therefore be resolved from the information of a single event.

As shown in Chapter 4.1.2, an impulsive signal allows to reconstruct the direction of origin in a single event depending on the timing resolution of the array. Synchronising the array with a sine beacon, any clock mismatch is discretized into a number of periods k . This allows to improve the reconstruction by iterating the discrete clock mismatches during reconstruction.

Of course, a limit on the number of periods is required to prevent over-optimisation. In general, they can be constrained using estimates of the accuracy of other timing mechanisms (see below).

¹For reference, GNSS timing is expected to be below 30 ns

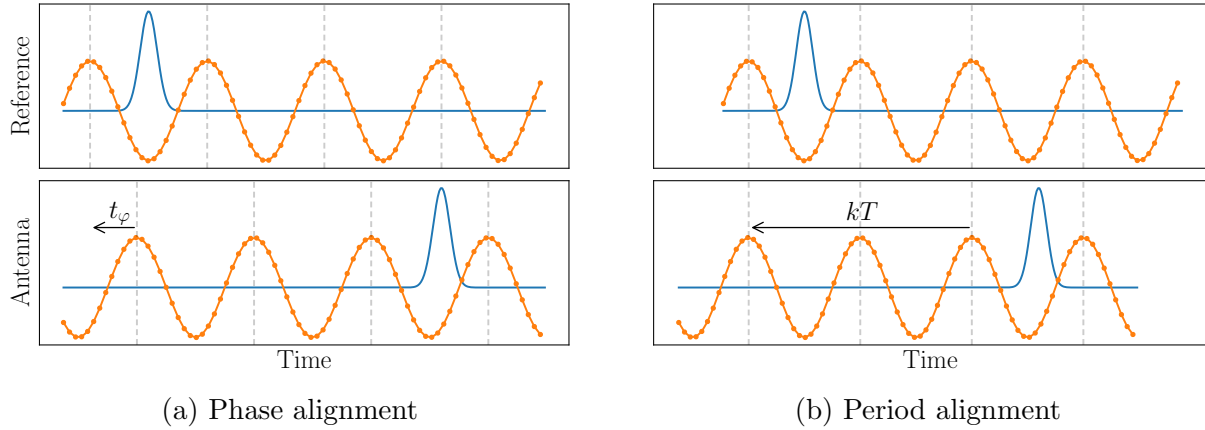


Figure 5.1: Synchronisation scheme for two antennas using a single sine wave beacon (orange) and an impulsive signal (blue). Vertical dashed lines indicate periods of the sine wave. (a) A small time delay t_φ is derived from the phase difference of the beacon as measured by the antennas. (b) The period mismatch k is determined from the overlap between the impulsive signals. Note that the impulsive signals do not coincide perfectly due to different propagation delays from the source to the antennas.

With a restricted set of allowed period shifts, we can alternate optimising the calibration signal's origin and optimising the set of period time delays of the array.

In the case of radio detection of air showers, the very signal of the air shower itself can be used as the calibration signal. This falls into the dynamic setup previously mentioned. The best period defects must thus be recovered from a single event.

When doing the interferometric analysis for a sine beacon synchronised array, waveforms can only be delayed by an integer amount of periods, thereby giving discrete solutions to maximising the interferometric signal.

5.1 Air Shower simulation

To test the idea of combining a single sine beacon with an air shower, we simulated a set of recordings of a single air shower that also contains a beacon signal.²

The air shower signal was simulated by ZHAireS[13] on a grid of 10x10 antennas with a spacing of 50 meters. Each antenna recorded a waveform of 500 samples with a samplerate of 1 GHz for each of the X,Y and Z polarisations. The air shower itself was generated by a 10^{16} eV proton coming in under an angle of 20° from zenith.

Figure 5.2 shows the maximum electric field measured at each of the antennas. The ring of antennas with maximum electric fields in the order of $25 \mu\text{V}/\text{m}$ at the center of the array is the Cherenkov-ring. The Cherenkov-ring forms due to the forward beaming of the radio emissions of the air shower. Outside this ring, the maximum electric field quickly falls with increasing distance to the array core. As expected for a vertical shower, the projection of the Cherenkov-ring on the ground is roughly circular.

A sine beacon ($f_{\text{beacon}} = 51.53 \text{ MHz}$) was introduced at a distance of approximately 75 km northwest of the array, primarily received in the X polarisation. The distance between the antenna and the transmitter results in a phase offset with which the beacon is received at each antenna.³ The beacon signal was recorded over a longer time (10240 samples), to be able to distinguish the beacon and air shower later in the analysis.

The final waveform of an antenna (see Figure 5.3) was then constructed by adding its beacon and air shower waveforms and band-passing with relevant frequencies (here 30 and 80 MHz are taken by default). Of course, a gaussian white noise component is introduced to the waveform as a simple noise model (see Figure 4.10 for a treatise on the timing accuracy of a sine beacon).

²https://gitlab.science.ru.nl/mthesis-edebone/m-thesis-introduction/-/tree/main/airshower_beacon_simulation or https://etdeboone.nl/masters-thesis/airshower_beacon_simulation

³The beacon's amplitude is also dependent on the distance. Although simulated, this effect has not been incorporated in the analysis as it is negligible for the considered grid and distance to the transmitter.

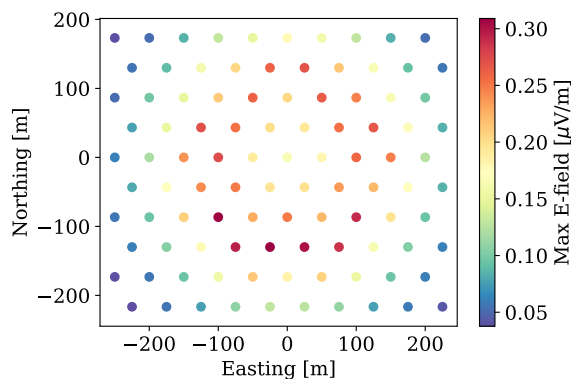


Figure 5.2: The 10x10 antenna grid used for recording the air shower. Colours indicate the maximum electric field recorded at the antenna. The Cherenkov-ring is clearly visible as a circle of radius 100 m centered at (0,0).

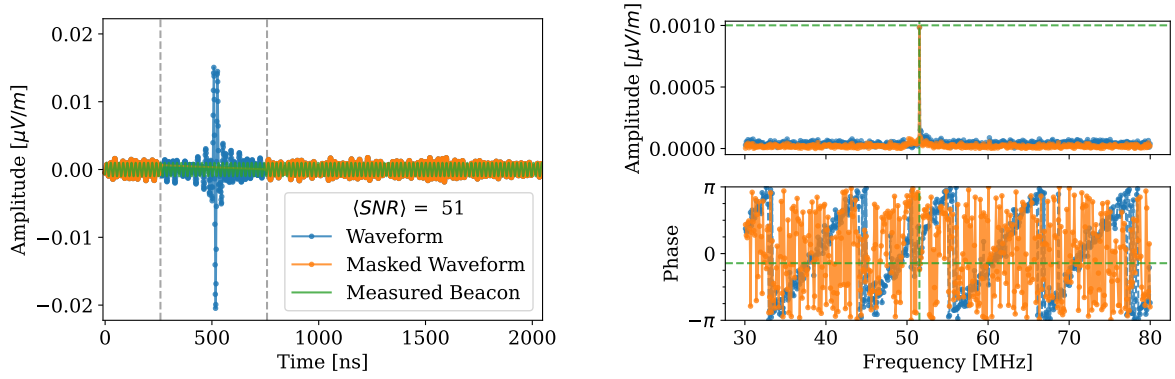


Figure 5.3: *Left:* Excerpt of a fully simulated waveform ($N = 10240$ samples) (blue) containing the air shower (a 10^{16} eV proton), the beacon (green, $f_{\text{beacon}} = 51.53$ MHz) and noise. The part of the waveform between the vertical dashed lines is considered air shower signal and masked before measuring the beacon parameters. *Right:* Fourier spectra of the waveforms. The green dashed lines indicate the measured beacon parameters. The amplitude spectrum clearly shows a strong component at roughly 50 MHz. The phase spectrum of the original waveform shows the typical behaviour for a short pulse.

After the creation of the antenna waveforms, the clocks are randomised by sampling a gaussian distribution with a standard deviation of 30 ns. At a beacon period of ~ 20 ns, this ensures that multiple antennas have clock defects of at least one beacon period. This in turn allows for synchronisation mismatches of more than one beacon period. Moreover, it falls in the order of magnitude of clock defects that were found in AERA[8].

To correctly recover the beacon from the waveform, it must be separated from the air shower. Typically, a trigger sets the location of the air shower signal in the waveform. In our case, the air shower signal is located at $t = 500$ ns (see Figure 5.3). Since the beacon can be recorded for much longer than the air shower signal, we mask a window of 500 samples around the maximum of the trace as the air shower's signal. The remaining waveform is fed into a DTFT (3.2) to measure the beacon's phase φ and amplitude. Note that due to explicitly including a time axis in a DTFT, a number of samples can be omitted without introducing artefacts.

With the obtained beacon parameters, the air shower signal is in turn reconstructed by subtracting the beacon from the full waveform in the time domain.

The small clock defect t_φ is then finally calculated from the beacon's phase φ by subtracting the phase introduced by the propagation from the beacon transmitter.

From the above, we now have a set of reconstructed air shower waveforms with corresponding clock defects smaller than one beacon period T . Shifting the waveforms to remove these small clock defects, we are left with resolving the correct number of periods k per waveform (see Figure 5.7b).

5.2 *k*-finding

Up until now, the shower axis and thus the origin of the air shower signal have not been resolved. This means that the unknown propagation time delays for the air shower (t_d) affect the alignment of the signals in Figure 5.1b in addition to the unknown clock period defects (kT). As such, both this origin and the clock defects have to be determined simultaneously.

If the antennas had been fully synchronised, radio interferometry as introduced in Chapter 2 can be applied to find the origin of the air shower signal, thus resolving the shower axis. Still, a (rough) first estimate of the shower axis might be made using this technique or by employing other detection techniques such as those using surface or fluorescence detectors.

On the true shower axis, the waveforms would sum most coherently when the correct k 's are used. Therefore, around the estimated shower axis, we define a grid search to both optimise this sum and the location of the maximum power. In this process each waveform of the array is allowed to shift by a restricted amount of periods with respect to a reference waveform (taken to be the waveform with the highest maximum).

Note that these grids are defined here in shower plane coordinates with \vec{v} the true shower axis and \vec{B} the local magnetic field. Searching a grid that is slightly misaligned with the true shower axis is expected to give comparable results.

The below *k*-finding algorithm is an iterative process where the grid around the shower axis is redefined on each iteration. Discussion is found in the next Chapter.

- Step 1. Define a grid around the estimated shower axis, zooming in on each iteration.
- Step 2. *k*-optimisation: per grid point, optimise the k 's to maximise the sum of the waveforms (see Figure 5.4).
- Step 3. *k*-finding: find the grid point with the maximum overall sum (see Figure 5.5a) and select its set of k 's.
- Step 4. Stop when the set of k 's is equal to the set of the previous iteration, otherwise continue.
- Step 5. Finally, make a power mapping with the obtained k 's to re-estimate the shower axis (location with maximum power) (see Figure 5.5b), and return to Step 1 for another iteration.

Here, Step 2 has been implemented by summing each waveform to the reference waveform (see above) with different time delays kT and selecting the k that maximises the amplitude of a waveform combination.⁴ As shown in Figure 5.4, the maximum possible period shift has been limited to ± 3 periods. This corresponds to the maximum expected time delay between two antennas with a clock randomisation up to 30 ns for the considered beacon frequency.⁵

⁴Note that one could use a correlation method instead of a maximum to select the best time delay. However, for simplicity and ease of computation, this has not been implemented. Other improvements are discussed in the next Section.

⁵Figure 5.6 shows this is not completely true. However, overall, it still applies.







Figure 5.4: Finding the maximum correlation for integer period shifts (best $k = -1$) between two waveforms recording the same (simulated) air shower. Randomising the antenna clocks up to 30 ns and $f_{\text{beacon}} = 51.53$ MHz corresponds to at most 3 periods of time difference between two waveforms.




(a) k -finding: optimise the k 's by shifting waveforms to find the maximum amplitude obtainable at each point.




(b) Power measurement with the k 's belonging to the overall maximum of the tested amplitudes.




(c) 2nd k -finding iteration: Zoom in on the location in (b) with the highest amplitude and repeat algorithm.



(d) Power measurement of the new grid.



(e) 3rd k -finding iteration: The same set of k 's has been found and we stop the algorithm.



(f) Final power measurement.

Figure 5.5: Iterative k -finding algorithm (see page 33 for explanation): First (a), find the set of period shifts k per point on a grid that returns the highest maximum amplitude (blue cross). The grid starts as a 8° wide shower plane slice at $X = 400 \text{ g/cm}^2$, centred at the true shower axis (red cross). Second (b), perform the interferometric reconstruction with this set of period shifts. Zooming on the maximum power (c),(d) repeat the steps until the k 's are equal between the zoomed grids (e),(f).

5.3 Strategy / Result

Figure 5.7 shows the effect of the various synchronisation stages on both the alignment of the air shower waveforms, and the interferometric power measurement near the true shower axis. Phase synchronising the antennas gives a small increase in observed power, while further aligning the periods after the optimisation process significantly enhances this power.

The initial grid plays an important role here in finding the correct axis. Due to selecting the highest maximum amplitude, and the process above zooming in aggressively, wrong candidate axes are selected when there is no grid-location sufficiently close to the true axis. Such locations are subject to differences in propagation delays that are in the order of a few beacon periods. The restriction of the possible delays is therefore important to limit the number of potential axis locations.

In this analysis, the initial grid is defined as a very wide 8° around the true axis. As the number of computations scales linearly with the number of grid points ($N = N_x N_y$), it is favourable to minimise the number of grid locations. Unfortunately, the above process has been observed to fall into local maxima when a too coarse and wide initial grid ($N_x < 13$ at $X = 400 \text{ g/cm}^2$) was used while restricting the time delays to $|k| \leq 3$.

As visible in the right side of Figure 5.7c, not all waveforms are in sync after the optimisation. In this case, the period defects have been resolved incorrectly for two waveforms (see Figure 5.6) due to too stringent limits on the allowable k 's. Looking at Figure 5.7b, this was to be predicted since there are two waveforms peaking at $k = 4$ from the reference waveform's peak (dashed line). As a result, the obtained power for the resolved clock defects is slightly less than the obtained power for the true clocks.

This does not impede resolving the shower axis. Figure 5.8 shows the power mapping at four different atmospheric depths for the resolved clock defects. Except for the low power case at $X = 800 \text{ g/cm}^2$, the shower axis is found to be $< 0.1^\circ$ of the true shower axis.

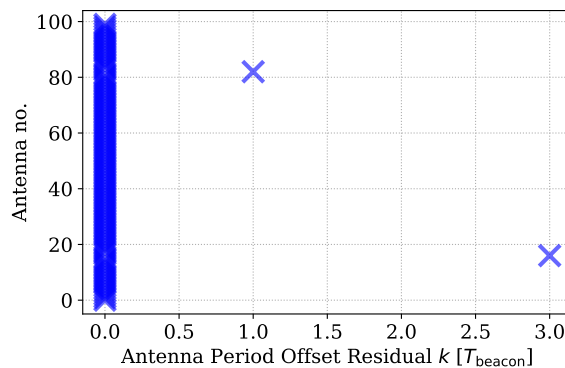


Figure 5.6: Errors in the resolved period defects with respect to the true period defects.

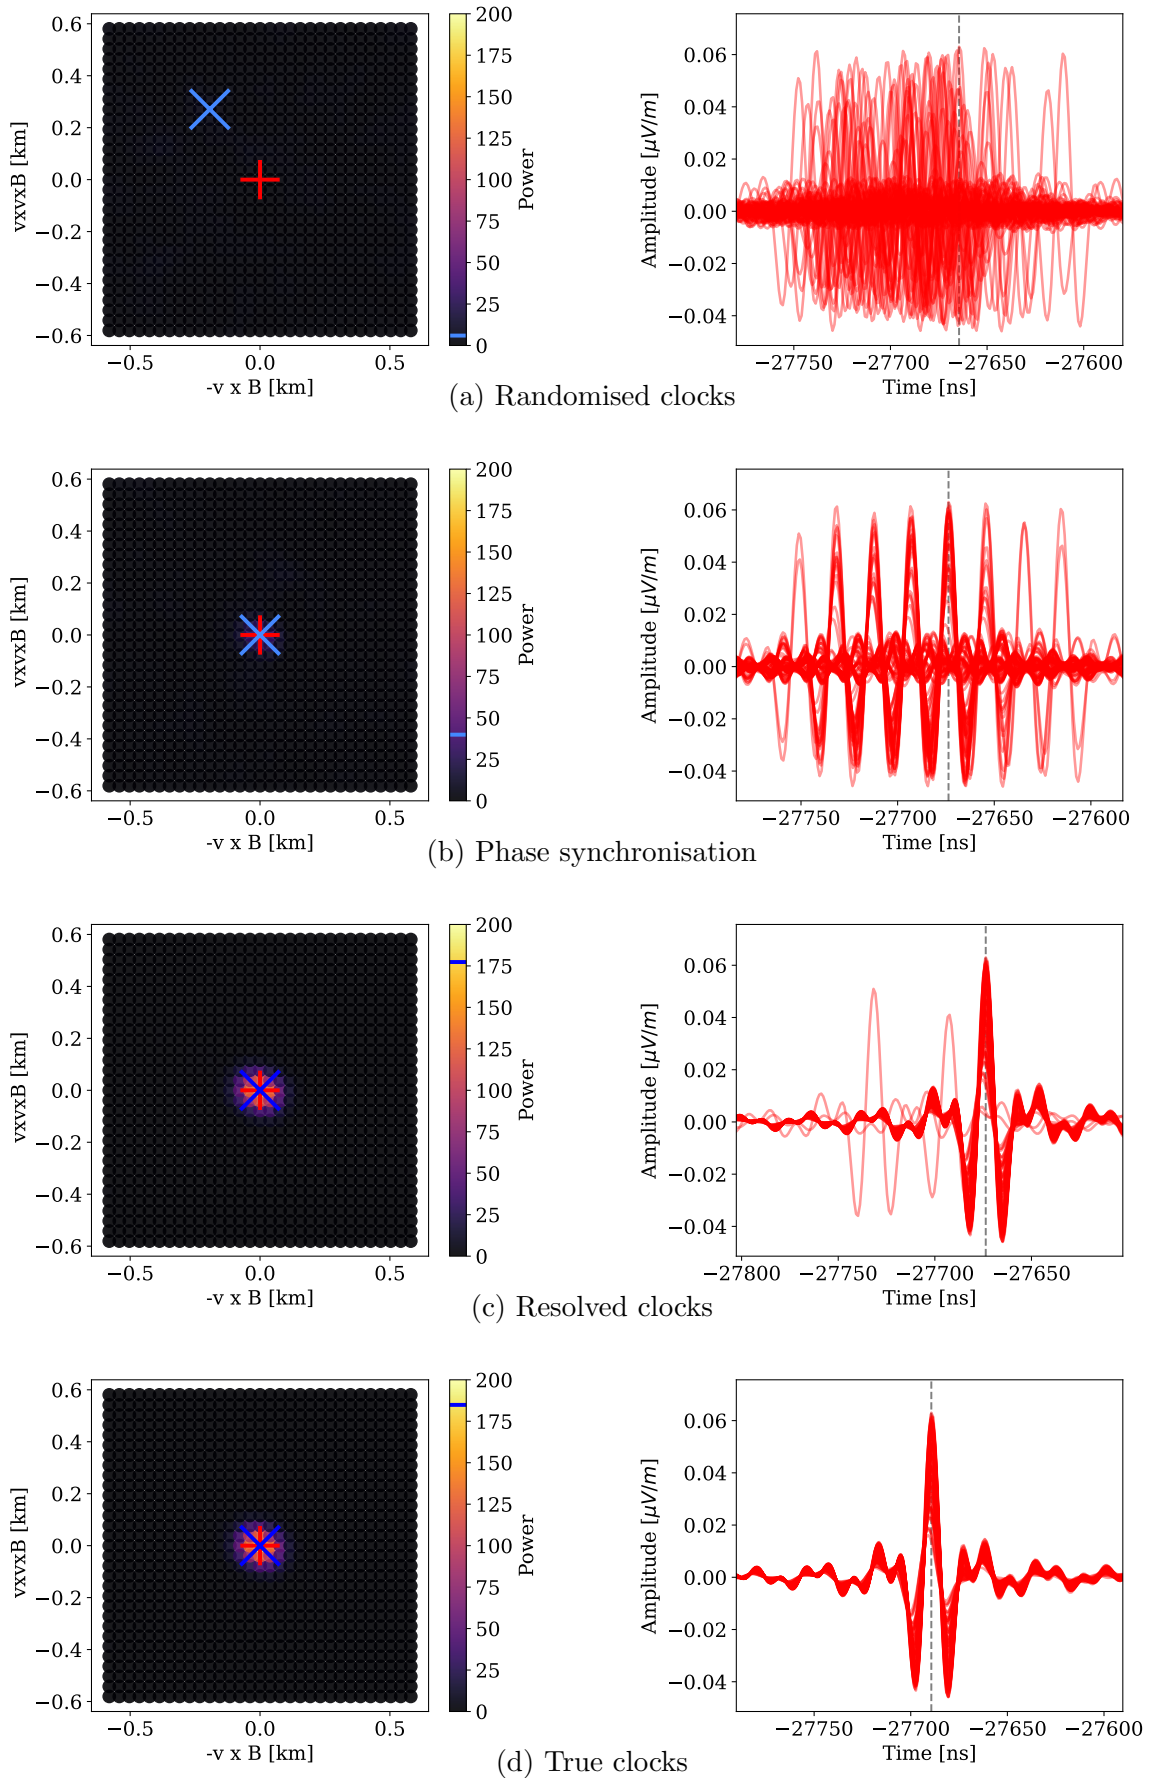


Figure 5.7: Different stages of array synchronisation (unsynchronised, beacon synchronised, k -resolved and true clocks) and their effect on (*right*) the alignment of the waveforms at the true axis and (*left*) the interferometric power near the simulation axis (red plus). The maximum power is indicated by the blue cross. In the right panes the vertical dashed line indicates the maximum of the reference waveform.

Of course, this algorithm must be evaluated at relevant atmospheric depths where the interferometric technique can resolve the air shower. In this case, after manual inspection, the air shower was found to have X_{\max} at roughly 400 g/cm^2 . The algorithm is expected to perform as long as a region of strong coherent power is resolved. This means that with the power in both Figure 5.8a and Figure 5.8c, the clock defects and air shower should be identified to the same degree.

Additionally, since the true period shifts are static per event, evaluating the k -finding algorithm at multiple atmospheric depths allows to compare the obtained sets thereof to further minimise any incorrectly resolved period defect.

Further improvements to the algorithm are foreseen in both the definition of the initial grid (Step 1) and the optimisation of the k 's (Step 2). For example, the k -optimisation step currently sums the full waveform for each k to find the maximum amplitude for each sum. Instead, the timestamp of the amplitude maxima of each waveform can be compared, directly allowing to compute k from the difference.

Finally, from the overlapping traces in Figure 5.7c, it is easily recognisable that some period defects have been determined incorrectly. Inspecting Figure 5.7b, this was to be expected as there are two waveforms with the peak at $|k| = 4$ from the reference waveform. Therefore, either the k -optimisation should have been run with a higher limit on the allowable k 's, or, preferably, these waveforms must be optimised after the algorithm is finished with a higher maximum k .

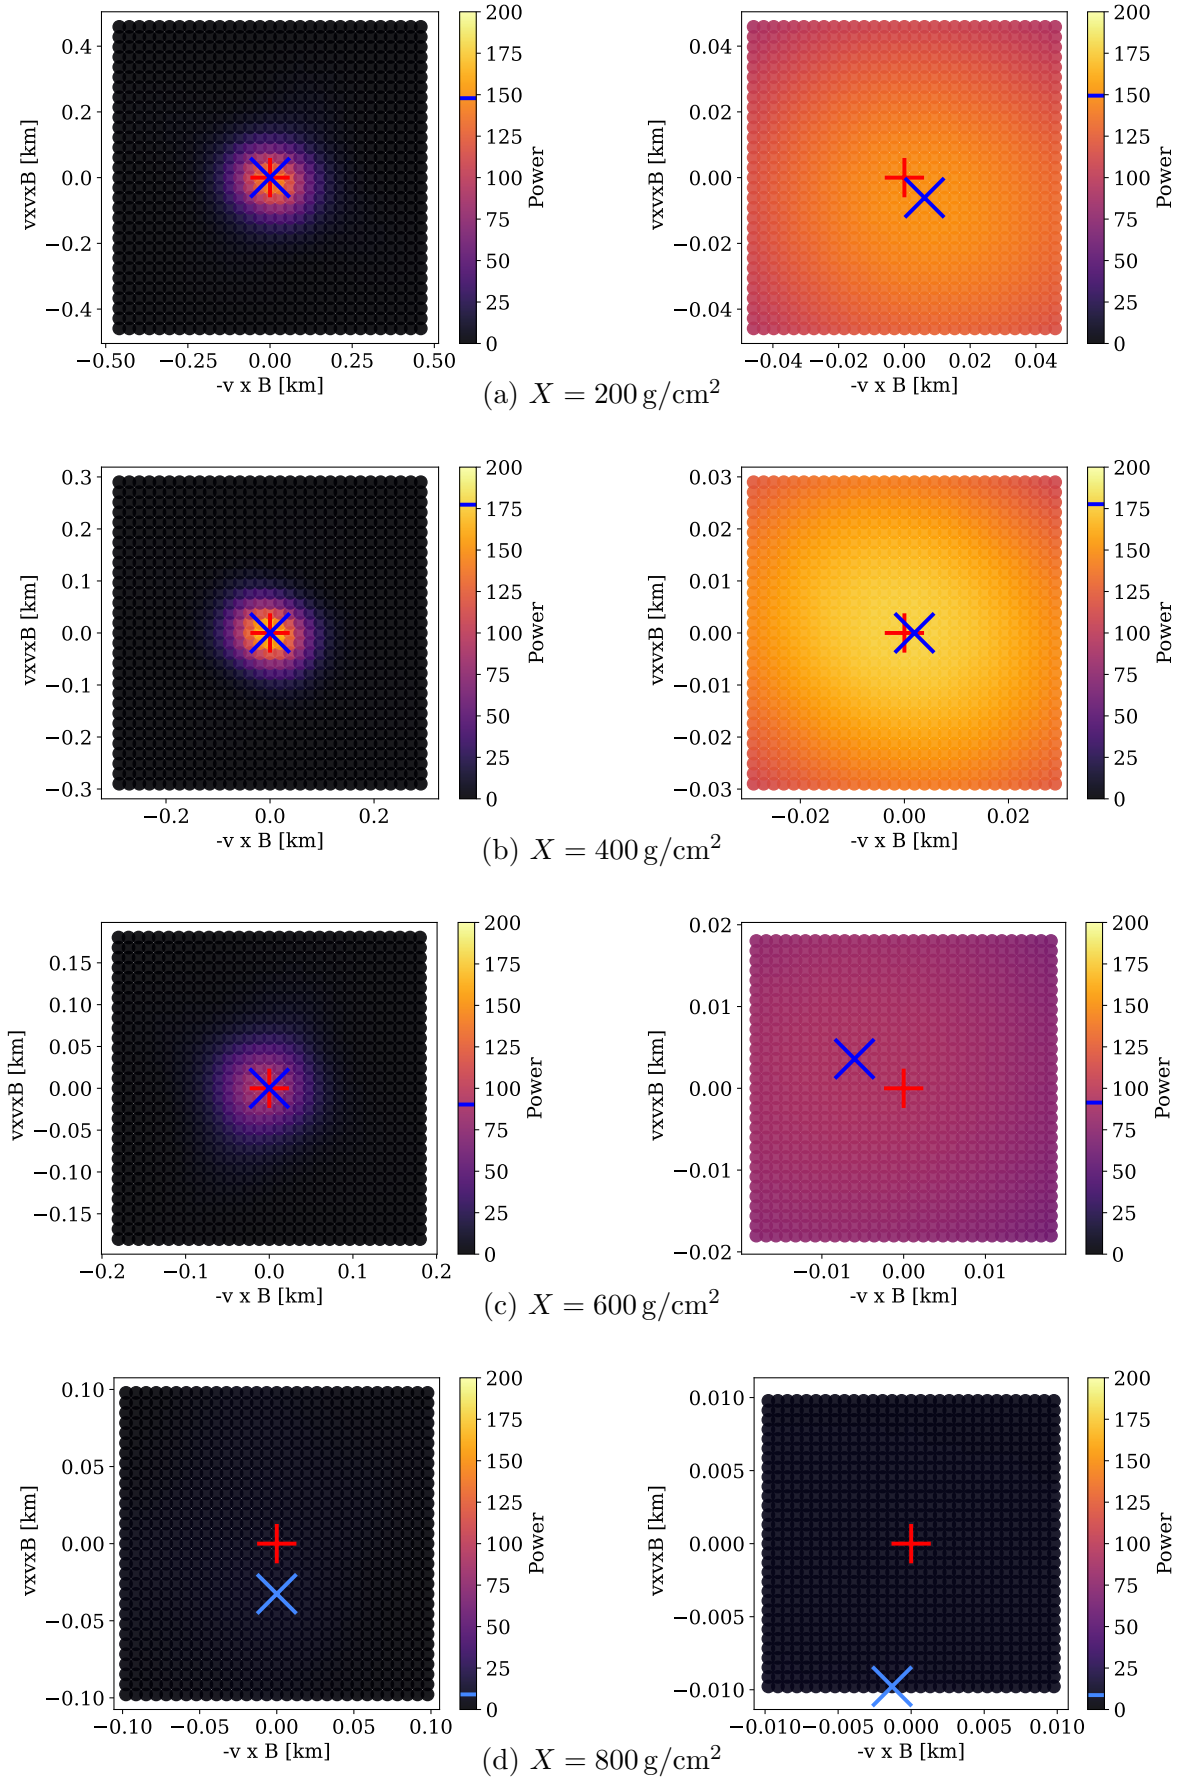


Figure 5.8: Interferometric power for the resolved clocks (from Figure 5.7c) at four atmospheric depths for an opening angle of 2° (left) and 0.2° (right). The simulation axis is indicated by the red plus, the maximum power is indicated by the blue cross. Except for (d) where there is no power, the shower axis is resolved within 0.1° of the true shower axis.

Chapter 6.

GRAND signal chain characterisation

The beacon synchronisation strategy hinges on the ability to measure the beacon signal with sufficient timing accuracy. In the previous chapters, the overall performance of this strategy has been explored by using simulated waveforms.

As mentioned in Chapter 3, the measured waveforms of a true detector will be influenced by characteristics of the antenna, the filter and the ADC. Especially the filter and ADC are important components to be characterised to compensate for possible systematic (relative) delays. This chapter starts an investigation into these systematic delays within GRAND’s Detector Unit (DU) V2.0[14].

At the base of every single antenna, a DU is mounted. Its protective encasing has three inputs to which the different polarisations of the antenna are connected. These inputs are connected to their respective filter chains, leaving a fourth filter chain as spare. Each filter chain band-passes the signal between 30 MHz and 200 MHz. Finally, the signals are digitised by a four channel 14-bit ADC sampling at 500 MHz. In our setup, the channels are read out after one of two internal “monitoring” triggers fire with the ten-second trigger (TD) linked to the 1 Pulse Per Second of the GNSS chip and the other (MD) a variable randomising trigger.

Both the ADC and the filter chains introduce systematic delays. Since each channel corresponds to a polarisation, it is important that relative systematic delays between the channels can be accounted for.




Figure 6.1: GRAND’s Detector Unit V2.0 inside its protective encasing.




Figure 6.2: Relative time delay experiment, a signal generator sends the same signal to two channels of the DU. The extra time delay incurred by the loop in the upper cable can be ignored by interchanging the cabling and doing a second measurement.

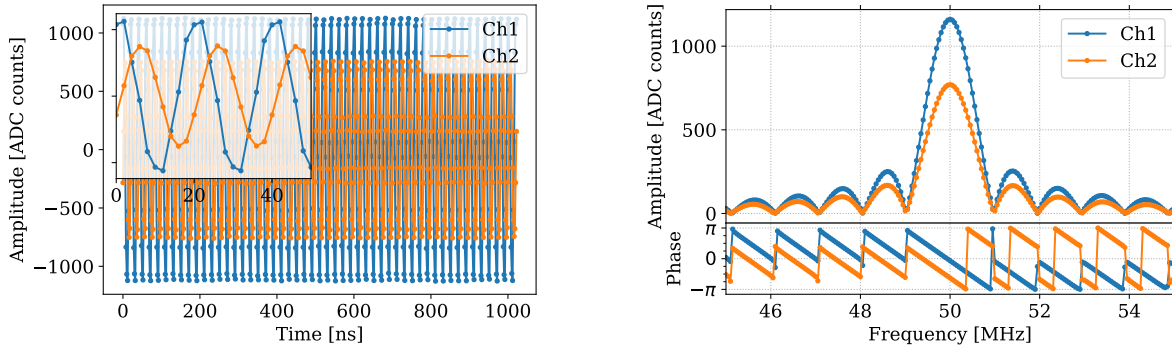


Figure 6.3: Waveforms of the sine wave measured in the “forward” setup and their spectra around the testing frequency of 50 MHz.. The sine wave was emitted at 200 mVpp.

Figure 6.2 illustrates a setup to measure the relative time delays of the filter chain and ADC. Two DU-channels receive the same signal from a signal generator where one of the channels takes an extra time delay Δt_{cable} due to extra cable length. In this “forward” setup, both channels are read out at the same time, and a time delay is derived from the channels’ traces. Afterwards, the cables are interchanged and a second (“backward”) time delay is measured.

The sum of the “forward” and “backward” time delays gives twice the relative time delay Δt without needing to measure the time delays due to the cable lengths t_{cable} separately since

$$\Delta t = (t_{\text{forward}} + t_{\text{backward}})/2 = ([\Delta t + t_{\text{cable}}] + [\Delta t - t_{\text{cable}}])/2. \quad (6.1)$$

We used a signal generator to emit a single sine wave at frequencies from 50 MHz to 200 MHz at 200 mVpp. Note that we measured the phases to determine the time delays for each channel. In Figure 6.3 the time delay between the channels is clearly visible in the measured waveforms as well as in the phase spectrum.

In our setup, the cable length difference was $3.17 - 2.01 = 1.06$ m, resulting in an estimated cable time delay of roughly 5 ns. At a frequency of 50 MHz, the difference between the forward and backward phase differences is thus expected to be approximately half a cycle. Figures 6.4 and 6.5 show this is in accordance with the measured delays.

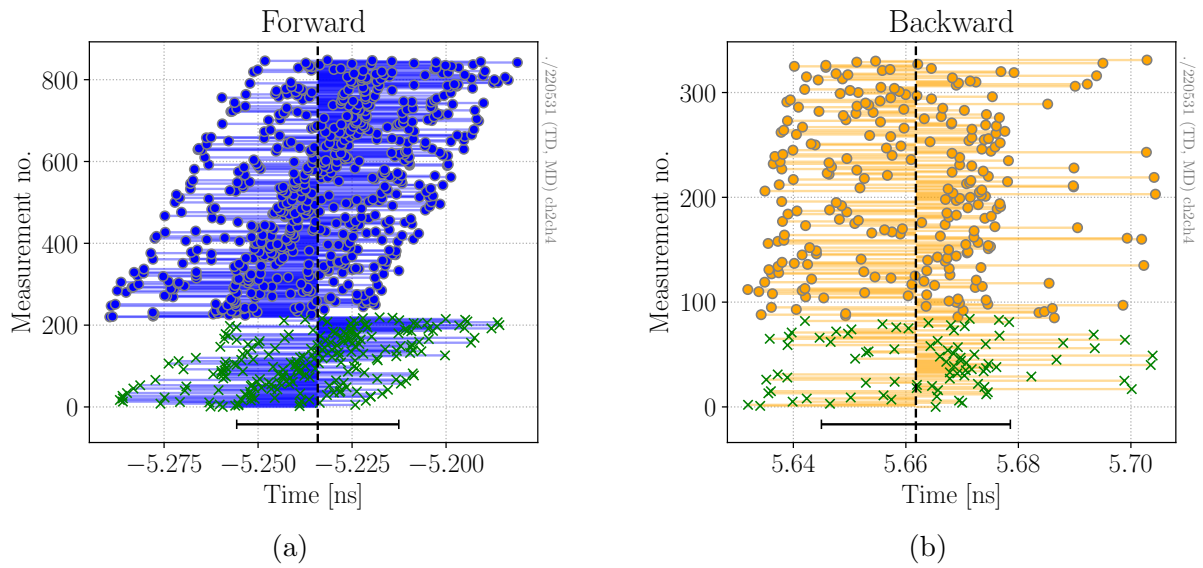


Figure 6.4: The measured phase differences between channels 2 and 4 at 50 MHz converted to a time delay for the (a) forward and (b) backward setups. The dashed vertical lines indicate the mean time delay, the errorbar at the bottom indicates the standard deviation of the samples. Crosses are TD-triggered events, circles are MD-triggered. The measurements are time-ordered within their trigger type.




Figure 6.5: Histogram of the measured phase differences in Figure 6.4. The relative signal chain time delay for the portrayed means is 0.2 ns.

Figure 6.6 shows the measured total time delays and the resulting signal chain time delays between both channels 1 and 2, and channels 2 and 4. Apart from two exceptional time delays up to 0.2 ns, the signal chain time delays are in general below 0.05 ns.

Note that the reported signal chain time delays must be taken to be indications due to systematic behaviours (see below).

Still, even when taking 0.2 ns as the upper limit of any relative signal chain time delay, the electric field at the antenna are reconstructable to a sufficient accuracy to use either the pulsed or sine beacon methods (see Figures 4.5 and 4.10 for reference) to synchronise an array to enable radio interferometry.

Note that at higher frequencies the phase differences are phase-wrapped due to contention of the used period and the cable time delay. Because it is symmetric for both setups, this should not affect the measurement of the signal chain time delay at the considered frequencies. Nevertheless, the result at these frequencies must be interpreted with some caution.

The time delays for both TD- and MD-triggered events in Figure 6.4 show a systematic behaviour of increasing total time delays for the forward setup. However, in the backward setup, this is not as noticeable.

This skewing of the channel time delays in one of the setups is also found at other frequencies (see Figures A.2 and A.3), raising questions on the stability of the setup. Unfortunately, it is primarily visible in the larger datasets which correspond to measurements over larger timescales. As the number of these large datasets is limited, further investigation with the current datasets is prohibited.

The skewing might also be an artefact of the short waveforms ($N \sim 500$ samples) the data acquisition system was able to retrieve at the time of measurement. Since the data acquisition system is now able to retrieve the maximum size waveforms, this systematic behaviour can be investigated in a further experiment.

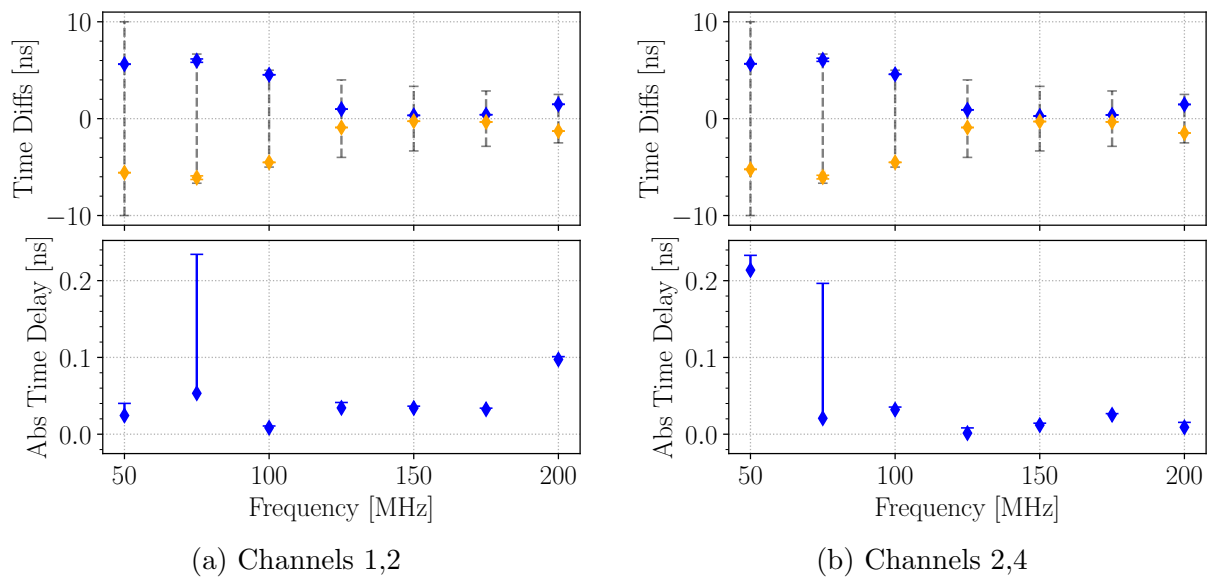


Figure 6.6: Total (*upper*) and signal chain (*lower*) time delays between (a) channels 1 and 2, and (b) 2 and 4. The dark grey vertical lines in the upper panes indicate the maximum measurable time delays at each frequency. Due to systematic effects in the measurements and a low number of samples at certain frequencies, the signal chain time delays depicted here must be taken as indicative. See text for discussion.

Chapter 7.

Conclusion

Using radio antennas to detect UHECRs has received much attention recently. The Pierre Auger Observatory is currently being upgraded to AugerPrime incorporating radio and scintillation detectors together with the already existing water-Cherenkov and fluorescence detectors. Other experiments, such as GRAND, envision to rely only on radio measurements of an EAS.

Time information in such large observatories is typically distributed using GNSSs, reaching up to 10 ns accuracy under very good conditions. For analysis using radio interferometry to be competitive, this timing accuracy needs to be improved towards the 1 ns mark.

A viable method to obtain this timing accuracy is to incorporate a beacon transmitter into the array. This (narrow-band) transmitter sends out a special signal that is picked up by the radio antennas in the array. With relatively simple techniques, the timing accuracy can be improved to below 1 ns. Thus, at a relatively low cost, the (relative) timing of radio arrays can be improved to enable radio interferometry.

In some circumstances, an external transmitter can be used as a beacon. For example, in Auger, a public TV broadcaster emits its signal at $f = 62.75$ MHz. With the source location and the frequency known, time delays can be calculated and this signal can be used to account for timing errors smaller than $T = 1/f \sim 16$ ns. Unfortunately, with the GNSS timing accuracy estimated in the same order of magnitude and the signal being periodic, the synchronisation of the antennas can be off by an integer amount of periods T .

Recording an air shower, in addition to such a narrow-band beacon, might provide a method to determine the correct beacon period. Radio interferometric analysis of the air shower depends on the coherence of the received signals. Any synchronicity problems in the radio antennas decrease the coherence and thus the power mapping used to derive properties of the air shower. With a limited set of periods to test, this power can be maximised while simultaneously inferring the correct beacon period.

The developed method to synchronise can be directly tested at Auger, both with data from AERA and the upcoming radio detectors from AugerPrime.

Appendix A.

Supplementary Figures



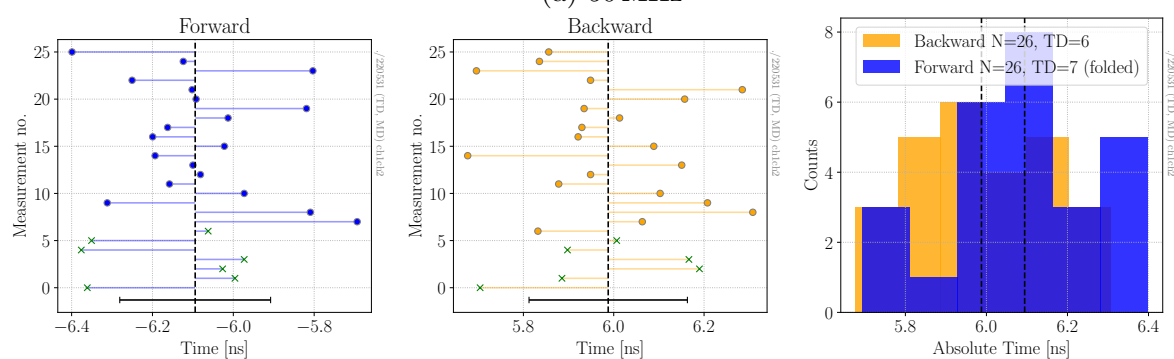



Figure A.1: Enhanced size of Figure 4.1.




(a) 50 MHz



(b) 75 MHz



(c) 100 MHz



(d) 125 MHz

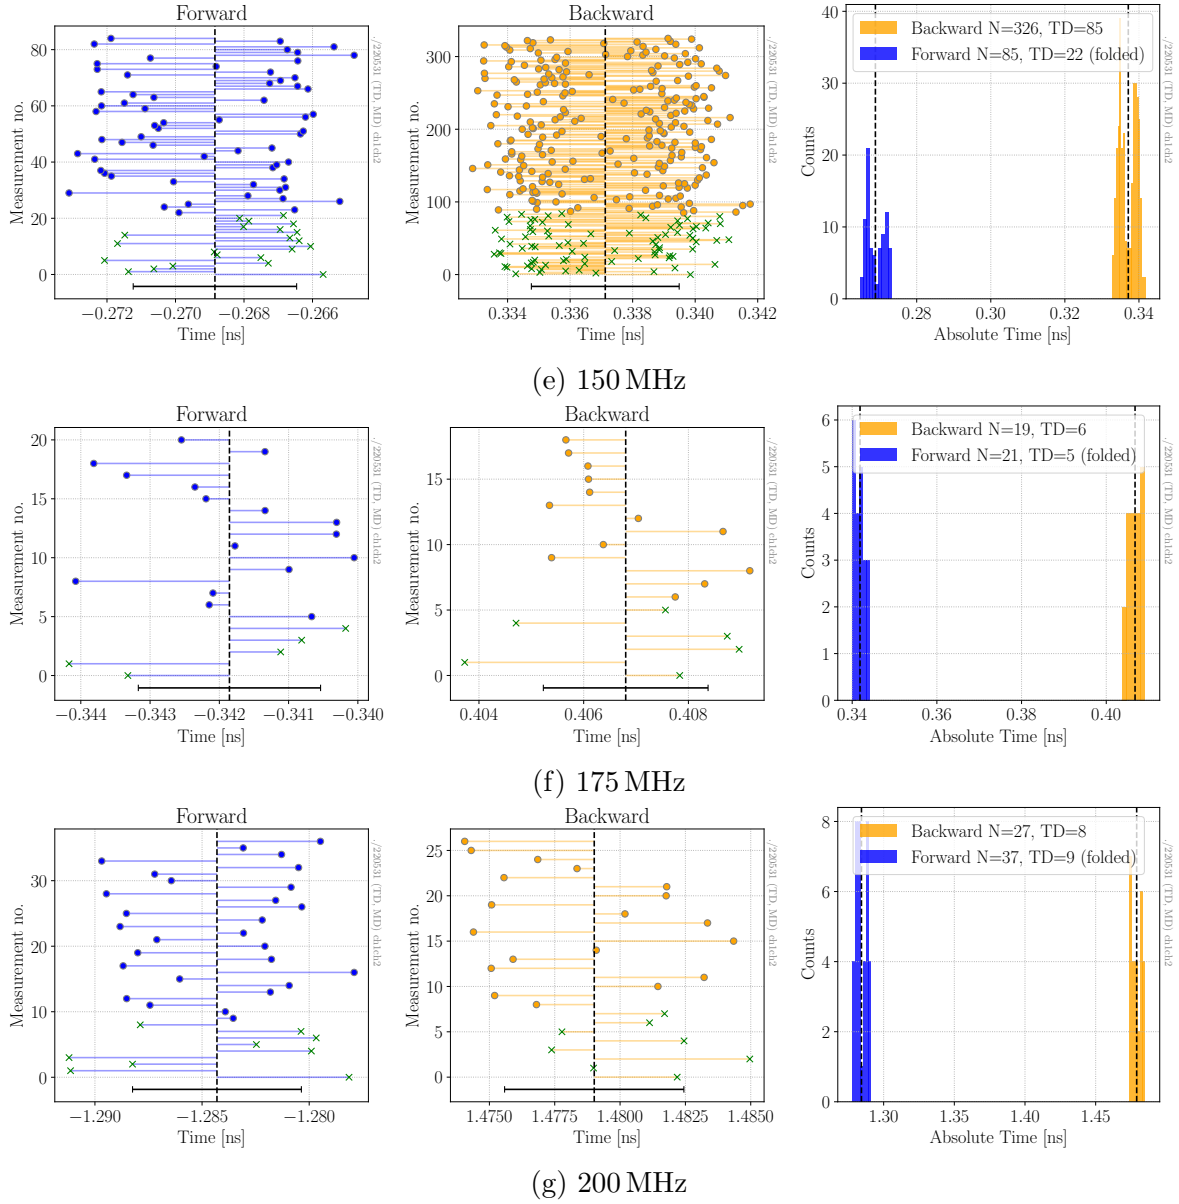

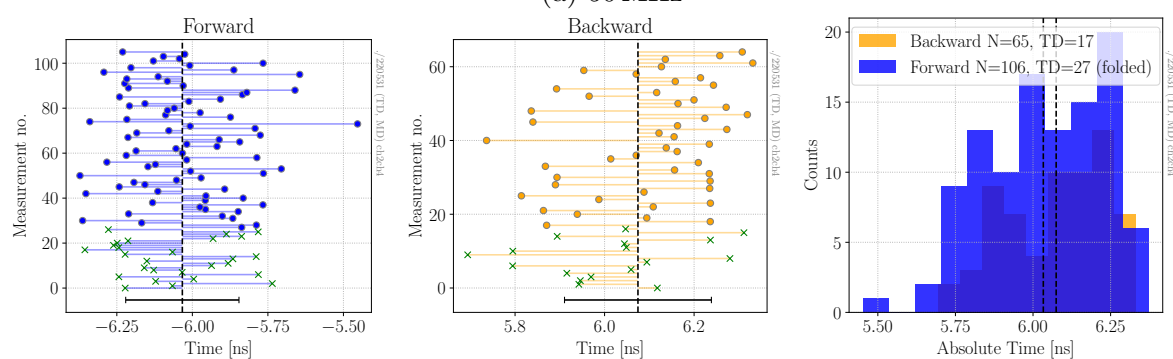



Figure A.2: The measured phase differences between channels 1 and 2. See Figure 6.4 for details.



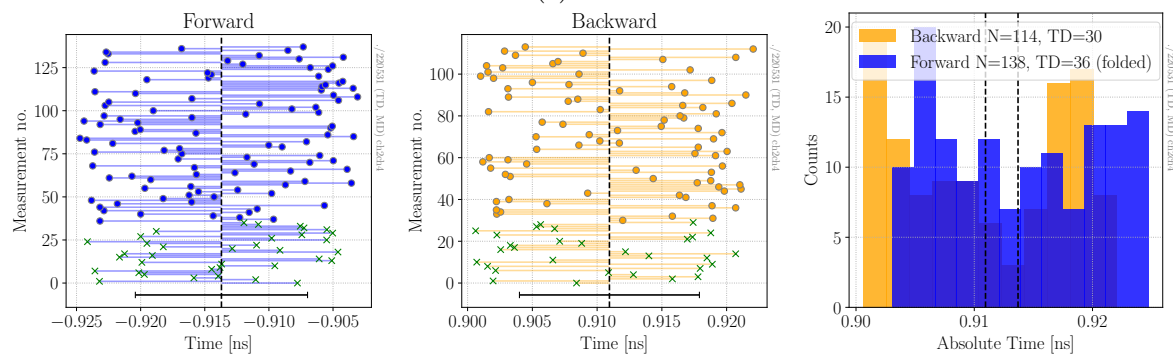
(a) 50 MHz



(b) 75 MHz



(c) 100 MHz



(d) 125 MHz

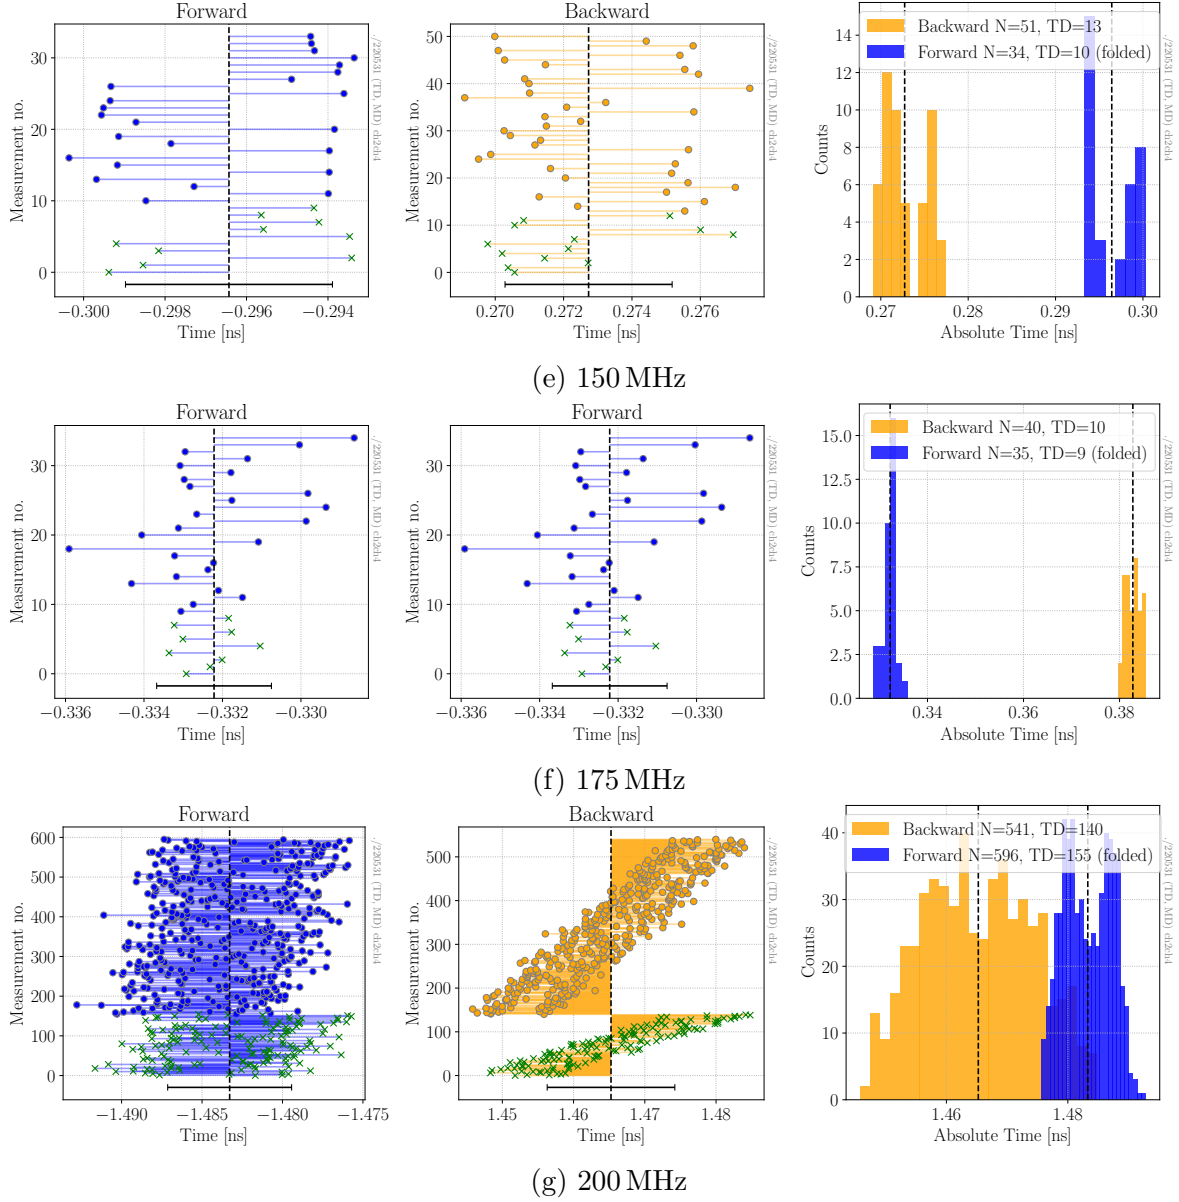


Figure A.3: The measured phase differences between channels 2 and 4. See Figure 6.4 for details.

Appendix B.

Random Phasor Sum Distribution

This section gives a short derivation of (4.6) using two frequency-domain phasors. Further reading can be found in Ref. [12, Chapter 2.9] under “Constant Phasor plus Random Phasor Sum”.

Write the noise phasor as $\vec{m} = a e^{i\phi}$ with phase $-\pi < \phi \leq \pi$ and amplitude $a \geq 0$, and the signal phasor as $\vec{s} = s e^{i\phi_s}$, but rotated such that its phase $\phi_s = 0$.

The noise phasor is fully described by the joint probability density function

$$p_{A\Phi}(a, \phi; \sigma) = \frac{a}{2\pi\sigma^2} e^{-\frac{a^2}{2\sigma^2}}, \quad (\text{B.1})$$

for $-\pi < \phi \leq \pi$ and $a \geq 0$.

Integrating (B.1) over the amplitude a , it follows that the phase is uniformly distributed. Likewise, the amplitude follows a Rayleigh distribution

$$p_A(a; \sigma) = \frac{a}{\sigma^2} e^{-\frac{a^2}{2\sigma^2}}, \quad (\text{B.2})$$

for which the mean is $\bar{a} = \sigma\sqrt{\frac{\pi}{2}}$ and the standard deviation is given by $\sigma_a = \sigma\sqrt{2 - \frac{\pi}{2}}$.

Adding the signal phasor, the mean in (B.1) shifts from $\vec{a}^2 = a^2(\cos\phi + \sin\phi)^2$ to $(\vec{a} - \vec{s})^2 = (a\cos\phi - s)^2 + (a\sin\phi)^2$, resulting in a new joint distribution

$$p_{A\Phi}(a, \phi; s, \sigma) = \frac{a}{2\pi\sigma^2} \exp\left[-\frac{(a\cos\phi - s)^2 + (a\sin\phi)^2}{2\sigma^2}\right]. \quad (\text{B.3})$$

Integrating (B.3) over ϕ one finds a Rice (or Rician) distribution for the amplitude,

$$p_A(a; s, \sigma) = \frac{a}{\sigma^2} \exp\left[-\frac{a^2 + s^2}{2\sigma^2}\right] I_0\left(\frac{as}{\sigma^2}\right), \quad (\text{B.4})$$

where $I_0(z)$ is the modified Bessel function of the first kind with order zero.

For the Rician distribution, two extreme cases can be highlighted (as can be seen in Figure B.1). In the case of a weak signal ($s \ll a$), (B.4) behaves as a Rayleigh distribution (B.2). Meanwhile, it approaches a gaussian distribution around s when a strong signal ($s \gg a$) is presented.

$$p_A(a; \sigma) = \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{(a - s)^2}{2\sigma^2}\right] \quad (\text{B.5})$$




Figure B.1: A signal phasor's amplitude in the presence of noise will follow a Rician distribution (B.4). For strong signals, this approximates a gaussian distribution, while for weak signals, this approaches a Rayleigh distribution.




Figure B.2: The Random Phasor Sum phase distribution (B.6). For strong signals, this approximates a gaussian distribution, while for weak signals, this approaches a uniform distribution.

Like the amplitude distribution (B.4), the marginal phase distribution of (B.3) results in two extreme cases; weak signals correspond to the uniform distribution for (B.1), while strong signals are well approximated by a gaussian distribution (see Figure B.2).

The analytic form takes the following complex expression,

$$p_{\Phi}(\phi; s, \sigma) = \frac{e^{-\left(\frac{s^2}{2\sigma^2}\right)}}{2\pi} + \sqrt{\frac{1}{2\pi}} \frac{s}{\sigma} e^{-\left(\frac{s^2}{2\sigma^2} \sin^2 \phi\right)} \frac{\left(1 + \operatorname{erf} \frac{s \cos \phi}{\sqrt{2}\sigma}\right)}{2} \cos \phi \quad (\text{B.6})$$

where

$$\operatorname{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z dt e^{-t^2}, \quad (\text{B.7})$$

is the error function.

Bibliography

- [1] Carmelo Evoli. *The Cosmic-Ray Energy Spectrum*. Dec. 2020. DOI: [10.5281/zenodo.4396125](https://doi.org/10.5281/zenodo.4396125). URL: <https://doi.org/10.5281/zenodo.4396125>.
- [2] Olivier Deligny. “The science of ultra-high energy cosmic rays after more than 15 years of operation of the Pierre Auger Observatory”. In: *Rencontres de Blois 2023*. Sept. 2023. arXiv: [2309.01259](https://arxiv.org/abs/2309.01259) [[astro-ph.HE](#)].
- [3] G. A. Askar’yan. “Excess negative charge of an electron-photon shower and its coherent radio emission”. In: *Zh. Eksp. Teor. Fiz.* 41 (1961), pp. 616–618.
- [4] Harm Schoorlemmer. “Tuning in on cosmic rays : polarization of radio signals from air showers as a probe of emission mechanisms”. PhD thesis. Nijmegen U., Nijmegen U., 2012.
- [5] Tim Huege. “Radio detection of extensive air showers”. In: *Nucl. Instrum. Meth. A* 876 (2017). Ed. by P. Krizan et al., pp. 9–12. DOI: [10.1016/j.nima.2016.12.012](https://doi.org/10.1016/j.nima.2016.12.012). arXiv: [1701.02995](https://arxiv.org/abs/1701.02995) [[astro-ph.IM](#)].
- [6] Jaime Álvarez-Muñiz et al. “The Giant Radio Array for Neutrino Detection (GRAND): Science and Design”. In: *Sci. China Phys. Mech. Astron.* 63.1 (2020), p. 219501. DOI: [10.1007/s11433-018-9385-7](https://doi.org/10.1007/s11433-018-9385-7). arXiv: [1810.09994](https://arxiv.org/abs/1810.09994) [[astro-ph.HE](#)].
- [7] Harm Schoorlemmer and Washington R. Carvalho. “Radio interferometry applied to the observation of cosmic-ray induced extensive air showers”. In: *Eur. Phys. J. C* 81.12 (2021), p. 1120. DOI: [10.1140/epjc/s10052-021-09925-9](https://doi.org/10.1140/epjc/s10052-021-09925-9). arXiv: [2006.10348](https://arxiv.org/abs/2006.10348) [[astro-ph.HE](#)].
- [8] A. Aab et al. “Nanosecond-level time synchronization of autonomous radio detector stations for extensive air showers”. In: *JINST* 11.01 (Jan. 2016), P01018. DOI: [10.1088/1748-0221/11/01/P01018](https://doi.org/10.1088/1748-0221/11/01/P01018). arXiv: [1512.02216](https://arxiv.org/abs/1512.02216) [[physics.ins-det](#)].
- [9] Tim Huege. “The Radio Detector of the Pierre Auger Observatory – status and expected performance”. In: *EPJ Web Conf.* 283 (2023), p. 06002. DOI: [10.1051/epjconf/202328306002](https://doi.org/10.1051/epjconf/202328306002). arXiv: [2305.10104](https://arxiv.org/abs/2305.10104) [[astro-ph.IM](#)].
- [10] J. Serrano et al. “The White Rabbit Project”. In: *ICALPECS*. Kobe, Japan, 2009.
- [11] Jane E. Gilligan et al. “White Rabbit Time and Frequency Transfer Over Wireless Millimeter-Wave Carriers”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 67.9 (Sept. 2020), pp. 1946–1952. ISSN: 1525-8955. DOI: [10.1109/TUFFC.2020.2989667](https://doi.org/10.1109/TUFFC.2020.2989667).
- [12] Joseph W. Goodman. *Statistical Optics*. en. 2nd ed. Wiley Series in Pure and Applied Optics. Nashville, TN: John Wiley & Sons, Apr. 2015. Chap. 2.9. ISBN: 978-1-119-00945-0.
- [13] Jaime Alvarez-Muniz et al. “Coherent Cherenkov radio pulses from hadronic showers up to EeV energies”. In: *Astropart. Phys.* 35 (2012), pp. 287–299. DOI: [10.1016/j.astropartphys.2011.10.002](https://doi.org/10.1016/j.astropartphys.2011.10.002). arXiv: [1005.0552](https://arxiv.org/abs/1005.0552) [[astro-ph.HE](#)].
- [14] Thei Wijnen. *The GRAND digitizer*. English. Version 1.0. Radboud University Nijmegen. Jan. 5, 2022.