UNIVERSITY OF GRONINGEN

BACHELOR THESIS

Detecting thunderstorm gamma radiation with LORA

Author:
Ewoud WEMPE

Supervisor:
Prof. Olaf SCHOLTEN

A thesis submitted in fulfillment of the requirements for the degree of Bachelor Physics

in the

KVI
Faculty of Science and Engineering

February 12, 2019

Abstract

Ewoud WEMPE

Detecting thunderstorm gamma radiation with LORA

Apart from the familiar lightning flashes, thunderstorms emit high-energy particles and gamma radiation. On ground level, the most important types of thunderstorm-related high-energy activity are Terrestrial Gamma-ray Flashes (TGFs) and gamma-ray glows. In this project, I looked for these phenomena in the data of the LORA cosmic ray detectors, that stand on the LOFAR core.

I estimated that it should be possible to observe TGFs in the data, if one occurred nearby. And although gamma-ray glows would not give a rise in the event rate, they should be measurable by looking at single-detector count rates.

It turned out, however, that the LORA data had many instrumental artifacts that made it unable to observe these thunderstorm-related events. In particular, the singles count rates showed many jagged features, and peaks that rose several orders of magnitude above background, even during fair weather. The rate of events that LORA measured also showed many spikes of instrumental nature.

One thunderstorm-correlated type of event was observed: many mysterious, non-causal, modulated, sinusoidal signals that varied over microsecond timescales showed up during thunderstorms. What caused them remains unclear.

Contents

Chapter 1

Introduction

Lightning has been observed for a long time and it is quite common, on average worldwide there are 6 discharges $\mathrm{km}^{-2}\,\mathrm{yr}^{-1}$(Dwyer and Uman, 2014). However, many questions remain unanswered when trying to understand the physics behind thunderstorms and lightning. In particular, the emission of high energy electrons and gamma rays (of MeV energies) is not well understood.

The physical mechanism for the production of high-energy particles in thunderclouds is currently explained by the Relativistic Runaway Electron Avalanche (RREA) model. When travelling through air (or other mediums), energetic electrons lose energy, mainly because of ionization losses (dominant at low energies) and radiative losses (i.e. Bremsstrahlung, dominant at high energies). The energy dependence of this energy loss is shown in Figure 1.1. One can regard this as a frictional force: eventually, high-energy electrons will slow down and be reabsorbed. However, if there is an electric field that causes a larger force (e.g. in thunderclouds), electrons can experience a runaway process. Then electrons are accelerated until the electric field is weaker than the frictional force again. Such electrons are called Wilson runaway electrons. When these Wilson runaway electrons scatter with other electrons, these can become new runaway electrons, and so forth. This creates an avalanche of electrons, and this process is called Relativistic Runaway Electron Avalanche (RREA). Such high-energy electrons will create gamma-rays and X-rays, by Bremsstrahlung, inverse compton scattering or synchrotron radiation, that can be observed.

There are also some secondary mechanisms to create new runaway electrons, for example by Compton scattering or absorption by X-ray or gamma-ray photons, and by positrons created in pair production by secondary gamma-rays. These feedback mechanisms are called relativistic feedback, and are most important when the electric fields are very high (Dwyer and Uman, 2014).

1.1 Gamma-ray glows

Gamma-ray glows are long-term excesses in gamma-ray intensity originating from thunderstorms. Because they are not nearly as luminous as Terrestrial Gamma-ray Flashes (TGFs), most measurements have been done close to the thundercloud. The first observations of gamma-ray glows were done by airplane flights carrying scintillators (e.g. by McCarthy and Parks (1985)). They observed glows that took some seconds, and the glows tended to end when lightning struck. The gamma-ray flux was an order of magnitude higher than the background. In the 90s, balloon observations by Eack, Suszcynsky, et al. (2000) showed that gamma-ray glows also occurred in 'thunderstorm anvils', as high as 14 km, and they showed a strong correlation between X-ray flux and electric field (Eack, Beasley, et al., 1996).

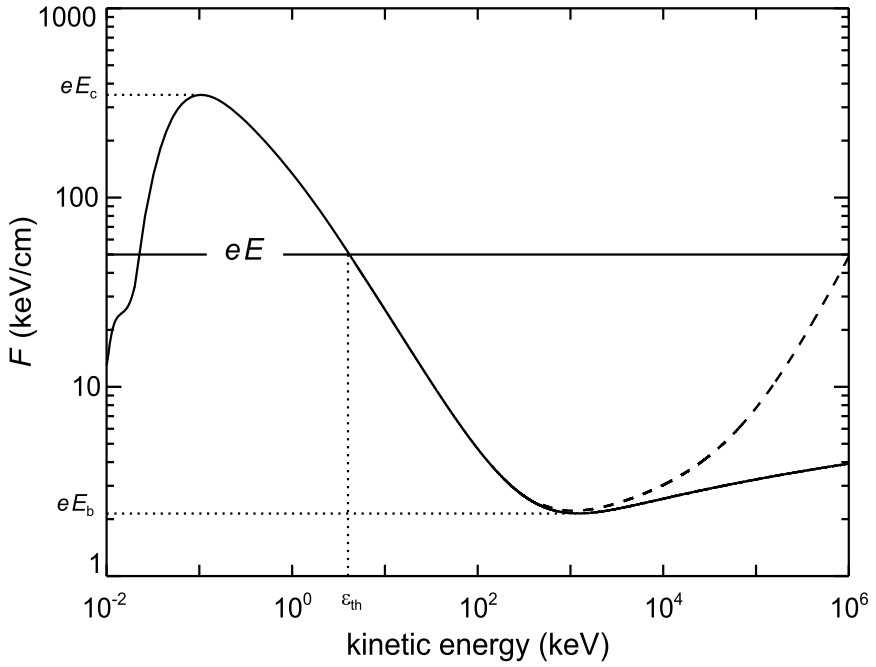

FIGURE 1.1: The energy loss per unit length that electrons with dependence on kinetic energy in air. An electron moving in a field with strength E, it will runaway if its energy is higher than a threshold energy ε_{th}. The critical field E_c is the field above which all electrons will runaway (also called a thermal runaway), and the breakeven field E_b is the minimum field in which the frictional force is countered. The dashed curve includes radiative losses (Brehmstrahlung). Figure from Dwyer and Uman (2014).

Ground based observations are harder, mainly because of the absorption by the air between thunderclouds and detectors, but also because of emission by radioactive emission of radon daughters, which is rainfall-correlated. To distinguish between emission of radon daughters and gamma ray glows, good spectra needed to be obtained, and to solve the problem of absorption, detections are preferably done close to the thundercloud (Dwyer, Smith, and Cummer, 2012).

The first convincing ground-based gamma-ray glow observations were done in Japan, where the charge centers of winter thunderstorms can be particularly low (Torii, Takeishi, and Hosono, 2002). Their observations showed a strong signal, and they obtained a spectrum that was consistent with simulated RREA spectra (although their energy range was not very large). This event also had the strongest signal of observations so far (70 times the background level). Some of the most elaborate ground-based gamma-ray glow observations were however done in the Armenian mountains, by Chilingarian et al. (2010). They detected over a 100 thunderstorm-correlated gamma-ray flux enhancements between 2003-2010, and they also meaured their spectra (much more precisely than the first Japanese observations). Additionally, they found evidence of a small increase in neutron flux. More recently, Chilingarian (2018) observed some gamma-ray glows lasting much longer than previously detected. They also found that gamma-ray glows had afterglows: 1 min to 10 min periods of high-energy radiation were followed by low-energy radiation that lasted for several hours.

The physical mechanism behind the emission of gamma-ray glows is reasonably well explained by RREA electrons (the spectra match well), but what the initial seed

particles are is not completely certain. Likely, secondary electrons/positrons of cosmic ray showers are the main source (Dwyer and Uman, 2014). The slow discharge that the gamma-ray glows are a sign of, can however have a large effect on the behaviour of thunderclouds. In Kelley (2014), it was estimated that a bright gamma-ray glow they observed could result in a discharge that is similar to the discharge provided by nearby lightning. Therefore, understanding gamma-ray glows is very important for the field of lightning research.

1.2 Terrestrial Gamma-ray Flashes

Compared to gamma-ray glows, terrestrial gamma-ray flashes (TGFs) happen on much short shorter timescales, with T_{50} (i.e. the time between 25 % and 75 % of the total number of counts) varying from $\sim$50 µs to $\sim$700 µs, with a median of 100 µs (Dwyer and Uman, 2014). The intensity is many orders of magnitude higher than the intensity of gamma-ray glows (Hare et al., 2016). They were first detected by orbiting satellites, because they were shorter and more intense than cosmic gamma ray flashes. Because the atmosphere quite quickly absorbs gamma rays, TGFs were originally thought to come from high-altitude discharges. But when in 2003 ground-level TGFs were observed, it started to become clear they originate from thunderstorms. Since then several other ground-level observations have been done, of which the measurement by Hare et al. (2016) has been the most extensive. TGFs correlate with lightning and are usually occur in upward positive intracloud lightning. The discharge that the runaway electrons from TGFs cause is significant, and may be as strong as lightning. However, much still remains unclear about the interpretation of the ground-based TGF observations.

1.3 This thesis

Firstly, in Chapter 2, the LORA detector setup and data structures will be shortly explained. Secondly, in Chapter 3 I will estimate what could be expected to be measured. In Chapter 4, the actual data is shown and discussed. Finally in Chapter 5, everything will be discussed more thoroughly, and reasons for disagreements between the expected and measured data. Additionally possible improvements to my work and to the LORA are suggested.

Chapter 2

Experimental setup

LORA (LOfar Radboud Air shower array) is an array of plastic scintillators designed for the detection of cosmic rays. It is divided into 5 stations, each containing 4 detector units. Each detector contains 4 plastic scintillators of $47.5 \times 47.5 \times 3\,\mathrm{cm}^3$, for a total area of about $1\,\mathrm{m}^2$ per detector. Photons emitted by the scintillators pass through a wavelength shifter to two photomultiplier tubes (PMTs), as shown in Figure 2.1. Each PMT's voltage is continuously calibrated to match the gain of the two PMTs in each detector module. The signals of the two PMTs are added up, and subsequently sent to the station electronics unit. Geometrically, the detectors are placed semi-randomly on the LOFAR core. Detectors close to each other are connected to the same stations, and a map is shown in Figure 2.2.

The electronic units were originally made for the HiSPARC experiment. Each station contains of 2 electronic units, each having a 'master' and a 'slave' detector, for a total of 4 detectors per station. The electronics first convert the PMT voltages to a digital ADC count x, according to $V/\mathrm{mV} = -0.57 \cdot x + 113$ (Fokkema, 2012)[1]. Then the 'master' detectors are continuously checked for an excess of about 5σ above background, and when this happens, all detectors in the station are checked for coincident events. The coincidence time window is 400 ns. If 3 or more detectors have a sufficiently strong signal, a station trigger is reached, and the traces (the ADC count time series) as well as some other data is saved, as described in Appendix A. This I will call a 'weak event'. If in a coincidence window of 500 ns some of the other stations also trigger, it is attributed to the same event. These events, of more than 8 detectors, I will call a 'strong event'. If there is a station trigger in 4 or more stations, LOFAR antenna data is also saved, and the events are processed further automatically to calculate some characteristics like cosmic ray energy and arrival direction.

The electronics have a buffer that is capable of storing up to 5 events at a time, and according to Fokkema (2012), this should eliminate dead time. And indeed, it turns out that in the data there exist some events that occur directly after a previous event ends.

[1] The offset may be different in LORA, but the relation that 1 ADC count $= 0.57$ mV is correct.

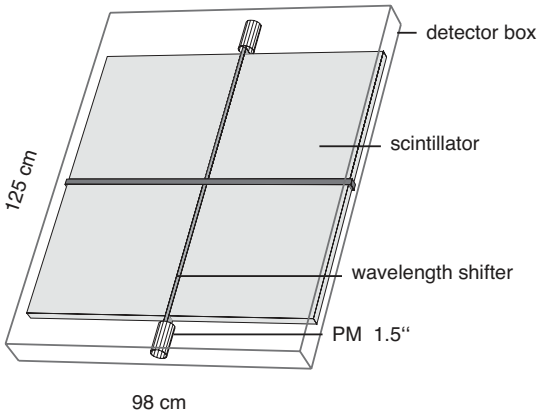

FIGURE 2.1: The layout of each detector unit. Figure from Thoudam et al. (2014).

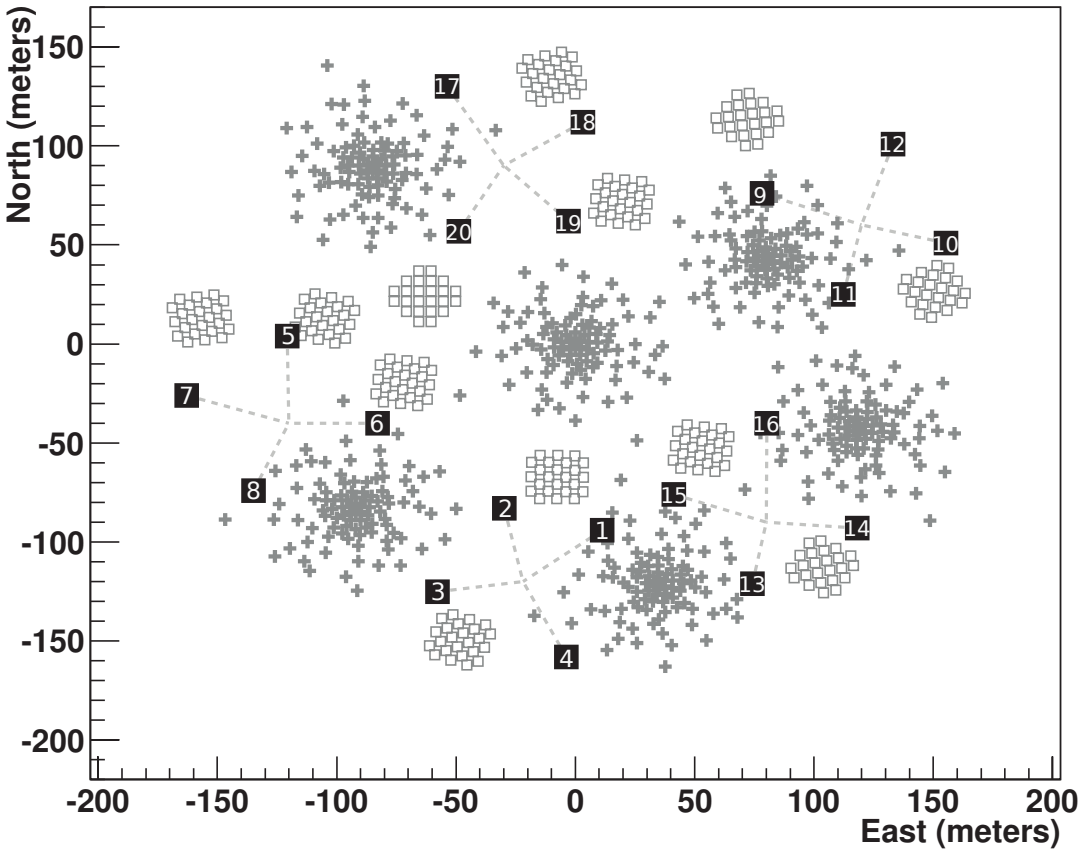

FIGURE 2.2: A map, showing the different detector positions. The black boxes indicate LORA detectors, and the numbers with in them the detector numbers. The dotted lines indicate the electrical connection between the stations and detectors. Figure adapted from Thoudam et al. (2014).

Chapter 3

Theory

3.1 Gamma-ray glows

One of the most elaborate ground-based observations of a gamma-ray glow has been done by Chilingarian et al. (2010). They measured a large, 13 min long, excess of high-energy electrons and gamma rays, whose spectra are shown in Figure 3.1. Note that this was the most intense event out of the more than 100 events that they measured in over the course of 5 years, so it is not representative of a normal gamma-ray glow. The spectra were described by the following fits:

$$\frac{I_\gamma}{\mathrm{m}^{-2}\,\mathrm{min}^{-1}\,\mathrm{MeV}^{-1}} = \begin{cases} 8.57 \times 10^6 \cdot E_{\mathrm{MeV}}^{-2.33} & \text{for} \quad E \leq 30\,\mathrm{MeV} \\ \mathrm{const} \cdot e^{-0.14 E_{\mathrm{MeV}}} & \text{for} \quad 30 \leq E_{\mathrm{MeV}} \leq 45 \end{cases} \tag{3.1}$$

$$\frac{I_{e^-}}{\mathrm{m}^{-2}\,\mathrm{min}^{-1}\,\mathrm{MeV}^{-1}} = 3.06 \times 10^5 e^{-0.18 E_{MeV}} \tag{3.2}$$

To measure these spectra, they used plastic scintillators similar to LORA's scintilla-

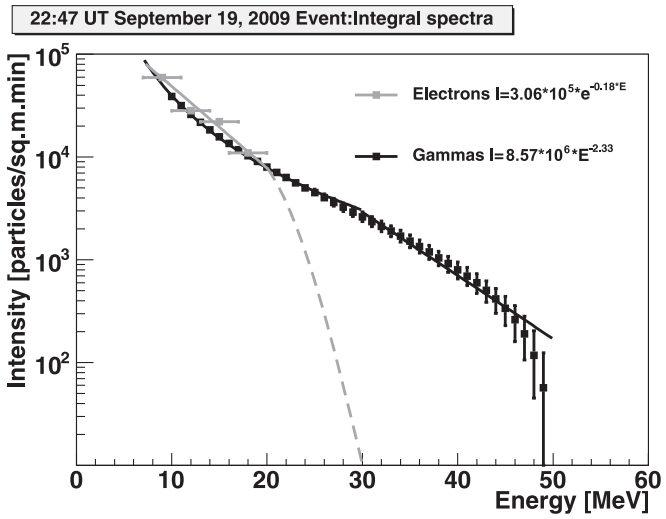

FIGURE 3.1: These gamma-ray glow spectra were measured in the mountains, close to the thunderstorm source. The bin size for gamma-rays was 1 MeV and for electrons the bin size is 3 MeV. The cut-off for the electrum spectrum reflects the fact that no electrons above 30 MeV were detected. Figure copied from Chilingarian et al. (2010).

tors, but the measurements were done in the mountains. That means they measured the fluxes much closer to the source, from a distacne of about 100 m to 200 m: they effectively measured inside of the cloud. For a detection by LORA, the spectrum

will be attenuated much more severely due to atmospheric losses. Typically, the clouds will be at least about 1 km above ground level (where the LORA detectors lie). To estimate the expected count rates, it is therefore necessary to calculate this attenuation.

Electrons

To calculate the evolution of the electron's spectrum through the atmosphere, I assumed that the electron energy decreases exponentially, as a function of the traversed column density x. This is a good approximation for the high energies (above 10 MeV), where Brehmstrahlung is dominant. I also assume that the number of energetic electrons stays conserved, but this is not entirely true: new electrons are created by photons in pair production, and electrons are lost in when combining with positrons. These effects have been neglected. Of course, electrons will eventually be captured by molecules, but this only occurs regularly at low energies. Since mainly particles above ∼5 MeV are detected by LORA, these effects have been ignored. Only including Brehmstrahlung, the energy follows:

$$\frac{1}{E_{e^-}}\frac{\mathrm{d}E_{e^-}}{\mathrm{d}X} = \frac{-1}{X_0} \approx \frac{-1}{37\,\mathrm{g\,cm^{-2}}} \implies E_{e^-}(X) = E_{e^-}(X=0)e^{-X/(37\,\mathrm{g\,cm^{-2}})} \tag{3.3}$$

In reality $X_0(E)$ has a strong energy dependence around these energies: at ∼10 MeV radiative losses equal ionization losses; below this ionization losses start to increase rapidly. So this calculation gives an upper limit on the electron flux. From this, one can calculate the spectrum of electrons when they reach the ground. Each electron will have its energy decrease according to Equation (3.3). That means that the spectrum per energy unit (particles $\mathrm{min^{-1}\,m^{-2}\,MeV^{-1}}$) will go up by this exponential factor, to conserve the total number of electrons. In reality, this conservation does not hold, but to get a upper limit it suffices. The result is that after some column density x the spectrum I_{e^-} will be[1]

$$I_{e^-}(E,x) = e^{x/X_0} I(e^{x/X_0}E, x=0). \tag{3.4}$$

To convert these column densities to physical heights, I assumed an isothermal exponential atmosphere, in which near sea level the atmospheric column density varies with height like (Grieder, 2010, p. 1029, p. 1076)

$$X/(\mathrm{g\,cm^{-2}}) = 1030e^{-h/(8.4\,\mathrm{km})}. \tag{3.5}$$

For various thunderstorm heights, I plotted the spectrum in the left panel of Figure 3.2. Already for a source about 0.5 km above the detectors at ground level, all of the electrons that reach the detectors are below 2 MeV[2]. Hence, the electrons will not be detectable at ground level by scintillators.

[1] Mathematically, from electron conversation, this result can be proven more rigorously by solving the transport equation: $0 = \frac{\partial I(E,x)}{\partial x} + \frac{\partial}{\partial E}\left(I(E,x)\frac{\partial E}{\partial x}\right)$.

[2] Actually, Chilingarian et al. (2010) did not find any electrons with energies higher than 30 MeV (this is reflected by the dashed line in Figure 3.1). This means that the actual incoming flux is even lower than what I estimated.

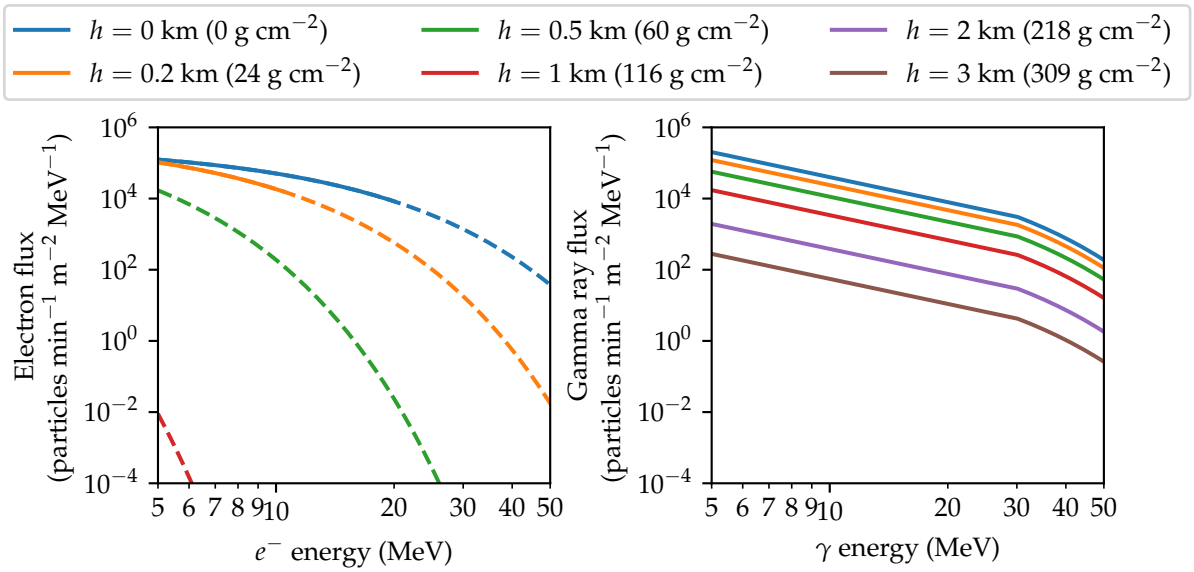

FIGURE 3.2: The expected electron and γ-ray spectrum at ground level for sources at various heights h. The indicated column densities are the ones between the ground and height h. The dashed lines indicate the part where the power-law spectrum was extrapolated. This is an overestimation, since Chilingarian et al. (2010) found that there actually was some cutoff at higher energies.

Gamma-rays

On the other hand, γ-rays do not lose energy in the same way electrons do. The attenuation of γ-rays arises mostly from two processes: pair production (at high energies) and Compton scattering (at lower energies). The sum of these two processes results in a attenuation that can be approximately described as:

$$\frac{N_\gamma}{N_{\gamma,0}} = e^{-\frac{7}{9}x/X_0(E)} \approx e^{-x/(47\,\mathrm{g\,cm^{-2}})} \tag{3.6}$$

This makes them much easier to detect than electrons, as the flux does not decrease as rapidly. The expected spectra are shown in Figure 3.2.

3.1.1 Coincidence rate increase

A rise in the event rate could perhaps be one way to detect gamma-ray glows, so to check this, the event rate is estimated here. For an uncorrelated flux of particles, the coincidence probability can then be calculated in the same way as in van Holten (2007). The expected number of 3-detector coincidences per unit time per station is

$$\langle c \rangle = N_{\text{permuts}} N_1 N_2 N_3 (\Delta\tau)^2 = 4N^3(\Delta\tau)^2, \tag{3.7}$$

where N is the expected number of single-detector events per unit time, $\Delta\tau$ is the coincidence time window, and N_{permuts} is the number of detector permutations leading to an event. Assuming that above some threshold energy E_{min} all events are counted,

one can calculate the total number of weak events:

$$\langle c_{\text{tot}} \rangle = 5 \cdot 4 \cdot \left(\int_{E_{\text{min}}}^{E_{\text{max}}} I_\gamma \, \mathrm{d}E \right)^3 (400\,\text{ns})^2 \tag{3.8}$$

Directly at the source, this would be $N = \int_{E_{\text{min}}}^{E_{\text{max}}} I_\gamma \, \mathrm{d}E = 7.09 \times 10^5\,\text{particles}\,\text{m}^{-2}\,\text{min}^{-1}$, giving an event rate of $\langle c_{\text{tot}} \rangle = 5.29\,\text{s}^{-1}$. For a source behind some column density x, this would go down by a factor of $\left(e^{x/(47\,\text{g}\,\text{cm}^{-2})}\right)^3$, since $N \propto e^{-x/(47\,\text{g}\,\text{cm}^{-2})}$. This rapid decrease results in rates of a few events per minute (the background level) at a height of already 500 m. Clearly the coincidence count rise from uncorrelated random events can not be detected.

But how many are generated in the cloud because of correlated events by this spectrum? For example, a gamma-ray could produce an electron-positron pair before colliding in the detector, and if these new particles hit separate detectors, this could more easily create an event. Presumably, such correlated events should not matter too much, because the detectors are far enough apart from each other and the angles between the velocities of the electron and positron would be minimal in the detector reference frame. However, in order to fully understand this behaviour, using a Monte Carlo air shower simulation code would be most practical, but that was not part of this thesis.

3.1.2 Single detector count rate increase

The LORA data also stores some threshold count rates for each station (see Table A.1). In practice, its background count varies quite chaotically (see Section 4.1.3), but here a rough estimate is made. For purely Gaussian noise, the probability of a 5σ-detection is $p((x-\mu) > 5\sigma) = 2.87 \times 10^{-7}$. The detectors take a measurement each $\Delta t = 2.5\,\text{ns}$, so the threshold level (which is about 5σ) would be reached with a frequency of $\langle c_t \rangle = \frac{p((x-\mu)>5\sigma)}{\Delta t} \approx 115\,\text{s}^{-1}$, for each detector. This background level would vary like a Poisson distribution, so the standard deviation is just $\sqrt{115}\text{s}^{-1} \approx 11\,\text{s}^{-1}$ The total number of incoming gamma-rays above 5 MeV is roughly $N = (1.18 \times 10^4\,\text{particles}\,\text{m}^{-2}\,\text{s}^{-1}) \cdot e^{-x/(47\,\text{g}\,\text{cm}^{-2})}$. This spectrum would thus be at the background noise level ($\sigma(c_t) \approx 11\,\text{s}^{-1}$) at a column depth of $x = (47\,\text{g}\,\text{cm}^{-2}) \ln\left(\frac{7.09 \times 10^5\,\text{particles}\,\text{m}^{-2}\,\text{min}^{-1}}{11\,\text{s}^{-1}}\right) = 329\,\text{g}\,\text{cm}^{-2}$, corresponding to a height of $\sim$3.3 km. That means that it should be possible to detect gamma-ray glows during thunderstorms, but one must keep in mind that this assumes the most intense glow that Chilingarian et al. (2010) observed in $\sim$7 years.

3.2 TGFs

TGF's (and X-ray flashes) are much shorter and stronger events. One of the most elaborate ground-level TGF observations so far has been done by Hare et al. (2016). This TGF happened 3.5 km high, geographically very close to the detectors. The event was so strong that magnetic field detectors in universities 250 km away measured it. The gamma radiation was strong enough to fully saturate an unshielded NaI scintillator, and for their plastic scintillator they obtained a digitizer volt signal ranging from about -1 V to 0.1 V. The trace of this TGF is shown in Figure 3.3.

Assuming that the efficiency of their plastic scintillators is reasonably similar to the LORA detectors, that means that you would expect that a similar source at

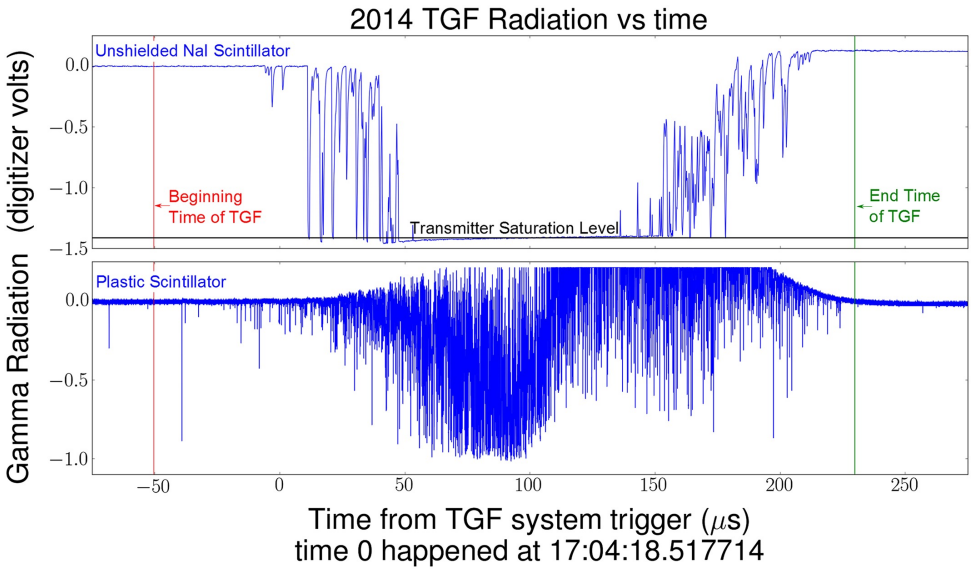

FIGURE 3.3: A trace of a NaI and plastic scintillator during the 2014 TGF. Figure copied from Hare et al. (2016).

similar distance would be clearly detectable in the LORA detectors. Although the sensitivity of the LORA detectors varies quite a bit, a typical signal has a noise level of about 4 ADC counts, or about $5 \cdot 4 \cdot 0.57\,\mathrm{mV} \sim 10\,\mathrm{mV}$. Assuming a linear response of digitizer volts with respect to gamma ray flux (this seems reasonable, it assumes that the PMT's are linear), one can estimate the maximum distance a TGF should be detectable. If one assumes a quadratic fall-off, a maximum distance would be of the order of $D_{\max} \approx 3.5\,\mathrm{km}\sqrt{\frac{1000\,\mathrm{mV}}{10\,\mathrm{mV}}} = 35\,\mathrm{km}$. However, it is more realistic to assume an exponential falloff with traversed column density x (i.e. $I_\gamma \propto e^{-x/(47\,\mathrm{g\,cm^{-2}})}$). Then, the maximum extra column density before the signal becomes undetectable is $x = (47\,\mathrm{g\,cm^{-2}}) \cdot \ln 100 = 216\,\mathrm{g\,cm^{-2}}$, approximately 1.5 km at ground level. This gives a maximum distance of about $D_{\max} = 5\,\mathrm{km}$.

There has definitely been some thunderstorm activity in a radius of 5 km around LORA during all the years it took data. So, if TGFs are not too rare and occur in a considerable fraction of thunderstorms, there should be some TGFs in the data. Of course, this only serves as an order-of-magnitude estimate: the LORA detectors might have less or more sensitive PMTs, or different electronics, causing that the digitizer volts units might be different.

Because the LORA detectors are built mostly for cosmic ray detection, the traces they save are only 10 µs long (4000 points with $\Delta t = 2.5\,\mathrm{ns}$). The 2014 TGF lasted for a total of ∼200 µs, with the plastic scintillator signal varying significantly in a time window of ∼100 µs. So if there was a TGF that LORA detected, LORA would not have saved the full TGF trace, but some separate 10 µs traces. These events would happen almost directly after each other. TGFs would leave very distinct signatures in the data: traces in which the signal changes over the whole timespan, and multiple events shortly after each other. The reality did not turn out to be so favourable, as is seen in the next chapter (Chapter 4)...

Chapter 4

Results

Since the data taking spanned over 6 years, I picked some specific days to look at. First, to check the goodness of LORA's data, I have looked at some days with both fair weather conditions, and some lightning activity. Then, I looked at LORA's behaviour during some intense thunderstorms that occurred above the LOFAR core.

4.1 LORA operation in normal conditions

I first looked at at data from 2016-07-10 to 2016-07-13. On 2016-07-12, there was some lightning activity in the area, and the LOFAR radio telescopes were used to image the radio emission of a lightning flash on 2016-07-12 at 17:34:55.100 (all times in this thesis are in UTC).

Radar reflectivities are converted to measure the precipitation flux, and although lightning does not necessarily have to coincide with rain, high values are a sign of lightning. In Figure 4.1 these measures are shown, to get an idea of the thunderstorm activity in the area of the LOFAR core. July 10 and 11 had fair-weather conditions, on July 12, there was some lightning activity, and on the 13th there was fair weather again. I have verified that there was no maintenance going on to the LORA detectors on any of these dates.

4.1.1 Event count rate

Although in Section 3.1.1 I estimated that a gamma-ray glow related increase in event count rate should be much below the background noise, it is still interesting to look at. In fact, looking at the weak event count rates revealed an issue with the LORA detectors. The weak event count rate includes all events where the photomultiplier tube (PMT) voltage reaches an excess of about 5σ above background level, for at least 3 detectors in coincidence within a window of 400 ns. To inspect the long term behaviour, the number of events in each time bin was counted, with a bin size of a few minutes. These event rates are plotted in Figure 4.2. Several features are visible, but the big spikes (for example on 2018-07-12 at 9:00) are most striking. They happen more or less randomly, mainly in station 4 (L4), and less commonly in station 2 (L2). There does not seem to be a temporal pattern to it, there are spikes at all times, day and night. On longer timescales (weeks, months, years), there is no pattern either. I verified that they are not correlated with lightning. In some cases, the spikes happen coincidentally, with both station 2 and 4 having one at the same time, but this does not occur consistently. Zooming in on specific spikes, it turns out they mostly last about half an hour, rising smoothly and falling off more steeply. To inspect if during these event rate spikes the total energy deposited in the LORA during events also rises, I counted the total number of ADC counts in events in time bins. This gives a measure for the total energy flux that each detector receives, and

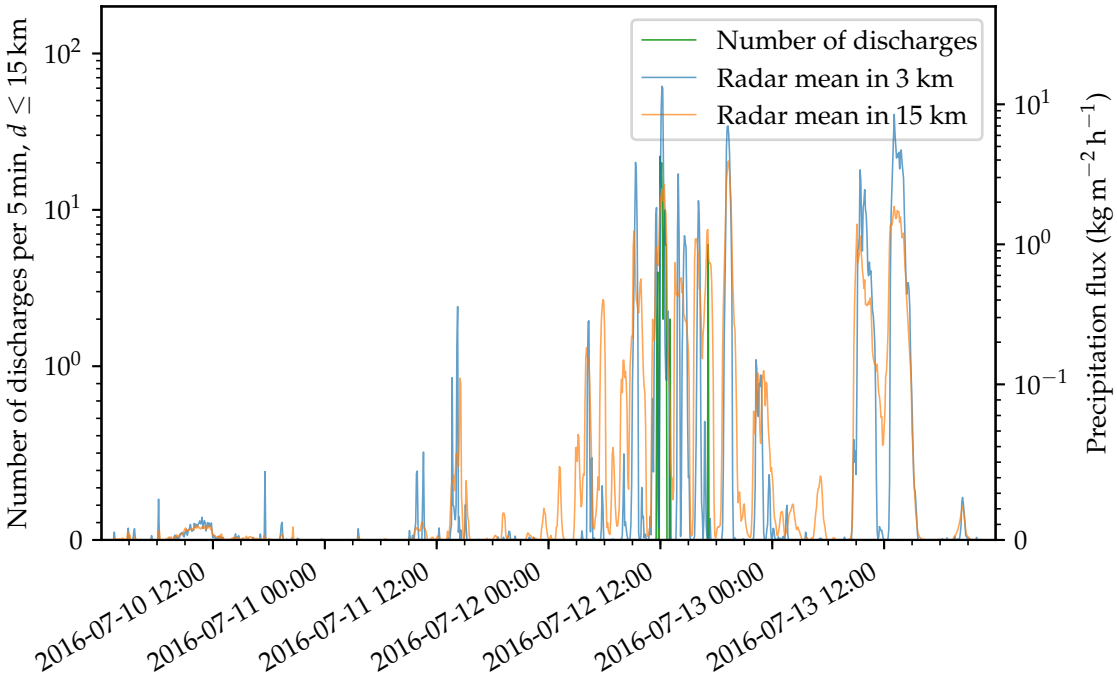

FIGURE 4.1: Thunderstorm and lightning activity data during these days. Data taken from KNMI and Buienradar.

it is plotted in Figure 4.3. In this figure, the signal seen in Figure 4.2 disappears: apparently, the spikes consist of very weak events.

4.1.2 Traces

For each event, LORA saves the photomultiplier signal in a time window of 10 µs around the event trigger. This signal is quantized by an Analog-to-Digital Converter (ADC), to some number of ADC counts. Although an ADC count is quite an arbitrary unit of itself, thankfully the LORA detectors are calibrated several times per year. This is done in such a way that the trigger energy of a the detector is the energy deposited by a 4 GeV muon (averaged over all angles), which is 6.7 MeV, as determined by GEANT simulations. Such an event will on average give 400 ADC counts when integrated over time, so this gives the relation that 400 ADC counts = 6.7 MeV deposited.

To further inspect the event rate spikes, I looked at the traces of some of the individual events, both during the spike and at normal times. An example of a trace of a normal cosmic ray event is shown in Figure 4.4. The trace starts recording 2 µs before the main peak, and keeps on recording for 8 µs after the main peak. Zooming into the main event at 2 µs shows that the main event takes of the order of ∼50 ns (although for strong events, this can be longer).

However, during the event rate spikes, the events show a very different behaviour. An example of such a trace is shown in Figure 4.5. Notice that only one detector clearly reaches the 5σ level, while for the normal events at least 3 of the detectors have a signal above the 5σ level. A possible cause for the detector to generate events anyway is that the traces show small peaks showing up in only a few (1-4) time bins for a pair of detectors. Somehow, this signal is sufficient to trigger a weak event, and the data is saved, even though there is not much there. The integrated number of ADC counts is also much less than those of normal cosmic ray events.

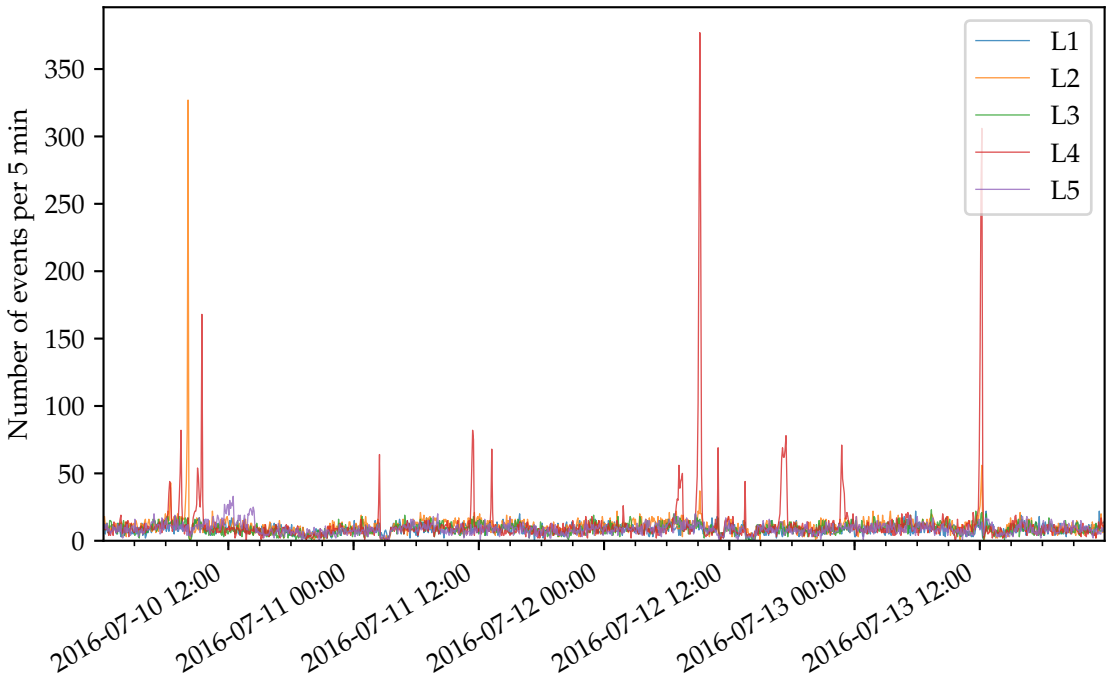

FIGURE 4.2: The weak event count rate. L1 refers to Lasa1, i.e. station 1, so detectors 1-4, etc. Notice the sharp peaks in both stations 4 and 2.

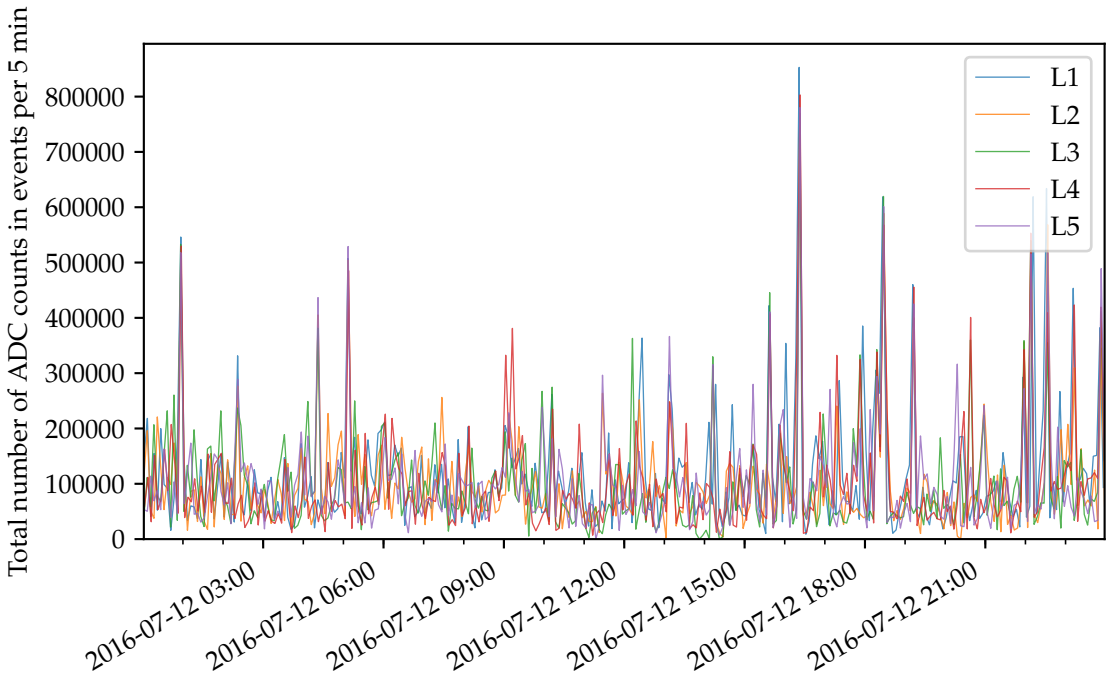

FIGURE 4.3: The total number of ADC counts that are registered in events. It is the same as Figure 4.2, but then each event is weighted by the integrated number of counts in the coincidence window, to give a measure for the energy flux that the detectors receive during events. The spikes that were previously visible disappear.

Usually, the integrated signal is 400 ADC counts (that is how the detectors are calibrated), but the bad pairs of detectors generally stayed below 100 ADC counts. The traces show clock-like ticks, that happen (3160 ± 10) ns apart, both for the Det15-Det16 pair and (more weakly) for the Det5-Det6 pair. The fact that it happens in detector pairs can mean that the PMT power source is faulty, or that something is wrong with the electronics, since the same pairs also share the same electronics (in particular, they share an ADC).

Because of the disagreement between stations, the random timing of the spikes, the fact that the traces were not reaching 5σ and the clock-like tick signal in the responsible pairs of detectors, the cause for these spikes is considered to be instrumental.

Separate from these ticks, a 100 MHz signal with an amplitude of a few ADC counts is also always present in the traces for all detectors. Because this signal is clearly just an instrumental effect, I removed it. Removing the signal was done by Fourier transforming the trace, setting the power of the 100 MHz bin to 0, and transforming back. This significantly reduced noise, by a factor of about 2, making the traces much nicer to inspect. One possible source is contamination from the LOFAR instrumentation: for their high-frequency antennas, they use a 100 MHz local oscillator. This 100 MHz signal might have some effect on the other data, because the station computer firmware does not seem to account for this. It could affect the threshold levels, because these are set by the noise, which in turn affects much more.

Another noteworthy feature is that detector 17 behaves differently from the rest (see Figure 4.4): there are small spikes present on top of the noise. This causes it to have a higher noise than the other detectors.

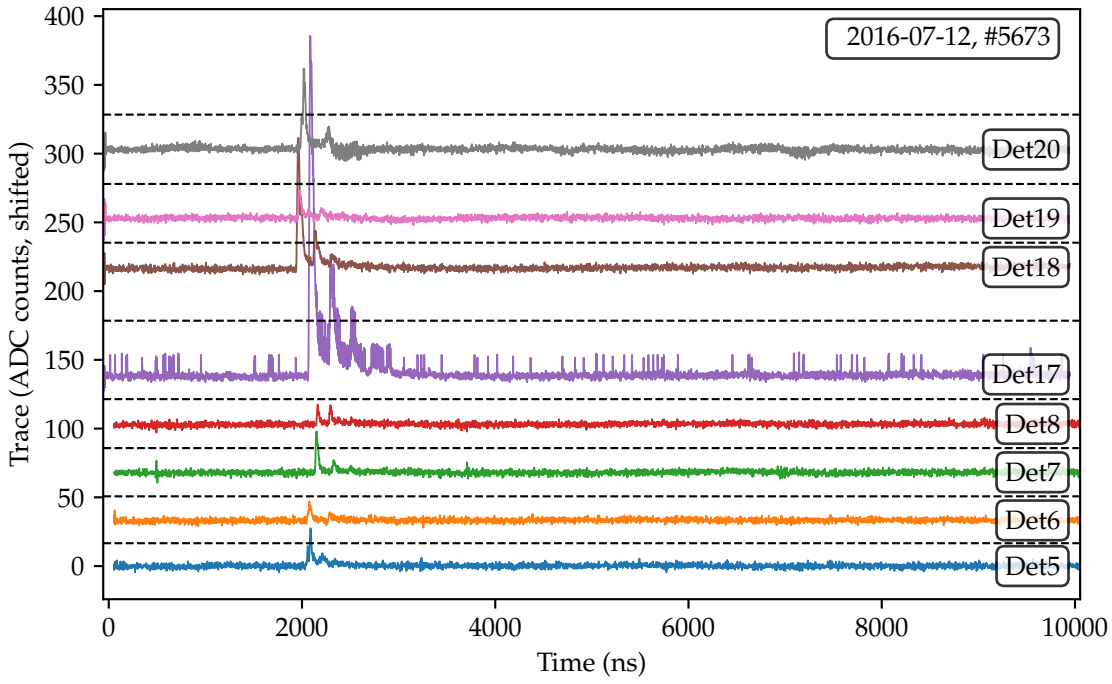

FIGURE 4.4: An example of a normal, respectable cosmic ray event. The y-axis is in units of ADC counts, but traces of different detectors were shifted to plot them side-by-side clearly. The dashed black lines indicate the 5σ above mean levels, when including the 100 MHz signal.

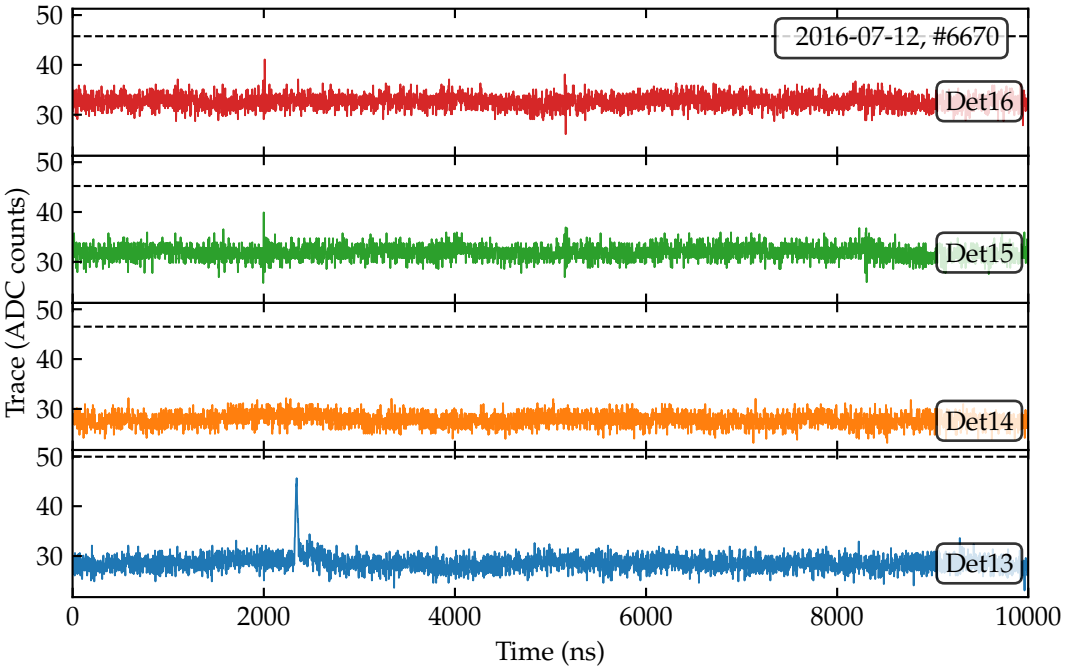

FIGURE 4.5: An example of a questionable event trace. In contrast to Figure 4.4, the signal does not reach the 5σ-level, and shows a tick pattern, at about 2000 ns, 5160 ns and 8330 ns. As in Figure 4.4, the traces of different detectors were shifted.

4.1.3 Singles count rate

Another interesting quantity are the singles count rate or, as it is called in the data, the threshold count rate. This registers the number of times the signal went above the required threshold level. Each station saves two columns of count rates, for two 'channels'. It is still unclear to us what these channels correspond to, and of what detectors it is saved. How exactly the threshold is chosen, is not completely clear, but it is approximately 5σ, it is the same level as the required level for a weak event. The data is saved each second, and long term variations are shown in Figure 4.6. The singles count rates generally range from $50\,s^{-1}$ to $150\,s^{-1}$. This is reasonably close to the $115\,s^{-1}$ that was calculated in Section 3.1.2. But many peculiar features are visible: firstly, in station 2 (L2, the red and green lines), there are many sharp rises and falls that give several orders of magnitude higher singles rates. Secondly, there are dips, in which many detectors report very low count rates (e.g. on 2016-07-11 at 3:00). Thirdly, around 2016-07-10 10:00, there also is a constant rise of channel 2 of L5, to extreme levels, lasting for several hours. But for all detectors, there discontinuities in the count rates. None of these features are thunderstorm-correlated, and I do not know their origin.

To try to understand the jagged features that all detectors have, I zoomed in to one detector on one day, next to the singles rate, I also plotted the mean and standard deviation that traces had on these days. On top of that, I plotted a threshold level that the data files specify. From this plot (Figure 4.7), it becomes clear that the discontinuities in the count rate come from a change in the threshold level: the

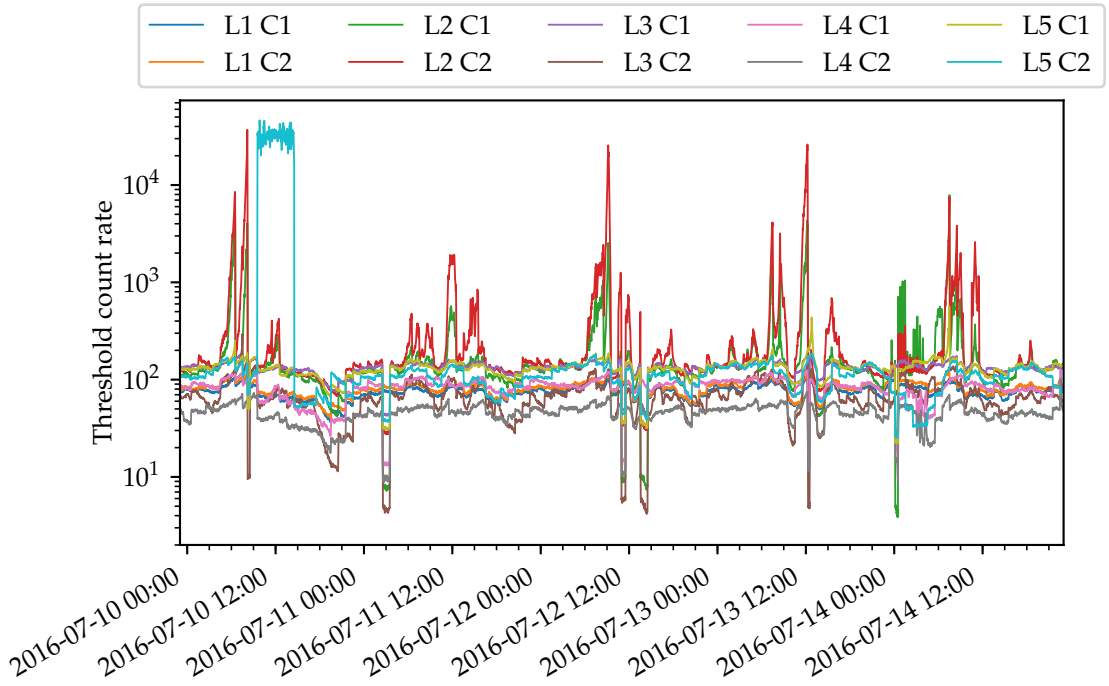

FIGURE 4.6: The singles count rate. For each station (L1 to L5 in the figure), for each second, the amount of times the photomultiplier tube signal went above a threshold (of about 5σ) is counted and saved. There are two channels for each station (C1 and C2 in the figure).

sudden rises and falls happen when the threshold level changes (at the blue dots)[1]. However, the relation between this threshold level, and the background mean and standard deviation remains unclear. One hypothesis is that a higher threshold level simply gives a lower count rate, and vice versa. This often happens, but about half of the time the singles count rate rises with a higher threshold, so the behaviour remains mysterious. This erratic behaviour is not limited to these couple of days: it happens all the time. Because of all this, unfortunately, the singles rates are not applicable for the detection of gamma-ray glows.

4.2 LORA during thunderstorm-rich days

Although quiet days already showed many features in the data, I also looked at some of the days where there was much lightning activity around LORA (as measured by the total number of discharges), namely 2018-05-28 and 2018-05-31. The number of discharges and the radar reflectivity are shown in Figure 4.8: some tens of lightning discharges were detected on 2018-05-29 at 14:00 and on 2018-06-01 at 14:00. The event rates and threshold count rates are shown in Figures 4.9 and 4.10. Again, even though these thunderstorms were some of the most intense in the last 7 years, there is no correlation between these quantities and thunderstorm activity: there is no event rate rise, and no special behaviour in the singles rates. However, another feature did show up: on 2018-05-31 at 12:00, the threshold count rate went haywire for a long time. And coincidentally with this, the event rate showed some

[1] Actually, there is a small time difference between the change in threshold and the discontinuity, of about a minute. But in all likelihood, this is because the station computer time is not synchronized to the GPS timing.

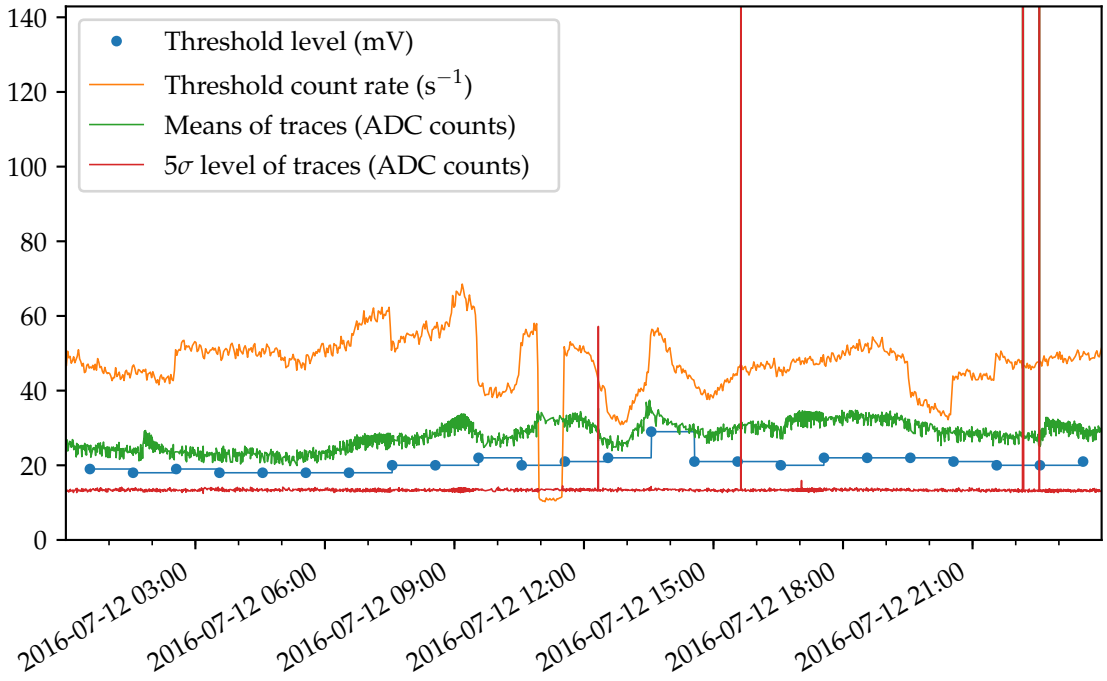

FIGURE 4.7: The threshold level, count rate, and the mean and σ of individual traces, zoomed in at one day for one detector (Det15). Notice that changes in the threshold level (the blue dots) coincide with discontiunities in the threshold count rate.

uncommon behaviour. Upon inspection of these events, it turned out that the traces were completely different than the normal traces: only 2 detectors had their trace written to the data files. An example of this behaviour is shown in Figure 4.11. This is different from the normal 3-detector trigger condition. Why during these times trigger conditions seem to differ, is not known. It is not lightning-correlated, since the same behaviour is visible on other days without any thunderstorms (they are easily found by inspecting the ROOT files with the largest file sizes). Probably, it is therefore just some instrumental quirk.

4.3 Flashes

To look for TGFs in the data, it was necessary to inspect the traces of all events in all of the data. In order to find events that look like TGFs, and not like normal cosmic ray showers, a suitable filter was required. In contrast to normal cosmic ray events, whose main pulse normally lasts much less than a microsecond, TGFs can last up to several hunderds of microseconds. Therefore, I selected the events whose traces varied longer than usual. Specifically: I selected the events in which the standard deviation of the signal in the last 3.75 µs is larger than 8 ADC counts for at least 5 detectors (this way, I selected a reasonable number of events that involve more than one station, but tuning them some more could give better results). If the traces of any TGFs behave similar to the trace shown in Figure 3.3, then this filter should match them.

It turned out that this filter selected two kinds of events. Firstly, there were events where the signal was very strong in the trigger time region and then settled down again to a more stable level. One of them is shown in Figure 4.12. These

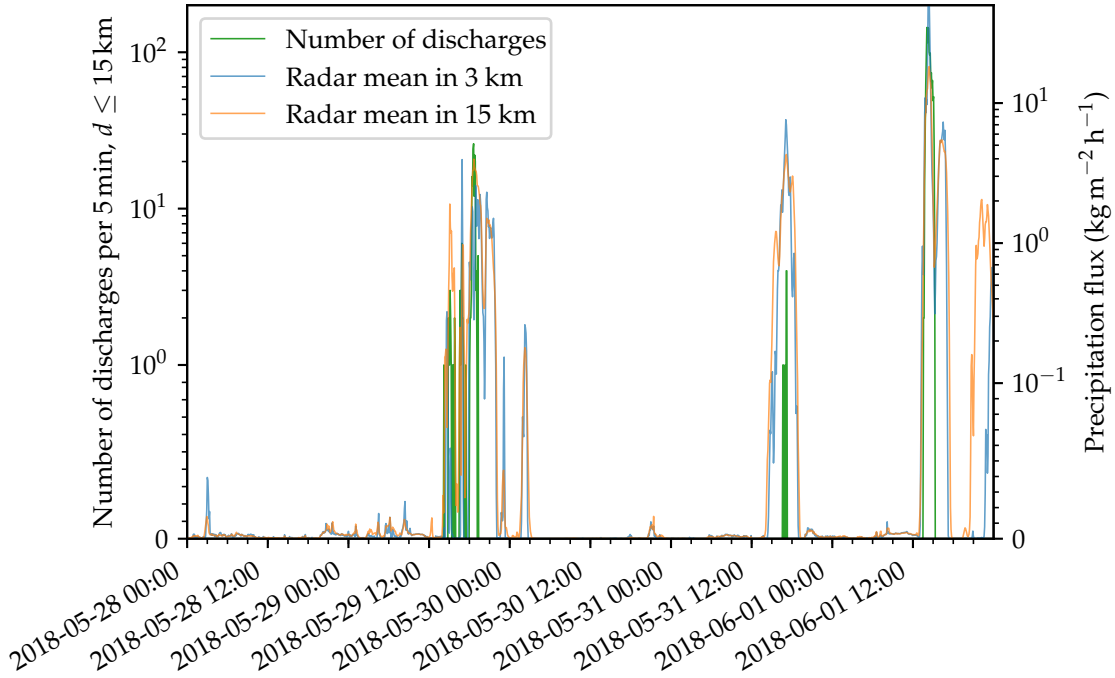

FIGURE 4.8: Thunderstorm and lightning activity data. Data taken from KNMI and Buienradar.

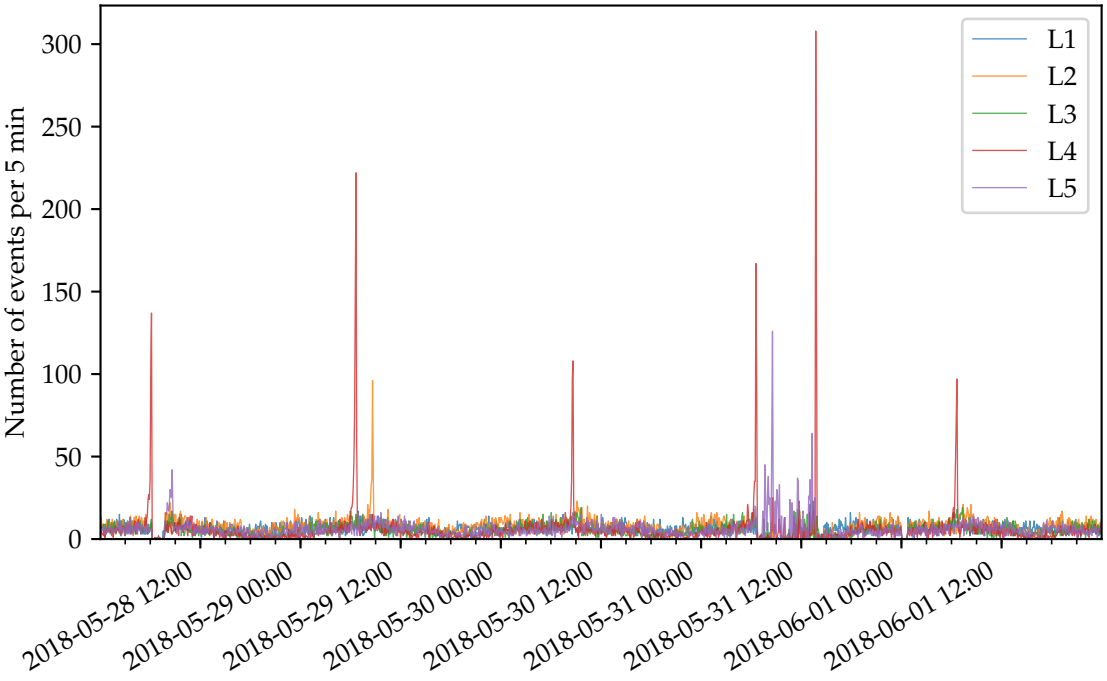

FIGURE 4.9: The event count rate during some days of thunderstorms. Notice the behaviour of the event rate at 2018-05-31 12:00, it is quite high, but vor a long time. Additionally, one can distinguish a day-night pattern in the event rate.

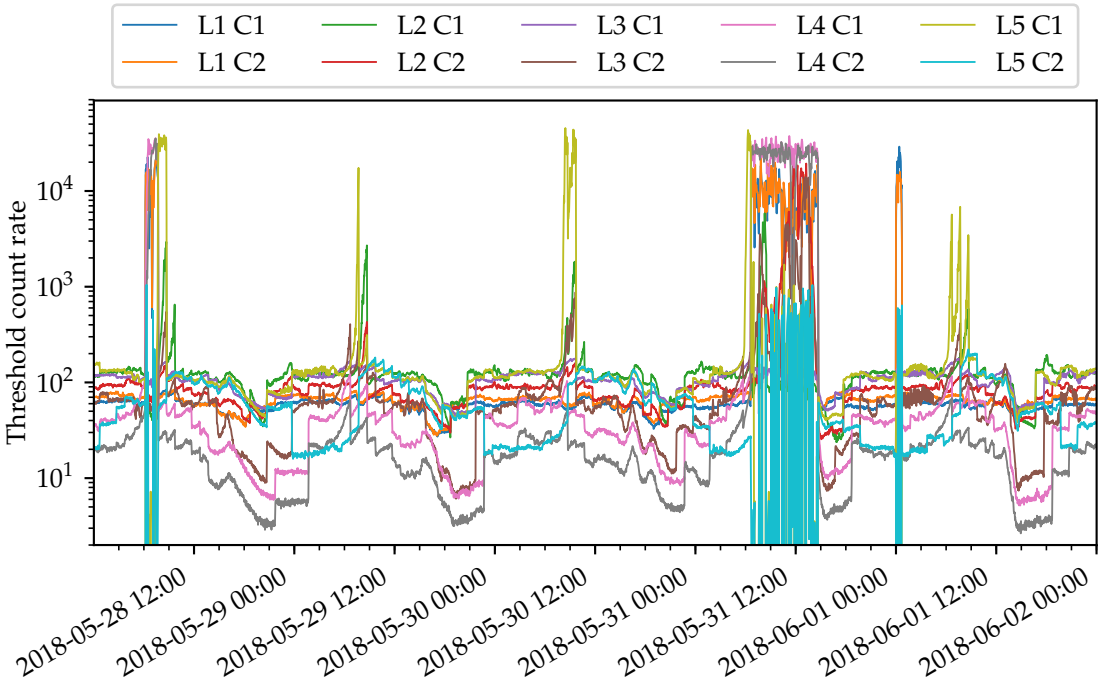

FIGURE 4.10: The threshold count rate during thunderstorm-rich days. Similar to Figure 4.6, L1-L5 refers to the station number, while C1-C2 refers to the channel. Although nothing uncommon takes place during thunderstorm times (2018-05-29 14:00 and 2018-06-01 14:00, see Figure 4.8), during quiet times, especailly on 2018-05-31 12:00, all stations report very high count rates.

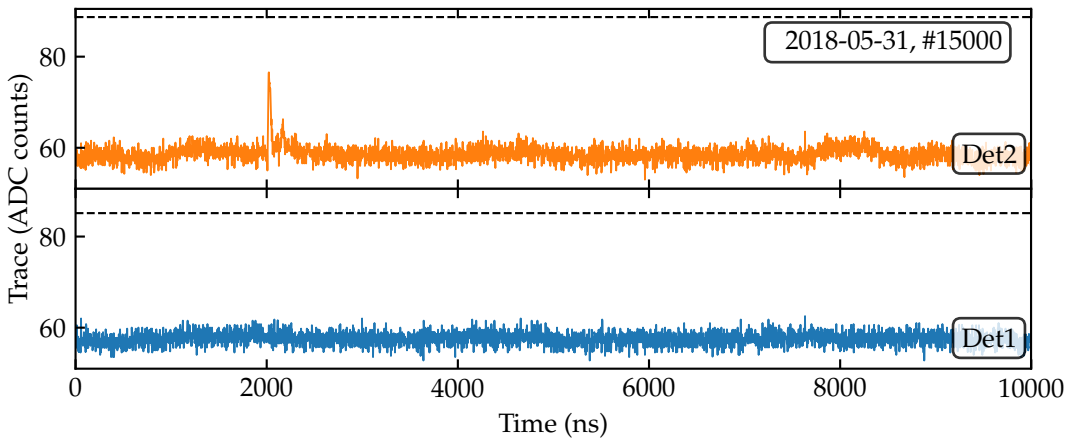

FIGURE 4.11: An example of a trace that only was only saved for two detectors.

events were often strong enough to fully saturate the detectors (a level of $2^{12} = 4096$ ADC counts). They are reminiscent of normal cosmic ray events, except for the fact that some detectors show exceptionally high background levels of thousands of ADC counts after the event. This caused them to satisfy the condition. Presumably this high background level is just some unphysical instrumental quirk. I have checked that they are not lightning-correlated.

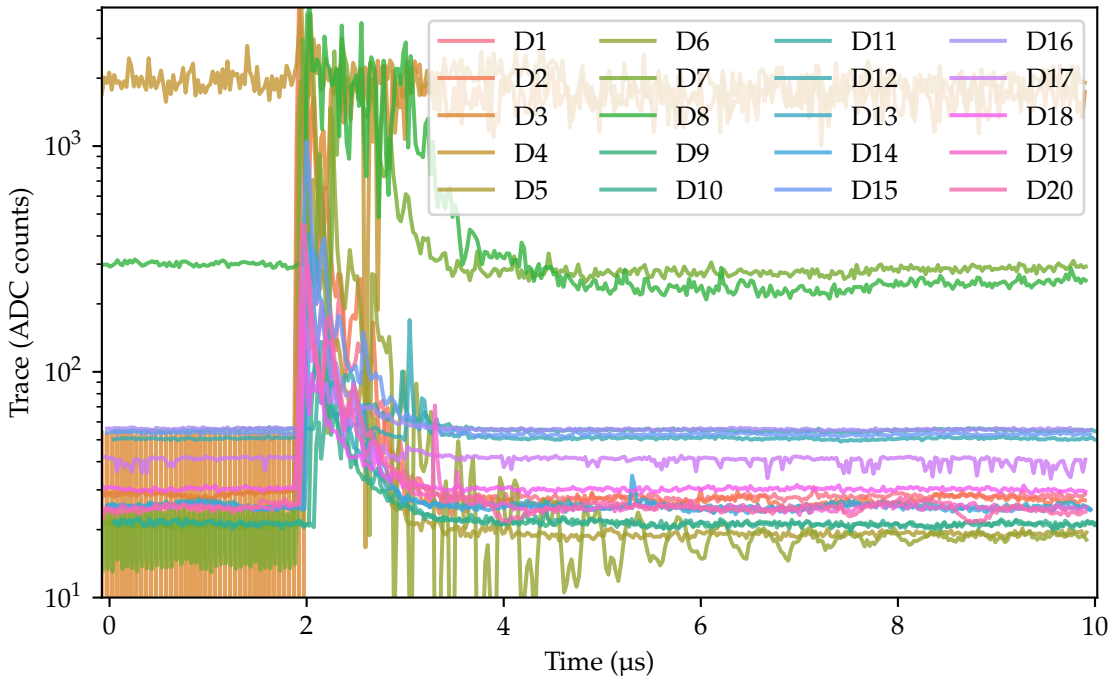

FIGURE 4.12: An example of a trace that showed up when searching for extensively varying signals. Notice the logarithmic axis: some of the detectors have extraordinarily high counts, also after the event. The signal was smoothed, by rebinning with a width of 10 bins.

But more interestingly, there are events where the signal stays relatively low (below several hundreds of ADC counts), but where the signal oscillates, as shown in Figure 4.13. This particular event was the strongest one I found. The signal itself is well fit by a sine-curve with some amplitude modulation. Frequencies these sine-curves vary been between 0.17 MHz to 0.22 MHz (corresponding to periods of 4 µs to 6 µs). The frequency also changes slightly during the 10 µs of one event, and they are different for different detectors in the same event at the same time. The ADC cannot go lower than 0 ADC counts, which causes a cutoff there, making the signal look like some hills. The amplitudes of the sine-signals vary up to some hundreds of ADC counts. Another example of such an event is in Figure 4.14. Notice that the event at least took some tens of microseconds, and possibly longer. Here, the signal is not strong enough to reach the 0 ADC counts cutoff. The events do not seem to be causally connected to a single continuous source: the times of the peaks (the phases) of the signals do not follow a clear pattern when considering the geometrical layout of the LORA detectors. When looking at the events that occurred directly before or after these, it turned out that this behaviour could in some cases last for some seconds. If the wave has some velocity, its apparent ground velocity is much less than c, because the time differences between peaks at different detectors are much more than light travel times between the detectors.

I at least found 35 such events on 18 different days. And these are just the strongest ones, there are much more of these events. The really interesting part? On all of these flashes, there was lightning activity close. Usually some tens to hundreds of discharges were detected in a 15 km radius. Radar reflectivities also indicated thunderstorm activity in all of them. Some possible causes are speculated in Chapter 5.

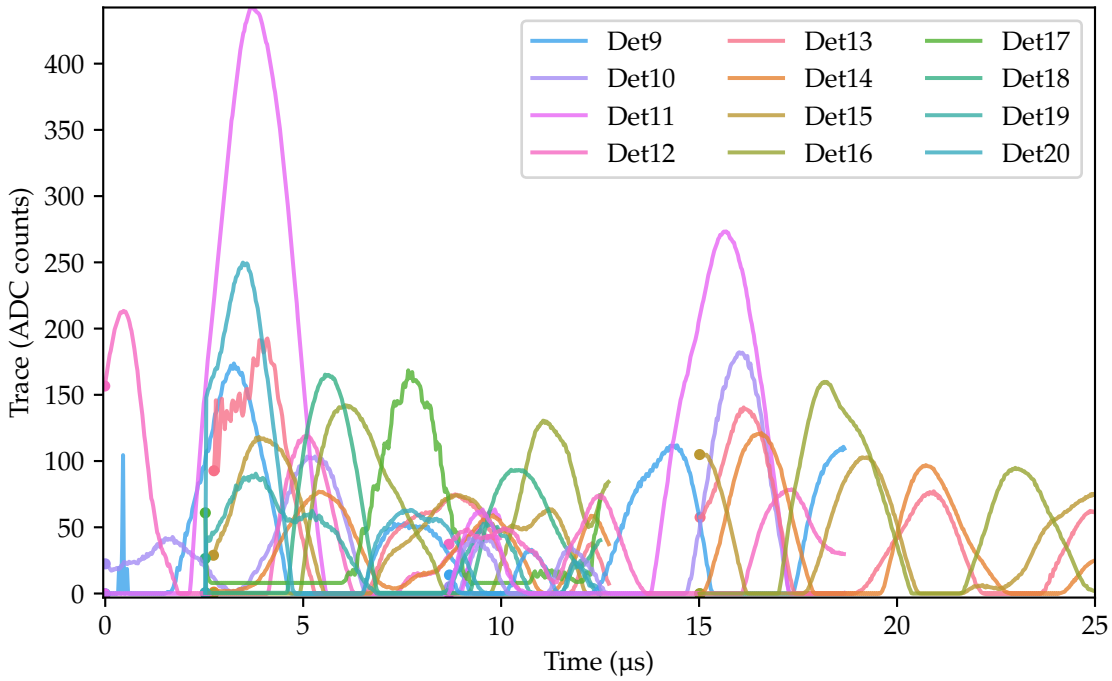

FIGURE 4.13: An example of thunderstorm-related traces, detected on 2016-08-28 at 00:16. In the half hour time window centered in the flash, there were 60 discharges in a 15 km radius. Figure 4.15 shows that there indeed was lightning activity on this time. These are several different events occurring close to each other. The signal was smoothed, by rebinning with a width of 10 bins. On this day, unfortunately only 3 out of 5 stations were operational.

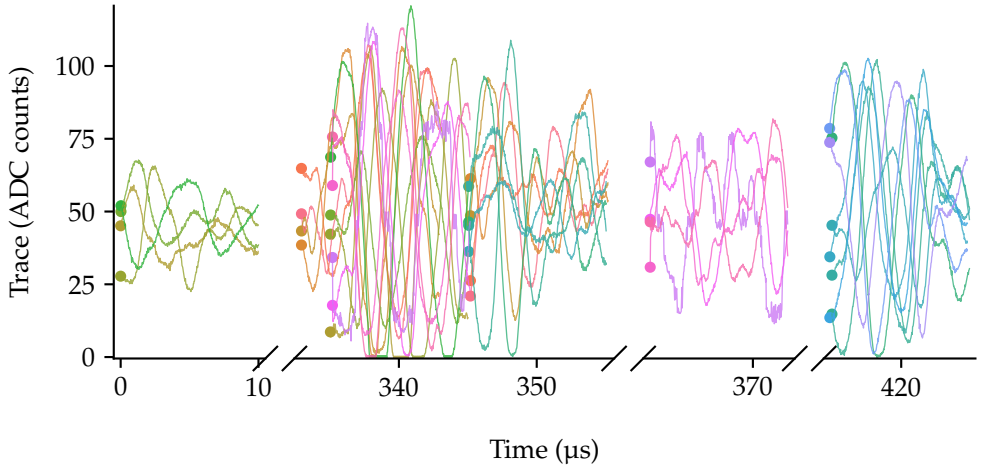

FIGURE 4.14: Another example of a thunderstorm-related trace, detected on 2018-06-01 at 13:51. Figure 4.8 shows that there indeed was lightning activity on this time. As with Figure 4.13, the data was rebinned with a width of 10 bins, but a broken axis was used to show all events occurring during this second.

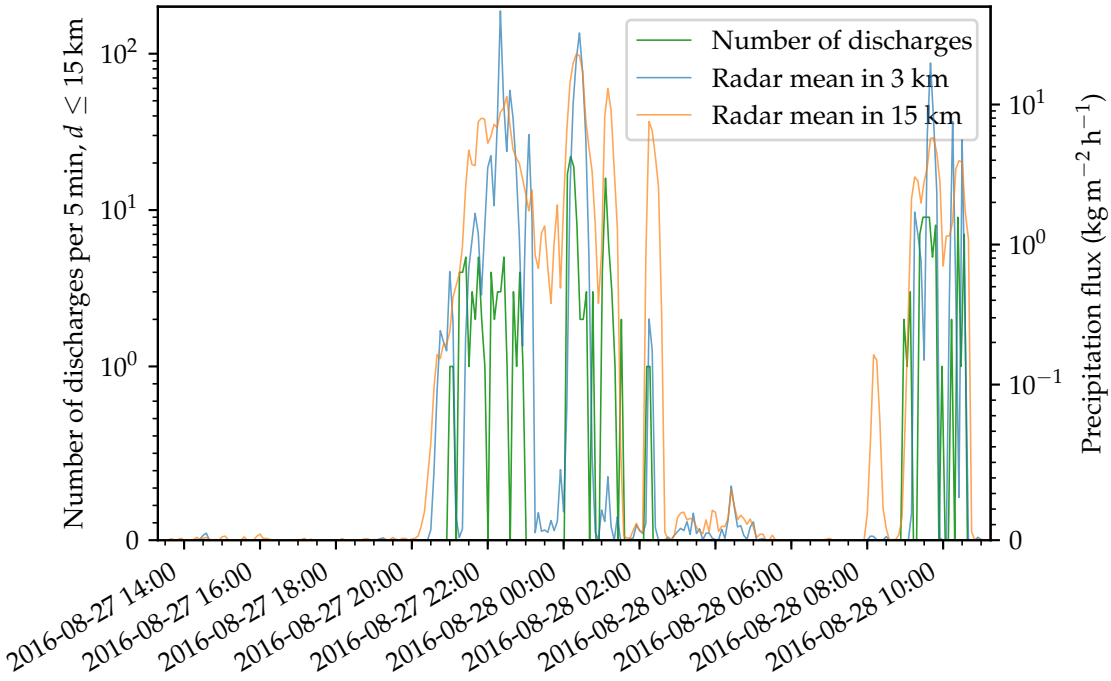

FIGURE 4.15: The lightning activity on the day of the flash shown in Figure 4.13.

Chapter 5

Discussion

In Chapter 3, we investigated the detectability of gamma-ray glows with LORA. Many assumptions and simplifications were made, but none of them should have resulted in an overestimate of what could be measurable: it is probably an underestimate. On the other hand, the spectrum from Chilingarian et al. (2010) was the strongest one they observed in many years of observations. I found that event rates should not change significantly due to gamma-ray glows. This is because an event requires a coincidence of 3 detectors, which is rather rare because the gamma-ray flux is uncorrelated. However, the single detector count rates should give a significant rise in flux because of gamma-ray glows.

It turned out however (Chapter 4), that the LORA event rates and threshold count rates showed many sharp instrumental features. In particular, the event rates showed spikes that came from an unknown instrumental defect. The cause of these spikes is still unclear. But they occur mainly in two pairs of detectors, and there is not a temporal pattern associated with them. The cause likely is related the clock-like ticks that the traces of events during spikes have, spaced 3.16 µs apart from each other.

Closely related to the single-detector count rates are the threshold count rates, which showed many jagged features with common unexplained spikes and dips. This is likely the result of many environmental and instrumental effects, because there are so many distinct features. What factors (instrumental or environmental) contribute to the final count rate is also not certain. I did check that they were not caused by maintenance of the detectors. Even when trying to ignore all the instrumental effects of the data, the data does not behave noticeably different in the case of thunderstorm activity. The conclusion we therefore reach is that the LORA data is not suited for detecting gamma-ray glows.

The other high-energy thunderstorm phenomenon that I looked for is evidence of TGFs in the data. I estimated the effect of a TGF on the LORA traces, and found that this is substatial. This seemed to be confirmed by some traces that I found showed long-term patterns of some tens of microseconds. These traces turned out to be clearly lightning-correlated. The structures I observe, however, do not agree with what one expects from a TGF. In the following, I present some speculation about their origin.

An attractive explanation would be that they are caused by TGFs. The timing and the duration of the events is reasonably consistent with this, but much does not make sense. For example, notice how different these signals (Figures 4.13 and 4.14) are from previously published TGF time traces (Figure 3.3). For one, the signals are much too smooth, and there is not a causal relation between the signals in different detectors. The frequency of the signals also differs across different detectors. The signals might still be the result of a TGF-like phenomenon, but the signals cannot be explained by gamma-rays creating this event.

The signal might also be caused by strong changes in the electromagnetic field that affect the operation of the LORA detectors. These fields could interfere with the electronics or with the photomultipliers inside the detector boxes. The fact that the signals have different frequencies could be explained by different response behaviours of the electronics of different detectors, since each detector is calibrated separately. But such interference is quite unexpected, since the scintillator plates are inside heavy metal boxes to obey the radio-quiteness conditions that are required at the LOFAR core, and the electronics are in cabinets that are very well shielded from the surroundings.

Or perhaps the lightning activity disturbed the local power grid, which influenced components that are not properly shielded against this. For instance, slightly different and voltages on the sensitive photomultipliers might explain it, as the PMT dark count rate (and thus the ADC counts of the traces) changes as a function of the voltage applied. But then, it is strange that the signals do not have amplitudes that are very different across detectors in different stations.

One of the major problems while trying to explain these events, is that the LORA detectors are not made to trigger on such events. They are made for cosmic rays, which requires a trigger on a short, intense spike. The observed sinusoidal signals can not be related to particles detected in the LORA detectors. The recorded data have large gaps (see Figures 4.13 and 4.14), which also makes it difficult to trace their origin.

As the result of this work, I suggest several possible improvements to LORA. First, there are some issues that would need to be solved: the clock-like tick pattern that was seen in Figure 4.5 is one clear example, since it creates a signal where there should not be one. Secondly, one of the detectors (Det17) has an issue that creates many random spikes on top of the data. This causes it to have a background noise standard deviation of at least twice the normal value. Additionally, it should be checked if the station computer has its clock synchronized: it looked like the station computer timestamps were off by about a minute, when compared to the GPS timestamps. Finally, the 100 MHz signal that is present on the traces of all detectors should be removed. That would considerably improve the signal-to-noise ratio. It should not affect the integrated number of counts, but a better S/N ratio would also propagate to more accurate count rates.

One simple, but helpful addition would be to add some thermometers that log the temperature at the scintillators/photomultipliers. In particular, the temperature affects the photomultiplier's performance, so this effect can then be properly accounted for. This would also help to distinguish lightning-related radiation from the effects related to changes in temperature and pressure.

In the near future, a complete revision of the LORA electronics is planned, and it would be nice to include there to include the possibility to store longer time traces, maybe in a special lightning-observation mode, that would facilitate the data analysis.

Appendix A

LORA data structure

Top tree	Column name	Description
`Tree_sec .Lasa[1-5]` (Saved each second)	`Lasa, YMD`	Station number, Date
	`Channel[1,2] _Thres_count_ [low,high]`	Number of times per second the analogue signal went over some high threshold level, which is about 5σ above the background level.
	`GPS_time_stamp`	Time stamp in unix format
	`sync`	GPS synchronization
	`CTP`	Count-ticks between PPS (the number of counts of the 200 MHz clock between each Pulse-per-second GPS message)
	`quant`	PPS quantization error
	`Satellite_info`	For 12 GPS satellites, the satellite number and levels
`Tree_noise .Det[1-20]` (Saved each hour)	`Detector`	Detector number
	`YMD, HMS`	Date and time
	`GPS_time_stamp`	GPS Time stamp in unix format
	`Mean`	Mean ADC count level (background)
	`Sigma`	Standard deviation of the ADC count level
`Tree_log .Det[1-20]` (Saved each hour)	`Detector, YMD, HMS, Time_stamp`	Detector number, date, time and timestamp
	`Channel_thres_ [low,high]`	Threshold levels
	`[Pre_coin,Coin, Post_coin]_time`	Coincidence time window and how much data to save
	`...`	Various other calibration parameters and settings
`Tree_event .Det[1-20]` (Saved for each weak event)	`Detector, YMD, GPS_ Time_stamp`	Detector number, date, and GPS timestamp
	`nsec`	Nanosecond timing (in units of 0.1 ns)
	`CTD`	200 MHz clock counter
	`counts`	The array containing the trace of the event (4000 entries, with $\Delta t = 2.5$ ns)
	`Trigg_condition, Trigger_pattern`	Information on the trigger conditions (usually just 3), and pattern (usually just 0)
	`Pulse_height`	Peak value of the trace (background subtracted)
	`Pulse_width`	Root mean square of the pulse

TABLE A.1: The structure of the LORA raw ROOT files. For the purposes of this project, the most important fields are the traces, event timings, and the threshold count rates. Sources: Verkooijen (2016) and K. Mulrey (personal communication)

Bibliography

Chilingarian, A. (July 2018). "Long lasting low energy thunderstorm ground enhancements and possible Rn-222 daughter isotopes contamination". In: Phys. Rev. D 98, 022007, p. 022007. DOI: 10.1103/PhysRevD.98.022007.

Chilingarian, A. et al. (Aug. 2010). "Ground-based observations of thunderstorm-correlated fluxes of high- energy electrons, gamma rays, and neutrons". In: Phys. Rev. D 82, 043009, p. 043009. DOI: 10.1103/PhysRevD.82.043009.

Dwyer, Joseph R., David M. Smith, and Steven A. Cummer (Nov. 2012). "High-Energy Atmospheric Physics: Terrestrial Gamma-Ray Flashes and Related Phenomena". In: Space Sci. Rev. 173, pp. 133–196. DOI: 10.1007/s11214-012-9894-0.

Dwyer, Joseph R. and Martin A. Uman (Jan. 2014). "The physics of lightning". In: Phys. Rep. 534, pp. 147–241. DOI: 10.1016/j.physrep.2013.09.004.

Eack, Kenneth B., William H. Beasley, et al. (Dec. 1996). "Initial results from simultaneous observation of X rays and electric fields in a thunderstorm". In: *Journal of Geophysical Research* 101, pp. 29, 637–29, 640. DOI: 10.1029/96JD01705.

Eack, Kenneth B., David M. Suszcynsky, et al. (Jan. 2000). "Gamma-ray emissions observed in a thunderstorm anvil". In: *Geophysical Research Letters* 27, pp. 185–188. DOI: 10.1029/1999GL010849.

Fokkema, David B.R.A. (2012). "The Hisparc cosmic ray experiment: data acquisition and reconstruction of shower direction". PhD thesis. URL: http://inspirehep.net/record/1203533/files/thesis_D_Fokkema.pdf.

Grieder, P. K. F. (2010). *Extensive Air Showers: High Energy Phenomena and Astrophysical Aspects - A Tutorial, Reference Manual and Data Book*. Springer. ISBN: 978-3-540-76940-8. DOI: 10.1007/978-3-540-76941-5.

Hare, B. M. et al. (June 2016). "Ground-level observation of a terrestrial gamma ray flash initiated by a triggered lightning". In: *Journal of Geophysical Research (Atmospheres)* 121, pp. 6511–6533. DOI: 10.1002/2015JD024426.

Kelley, Nicole A. (Jan. 2014). "Long duration gamma-ray emission from thunderclouds". PhD thesis. University of California, Santa Cruz. URL: https://escholarship.org/uc/item/6009v2kg.

McCarthy, M. and G. K. Parks (June 1985). "Further observations of X-rays inside thunderstorms". In: *Geophysical Research Letters* 12, pp. 393–396. DOI: 10.1029/GL012i006p00393.

Thoudam, S. et al. (Dec. 2014). "LORA: A scintillator array for LOFAR to measure extensive air showers". In: *Nuclear Instruments and Methods in Physics Research A* 767, pp. 339–346. DOI: 10.1016/j.nima.2014.08.021. arXiv: 1408.4469.

Torii, Tatsuo, Minoru Takeishi, and Teruo Hosono (Sept. 2002). "Observation of gamma-ray dose increase associated with winter thunderstorm and lightning activity". In: *Journal of Geophysical Research (Atmospheres)* 107, 4324, p. 4324. DOI: 10.1029/2001JD000938.

van Holten, J. W. (Oct. 2007). *Satistics of coincidences*. URL: http://www.hisparc.nl/fileadmin/HiSPARC/Lesmateriaal_fysica__jan-willem_/coinc.pdf (visited on Nov. 14, 2018).

Verkooijen, Hans (Feb. 2016). *Message Structures HiSPARC*. URL: https://docs.hisparc.nl/firmware/messages.html (visited on Nov. 11, 2018).