

Detecting cosmic rays with the LOFAR radio telescope

P. Schellart¹, A. Nelles^{1,2}, S. Buitink^{5,1}, A. Corstanje¹, J. E. Enriquez¹, H. Falcke^{1,2,3,4}, W. Frieswijk³, J. R. Hörandel^{1,2}, A. Horneffer⁴, C. W. James⁶, M. Krause¹, M. Mevius^{3,8}, O. Scholten⁵, S. ter Veen¹, S. Thoudam¹, M. van den Akker¹, A. Alexov^{33,7}, J. Anderson⁴, I. M. Avruch²³, L. Bähren¹, R. Beck⁴, M. E. Bell²⁶, P. Bennema³, M. J. Bantum³, G. Bernardi^{29,30}, P. Best¹⁶, J. Bregman³, F. Breitling¹⁹, M. Brentjens³, J. Broderick¹², M. Brügger²⁷, B. Ciardi¹³, A. Coolen³, F. de Gasperin²⁷, E. de Geus³, A. de Jong³, M. de Vos³, S. Duscha³, J. Eisloffel²⁰, R. A. Fallows³, C. Ferrari³², M. A. Garrett^{3,9}, J. Griebmeier²², T. Grit³, J. P. Hamaker³, T. E. Hassall^{12,10}, G. Heald³, J. W. T. Hessels^{3,7}, M. Hoefft²⁰, H. A. Holties³, M. Iacobelli⁹, E. Juette²¹, A. Karastergiou¹¹, W. Klijn³, J. Kohler⁴, V. I. Kondratiev^{3,28}, M. Kramer^{4,10}, M. Kuniyoshi⁴, G. Kuper³, P. Maat³, G. Macario³², G. Mann¹⁹, S. Markoff⁷, D. McKay-Bukowski^{34,15}, J. P. McKean³, J. C. A. Miller-Jones^{14,7}, J. D. Mol³, D. D. Mulcahy⁴, H. Munk³, R. Nijboer³, M. J. Norden³, E. Orru³, R. Overeem³, H. Paas²⁴, M. Pandey-Pommier²⁵, R. Pizzo³, A. G. Polatidis³, A. Renting³, J. W. Romein³, H. Röttgering⁹, A. Schoenmakers³, D. Schwarz³⁶, J. Sluman³, O. Smirnov^{29,30}, C. Sobey⁴, B. W. Stappers¹⁰, M. Steinmetz¹⁹, J. Swinbank⁷, Y. Tang³, C. Tasse¹⁷, C. Toribio³, J. van Leeuwen^{3,7}, R. van Nieuwpoort^{35,3}, R. J. van Weeren³¹, N. Vermaas³, R. Vermeulen³, C. Vocks¹⁹, C. Vogt³, R. A. M. J. Wijers⁷, S. J. Wijnholds³, M. W. Wise^{3,7}, O. Wucknitz^{18,4}, S. Yatawatta³, P. Zarka¹⁷, and A. Zensus⁴

(Affiliations can be found after the references)

Received 16 September 2013 / Accepted 4 November 2013

ABSTRACT

The low frequency array (LOFAR), is the first radio telescope designed with the capability to measure radio emission from cosmic-ray induced air showers in parallel with interferometric observations. In the first ~ 2 years of observing, 405 cosmic-ray events in the energy range of 10^{16} – 10^{18} eV have been detected in the band from 30–80 MHz. Each of these air showers is registered with up to ~ 1000 independent antennas resulting in measurements of the radio emission with unprecedented detail. This article describes the dataset, as well as the analysis pipeline, and serves as a reference for future papers based on these data. All steps necessary to achieve a full reconstruction of the electric field at every antenna position are explained, including removal of radio frequency interference, correcting for the antenna response and identification of the pulsed signal.

Key words. astroparticle physics – methods: data analysis – instrumentation: interferometers

1. Introduction

With the development of ever faster electronics and the increase in computational power, the construction of radio telescopes as large interferometric arrays of rather simple antennas opens a new window for observations. The low frequency array (LOFAR; van Haarlem et al. 2013), is the first large-scale implementation of this technique. In addition to producing the first high quality images at these low frequencies of 10–240 MHz, LOFAR was designed to study short, pulsed signals in the time-domain. With a vast array of antennas observing the whole sky simultaneously, observations are not limited to a predefined direction, therefore providing optimal conditions for cosmic-ray detection.

Cosmic rays, accelerated charged particles from astrophysical sources, can be observed over several decades of energy. When cosmic rays of high energies reach the Earth, they do not reach the surface as primary particles, but instead interact with atmospheric nuclei. Thereby, a cascade of particles is created, consisting mostly of photons and a significant fraction of charged particles. While propagating through the atmosphere, the charged particles of this extensive air shower emit electromagnetic radiation, which adds up coherently for wavelengths comparable to the dimensions of the shower front (Huege 2013).

Already in the 1960s it was proven that cosmic ray-induced air showers emit nanosecond duration pulses with significant power in the MHz radio frequency range (Jelley et al. 1965; Allan & Jones 1966), but due to lack of sufficiently sophisticated and fast electronics the technique was not pursued further. Only in the past decade interest in the detection technique was rekindled and successfully applied (Falcke et al. 2008). The proof of principle and large progress in the understanding of the emission was made at the LOFAR Prototype Experimental Station (LOPES; Falcke et al. 2005; Huege et al. 2012) and further refined by measurements at the CODALEMA experiment (Ardouin et al. 2005).

Similar to optical measurements of the fluorescence emission from atoms excited by interaction with the air shower, radio emission directly traces the longitudinal shower development, which is closely related to the type of the primary particle. Unlike optical fluorescence measurements, radio emission measurements are less dependent on observing conditions and can operate day and night matching the duty cycle of particle detector measurements.

Due to the very steep energy spectrum, measuring the highest-energy cosmic rays requires vast detector areas. Cost constraints therefore limit the density of detectors within this area giving a wide spacing between the individual antennas.

Theoretical models describing the different emission mechanisms at play point to a very detailed and non-symmetrical emission pattern at ground level (Werner et al. 2012; Alvarez-Muñiz et al. 2012; Marin & Revenu 2012; Huege et al. 2013). Testing these models therefore requires dense sampling of the electric field over a sufficiently large area.

LOFAR offers a high number of antennas clustered on an irregular grid, with increasingly large spacing between antenna clusters further away from the center. In the core of the array about 2300 antennas are installed within about 4 km², which allows air showers to be measured with unprecedented spatial resolution. These measurements will contribute significantly to conclusively confirm theoretical models for the radio emission on a shower by shower basis, a goal previously unattainable due to lack of sufficiently high quality data.

Measurements and converging theoretical predictions of the expected radio signal from a cosmic-ray induced air shower give a short, nanosecond time-scale bi-polar pulse, which is mostly linearly polarized. This article describes the detection set-up and automated processing pipeline used at LOFAR to measure and identify these signals.

Starting with a description of the instrumental set-up at LOFAR in Sect. 2, an overview of the data reduction pipeline is given in Sect. 3. Finally, Sect. 4 describes the characteristics of the dataset obtained between June 2011 and April 2013.

The LOFAR dataset will be used in forthcoming publications to verify existing models for radio emission from air showers and to develop new techniques that use radio emission to measure important characteristics of the incoming particle, such as energy and mass.

2. LOFAR

LOFAR is a distributed radio telescope. Its antennas are distributed over northern Europe with the densest concentration in the north of the Netherlands, in the Province of Drenthe. The observation support center and processing facilities are also located near this central core. The antennas of LOFAR are grouped into *stations*, each station taking the role of a single dish in a traditional radio interferometer array. A station consists of a number of low-band antennas (LBAs, 10–90 MHz) and high-band antennas (HBAs, 110–240 MHz). The 24 stations within the ~2 km wide core are distributed in an irregular pattern that maximizes *uv*-coverage, or spatial frequencies for standard interferometric observations. The 16 additional Dutch remote stations are distributed with increasing distance to the core. International stations are currently located in Germany, France, the United Kingdom, and Sweden, giving LOFAR a maximum baseline of 1292 km for interferometric observations. Core stations and remote stations consist of 96 LBAs plus 48 HBAs. International stations have 96 LBAs and 96 HBAs. At the center of the LOFAR core six stations are located in a roughly 320 m diameter area, called the *Superterp*, providing both the shortest baselines for interferometric observations and the densest population of antennas ideal for cosmic-ray observations. While every LOFAR station is equipped with the necessary electronics to observe cosmic rays, the current data set is taken with the central 24 stations, where additional information from particle detectors is available (see Sect. 2.3). The positions of the antennas of the seven most central LOFAR stations are shown in Fig. 1.

2.1. The antennas

The LBAs are the main tool for cosmic-ray detection. An LBA consists of two orthogonal inverted V-shaped dipoles, each with

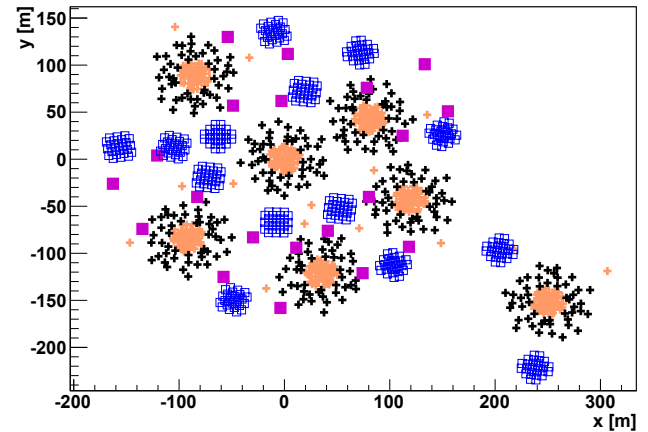


Fig. 1. Layout of the center of LOFAR. The six stations to the left form the Superterp. The crosses indicate the LBA inner and outer antenna sets, respectively. The open squares show the positions of the HBA tiles, which are split into two groups per station. The filled squares indicate the positions of the LORA particle detectors.



Fig. 2. Low-band antennas at the central core of LOFAR, the Superterp. In the background the black box of a LORA particle detector can be seen.

a length of 1.38 m. These are supported by a central PVC¹ pole, which holds the low-noise amplifier and guides the signal cables, as shown in Fig. 2. The dipoles *X* and *Y*, that make up each antenna, are oriented northeast to southwest (NE-SW) and northwest to southeast (NW-SE), as can be seen in Fig. 3.

The low-noise amplifier has an intentional impedance mismatch with the antenna. This mismatch, combined with the characteristic length of the dipoles, makes the system sensitive in a broad band from 10–90 MHz. In principle, this allows observations from the ionospheric cutoff up to the start of the commercial FM radio band. For most observations the frequency range is limited by a combination of selectable hardware and software filters to 30–80 MHz to suppress strong Radio Frequency Interference (RFI) in the outer bands. The LBAs are designed to be sky noise limited after RFI has been removed (van Cappellen et al. 2007). After amplification the signals from the individual dipoles are transmitted through coaxial cables to the electronics cabinet located at every station.

The HBAs have been optimized for a frequency band of 110–240 MHz. The design clusters 16 antenna elements into a *tile*, the signals from these elements are amplified and combined in an analog beam-former. This means that while the LBAs

¹ Polyvinyl chloride.

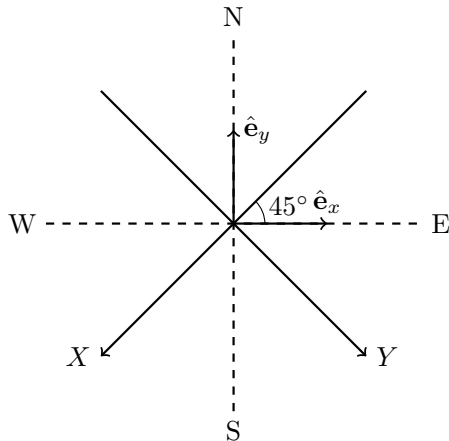


Fig. 3. Geometry of the LBA. The X and Y dipoles are oriented NE-SW and NW-SE respectively. This is rotated by 225 degrees with respect to the standard local Cartesian coordinate system used in Sect. 3.4.

are sensitive to the whole sky the HBAs are most sensitive within the $\sim 20^\circ$ of the tile-beam, of which the direction is chosen at the start of every observation. This results in a smaller effective area for cosmic-ray observations, as the measurement will only be optimal if the direction of the cosmic ray happens to coincide with the beam direction of the observation. Therefore, the analysis of HBA data and their interesting higher frequency range requires a different approach for cosmic-ray studies. Results of these measurements will be described in a later publication.

2.2. The transient buffer boards

After being forwarded to the electronics cabinet the signals of the LBAs are again amplified, filtered, and digitized by a 12 bit A/D converter with a sampling frequency of 200 MHz². Due to signal path limitations in the Dutch stations only 48 dual-polarized or 96 single-polarized antennas can be processed at a given time. For the dual-polarized option the antennas are grouped into an inner and an outer set, which has to be chosen before an observation.

For astronomical observations the data are then beamformed and sent to the central processing facility. In addition, there is the possibility to store a snapshot of the original data. Every station is equipped with ring-buffers, the so called Transient Buffer Boards (TBBs). These continuously store the last 1.3 s of raw data (an extension to 5 s is currently being deployed). When triggered, the contents of the TBBs are frozen, read out via the Wide Area Network and stored on disk for further analysis. The trigger can be generated based on various parameters in an FPGA³ at the local receiver unit. Alternatively, the trigger can be generated by an array of particle detectors (see Sect. 2.3) or received from outside of LOFAR. Currently, the main trigger for cosmic-ray observation is provided by the particle detectors. Later, a radio self-trigger will be implemented, using the current dataset as a training set to deduce trigger criteria, so that the FPGA trigger can be run independently at every LOFAR station. These criteria have to reduce false triggers to limit the data rate. Using every LOFAR station individually will dramatically increase the effective area.

Essential for measuring cosmic rays with LOFAR as a radio telescope is that the whole process of triggering and storing

² A 160 MHz clock is also available.

³ Field Programmable Gate Array.

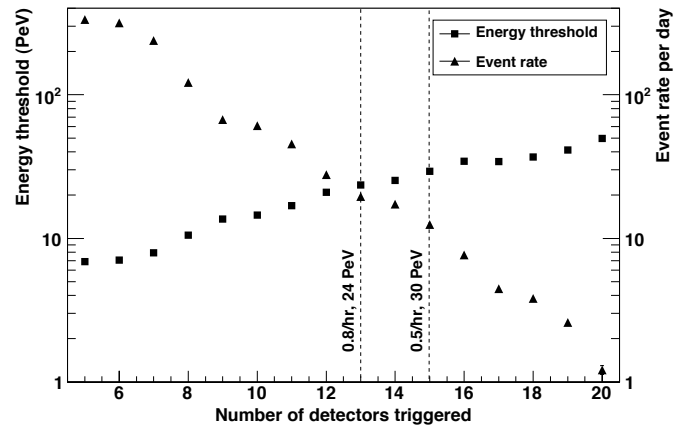


Fig. 4. Energy threshold in PeV (*left*) and the event rate per day (*right*) are shown as a function of the number of triggered particle detectors. Two possible trigger conditions are indicated with the dotted lines.

radio-pulse data can take place without interfering with the ongoing observations.

2.3. The LOFAR Radboud Air Shower Array

LORA, the LOFAR Radboud Air Shower Array, is an array of particle detectors co-located with the center of LOFAR. The array provides a reconstruction of basic parameters of recorded air showers, such as the direction and the position of impact, as well as the energy of the incoming cosmic ray (Thoudam et al., in prep.). It also provides the time of arrival, which is used to trigger the read-out of the radio antennas.

LORA consists of 20 detector units distributed on the Superterp, as shown in Fig. 1. Each detector contains two scintillators (0.45 m^2 , type: NE 114), which are individually read out through a photomultiplier tube. The detectors are inside weather-proof shelters and have been tested to not create any interference at radio frequencies.

Conditions at which triggers are sent to LOFAR can be adjusted to match the desired energy threshold. There are two constraints on the desired rate: the rate of events interesting for radio observations has to be maximized, while the network load on the LOFAR system has to be kept low in order to avoid interfering with the primary observation. A trigger in a single detector is generated when a particle signal of more than 4σ above the noise is registered. In order to only detect air showers a coincidence of several detectors is needed. Events of less than 10^{16} eV have a very low probability to be observable in radio above the sky-noise level. The energy threshold and the corresponding event rate are shown in Fig. 4 as the function of the number of triggered detectors. Requiring triggers in 13 detectors yields a threshold energy of $2.4 \times 10^{16} \text{ eV}$, with an average trigger-rate of 0.8 events/h. This trigger rate has been selected as the optimal setting for the observations.

2.4. Observations

After the commissioning phase LOFAR is to be used on a proposal-based schedule. Proposals are open to the community for imaging or beamformed observations, as well as TBB observations. Some fraction of the observing time is reserved for participating consortia and key science projects. The LOFAR cosmic ray key science project (CRKSP) is one of six LOFAR key science projects.

To maximize the duty cycle TBB observations can be run in the background of all other observations that do not need the full network bandwidth. This does however mean that the array configuration is determined by the primary observation, therefore the amount of data in a specific array configuration (such as the selection of LBA or HBA antenna type) available for analysis is not determined by the cosmic-ray project itself, except when LOFAR is otherwise idle and the observing configuration can be chosen freely.

During the observation, triggers from LORA are received by the LOFAR control system. The system checks whether a dump from the TBBs is allowed. If so, the ring-buffers are frozen and a specified block of data around the trigger time is dumped to disk. For each cosmic-ray event 2.1 ms of radio data are stored, which corresponds to 77 MB per station. This provides sufficient frequency resolution for high quality RFI cleaning while minimizing data transfer and storage requirements.

Every evening, the data-files are archived at LOFAR and compressed for transport. They are stored in the Long Term Archive (van Haarlem et al. 2013), from where they can be retrieved for data analysis.

3. Reconstruction of cosmic-ray data

All newly recorded data are processed every evening, after having been copied via the network to the processing cluster of the Astrophysics department at the Radboud University Nijmegen. In addition to the HDF5⁴ files, containing the data of one LOFAR station each, the recorded data from the particle detectors and a trigger log file are transferred. With this information an automated pipeline is run. The pipeline is based on the task oriented PyCRTools framework consisting of fast low-level C++ routines embedded in Python for maximum flexibility. All results are stored in a PostgreSQL database for subsequent data mining analysis. The goal of the processing pipeline is to autonomously identify a full set of physics quantities for each air shower detected with LOFAR. The pipeline is optimized to identify those nanosecond pulses that are not generated by terrestrial sources.

All data are first processed per station, i.e. per file. The set of files received for a single trigger form an *event*. When the data from one station pass the criteria for containing a cosmic-ray signal (see Sect. 3.3), the corresponding event is called a *cosmic-ray event*. It is not necessary to observe a pulse in all stations, only the stations with a significant signal are used in a combined analysis.

3.1. Pipeline structure

The reconstruction pipeline comprises a number of steps that will be individually explained in the following sections. An overview of the steps and the overall structure is depicted in Fig. 5.

3.2. Preparing the data

Before proceeding to extract the cosmic-ray signal from the data, some preparatory steps have to be performed. Knowledge about the system is applied in the form of calibration procedures, the data are cleaned of narrowband-transmitters, and antennas that show malfunctions are flagged.

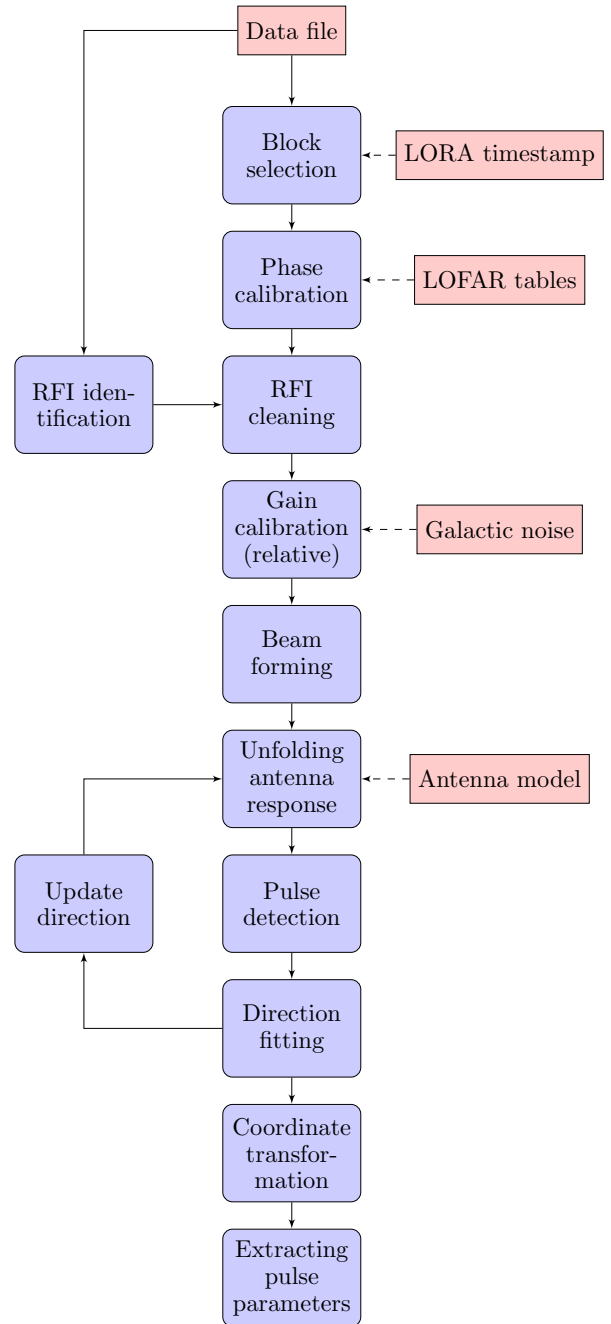


Fig. 5. General structure of the analysis pipeline. Rectangles represent input and rounded squares are processing steps.

3.2.1. Timing offsets and phase calibration

There are known signal path differences between the LOFAR antennas. Measured differences of cable lengths between the antennas are corrected for up to the 5 ns sample level already at the stations before the data are written to disk. Additionally, relative time offsets between the antennas are corrected for at sub-sample accuracy using standard LOFAR calibration tables. These tables are generated by phase-calibrating on the strongest astronomical radio sources and are regularly tested and updated if necessary (van Haarlem et al. 2013). Sub-sample corrections are applied as phase offsets to the Fourier transformed signal in the cosmic-ray pipeline, before processing it in the data analysis.

⁴ A tree-like file format (Alexov et al. 2012).

3.2.2. RFI cleaning

Narrow-band RFI in the time series signal can be revealed by making an average power spectrum. An example is shown in the top panel of Fig. 6, where most of the strong RFI is visible outside the 30–80 MHz range. The average power spectrum is created by averaging the square of the absolute value of the Fourier transform over several blocks of data. The block size can be freely chosen within the full data length to obtain a desired frequency resolution; here 2^{16} samples are used, giving a resolution of ~ 3 kHz, enough to resolve most RFI lines. A reasonable data length is needed for this procedure to produce a stable average, which sets the limit for the chosen block length to be stored from the TBBs, as mentioned in Sect. 2.4. In order to minimize artificial side lobes a half-Hann window is applied to the first and last 10% of each trace prior to the Fourier transformation.

The standard approach to RFI cleaning (or RFI flagging) is to identify peaks sticking out significantly above the overall spectral shape, also called the *baseline*, and set the corresponding Fourier component amplitudes to zero. However, this requires a priori knowledge of the baseline. While the baseline can be obtained through a smoothing or fitting procedure, this is often not stable in the presence of strong RFI, requiring an iterative approach.

An alternative approach to RFI cleaning uses the phase information in the complex-valued spectrum instead. If an RFI transmitter is measured in all antennas, the phase difference, or relative phase, between each pair of antennas will be a constant value as function of time with a small non constant random noise contribution. Note that the exact value of the constant, which only depends on the geometric delay between antennas, is not relevant, only its non time-varying nature. When no transmitter is present, the relative phase is expected to be both random and time varying, as the signal then consists of the added signals from many incoherent sources on the sky with additional random noise. Therefore, RFI can be identified by looking at the stability of phase differences between antennas over time. For each antenna-dipole $j = 0, 1, \dots, 95$ in a station and data block k , the phase spectrum is calculated as

$$\phi_{j,k}(\omega) = \arg(x_{j,k}(\omega)), \quad (1)$$

where $x_{j,k}(\omega)$ is the complex frequency component ω of the spectrum.

Subtracting the phase of one of the antennas as reference antenna gives the relative phases and results in a set of phases for every frequency channel, one for each block of data. Only one reference antenna is used and this is taken to be the one with median power to avoid selecting a broken antenna.

The average phase is defined as

$$\bar{\phi}_j(\omega) = \arg\left(\sum_{k=0}^{N-1} \exp(i\phi_{j,k}(\omega))\right), \quad (2)$$

and the phase variance as

$$s_j(\omega) = 1 - \frac{1}{N} \left| \sum_{k=0}^{N-1} \exp(i\phi_{j,k}(\omega)) \right|, \quad (3)$$

where N is the number of data blocks.

For completely random phases one expects $s_j(\omega) \approx 1$ as opposed to $s_j(\omega) = 0$ when all phases are equal. The phase variance per frequency channel will now either be at a value close to 1, including some random noise, or at a significantly lower level.

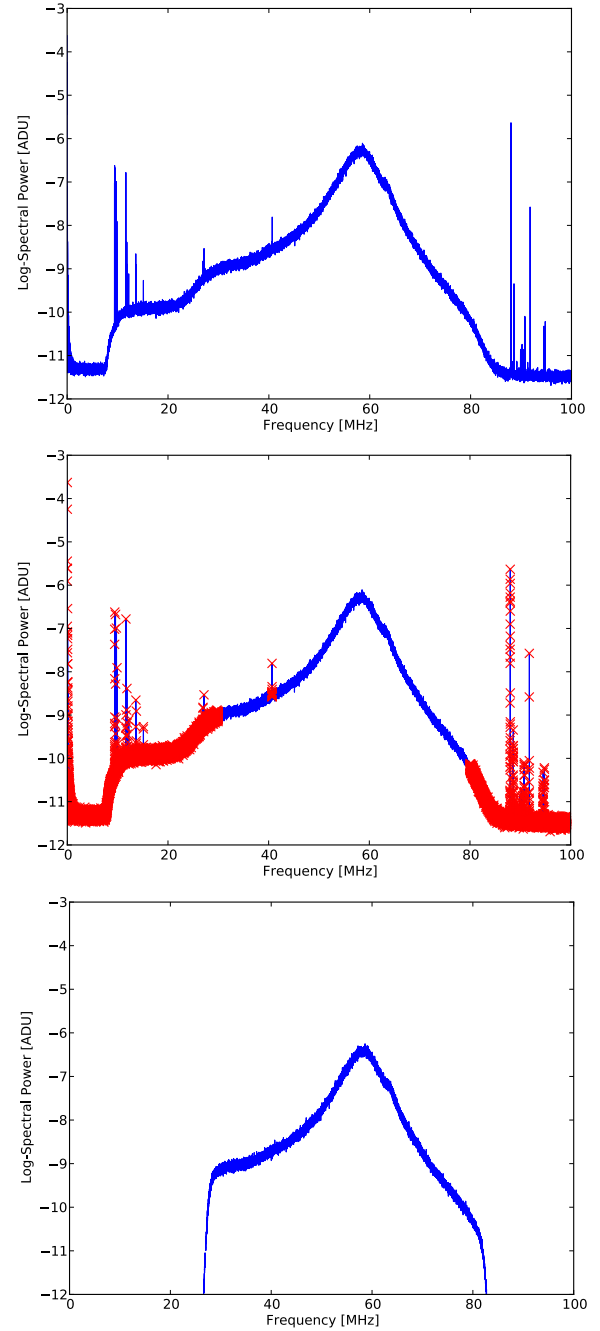


Fig. 6. The average spectrum of a typical LBA event. The raw data (*top*), with flagged contaminated channels (*middle*), cleaned and clipped to 30–80 MHz (*bottom*).

The latter reveals the presence of a radio transmitter, as shown in Fig. 7, where a contaminated part of the spectrum is shown with the corresponding phase variance.

Since RFI lines will result in peaks toward smaller values of the phase variance, and noise has no preferred peak direction, calculating the standard deviation σ in this plot only for values above the median will ensure a stable result. All frequencies that have a phase variance of at least 5σ below the median⁵ are flagged as containing RFI. Additionally a 30–80 MHz band-pass filter is applied, flagging the most heavily RFI polluted

⁵ Assuming a Gaussian distribution, σ can be estimated by sorting the data points, and comparing the value at 95 percentile to the median. This difference amounts to $\sim 1.64\sigma$.

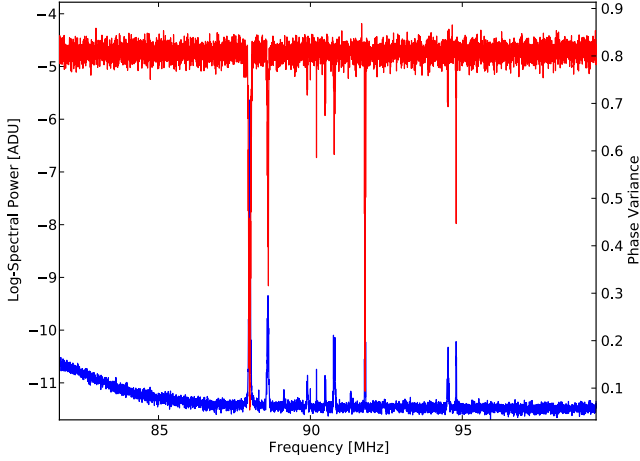


Fig. 7. Average LBA spectrum (bottom, left axis) with the corresponding phase variance (top, right axis). RFI lines can clearly be identified in the phase variance with peaks toward lower values, representing more stable phase differences between antennas over time.

low and high frequency parts of the bandwidth by default. To prevent pulse-ringing the 30–80 MHz block filter is first convolved with a, $\sigma_{\text{tapering}} = 2.5$ MHz, Gaussian⁶. After removing the flagged channels, the resulting cleaned spectrum is shown in the bottom panel of Fig. 6.

In general, there is very little RFI at the LOFAR Superterp. A lot of effort has been made to remove local sources that could disturb the LOFAR measurements and a protected zone has even been established (Offringa et al. 2013). This relative quietness is illustrated in Fig. 8. It shows the result of the RFI cleaning for all events for all frequencies. While every event has some RFI, no single RFI line is present in every event. Within the 30–80 MHz band, there are only two lines that are present in more than 40% of the events. In total there are rarely events with more than 2% flagged channels out of the more than 32 000 frequency channels in a block of data. This is shown in Fig. 9, where the total fraction of events is plotted against the number of flagged channels.

3.2.3. Flagging bad antennas

Occasionally, one or more antennas give invalid signals, e.g. due to hardware malfunction. To identify these *bad antennas* the integrated spectral power is calculated

$$P = \int_{30\text{ MHz}}^{80\text{ MHz}} |x(\omega)|^2 d\omega, \quad (4)$$

where $x(\omega)$ is the ω frequency component of the cleaned spectrum. The power in every antenna is required to be in the range of one half to two times the median power from all antennas. Antennas outside this range are marked as bad and excluded from further analysis.

3.2.4. Absolute gain calibration

There are ongoing efforts for an absolute calibration of the voltage traces of LOFAR and therefore the reconstructed electric

⁶ This effect also occurs when flagging large blocks of RFI but this does not happen in practice and so no tapering window is applied for this case.

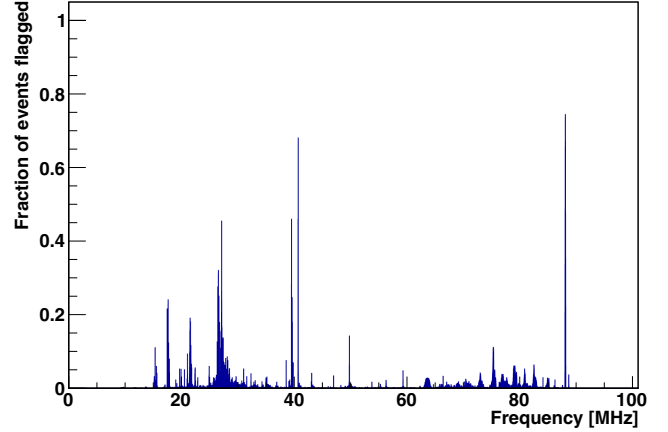


Fig. 8. Fraction of events that is affected by narrow-band RFI in each of the ~ 3 kHz frequency channels as function of frequency.

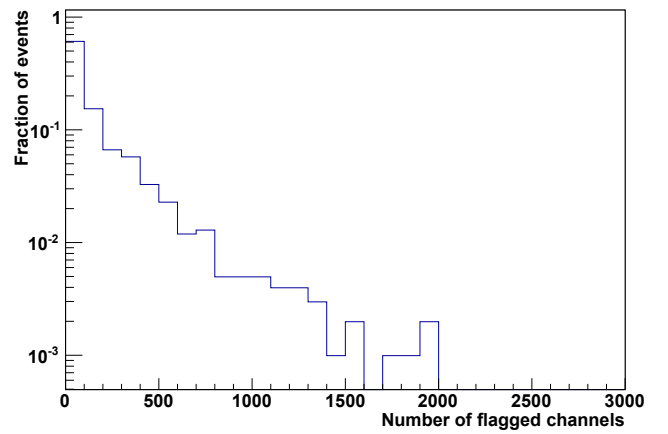


Fig. 9. Relative fraction of events with a certain number of flagged channels. Over 60% of the events have less than 100 channels (≈ 300 kHz) flagged out of the full used bandwidth of more than 16 000 channels.

field. Those efforts will be described in a forthcoming publication and include calibration on astronomical sources, terrestrial transmitters, and already conducted dedicated measurement campaigns, similar to those performed at other experiments, e.g. Nehls et al. (2008). Once implemented, the reconstruction pipeline will deliver calibrated electric field strengths and their polarization components for all events. However, significant progress in understanding the mechanisms of radio emission in air showers can already be made with a relative calibration.

3.2.5. Relative gain calibration

The LBA measurement is dominated by sky noise, which in turn is dominated by the Galaxy moving through the antenna beam pattern. Therefore, the noise as seen by each antenna is a function of the Local Sidereal Time (LST) and can be used to correct for differences in gain between antennas. Instead of correcting all antennas at all times to a fixed value, which would be over- or underestimating the noise at certain times, the received power can be normalized to a LST-dependent reference value. In Fig. 10 the integrated spectral power (Eq. (4)), after RFI cleaning, is given as a function of LST for the instrumental polarization X and Y . The data have been retrieved from all cosmic-ray events measured within the first year of data-taking. One can define a reference value for the integrated spectral power as a

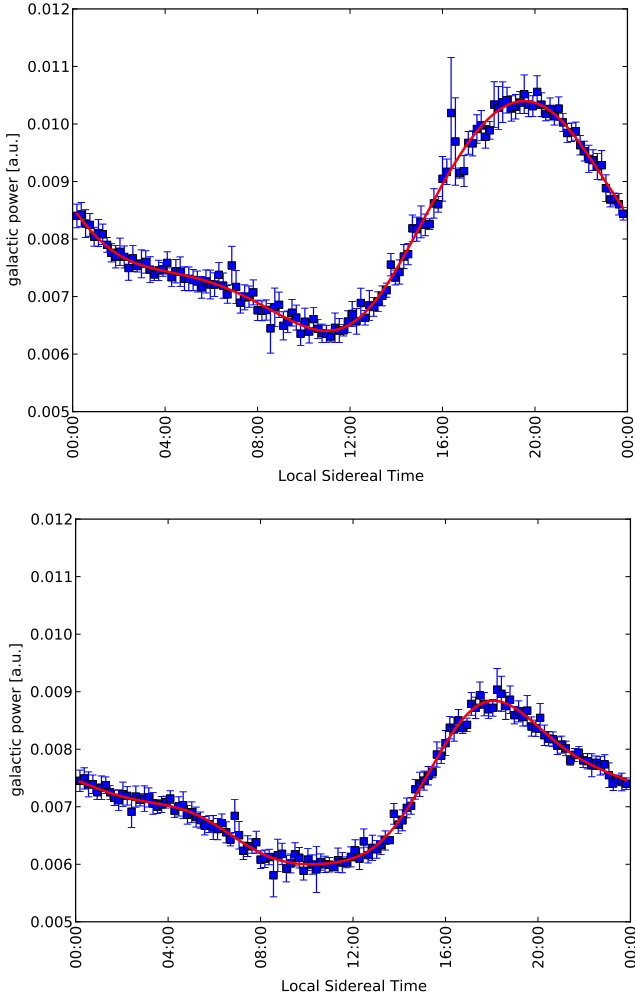


Fig. 10. Integrated spectral power normalized to the bandwidth, after RFI cleaning, as a function of local sidereal time for the X (NE-SW) (top) and Y (NW-SE) (bottom) instrumental antenna polarizations. Also shown is the fitted second order Fourier transform (solid line). The uncertainties on the data still include systematic effects due to the set-up itself, as well as possible artifacts of the RFI cleaning, when having certain frequencies that are contaminated in a significant fraction of the data.

function of LST by fitting a function to these data points. Since the movement of the Galaxy through the antenna beam pattern is periodic by nature it is fitted with the 2nd order Fourier series

$$P_{\text{ref}}(t) = \frac{a_0}{2} + \sum_{n=1}^2 a_n \sin(nt) + b_n \cos(nt), \quad (5)$$

thereby avoiding artificial jumps in the fit at 0:00 LST. The time t is given in units of radian here. This results in a gain correction for each antenna as

$$x'(\omega) = \sqrt{\frac{P_{\text{ref}}(t)}{P(t)}} x(\omega), \quad (6)$$

where the square root is needed, because the correction is applied to the amplitude spectrum.

3.3. Identifying cosmic-ray signals

After cleaning and calibration of the data, the central element of the pipeline is the identification and characterization of the radio pulse as the signal of the air shower.

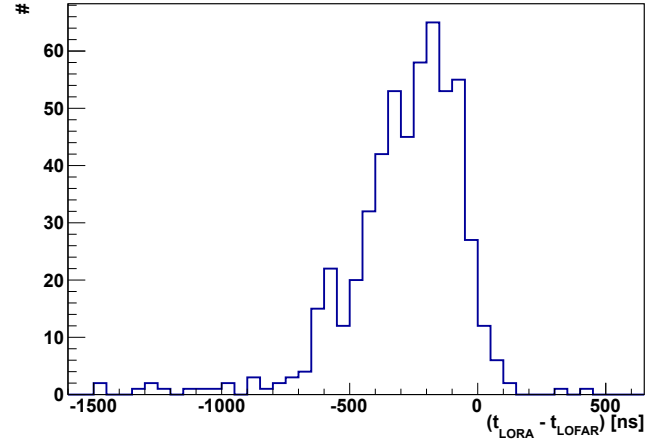


Fig. 11. Difference in time between the time of a pulse identified in the radio signal and the trigger time set by the signal in the particle detectors. This plot shows the distribution summed over all Superterp stations.

3.3.1. Using information from the particle detectors

In order to restrict the search for the radio pulse to a smaller region in the trace, the information from the trigger time of the particle detectors is used. Figure 11 shows the difference in time between the trigger from the particle detectors and the pulse location in the radio data obtained from a search with a large window. The distribution shows a clear peak at the region of the coincidences at an offset of 253 ± 168 ns. In absolute timing the offset between LORA and LOFAR is 10 253 ns, of which 10 000 ns are already accounted for in the triggering system.

Average offset is obtained by fitting a Gaussian to the distribution of pulse positions with respect to the trigger time. This is only an approximation, as the real offset per event depends on the position of the core and the incoming direction of the air shower. Also, effects due to the propagation of particles and radiation in the atmosphere can play a role. The overall difference is due to the fact that both detectors operate independently on different timing systems. Both are based on GPS timing, but correct for drifts (<20 ns) in different ways and have a differing absolute time. The spread on the differences is however sufficiently small for Superterp stations to not require additional synchronization of the two systems. Stations further away can have larger offsets due to the signal travel time, which can be corrected for after a reconstruction of the shower.

These measurements allow for the pulse search to be restricted to a small fraction of the full time trace, limiting the chance to pick up random noise fluctuations.

3.3.2. Finding candidate events

The trigger threshold of the scintillator array is chosen to be lower than the threshold to detect a radio signal. This ensures a full sample, but also makes it necessary to identify in a first quality check whether there is a detectable signal present. Therefore, per antenna polarization, the signals are first *beamformed* in the direction reconstructed from the data of the particle detectors. This direction is given in the local Cartesian coordinate frame of the station by \mathbf{n} and the position of each antenna j is given by \mathbf{r}_j . A planar wavefront arriving at the phase center $(0, 0, 0)$ at time $t = 0$ will arrive at antenna j with a delay given by

$$\Delta t_j = -\frac{1}{c} \frac{\mathbf{n} \cdot \mathbf{r}_j}{|\mathbf{n}|} = -\frac{1}{c} \hat{\mathbf{e}}_n \cdot \mathbf{r}_j, \quad (7)$$

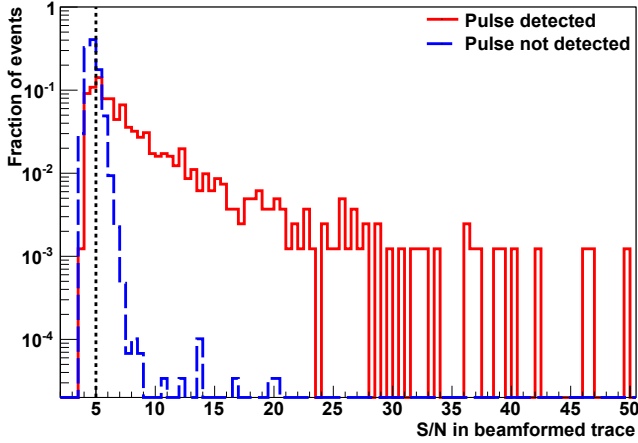


Fig. 12. Distribution of the signal-to-noise ratio (S/N) in initial beamforming. The S/N is defined the ratio of the peak amplitude of the beamformed trace and the RMS of this trace. Two cases are separated: a cosmic-ray event was ultimately detected by the pipeline (solid line) or not (dashed line). The initial cut, which is applied in the pipeline, is indicated by the dotted line.

where c is the speed of light. The beamformed signal, in frequency space, in this direction is then given by

$$x_{\text{br}}(\omega) = \sum_{j=0}^{N_a} x_j(\omega) e^{2\pi i \omega \Delta t_j}, \quad (8)$$

where $x_j(\omega)$ is frequency component ω of the Fourier transform of the signal from antenna j and N_a is the number of antennas. The inverse Fourier transform gives the beamformed time series signal. Due to beamforming any signal coming from the direction of the air shower is amplified by a factor N_a in amplitude while uncorrelated noise is only amplified by a factor $\sqrt{N_a}$. Therefore, if no significant signal is detected in the beamformed trace, the event very unlikely contains a cosmic-ray signal strong enough to be detected at single dipole level by the rest of the pipeline. Thus, the analysis of the data of that station is aborted.

To test this assumption, Fig. 12 shows the distribution of the peak amplitude in the beamformed signal per station, distinguishing between events in which ultimately a cosmic-ray was identified and those in which there was not. The peak amplitude is normalized by the root mean square of the trace, as a proxy for the noise contribution. From this it can be seen that the fraction of events where a strong signal is observed in the beamformed trace is significantly higher for stations where eventually a cosmic-ray signal is detected. All events in the tail of the non-detected distribution were visually inspected and identified as broad-band RFI, with pulses differing significantly in shape from those of cosmic rays and directions ultimately deviating significantly from the direction as measured with the particle detectors. This distribution shows that an initial filtering based on a moderate signal-to-noise of beamformed pulses is a quick and effective way to filter out those events that are potentially interesting, as well as further narrowing the search window per antenna reducing false positives for pulse detection.

3.3.3. Correction for the antenna response

The sensitivity of the LOFAR LBA is a complex function of both frequency and direction. Correcting for this *antenna pattern*, i.e. unfolding, requires an initial guess for the pulse direction and in turn may influence the position of the pulse in time and thus

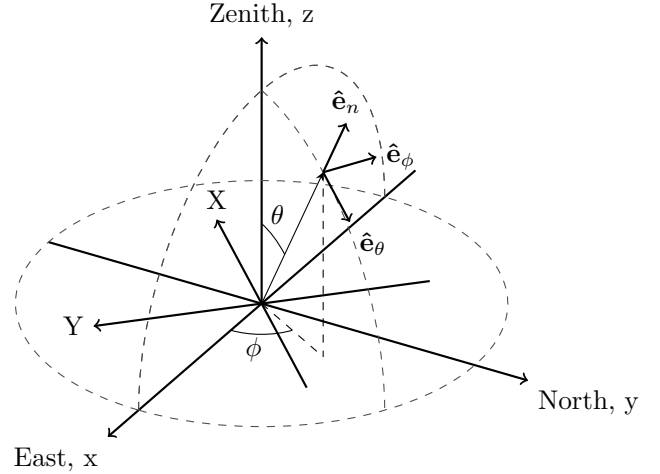


Fig. 13. On-sky polarization coordinate frame ($\hat{e}_\theta, \hat{e}_\phi, \hat{e}_n$). Also depicted is the (north, east, zenith) coordinate frame of the simulations, where the unit vectors ($\hat{e}_x, \hat{e}_y, \hat{e}_z$) correspond to the x, y and z -axis, respectively. Furthermore the dipole antennas X and Y are shown.

the direction by changing the phase at which each frequency arrives. Therefore the correction has to be done in an iterative loop as indicated in Fig. 5. Each iteration starts with an increasingly accurate signal direction and proceeds by unfolding the antenna pattern, pulse detection, and direction fitting. The loop is concluded when the direction no longer significantly changes, which usually happens in less than ~ 5 iterations.

For the antenna pattern of the LBA a simulation is used, which is made using the software WIPL-D (Kolundzija 2011) and a customized software model of the electronics chain.

From the impedance and radiation pattern in a transmitting situation the open circuit voltage is calculated as a function of frequency and direction for an incoming plane wave with an electric field strength of 1 V/m. The equivalent circuit of the antenna in a receiving situation is a voltage source with an internal resistance equal to the antenna impedance. This is combined with measured data of the amplifier directly behind the antenna. The result of the model is the output voltage of the amplifier over a 75Ω resistor⁷.

Any wave coming from a direction \hat{e}_n can be seen as a linear superposition of monochromatic plane waves, polarized in the \hat{e}_ϕ and \hat{e}_θ direction. Here ϕ and θ are the standard spherical coordinate angles with the x and z axis respectively, e.g.

$$\mathbf{E}(t) = \sum_{\omega} (E_{\theta, \omega} \hat{e}_\theta + E_{\phi, \omega} \hat{e}_\phi) e^{-i(kn \cdot x + \omega t)}. \quad (9)$$

This geometry can be seen in Fig. 13.

These terms are related to the output voltage of the amplifier for each dipole, and for each frequency, via the Jones matrix \mathbf{J} (Jones 1941; Hamaker et al. 1996) of the antenna model

$$\begin{pmatrix} V_X \\ V_Y \end{pmatrix} = \begin{pmatrix} J_{X\theta} & J_{X\phi} \\ J_{Y\theta} & J_{Y\phi} \end{pmatrix} \begin{pmatrix} E_\theta \\ E_\phi \end{pmatrix}, \quad (10)$$

where $J_{X\theta}$ is the complex response of the antenna and amplifier of the X -dipole to a wave purely polarized in the \hat{e}_θ direction.

Therefore, in order to both correct for the antenna response and convert from output voltage to electric field strength in the on-sky frame (see Sect. 3.4), each pair of Fourier components

⁷ Matched to the impedance of the coaxial cables connecting the antenna to the station electronics cabinet.

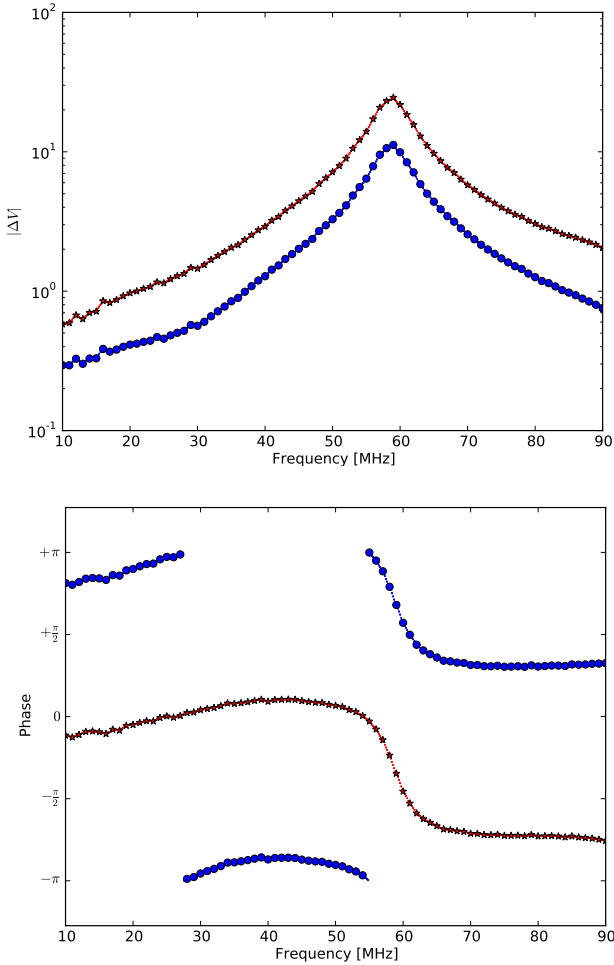


Fig. 14. Jones matrix components of the antenna model amplitudes (*top*) and phases (*bottom*) for a dipole receiving a wave polarized in the \hat{e}_θ direction (circles) and a wave polarized in the \hat{e}_ϕ direction (stars) for an arrival direction of $\phi = 345^\circ$, $\theta = 50^\circ$. Also plotted, as the dotted line, are the interpolated values.

from the signal in the two instrumental polarizations (X, Y) is multiplied by the inverse Jones matrix, followed by an inverse Fourier transform back to the time domain.

The components of the Jones matrix of the antenna model are simulated on a grid with steps of 1 MHz in frequency, 5° in θ and 10° in ϕ . In order to obtain the components at the frequency and direction of observation, trilinear interpolation is performed on the real and imaginary parts of the complex table when needed. Examples of the response are depicted as a function of frequency in Fig. 14 and as a function of direction in Fig. 15.

3.3.4. Pulse detection

Estimating the direction of the incoming air shower, see Sect. 3.3.5, can either be done using beamforming or through pulse timing. Beamforming was found to be very sensitive to the optimization algorithm used, essentially requiring a grid search to avoid getting stuck in a local minimum. This is computationally very expensive, moreover it only provides relative time differences between any two antennas rather than an absolute time needed for extraction of relevant physical parameters (see Sect. 3.4).

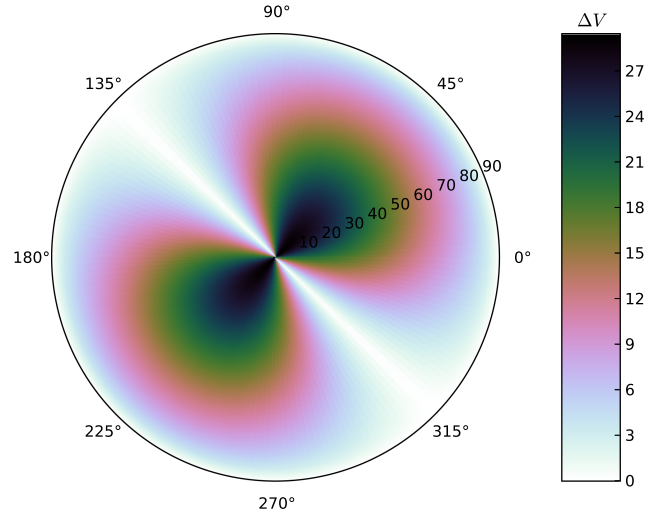


Fig. 15. An example Jones matrix component describing the dipole response, at 60 MHz, $|J_{X,\theta}|$ in the form of the output Voltage (ΔV) as a function of direction for an incoming wave that is purely linearly polarized in the \hat{e}_θ direction.

In order to use pulse timing, individual pulses have to be identified. This can be done by using the cross-correlation method, where one looks for the maximum in the cross correlation of the signals between all antennas. This however has the same drawback as beamforming, as only relative timing is calculated. A method to retrieve the absolute pulse timing is through the use of the *Hilbert envelope*, which is used in this pipeline. A detailed comparison of the methods is given in Sect. 3.3.5.

A sensible definition of the pulse arrival time is the measured arrival time of the maximum of the electric field strength. In practice, however, using directly $\max(|x(t)|^2)$ is highly dependent on the filter characteristics of the receiving system and the sampling used. Therefore, the arrival time is defined as the position of the maximum in the amplitude envelope of the analytic signal, also called the Hilbert envelope

$$A(t) = \sqrt{x^2(t) + \hat{x}^2(t)} \quad (11)$$

where $\hat{x}(t)$ is the *Hilbert transform*, or imaginary propagation, of the signal $x(t)$ defined by

$$\mathcal{F}(\hat{x}(t))(\omega) = -i \cdot \text{sgn}(\omega) \cdot \mathcal{F}(x(t))(\omega) \quad (12)$$

where \mathcal{F} denotes the Fourier transform.

In order to find the pulse maximum with subsample precision, the signal is first up-sampled by a factor 16, such that the maximum search will not be the limiting factor in the timing resolution. Subsequently, a simple maximum search is performed on the envelope. In addition, the signal-to-noise ratio is calculated where the signal is defined as the maximum and the noise as the root mean square of the envelope. An example can be seen in Fig. 16.

This maximum search is performed on each of the on-sky polarizations $E_\theta(t)$ and $E_\phi(t)$ separately and any pulse with a signal to noise greater than three is marked as a possible cosmic-ray signal to be used for direction fitting. Because the pulse is expected to be intrinsically stronger in one of the two polarizations, depending on the angle between the shower axis and the geomagnetic field, the polarization with the highest average signal to noise (over all antennas) is first identified and only its maximum positions are used for the subsequent direction fit.

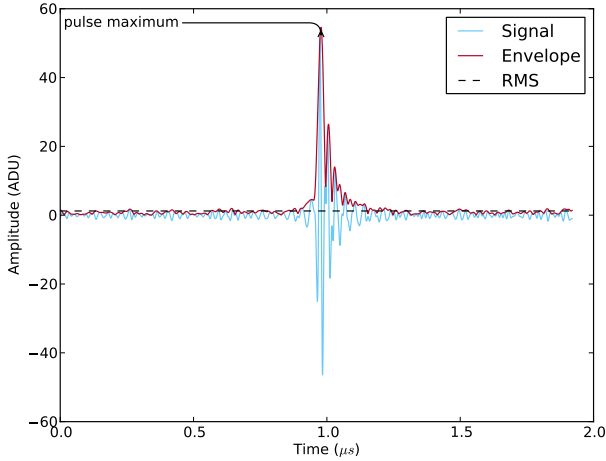


Fig. 16. The solid light line shows the up-sampled signal. Overlaid is the Hilbert envelope and the RMS noise in black dashes. A pulse is accepted whenever the signal to noise ratio exceeds three.

3.3.5. Arrival direction fitting

As described above, every station is processed separately, meaning that the data do not provide a large lever arm for direction fitting. However, it also means that the actual shape of the shower front is an insignificant factor in the direction fitting. For a measurement with a single station, which has a maximum baseline of 80 m, the shower front can be approximated by a plane wave. Thus, to determine the arrival direction of the cosmic ray a planar wavefront is fit to the arrival times of the pulses.

This method assumes that essentially all antennas are on a single plane, which certainly holds for all LOFAR stations as the ground was flattened during construction. Given a vector of arrival times t , and the vectors x and y for the coordinates of the antennas, the best fitting solution for a plane wave:

$$ct = Ax + By + C, \quad (13)$$

can be found using a standard least squares approach. From A and B the Cartesian directions ϕ , θ can be extracted as:

$$A = \sin(\theta) \sin(\phi), \quad (14)$$

$$B = \sin(\theta) \cos(\phi). \quad (15)$$

The plane wave fit itself is done in several iterations. After a fit is performed the residual delays are investigated and those antennas that have residual delays larger than 3 times the standard deviation on the residual delays, are removed from the set and the data are refitted. The fit is terminated when there are less than four antennas left in the set or if no further antennas need to be removed. For this best direction all residual delays, including those of removed antennas, are calculated again and used for quality cuts later.

There are several quality criteria in the pipeline related to the plane wave fit. If the fit fails, a station is not considered further. In addition, a cut is made on the remaining average residual delays with respect to the expectation of the best fit. This cut can be derived from the distribution of all occurring plane wave residual delays, as shown in the top panel of Fig. 17. The first peak with events of an average residual delay of less than 10 ns corresponds to excellent events, in which a clear cosmic-ray pulse can be identified in all antennas. The largest peak corresponds to all those events in which random noise fluctuations are identified as a pulse. This can be illustrated by a small Monte Carlo simulation. A random sample is picked from the range of the search window and its residual to the middle of the search window

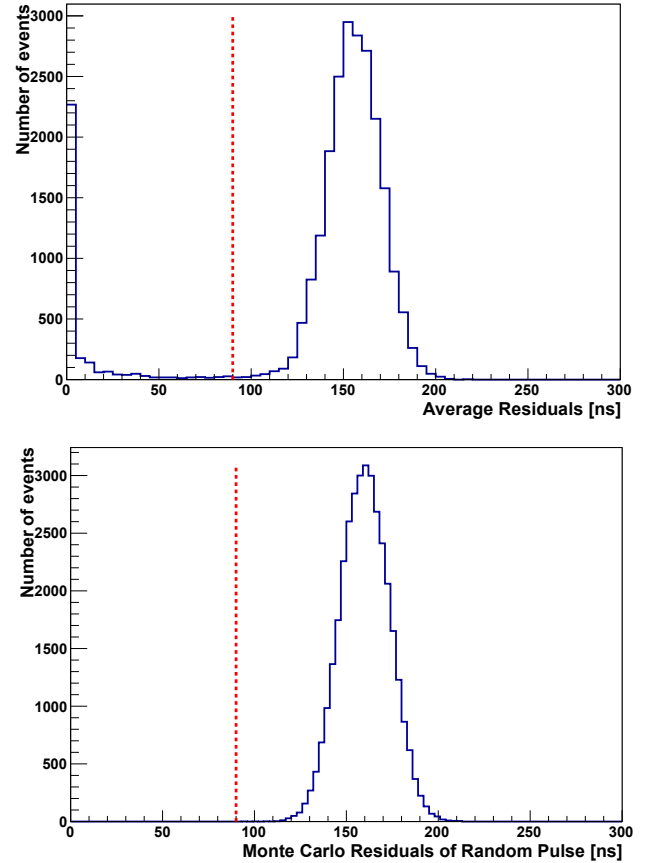


Fig. 17. Average residual delays derived from a plane-fit to data (*top*) and from random samples in the search window with respect to a horizontal shower front (*bottom*). The vertical line indicates the cut value derived from the simulated distribution, which is applied to the data.

(corresponding to a vertical shower) is calculated. This results in the distribution in the bottom panel of Fig. 17. The peak in the distribution obtained from data and the Monte Carlo distribution are centered around the same value and can therefore be identified with each other. Second order effects, being the directions of the air showers and non-infinite sampling, can influence the shape of the peak. The longer tail of the first peak (up to about 50 ns) corresponds to events that have some antennas with correctly identified pulses and varying numbers of outliers, i.e. antennas where a random pulse is identified.

Therefore one can safely choose the value 90 ns as a first cut for good cosmic-ray events. Further cuts for higher quality events or stations can be applied in later analyses.

The plane wave fit results now also allow for a justification of the choice of the Hilbert envelope as the method for pulse timing, as opposed to cross correlation. Figure 18 shows the ratio of the number of antennas in which a pulse has been identified by either method with respect to the remaining residuals on a test-set of randomly chosen events that contain a cosmic-ray signal. The distribution clearly shows that the Hilbert envelope finds significantly more signals in the first bin, i.e. in the correct bin with small residuals. In general, cross correlating is expected to be better for pulses with lower signal-to-noise ratio. For pulses with a high signal-to-noise, however, the Hilbert transform performs more accurately. When using the Hilbert envelope, the position of the maximum is only determined by the recorded individual pulse, whereas the peak of the cross correlation is determined by the degree to which two signals correlate. This degree of correlation can be influenced by correlations in the noise (for instance

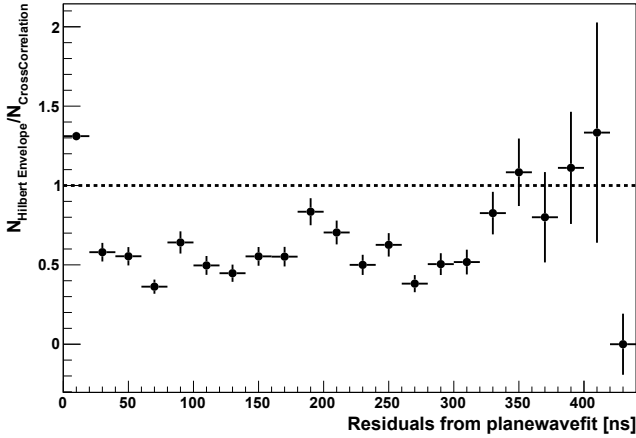


Fig. 18. Difference in reconstruction between Hilbert Envelope and Cross Correlation. The different quality of the reconstruction is illustrated by plotting the fraction of the numbers of antennas N , identified by each method, with respect to the residual that was found in the plane wave reconstruction. For values above one the Hilbert Envelope identified more antennas, which is the case for the desired correctly identified signals, which can be found below 20 ns.

residual RFI) or lacking similarity of the pulse shape between antennas, thereby making the cross-correlation less accurate for timing of pulses with a high signal-to-noise ratio.

3.4. Coordinate transformation

After the antenna pattern unfolding cycle completes with a successful direction fit for a given station, the electric field components in on-sky polarizations $E_\theta(t)$, $E_\phi(t)$, and the shower arrival direction \mathbf{n} are known. However, to compare measured data to air-shower simulations the three-dimensional electric field at ground level

$$\mathbf{E}(t) = E_x(t)\hat{\mathbf{e}}_x + E_y(t)\hat{\mathbf{e}}_y + E_z(t)\hat{\mathbf{e}}_z \quad (16)$$

is needed, where $\hat{\mathbf{e}}_x$, $\hat{\mathbf{e}}_y$ and $\hat{\mathbf{e}}_z$ form the right handed coordinate system pointing east, north and up, respectively. This geometry can also be seen in Fig. 13.

Assuming the signal has no electric field component in the propagation direction $-\hat{\mathbf{e}}_n$, this follows from a simple rotation $(E_x, E_y, E_z)^T = \mathbf{R} \cdot (E_\theta, E_\phi, 0)^T$, with the rotation matrix

$$\mathbf{R} = \begin{pmatrix} \cos \theta \cos \phi & -\sin \phi & \sin \theta \cos \phi \\ \cos \theta \sin \phi & \cos \phi & \sin \theta \sin \phi \\ -\sin \theta & 0 & \cos \theta \end{pmatrix}. \quad (17)$$

Note that this assumption is only an approximation, since the signal is measured in the near field of the shower and the source is moving. However, these are second order effects and the $E_\theta(t)$ and $E_\phi(t)$ components are expected to dominate over the $E_n(t)$ component (Huege 2013). Moreover, since LOFAR uses a dual polarization set-up it is not possible to extract the $E_n(t)$ component of a linearly polarized signal.

The pipeline concludes by storing pulse parameters for each antenna in the projected directions.

3.5. Extracting pulse parameters

In addition to the shower arrival direction, obtained from pulse timing, two more parameters are extracted: for each antenna the peak amplitude and integrated power of the pulse are calculated.

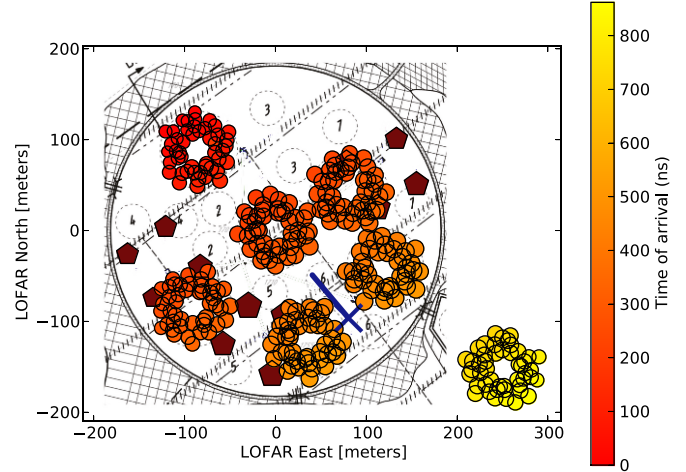


Fig. 19. Footprint of an air shower measured with LOFAR. The signal strength (peak amplitude of the radio signal) is encoded logarithmically in the size of the marker and the color shows the time of arrival. The pentagons represent the positions of the particle detectors, their size is proportional to the number of recorded particles. The reconstructed shower axis is indicated by the blue cross for the core position and the line for the projected arrival direction.

Without multiplicative unit conversion factors, ignored for current lack of absolute calibration, the integrated pulse power is defined through the integration of the instantaneous Poynting vector and the electric field strength as:

$$P = \sum_k P_k \propto \sum_k \int_{\Delta t} |E_k(t)|^2 dt, \quad (18)$$

where $k = (x, y, z)$ are the polarization components of the electric field and Δt is taken as a symmetric window around the pulse maximum.

This is calculated in discrete sampling x_i as

$$P_k = \frac{1}{f} \left(\sum_{\text{signal}} |x_i|^2 - \frac{N_{\text{signal}}}{N_{\text{noise}}} \sum_{\text{noise}} |x_i|^2 \right), \quad (19)$$

where $f = 200$ MHz is the sampling frequency and N_{signal} and N_{noise} are the number of samples in the signal and noise windows respectively. The noise window consists of the full 327 680 ns block excluding the pulse window.

3.6. A measured air shower

The result of the reconstruction pipeline is a full three-dimensional electric field vector per antenna position as a function of time. There are various ways in which this result can be visualized. The shower footprint, Fig. 19, shows the signal strength (peak amplitude of the radio signal) at the measured antenna locations as well as the time of arrival. Here, one can see that both the radio signal strength and the arrival times are consistent with the air-shower direction and core position as determined by the scintillator array. Both effects are distinctive properties of radio emission from air showers and are not produced by RFI.

Another common way to visualize the result is in the form of the lateral distribution, shown in Fig. 20. Here the radio signal strength, in all three polarization components, is shown as a function of projected distance to the shower axis. This projection retains the spatial distribution of the antennas (i.e. stations

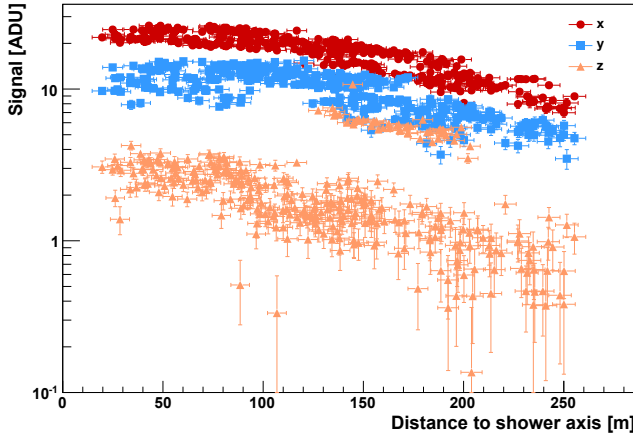


Fig. 20. Distribution of radio signals (peak amplitude in arbitrary units) with respect to the distance from the shower axis as reconstructed from the scintillator data. Shown are the three components of the reconstructed electric field.

can be seen as groups), but azimuthal symmetry in the shower plane is assumed. This rather complicated looking distribution can be explained using detailed models of the radio emission, which also include non-rotational symmetrical effects. Further details of event by event characteristics will be reported in forthcoming publications.

4. Properties of reconstructed air showers

In order to verify the data quality and the method of reconstruction a short overview of the first data taken with LOFAR is given. The data set used here (June 2011 until April 2013) contains 3341 recorded triggers, of which 1597 pass the strict quality cut for a good data reconstruction of the particle measurement. Of all triggers, 405 events contain signals of cosmic rays as identified by the pipeline, with a threshold energy of 5×10^{15} eV.

4.1. Triggers from the array of particle detectors

On the reconstruction of air showers from the particle data quality cuts are applied. The reconstruction is considered reliable, when the reconstructed shower core is contained within the array, the shower is not too horizontal ($\theta < 50^\circ$) and the reconstructed Molière radius⁸ falls in the range of 20–100 m. After cuts, the lowest energy of a shower that triggered a read-out of the LOFAR buffers is 1.8×10^{15} eV and the highest is 1.9×10^{18} eV. The LORA scintillator array becomes fully efficient above 2×10^{16} eV.

All triggers sent by the scintillator array follow a nearly uniform distribution in azimuth and a $\sin(\theta)\cos(\theta)$ -distribution in zenith angle as it is expected from the geometry for a horizontal array with flat detectors.

The number of events with a detectable radio signal increases with the number of triggered particle detectors, as can be seen in Fig. 21, where the fraction of triggered events, with and without a detected radio signal, is plotted against the number of particle detectors per event. The fraction is clearly increasing with the number of triggered detectors, as shown by a fitted straight line. According to this fit, at a threshold of 13 detectors about 10% of the events contain a cosmic-ray signal.

⁸ Characteristic transverse size of an air shower.

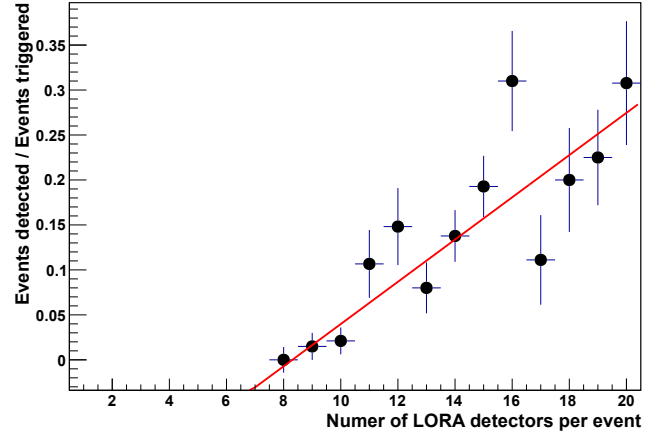


Fig. 21. Fraction of air showers with a detectable radio signal over the number of air showers triggered with a scintillator signal is plotted against the number of particle detectors above threshold in an event. The red straight line is a fit to the data.

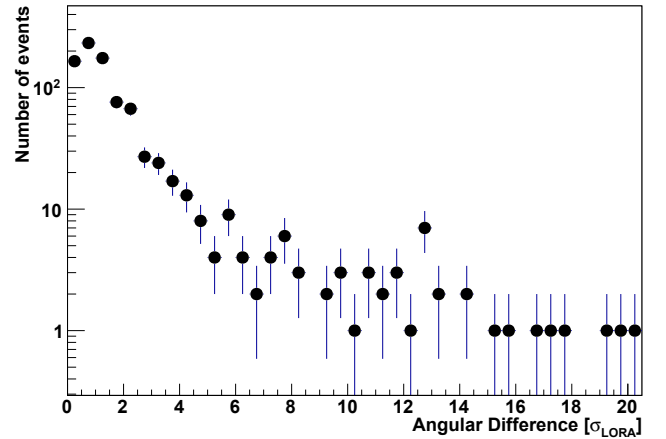


Fig. 22. Angular difference between the shower axis reconstructed from the particle data and the direction estimate from the radio signal. To make the events comparable, the difference is scaled with the uncertainty of the individual reconstruction σ_{LORA} .

4.2. Event rates and sensitivity

For a first estimate all reconstructed triggers are considered valid events which show radio pulses coming from a direction that agrees to 10° angular distance with the direction that was reconstructed from the arrival times measured with the particle detectors. This choice is based on the results shown in Fig. 22. This figure shows the angular difference between the two reconstructed axes for all events. A steep fall-off in number of events with an increasing angular difference can be seen. Any event that deviates more than $10\sigma_{\text{LORA}}$ certainly lies outside the correct distribution. The shower axis is on average reconstructed with an uncertainty $\sigma_{\text{LORA}} \sim 1^\circ$ from the data of the particle detectors. Thus, a quality cut of 10° is chosen.

Figure 23 shows all 405 cosmic-ray events successfully detected with the LBAs as distributed on the local sky. Visible is a clear north-south asymmetry, where 276 events arrive from the northern hemisphere. This corresponds to a probability $p = 0.68 \pm 0.02$ for a detected event to arrive from the north. As the magnetic field at LOFAR is parallel to the north-south axis this is expected, if the main contribution to the signal is of geomagnetic origin (Falcke et al. 2005; Ardouin et al. 2009).

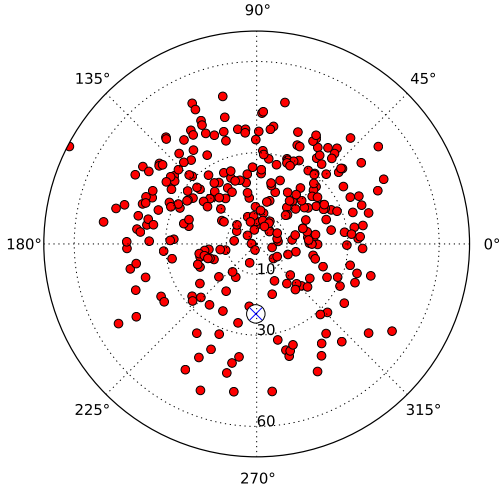


Fig. 23. Arrival directions of the cosmic-ray events detected with LOFAR from June 2011 until April 2013. East is 0° and north corresponds to 90°. Also indicated (cross) is the direction of the magnetic field at LOFAR.

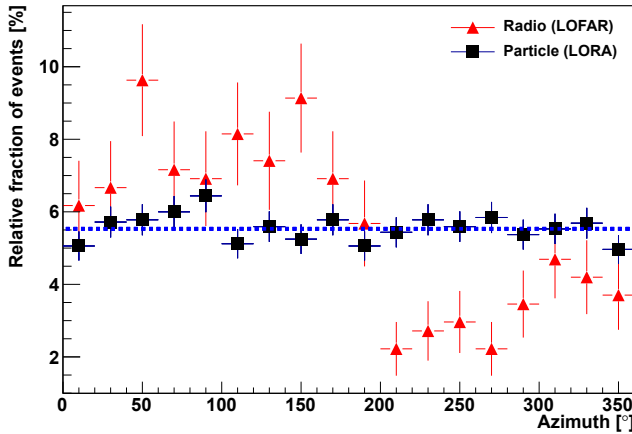


Fig. 24. Binned distribution of the azimuth angles of all events measured with the particle detectors (black squares) and those in coincidence of particle detectors and radio antennas (red triangles). The best fit of a straight line to the particle data is also shown. The fit has a $\chi^2/\text{nd.o.f.} = 0.9$.

The effect is also illustrated in Fig. 24, which shows the fraction of detected air showers as a function of azimuth angle for the events with radio signal, as well as for all LORA triggers sent. While the events registered with the LORA detectors are uniformly distributed in azimuth, the radio events show a clear deficit from the south. Due to the orientation of the LOFAR antennas and thereby the reduced sensitivity for purely east-west polarized signals, events arriving directly from the north are not necessarily preferred, as their signal is expected to be mainly polarized in the east-west direction (Huege 2013). The detection efficiency as a function of direction follows from a deconvolution of the expected emission strength with the antenna pattern and will not be discussed in detail here.

The energies of the air showers with a detectable radio signal are shown in Fig. 25. The depicted energy is the one reconstructed from the corresponding particle data. This reconstruction has an overall systematic uncertainty of 27% and varying event by event uncertainties (Thoudam et al., in prep.). One clearly sees that below $\sim 10^{17}$ eV the detection of air showers through their radio signal is not fully efficient, as the strength

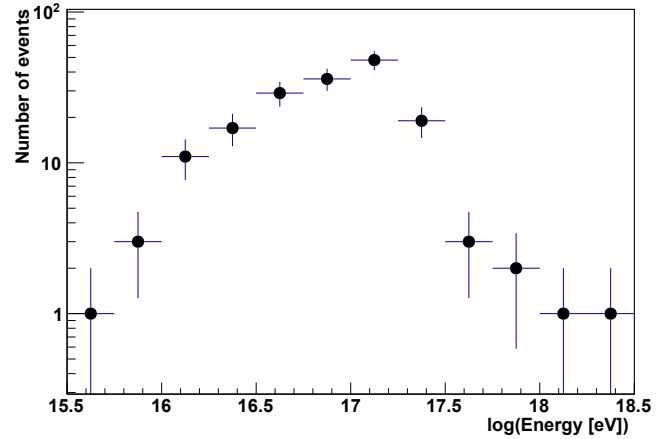


Fig. 25. Distribution of the energies of the cosmic rays which had a measurable radio signal in the LOFAR data. The depicted energy is the one reconstructed from the corresponding particle data. The quality cuts, as described in Sect. 4.1, are applied.

of the radio signal scales with the energy of the shower. Higher energies in this distribution are constrained by the steeply falling cosmic-ray energy spectrum and limited size of the detector array, which leads to limited event statistics at the highest energies. There are significant hints that showers of higher energies have been measured with LOFAR (especially when including the stations outside the Superterp), but these events are not well enough constrained by the data from the particle detectors in order to have a reference energy of the necessary accuracy. After a calibration of the energy of the radio measurements, those events will be used in a radio-stand-alone reconstruction.

5. Conclusions

At LOFAR cosmic-ray induced air showers are regularly measured with an array of particle detectors, LORA, and a large array of radio antennas. The cosmic-ray pipeline is routinely finding their distinctive radio signatures in the measurements and a full three-dimensional electric field vector is reconstructed for every antenna position.

A large dataset has been gathered with hundreds of identified cosmic-ray events in data from the LBAs. With up to a thousand antennas per events, these are the first highly detailed measurement of the radio signal of air showers. These measurements will be used for a detailed characterization of the shower shape and will be the benchmark data for comparison with models of radio emission in air showers.

Acknowledgements. The LOFAR cosmic ray key science project very much acknowledges the scientific and technical support from ASTRON, especially in constructing the LORA particle detectors. We thank the KASCADE Collaboration for providing the scintillator detectors. Furthermore, we acknowledge financial support from the Netherlands Research School for Astronomy (NOVA), the Samenwerkingsverband Noord-Nederland (SNN) and the Foundation for Fundamental Research on Matter (FOM) as well as support from the Netherlands Organization for Scientific Research (NWO), VENI grant 639-041-130. We acknowledge funding from an Advanced Grant of the European Research Council under the European Union's Seventh Framework Program (FP/2007-2013) / ERC Grant Agreement n. 227610. LOFAR, the Low Frequency Array designed and constructed by ASTRON, has facilities in several countries, that are owned by various parties (each with their own funding sources), and that are collectively operated by the International LOFAR Telescope (ILT) foundation under a joint scientific policy. The authors would like to thank both the internal and external referees for carefully reading the manuscript. Chiara Ferrari and Giulia Macario acknowledge financial support by the "Agence Nationale de la Recherche" through grant ANR-09-JCJC-0001-01.

References

- Alexov, A., Schellart, P., ter Veen, S., et al. 2012, in *Astronomical Data Analysis Software and Systems XXI*, eds. P. Ballester, D. Egret, & N. P. F. Lorente, ASP Conf. Ser., 461, 283
- Allan, H., & Jones, J. 1966, *Nature*, 212, 129
- Alvarez-Muñiz, J., Carvalho, W. R., & Zas, E. 2012, *Astropart. Phys.*, 35, 325
- Ardouin, D., Belletoile, A., Charrier, D., et al. 2005, *Int. J. Mod. Phys. A*, 20, 6869
- Ardouin, D., Belletoile, A., Berat, C., et al. 2009, *Astropart. Phys.*, 31, 192
- Falcke, H., Apel, W. D., Badea, A. F., et al. 2005, *Nature*, 435, 313
- Falcke, H., et al. 2008, in *ICRC Merida*, Vol. Rapporteur, 30th Int. Cosmic Ray Conf.
- Hamaker, J. P., Bregman, J. D., & Sault, R. J. 1996, *A&AS*, 117, 137
- Huege, T. 2013, *ARENA 2012*, AIP Conf. Proc. 1535, 121
- Huege, T., Apel, W. D., Arteaga, J. C., et al. 2012, *Nucl. Instr. Meth. Phys. Res. A*, 662, 72
- Huege, T., Ludwig, M., & James, C. W. 2013, *ARENA 2012*, AIP Conf. Proc. 1535, 128
- Jelley, J., Fruin, J., Porter, N., et al. 1965, *Nature*, 205, 327
- Jones, R. C. 1941, *Journal of the Optical Society of America*, 31, 488
- Kolundzija, B. 2011, in *Proc. of the 5th European Conference on Antennas and Propagation (EUCAP)*
- Marin, V., & Revenu, B. 2012, *Astropart. Phys.*, 35, 733
- Nehls, S., Hakenjos, A., Arts, M. J., et al. 2008, *Nucl. Inst. Meth. A*, 589, 350
- Offringa, A. R., de Bruyn, A. G., Zaroubi, S., et al. 2013, *A&A*, 549, A11
- van Cappellen, W., Rüter, M., & Kant, G. 2007, *LOFAR-ASTRON-ADD-009*
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, *A&A*, 556, A2
- Werner, K., de Vries, K. D., & Scholten, O. 2012, *Astropart. Phys.*, 37, 5
- ¹³ Max Planck Institute for Astrophysics, Karl Schwarzschild Str. 1, 85741 Garching, Germany
- ¹⁴ International Centre for Radio Astronomy Research – Curtin University, GPO Box U1987, WA 6845 Perth, Australia
- ¹⁵ STFC Rutherford Appleton Laboratory, Harwell Science and Innovation Campus, Didcot OX11 0QX, UK
- ¹⁶ Institute for Astronomy, University of Edinburgh, Royal Observatory of Edinburgh, Blackford Hill, Edinburgh EH9 3HJ, UK
- ¹⁷ LESIA, UMR CNRS 8109, Observatoire de Paris, 92195 Meudon, France
- ¹⁸ Argelander-Institut für Astronomie, University of Bonn, Auf dem Hügel 71, 53121 Bonn, Germany
- ¹⁹ Leibniz-Institut für Astrophysik Potsdam (AIP), An der Sternwarte 16, 14482 Potsdam, Germany
- ²⁰ Thüringer Landessternwarte, Sternwarte 5, 07778 Tautenburg, Germany
- ²¹ Astronomisches Institut der Ruhr-Universität Bochum, Universitätsstrasse 150, 44780 Bochum, Germany
- ²² Laboratoire de Physique et Chimie de l'Environnement et de l'Espace, LPC2E UMR 7328 CNRS, 45071 Orléans Cedex 02, France
- ²³ SRON Netherlands Institute for Space Research, PO Box 800, 9700 AV Groningen, The Netherlands
- ²⁴ Center for Information Technology (CIT), University of Groningen, The Netherlands
- ²⁵ Centre de Recherche Astrophysique de Lyon, Observatoire de Lyon, 9 av Charles André, 69561 Saint Genis Laval Cedex, France
- ²⁶ ARC Centre of Excellence for All-sky astrophysics (CAASTRO), Sydney Institute of Astronomy, University of Sydney, Australia
- ²⁷ University of Hamburg, Gojenbergsweg 112, 21029 Hamburg, Germany
- ²⁸ Astro Space Center of the Lebedev Physical Institute, Profsoyuznaya str. 84/32, 117997 Moscow, Russia
- ²⁹ Centre for Radio Astronomy Techniques & Technologies (RATT), Department of Physics and Electronics, Rhodes University, PO Box 94, 6140 Grahamstown, South Africa
- ³⁰ SKA South Africa, 3rd Floor, The Park, Park Road, 7405 Pinelands, South Africa
- ³¹ Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA
- ³² Laboratoire Lagrange, UMR7293, Université de Nice Sophia-Antipolis, CNRS, Observatoire de la Côte d'Azur, 06300 Nice, France
- ³³ Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA
- ³⁴ Sodankylä Geophysical Observatory, University of Oulu, Tähteläntie 62, 99600 Sodankylä, Finland
- ³⁵ Netherlands eScience Center, Science Park 140, 1098 XG Amsterdam, The Netherlands
- ³⁶ Fakultät für Physik, Universität Bielefeld, Postfach 100131, 33501 Bielefeld, Germany