

LOFAR: The LOw-Frequency ARray

M. P. van Haarlem¹, M. W. Wise^{*1,2}, A. W. Gunst¹, G. Heald¹, J. P. McKean¹, J. W. T. Hessels^{1,2}, A. G. de Bruyn^{1,3}, R. Nijboer¹, J. Swinbank², R. Fallows¹, M. Brentjens¹, A. Nelles⁵, R. Beck⁸, H. Falcke^{5,1}, R. Fender⁹, J. Hörandel⁵, L. V. E. Koopmans³, G. Mann¹⁷, G. Miley⁴, H. Röttgering⁴, B. W. Stappers⁶, R. A. M. J. Wijers², S. Zaroubi³, M. van den Akker⁵, A. Alexov², J. Anderson⁸, K. Anderson², A. van Ardenne^{1,29}, M. Arts¹, A. Asgekar¹, I. M. Avruch^{1,3}, F. Batejat¹¹, L. Bähren², M. E. Bell⁹, M. R. Bell¹⁰, I. van Bemmel¹, P. Bennema¹, M. J. Bentum¹, G. Bernardi³, P. Best¹⁴, L. Bîrzan⁴, A. Bonafede²¹, A.-J. Boonstra¹, R. Braun²⁷, J. Bregman¹, F. Breitling¹⁷, R. H. van de Brink¹, J. Broderick⁹, P. C. Broekema¹, W. N. Brouw^{1,3}, M. Brügger²⁰, H. R. Butcher^{1,26}, W. van Cappellen¹, B. Ciardi¹⁰, T. Coenen², J. Conway¹¹, A. Coolen¹, A. Corstanje⁵, S. Damstra¹, O. Davies¹³, A. T. Deller¹, R.-J. Dettmar¹⁹, G. van Diepen¹, K. Dijkstra²³, P. Donker¹, A. Doorduyn¹, J. Dromer¹, M. Drost¹, A. van Duin¹, J. Eislöffel¹⁸, J. van Enst¹, C. Ferrari³⁰, W. Frieswijk¹, H. Gankema³, M. A. Garrett^{1,4}, F. de Gasperin¹⁰, M. Gerbers¹, E. de Geus¹, J.-M. Grießmeier^{22,1}, T. Grit¹, P. Gruppen¹, J. P. Hamaker¹, T. Hassall⁶, M. Hoeft¹⁸, H. A. Holties¹, A. Horneffer^{8,5}, A. van der Horst², A. van Houwelingen¹, A. Huijgen¹, M. Iacobelli⁴, H. Intema^{4,28}, N. Jackson⁶, V. Jelic¹, A. de Jong¹, E. Juette¹⁹, D. Kant¹, A. Karastergiou⁷, A. Koers¹, H. Kollen¹, V. I. Kondratiev¹, E. Kooistra¹, Y. Koopman¹, A. Koster¹, M. Kuniyoshi⁸, M. Kramer^{8,6}, G. Kuper¹, P. Lambropoulos¹, C. Law^{24,2}, J. van Leeuwen^{1,2}, J. Lemaître¹, M. Loose¹, P. Maat¹, G. Macario³⁰, S. Markoff², J. Masters^{28,2}, R. A. McFadden¹, D. McKay-Bukowski¹³, H. Meijering¹, H. Meulman¹, M. Mevius³, E. Middelberg¹⁹, R. Millenaar¹, J. C. A. Miller-Jones^{12,2}, R. N. Mohan⁴, J. D. Mol¹, J. Morawietz¹, R. Morganti^{1,3}, D. D. Mulcahy⁸, E. Mulder¹, H. Munk¹, L. Nieuwenhuis¹, R. van Nieuwpoort^{1,32}, J. E. Noordam¹, M. Norden¹, A. Noutsos⁸, A. R. Offringa³, H. Olofsson¹¹, A. Omar¹, E. Orrù^{5,1}, R. Overeem¹, H. Paas²³, M. Pandey-Pommier^{4,25}, V. N. Pandey³, R. Pizzo¹, A. Polatidis¹, D. Rafferty⁴, S. Rawlings⁷, W. Reich⁸, J.-P. de Reijer¹, J. Reitsma¹, G. A. Renting¹, P. Riemers¹, E. Rol², J. W. Romein¹, J. Roosjen¹, M. Ruiter¹, A. Scaife⁹, K. van der Schaaf¹, B. Scheers^{2,33}, P. Schellart⁵, A. Schoenmakers¹, G. Schoonderbeek¹, M. Serylak^{31,22}, A. Shulevski³, J. Sluman¹, O. Smirnov¹, C. Sobey⁸, H. Spreeuw², M. Steinmetz¹⁷, C. G. M. Sterks²³, H.-J. Stiepel¹, K. Stuurwold¹, M. Tagger²², Y. Tang¹, C. Tasse¹⁵, I. Thomas¹, S. Thoudam⁵, M. C. Toribio¹, B. van der Tol⁴, O. Ussov⁴, M. van Veelen¹, A.-J. van der Veen¹, S. ter Veen⁵, J. P. W. Verbiest⁸, R. Vermeulen¹, N. Vermaas¹, C. Vocks¹⁷, C. Vogt¹, M. de Vos¹, E. van der Wal¹, R. van Weeren^{4,1}, H. Weggemans¹, P. Weltevrede⁶, S. White¹⁰, S. J. Wijnholds¹, T. Wilhelmsson¹⁰, O. Wucknitz¹⁶, S. Yatawatta³, P. Zarka¹⁵, A. Zensus⁸, and J. van Zwieten¹

(Affiliations can be found after the references)

Received 7 December 2012 / Accepted 9 May 2013

ABSTRACT

LOFAR, the LOw-Frequency ARray, is a new-generation radio interferometer constructed in the north of the Netherlands and across Europe. Utilizing a novel phased-array design, LOFAR covers the largely unexplored low-frequency range from 10–240 MHz and provides a number of unique observing capabilities. Spreading out from a core located near the village of Exloo in the northeast of the Netherlands, a total of 40 LOFAR stations are nearing completion. A further five stations have been deployed throughout Germany, and one station has been built in each of France, Sweden, and the UK. Digital beam-forming techniques make the LOFAR system agile and allow for rapid repointing of the telescope as well as the potential for multiple simultaneous observations. With its dense core array and long interferometric baselines, LOFAR achieves unparalleled sensitivity and angular resolution in the low-frequency radio regime. The LOFAR facilities are jointly operated by the International LOFAR Telescope (ILT) foundation, as an observatory open to the global astronomical community. LOFAR is one of the first radio observatories to feature automated processing pipelines to deliver fully calibrated science products to its user community. LOFAR's new capabilities, techniques and modus operandi make it an important pathfinder for the Square Kilometre Array (SKA). We give an overview of the LOFAR instrument, its major hardware and software components, and the core science objectives that have driven its design. In addition, we present a selection of new results from the commissioning phase of this new radio observatory.

Key words. instrumentation: interferometers – radio continuum: general – radio lines: general – dark ages, reionization, first stars – telescopes

1. Introduction

During the last half century, our knowledge of the Universe has been revolutionized by the opening of observable windows outside the narrow visible region of the electromagnetic spectrum. Radio waves, infrared, ultraviolet, X-rays, and most recently γ -rays have all provided new, exciting, and completely

* Corresponding author e-mail: wise@astron.nl

unexpected information about the nature and history of the Universe, as well as revealing a cosmic zoo of strange and exotic objects. One spectral window that as yet remains relatively unexplored is the low-frequency radio domain below a few hundred MHz, representing the lowest frequency extreme of the accessible spectrum.

Since the discovery of radio emission from the Milky Way (Jansky 1933), now 80 years ago, radio astronomy has made a continuous stream of fundamental contributions to astronomy. Following the first large-sky surveys in Cambridge, yielding the 3C and 4C catalogs (Edge et al. 1959; Bennett 1962; Pilkington & Scott 1965; Gower et al. 1967) containing hundreds to thousands of radio sources, radio astronomy has blossomed. Crucial events in those early years were the identifications of the newly discovered radio sources in the optical waveband. Radio astrometric techniques, made possible through both interferometric and lunar occultation techniques, led to the systematic classification of many types of radio sources: Galactic supernova remnants (such as the Crab Nebula and Cassiopeia A), normal galaxies (M 31), powerful radio galaxies (Cygnus A), and quasars (3C 48 and 3C 273).

During this same time period, our understanding of the physical processes responsible for the radio emission also progressed rapidly. The discovery of powerful very low-frequency coherent cyclotron radio emission from Jupiter (Burke & Franklin 1955) and the nature of radio galaxies and quasars in the late 1950s was rapidly followed by such fundamental discoveries as the Cosmic Microwave Background (Penzias & Wilson 1965), pulsars (Bell & Hewish 1967), and apparent superluminal motion in compact extragalactic radio sources by the 1970s (Whitney et al. 1971).

Although the first two decades of radio astronomy were dominated by observations below a few hundred MHz, the prediction and subsequent detection of the 21 cm line of hydrogen at 1420 MHz (van de Hulst 1945; Ewen & Purcell 1951), as well as the quest for higher angular resolution, shifted attention to higher frequencies. This shift toward higher frequencies was also driven in part by developments in receiver technology, interferometry, aperture synthesis, continental and intercontinental very long baseline interferometry (VLBI). Between 1970 and 2000, discoveries in radio astronomy were indeed dominated by the higher frequencies using aperture synthesis arrays in Cambridge, Westerbork, the VLA, MERLIN, ATCA and the GMRT in India as well as large monolithic dishes at Parkes, Effelsberg, Arecibo, Green Bank, Jodrell Bank, and Nançay.

By the mid 1980s to early 1990s, however, several factors combined to cause a renewed interest in low-frequency radio astronomy. Scientifically, the realization that many sources have inverted radio spectra due to synchrotron self-absorption or free-free absorption as well as the detection of (ultra-) steep spectra in pulsars and high redshift radio galaxies highlighted the need for data at lower frequencies. Further impetus for low-frequency radio data came from early results from Clark Lake (Erickson & Fisher 1974; Kassim 1988), the Cambridge sky surveys at 151 MHz, and the 74 MHz receiver system at the VLA (Kassim et al. 1993, 2007). In this same period, a number of arrays were constructed around the world to explore the sky at frequencies well below 1 GHz (see Table 2 in Stappers et al. 2011, and references therein).

Amidst all this progress, radio astronomers nonetheless began to look toward the future and one ambition that emerged was the proposed construction of an instrument capable of detecting neutral hydrogen at cosmological distances. A first order analysis, suggested that a telescope with a collecting area of about one square kilometer was required (Wilkinson 1991). The project,

later to be known as the Square Kilometre Array (SKA; Ekers 2012), was adopted by the community globally, and around the world various institutes began to consider potential technologies that might furnish such a huge collecting area at an affordable cost.

At ASTRON in the Netherlands, the concept of Phased or Aperture Arrays was proposed as a possible solution to this problem (van Ardenne et al. 2000), and in the slip-stream of those early developments, the idea of constructing a large low frequency dipole array also emerged (Bregman 2000; Miley 2010). The concept of a large, low frequency array had arisen previously (Perley & Erickson 1984), and been revisited several times over the years (e.g., see Kassim & Erickson 1998). These plans were greatly aided by the revolution then taking place in other fields, in particular major advances in digital electronics, fibre-based data networks, high performance computing and storage capacity, made it possible to consider the construction of a transformational radio telescope design that would operate between 10–200 MHz with unprecedented sensitivity and angular resolution. This telescope would be a major scientific instrument in its own right, bridging the gap to the even more ambitious SKA (Miley 2010). This international initiative became known as the LOw Frequency ARray or LOFAR (Bregman 2000; Kassim et al. 2003; Butcher 2004).

As originally envisioned, LOFAR was intended to surpass the power of previous interferometers in its frequency range by 2–3 orders of magnitude providing a square kilometer of collecting area at 15 MHz, millijansky sensitivity, and arcsecond resolution (Kassim et al. 2003). Due to funding constraints, the original collaboration split in 2004 resulting in three currently ongoing low-frequency array development projects: the European LOFAR project described here; the US-led, Long Wavelength Array (LWA; Ellingson et al. 2009, 2013); and the international Murchison Widefield Array (MWA) collaboration (Lonsdale et al. 2009; Tingay et al. 2013a).

The scientific motivation for the construction of these arrays has become very broad. Among the most interesting application of the low-frequency arrays is the detection of highly redshifted 21 cm line emission from the epoch of reionization (HI redshifts $z = 6$ to 20) and a phase called cosmic dawn (HI redshifts from $z = 50$ to $z = 20$; see Zaroubi et al. 2012). However, the science case for LOFAR has continued to broaden since 2000 to include the detection of nanosecond radio flashes from ultra-high energy cosmic rays (CRs; Falcke et al. 2005), deep surveys of the sky in search for high redshift radio sources (Röttgering et al. 2011), surveys of pulsars and cosmic radio transients (Stappers et al. 2011), or the radio detection of exoplanets (Zarka 2011). The great sensitivity and broad low-frequency bandwidth may also prove crucial for studies of cosmic magnetic fields (see Sect. 13.6).

In this paper, we present an overview and reference description of the LOFAR telescope. We aim to give the potential LOFAR user a general working knowledge of the main components and capabilities of the system. More detailed descriptions of individual components or subsystems will be published elsewhere. The paper continues in Sect. 2 with a general overview of the system and descriptions of the overall layout of the array and the antenna fields themselves in Sect. 3 and Sect. 4. The LOFAR processing hardware and data-flow through the system are summarized in Sect. 5 and Sect. 6. An overview of the software infrastructure including a description of LOFAR's primary observational modes and science pipelines is given in Sect. 9, Sect. 10, and Sect. 11, respectively. In Sect. 12, an initial set of performance metrics are presented. LOFAR's key science drivers

are reviewed in Sect. 13 along with examples of recent results that demonstrate the potential of this new facility. A discussion of ongoing construction plans and possible future enhancements to the system are given in Sect. 14. Lastly, Sect. 15 offers some brief conclusions.

2. System overview

LOFAR, the LOW-Frequency ARray, is a new and innovative radio telescope designed and constructed by ASTRON to open the lowest frequency radio regime to a broad range of astrophysical studies. Capable of operating in the frequency range from 10–240 MHz (corresponding to wavelengths of 30–1.2 m), LOFAR consists of an interferometric array of dipole antenna stations distributed throughout the Netherlands and Europe. These stations have no moving parts and, due to the effectively all-sky coverage of the component dipoles, give LOFAR a large field-of-view (FoV). At station level, the signals from individual dipoles are combined digitally into a phased array. Electronic beam-forming techniques make the system agile and allow for rapid repointing of the telescope as well as the simultaneous observation of multiple, independent areas of the sky. Brief descriptions of the LOFAR system have been presented previously in Bregman (2000); Falcke (2006); Falcke et al. (2007); de Vos et al. (2009).

In the Netherlands, a total of 40 LOFAR stations are being deployed with an additional 8 international stations currently built throughout Europe. The densely sampled, 2-km-wide, core hosts 24 stations and is located ~30 km from ASTRON's headquarters in Dwingeloo. The datastreams from all LOFAR stations are sent via a high-speed fiber network infrastructure to a central processing (CEP) facility located in Groningen in the north of the Netherlands. At the computing center of the University of Groningen, data from all stations are aligned, combined, and further processed using a flexible IBM Blue Gene/P supercomputer offering about 28 Tflop/s of processing power. In the Blue Gene/P, a variety of processing operations are possible including correlation for standard interferometric imaging, tied-array beam-forming for high time resolution observations, and even real-time triggering on incoming station data-streams. Combinations of these operations can also be run in parallel.

After processing in the Blue Gene/P, raw data products are written to a storage cluster for additional post-processing. This cluster currently hosts 2 Pbyte of working storage. Once on the storage cluster, a variety of reduction pipelines are then used to further process the data into the relevant scientific data products depending on the specific type of observation. In the case of the standard imaging pipeline, subsequent processing includes flagging of the data for the presence of radio frequency interference, averaging, calibration, and creation of the final images. This and other science-specific pipelines run on a dedicated compute cluster with a total processing power of approximately 10 Tflop/s. After processing, the final scientific data products are transferred to the LOFAR long-term archive (LTA) for cataloging and distribution to the community.

In order to fully exploit this new wavelength regime with unprecedented resolution and sensitivity, LOFAR must meet several non-trivial technical challenges. For example, the meter-wave wavelength regime is prone to high levels of man-made interference. Excising this interference requires high spectral and time resolution, and high dynamic range analog to digital (A/D) converters. Furthermore, for the typical sampling rate of 200 MHz, the raw data-rate generated by the entire LOFAR array is 13 Tbit/s, far too much to transport in total.

Even utilizing beam-forming at the station level, the long range data transport rates over the array are of order 150 Gbit/s requiring dedicated fibre networks. Such large data transport rates naturally also imply data storage challenges. For example, typical interferometric imaging observations can easily produce 35 Tbyte/h of raw, correlated visibilities. LOFAR is one of the first of a number of new astronomical facilities coming online that must deal with the transport, processing, and storage of these large amounts of data. In this sense, LOFAR represents an important technological pathfinder for the SKA and data intensive astronomy in the coming decade.

In addition to hardware and data transport challenges, LOFAR faces many technical challenges that are conceptual or algorithmic in nature. Low-frequency radio signals acquire phase-shifts due to variations in the total electron content of the ionosphere. For baselines longer than a few kilometers, the dynamic and non-isoplanatic nature of the ionosphere has a dramatic impact on the quality of the resulting scientific data. Correcting for these effects in LOFAR data has required improving existing calibration techniques that can simultaneously determine multi-directional station gain solutions to operate in the near, real-time regime. Likewise, LOFAR's huge FoV means the traditional interferometric assumption of a coplanar array is no longer valid. Consequently, highly optimized versions of imaging algorithms that recognize that the interferometric response and the sky brightness are no longer related by a simple 2D Fourier transform were required. These and similar issues have driven much of the design for LOFAR's processing software and computational architecture.

Scientifically, this new technology makes LOFAR a powerful and versatile instrument. With the longer European baselines in place, LOFAR can achieve sub-arcsecond angular resolution over most of its 30–240 MHz nominal operating bandpass, limited primarily by atmospheric effects and scattering due to interplanetary scintillation (IPS). This resolution, when combined with the large FoV, makes LOFAR an excellent instrument for all-sky surveys. Exploiting this potential has been one of LOFAR's key science drivers from its inception. The large effective area of LOFAR's densely populated core, support for multi-beaming, and inherent high time resolution also make LOFAR a breakthrough instrument for the detection and all-sky monitoring of transient radio sources. Finally, the ability to buffer large amounts of data at the dipole level provides a unique capability to perform retrospective imaging of the entire sky on short timescales. Among other applications, these buffers are used to detect radio emission from CR air showers. As discussed later, this versatility is apparent in the wide array of key science projects (KSPs) that have driven the initial design and commissioning phase.

3. Array configuration

The fundamental receiving elements of LOFAR are two types of small, relatively low-cost antennas that together cover the 30–240 MHz operating bandpass. These antennas are grouped together into 48 separate stations distributed over the northeastern part of the Netherlands as well as in Germany, France, the UK, and Sweden. The majority of these stations, 40 in total, are distributed over an area roughly 180 km in diameter centered near the town of Exloo in the northeastern Dutch province of Drenthe. This area was chosen because of its low population density and relatively low levels of radio frequency interference (RFI). The feasibility of obtaining the land required to build the stations (~20 000 m² per station) also played an important part in the final decision to site the array here.



Fig. 1. Aerial photograph of the Superterp, the heart of the LOFAR core, from August 2011. The large circular island encompasses the six core stations that make up the Superterp. Three additional LOFAR core stations are visible in the upper right and lower left of the image. Each of these core stations includes a field of 96 low-band antennas and two sub-stations of 24 high-band antenna tiles each.

For the majority of the array located in the Netherlands, the geographic distribution of stations shows a strong central concentration with 24 stations located within a radius of 2 km referred to as the “core”. Within the core, the land was purchased to allow maximum freedom in choosing station locations. This freedom allowed the core station distribution to be optimized to achieve the good instantaneous uv coverage required by many of the KSPs including the epoch of reionization (EoR) experiment and radio transients searches (see Sect. 13). At the heart of the core, six stations reside on a 320 m diameter island referred to as the “Superterp”; “terp” is a local name for an elevated site used for buildings as protection against rising water. These Superterp stations, shown in Fig. 1, provide the shortest baselines in the array and can also be combined to effectively form a single, large station as discussed in Sect. 12.10.

Beyond the core, the 16 remaining LOFAR stations in the Netherlands are arranged in an approximation to a logarithmic spiral distribution. Deviations from this optimal pattern were necessary due to the availability of land for the stations as well as the locations of existing fiber infrastructure. These outer stations extend out to a radius of 90 km and are generally classified as “remote” stations. As discussed below, these remote stations also exhibit a different configuration to their antenna distributions than core stations. The full distribution of core and remote stations within the Netherlands is shown in Fig. 2.

For the 8 international LOFAR stations, the locations were provided by the host countries and institutions that own them. Consequently, selection of their locations was driven primarily

by the sites of existing facilities and infrastructure. As such, the longest baseline distribution has not been designed to achieve optimal uv coverage, although obvious duplication of baselines has been avoided. Figure 3 shows the location of the current set of international LOFAR stations. Examples of the resulting uv coverage for the array can be found in Sect. 12.

4. Stations

LOFAR antenna stations perform the same basic functions as the dishes of a conventional interferometric radio telescope. Like traditional radio dishes, these stations provide collecting area and raw sensitivity as well as pointing and tracking capabilities. However, unlike previous generation, high-frequency radio telescopes, the antennas within a LOFAR station do not physically move. Instead, pointing and tracking are achieved by combining signals from the individual antenna elements to form a phased array using a combination of analog and digital beam-forming techniques (see Thompson et al. 2007). Consequently, all LOFAR stations contain not only antennas and digital electronics, but significant local computing resources as well.

This fundamental difference makes the LOFAR system both flexible and agile. Station-level beam-forming allows for rapid repointing of the telescope as well as the potential for multiple, simultaneous observations from a given station. The resulting digitized, beam-formed data from the stations can then be streamed to the CEP facility (see Sect. 6) and correlated to produce visibilities for imaging applications, or further combined



Fig. 2. Geographic distribution of LOFAR stations within the Netherlands. *Left:* this panel shows the distribution for the majority of the stations within the LOFAR core. The central, circular area contains the six Superterp stations described in the text. The white, polygonal areas mark the location of LOFAR core stations. In addition to the Superterp stations, 16 of the remaining 18 core stations are shown. *Right:* this panel shows the distribution of remote stations within the Netherlands located at distances of up to 90 km from the center of the array. Stations shown in green are complete and operational while yellow depicts stations that are under construction as of March 2013 (see Sect. 14.1).

Table 1. Overview of stations and antennas.

Station configurations	Number of stations	LBA dipoles	HBA tiles	Signal paths	Min. baseline (m)	Max. baseline (km)
Superterp	6	2×48	2×24	96	68	0.24
NL Core Stations	24	2×48	2×24	96	68	3.5
NL Remote Stations	16	2×48	48	96	68	121.0
International Stations	8	96	96	192	68	1158.0

Notes. The 6 stations comprising the central Superterp are a subset of the total 24 core stations. Please note that the tabulated baseline lengths represent unprojected values. Both the LBA dipoles and the HBA tiles are dual polarization.

into array beams (i.e. the sum of multiple stations) to produce high resolution time-series (e.g. for pulsar, CR, and solar studies). In effect, each individual LOFAR station is a fully functional radio telescope in its own right and a number of the main science drivers exploit this flexibility (e.g., see Sect. 5.3 of Stappers et al. 2011). In the following section, we review the major hardware and processing components of a LOFAR station.

4.1. Station configurations

As discussed in Sect. 3, LOFAR stations are classified as either core, remote, or international, nominally corresponding to their distance from the center of the array. More fundamentally, each of these three types of stations have different antenna field configurations. In its original design, all LOFAR stations were envisioned to be identical to simplify both construction and deployment as well as subsequent calibration. Due to funding considerations, this design was altered in 2006 to reduce costs while preserving the maximum number of stations possible and the corresponding quality of the uv coverage. This decision led to different choices for the antenna configurations and underlying electronics in the core, remote, and international stations. Consequently all LOFAR stations in the Netherlands have 96 signal paths that can be used to simultaneously process signals from either 48 dual-polarized or 96 single-polarized antennas.

To provide sufficient sensitivity on the longest baselines, international LOFAR stations are equipped with 192 signal paths. These three station types are summarized in Table 1.

The geometric distribution of low-band antennas (LBAs) and high-band antennas (HBAs) for each of the three LOFAR station configurations is shown in Fig. 4. All stations in the Netherlands have 96 LBAs, 48 HBAs, and a total of 48 digital receiver units (RCUs). These RCUs represent the beginning of the digital signal path and feature three distinct inputs per board (see Sect. 4.4 below). For core and remote stations in the Netherlands, two of these inputs are assigned to the 96 LBA dipoles while the remaining input is used for the 48 HBA tiles. Only one of these three RCU inputs, however, can be active at any one time. As a result, whereas all 48 HBA tiles can be used at once, only half the 96 signals coming from the LBA dipoles can be processed at any given time. Operationally, the LBA dipoles are grouped into an inner circle and an outer annulus each consisting of 48 dipoles and identified as the “LBA Inner” and “LBA Outer” configurations, respectively. These two configurations result in different FoVs, and potentially sensitivity (due to mutual coupling of closely spaced antennas), and can be selected by the observer during the observation specification process.

As Fig. 4 illustrates, a further distinction is apparent in the layout of HBA tiles within the core and remote stations in the Netherlands. In contrast to remote stations, where the HBAs are



Fig. 3. Current distribution of the European LOFAR stations that have been built in Germany (5), France (1), Sweden (1) and the UK (1). The color scheme for the stations is the same as in Fig. 2. A sixth German station located near Hamburg (shown in yellow) has recently begun construction and is expected to be online by the end of 2013. Data from all international stations is routed through Amsterdam before transfer to CEP in Groningen, NL. For the German stations, data are first routed through Jülich before transfer on to Amsterdam (see Sect. 5).

contained within a single field, the HBA dipoles in LOFAR core stations are distributed over two sub-stations of 24 tiles each. These core HBA sub-stations can be used in concert as a single station or separately as independent LOFAR stations. The latter option has the advantage of providing many more short baselines within the core and by extension a significantly more uniform uv coverage. In addition, many of the short baselines that result from the dual HBA sub-stations are redundant and therefore yield additional diagnostics for identifying bad phase and gain solutions during the calibration process. These advantages are especially important for science cases that depend critically

on the use of the LOFAR core such as the EoR experiment or the search for radio transients.

Since the stations are constructed with a finite number of individual elements, the digitally formed station beams have non-negligible sidelobe structure. The sidelobe pattern is particularly strong for the HBA stations, because the tiles are laid out on a uniform grid. In order to reduce the effect of bright off-axis sources contributing strongly to the visibility function when located in a sidelobe, the layout of each individual station is rotated by a particular angle. This rotation in turn causes the sidelobe pattern of each station to be projected differently on the sky




Fig. 4. Station layout diagrams showing core, remote and international stations. The large circles denote the LBA antennas while the arrays of small squares indicate the HBA tiles. Note that the station layouts are not shown on the same spatial scale.

from the others, so that the sensitivity to off-axis sources is reduced on any particular baseline. Note that only the station *layout* is rotated. Each of the individual dipole pairs are oriented at the same angle with respect to a commonly defined polarimetric axis.

Unlike stations in the Netherlands, international LOFAR stations are uniform and most closely follow the original station design. These stations consist of a full complement of 96 HBAs, 96 LBAs, and 96 RCUs. The additional RCUs in these stations provide a total of 192 digital signal paths such that the full set of HBA tiles or LBA dipoles are available during any given observation. In these stations, the third RCU input is currently not used and therefore available for possible future expansion. Several proposals are already under consideration that would take advantage of this unused capacity in order to expand the capabilities of the international LOFAR stations (see Sect. 14.2).

4.2. Low-band antenna

At the lowest frequencies, LOFAR utilizes the LBAs, which are designed to operate from the ionospheric cutoff of the “radio window” near 10 MHz up to the onset of the commercial FM radio band at about 90 MHz. Due to the presence of strong RFI at the lowest frequencies and the proximity of the FM band at the upper end, this range is operationally limited to 30–80 MHz by default. An analog filter is used to suppress the response below 30 MHz, although observers wishing to work at the lowest frequencies have the option of deselecting this filter (see van Weeren et al. 2012). In designing the LOFAR LBAs, the goal was to produce a sky-noise dominated receiver with all-sky sensitivity and that goal has largely been achieved over ~70% of the bandpass (see Sect. 12.6). At the same time, the resulting antenna needed to be sturdy enough to operate at least 15 years in sometimes harsh environmental conditions as well as be of sufficiently low cost that it could be mass produced. The resulting LBA is shown in Fig. 5.

The LBA element, or dipole, is sensitive to two orthogonal linear polarizations. Each polarization is detected using two

copper wires that are connected at the top of the antenna to a molded head containing a low-noise amplifier (LNA). At the other end, these copper wires terminate in either a synthetic, rubber spring or a polyester rope and are held in place by a ground anchor. The molded head of the LBA rests on a vertical shaft of PVC pipe. The tension of the springs and the ground anchor hold the antenna upright and also minimize vibrations in the wires due to wind loading. The dipole itself rests on a ground plane consisting of a metal mesh constructed from steel concrete reinforcement rods. A foil sheet is used to minimize vegetation growth underneath the antenna. Each polarization has its own output and hence two coaxial cables per LBA element run through the vertical PVC pipe. Power is supplied to the LNA over these same coaxial cables. The dipole arms have a length of 1.38 meter corresponding to a resonance frequency of 52 MHz. The additional impedance of the amplifier shifts the peak of the response curve to 58 MHz, however, as shown in the right panel of Fig. 5.

Despite the deceptively simple design, when coupled with digital beam-forming techniques, the LOFAR LBA dipole provides a powerful detection system at low frequencies. In particular, the omnidirectional response of the LBA antennas allows for the simultaneous monitoring of the entire visible sky. The LBA dipoles in a given LOFAR station can easily be correlated to provide all-sky maps on timescales of seconds (see Fig. 6). This novel capability is useful for a number of scientific objectives including studies of the large scale Galactic emission from the Milky Way and all-sky monitoring for radio transients.

4.3. High-band antenna

To cover the higher end of the LOFAR spectral response, an entirely different mechanical design has been utilized. The LOFAR HBA has been optimized to operate in the 110–250 MHz range. In practice, the frequency range above 240 MHz is heavily contaminated by RFI so operationally the band is limited to 110–240 MHz. At these frequencies, sky noise no longer dominates the total system noise as is the case for the LBAs. Consequently, another design topology for the HBA antennas

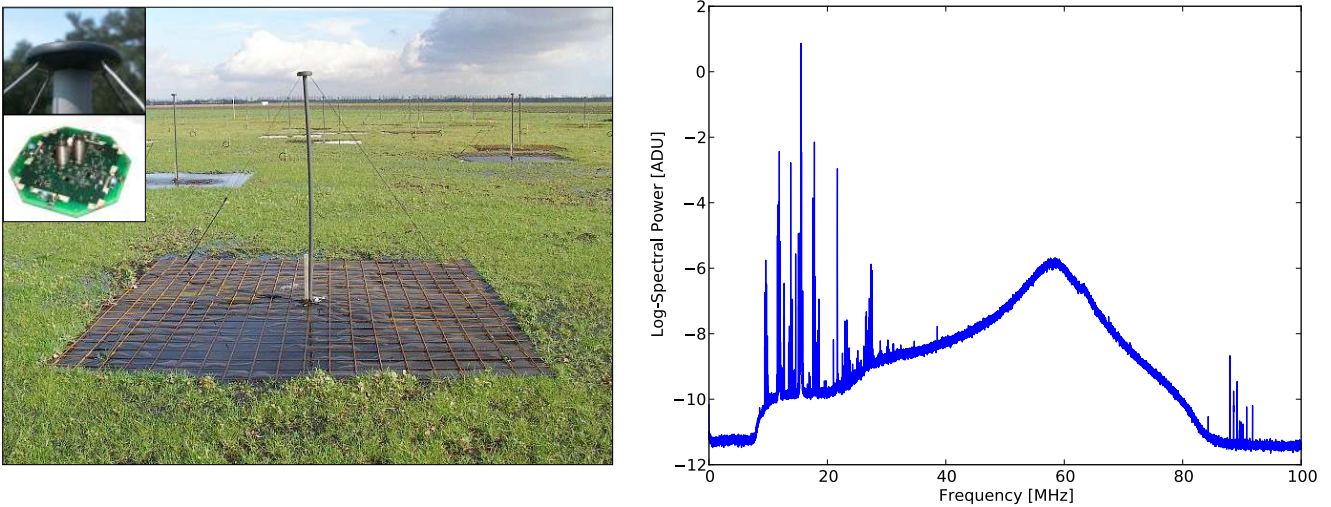


Fig. 5. *Left:* image of a single LOFAR LBA dipole including the ground plane. The inset images show the molded cap containing the LNA electronics as well as the wire attachment points. *Right:* median averaged spectrum for all LBA dipoles in station CS003. The peak of the curve near 58 MHz is clearly visible as well as strong RFI below 30 MHz, partly because of ionospheric reflection of sub-horizon RFI back toward the ground, and above 80 MHz, due to the FM band.

was required in order to minimize contributions to the system noise due to the electronics. Nonetheless, the HBA design was of course subject to the similar constraints on environmental durability and low manufacturing cost as the LBA design. An image of the final HBA tile is shown in Fig. 7.

In order to minimize cost while maintaining adequate collecting area, the HBA design clusters 16 antenna elements together into “tiles” that include initial analog amplification and a first stage of analog beam-forming. A single “tile beam” is formed by combining the signals from these 16 antenna elements in phase for a given direction on the sky. Hence, while the LBAs are effectively passive (requiring power but no active control and synchronization), the HBAs contain tile-level beam forming and are subject to control by the Monitoring and Control system MAC (see Sect. 9.1).

A single HBA tile consists of a square, 4×4 element (dual polarized) phased array with built-in amplifiers and an analog beam-former consisting of delay units and summatoms. The 5 bit time delay can be up to 15 ns long and is set by a signal received from the MAC system. Each 16 element tile measures 5×5 meter and is made of an expanded polystyrene structure which supports the aluminum antenna elements. The distance between tile centers is 5.15 m resulting in a spacing between tiles of 15 cm. The contents of the tile are protected from weather by two overlapping flexible polypropylene foil layers. A light-weight ground plane consisting of a 5×5 cm wire mesh is integrated into the structure. As with the LBAs, the resulting signals are transported over coaxial cables to the receiver unit in the electronics cabinet.

4.4. Receiver unit

At the receiver unit (RCU), the input signals are filtered, amplified, converted to base-band frequencies and digitized. A sub-sampling architecture for the receiver is used. This choice implies a larger required analog bandwidth and multiple band-pass filters to select the frequency band of interest. The receiver is designed to be sky noise limited so a 12 bit A/D converter is used with 3 bits reserved to cover the anticipated range of sky noise and the rest available for RFI headroom. This number of bits is

sufficient to observe signals, including strong RFI sources, with strengths up to 40 dB over and above the integrated sky noise in a bandwidth of at least 48 MHz.

Because observing in the FM band is not feasible, a sampling frequency of 200 MHz has been chosen for most of the receiver modes. This sampling results in a Nyquist edge almost at the center of the FM band. To cover the region around 200 MHz in the HBA band, which will suffer from aliasing due to the flanks of the analog filter, an alternative sampling frequency of 160 MHz is also supported. These choices result in several possible observing bands to cover the total HBA frequency range. The available frequency bands are summarized in Table 2.

As discussed above in Sect. 4.1, three main signal paths can be distinguished in the RCU. For stations in the Netherlands, two of these are allocated to the two sets of LBAs, although only one can be used at any given time. One of these signal paths was originally intended for a (not currently planned) low-band antenna optimized for the 10–30 MHz frequency range. For the present LBA, either a 10-MHz or 30-MHz high-pass filter can be inserted to suppress the strong RFI often encountered below 20 MHz. The remaining signal path is used for the HBA. It is first filtered to select the 110–250 MHz band and then again by one of three filters that select the appropriate Nyquist zones listed in Table 2.

4.5. Digital signal processing

Both the LBA and HBA antennas are connected via coaxial cables to the electronics housed in a cabinet located on the edge of each LOFAR station. This cabinet is heavily shielded and contains the RCUs, digital signal processing (DSP) hardware, local control unit (LCU), and other equipment used to perform the first data processing stage. After digitization by the RCUs, the datastreams enter the digital electronics section. This section is mainly responsible for beam-forming although either raw or filtered signals can also be stored in a circular buffer in order to trap specific events (see Sect. 4.6 below). Further processing is done by the remote station processing (RSP) boards utilizing low-cost, field programmable gate arrays (FPGAs). These FPGAs provide

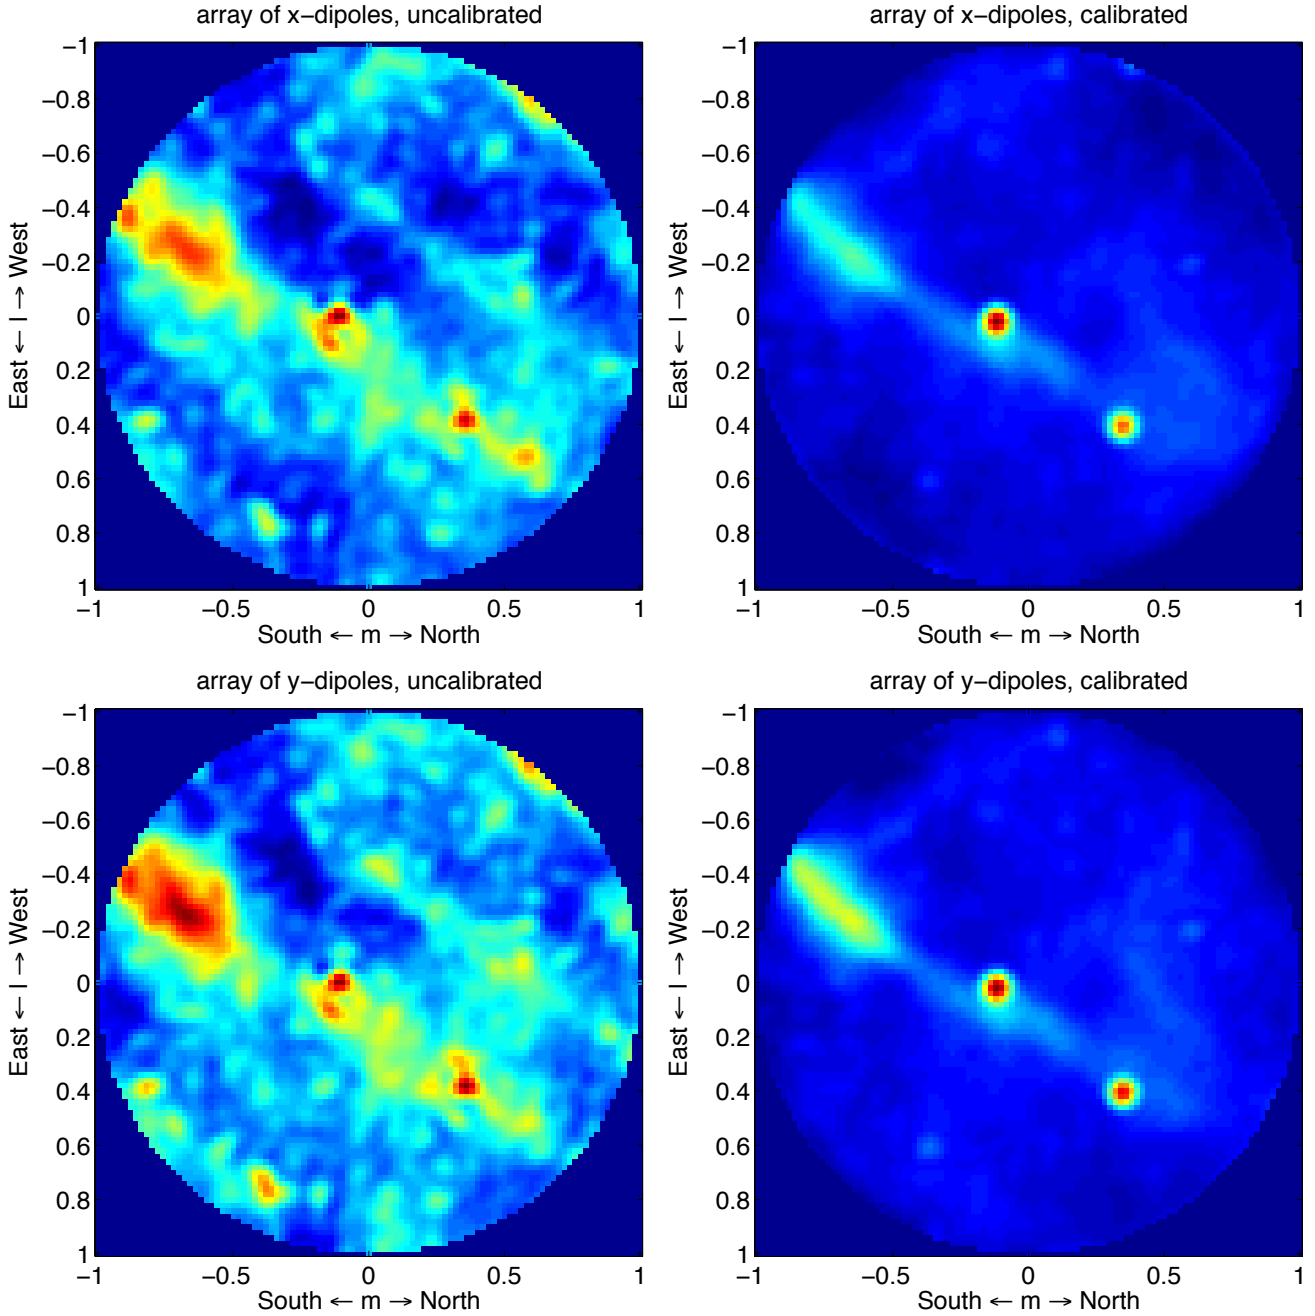


Fig. 6. All-sky observation produced by a single LOFAR station (station FR606 in Nançay, France) and created offline by correlating the signals from each of the individual dipoles in the station. The station level data collection and processing is described in Sect. 4.4–4.7. The observation was taken at a frequency of 60 MHz, with a bandwidth of only 195.3125 kHz (1 subband). The integration time was 20 s. Even with this limited dataset, Cassiopeia A, Cygnus A, and the Galactic plane are all clearly visible. *The left panels* show images made from uncalibrated station data while the calibrated images are shown on *the right*. *The upper and lower panels* give images for the X and Y polarizations, respectively.

sufficient computing power to keep up with the datastream and can also be updated remotely allowing for easy patches and enhancements to be applied. Following the beam-forming step, the data packets are streamed over the wide-area network (WAN) to the CEP facility in Groningen. A schematic of this data flow is given in Fig. 8.

Once digitized, the RSP boards first separate the input signals from the RCUs into 512 sub-bands via a polyphase filter (PPF). Further processing is done per sub-band. The sub-bands have widths of 156 kHz or 195 kHz depending on whether the 160 MHz or 200 MHz sampling clock is selected, respectively. By default sample values are stored using 16 bit floating

point representations allowing up to 244 of these sub-bands to be arbitrarily distributed over the bandpass for a total bandwidth of 48 MHz per polarization. Alternatively, the station firmware may be configured to utilize an 8 bit representation for the sample values yielding up to 488 sub-bands for a total bandwidth of 96 MHz per polarization. Although providing increased bandwidth, this 8 bit mode is potentially more vulnerable to periods of strong RFI. The frequency selection can vary for each station and is configured by the user during the initial observation specification.

After formation of the sub-bands, the primary processing step is the digital, phase rotation-based beam-former. This




Fig. 7. *Left:* closeup image of a single LOFAR HBA tile. The protective covering has been partially removed to expose the actual dipole assembly. The circular dipole rotation mechanism is visible. *Right:* median averaged spectrum for all HBA tiles in station CS003. Various prominent RFI sources are clearly visible distributed across the band including the strong peak near 170 MHz corresponding to an emergency pager signal.

Table 2. Overview of LOFAR system parameters.

System characteristic	Options	Values	Comments	
Frequency range	Low-band Antenna	10–90 MHz	With analog filter	
		30–90 MHz		
	High-band Antenna	110–190 MHz		200 MHz sampling (2nd Nyquist zone)
		170–230 MHz		160 MHz sampling (3rd Nyquist zone)
		210–250 MHz	200 MHz sampling (3rd Nyquist zone)	
Number of polarizations		2		
Bandwidth	Default	48 MHz	16-bit mode	
	Maximum	96 MHz	8-bit mode	
Number of simultaneous beams	Minimum	1		
	Maximum	244	16 bit mode, one per sub-band	
	Maximum	488	8 bit mode, one per sub-band	
Sample bit depth		12		
Sample rate	Mode 1	160 MHz		
	Mode 2	200 MHz		
Beamformer spectral resolution	Mode 1	156 kHz		
	Mode 2	195 kHz		
Channel width (raw correlator resolution)	Mode 1	610 Hz		
	Mode 2	763 Hz		

beam-former sums the signals from all selected RCUs after first multiplying them by a set of complex weights that reflect the phase rotation produced by the geometrical and other delays toward a certain direction. The weights are calculated in the local control unit (see Sect. 4.7) and sent to the RSP boards during the observation. The update rate of the beam-former is set to 1 s by default resulting in about 0.3% gain variation for a station beam of 3° in diameter. The beam-forming is done independently per sub-band and the resulting beam for each sub-band is referred to as a “beamlet”. Multiple beamlets with the same pointing position can be combined to produce beams with larger bandwidth.

The number of simultaneous beams that may be constructed can in principle be as high as the number of beamlets since all operate independently of each other. Operationally, the number of independent beams per station is currently limited to 8, although this limit will ultimately increase. Successful experiments utilizing the maximum 244 beams available in 16 bit

mode have already been conducted. Similarly, for 8 bit observations, a maximum of 488 beams can be formed. The 48 MHz (16 bit mode) or 96 MHz (8 bit mode) total bandwidth can be distributed flexibly over the number of station beams by exchanging beams for bandwidth. In the case of the LBA, simultaneous beams can be formed in any combination of directions on the sky. While strictly true for the HBA as well, HBA station beams can only usefully be formed within pointing directions covered by the single HBA tile beam, corresponding to a FWHM of $\sim 20^\circ$ at 140 MHz.

4.6. Transient buffer boards

In addition to the default beam-forming operations, the LOFAR digital processing also provides the unique option of a random access memory (RAM) buffer at station level. These RAM




Fig. 8. Schematic illustrating the signal connections at station level as well as the digital processing chain. After the beam-forming step, the signals are transferred to the correlator at the CEP facility in Groningen.

buffers provide access to a snapshot of the running data-streams from the HBA or LBA antennas. As depicted in Fig. 8, a dedicated transient buffer board (TBB) is used that operates in parallel with the normal streaming data processing. Each TBB can store 1 Gbyte of data for up to 8 dual-polarized antennas either before or after conversion to sub-bands. This amount is sufficient to store 1.3 s of raw data allowing samples to be recorded at LOFAR’s full time resolution of 5 ns (assuming the 200-MHz sampling clock). Following successful tests for various science cases (see Sect. 11.3), an upgrade of the RAM memory to store up to 5 s of raw-data has been approved and is currently being installed. The temporal window captured by the TBBs can be further extended by up to a factor of 512 by storing data from fewer antennas or by storing sub-band data. We note that while the TBBs may operate in either raw timeseries or sub-band mode, they can not operate in both at the same time.

Upon receiving a dump command, the TBB RAM buffer is frozen and read out over the WAN network directly to the storage section of the CEP post-processing cluster (see Sect. 6.2). These commands can originate locally at the station level, from the system level, or even as a result of triggers received from other telescopes or satellites. At the station level, each TBB is constantly running a monitoring algorithm on the incoming data-stream. This algorithm generates a continuous stream of event data that is received and processing by routines running on the local control unit (LCU). If the incoming event stream matches the pre-defined criteria, a trigger is generated and the TBBs are read out. As discussed in Sect. 11.3, this local trigger mechanism gives LOFAR the unique ability to respond to ns-scale events associated with strong CRs. The Transients KSP also intends to utilize this functionality to study fast radio transients (see Sect. 11.4).

4.7. Local control unit

Each LOFAR station, regardless of configuration, contains computing resources co-located adjacent to the HBA and LBA antenna fields. This local control unit (LCU) is housed inside the RF-shielded cabinet containing the other digital electronics and consists of a commodity PC with dual Intel Xeon 2.33 GHz quad-core CPUs, 8 Gbyte of RAM, and 250 Gbyte of local disk

storage. The station LCUs run a version of Linux and are administered remotely over the network from the LOFAR operations center in Dwingeloo. Processes running on the LCU can include control drivers for the TBBs, RCUs, and other hardware components as well as additional computational tasks. All processes running on the LCUs are initialized, monitored, and terminated by the MAC/SAS control system discussed below in Sect. 9.

Computationally the LCU provides several crucial computing tasks at the station level. Chief among these are the beam-former computations mentioned previously in Sect. 4.5. The number of independent beams that may be supported is limited by the processing power of the LCU since it must calculate the appropriate weights for each direction on the sky every second.

Equally important, the LCU runs a station-level calibration algorithm to correct for gain and phase differences in all the individual analog signal paths. The correlation matrix of all dipoles in the station is calculated for one sub-band each second as input to this calibration and the procedure runs in real-time during an observation (Wijnholds & van der Veen 2009, 2010; Wijnholds et al. 2010). The algorithm cycles through the selected sub-bands, with a new sub-band calibrated each second, resulting in an updated calibration for the complete band every 512 s. This active calibration is necessary to compensate for environmental temperature variations that cause gain and phase drifts in the signal paths (see the discussion in Sect. 12.1). The array correlation matrix can also be used for RFI detection and mitigation (Boonstra & van der Tol 2005).

Additional computational tasks can also be run on the LCU subject to the constraint that they do not impact the performance of the core calibration and beam-forming capabilities. Current examples of these station-level applications include the TBB trigger algorithms discussed previously in Sect. 4.6. We note that adding additional compute capacity to the LCU is a fairly straightforward way to expand the capabilities of the LOFAR array (see Sect. 14.2 for some currently planned enhancements).

5. Wide-area network

The function of the LOFAR Wide-Area Network (WAN) is to transport data between the LOFAR stations and the central

processor in Groningen. The main component is the streaming of observational data from the stations. A smaller part of the LOFAR datastream consists of Monitoring And Control (MAC) related data and management information of the active network equipment. Connections of the LOFAR stations in the Netherlands to Groningen run over light-paths (also referred to as managed dark fibers) that are either owned by LOFAR or leased. This ensures the required performance and security of the entire network and the equipment connected to it. Signals from all stations in the core and an area around it are first sent to a concentrator node and subsequently patched through to Groningen.

The LOFAR stations outside the Netherlands are connected via international links that often involve the local NRENs (National Research and Education Networks). In some cases, commercial providers also play a role for part of the way.

For the communication over the light-paths 10 Gigabit Ethernet (GbE) technology has been adopted. The high bandwidth connection between the concentrator node in the core and Groningen uses Course Wavelength Division Multiplexing (CWDM) techniques to transfer multiple signals on a single fiber, thereby saving on costs. Since the availability requirement for LOFAR is relatively low (95%), when compared with commercial data communication networks, redundant routing has not been implemented.

6. Central processing (CEP)

LOFAR's CEP facility is located at the University of Groningen's Centre for Information Technology (CIT). The CIT houses the hardware for the CEP system but also part of the distributed LTA discussed in Sect. 7. With the exception of standalone operation where a given LOFAR station can be used locally independent from the rest of the array, data from all LOFAR stations, including the international stations, is received at CEP in a streaming mode. At CEP these raw datastreams are subsequently processed into a wide variety of data products as discussed in Sect. 11 below.

The CEP facility can be broadly divided into two essentially autonomous sections. The "online" section collects and processes the incoming station datastreams in real-time and all operations on the data are completed before it is written to disk. Once the initially processed data-streams are stored, additional, less time-critical processing is done on the "offline" section to produce the final set of LOFAR data products. A large storage cluster connects these two distinct processing phases. The same Monitoring and Control system discussed in Sect. 9 and used to operate the stations themselves also manages the allocation of processing and storage resources at CEP. Multiple observations and processing streams on both the online and offline sections can be performed in parallel. In the following, we briefly review the major features of these two components.

6.1. Online central processing

The online processing section handles all real-time aspects of LOFAR and is built around a three-rack IBM Blue Gene/P (BG/P) supercomputer. Current LOFAR operations are limited to one rack of the three available. Each rack of the BG/P is equipped with 64 individual 10 GbE interfaces (I/O nodes). A single LOFAR station can be mapped to one I/O node. The peak performance of each rack is 14 Tflop/s. The processing power and I/O bandwidth of one rack is sufficient to correlate 2048 baselines at full-polarization for the maximum bandwidth of 48 MHz with an integration time of one second.

Each BG/P I/O node receives data from a station and runs a data-handling application that buffers the input data and synchronizes its output stream with the other input nodes based on the timestamps contained in the data. For imaging observations, the BG/P performs its main function as the correlator of the array. As Fig. 9 shows, it can also support a variety of other processing streams including the formation of tied-array beams and real-time triggering. Combinations of these processing streams can be run simultaneously subject to resource constraints.

The current set of supported online processing streams is depicted in Fig. 9. Most of these represent the initial processing stages in the observing modes discussed in Sect. 10. Several common transformations are applied to all incoming station datastreams regardless of subsequent processing. For example, time offsets are applied to each incoming datastream to account for geometric delays caused by differing station distances from the array phase center. These offsets must be calculated on-the-fly since the rotation of the Earth alters the orientation of the stations continuously with respect to the sky. For observations with multiple beams, unique delays must be calculated for each beam.

Once the geometric delays are applied, a transpose operation is performed to reorder the now aligned station data packets. Incoming data packets from the stations are grouped as a set of sub-bands per station. After the transpose, the data are rearranged such that all station data for a given sub-band is grouped. At this point, a second polyphase filter is applied to resample the data to the kHz level. The filter-bank implemented on the BG/P splits a 195 kHz (or 156 kHz) sub-band datastream into, typically, 256 frequency channels of 763 Hz (or 610 Hz) each. Splitting the data into narrow frequency channels allows the offline processing to flag narrow-band RFI, so that unaffected channels remain usable.

In classical radio telescopes an XF correlator was generally used, meaning that first the correlation and integration of the signals was done in the time domain (X) and afterwards the Fourier transform (F) was accomplished to get a cross power spectrum out of the correlator (Romney 1999). This option is still an economically attractive technique for radio telescopes with a limited number of antennas (input signals to the correlator). However, for LOFAR an FX correlator (first Fourier transform and then correlating the resulting channels) is favorable in terms of processing at the expense of data transport (the signals must be regrouped per channel instead of per antenna, resulting in a transpose operation). Using only a Fourier transform in the FX correlator leads to a significant amount of leakage between the channels. Therefore it was chosen to use filter banks before the correlator. This architecture is also known as an HFX (Hybrid FX correlator) architecture (Romney 1999).

The correlator calculates the auto and cross correlations between all pairs of stations, for each channel and for each polarization (XX, XY, YX, and YY). A correlation is the complex product of a sample from one station and the complex conjugate of a sample from the other station. By default, the results are integrated (accumulated) over one second of data; however, smaller integration times are possible for applications such as full-field imaging with the international stations or fast solar imaging. Since the correlation of station S1 and S2 is the conjugate of the correlation of station S2 and S1, we only compute the correlations for $S1 \leq S2$. The output data rate of the correlator is significantly lower than the input data rate. To achieve optimal performance, the correlator consists of a mix of both C++ and assembler code, with the critical inner loops written entirely in assembly language (Romein et al. 2006, 2010).

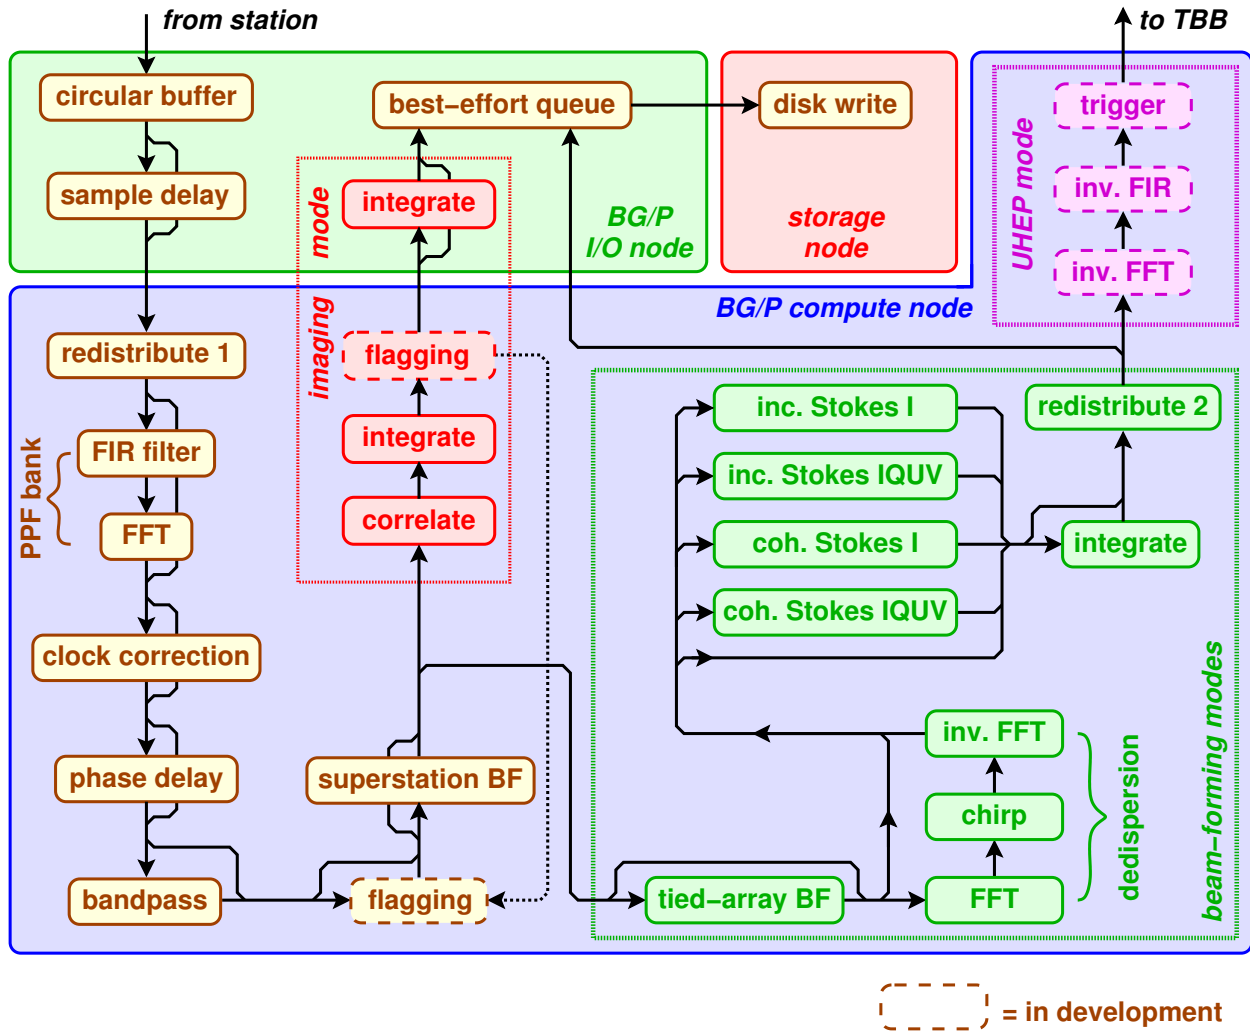


Fig. 9. Schematic showing the possible online data processing paths currently available or under development. These pipelines run in real-time on the IBM Blue Gene/P supercomputer that comprises the core of LOFAR’s online processing system (see Sect. 6.1). This schematic illustrates that many processing steps can be selected or deselected as necessary. Pipelines can also be run in parallel with, for example, the incoming station datastream being split off to form both correlated and beam-formed data simultaneously. The imaging and beam-formed data pipelines are indicated separately. The online triggering component of the CR UHEP experiment currently under development is also shown (see Sect. 13.5).

6.2. Offline central processing

The offline central processing cluster provides disk space for the collection of datastreams and storage of complete observation datasets for offline processing. This storage is intended for temporary usage (typically a week) until the final data products are generated and archived or the raw data themselves are exported or archived. In addition to the storage part the offline cluster offers general-purpose compute power and high bandwidth interconnections for the offline processing applications.

The offline cluster is a Linux cluster that is optimized for cost per flop and cost per byte. The cluster consists of 100 hybrid storage/compute nodes. Each node has 12 disks of 2 Tbyte each providing 20 Tbyte of usable disk space per node. Furthermore, each node contains 64 Gbyte of memory and 24, 2.1 GHz cores. Thus the cluster has 2 Pbyte of storage capacity total and 20.6 Tflop/s peak performance. The offline tasks differ depending on the application at hand. For example for the imaging application the offline tasks are typically flagging of bad data, self-calibration and image creation.

In addition to the offline cluster extra processing power will be available in GRID networks in Groningen or at remote sites.

GRID networks also provide the basic infrastructure for the LOFAR archive enabling data access and data export to users.

7. LOFAR long-term archive

The LOFAR long-term archive (LTA) is a distributed information system created to store and process the large data volumes generated by the LOFAR radio telescope. When in full operation, LOFAR can produce observational data at rates up to 80 Gbit/s. Once analyzed and processed, the volume of data that are to be kept for a longer period (longer than the CEP storage is able to support) will be reduced significantly. These data will be stored in the LTA and the archive of LOFAR science data products is expected to grow by up to 5 Pbyte per year. The LTA currently involves sites in the Netherlands and Germany.

For astronomers, the LOFAR LTA provides the principal interface not only to LOFAR data retrieval and data mining but also to processing facilities for this data. Each site involved in the LTA provides storage capacity and optionally processing capabilities. To allow collaboration with a variety of institutes and

projects, the LOFAR LTA merges different technologies (EGI, global file systems, Astro-WISE dataservers). Well-defined interfaces ensure that to both the astronomer and the LOFAR observatory the LTA behaves as a coherent information system. Access and utilization policies are managed via the central LOFAR identity management system that is designed to allow federation with organizational user administrations. The network connecting LOFAR to the LTA sites in Groningen, Amsterdam and Jülich, Germany consists of light-path connections, currently utilizing 10GbE technology, that are shared with LOFAR station connections and with the European eVLBI network (e-EVN; Szomoru 2008).

The 10 Gbit/s bandwidth between the sites is sufficient for regular one-time LTA data transports but to allow transparent processing within the LTA it may grow to 60–80 Gbit/s bandwidth in the future. Such bandwidths will enable two major new processing modes: 1) Streaming of realtime or buffered observation data to a remote HPC system; 2) Streaming of stored data from one LTA site to a compute cluster located at another site. With these modes an optimal utilization of storage and processing facilities can be realized. If additional processing capacity is required for a given observing mode or for large-scale data processing, existing resources at partner institutes can be brought in without having to store (multiple copies of) datasets before processing can commence. For LOFAR datasets, which can grow up to hundreds of Tbyte, this capability will be essential.

8. Operations and management

Everyday LOFAR operations are coordinated and controlled from ASTRON's headquarters in Dwingeloo. Operators perform the detailed scheduling and configuration of the instrument, which includes setting up the appropriate online processing chain and destination of the data. The proper functioning of the stations, WAN and CEP system can be verified remotely. The monitoring and control system also collects and analyses the meta data gathered throughout the system in order to trace (impending) problems. Maintenance and repair of systems in the field is carried out under supervision of ASTRON personnel or the staff of an international station owner. Central systems maintenance is performed by staff of the University of Groningen's Centre for Information Technology. Advice and support is also given to the staff of the international partners who retain overall responsibility for their stations.

The International LOFAR Telescope (ILT) is a foundation established in Dwingeloo, the Netherlands, to coordinate the exploitation of the LOFAR resources under a common scientific policy. ASTRON provides the central operational entity for the ILT and the foundation is governed by a board consisting of delegates from each of the national consortia as well as a separate delegate from ASTRON itself. In relative proportion to their number of stations, the national owners put together the central exploitation budget. All observing proposals utilizing ILT facilities are reviewed on scientific merit by an independent ILT programme committee (PC). In the first and second years of operation, 10% and 20%, respectively, of the LOFAR observing and processing capacity will be distributed directly under Open Skies conditions and available to the general astronomical community. For the remainder, the national consortia each play a role in distributing reserved access shares, partly following national priorities, and partly taking into account the PC rankings. The fractions of time for open and reserved access in later years will be set by the ILT board.

9. Software control infrastructure

9.1. Monitoring and control system

In the data processing pipeline of LOFAR, real time control is required to set the instrument in a certain state at a defined time. Furthermore, the instrument needs to be able to quickly switch between observing modes and be able to track sources. Hence, a distinction is made between real-time control during data taking and processing on the one hand, and control prior to this phase (mainly specification) and after that phase (mainly inspection) on the other hand. This separation is motivated by the different types of database technology and software design issues related to real-time operation requirements.

The Specification, Administration and Scheduling (SAS) subsystem takes care of the specification and configuration of all observations and instrument settings. In contrast, the Monitoring and Control (MAC) subsystem is responsible for the operation of the instrument and the execution of observations, while collecting meta-data about those operations and observations. All user interaction is through the SAS and MAC systems. MAC is used to interface to running observations or processing pipelines. SAS is used for all other interaction prior to execution and afterwards. There is no “direct” interaction with applications. Interfaces to specific processing applications are implemented through the MAC layer and via a set of SAS GUIs. Finally, the SAS subsystem is used to provide an interface to the users for the collected meta-data and possible snapshots to inspect the observation performance and quality.

At system level the choice has been made to control LOFAR centrally so that information is collected (and accessible) in a single place as much as possible. However, one of the design requirements is that the stations should be able to function for at least one hour autonomously. Hence, in each station a Local Control Unit (LCU) is present which controls the complete station (see Sect. 4.7). In practice, the stations can operate autonomously indefinitely. All LCU functions are controlled remotely from the LOFAR operations center via the MAC system.

9.2. System health monitoring (SHM)

The percentage of time during which the LOFAR system is effectively operational, i.e. the system uptime, is an important issue that warrants considerable attention. Due to the complexity of the LOFAR system and the harsh operating environment, it is almost certain that at any moment in time several of LOFAR's components (antennas, amplifiers, network links, computing nodes, etc.) will be non-functional. In the Netherlands, for example, the moisture levels and high humidity can lead to higher rates of component failure. Within reasonable bounds, this fact should not impact the usability of LOFAR for performing useful scientific measurements; rather, the system performance should gracefully degrade with each failing component until repairs can be effected.

Any faulty component may affect the quality of the measurements in a negative way, and may also jeopardize the operational capabilities of the LOFAR network. The objective of the SHM module of the LOFAR system is to support the efforts to maximize the system uptime. The main functions of the module will be the early detection of system failure, the accurate identification of failing components, and the support for remedial actions.

Daily or weekly on-site inspections cannot be performed in an economically viable way (at least for the remote stations).

Hence the SHM subsystem will primarily be guided by the data that is generated by the LOFAR system in an automated fashion. This information consists of both the scientific data (generated by the antennas) and the “housekeeping” data of the equipment that controls the sensor network. Deviations in system health will be reflected in sensor data that deviate from the normative measurements. These deviations are called symptoms and are used by the SHM module to detect system failure and identify the responsible system component.

9.3. Event triggers

LOFAR’s digital nature makes it an inherently responsive telescope. With few moving parts, the ability to observe multiple targets simultaneously, and a software-driven control system, it is possible to make the telescope react intelligently to events (such as the detection of a fast transient, Sect. 13.3, or CR, Sect. 13.5) as they happen, enabling the full capabilities of the telescope (long baselines, TBBs) to be rapidly brought to bear and ultimately maximizing scientific output.

The LOFAR pipeline system will make it possible to generate triggers as part of regular data processing, or in response to notifications from other facilities. The scheduler and control systems will then be able to insert appropriate follow-up actions into the schedule on the fly. Such actions will include, for example, reconfiguring the array, performing a new observation, or re-running a data processing pipeline with modified parameters.

For exchanging information about transient events with other facilities, LOFAR has standardized on the International Virtual Observatory Alliance (IVOA) VOEvent system (Seaman et al. 2008, 2011). VOEvent provides a convenient and flexible way of representing and publishing information about events in a structured form that is well suited for machine processing. Although the full LOFAR VOEvent system is still under development, a VOEvent-based trigger has already been used to initiate LOFAR follow-up observations of gravitational wave event candidates detected by LIGO during September and October 2010 (LIGO Scientific Collaboration et al. 2012).

10. Observing modes

10.1. Interferometric imaging

The interferometric imaging mode provides correlated visibility data, just like traditional aperture synthesis radio telescope arrays consisting of antenna elements. The goal of the LOFAR imaging mode is to achieve high fidelity, low noise images of a range of astronomical objects, using customizable observing parameters. In this operating mode, station beams are transferred to the CEP facility where they are correlated to produce raw visibility data. The raw uv data are stored on the temporary storage cluster. Further processing, which consists of calibration and imaging (see Sect. 11.1), is handled off-line. Calibration is an iterative process of obtaining the best estimates of instrumental and environmental effects such as electronic station gains and ionospheric delays.

The final data products for this mode include the calibrated uv data, optionally averaged in time and frequency, and corresponding image cubes. The visibility averaging is performed to a level which reduces the data volume to a manageable level, while minimizing the effects of time and bandwidth smearing. It will be possible to routinely export datasets to investigators for reduction and analysis at their Science Centre or through the use of suitable resources on the GRID.

For imaging observations, a wide range of user interaction will be supported. Experienced users will require control over the calibration and imaging stages of data reduction, while more typical users will not wish to recalibrate the visibility data, but may need to control imaging parameters. Many users may require only a fully processed image. The MAC system will provide personalized control over key aspects of the calibration and imaging pipelines. For expert users, interactive control of this processing will be available using the SAS and MAC GUIs over the network.

This mode requires medium to long-term storage of uncalibrated or partially calibrated data at the CEP facility, to allow reprocessing of data following detailed inspection of results by the user. The resulting storage and processing requirements may impose limits on the amount of such customized reprocessing which may be conducted in the early years of LOFAR operation.

10.2. Beam-formed modes

Instead of producing interferometric visibilities, LOFAR’s beam-formed modes can either combine the LOFAR collecting area into “array beams”- i.e. the coherent or incoherent sum of multiple station beams - or return the un-correlated station beams from one or more stations (see also Stappers et al. 2011; Mol & Romein 2011). These data are used to produce time-series and dynamic spectra for high-time-resolution studies of, e.g., pulsars, (exo)planets, the Sun, flare stars, and CRs. These modes are also useful for system characterization and commissioning (e.g. beam-shape characterization, offline correlation at high time resolution, etc.). In the current implementation, there are several beam-formed sub-modes: i) Coherent Stokes; ii) Incoherent Stokes; and iii) Fly’s Eye. These can all be run in parallel in order to produce multiple types of data products simultaneously.

The Coherent Stokes sub-mode produces a coherent sum of multiple stations (also known as a “tied-array” beam) by correcting for the geometric and instrumental time and phase delays. This produces a beam with restricted FoV, but with the full, cumulative sensitivity of the combined stations. This sub-mode can currently be used with all 24 LOFAR core stations which all receive the same clock signal and hence do not require a real-time clock calibration loop for proper phase alignment. This coherent summation results in a huge increase in sensitivity, but with a limited FoV of only $\sim 5'$ (see Fig. 10). The Superterp and in fact the entire 2-km LOFAR core are compact enough that ionospheric calibration is also not likely to be a major limitation to coherently combining these stations, at least not for the high band. In practice, experience has shown that the calibration tables used to correct for the delays are stable over timescales of many months and need only be updated occasionally.

In this mode, one can write up to ~ 300 simultaneous, full-bandwidth tied-array beams as long as the time and frequency resolution are modest (for limitations and system benchmarking results, see Mol & Romein 2011). Note that the Superterp tied-array beams have a FWHM of $\sim 0.5^\circ$ and roughly 127 are required to cover the full single station FoV (see Fig. 11). Depending on the scientific goal of the observations, either Stokes I or Stokes I, Q, U, V can be recorded, with a range of possible frequency (0.8–195 kHz) and time ($\geq 5.12 \mu\text{s}$) resolutions. It is also possible to record the two Nyquist-sampled linear polarizations separately, which is referred to as “Complex Voltage” mode. This mode is necessary for applications such as offline coherent dedispersion, fast imaging, or inverting the initial, station-level poly-phase filter to achieve the maximum possible time resolution.

Table 3. Current LOFAR observing modes.

Type	Mode	Outputs	Description
Interferometric	Correlated	Visibilities	Arbitrary number of stations, 8 beams per station, full Stokes
Beam-formed	Incoherent stokes	BF data file	Incoherent summation, arbitrary stations, 8 station beams, full Stokes
	Coherent stokes	BF data file	Coherent summation, Superterp only, 20 full-resolution beams, full Stokes
	Complex voltage	BF data file	Coherent summation, Superterp only, bypasses 2nd PPF, raw voltage output
	Station level	BF data file	Arbitrary stations, individual pointing and frequency settings per station
Direct storage	Raw voltage	TBB data file	8 station beams, Stokes I
			Station level triggering of TBB dumps, direct storage to CEP cluster




Fig. 10. Increase in the signal-to-noise ratio (S/N) of the pulsar PSR B1530+27 as a function of the number of coherently (squares) and incoherently (triangles) added HBA sub-stations in the LOFAR core. The S/N is seen to increase linearly (solid line) in the case of coherent addition and as the square-root (dashed line) of the number of stations in the case of incoherent addition – as expected for sources that do not contribute significantly to the system temperature. Coherently and incoherently summed data were acquired simultaneously in 11 separate observations that summed between 1 to 42 HBA sub-stations. Note that the typical error on the S/N measurements is $\sim 10\%$ and these measurements are also systematically affected by the intrinsic brightness of the source (pulse-to-pulse brightness variations) as well as RFI.

The Incoherent Stokes sub-mode produces an incoherent combination of the various station beams by summing the powers after correcting for the geometric delay. This produces beams with the same FoV as a station beam, but results in a decrease in sensitivity compared with a coherently added tied-array beam – i.e. the gain in sensitivity scales with the square-root of the number of stations as opposed to linearly (see Fig. 10). One incoherent array beam can be formed for each of the beams created at station level – e.g., if all the stations being summed split their recorded bandwidth across 8 pointing directions, then 8 incoherent array beams can also be formed from these. All LOFAR stations, including the international stations, can be summed in this sub-mode, which can be run in parallel with the Coherent Stokes sub-mode. As in Coherent Stokes, one can record either

Stokes I or Stokes I, Q, U, V with the same range of frequency and time resolutions.

The Fly’s Eye sub-mode records the individual station beams (one or multiple per station) *without* summing. As with the Coherent and Incoherent Stokes modes, the normal online BG/P processing steps (e.g. channelization and bandpass correction) are still applied. This mode is useful for diagnostic comparisons of the stations, e.g. comparing station sensitivities, but can also be used for extremely wide-field surveys if one points each station in a different direction. In combination with the Complex Voltage sub-mode, Fly’s Eye can also be used to record the separate station voltages as input for offline fast-imaging experiments. It is also possible to simultaneously record a coherent and incoherent sum of all the stations used in this mode.




Fig. 11. Cumulative S/N map from a 127-beam tied-array observation of pulsar B2217+47, using the 6 HBA Superterp stations. These beams have been arranged into a “honeycomb” pattern in order to completely cover the 5.5° station beam FoV. The circle sizes represent the tied-array beam full-width half maxima (roughly 0.5°), and the color scale reflects the S/N of the pulsar in each beam.

These modes, and in some cases even a combination of these modes, can be run in parallel with the standard imaging mode described above. This allows one to simultaneously image a field while recording high-time-resolution dynamic spectra to probe sub-second variations of any source in the field (see for example Fig. 11 in [Stappers et al. 2011](#)).

10.3. Direct storage modes

Direct storage modes refer to observing modes that bypass the BG/P and deliver station data directly to the storage nodes of the offline cluster. These modes typically correspond either to triggered, short-term observations run in parallel with other observing modes, such as dumping the TBB boards following a CR or transient detection, or data taken by a single station in standalone mode. Types of data that can currently be stored in this manner include TBB data dumps using either full resolution or sub-band mode, station level beamformed data, and station level metadata. A variety of metadata are produced on the stations such as event triggers from the TBB boards (see Sect. 4.6) as well as diagnostic output from the calibration algorithm running on the LCU. Any or all of these data and metadata may be streamed directly from the stations to the storage nodes where they are incorporated into LOFAR standard data products. Once on the offline cluster, these data products can then be archived or further processed depending upon the scientific objective as with all LOFAR outputs. Example astronomical applications that utilize data from direct storage modes include all-sky imaging using intra-station baselines, single station observations of bright pulsars, dynamic

spectral monitoring of the Sun or planets, and the detection of CR air showers.

11. Processing pipelines

11.1. Standard imaging

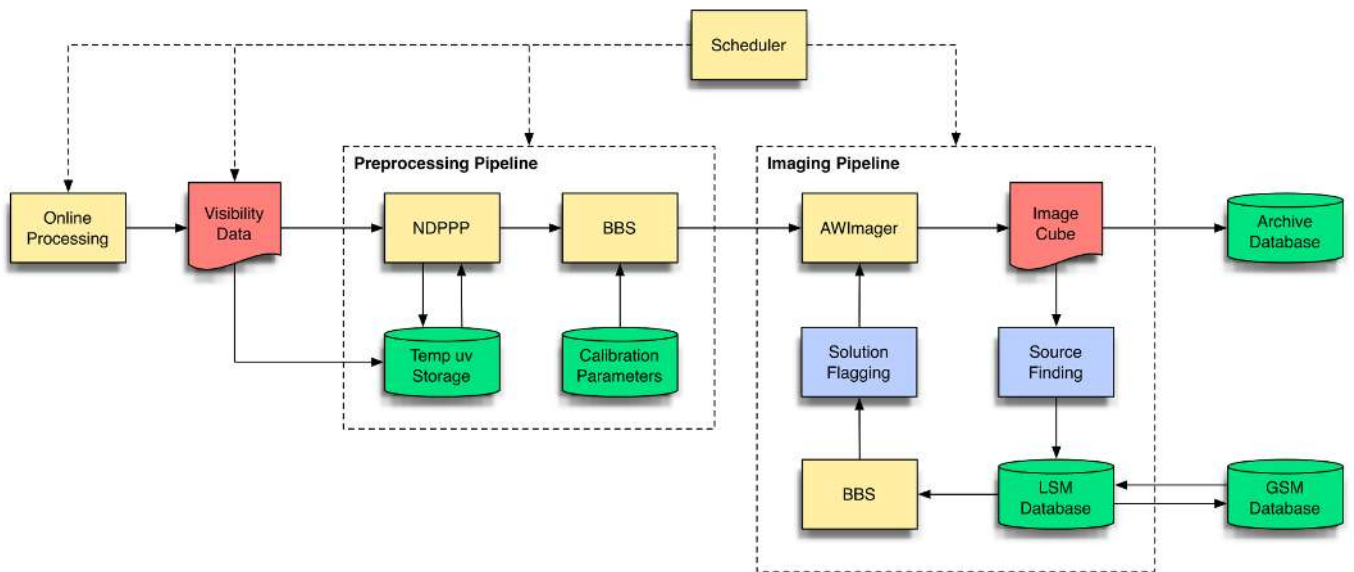
The standard imaging pipeline (SIP) is shown schematically in Fig. 12. A short overview of the pipeline is given by [Heald et al. \(2011\)](#), and a more in-depth description of the pipeline, its components, and the intermediate data products is in preparation ([Heald et al., in prep.](#)). Here, we outline the main pipeline features.

Following the data path from the left, visibility data are produced in the form of measurement sets at CEP, and recorded to multiple nodes in the LOFAR offline CEP cluster. The first standard data processing steps are encapsulated within a sub-pipeline called the pre-processing pipeline. Its role is to flag the data in time and frequency, and optionally to average the data in time, frequency, or both. The software that performs this step is labelled new default pre-processing pipeline, or NDPPP, and includes flagging using the `A0Flagger` routine (see Sect. 12.8).

This first stage of the processing also includes a subtraction of the contributions of the brightest sources in the sky (Cygnus A, Cassiopeia A, etc.) from the visibilities, using the demixing technique described by [van der Tol et al. \(2007\)](#) and implemented in NDPPP. Next, an initial set of calibration parameters is applied. In the current system, the initial calibration comes from

Table 4. Current LOFAR processing pipelines.

Pipeline	Mode	Inputs	Outputs	Description
Standard	Interferometric	Visibilities	Image cubes, source lists sky models, quality metrics	Limited angular resolution, full FOV
Long-baseline	Interferometric	Visibilities	Image cubes, source lists sky models, quality metrics	Highest angular resolution, limited FOV
Known pulsar	Beam-formed	BF data file	Folded pulse profiles de-dispersed time series	Arbitrary number of stations, 8 beams per station, full Stokes
CR event	Direct storage	TBB data file	CR characteristics event database	Single or multiple station event triggering
Transient detection	Interferometric	Image cubes	Source lists, light curves classifications, triggers	Can run in dedicated mode or commensal with other imaging observations

**Fig. 12.** The LOFAR imaging pipeline presented in schematic form (Heald et al. 2010). See the text for a description of the various software components and the data path.

an observation of a standard flux reference source (as characterized by Scaife & Heald 2012) which may have been performed in parallel with, or immediately preceding, the main observation. An initial phase calibration is achieved using the BlackBoard Selfcal (BBS) package developed for LOFAR.

The local sky model (LSM) used for the phase calibration is generated from the LOFAR Global Sky Model (GSM) that is stored in a database. The LOFAR GSM contains entries from the VLA Low-frequency Sky Survey (VLSS and VLSSr; Cohen et al. 2007; Lane et al. 2012), the Westerbork Northern Sky Survey (WENSS; Rengelink et al. 1997), and the NRAO VLA Sky Survey (NVSS; Condon et al. 1998) catalogs, and is being supplemented with entries from the Multifrequency Snapshot Sky Survey (MSSS, see Sect. 12.9). Finally, additional flagging and filtering operations (not shown in the figure) are performed in order to remove any remaining RFI or bad data.

Following the pre-processing stage, the calibrated data are further processed in the Imaging Pipeline, which begins with an imaging step that uses a modified version of the CASA imager (Tasse et al. 2013). This imager applies the w-projection algorithm (Cornwell et al. 2008) to remove the effects of non-coplanar baselines when imaging large fields and the new A-projection algorithm (Bhatnagar et al. 2008) to take into account the varying primary beam during synthesis observations.

Source finding software is used to identify the sources detected in the image, and generate an updated LSM. One or more “major cycle” loops of calibration (with BBS), flagging, imaging, and LSM updates are performed. At the end of the process, the final LSM will be used to update the GSM, and final image products will be made available via the LTA.

The Scheduler oversees the entire end-to-end process, from performing the observation through obtaining the final images. In addition to scheduling the observing blocks at the telescope level, it keeps an overview of the storage resources in order to decide where to store the raw visibilities. It also keeps an overview of the computational resources on the cluster, so that runs of the Pre-processing Pipeline and Imaging Pipeline can be scheduled and distributed over cluster nodes with available processing power.

11.2. Pulsar processing

The raw beam-formed data written by BG/P (see Sect. 10.2) are stored on the LOFAR offline processing cluster and LTA in the HDF5 format (hierarchical data format). The exact structure of these files as well as the metadata are fully described in the appropriate LOFAR interface control document (ICD) available from the LOFAR website.

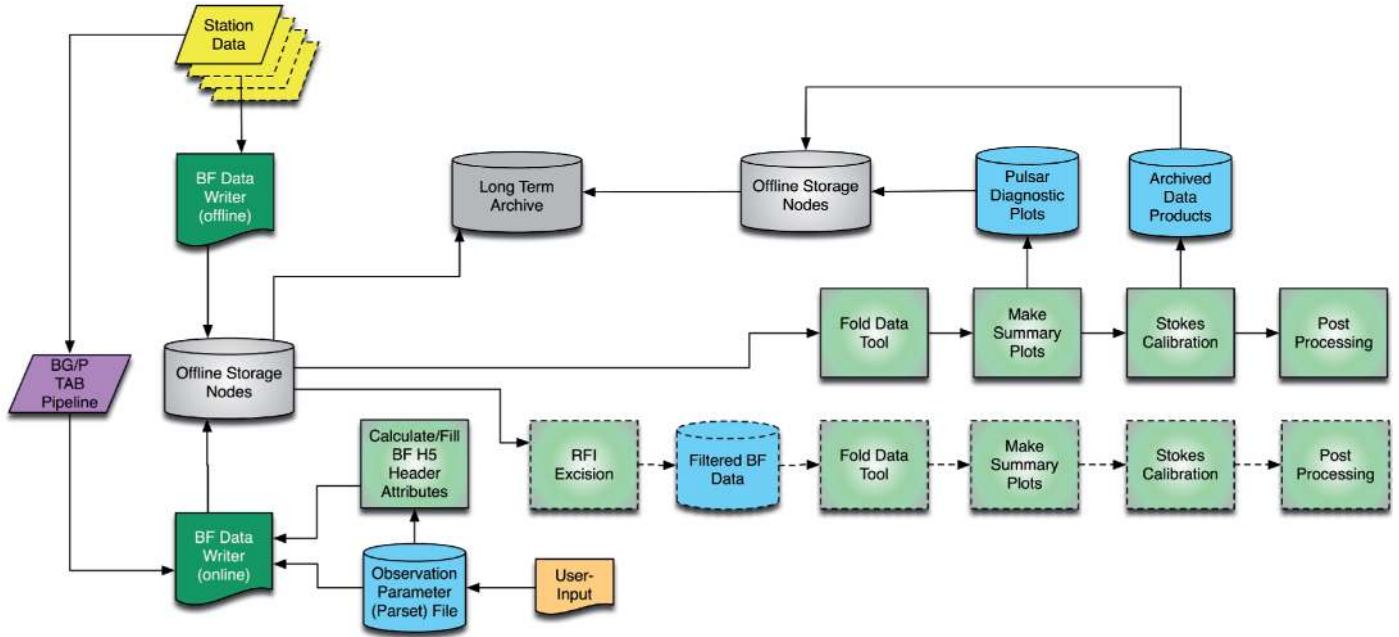


Fig. 13. Schematic diagram of the overall Pulsar Pipeline, as it runs online on the BG/P (see also Fig. 9) followed by offline scientific processing on the offline cluster. Offline pipeline processing can be run on data directly out of the BG/P or on RFI-filtered data.

Since the beam-formed data serve a variety of different science cases, several pipelines exist, e.g., to create dynamic spectra, search in real-time for fast transients, and for performing standard pulsar processing. The most advanced of these pipelines is the standard pulsar pipeline, “Pulp”, which is shown schematically in Fig. 13, and is described in more detail by Stappers et al. (2011). Pulp is currently implemented within a python-based framework that executes the various processing steps. The framework is sufficiently flexible that it can be extended to include other processing steps in the future.

Several conversion tools have been developed to convert these data into other formats, e.g. PSRFITS (Hotan et al. 2004), suitable for direct input into standard pulsar data reduction packages, such as PSRCHIVE (Hotan et al. 2004), PRESTO (Ransom 2011), and SIGPROC (Lorimer 2011). The long-term goal is to adapt these packages to all natively read HDF5, using the LOFAR Data Access Layer (DAL) for interpreting the HDF5 files. We have already successfully done this adaptation with the well-known program DSPSR (van Straten & Bailes 2011), which now natively reads LOFAR HDF5. Among other things, these reduction packages allow for RFI masking, dedispersion, and searching of the data for single pulses and periodic signals. Already, a test-mode exists to perform coherent dedispersion online, also for multiple beams/dispersion measures. Likewise, online RFI excision is also being implemented in order to excise corrupted data from individual stations before it is added in to form an array beam.

11.3. CR event processing

The high digital sampling rate of LOFAR (5 ns or 6.25 ns for the 200 MHz or 160 MHz clock, respectively) combined with the wide-field nature of its receivers make it a uniquely powerful instrument for the detection and study of CRs. Air showers of charged particles produced by CRs striking the Earth’s upper atmosphere can generate bright, extremely short duration radio pulses (Falcke & Gorham 2003). Depending on the energy and

direction of the incident CR, these pulses can be detected by the antennas in one or more LOFAR stations, as shown in Fig. 15.

Due to their short duration, in order to measure radio pulses from CRs, LOFAR must be triggered. When triggered the TBB RAM buffers in the station are frozen and the data are transferred directly to the CEP post-processing cluster (see Sect. 4.6 and Sect. 10.3). Such a trigger can be initiated in several ways. Either a pulse-finding algorithm is run on the FPGA and if a pulse is recorded by multiple dipoles simultaneously within a specified time window a dump is initiated. Alternatively, a dump of the TBB RAM buffers in a given station (or stations) can be triggered from the system level by triggers external to the station itself. These external triggers may come from outside LOFAR, as in the case of VOEvents from other observatories (see Sect. 9.3), or from within the LOFAR system.

As a CR produces its signal in the atmosphere a single CR pulse in an individual antenna does not largely differ from a RFI pulse. A large training set of detected CRs is needed in order to program the pulse-finding algorithm to only send a minimal amount of false triggers. In order to achieve this, one of the internal triggers is sent to LOFAR by an array of particle detectors, which is set up at the Superterp (Thoudam et al. 2011). These detectors trigger LOFAR only on CRs and also allow a cross-calibration of the measured characteristics of the air shower. Other types of internal triggers from for example other processing pipelines are foreseen. All CR detection modes place a strong constraint on the response time of the LOFAR system. The system must be able to process a detected pulse and freeze the contents of the TBB RAM buffer within a time interval which is smaller than the length of the buffer itself, otherwise the data from the event will be lost.

Following a trigger, raw voltage time series data from the TBB RAM buffers are stored in the HDF5 format using a data structure similar to the beam-formed data files mentioned in Sect. 11.2. The relevant ICD describing this TBB format is available from the LOFAR website. From the CEP post-processing cluster, these TBB data files are sent to the LTA (see

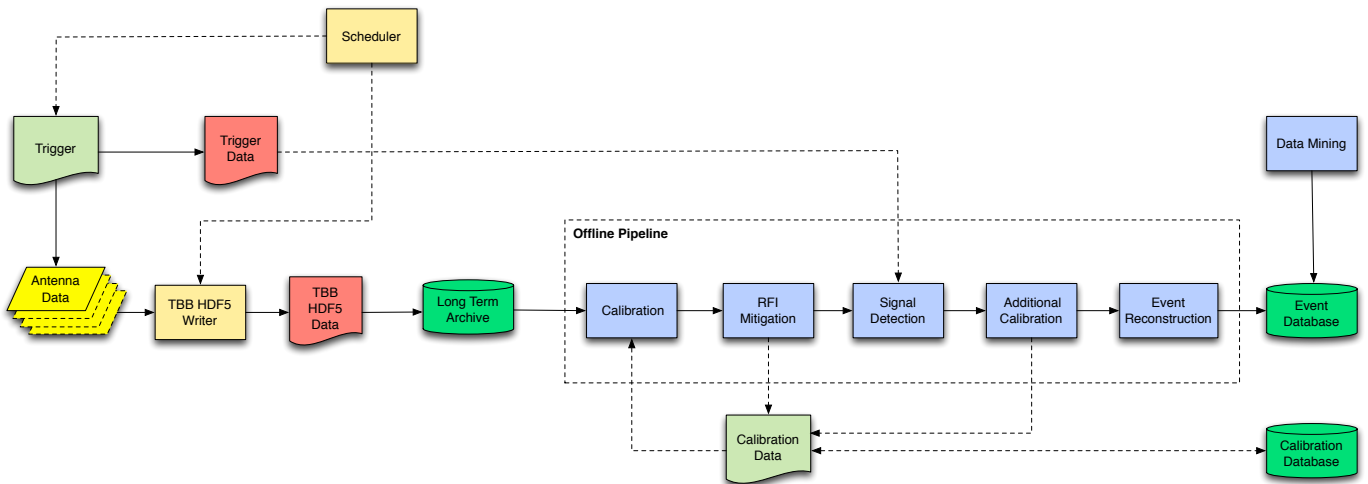


Fig. 14. Schematic view of the CR pipeline. The HDF5 data are the standardized output. The offline pipeline can be adapted to the purpose and type of the observation.




Fig. 15. Illustrative results of the CR pipeline. *Left:* CR pulse as recorded by one LBA antenna along with the reconstructed Hilbert envelope. The square of the Hilbert envelope corresponds to the sum of the squares of the original signal and the squares of the Hilbert transform. The Hilbert envelope is the amplitude of the analytic signal and essentially captures the amplitude of the pulse. *Right:* Hilbert envelopes for all antennas of one station ordered by their RCU number. One can clearly see the time delay of the air shower signal between different antennas as the antennas are numbered in a circular layout (Nelles et al. 2013).

Sect. 7) where they can then be accessed for offline processing as described in Fig. 14. The processing pipeline applies filtering and calibration corrections and characterizes the original CR event itself (in terms of direction and signal distribution). Finally the event information is stored in an SQL database in order to provide fast access for further study. The pipeline is implemented as a mixture of C++ libraries with Python bindings and Python scripts and accounts for the fact that the source is in the near-field, not at infinity as is assumed in other LOFAR processing pipelines. The pipeline can be run as a post-processing step performed automatically following any CR observations that result in data dumps from the TBBs, as well as interactively. The pipeline will be described in detail in a forthcoming publication.

Currently, the LOFAR system supports two types of CR observing modes differing only in whether the detection trigger is generated at the station level or by an external event at the system level. A number of additional modes are, however, envisioned for future development. These include modifications to the station level detection algorithm to tune the trigger mechanism to characteristic pulse profiles from different phenomena

such as lightning for example. Furthermore, this type of pulse search can be tuned to detect single dispersed pulses originating from fast radio transients like pulsars or other astronomical objects (Falcke 2008). This method has already been successfully tested at LOFAR by detecting a giant pulse from the Crab Nebula (ter Veen et al. 2012).

Similarly for pulses too faint to be detected by individual antennas, a trigger mechanism employing an anti-coincidence check between multiple on/off-source tied-array beams is envisioned. Such a trigger mechanism could in principle be used to detect faint radio flashes due to neutrinos interacting with the lunar regolith (see Sect. 13.5). Although much more sensitive, tied-array beam-based trigger mechanisms will also necessarily have more limited fields of view as opposed to dipole-based triggering that is essentially omnidirectional.

11.4. Transient detection

Beyond the pipelines already deployed as part of the operational LOFAR system, an additional science pipeline is currently

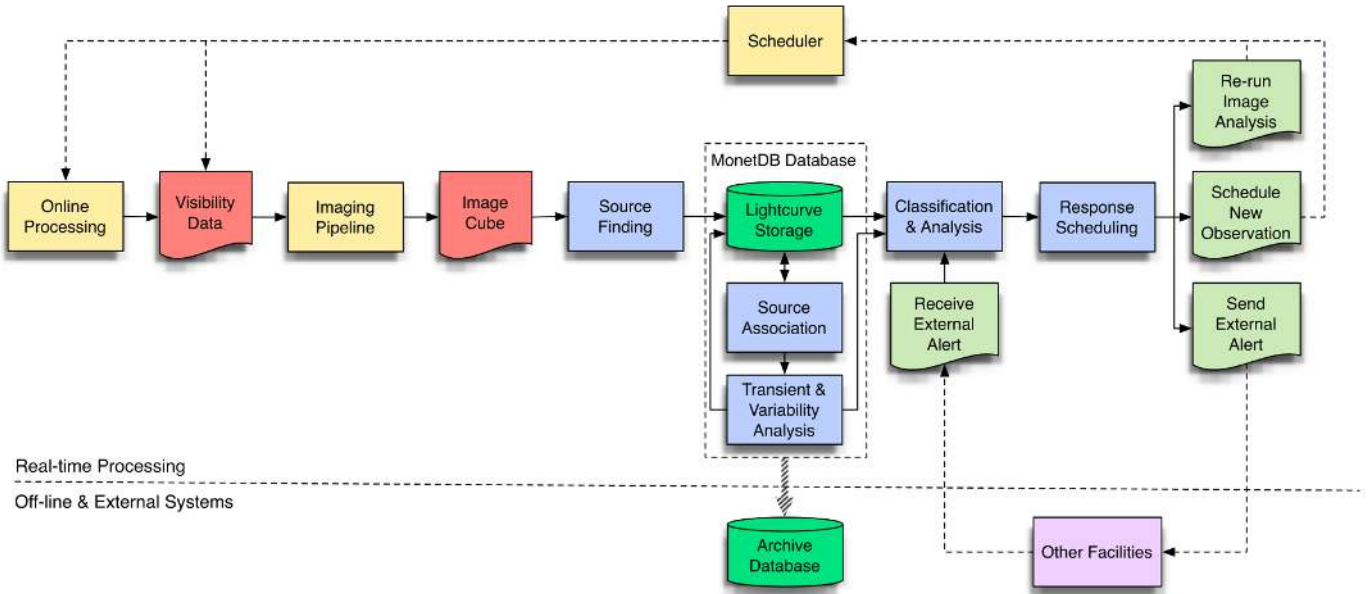


Fig. 16. Schematic outline of the main components in the LOFAR transients detection pipeline. Data are ingested from a modified version of the standard imaging pipeline (Sect. 11.1), while transients analysis is performed using a combination of custom source-finding and analysis routines and a high-performance MonetDB database.

under development tailored to detect transient and variable radio sources. The digital nature of the LOFAR system makes it inherently agile and an ideal instrument for detecting and, perhaps more crucially, responding to transient sources. Unlike the modes discussed previously, the transient detection pipeline will consist of a near real-time imaging pipeline that monitors the incoming stream of correlated data for both known variable sources and previously unknown transients. When a new source is detected, or a known source undergoes a rapid change in state, this mode will make it possible to respond on short timescales. Through a mixture of processing performance improvements and data buffering, the ultimate goal is to deploy a system capable of detecting radio transients down to timescales ~ 1 s with a response time latency of order ~ 10 s.

Possible responses include triggering actions within the LOFAR system such as switching to a different, targeted observational mode, adjusting the sub-band selection for an optimal frequency coverage, or dumping the data from the transient buffer; or, alternatively sending notifications to other observatories to initiate coordinated observations. LOFAR will also be capable of receiving and responding to triggers from external facilities in much the same way. This section presents a brief overview of the main components of the pipeline. A more comprehensive description is available in Swinbank (2011) and Swinbank et al. (in prep.).

An overview of the design of the transients detection pipeline is shown in Fig. 16. Image cubes produced by a variant of the standard imaging pipeline (Sect. 11.1) are ingested into the system, which identifies transients both by image plane analysis and by comparing the list of sources found in the images with those in previous LOFAR observations and other catalogs. Measurements from individual images are automatically associated across time and frequency to form light-curves, which are then analyzed for variability. Cross-catalog comparison, light-curve construction and variability analysis take place within a high-performance MonetDB database (Boncz et al. 2006; Ivanova et al. 2007). A classification system, based initially on simple,

astronomer-defined decision trees, but later to be augmented by machine learning-based approaches, is then used to identify events worthy of response. The primary data products of this mode will be rapid notifications to the community of transient events and a database of light-curves of all point sources observed by LOFAR (around $50\text{--}100$ Tbyte year^{-1}). In addition, snapshot images integrated over different time-scales as well as uv datasets suitably averaged in frequency will be archived.

High-speed response to transients is essential for the best scientific outcome so low-latency operation is therefore crucial in this mode. Achieving these low system latencies and response times may ultimately require adaptation of new, non-imaging algorithms for transient detection that rely on phase closure quantities (Law & Bower 2012; Law et al. 2012). Experiments employing these new algorithms are already underway as part of the commissioning process. If sufficiently low latency can be achieved, the TBBs (Sect. 4.6) can be dumped in response to a new transient, providing a look-back capability at the highest possible time, frequency and angular resolution. These requirements preclude human intervention, so all processing is fully automated. Efforts are also underway to minimize the time taken to transport and process data within the LOFAR system.

12. System performance

12.1. System stability

There are several effects that can lead to the deterioration of the phase and amplitude stability of LOFAR stations. For example, the ionospheric phase above the Netherlands often changes by one radian per 15 s in the 110–190 MHz band, and one radian per 5 s around 50 MHz on LOFAR NL baselines. Because typical ionospheric disturbances have scales of order ~ 100 km, the ionospheric phase on EU baselines fluctuates similarly. The GPS-corrected rubidium clocks at Dutch remote stations and most international stations can typically drift

up to 20 ns per 20 min, which corresponds to about a radian per minute at 150 MHz. This drift is much less than the ionospheric changes under solar maximum conditions, but comparable at solar minimum.

The HBA amplitudes are generally very stable. Although early experiments with prototype tiles showed up to 10–30% reduction in gain when the tiles were covered with a few cm of water and held in place by improvised edges, in practice these circumstances never occur in reality with the production tiles. The LBA antennas on the other hand, *are* sensitive to water under fairly normal operating conditions. If wet, and covered in water droplets, the resonance frequency can shift by several MHz, increasing the gain on one side of the peak by of order $\sim 10\%$, and decreasing the gain on the other side by a similar amount.

Fortunately, these effects are all station-based, hence easily corrected by self calibration given sufficient flux in the FoV and enough equations per unknown. LOFAR's tremendous sensitivity and large number of stations are therefore key. The MSSS surveys in both LBA and HBA clearly show that there is more than enough flux to calibrate within an ionospheric coherence time in the vast majority of fields, even during solar maximum (see Sect. 12.9 and Heald et al., in prep.). The EoR group has furthermore demonstrated image sensitivities better than $100 \mu\text{Jy}$ in the 110–190 MHz range, thereby demonstrating that system stability issues are not the limiting factor in achieving quality images (see Sect. 12.9).

12.2. *uv*-coverage

The image fidelity of an aperture synthesis array is dependent on how well the *uv*-plane is sampled. A poorly sampled *uv*-plane can result in strong side-lobes in the synthesized beam that will limit the overall dynamic range of an image. Also, since an interferometer samples discrete points in the *uv*-plane, incomplete *uv*-coverage can result in a loss of information on particular angular scales in the sky brightness distribution, which is important for imaging extended radio sources. For LOFAR, the *uv*-coverage has been optimized by choosing suitable locations for the stations throughout the Netherlands and by taking advantage of the large fractional bandwidth that is available. The positions of the international stations have not been chosen to maximize the filling of the *uv*-coverage, but care has been taken to avoid duplicate baseline lengths.

In Fig. 17, the *uv*-coverage for the completed LOFAR has been simulated using the known and expected positions of the 40 core and remote stations in the Netherlands and the 8 currently existing international stations. This simulation is based on a hypothetical 6 h observation of a radio source at declination 48° between 30 and 78 MHz and uses a single beam with a total contiguous bandwidth of 48 MHz. The *uv*-coverage for an array comprised of only the core stations, only the core and remote stations, and all of the LOFAR stations are shown. For clarity, the *uv*-distances are given in meters, since for *uv*-distances shown in λ the *uv*-coverage is densely sampled due to the large fractional bandwidth (~ 0.88 between 30 and 78 MHz; ~ 0.33 between 120 and 168 MHz). Also shown in Fig. 17 are the synthesized beams for each of the different array configurations between 30 and 78 MHz using uniform weighting, which show the side-lobe response pattern. The excellent *uv*-coverage results in first side-lobes that are $\sim 5\%$, $\sim 5\%$ and $\sim 7\%$ of the synthesized beam peak for the core, core and remote, and the full array, respectively. Similar values are obtained for a simulation that is carried out with the HBA frequencies between 120 and 168 MHz, and between 210 and 250 MHz. As normal,

sidelobe levels can be reduced at the expense of angular resolution through the use of other visibility weighting schemes.

The simulations above are for a standard long-track observation. The sensitivity and the large FoV of LOFAR will also allow surveys of the radio sky to be carried out efficiently using snapshot observations. The point source response of LOFAR in snapshot mode has also been simulated by calculating the instantaneous *uv*-coverage for the hypothetical observation described above, for a radio source at 0 h angle (transit). The resulting instantaneous *uv*-coverages for the core, core and remote, and full LOFAR arrays are also shown in Fig. 18. Note that for these simulations the data for the full 48 MHz bandwidth are presented, which highlights the excellent large fractional bandwidth of LOFAR. For the core, the *uv*-plane is well sampled, but for the core and remote, and for the full array, multiple snapshot observations over several hour-angles are needed to fill the gaps in the *uv*-coverage for the >5 km baselines.

12.3. Angular resolution

The ability to identify and characterize structures of different angular sizes depends on the angular resolution of an interferometric array. One of the transformational aspects of LOFAR is the unprecedented range of angular scales that are achievable at low observing frequencies. In Table B.2, and also shown in Fig. 19, the angular resolution for various baseline lengths as a function of frequency are given. These angular resolutions have been calculated using,

$$\theta_{\text{res}} = \alpha \frac{\lambda}{D}, \quad [\text{rad}] \quad (1)$$

where θ_{res} is the full width at half maximum (FWHM) of the synthesized beam in radians, λ is the observing wavelength, D is the maximum baseline length and α depends on the array configuration and the imaging weighting scheme (natural, uniform, robust, etc). The angular-resolutions given in Table B.2 are based on a value of $\alpha = 0.8$, corresponding to a uniform weighting scheme for the Dutch array. Note that this is for the ideal case of a source that has a maximum projected baseline length.

In reality, the angular resolution of an interferometric observation will be dependent on the declination of the source, the composition of the array, the observing frequency and the visibility weighting that is used. With baseline lengths ranging from a few tens of meters to over one thousand kilometers, the angular resolution of LOFAR extends from 0.5° to sub-arcsecond scales.

12.4. Bandpass

There are several contributions to the frequency dependent sensitivity of LOFAR to incoming radiation (the bandpass). At the correlator, a digital correction is applied within each 0.2 MHz subband to remove the frequency-dependent effects of the conversion to the frequency domain. The station beam is also strongly frequency dependent, except at the beam pointing center. Finally, the physical structure of the individual receiving elements (described in Sect. 4.2 and Sect. 4.3) causes a strongly peaked contribution to the bandpass near the resonance frequency of the dipole. In the case of the LBA dipoles, the nominal resonance frequency is at 52 MHz. However, as can already be seen in Fig. 5 and Fig. 20, the actual peak of the dipole response is closer to 58 MHz in dry conditions (see also Sect. 4.2). This shift in the peak is caused by the interaction between the low-noise amplifiers (LNAs) and the antenna.




Fig. 17. Sample uv coverage plots (*left*) and synthesized beams (*right*) including all present and planned LOFAR stations. The uv coverage is calculated for a source at declination 48° , and covers a 6 h track between hour angles of approximately -3 to $+3$ hours. One point is plotted every minute. Synthesized beams are calculated using uniform weighting, and using multi-frequency synthesis over the full LBA frequency range from 30–78 MHz. The *top frames* are for the 24 core stations only, *middle frames* include all 40 core and remote stations, and the *bottom frames* include all 48 core, remote, and international stations.




Fig. 18. As in Fig. 17, but illustrating the instantaneous uv coverage (near transit at the same declination) for the same complement of stations, and showing the effect of the full 48 MHz bandwidth from 30–78 MHz on the uv coverage. In the left frames, one point is plotted every 0.2 MHz.




Fig. 19. *Left:* full-width half-maximum (FWHM) of a LOFAR Station beam as a function of frequency for the different station configurations. The curves are labeled by the type of LOFAR antenna field, whether LBA or HBA, as well as by either core, remote, or international. The effective size of each field is also indicated in meters. *Right:* effective angular resolution as a function of frequency for different subsets of the LOFAR array.

Determining the combined or “global” bandpass can be achieved during the calibration step post-correlation. This global bandpass combines all frequency dependent effects in the system that have not already been corrected following correlation. To illustrate this point, the bright quasar 3C 196 has been observed using the core and remote stations in the LBA and all three HBA bands. 3C 196 is unresolved on the angular scales sampled by those baselines at LBA frequencies. 3C 196 has a known spectral energy distribution down to the lowest LOFAR observing frequencies (Scaife & Heald 2012). Moreover, 3C 196 is the dominant source in its field. Observations of 3C 196 between 15–78 MHz were obtained in two observing sessions (15–30 MHz in one session, and 30–78 MHz in the other). BBS was used to calibrate the data.

For each subband, a system gain was determined, and the gain amplitude was taken as the value of the global bandpass at the frequency of the particular subband. The median of all stations is shown in Fig. 20 for a typical 10 min of data. Curves are shown for all four LOFAR observing bands. While most stations exhibit individual bandpasses which are similar to the median value, some stations deviate significantly due to RFI or incomplete calibration information. In the future, station-level calibration information will be updated in near-realtime during observations.

12.5. Beam characterization

The response pattern, or beam, of a LOFAR observation is determined by the combination of several effects. The first is the sensitivity pattern of the individual dipoles themselves. This pattern changes relatively slowly across the sky. Electromagnetic simulations of the LBA and HBA dipoles have been performed, and parameterized descriptions of the results of these simulations form the basis of the dipole beam model used in the calibration of LOFAR data.

For imaging observations, the dominant effect in determining the sensitivity within the FoV (analogous to the “primary beam” of traditional radio telescopes) is the electronically

formed station beam. The size of the LOFAR primary beam determines the effective FoV for a given observation. The pointing of the station beam is determined by digital delays applied to the elements that make up an individual station. In the LBA, the elements are the dipoles themselves. In the HBA, groups of 16 dipole pairs are combined into HBA tiles. Each tile contains an analog beam former, which adds physical delay lines to each dipole and thus “points” the tile in a particular direction (as described in Sect. 4.3). The HBA station beam is formed from the combined signal from the tiles rather than directly from the dipoles. A description of the full HBA beam thus includes a term for the tile response pattern, which is of intermediate angular scale when compared to the dipole beam and the station beam.

Delays within the station (i.e. delays between individual dipoles or tiles) are calibrated by observing a bright source and determining the complex gain for each element that maximizes the response toward that bright source. These complex gains are stored as a calibration table at the station level and applied when forming the station beam. In future, it will be possible to control the shape of the station beam by applying a tapering function to the individual elements that are combined.

The nominal FWHM of a LOFAR station beam is determined using Eq. (1), where D is now the diameter of the station and the value of α will depend on the tapering intrinsic to the layout of the station, and any additional tapering which may be used to form the station beam. No electronic tapering is presently applied to LOFAR station beamforming. For a uniformly illuminated circular aperture, α takes the value of 1.02, and the value increases with tapering (Napier 1999). The FoV of a LOFAR station can then be approximated by

$$\text{FoV} = \pi \left(\frac{\text{FWHM}}{2} \right)^2. \quad (2)$$

An overview of the expected beam sizes for the various LOFAR station configurations is presented in Table B.1 and is also shown in Fig. 19. In the Dutch LBA stations, 48 dipoles must be selected out of the total 96. Selecting the innermost dipoles results

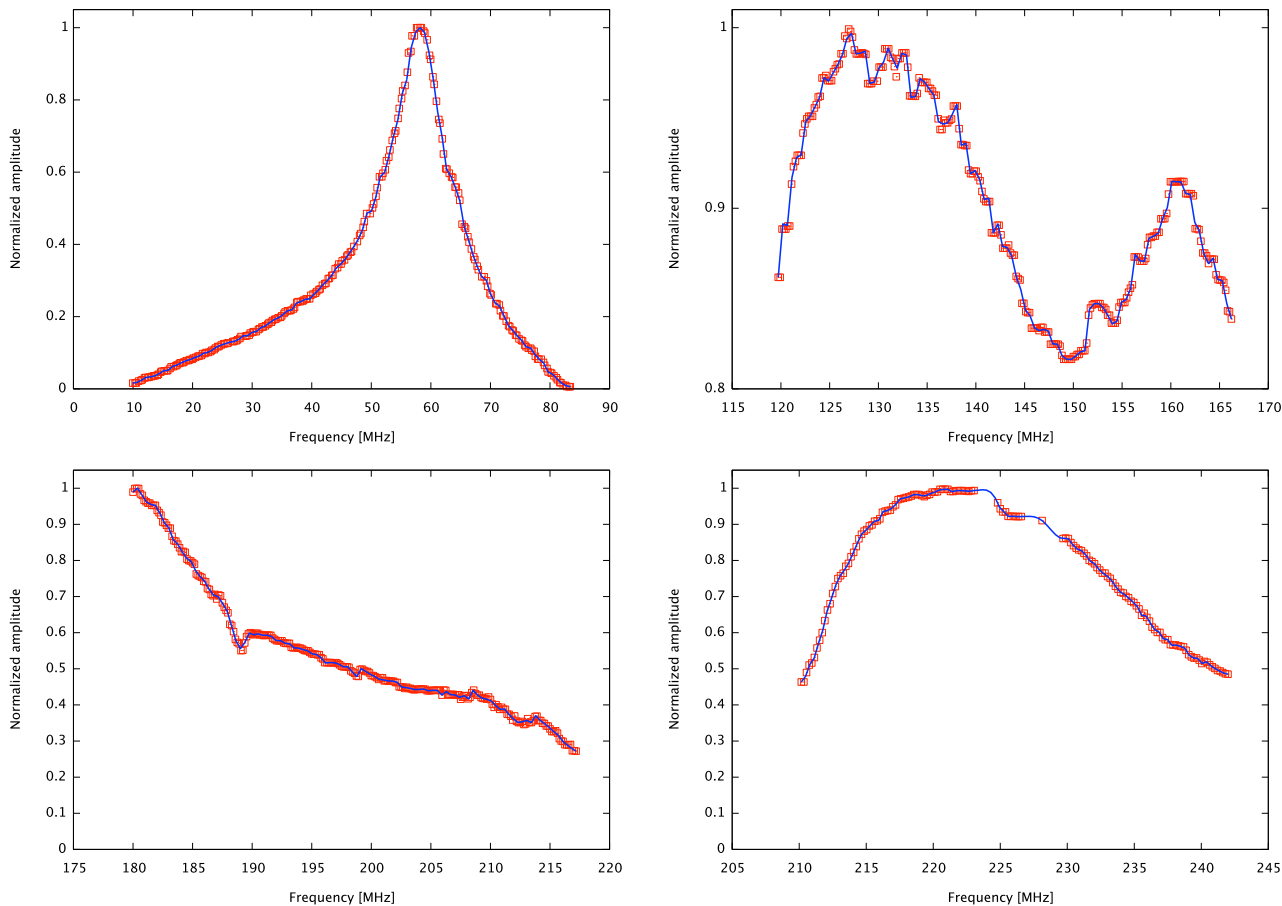


Fig. 20. Normalized global bandpass for several of the LOFAR bands. The global bandpass is defined as the total system gain converting measured voltage units to flux on the sky. These bandpass measurements were determined using short observations of the bright source 3C 196 and calculating the mean gain for all stations in each sub-band. The values have been corrected for the intrinsic spectrum of 3C 196 assuming a spectral index of -0.70 (Scaife & Heald 2012). Curves are shown for (*upper left*) LBA from 10–90 MHz with 200 MHz clock sampling, (*upper right*) HBA low from 110–190 MHz with 200 MHz clock sampling, (*lower left*) HBA mid from 170–230 MHz with 160 MHz clock sampling, and (*lower right*) HBA high from 210–250 MHz with 200 MHz clock sampling.

in a large-FoV configuration with a diameter of 32.25 m. Selecting the outermost dipoles results in a small-FoV configuration with a diameter of 81.34 m. The European stations always use all 96 dipoles in the low-band, which corresponds to a station diameter of 65 m. In the high-band, the core stations are split into two sub-stations, each with 24 tiles and a diameter of 30.75 m. The Dutch remote stations have 48 tiles and a diameter of 41.05 m. The European stations consist of 96 tiles and have a diameter of 56.5 m. In addition, individual stations are rotated relative to one another in order to suppress sensitivity in the sidelobes (see Sect. 4.1). As a result of this wide variation in station configurations and orientations, the beam modeling software is required to treat each station independently. The FoV of LOFAR imaging observations can range from ~ 2 –1200 deg², depending on the observing frequency and the station configuration.

We have observationally verified the station beam shapes and diameters using a strategy that takes advantage of LOFAR’s multi-beaming capabilities. A grid of 15×15 pointings, centered on Cygnus A, was observed simultaneously in interferometric imaging mode for 2 min at each of a sequence of frequencies in LBA and HBA. Calibration solutions were determined independently in each of the 225 directions to Cygnus A. Since Cygnus A is so bright, it dominates the visibility function in all grid points, and allows a good calibration solution. The influence of the distant bright source Cassiopeia A was overcome by using a

long solution interval (in both frequency and time). All baselines were used to determine the gain solutions, since removal of the long and/or short baselines was found to have no effect on the quality of the output.

The resulting gain amplitudes were mapped onto a complex gain pattern that was in turn used to derive a “power beam” (the square of the complex beam pattern) for each station. Examples of the power beams observed at 60 MHz and 163 MHz, for the core station CS004, are shown in the top panels of Fig. 21. Gaussians were fitted to vertical cuts through the center of the power beams, and resulted in FWHM values shown in the bottom panels of Fig. 21. By fitting Eq. (1) for α , we determined that the actual values are 1.02 ± 0.01 for HBA (core stations) and 1.10 ± 0.02 for LBA (in the LBA_INNER configuration). Since the LBA stations are less uniformly distributed than the HBA stations, the value of α is expected to be larger, as the observations confirm. We note that these values are only indicative since the stations are not circular apertures. The LOFAR processing software includes a beam model that directly computes the instantaneous station beam pattern for each station using the appropriate pointing direction and observing frequency.

Although the beam mapping observations discussed here were not specifically designed to carefully study the sidelobe pattern, they were sufficient to quantify the strength of the innermost sidelobes. Typical sidelobes levels of 20–25% were found




Fig. 21. The upper row of figures shows the observationally determined station beam patterns as described in the text. *On the left*, the beam is shown for core station CS004 in its LBA_INNER configuration. *On the right*, the beam is shown for a single HBA ear of the same core station. The peak response is normalized to unity in both plots. Note the low-level sidelobe structure apparent in the figures. *The bottom row* of plots gives the FWHM for a Gaussian fit to the main station beam lobe, plotted as a function of frequency. Solid lines indicate the fitted $\alpha\lambda/D$ relations explained in the text.

for both the LBA and HBA with structure consistent with a Bessel sinc function (see Napier 1999). A more detailed discussion of the beam structure can be found in Heald et al. (in prep.).

12.6. Sensitivity

Given estimates for the system equivalent flux density (SEFD) of a LOFAR station, one can calculate the expected sensitivity for different configurations of the array. The SEFD of a LOFAR station, in turn, depends on the ratio of the system noise temperature (T_{sys}) and the total effective area (A_{eff}) (Taylor et al. 1999). Since LOFAR consists of stations with different numbers of receiving elements, A_{eff} differs for the various types of stations and hence their SEFD also varies. The adopted values of the effective area, A_{eff} , were obtained from numerical simulations that account for the overlap of dipoles in the different station layouts and are provided in Table B.1.

To obtain empirical SEFD values for the Dutch stations, we have utilized 2-min imaging-mode observations of 3C295, taken near transit. The visibilities were flagged to remove RFI, and the contributions of Cygnus A and Cassiopeia A were modeled and removed (in the case of the LBA). From these pre-processed data, we determined the S/N of the visibilities for each baseline between similar stations (i.e., core-core and remote-remote baselines for the HBA). The S/N was defined as the mean of the parallel-hand (XX, YY) visibilities, divided by the standard deviation of the cross-hand (XY, YX) visibilities. These S/N values were then combined with the spectral model of 3C 295 from Scaife & Heald (2012), and taking the bandwidth and integration time of the individual visibilities into account, we directly obtain an estimate of the SEFD for the type of station comprising this baseline selection. The median contribution of all stations is plotted in Fig. 22. The most distant remote stations are excluded from this analysis as 3C 295 is resolved on all baselines to those




Fig. 22. *Top and bottom-left panels:* plots of the SEFD as a function of frequency for the various LOFAR operating bands and station configurations. These curves are derived using the procedure described in Sect. 12.6. The grayed regions are excluded from plotting due to strong post-flagging RFI contamination. In the case of HBA, the circles are for core stations and squares are for remote stations. In the LBA, the circles are LBA_INNER core stations and the squares are LBA_OUTER core stations. *Bottom-right panel:* contribution of sky temperature to the total system temperature of the LOFAR LBA system. Values are plotted separately for the LBA_INNER (circles) and LBA_OUTER (squares) station configurations. Empirical system temperature values and sky temperatures were determined as described in the text. The LBA system is clearly sky dominated at frequencies below ≈ 65 MHz. Above that frequency, the instrumental noise term dominates.

stations, thereby invalidating our S/N proxy. For the same reason, we have not attempted to determine empirical SEFDs for international stations using this procedure.

Starting with the empirical LBA SEFDs shown in the top-left panel of Fig. 22, we have derived the corresponding T_{sys} values, using $\text{SEFD} = 2760 T_{\text{sys}}/A_{\text{eff}}$ and the adopted values for the effective area, A_{eff} , given in Table B.1. To determine what fraction of the T_{sys} of the LBA system can be attributed to sky flux, we have compared our measured system temperatures with the standard equation from Thompson et al. (2007),

$$T_{\text{sky}} = 60 \lambda^{2.55}, \quad (3)$$

where T_{sky} is in K, and λ in meters. This expression corresponds to the average sky contribution. In the Galactic plane, the value of T_{sky} will be higher, and lower at the Galactic pole. The result is shown in the bottom-right panel of Fig. 22 and clearly illustrates that the LOFAR LBA system is sky-noise dominated below 65 MHz, in parts of the sky where the adopted sky spectrum is appropriate or an overestimate.

With these derived SEFD values, one can now compute the expected sensitivity of the array during a typical observation

(Taylor et al. 1999). In Table B.3, sensitivities are quoted for an 8 hour integration time and an effective bandwidth of 3.66 MHz (20 subbands) for the cases of a 6-station Superterp, a 24-station core array, a 40-station Dutch array, and a 48-station full array. The quoted sensitivities are for image noise and assume a factor of 1.3 loss in sensitivity due to time-variable station projection losses for a declination of 30 degrees, as well as a factor 1.5 to take into account losses for “robust” weighting of the visibilities, as compared to natural weighting. Note that this robustness factor is very strongly dependent on how the various stations, which all have different sensitivity, are weighted during the imaging process.

The values quoted for the HBA in Table B.3 agree with empirical values derived from recent observations on 3C 196 and the North Celestial Pole (NCP) where all NL remote stations were tapered to match 24-tile core stations. With improved station calibration, these estimates can likely be improved in the future by a factor of about 1.2. For the more compact LOFAR configurations, confusion noise will exceed the quoted values (see Sect. 12.7). The quoted sensitivities for the lower LBA frequencies have not yet been achieved in practice. At the lowest

frequencies below 30 MHz, values have not yet been determined awaiting a final station calibration. Similarly, the quoted values at 200, 210 and 240 MHz should be viewed as preliminary and are expected to improve with revised station calibration as well. For more recent values of the estimated array sensitivities and updates on the status of the station calibration, the reader is referred to the online documentation¹.

12.7. Confusion noise

The presence of faint, unresolved extragalactic sources in the synthesized beam produces “confusion” fluctuations in deep radio maps and represents a fundamental limit to the achievable sensitivity of a radio telescope. Confusion is normally said to occur when more than one source falls within the telescope beam and the classical confusion limit, σ_c , is defined as the flux density level where this condition is met taking into account the underlying population of faint sources. Formally, this condition can be written

$$M \Omega_b N(\sigma_c) = 1 \quad (4)$$

where $N(S)$ specifies the number of sources per steradian with a flux density greater than S and Ω_b represents the solid angle of the synthesized beam. The parameter M represents the number of beam solid angles per source and depends on the assumed form for the underlying distribution of sources.

In order to estimate σ_c , we first adopt a parameterization for $N(S)$ determined from the VLSS sky survey at 74 MHz (Cohen et al. 2007; Lane et al. 2012) and given by

$$N(>S) = A S^\beta \left(\frac{\lambda}{\lambda_0}\right)^{\alpha\beta} = 1.14 S^{-1.30} \left(\frac{\lambda}{4\text{ m}}\right)^{0.91} \quad (5)$$

where β represents the intrinsic slope of the underlying source distribution as a function of flux density S and α is the mean spectral index of a source at these wavelengths (Cohen 2004, 2006). Based on the VLSS catalog, Cohen (2006) estimates values of -0.7 and -1.30 for α and β , respectively. The normalization constant is $A = 1.14 \text{ Jy beam}^{-1}$ where the beam size is given in degrees.

Following Condon (1974), the solid angle for a Gaussian beam with FWHM, θ , is given by $\Omega_b = \pi\theta^2/[4 \ln(2)] \sim 1.133 \theta^2$. The final term, M , corresponds to the number of synthesized beams per source assuming a given flux density limit cutoff of $S = q\sigma_c$ and is given by $M = q^2/(2 + \beta)$ (Condon 1974). In the following, we have selected a cutoff of $q = 3$ yielding a value for $M = 12.8571$.

Combining these expressions with Equation 4, we can derive an expression for the expected confusion limit in the LOFAR band for different array configurations. Substituting these values, we obtain

$$\sigma_c = 30 \left(\frac{\theta}{1''}\right)^{1.54} \left(\frac{\nu}{74 \text{ MHz}}\right)^{-0.7} [\mu\text{Jy beam}^{-1}] \quad (6)$$

for the classical confusion limit. To put this expression in the context of LOFAR, for a frequency of 60 MHz near the peak of the LBA band, we would estimate values for σ_c of 150 mJy, 0.7 mJy, and 20 μJy for observations using the NL core, full NL array, and full European array, respectively. Similarly

for 150 MHz in the HBA band, we would estimate confusion limits of 20 mJy, 80 μJy , and 3 μJy for the core, NL, and international baseline configurations, respectively.

It is worth noting that these estimates rely on the VLSS source counts which are 100% complete down to flux density levels of only 1 Jy (see Fig. 15 in Cohen et al. 2007). The source distribution at much lower flux densities and lower frequencies may be significantly different than seen by the VLSS. Ultimately the source catalog produced by LOFAR’s first all-sky, calibration survey, the Multifrequency Snapshot Sky Survey (MSSS) (see Sect. 12.9 below and Heald et al., in prep.), will provide better constraints on the actual degree of source confusion in LOFAR images.

12.8. RFI environment

A possible concern with the construction of LOFAR in the high population density environment of the Netherlands and surrounding countries is terrestrial RFI in the local low-frequency radio spectrum. To overcome this, LOFAR has been designed to provide extremely high frequency- and time-resolution data during normal interferometric operations. The default frequency resolution is 610 or 763 Hz (each subband is subsequently divided into 256 channels), depending on the clock setting, and the typical visibility integration times are either 1 s in the low-band (10–80 MHz) or 3 s in the high-band (120–240 MHz). Even though 256 channels are available, in practice typical observations are performed using only 64 channels per subband. This choice lowers the resulting data volume by a factor of 4 without additional loss of data due to RFI flagging.

The flagging of the full resolution data in both time and frequency is carried out using the AOFlogger, a post-correlation RFI mitigation pipeline developed by Offringa et al. (2010, 2012a,b). This routine uses an iterative method to determine the true sky brightness by applying a high-pass filter to the visibility amplitudes in the time–frequency plane. Subsequently, it flags line-shaped features with the SumThreshold method (Offringa et al. 2010). Finally, the scale-invariant rank operator, a morphological technique to search for contaminated samples, is applied on the two-dimensional flag mask (Offringa et al. 2012b). Additional developments, for example, pre-correlation RFI mitigation, will be incorporated into the LOFAR analysis routines in the future.

As an example of the RFI environment of LOFAR, the percentages of RFI that have been identified and removed from 24 h datasets taken with the LBA and HBA-low systems are shown in Fig. 23. For the low-band system, the median level of RFI is estimated to be around 2% of the data, although this increases to around 10% at the lowest frequencies. These values represent the maxima per sub-band since within a sub-band any given channel can be 100% contaminated. For some sub-bands in the 30–80 MHz range, the median level of RFI can spike to as high as 7 to 20% of the data. For the HBA-low case, the median level of RFI that is identified over a 24 h period is around 1% of the data. However, there are several frequencies that show RFI spikes between 5 and 17% of the data across the band. For the HBA system, the median baseline RFI is again around 1 to 2% of the data, but there are also significant broadband RFI signals that range from 5 to above 50% of the data. In general, the level of RFI that is identified and removed from LOFAR datasets during the commissioning phase is not severe over the standard 30–80 MHz and 110–240 MHz observing frequencies in which LOFAR operates. Additional results of the automated flagging algorithm are shown for a more extreme case of RFI contamination in Fig. 24.

¹ See <http://www.astron.nl/radio-observatory/astronomers> for current updates on the calibration status of the LOFAR array, including up-to-date estimates for the achievable sensitivities.




Fig. 23. Dynamic spectrum of RFI occupancy during the LBA and HBA survey. The median baseline RFI is around 1 to 2% across the bands, although there are regions of the spectrum with significant narrow and broad band RFI features (Offringa et al. 2013).

12.9. Image quality

The calibration step is performed using BlackBoard Selfcal (BBS). This calibration package is based on the Hamaker-Bregman-Sault measurement equation (ME; see Hamaker et al. 1996; Sault et al. 1996; Hamaker & Bregman 1996; Hamaker 2000, 2006; Smirnov 2011) which expresses the instrumental response to incoming electromagnetic radiation within the framework of a matrix formalism. Here, the various instrumental

effects are identified, their effect on the signal is characterized in full polarization, and are quantified and parameterized as separate Jones matrices. Each of these terms may depend on different dimensions: frequency (e.g. the bandpass); time (e.g. the station gains); or direction (e.g. the station beam). Because it is based on the general form of the ME, BBS can natively handle difficult problems such as direction dependent effects and full polarization calibration, using parameterized models based on the physics of the signal path.

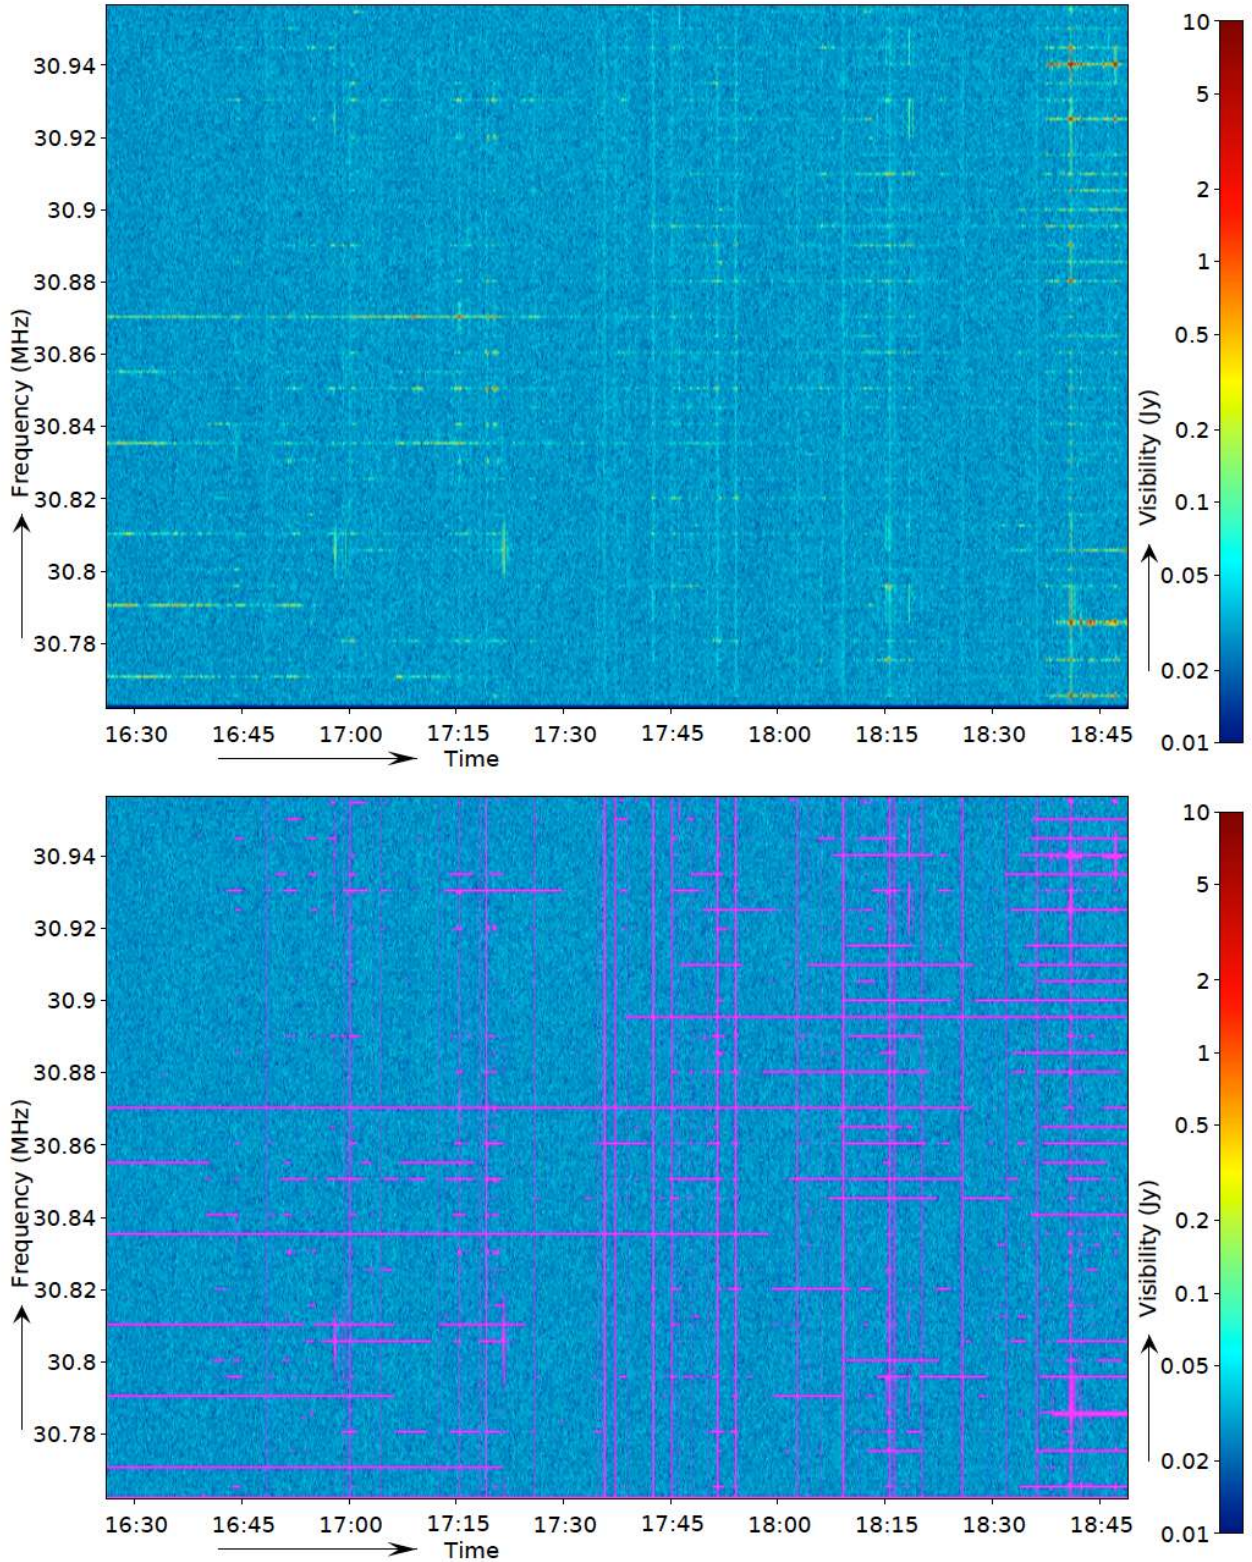


Fig. 24. Dynamic spectrum of data from one sub-band of the LBA survey, formed by the correlation coefficients of baseline CS001–CS002 at the original frequency resolution of 0.76 kHz. The displayed sub-band is one of the worst sub-bands in terms of the detected level of RFI. *The top image* shows the original spectrum, while *the bottom image* shows with purple what has been detected as interference (Offringa et al. 2013).

A critical input to BBS is the sky model that is used to predict the visibilities. Early in the commissioning process, this input to the BBS stage of the pipeline was a hand-crafted listing of the brightest sources in the field of interest. The current SIP (see Sect. 11.1) automatically constructs an initial sky model

based on cataloged values from the VLSS, WENSS, and NVSS. Note that this should be considered the “Mark-0” LOFAR GSM; the “Mark-1” LOFAR GSM is being generated by the Multi-frequency Snapshot Sky Survey (MSSS). MSSS is a broadband survey of the northern ($\delta > 0^\circ$) sky, using multiple simultaneous




Fig. 25. LOFAR wide-field image of the area around M51. This image represents a 6 h integration using 34 MHz (204 sub-bands) of bandwidth centered on 151 MHz taken with the HBA. The noise in the image is $\sigma \sim 1$ mJy beam⁻¹ in the vicinity of M51 and $\sigma \sim 0.6$ mJy beam⁻¹ away from bright sources with an effective beamsize of 20'' (Mulcahy et al., in prep.).

station beams to increase the survey speed. MSSS provides a higher areal density of sources than the VLSS catalog, and more importantly includes well-sampled spectral information in 16 bands spanning 30 MHz to 160 MHz. The primary goal of the survey is to provide a broadband catalog of the brightest population of sources in the LOFAR sky, creating a low-frequency calibration database for future imaging observations. MSSS observations began in autumn 2011 and were nearly half completed during 2012. A detailed description of the survey setup, data processing, and results is in preparation (Heald et al., in prep.).

The quality of images produced by LOFAR's interferometric imaging mode is dependent on many factors. Novel techniques must be brought to bear in order to achieve imaging at high dynamic range and fidelity over a large FoV. There are many factors that may limit the achievable dynamic range (DR) in LOFAR images. In fields where there are no unusually bright sources (see below) the main limiting factors are direction-dependent effects, namely issues related to variable beam response as a function of time, and the ionosphere. The former is being handled by the inclusion of a comprehensive beam modeling library in both the calibration and imaging software, while techniques to address the latter are based on the method used by Intema et al. (2009). Polarization calibration will include the prediction

(to within ~ 0.1 rad m⁻²) and application of ionospheric rotation measure values as described by Sotomayor-Beltran et al. (2013). Fig. 25 shows a relatively wide-field ($\sim 4 \times 4$ square degrees) HBA image of the field surrounding the bright galaxy M51 (Mulcahy et al., in prep.).

A large number of commissioning observations have now been obtained by the LOFAR EoR project team on fields containing the bright compact sources 3C 196 and 3C295, which have a flux density of 100 Jy at 115 and 144 MHz, respectively. The dynamic range achieved in these fields exceeds more than 500 000:1 at distances at least several arcmin from the sources. In the neighbourhood of these sources, the dynamic range is currently still restricted to 10 000–100 000:1, depending on the PSF used, and appears to be limited only by our imperfect knowledge of the (sub-)arcsecond structure of the sources themselves. We note that this knowledge will likely improve in the very near future with the inclusion of structural information obtained using LOFAR's international baselines. The correlator itself, therefore, does not appear to introduce any errors that limit the LOFAR's achievable dynamic range.

Although the dynamic ranges already achieved in images of select LOFAR fields are impressive, the more relevant number characterizing the quality of LOFAR images is the achievable

noise level as given in Table B.3, with the caveats listed in Sect. 12.6. The many factors influencing the actual noise in real observations were already listed above. In practice, deep observations (3 nights, of 6 h each) of the NCP field have reached noise levels of $100 \mu\text{Jy}$ or better corresponding to a factor of only 1.4 above the thermal limit set by the noise from our Galaxy and the receivers (Yatawatta et al. 2013). For more complicated fields of course, the noise levels can be higher.

The uv -coverage of the LOFAR array is also designed to provide excellent imaging of extended sources. When combined with its high sensitivity, LOFAR can deliver high quality images of faint, diffuse objects (van Weeren et al. 2012; de Gasperin et al. 2012). For example, Fig. 26 shows an HBA image of the diffuse emission associated with the bright AGN Cygnus A obtained during the commissioning phase (McKean et al., in prep.). During the commissioning period, the available calibration and imaging software has been shown to deliver on-axis image rms levels near the expected thermal noise. Predicted values for the achievable sensitivities are given in Table B.3. Significantly deeper images are achievable by utilizing the full bandwidth (Yatawatta et al. 2013). A discussion of the trade-off between sensitivity and resolution is given by Heald et al. (2011).

12.10. Beam-formed modes

LOFAR's beam-formed modes share many of the same system requirements imposed by the interferometric imaging mode, but they also have some unique requirements of their own. As such, beam-formed observations provide a complementary, and sometimes orthogonal, means with which to test both generic and mode-specific system performance (see Hessels et al. 2010; Stappers et al. 2011, for several examples).

Wide-band observations of continuum sources like pulsars are well-suited to measuring the instrumental bandpass and overall sensitivity (e.g., Fig. 10). They also serve as important polarimetric calibration sources, both using interferometric imaging and beam-formed data. For example, Fig. 32 shows the rotation measure spread functions (RMSF) for two pulsars in the HBA and LBA bands.

The microsecond to millisecond time resolution typically used in the beam-formed modes also probes a different RFI regime than that apparent from the >1 s time resolution used to record visibilities. RFI occupancy histograms for LOFAR beam-formed data can be found in Stappers et al. (2011).

The various possible beam shapes (e.g. element beam, station beam, and tied-array beam) can also be mapped by the beam-formed mode by scanning across relatively bright point sources (e.g. see Stappers et al. 2011). This knowledge can then be used as an input for imaging calibration. Alternatively, using LOFAR's multi-beam tied-array survey mode, one can map the instantaneous beam in two dimensions. Figure 27 shows the result of an observation in which 217 simultaneous tied-array beams were pointed in a honey-comb pattern around the phase center in order to map the shape of the coherent Superterp beam. The observed beam shape agrees well with that predicted by a simple model that takes into account station sizes and positions.

Clock calibration is vital to interferometric imaging, but these corrections can largely be made in the post-processing. For tied-array observations, however, the instrumental, geometrical, and environmental (e.g. ionospheric) delays must be applied in real time in order to form a properly coherent sum of the station beams (e.g., see Stappers et al. 2011, and Fig. 11 and 10 for demonstrations of tied-array beams using the LOFAR

Superterp). The implementation of a single clock signal shared between all 24 Core stations has greatly simplified this process.

Furthermore, for applications like long-term, phase-coherent pulsar timing, the clock reference standard needs to be maintained between observations and systematic offsets and changes to the observatory's time standard need to be well-documented. LOFAR's station-level Rubidium clocks are guided by local GPS receivers, and thus the long-term LOFAR time standard is automatically in sync with this system. Comparison of the time stamps between observations is greatly simplified by the use of a consistent geographical reference. For LOFAR the phase center of all observations is the geographical center of the LBA field of station CS002, regardless of whether this particular station is being used or not (see Table A.1). This reference position is used, e.g., for barycentering, pulsar timing, and phasing-up the array. Figure 28 shows an example of phase-coherent timing of a millisecond pulsar using LOFAR.

13. Key science drivers

13.1. Epoch of reionization

The formation of the first stars and galaxies marks a major transition in the evolution of structure in the Universe. These galaxies with their zero-metallicity Population-III and second-generation Population-II stars and black-hole driven sources (e.g., mini-quasars, X-ray binaries, etc.) first heated and subsequently transformed the intergalactic medium from neutral to ionized. This period is known as the Cosmic Dawn and epoch of reionization (EoR). Observing and quantifying this poorly observed and little understood process is the main aim of the LOFAR EoR Key-Science Project (KSP).

The last thirty years have witnessed the emergence of an overarching paradigm, the Λ CDM model, that describes the formation and evolution of the Universe and its structure. The Λ CDM model accounts very successfully for most of the available observational evidence on large-scales. According to this paradigm about 400 000 years after the Big Bang ($z \approx 1100$), the temperature and density decreased enough to allow ions and electrons to recombine and the Universe to become neutral. As a result, the Universe became almost transparent leaving a relic radiation, known as the cosmic microwave background (CMB) radiation (for recent results see the WMAP papers, e.g., Spergel et al. 2007; Page et al. 2007; Komatsu et al. 2011). The matter-radiation decoupling has ushered the Universe into a period of *darkness* as its temperature dropped below 3000 K and steadily decreased with the Universe's expansion. These *Dark Ages* ended about 400 million years later, when the first radiation emitting objects (stars, black-holes, etc.) were formed and assembled into protogalaxies during the Cosmic Dawn (CD).

The most accepted picture on how the cosmic dawn and reionization unfolded is simple. The first radiation-emitting objects heated and subsequently ionized their immediate surroundings, forming ionized bubbles that expanded until the neutral intergalactic medium consumed all ionizing photons. As the number of objects increased, so did the number and size of their ionization bubbles. These bubble gradually percolated until eventually they filled the whole Universe.

Most of the details of this scenario, however, are yet to be clarified. For example: what controlled the formation of the first objects and how much ionizing radiation did they produce? How did the bubbles expand into the intergalactic medium and what did they ionize first, high-density or low density regions? The answers to these questions and many others that arise in the




Fig. 26. Example of LOFAR single sub-band (0.2 MHz bandwidth) imaging of the radio galaxy Cygnus A with the HBA system at 150 MHz (*top*) and the LBA system at 74 MHz (*bottom*), made using 24 core stations and 9 remote stations during the commissioning phase (McKean et al., in prep.). Both datasets consist of 12 h synthesis observations. These images show the expected combination of compact and extended structure that has previously been seen in this source at these frequencies using the VLA and MERLIN, cf. with the images of [Lazio et al. \(2006\)](#) and [Leahy et al. \(1989\)](#), respectively. The beam-size of the images are 5.7×3.5 arcsec and 11.7×7.4 arcsec, respectively, and are shown as the white ellipses in the bottom left corner of each image. The dynamic ranges are ~ 3500 and ~ 2000 for the 74 MHz and 150 MHz maps, respectively.




Fig. 27. Maps of the coherent Superterp beam, also known as a tied-array beam. These were made by simultaneously recording 217 coherent and 1 incoherent beam for all 12 Superterp HBA sub-stations from 120–168 MHz. The bright pulsar B0329+54 was observed twice for 5 min, near zenith and the color maps reflect the S/N of the dedispersed and folded signal in various directions. The background color reflects the S/N of the simultaneously acquired incoherent beam. *Left:* observation L57554 in which the tied-array beams were arranged in a hexagonal grid with spacing 0.05° . This densely samples the main lobe of the Superterp beam. *Right:* observation L57553 in which the tied-array beams were more widely spaced (0.15° apart) in order to probe the sidelobes. Asymmetry in the sidelobe pattern is due to imperfect phasing of the coherent beam.

context of studying the CD and EoR touch upon many fundamental questions in cosmology, galaxy formation, quasars and very metal-poor stars; all are foremost research topics in modern astrophysics. A substantial theoretical effort is currently dedicated to understanding the physical processes that triggered this epoch, governed its evolution, and the ramifications it had on subsequent structure formation (cf., Barkana & Loeb 2001; Bromm & Larson 2004; Ciardi & Ferrara 2005; Choudhury & Ferrara 2006; Furlanetto et al. 2006; Zaroubi 2013). Observationally however, this epoch is poorly studied. Still the current constraints strongly suggest that the EoR roughly straddled the redshift range of $z \sim 6$ –12 (Komatsu et al. 2011; Fan, et al. 2003, 2006; Bolton & Haehnelt 2007; Theuns et al. 2002; Bolton et al. 2010; Oesch et al. 2010; Bunker et al. 2010).

It is generally acknowledged that the 21-cm emission line from neutral hydrogen at high redshifts is the most promising probe for studying the *Cosmic Dawn* and the EoR in detail (Field 1958; Madau et al. 1997; Ciardi & Madau 2003). HI fills the IGM except in regions surrounding the ionizing radiation of the first objects to condense out of the cosmic flow. Computer simulations suggest that we may expect an evolving complex patch work of neutral (HI) and ionized hydrogen (HII) regions (Gnedin & Abel 2001; Ciardi et al. 2003; Whalen & Norman 2006; Mellema et al. 2006; Zahn et al. 2007; Mesinger & Furlanetto 2007; Thomas et al. 2009; Thomas & Zaroubi 2011; Ciardi et al. 2012).

LOFAR with its highly sensitive HBA band is the best available instrument to-date to probe this process from $z = 11.4$ (115 MHz) down to $z = 6$ (203 MHz). At lower frequencies, both the sensitivity of LOFAR drastically decrease and the sky noise dramatically increase making it very hard to use the telescope for such a measurement. Given the very low brightness temperature and the angular scale of the expected EoR signal, the only part of LOFAR that has the sensitivity to detect the EoR redshifted 21-cm signal is the LOFAR core. The LOFAR core stations give a resolution of about 3 arcmin over a FoV, given by the HBA LOFAR station size, of about 5° corresponding, at $z = 9$, to ≈ 8

and 800 comoving Mpc, respectively. A number of fields (~ 5) with minimal Galactic foreground emission and polarization will be observed for a total of several thousands of hours, reaching brightness temperatures of 50–100 mK per resolution element per MHz bandwidth, close to that of the redshifted 21-cm emission from the EoR. Ultimately, the EoR KSP hopes to achieve a noise level of approximately 60 mK per resolution element per MHz after 600 h.

Studying the power-spectra as function of redshift (or frequency) allows us to probe the EoR as it unfolded over cosmic time. The EoR power spectrum can be observed over about two orders of magnitude in wave numbers and other higher-order statistical measures can be obtained as well (Jelić et al. 2008; Harker et al. 2009, 2010; Labropoulos et al. 2009). It might even be possible to image the EoR as it unfolds on very large scales after several thousand hours of integration time on a single field (Zaroubi et al. 2012). Finally, using total-power measurements, LOFAR might also be able to probe the total (i.e. global) intensity signal from neutral hydrogen to even higher redshifts with the LBA, complementing interferometric measurements at lower redshifts with the HBA. The LOFAR EoR KSP is currently investigating, using the LOFAR LBA system in different beam-forming modes, whether the system is suitable to detect, or place stringent upper limits, on the global redshifted 21-cm signal from the Cosmic Dawn around $z \sim 20$.

Summarizing, LOFAR observations will allow detection and quantification of the Cosmic Dawn and EoR over wide range in angular scales and redshifts. Such measurement will help answer the main questions surrounding the earliest phases of the formation of the Universe: What is the nature of the first objects that ended the Dark Ages, ushering in the Cosmic Dawn and the reionization of the high-redshift IGM? What is the relative role of galaxies and AGN, of UV-radiation and X-rays? When did the EoR start and how did it percolate through the IGM? Which regions re-ionized first: low or high density regions (inside-out versus outside-in scenario)? What are the detectable imprints that the re-ionization process left on the 21-cm signal? What can we




Fig. 28. Pre-fit timing residuals from a 1-month LOFAR campaign on PSR J0034-0534. The individual time-of-arrival measurements reflect the deviations from a Westerbork-derived timing model for this pulsar (which includes astrometric, spin, and orbital parameters), with no additional re-fitting of these parameters. The measurements have an rms of $11 \mu\text{s}$, and it is possible that the deviations of individual points reflect small changes in the dispersion and/or scattering measure of the pulsar. In any case, these data clearly demonstrate LOFAR’s capability to do precision pulsar timing.

learn from 21-cm measurements about the matter density fluctuations on the conditions prior to the EoR? What can we learn about the formation of (supermassive) black holes and the duration of their active phases?

The LOFAR EoR KSP plans to address all these questions over the next years, playing an important role as well in paving part of the way for future more sensitive observations of both the Cosmic Dawn and EoR with the SKA.

13.2. Surveying the low-frequency sky

An important goal that has driven the development of LOFAR since its inception is to explore the low-frequency radio sky through several dedicated surveys. The main science driving the design of these surveys will use the unique aspects of LOFAR to advance our understanding of the formation and evolution of galaxies, AGNs and galaxy clusters over cosmic time. Since LOFAR will open a new observational spectral window and is a radio “synoptic” telescope, the surveys will explore new parameter space and are well-suited for serendipitous discovery.

Furthermore, a carefully designed and easily accessible LOFAR data archive will provide the maximum scientific benefit to the broader astronomical community.

Due to LOFAR’s low operating frequencies, and the resultant large beam size on the sky, this radio telescope is an ideal survey facility. For example, at 50 MHz, each beam typically has a FoV of 7–8 deg. With theoretical LOFAR sensitivities and feasible observing times, such a field will typically contain 1 radio galaxy at $z > 6$, 5 Abell clusters, 5 NGC galaxies, 5 lensed radio sources and several giant (>1 Mpc) radio galaxies. The aimed legacy value of the LOFAR surveys will be comparable to previous high-impact surveys (e.g. Palomar, IRAS, SDSS, GALEX, *Spitzer*, NVSS) and will also complement currently planned surveys in other wavebands (e.g. JEDAM, Euclid, Pan-STARRS, *Herschel*, *Planck*, VISTA, VST, JVLA, ASKAP, MeerKAT, ATA). The surveys described here will provide meter-wave data on up to 10^8 galaxies and 10^4 clusters out to $z \sim 8$ and will address a wide range of topics from current astrophysics. The LOFAR survey key Project has, from the outset, been driven by four key topics. The first three are directly related to the formation of massive black holes, galaxies and clusters. The fourth

is the exploration of parameter space for serendipitous discovery. The four key topics are:

- (1) *High-redshift radio galaxies* (HzRGs, $z > 2$) are unique laboratories for studying the formation and evolution of massive galaxies, rich clusters and massive black holes (see review by Miley & de Breuck 2008). Presently, the most distant HzRG has a redshift of $z = 5.1$ (van Breugel et al. 1999). However, due to its low operating frequency, LOFAR will detect about 100 radio galaxies at $z > 6$, enabling robust studies of massive galaxies and proto-clusters at formative epochs, and provide sufficient numbers of radio sources to probe structure in the neutral IGM near or even within the EoR through HI absorption studies.
- (2) *Clusters of galaxies* are the most massive gravitationally bound structures in the Universe, and drive galaxy evolution through mergers and interactions. However, approximately 40 clusters are known to also contain Mpc-sized, steep spectrum synchrotron radio sources that are not clearly associated with individual cluster galaxies. These are classified either as radio halos or radio relics, depending on their location, morphology and polarization properties (Ferrari et al. 2008). The LOFAR surveys will allow detailed studies of about 15 local clusters in unprecedented detail, detect about 100 clusters at $z \gtrsim 0.6$, and will contain thousands of diffuse cluster radio sources out to $z \sim 1$. These surveys will enable the characteristics of the magnetic fields (strength, topology) in clusters to be determined and test models for the origin and amplification of these fields. Also, the origin and properties of the CR acceleration and evolution within clusters will be studied in detail.
- (3) Determining the *cosmic star-formation history* of the Universe is a key goal of the LOFAR surveys, the deepest of which will detect radio emission from millions of regular star-forming galaxies at the epoch when the bulk of galaxy formation occurred. The combination of LOFAR and infrared surveys will yield radio-IR photometric redshifts, enabling studies of the volume-averaged star formation rate as a function of epoch, galaxy type and environment. These studies will cover a sky area large enough to sample diverse environments (from voids to rich proto-clusters) and over a wide range of cosmic epochs.
- (4) One of the most exciting aspects of LOFAR is the potential of *exploring new parameter space for serendipitous discovery*. The uncharted parameter space with the highest probability of serendipitous discovery is at frequencies < 30 MHz, where the radiation mechanisms being probed are not observable at higher radio frequencies, such as coherent plasma emission.

These four key topics drive the areas, depths and frequency coverage of the LOFAR surveys. In addition to the key topics (1–4), the LOFAR surveys will provide a wealth of unique data for the following additional science topics; (5) detailed studies of AGN and AGN physics, (6) AGN evolution and black hole accretion history studies, (7) observations of nearby radio galaxies, (8) strong gravitational lensing, (9) studies of the cosmological parameters and large-scale structure, and (10) observations of Galactic radio sources. Furthermore, to maximise the usefulness of the survey data for the Magnetism key project, the LOFAR survey data will be taken with sufficient bandwidth so that the technique of rotation measure synthesis can be applied. In collaboration with members of the Transient key project, the survey observations will be taken in several passes, to facilitate searches for variable sources on various timescales.

To achieve these science goals, a three-tier approach to the LOFAR surveys has been adopted, using five different frequency setups. For each part of the surveys a minimum total bandwidth of 24 MHz will be used to improve the uv -coverage through multi-frequency-synthesis, as well as offering a significant benefit for polarization studies.

1. **Tier 1: The “large area” 2π steradian surveys:** These shallow wide area surveys will be carried out at 15–40, 40–65 and 120–180 MHz, and will reach an rms of 2, 1 and 0.07 mJy, respectively. It is expected that up to 3×10^7 radio sources will be detected by these three surveys, including, ~ 100 cluster halos at $z > 0.6$ and ~ 100 radio galaxies at $z > 6$ (cf. Enßlin & Röttgering 2002; Cassano 2010; Cassano et al. 2010). The sensitivity and multi-frequency nature of the wide-area surveys will allow the low-frequency spectral shape of distant galaxy candidate sources with at least $\alpha = -2.0$ to be measured.
2. **Tier 2: The “deep” surveys:** These surveys will be carried out at the same frequencies as the shallow all-sky surveys, but will be substantially deeper and over a smaller sky area. The HBA part will cover around 550 square degrees to provide a representative volume of the Universe. To maximise the additional science, the pointings will be centered on 25 well-studied extragalactic-fields that already have excellent degree-scale multi-wavelength data, 15 fields centered on clusters or super-clusters and 15 fields centered on nearby galaxies. The HBA survey will reach an rms of $15 \mu\text{Jy}$ at 150 MHz, which will be sensitive enough to detect galaxies with a star formation rate $SFR > 10$ and $> 100 M_{\odot} \text{yr}^{-1}$ (5σ) out to $z = 0.5$ and 2.5 , respectively (Carilli & Yun 1999). The LBA part of the deep survey will be over 1000–1500 square degrees and will reach a depth of 0.3–1.0 mJy. The pointings will be centered on 6 of the best-studied extragalactic fields and 9 of the most important nearby galaxies and clusters.
3. **Tier 3: The “ultra-deep” survey:** Finally, there will be 5 fields which will be observed with the HBA covering a sky-area of 83 square degrees. This part of the survey will be ultra-deep, reaching an rms of $7 \mu\text{Jy}$ at 150 MHz. This sensitivity will be sufficient to detect 50 proto-clusters at $z > 2$, and detect galaxies with a SFR of 10 and $100 M_{\odot} \text{yr}^{-1}$ at $z \sim 1.5$ and ~ 5 , respectively. These sensitivities are similar to that needed for the EoR Project.

The LOFAR surveys will not only be unique due to their low frequencies, but will also reach 2–3 orders of magnitude deeper in sensitivity than existing large-sky radio surveys, as is illustrated in Fig. 29. They will permit a wide range of science goals to be attained and provide a legacy value data set for studies of the low frequency radio sky.

13.3. The transient radio sky

LOFAR’s ability to image very wide fields with good sensitivity, and to eventually do so in nearly real time (see Sect. 11.4), opens up a very large discovery space in time-domain astronomy. The known and suspected transients already span a very large range of properties. On the shorter timescales, coherent emitters are the only ones reaching detectable fluxes: Jupiter’s radio outbursts reach fluxes of thousands of kJy, with substructure down to below milliseconds, rich polarization variations and narrow structures in frequency (Zarka 2004). Giant pulses of regular radio pulsars can reach upwards of 100 Jy over microseconds, and millisecond single pulses of RRATS are at a wide range of flux




Fig. 29. Flux limits (5σ) of the proposed LOFAR surveys compared to other existing radio surveys. The triangles represent existing large area radio surveys. The lines represent different power-laws ($S \sim \nu^\alpha$, with $\alpha = -1.6$ and -0.8) to illustrate how, depending on the spectral indices of the sources, the LOFAR surveys will compare to other surveys.

levels up to 1 Jy. Stellar radio flares have similarly rich structures in polarization, time and frequency as Jupiter, but at lower fluxes and with durations from minutes to hours. Elusive Jupiter-like signals from exoplanets, still to be discovered, are expected to be much weaker (Zarka 2011, and references therein).

At longer timescales, a wide variety of jet sources produce incoherent synchrotron emission with a large range of variability timescales: Galactic microquasars have outbursts that may last from days to months, but with rise times and substructures that can be very much shorter, at flux levels from hundreds of Jy down to the mJy level. AGN flares typically have very much longer time scales due to the scaling of black-hole phenomena with mass. On the patient end of the range, the variability of radio supernovae and gamma-ray bursts (GRBs) at LOFAR frequencies is measured in years to decades, with peak fluxes below a mJy in many cases; here the challenge changes from rapid response and high-volume fast data processing to careful analysis of deep images, and good use of supporting data from other instruments to tell the different types of slow radio transient apart.

Besides its potential as a discovery tool, the fully electronic operation of LOFAR makes it an excellent followup response machine for rapidly variable phenomena (see Sect. 9.3). LOFAR's electronic repointing capability enables it to start a completely new observation (new settings of observing mode and pointing direction) in well under a minute, and to do so fully automatically upon receipt of external triggers from any telescope using, e.g., VOEvent protocols (Seaman et al. 2011). It can thus play a prominent role in the emerging network of

wide-field sky monitors at many wavelengths, as well as in the multi-messenger world of gravity-wave and particle telescopes.

That there is still much to discover in the transient radio sky may perhaps be best illustrated by some recent examples of radio transient discoveries that have already greatly expanded the range of known phenomena. With some of these discoveries has come the realization that many previously known extreme objects also emit radio flares –such as those associated with giant flares of Soft Gamma Repeaters. In other cases, radio transient searches have produced serendipitous finds of completely new types of object. For example, one of these is the discovery of so-called RRATs (McLaughlin et al. 2006), apparently pulsars that only emit a pulse of radio emission once per very many rotation periods. Another example is the discovery of an enigmatic radio source close to a supernova remnant near the Galactic Center (Hyman et al. 2005). This source emits 10 min bursts in the radio with a very precisely constant 77 min period (Spreeuw et al. 2009). It is only detected in about 10% of the previous attempts to observe it; however, emphasizing the importance of long-term monitoring of such objects. Most strange perhaps was the discovery of a millisecond dispersed radio burst from near the direction of the Small Magellanic Cloud (Lorimer et al. 2007), which –if astrophysical in origin at all– certainly implies most strange and extreme compact-object astrophysics. Strange transients discovered at other wavelengths (e.g., the puzzling Swift J1955 – Castro-Tirado et al. 2008) will often require radio observations to help elucidate their nature.

One of the dominant problems in making progress in understanding what these strange objects are, and by what mechanism

they radiate, is the great sparsity of observations we have of each one, and the fact that only one or a few sources of each class are known. LOFAR's Transient capability will do much to remedy both of these: its very wide FoV combined with good sensitivity will make it likely that many representatives of these classes of object will be found in systematic transient surveys. The trigger and response capability will reliably provide fast (within seconds to minutes) radio data on sources newly discovered by other telescopes.

Due to the effects of dispersion, the signal from a fast transient in the range of 0.1–10 s will be spread over a large bandwidth in the LOFAR frequency range making it more difficult to detect. In such cases, fast-imaging modes will naturally give way to beam-formed modes as the method of choice for exploring the transient radio sky due to the increased sensitivity albeit at the expense of angular resolution. Consequently our standard transient imaging pipelines will target 1 s as the shortest timescale to detect and characterize. Since crude dedispersion is also possible on sets of narrow-band images at least when dispersion is still modest, the optimal boundary between the two techniques will have to be explored once a specific set of algorithms and computing platforms is in place.

13.4. Pulsar studies and surveys

Although pulsars were discovered at 82 MHz (Hewish et al. 1968), the majority of pulsar studies have been at frequencies >300 MHz, and often at ~1.4 GHz, because effects in the interstellar medium (ISM, e.g. dispersion and scattering), coupled with the Galactic synchrotron background and the steep power-law spectra of most pulsars, combine to make these frequencies well suited for the study of typical radio pulsars. Nonetheless, pulsar observations in the LOFAR frequency range of 10–240 MHz are also very interesting for addressing some long-standing issues about the pulsar emission mechanism, and for studying the ISM. Furthermore, LOFAR's high sensitivity, flexible beam-formed observing modes, multi-beaming, and large FoV are well-suited for pulsar and fast transient searches (e.g. Fig. 11). Here we give a very brief overview of the expected studies of known pulsars and searches for new pulsars and fast transients. More details and early LOFAR pulsar results can be found in Stappers et al. (2011), Hassall et al. (2012), Hessels et al. (2010), van Leeuwen & Stappers (2010), and Hermsen et al. (2013).

LOFAR will study the pulsar radio emission mechanism by providing wide-bandwidth, low-frequency spectra at high time resolution. It is believed that the power-law spectra of most pulsars turn over somewhere in the 10–240 MHz frequency range, making LOFAR an ideal instrument to study this important aspect of the pulsar emission mechanism. The roughly 4 octaves of frequency coverage provided by LOFAR allow very detailed studies of profile morphology as a function of observing frequency (see Fig. 30), e.g. the so-called “radius-to-frequency mapping” phenomenon (Cordes 1978; Hassall et al. 2012). Most of the known radio pulsars in the northern hemisphere will be detectable with LOFAR (by summing many hundreds of individual pulses), and we expect to detect single pulses from one-third of the visible pulsars in the high-band and ten percent of visible pulsars in the low-band. Though low radio frequencies are poorly suited to precision timing tests, LOFAR will allow the frequent monitoring of many pulsars to look for timing anomalies such as glitches (e.g. Espinoza et al. 2011) and sudden profile and spin-down changes (Lyne et al. 2010; Kramer et al. 2006). There is also the possibility that the radio emission from some neutron

stars may *only* be detectable at the lowest radio frequencies (e.g. PSR B0943+10 Deshpande & Radhakrishnan 1994).

Targeted surveys of, e.g., nearby galaxies, globular clusters, supernova remnants, and the large population of γ -ray sources recently found with *Fermi* are likely to find interesting new radio pulsars. These sources have a small extent on the sky and can be observed using either one or a few tied-array beams simultaneously (for reference a tied-array beam made from the Superterp/entire LOFAR core has a FWHM of 30/5 arcmin). Efficient all-sky surveys have already begun and can be done either using hundreds of tied-array beams (which provides high sensitivity and excellent source location, but produces a large data rate) or with the incoherent sum of the station beams (lower raw sensitivity and poorer source localization, but the single beam FoV is 5.5 deg across and the data rate is low). These surveys will also search for generic fast transients (e.g., Burke-Spolaor et al. 2011), and aim to eventually trigger on transient bursts in real time in order to dump the TBBs for better source localization.

13.5. Astroparticle physics

CR atomic nuclei and electrons have been detected with various methods on Earth, and indirectly in our Galaxy as well as other galaxies from their electromagnetic signature in gamma rays down to the radio domain, e.g. in supernova remnants (SNRs), radio pulsars, protostars, planetary magnetospheres, X-ray binaries, jets in radio galaxies and quasars, active galactic nuclei (AGN), and GRBs.

Over a wide range of energies (E) the primary CR flux follows a simple power law $dN/dE \propto E^{-\gamma}$, as shown in Fig. 31. At 10^{11} eV about one particle per second per square meter hits the Earth on average. This number changes to approximately one particle $\text{yr}^{-1} \text{m}^{-2}$ at 5×10^{15} eV, and above 10^{19} eV only about one particle per century per square kilometer hits the Earth. These low fluxes require experiments with large effective areas in order to collect sufficient statistics.

At an energy of $E \sim 5 \times 10^{15}$ eV the spectrum shows a turn-over where the power law index γ changes from $\gamma \approx 2.7$ to $\gamma \approx 3.1$. This feature is called the knee of the CR spectrum. Up to the knee in the spectrum the composition of the primary CRs is dominated by protons, but at higher energies the composition still needs clarification (Blümer et al. 2009; Kampert & Unger 2012). The question about the composition of these ultra-high energy CRs will be crucial for the understanding of acceleration and propagation mechanisms. At the highest energies above 10^{19} eV there is a flattening of the spectrum, the so-called ankle which could be caused by the Greisen-Zatsepin-Kuz'min (GZK) effect (Greisen 1966; Zatsepin & Kuz'min 1966). Ultra High Energy protons above $\sim 5 \times 10^{19}$ eV lose their energy quickly by producing pions in collisions with photons from the CMB. This effect accumulates protons that had been accelerated to higher energies at energies below the reaction threshold, and implies that any observed CR of this energy finds its origin in the near Universe (<50 Mpc).

Because of the smoothness of the spectrum, much effort has gone into identifying a universal acceleration process. It is believed that diffusive shock acceleration – a first-order Fermi-type acceleration process – is this universal mechanism. It operates in strong collisionless shocks such as occur in a multitude of explosive objects in the Universe and produces a differential power law spectrum in energy with power law index of -2 , close to and somewhat flatter than is observed, for any shock as long as it is both strong and non-relativistic. Up to the knee, diffusive shock acceleration in SNRs is believed to be the main




Fig. 30. A 1 h LBA observation (L77295) of pulsar B0809+74 using a coherent addition of all 24 LOFAR core stations from 15–93 MHz. The data has been dedispersed and folded using a rotational ephemeris to produce a cumulative pulse profile as a function of frequency. Given that the central observing frequency is 53.8 MHz, the fractional bandwidth is 145%. This wide bandwidth is key to following the drastic evolution of the cumulative profile with frequency. At the bottom of the LBA band there are two distinct pulse components that almost completely merge toward the top of the band. Such data are being used to constrain properties of the emitting region in the pulsar magnetosphere (Hassall et al. 2012).

acceleration process. Above the knee, possible candidate sources of high energy CRs are shocks in radio lobes of powerful radio galaxies, intergalactic shocks created during the epoch of galaxy formation, magnetars, so-called hyper-novae, and GRBs. So far no conclusive evidence has been found that clearly identifies the source of the highest energy CRs.

The identification of the sources of CRs is not only hindered by the low statistics of events measured at Earth, but also by the lack of knowledge of the Galactic and intergalactic magnetic fields. CRs will propagate a considerable amount of time through the Galaxy and intergalactic space before finally reaching Earth. Magnetic fields of different strengths and degrees of turbulence will obscure their original direction. As this effect is energy dependent, there is hope that the most energetic particles will still indicate their sources and their paths will then provide information about the magnetic field structure. But also in reverse: an improved knowledge about the magnetic fields in the Universe will help to solve the open questions about the origin of the CRs.

When a CR hits the nucleus of an atom in the terrestrial atmosphere it undergoes a nuclear reaction and produces several secondary particles. These secondary particles again react with atmospheric nuclei and produce more secondary particles.

Together these particles form an extensive air shower. If the energy of the primary particle was high enough this air shower can be measured at ground level. The highest energies observed fall outside the domain which is currently being studied with Earth bound particle accelerators. Thus air showers not only are messengers from the distant Universe but also form a laboratory to study new particle physics (The Pierre Auger Collaboration 2012).

LOFAR will observe CRs above 10^{16} eV up to 10^{19} eV from their bright radio flashes in the terrestrial atmosphere. These flashes are caused by the deflection of particles in the Earth magnetic field and charged processes within the development of the air shower. The theories explaining this phenomenon have developed rapidly during the last ten years, e.g. Huege & Falcke (2005), Ludwig & Huege (2011), Scholten et al. (2008) or Werner et al. (2012). They indicate that the radio emission will also be sensitive to the height of the shower development and thereby able to identify the particle type of the primary CR. However, to really confirm the predictions, data of higher quality and abundance is needed. The high numbers of antennas at LOFAR are essential to measure every shower in highest possible detail.




Fig. 31. High energy end of the spectrum of the CR flux as measured by a number of current experiments. The flux has been multiplied by a factor of $E^{2.5}$ to better show features in the spectrum, which are related to acceleration and propagation mechanisms. The gray bars indicate the energy range in which LOFAR will be sensitive to CRs. Furthermore, the energy ranges of other experiments detecting radio emission of CRs are shown. Among those are the Lofar Prototype Station (LOPES; Falcke et al. 2005; Apel et al. 2012) and the Auger Engineering Radio Array (AERA; Kelley & The Pierre Auger Collaboration 2011).

In addition to conclusively explaining the exact mechanisms of these radio emissions, the observations with LOFAR will aim to answer a number of fundamental questions in astroparticle physics such as the composition and origin of these particles. The measurements of the radio emission of air showers will be complemented with LOFAR observations of the Moon to detect (or put upper limits to) Cherenkov flashes from CRs from 10^{21} to 10^{22} eV in the lunar regolith (LOFAR Nu Moon). For a more detailed description refer to Scholten et al. (2009), Buitink et al. (2010) and Mevius et al. (2012). Together with LOFAR observations of supernova remnants and other explosive events in the Universe, the study of magnetic fields, and, ultimately, with radio observations done at the Pierre Auger Observatory in Argentina at (well- calibrated) energies above a few 10^{18} eV one aims at a full picture of CR physics.

13.6. Magnetic fields in the universe

Understanding the Universe is impossible without understanding magnetic fields. Magnetic fields are present in almost every place in the Universe but in spite of their importance the evolution, structure and origin of magnetic fields all remain open fundamental problems. Most of our knowledge of astrophysical magnetic fields has come from radio-frequency observations of synchrotron radiation from relativistic cosmic-ray leptons (mostly electrons). These observations trace the total field strength (from the synchrotron intensity), the orientation and degree of ordering of fields in the plane of the sky (from the polarized component of the radiation), and the component of ordered fields along the line of sight (via Faraday rotation).

LOFAR's exceptionally wide bandwidth at low frequencies is extremely useful for the study of magnetic fields, for several

complementary reasons: (i) it provides excellent leverage on the spectral characteristics of the synchrotron radiation, which allows study of the synchrotron losses of the emitting electrons; (ii) low energy synchrotron-emitting electrons are detectable *only* at low frequencies, so LOFAR can uniquely trace magnetic fields far away from CR acceleration sites; and (iii) studies of Faraday rotation have the best precision when the range of measured wavelength is wide – LOFAR will thus trace weak magnetic fields (Beck 2010). A very powerful tool for detection and characterization of polarized emission with LOFAR will be the Rotation Measure Synthesis (RM Synthesis; Brentjens & de Bruyn 2005) technique. This provides the Faraday dispersion function, or, in short, the Faraday spectrum (Fig. 32) that gives information about the structure of the magneto-ionic medium along the line of sight.

The magnetism key science project (MKSP) aims to investigate cosmic magnetic fields in a variety of astrophysical sources, including an initial target list of galaxies, followed by deep observations of galaxies and galaxy groups. These deep fields will also serve as targets to investigate magnetic fields in the Milky Way foreground. The structure of small-scale magnetic fields will be studied the lobes of giant radio galaxies. Polarized synchrotron emission and rotation measures from pulsars and polarized jets from young stars will be observed.

At high latitudes above the Milky Way plane, LOFAR will be uniquely sensitive to synchrotron emission of low-energy electrons in the Galactic halo, which will allow investigations of the propagation and evolution of matter and energy far from the Galactic disk. Weak magnetic fields or small density fluctuations of thermal electrons will become visible through Faraday rotation, leading to a better understanding of the turbulent ISM, and allowing a three-dimensional model of the gas and magnetic fields in the solar neighborhood to be constructed.




Fig. 32. Faraday dispersion functions (FDFs, or Faraday spectra) obtained from LOFAR observations of the polarised pulsar B0950+08 (normalized absolute value). *Left:* 27-min LBA tied-array beam-formed observation using coherent addition of the six LOFAR Superterp stations with a center frequency of 56 MHz, 10 MHz bandwidth, MJD 55 901, and using data from obsID L36787. *Right:* 10-min HBA tied-array beam-formed observation using coherent addition of 20 LOFAR core stations with a center frequency of 150 MHz, 90 MHz bandwidth, MJD 56 260, and using data from obsID L78234. The narrow FWHM of the functions allows the peaks associated with the pulsar (2.373 ± 0.011 and 2.136 ± 0.061 rad m⁻², respectively) and instrumental response (~ 0 rad m⁻²) to be individually resolved, despite the very low absolute rotation measure (RM). These RMs were corrected for ionospheric Faraday rotation (0.899 ± 0.042 and 0.665 ± 0.059 rad m⁻², respectively) using the `ionFR` code which employs International GNSS Service vertical total electron content (VTEC) maps (Hernández-Pajares et al. 2009) and data from the International Geomagnetic Reference Field (Finlay et al. 2010) (see Sotomayor-Beltran et al. 2013). The resulting RM of the ISM toward B0950+08 was determined to be 1.47 ± 0.04 and 1.47 ± 0.08 rad m⁻², from LBA and HBA observations respectively. These results are significantly more precise and in good agreement with the value of 1.35 ± 0.15 rad m⁻² previously measured (Taylor et al. 1993).

With present-day radio telescopes, GHz synchrotron emission from electrons in a $5 \mu\text{G}$ magnetic field can be detected in external galaxies, or a $1 \mu\text{G}$ field in clusters. The minimum detectable magnetic field strength varies with $\nu^{-\alpha/(3-\alpha)}$ (where α is the synchrotron spectral index, $\alpha \simeq -0.8$) so that all else being equal, observing at a $10\times$ lower frequency permits the detection of $\simeq 2\times$ weaker magnetic fields. The observable extent of radio emitters is limited by the propagation speed of CRs away from their sources and by the extent of the magnetic fields. At high radio frequencies (1–10 GHz) the radio emission from disks of star-forming galaxies is restricted to about 1 kpc from the sources of CRs. Low-frequency radio emission traces low-energy CRs which suffer less from energy losses and hence can propagate further away from their sources into regions with weak magnetic fields. The lifetime of CRs due to synchrotron losses increases with decreasing frequency and decreasing total field strength. In a $5 \mu\text{G}$ field electrons emitting in the LOFAR bands have a lifetime of $2\text{--}5 \times 10^8$ yr and can travel several tens of kpc in a magnetic field of about $3 \mu\text{G}$.

Many of the traditional depolarization effects expected from the technical limitations of low frequency radio observing are mitigated by the long baselines and high spectral resolution of the LOFAR instrument. As at higher frequencies, beam depolarization will limit polarization studies of sources with rapid image plane field reversals for any instrument of finite resolution. More important at these frequencies, however, is the internal depolarization of radio sources caused by field fluctuations along the line of sight. The effect of this internal Faraday dispersion has been discussed in the case of specific source morphologies by a number of authors (e.g., Cioffi & Jones 1980; Laing 1981) as well as for analytic geometries and random fluctuations (e.g., Tribble 1991; Sokoloff et al. 1998). Although these effects are expected to become increasingly important at longer

wavelengths a directed use of Faraday Rotation Measure Synthesis (Brentjens & de Bruyn 2005) will mitigate them in many circumstances. Indeed this technique is now becoming standard not only for recovering widespread polarized emission in a variety of environments but also for characterizing the line of sight medium itself (e.g., de Bruyn & Brentjens 2005; Heald et al. 2009).

LOFAR’s sensitivity to regions of low density and weak field strengths will allow us to measure the magnetic structure in the halos of galaxy clusters, in the intergalactic medium of galaxy groups, in wider halos and in outer disks of spiral galaxies. It is here that star formation activity is low, and processes additional to dynamo action, such as gas outflows from the inner disk, the magneto-rotational instability, gravitational interaction and ram pressure by the intergalactic medium are imprinted on this magnetic structure. The low frequencies provided by LOFAR will be highly sensitive to such steep-spectrum shock-like features, resembling relics in clusters, and knowledge of their 3D magnetic field structures from RM Synthesis will allow us vastly improved understanding of intergalactic gas dynamics.

Grids of RM measurements of polarized background sources are powerful tools to study magnetic field patterns in foreground galaxies and clusters of galaxies (Stepanov et al. 2008). Greater leverage on Faraday RM values is expected at lower frequencies; thus, LOFAR will observe tenuous ionized gas and/or very weak magnetic fields. It is even possible that LOFAR will directly detect magnetic fields in the filamentary intergalactic medium of the cosmic web. Unlike evolved clusters of galaxies, where highly efficient turbulent amplification is expected to have lead to saturation of the fields, the filamentary structures of the cosmic web are anticipated to be far more sensitive to the original seed mechanism responsible for cosmic magnetism. Importantly, detection of this field, or placing stringent upper limits on it, will

provide powerful observational constraints on the origin of cosmic magnetism.

LOFAR pulsar searches will benefit from both high sensitivity and an increasing pulsar brightness at low frequencies. This is expected to result in the discovery of a new population of dim, nearby and high-latitude pulsars too weak to be found at higher frequencies: roughly 1000 pulsar discoveries are expected from LOFAR. Polarization observations of these pulsars will approximately double the current sample of Faraday rotation measures (RMs) (see Fig. 32). This will provide the strength and direction of the regular magnetic field in previously unexplored directions and locations in the Galaxy; e.g. very little is known about the magnetic field properties of the Milky Way beyond a few hundred parsecs from the Galactic plane. RMs of high-latitude pulsars and extragalactic sources are crucial for determining fundamental properties such as the scale height and geometry of the magnetic field in the thick disk and halo, as well as providing the exciting prospect of discovering magnetic fields in globular clusters.

13.7. Solar physics and space weather

The Sun is an active star which exerts a strong influence on the space environment around Earth. This *Space Weather* can strongly affect global communication technology on which we increasingly rely. The Sun is an intense and variable source of radio emission. The strong thermal radiation of the quiet Sun is interspersed with intense radio bursts associated with solar activity such as flares and coronal mass ejections (CMEs). By combining the imaging and beam-forming observational modes, LOFAR can serve as a highly effective solar monitoring and imaging system. Thus, the study of the Sun by LOFAR is of great interest in solar physics and space weather.

The nonthermal radio radiation of the Sun is generated by energetic electrons produced by flares and/or CMEs. These energetic electrons excite high-frequency plasma waves (Langmuir and/or upper-hybrid waves) leading to the emission of radio waves by nonlinear plasma processes near the local electron plasma frequency and/or its harmonics (Melrose 1985). Since the plasma frequency only depends on the electron number density, and due to the gravitational stratification of the corona, each frequency corresponds to a certain height level in the corona (Mann et al. 1999). Thus, LOFAR enables the study of plasma processes associated with energetic electrons at different heights in the corona.

In March 2011, LOFAR observed its first solar radio burst, a so-called type I burst (McLean 1985), seen at 150 MHz on the West limb of the Sun (see Fig. 33). A detailed study of the radio morphology indicates that the burst is located above an active region. During the flare, as a signature of magnetic reconnection, a hot plasma jet is injected into the corona leading to the acceleration of electrons and subsequent radio emission (Miteva et al. 2007, Mann et al., in prep.).

Solar type III radio bursts were observed by LOFAR in October 2011, in both radio images and dynamic spectra. They manifest as a rapid drift from high to low frequencies in the dynamic radio spectrum (McLean 1985; Breitling et al. 2011). In this case, the type III burst source is located at the east limb above an active region (Mann et al., in prep.). Type III radio bursts are signatures of electron beams propagating along open magnetic field lines through the solar corona and sometimes into interplanetary space. They arise from electrons accelerated within a solar flare being injected into open magnetic field geometries.

In the corona, a shock wave is produced by a flare and/or driven by a CME. Signatures of such shock waves appear as type II radio bursts in solar dynamic radio spectra (Mann 1995; Aurass 1997). These shock waves are able to accelerate electrons up to supra-thermal velocities, resulting in type II radio bursts. Such bursts have also been detected in LOFAR observations.

Only a few instruments currently observe the Sun at low radio frequencies. The radioheliograph at Nançay (Kerdran & Delouis 1997), for example, is one of the few solar observatories currently operated at selected frequencies ranging from 150 to 432 MHz and yields a typical image resolution of 2–5 arcmin at the lower end of its frequency band. In recent years, however, a number of new instruments well suited for low frequency solar observations have come online. Besides LOFAR itself, these new instruments include the recently commissioned LWA (10–88 MHz) (see Kassim et al. 2010; Lazio et al. 2010; Taylor et al. 2012, and examples therein) and MWA (80–300 MHz) (see Rightley et al. 2009; Oberoi et al. 2011; Bowman et al. 2013; Tingay et al. 2013b, and examples therein), both of which provide powerful solar observing capabilities. Although not specifically intended for solar observations, the PAPER (see Parsons et al. 2010; Stefan et al. 2013; Pober et al. 2013, for descriptions of the PAPER instrument) operates in the 100–200 MHz range and also provides similar potential capabilities. In the Ukraine, the radio telescopes UTR2 (Sidorchuk et al. 2005) and URAN2 (Brazhenko et al. 2005) also work at low frequencies.

Like these other instruments, LOFAR expands the observable solar frequency range down to 10 MHz. Unlike these other facilities, however, it can also achieve angular resolutions of a few 10 arcsec, scattering in the solar corona becoming the limiting factor for resolution rather than baseline length. The broad low-frequency coverage combined with high resolution imaging makes LOFAR a powerful tool for probing previously unexplored solar coronal structures.

The excellent uv-coverage available using the Superterp stations also enable direct snapshot imaging of radio emission from CMEs. Further, it is possible to use a grid of many tied-array beams (as illustrated in Fig. 27) to form “maps” of radio emission covering a broad area around the Sun. These maps will in practice have a lower spatial resolution than those obtained by direct imaging, but have the advantage of the high time and frequency resolution available using the beam-formed mode and enable the direct deduction via dynamic spectra of the types of radio burst formed within the CME.

The solar wind, the expansion of the solar corona through interplanetary space, can be probed by observing the interplanetary scintillation (IPS) of compact radio sources (Hewish et al. 1964). Observations of IPS provide the ability to systematically study the solar wind at nearly all heliographic latitudes over a wide range of distances from the Sun. Various analysis techniques are used to probe different aspects of solar wind structure:

- Cross-correlation of the signals from two antennas, taken at times of suitable geometrical alignment between the antennas, Sun and radio source, can be used to resolve multiple solar wind streams in the lines of sight between antennas and radio source, yielding solar wind speeds and flow direction (e.g. Breen et al. 2006; Fallows et al. 2013).
- Multiple IPS observations over several days or more can be combined to produce three-dimensional reconstructions of solar wind speed and density throughout the inner heliosphere (e.g. Jackson & Hick 2005; Bisi et al. 2010).
- The combination of IPS and white light observations of the solar corona permit a far improved understanding of solar




Fig. 33. LOFAR image of a Type I solar radio burst, observed at 150 MHz on 17 March 2011. The white circle indicates the edge of the solar photosphere. Further study has revealed that this radio burst was located above an active region on the Sun (Mann et al., in prep.).

wind processes than would be possible with either technique alone (Dorrian et al. 2010; Hardwick et al. 2011).

- The high bandwidth capabilities of LOFAR enable the dynamic spectrum of IPS to be studied (see Fig. 34). The opportunity for such studies is available on few other instruments and may yield further information on solar wind micro-structure (Fallows et al. 2013).

The low-frequency bands available with LOFAR are best suited to using observations of interplanetary scintillation to study the solar wind from the orbit of Mercury out to beyond Earth orbit. This region is of particular interest in space weather as this is where much of the evolution of solar wind and CME structure occurs – information on which is essential for the accurate timing of the impact of such structures on the space environment around Earth.

Thus, with LOFAR, both the corona of the Sun and interplanetary space can be observed with unprecedented spatial and temporal resolution. This allows plasma processes in the corona and the solar wind to be studied in a manner that could not be achieved with other (e.g. optical) instruments. Both the imaging of the corona and the observation of IPS can contribute to the investigation of processes from initiation and launch of a CME to

its subsequent development and propagation through interplanetary space, topics of great importance for understanding many aspects of space weather.

14. Current and future developments

14.1. Final construction

Construction of the LOFAR array has been underway for the past five years and began in 2006 with the placement of several test stations on the site of the array core near Exloo in the Netherlands. This core is located in an area rich in peat deposits that were extensively harvested between 1850 and 1950 leaving behind a landscape used primarily for starch production from potato farming. Starting in 2008, the core area was extensively reshaped and established as a nature reserve with dedicated locations for the LOFAR stations.

Due to its agricultural history and the extensive landscaping required to establish the nature area, a large effort was required to stabilize the soil in the region. This work combined with the required ± 3 –6 cm tolerances on the flatness of the fields delayed the start of the large-scale civil engineering effort until the spring of 2009. Once begun however, progress has been rapid since with




Fig. 34. Dynamic spectrum of LOFAR LBA data from an observation of 3C48 (J0137+331) at an elongation of 50 degrees from the Sun (*top panel*) taken on 07 March 2012. These data have been averaged to a time resolution of 1.24 s and a frequency resolution of 180 kHz, matching the full-resolution data for the corresponding time period from the Nançay Decametric Array (*lower panel*). The Nançay Decametric Array is an array dedicated to solar observation (data courtesy A. Kerdraon, Meudon-Paris). The radio emission appears to be consistent with a Type II radio burst from solar flare activity.

the deployment of 22 stations in 2009 and a further 11 stations in 2010. For the remaining 7 remote NL stations, additional effort was required to obtain the necessary planning permission and building permits. Nonetheless, construction for the majority of these remaining stations was completed in 2012. At the time of writing, only one final station is as yet unfinished and, subject to obtaining final building permits, the NL array should be fully complete in 2013.

The international LOFAR stations have been built in parallel with those in the Netherlands beginning with the construction of the LBA field of the first German station near Effelsberg in 2007. The Effelsberg station was augmented with HBA tiles in 2009, and additional stations in Germany, each under different ownership, were completed throughout 2009–2011 in Unterweilbach, Tautenburg, Potsdam, and Jülich. A further station near Hamburg has recently been funded and construction is planned to start in early 2013. In 2010, a station was also built in Nançay, France, and in Chilbolton in the UK. Finally, a station near Onsala, Sweden was completed in 2011. Further expansion of the international array, in particular with a view to filling gaps in the uv -plane or extending the array, resulting in higher resolution, is currently under study.

14.2. Functionality enhancements

One of the great strengths of the LOFAR system is its capacity for enhancement. It is of course common for astronomical

facilities to increase their capabilities through continued software development. For LOFAR however, the system design is sufficiently flexible that scientific capacity can be added relatively straightforwardly at both the software and hardware levels. In the simplest case, this capacity increase can be achieved through the addition of more stations to the array resulting in improved uv coverage, longer baselines, and increased sensitivity. Such extensions to the array would, however, also require the addition of significant additional compute capacity. Due to practical I/O limits set by the BG/P configuration, only 64 stations of the 72² total possible can be correlated any given time with the current configuration. Increasing this number would require substantial changes to the current LOFAR computing infrastructure (see 6.1). We note that even adding smaller numbers of additional stations would require increasing the computing and storage capacity of the post-processing cluster in order to keep up with the increased data flow.

Similarly, the capabilities of individual LOFAR stations can also be expanded. With minimal modifications, the data-stream from a given station can be replicated and processed independently of the standard LOFAR processing. A number of such “stand-alone” or single station enhancements are already in development. The first of these, called ARTEMIS, implements a real-time dedispersion search engine to detect pulsars using the

² Assuming all core HBA sub-stations are being correlated as independent stations.

data-streams from one or more LOFAR stations (Serylak et al. 2012; Armour et al. 2012, and Karastergiou et al., in prep.). A second, EU-funded project named AARTFAAC will expand upon LOFAR's ability to monitor radio transients by correlating the signals from all dipoles on the Superterp in real-time (Prasad & Wijnholds 2012). Finally, a design for an expanded station concept has been proposed by the French LOFAR consortium. This design would add significant numbers of additional dipoles as well as computing capability to the current French station at Nançay resulting in a "SuperStation" optimized for beam-formed observations with high instantaneous sensitivity in the 10–80 MHz range (Zarka et al. 2012). Although the design has yet to be finalized, the proposed SuperStation would provide a factor of ~ 20 increase in effective area relative to a standard international LOFAR station. For comparison, the French SuperStation would deliver ~ 7 times the effective area of the current LWA station (Taylor et al. 2012).

15. Conclusions

In this paper, we have presented an overview and brief introduction to the LOFAR telescope. LOFAR represents a step-change in the evolution of radio astronomy technology. As one of the first of a new generation of radio instruments, LOFAR provides a number of unique capabilities for the astronomical community. These include among others remote configuration and operation, data processing that is both distributed and parallel, buffered retrospective all-sky imaging, dynamic real-time system response, and the ability to provide multiple simultaneous streams of data to a community whose scientific interests run the gamut from radio aurorae in the magnetospheres of distant planets to the origins of the Universe itself. Due to the tremendous data rates generated, LOFAR will also be one of the first radio observatories to feature automated processing pipelines to deliver fully calibrated scientific products to the community. Many of the technological solutions developed for LOFAR, in particular the calibration of phased-arrays as well as large-scale data transport and processing, will be highly relevant for future radio telescope projects such as the SKA.

Acknowledgements. The authors would like to thank the referee for the careful reading and many constructive comments that helped improve the paper. The LOFAR facilities in the Netherlands and other countries, under different ownership, are operated through the International LOFAR Telescope foundation (ILT) as an international observatory open to the global astronomical community under a joint scientific policy. In the Netherlands, LOFAR is funded through the BSIK program for interdisciplinary research and improvement of the knowledge infrastructure. Additional funding is provided through the European Regional Development Fund (EFRO) and the innovation program EZ/KOMPAS of the Collaboration of the Northern Provinces (SNN). ASTRON is part of the Netherlands Organization for Scientific Research (NWO). C. Ferrari and G. Macario acknowledge financial support by the "Agence Nationale de la Recherche" through grant ANR-09-JCJC-0001-01.

References

- Apel, W., Arteaga, J., Bähren, L., et al. 2012, Nucl. Inst. Meth. A, 696, 100
- Armour, W., Karastergiou, A., Giles, M., et al. 2012, in *Astronomical Data Analysis Software and Systems XXI*, eds. P. Ballester, D. Egret, & N. P. F. Lorente, ASP Conf. Ser., 461, 33
- Aurass, H. 1997, in *Coronal Physics from Radio and Space Observations*, ed. G. Trotter (Berlin: Springer Verlag), Lecture Notes in Physics, 483, 135
- Barkana, R., & Loeb, A. 2001, Phys. Rep., 349, 125
- Beck, R. 2010, in *ISKAF2010 Science Meeting*
- Bell, S. J., & Hewish, A. 1967, Nature, 213, 1214
- Bennett, A. S. 1962, MNRAS, 125, 75
- Bhatnagar, S., Cornwell, T. J., Golap, K., & Uson, J. M. 2008, A&A, 487, 419
- Bisi, M. M., Jackson, B. V., Hick, P. P., et al. 2010, ApJ, 715, L104
- Blümer, J., Engel, R., & Hörandel, J. R. 2009, Prog. Part. Nucl. Phys., 2, 293
- Bolton, J. S., & Haehnelt, M. G. 2007, MNRAS, 382, 325
- Bolton, J. S., Becker, G. D., Wyithe, J. S. B., Haehnelt, M. G., & Sargent, W. L. W. 2010, MNRAS, 771
- Boncz, P., Grust, T., van Keulen, M., et al. 2006, in the 2006 ACM SIGMOD Int. Conf. (New York, USA: ACM Press), 479
- Boonstra, A. J., & van der Tol, S. 2005, Radio Science, 40, RS5S09
- Bowman, J. D., Cairns, I., Kaplan, D. L., et al. 2013, PASA, 30, 31
- Brazhenko, A. I., Bulatsen, V. G., Vashchishin, R. V., et al. 2005, Kinematika i Fizika Nebesnykh Tel Supplement, 5, 43
- Breen, A., Fallows, R., Bisi, M., et al. 2006, J. Geophys. Res., 111, A08104
- Bregman, J. D. 2000, in SPIE Conf. Ser. 4015, ed. H. R. Butcher, 19
- Breitling, F., Mann, G., & Vocks, C. 2011, Planetary, Solar and Heliospheric Radio Emissions (PRE VII), 373
- Brentjens, M. A., & de Bruyn, A. G. 2005, A&A, 441, 1217
- Bromm, V., & Larson, R. B. 2004, ARA&A, 42, 79
- Buitink, S., Scholten, O., Bacelar, J., et al. 2010, A&A, 521, A47
- Bunker, A. J., Wilkins, S., Ellis, R. S., et al. 2010, MNRAS, 409, 855
- Burke, B. F., & Franklin, K. L. 1955, J. Geophys. Res., 60, 213
- Burke-Spolaor, S., Bailes, M., Johnston, S., et al. 2011, MNRAS, 416, 2465
- Butcher, H. R. 2004, in SPIE Conf. Ser. 5489, ed. J. M. Oschmann, Jr., 537
- Carilli, C. L., & Yun, M. S. 1999, ApJ, 513, L13
- Cassano, R. 2010, A&A, 517, A10
- Cassano, R., Brunetti, G., Röttgering, H. J. A., & Brügger, M. 2010, A&A, 509, A68
- Castro-Tirado, A. J., de Ugarte Postigo, A., Gorosabel, J., et al. 2008, Nature, 455, 506
- Choudhury, T. R., & Ferrara, A. 2006, MNRAS, 371, L55
- Ciardi, B., & Ferrara, A. 2005, Space Sci. Rev., 116, 625
- Ciardi, B., & Madau, P. 2003, ApJ, 596, 1
- Ciardi, B., Ferrara, A., & White, S. D. M. 2003, MNRAS, 344, L7
- Ciardi, B., Bolton, J. S., Maselli, A., & Graziani, L. 2012, MNRAS, 423, 558
- Cioffi, D. F., & Jones, T. W. 1980, AJ, 85, 368
- Cohen, A. S. 2004, LWA Memo Series, 17
<http://www.ece.vt.edu/swe/lwa/memo/lwa0017.pdf>
- Cohen, A. S. 2006, LWA Memo Series, 52
<http://www.ece.vt.edu/swe/lwa/memo/lwa0052.pdf>
- Cohen, A. S., Lane, W. M., Cotton, W. D., et al. 2007, AJ, 134, 1245
- Condon, J. J. 1974, ApJ, 188, 279
- Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, AJ, 115, 1693
- Cordes, J. M. 1978, ApJ, 222, 1006
- Cornwell, T. J., Golap, K., & Bhatnagar, S. 2008, IEEE J. Sel. Top. Sign. Proc., 2, 647
- de Bruyn, A. G., & Brentjens, M. A. 2005, A&A, 441, 931
- de Gasperin, F., Orrù, E., Murgia, M., et al. 2012, A&A, 547, A56
- de Vos, M., Gunst, A. W., & Nijboer, R. 2009, IEEE Proc., 97, 1431
- Deshpande, A. A., & Radhakrishnan, V. 1994, J. Astrophys. Astron., 15, 329
- Dorrian, G., Breen, A., Davies, J., et al. 2010, Sol. Phys., 265, 207
- Edge, D. O., Shakeshaft, J. R., McAdam, W. B., Baldwin, J. E., & Archer, S. 1959, MNRAS, 68, 37
- Ekers, R. 2012, in *Proceedings of Resolving The Sky – Radio Interferometry: Past, Present and Future* [arXiv:1212.3497]
- Ellingson, S. W., Clarke, T. E., Cohen, A., et al. 2009, IEEE Proc., 97, 1421
- Ellingson, S. W., Taylor, G. B., Craig, J., et al. 2013, Antennas and Propagation, IEEE Trans., 61, 2540
- Enßlin, T. A., & Röttgering, H. 2002, A&A, 396, 83
- Erickson, W. C., & Fisher, J. R. 1974, Radio Science, 9, 387
- Espinoza, C. M., Lyne, A. G., Stappers, B. W., & Kramer, M. 2011, MNRAS, 414, 1679
- Ewen, H. I., & Purcell, E. M. 1951, Nature, 168, 356
- Falcke, H. D. E. 2006, in *IAU Joint Discussion*, 12
- Falcke, H. 2008, in *Frontiers of Astrophysics: A Celebration of NRAO's 50th Anniversary*, eds. A. H. Bridle, J. J. Condon, & G. C. Hunt, ASP Conf. Ser., 395, 271
- Falcke, H., & Gorham, P. 2003, Astropart. Phys., 19, 477
- Falcke, H., Apel, W. D., Badae, A. F., et al. 2005, Nature, 435, 313
- Falcke, H. D., van Haarlem, M. P., de Bruyn, A. G., et al. 2007, Highlights of Astronomy, 14, 386
- Fallows, R. A., Asgekar, A., Bisi, M. M., Breen, A. R., & Ter-Veen, S. 2013, Sol. Phys., 285, 127
- Fan, X., Strauss, M. A., Richards, G. T., et al. 2003, AJ, 125, 1649
- Fan, X., Strauss, M. A., Schneider, D. P., et al. 2006, AJ, 131, 1203
- Ferrari, C., Govoni, F., Schindler, S., Bykov, A. M., & Rephaeli, Y. 2008, Space Sci. Rev., 134, 93
- Field, G. B. 1958, Proc. IRE, 46, 240
- Finlay, C. C., Maus, S., Beggan, C. D., et al. 2010, Geophys. J. Int., 183, 1216
- Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, Phys. Rep., 433, 181
- Gnedin, N. Y., & Abel, T. 2001, New Astron., 6, 437

- Gower, J. F. R., Scott, P. F., & Wills, D. 1967, MNRAS, 71, 49
- Greisen, K. 1966, Phys. Rev. Lett., 16, 748
- Hamaker, J. P. 2000, A&AS, 143, 515
- Hamaker, J. P. 2006, A&A, 456, 395
- Hamaker, J. P., & Bregman, J. D. 1996, A&AS, 117, 161
- Hamaker, J. P., Bregman, J. D., & Sault, R. J. 1996, A&AS, 117, 137
- Hardwick, S. A., Breen, A., Bisi, M. M., et al. 2011, AGU Fall Meeting Abstracts, 7
- Harker, G. J. A., Zaroubi, S., Thomas, R. M., et al. 2009, MNRAS, 393, 1449
- Harker, G., Zaroubi, S., Bernardi, G., et al. 2010, MNRAS, 405, 2492
- Hassall, T. E., Stappers, B. W., Hessels, J. W. T., et al. 2012, A&A, 543, A66
- Heald, G., Braun, R., & Edmonds, R. 2009, A&A, 503, 409
- Heald, G., McKean, J., Pizzo, R., et al. 2010, in Proc. ISKAF2010 Science Meeting, Published online at <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=112>, 57
- Heald, G., Bell, M. R., Horneffer, A., et al. 2011, J. Astrophys. Astron., 32, 589
- Hermesen, W., Hessels, J. W. T., Kuiper, L., et al. 2013, Science, 339, 436
- Hernández-Pajares, M., Juan, J. M., Sanz, J., et al. 2009, J. Geodesy, 83, 263
- Hessels, J., Stappers, B., Alexov, A., et al. 2010, in Proc. ISKAF2010 Science Meeting, Published online at <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=112>, 25
- Hewish, A., Scott, P., & Wills, D. 1964, Nature, 203, 1214
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, Nature, 217, 709
- Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, PASA, 21, 302
- Huege, T., & Falcke, H. 2005, A&A, 430, 779
- Hyman, S. D., Lazio, T. J. W., Kassim, N. E., et al. 2005, Nature, 434, 50
- Intema, H. T., van der Tol, S., Cotton, W. D., et al. 2009, A&A, 501, 1185
- Ivanova, M., Nes, N., Goncalves, R., & Kersten, M. 2007, in 19th International Conference on Scientific and Statistical Database Management (SSDBM 2007), IEEE, 13
- Jackson, B. V. & Hick, P. P. 2005, in Solar and Space Weather Radiophysics, eds. D. E. Gary, & C. U. Keller (Netherlands: Springer), A&SS Lib., 314, 355
- Jansky, K. G. 1933, Nature, 132, 66
- Jelić, V., Zaroubi, S., Labropoulos, P., et al. 2008, MNRAS, 389, 1319
- Kampert, K.-H., & Unger, M. 2012, Astro. Part. Phys., 10, 660
- Kassim, N. E. 1988, ApJS, 68, 715
- Kassim, N. E., & Erickson, W. C. 1998, in SPIE Conf. Ser. 3357, ed. T. G. Phillips, 740
- Kassim, N. E., Perley, R. A., Erickson, W. C., & Dwarakanath, K. S. 1993, AJ, 106, 2218
- Kassim, N. E., Lazio, T. J., Crane, P. C., et al. 2003, in EGS – AGU – EUG Joint Assembly, 3523
- Kassim, N. E., Lazio, T. J. W., Erickson, W. C., et al. 2007, ApJS, 172, 686
- Kassim, N., White, S., Rodriguez, P., et al. 2010, in Advanced Maui Optical and Space Surveillance Technologies Conference, held in Wailea, Maui, Hawaii, September 14–17, ed. S. Ryan, The Maui Economic Development Board, E59
- Kelley, J., & The Pierre Auger Collaboration 2011, in Proc. 32nd Int. Cosmic Ray Conference, Beijing
- Kerdran, A., & Delouis, J.-M. 1997, in Coronal Physics from Radio and Space Observations, ed. G. Trottet, Lect. Notes Phys. (Berlin: Springer Verlag), 483, 192
- Komatsu, E., Smith, K. M., Dunkley, J., et al. 2011, ApJS, 192, 18
- Kramer, M., Lyne, A. G., O'Brien, J. T., Jordan, C. A., & Lorimer, D. R. 2006, Science, 312, 549
- Labropoulos, P., Koopmans, L. V. E., Jelic, V., et al. 2009, ArXiv e-prints, [[arXiv:0901.3359](https://arxiv.org/abs/0901.3359)]
- Laing, R. A. 1981, ApJ, 248, 87
- Lane, W. M., Cotton, W. D., Helmboldt, J. F., & Kassim, N. E. 2012, Radio Science, 47, Issue 3, CiteID RS0K04
- Law, C. J., & Bower, G. C. 2012, ApJ, 749, 143
- Law, C. J., Bower, G. C., Pokorný, M., Rupen, M. P., & Sowiński, K. 2012, ApJ, 760, 124
- Lazio, T. J. W., Cohen, A. S., Kassim, N. E., et al. 2006, ApJ, 642, L33
- Lazio, T. J. W., Clarke, T. E., Lane, W. M., et al. 2010, AJ, 140, 1995
- Leahy, J. P., Muxlow, T. W. B., & Stephens, P. W. 1989, MNRAS, 239, 401
- LIGO Scientific Collaboration, Virgo Collaboration, Abadie, J., et al. 2012, A&A, 539, A124
- Lonsdale, C. J., Cappallo, R. J., Morales, M. F., et al. 2009, IEEE Proc., 97, 1497
- Lorimer, D. R. 2011, Astrophysics Source Code Library, 7016
- Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, Science, 318, 777
- Ludwig, M., & Huege, T. 2011, Astropart. Phys., 34, 438
- Lyne, A., Hobbs, G., Kramer, M., Stairs, I., & Stappers, B. 2010, Science, 329, 408
- Madau, P., Meiksin, A., & Rees, M. J. 1997, ApJ, 475, 429
- Mann, G. 1995, in Coronal Magnetic Energy Releases, eds. A. O. Benz, & A. Krüger, Lecture Notes in Physics (Berlin: Springer Verlag), 444, 183
- Mann, G., Jansen, F., MacDowall, R. J., Kaiser, M. L., & Stone, R. G. 1999, A&A, 348, 614
- McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., et al. 2006, Nature, 439, 817
- McLean, D. J. 1985, in Solar radiophysics, eds. D. J. McLean, & N. R. Labrum (Cambridge University Press), 37
- Mellema, G., Iliev, I. T., Alvarez, M. A., & Shapiro, P. R. 2006, New Astron., 11, 374
- Melrose, D. B. 1985, in Solar radiophysics, eds. D. J. McLean, & N. R. Labrum (Cambridge University Press), 177
- Mesinger, A., & Furlanetto, S. 2007, ApJ, 669, 663
- Mevius, M., Buitink, S., Falcke, H., et al. 2012, Nucl. Inst. Meth. A, Suppl. 1, 662, S26
- Miley, G., & de Breuck, C. 2008, A&ARv, 15, 67
- Miley, G. G. K. 2010, in Proc. of the ISKAF2010 Science Meeting, Published online at <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=112>, 40
- Miteva, R., Mann, G., Vocks, C., & Aurass, H. 2007, A&A, 461, 1127
- Mol, J. D., & Romein, J. W. 2011 [[arXiv:1105.0661](https://arxiv.org/abs/1105.0661)]
- Napier, P. J. 1999, in Synthesis Imaging in Radio Astronomy II, eds. G. B. Taylor, C. L. Carilli, & R. A. Perley, ASP Conf. Ser., 180, 37
- Nelles, A., Buitink, S., Corstanje, A., et al. 2013, in Proc. the 5th international workshop on Acoustic and Radio EeV Neutrino Detection Activities (ARENA), Erlangen, Germany, AIP Conf. Ser. (in press)
- Oberoi, D., Matthews, L. D., Cairns, I. H., et al. 2011, ApJ, 728, L27
- Oesch, P. A., Bouwens, R. J., Illingworth, G. D., et al. 2010, ApJ, 709, L16
- Offringa, A. R., de Bruyn, A. G., Biehl, M., et al. 2010, MNRAS, 405, 155
- Offringa, A. R., de Bruyn, A. G., & Zaroubi, S. 2012a, MNRAS, 422, 563
- Offringa, A. R., van de Gronde, J. J., & Roerdink, J. B. T. M. 2012b, A&A, 539, A95
- Offringa, A. R., de Bruyn, A. G., Zaroubi, S., et al. 2013, A&A, 549, A11
- Page, L., Hinshaw, G., Komatsu, E., et al. 2007, ApJS, 170, 335
- Parsons, A. R., Backer, D. C., Foster, G. S., et al. 2010, AJ, 139, 1468
- Penzias, A. A., & Wilson, R. W. 1965, ApJ, 142, 419
- Perley, R. A., & Erickson, W. C. 1984, VLA Scientific Memorandum, 146 <http://www.vla.nrao.edu/memos/sci/146.pdf>
- Pilkington, J. D. H. & Scott, P. F. 1965, MNRAS, 69, 183
- Pober, J. C., Parsons, A. R., Aguirre, J. E., et al. 2013, ApJ, 768, L36
- Prasad, P., & Wijnholds, S. J. 2012, in New Windows on Transients across the Universe: A Discussion meeting of the Royal Society, London, April [[arXiv:1205.3056](https://arxiv.org/abs/1205.3056)]
- Ransom, S. 2011, Astrophysics Source Code Library, 7017
- Rengelink, R. B., Tang, Y., de Bruyn, A. G., et al. 1997, A&AS, 124, 259
- Rightley, S. M., Matthews, L. D., & Oberoi, D. 2009, AGU Fall Meeting Abstracts, 1548
- Romein, J. W., Broekema, P. C., van Meijeren, E., van der Schaaf, K., & Zwart, W. H. 2006, in Proc. 18th Annual ACM Symp. on Parallelism in Algorithms and Architectures, SPAA '06 (New York, NY, USA: ACM), 59
- Romein, J. W., Broekema, P. C., Mol, J. D., & van Nieuwpoort, R. V. 2010, in ACM Symp. on Principles and Practice of Parallel Programming (PPoPP'10), Bangalore, India, 169
- Romney, J. D. 1999, in Synthesis Imaging in Radio Astronomy II, eds. G. B. Taylor, C. L. Carilli, & R. A. Perley, ASP Conf. Ser., 180, 57
- Röttgering, H., Afonso, J., Barthel, P., et al. 2011, J. A&A, 32, 557
- Sault, R. J., Hamaker, J. P., & Bregman, J. D. 1996, A&AS, 117, 149
- Scaife, A. M. M. & Heald, G. H. 2012, MNRAS, 423, L30
- Scholten, O., Werner, K., & Rusydi, F. 2008, Astropart. Phys., 29, 94
- Scholten, O., Buitink, S., Bacelar, J., et al. 2009, Phys. Rev. Lett., 103, 191301
- Seaman, R., Allan, A., & Williams, R. 2008, in Astronomical Data Analysis Software and Systems XVII, eds. R. W. Argyle, P. S. Bunclark, & J. R. Lewis, ASP Conf. Ser., 394, 213
- Seaman, R., Williams, R., Allan, A., et al. 2011, IVOA Recommendation, 27 [[arXiv:1110.0523](https://arxiv.org/abs/1110.0523)]
- Serylak, M., Karastergiou, A., Williams, C., et al. 2012, in Neutron Stars and Pulsars: Challenges and Opportunities after 80 years, ed. J. van Leeuwen, in press, IAU Symp., 291, 492
- Sidorchuk, K. M., Konvalenko, A. A., Rucker, H. O., et al. 2005, Kinematika i Fizika Nebesnykh Tel Supplement, 5, 57
- Smirnov, O. M. 2011, A&A, 527, A106
- Sokoloff, D. D., Bykov, A. A., Shukurov, A., et al. 1998, MNRAS, 299, 189
- Sotomayor-Beltran, C., Sobey, C., Hessels, J. W. T., et al. 2013, A&A, 552, A58
- Spergel, D. N., Bean, R., Doré, O., et al. 2007, ApJS, 170, 377
- Stappers, B. W., Hessels, J. W. T., Alexov, A., et al. 2011, A&A, 530, A80
- Stefan, I. I., Carilli, C. L., Green, D. A., et al. 2013, MNRAS, 432, 1285
- Stepanov, R., Arshakian, T. G., Beck, R., Frick, P., & Krause, M. 2008, A&A, 480, 45
- Swinbank, J. 2011, in Astronomical Data Analysis Software and Systems XX, eds. I. N. Evans, A. Accomazzi, D. J. Mink, & A. H. Rots, ASP Conf. Ser., 442, 313

- Szomoru, A. 2008, in Proc. 9th European VLBI Network Symposium on The role of VLBI in the Golden Age for Radio Astronomy and EVN Users Meeting, September 23–26, Bologna, Italy. Available online at <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=72>
- Tasse, C., van der Tol, B., van Zwieteren, J., van Diepen, G., & Bhatnagar, S. 2013, *A&A*, 553, A105
- Taylor, J. H., Manchester, R. N., & Lyne, A. G. 1993, *ApJS*, 88, 529
- Taylor, G. B., Carilli, C. L., & Perley, R. A., eds. 1999, *Synthesis Imaging in Radio Astronomy II*, ASP Conf. Ser., 180
- Taylor, G. B., Ellingson, S. W., Kassim, N. E., et al. 2012, *J. Astron. Instrumentation*, 1, 50004
- ter Veen, S., Schellart, P., & Falcke, H. 2012, in *IAU Symp.* 285, eds. E. Griffin, R. Hanisch, & R. Seaman, 411
- The Pierre Auger Collaboration. 2012, *Phys. Rev. Lett.*, 109, 2002
- Theuns, T., Schaye, J., Zaroubi, S., et al. 2002, *ApJ*, 567, L103
- Thomas, R. M., & Zaroubi, S. 2011, *MNRAS*, 410, 1377
- Thomas, R. M., Zaroubi, S., Ciardi, B., et al. 2009, *MNRAS*, 393, 32
- Thompson, A. R., Moran, J. M., & Swenson, G. W. 2007, *Interferometry and Synthesis in Radio Astronomy*, 2nd Edition (Wiley)
- Thoudam, S., Aar, G. V., Akker, M. V. D., et al. 2011, *Ap&SS Trans.*, 7, 195
- Tingay, S. J., Goeke, R., Bowman, J. D., et al. 2013a, *PASA*, 30, 7
- Tingay, S. J., Oberoi, D., Cairns, I., et al. 2013b, *J. Phys. Conf. Ser.*, 440, id. 012033
- Tribble, P. C. 1991, *MNRAS*, 250, 726
- van Ardenne, A., Smolders, B., & Hampson, G. 2000, in *SPIE Conf. Ser.* 4015, ed. H. R. Butcher, 420
- van Breugel, W., de Breuck, C., Stanford, S. A., et al. 1999, *ApJ*, 518, L61
- van de Hulst, H. C. 1945, *Ned. Tijdschr. Natuurk.*, 11, 210
- van der Tol, S., Jeffs, B. D., & van der Veen, A.-J. 2007, *IEEE Trans. Sig. Process.*, 55, 4497
- van Leeuwen, J., & Stappers, B. W. 2010, *A&A*, 509, A7
- van Straten, W., & Bailes, M. 2011, *PASA*, 28, 1
- van Weeren, R. J., Röttgering, H. J. A., Rafferty, D. A., et al. 2012, *A&A*, 543, A43
- Werner, K., de Vries, K. D., & Scholten, O. 2012, *Astropart. Phys.*, 37, 5
- Whalen, D., & Norman, M. L. 2006, *ApJS*, 162, 281
- Whitney, A. R., Shapiro, I. I., Rogers, A. E. E., et al. 1971, *Science*, 173, 225
- Wijnholds, S. J., & van der Veen, A.-J. 2009, *IEEE Trans. Sig. Process.*, 57, 3512
- Wijnholds, S. J., & van der Veen, A.-J. 2010, in Proc. 17th European Sig. Process. Conf. (EuSIPCo), Glasgow, UK, 24–28 August 2009 [[arXiv:1003.2497](https://arxiv.org/abs/1003.2497)]
- Wijnholds, S., van der Tol, S., Nijboer, R., & van der Veen, A.-J. 2010, *IEEE Sig. Process. Mag.*, 27, 30
- Wilkinson, P. N. 1991, in *IAU Colloq.* 131, *Radio Interferometry. Theory, Techniques, and Applications*, eds. T. J. Cornwell, & R. A. Perley, ASP Conf. Ser., 19, 428
- Yatawatta, S., de Bruyn, A. G., Brentjens, M. A., et al. 2013, *A&A*, 550, A136
- Zahn, O., Lidz, A., McQuinn, M., et al. 2007, *ApJ*, 654, 12
- Zarka, P. 2004, *Planet. Space Sci.*, 52, 1455
- Zarka, P. 2011, *Planetary, Solar and Heliospheric Radio Emissions (PRE VII)*, 287
- Zarka, P., Girard, J. N., Tagger, M., & Denis, L. 2012, in SF2A-2012: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics, eds. S. Boissier, P. de Laverny, N. Nardetto, et al., 687
- Zaroubi, S. 2013, in *A&SS Lib.* 396, eds. T. Wiklind, B. Mobasher, & V. Bromm, 45
- Zaroubi, S., de Bruyn, A. G., Harker, G., et al. 2012, *MNRAS*, 425, 2964
- Zatsepin, G., & Kuz'min, V. 1966, *JETP Lett.*, 4, 78
- ⁴ Leiden Observatory, Leiden University, PO Box 9513, 2300 RA Leiden, The Netherlands
- ⁵ Department of Astrophysics/IMAPP, Radboud University Nijmegen, PO Box 9010, 6500 GL Nijmegen, The Netherlands
- ⁶ Jodrell Bank Center for Astrophysics, School of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK
- ⁷ Astrophysics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK
- ⁸ Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany
- ⁹ School of Physics and Astronomy, University of Southampton, Southampton, SO17 1BJ, UK
- ¹⁰ Max Planck Institute for Astrophysics, Karl Schwarzschild Str. 1, 85741 Garching, Germany
- ¹¹ Onsala Space Observatory, Dept. of Earth and Space Sciences, Chalmers University of Technology, 43992 Onsala, Sweden
- ¹² International Centre for Radio Astronomy Research – Curtin University, GPO Box U1987, WA 6845 Perth, Australia
- ¹³ STFC Rutherford Appleton Laboratory, Harwell Science and Innovation Campus, Didcot OX11 0QX, UK
- ¹⁴ Institute for Astronomy, University of Edinburgh, Royal Observatory of Edinburgh, Blackford Hill, Edinburgh EH9 3HJ, UK
- ¹⁵ LESIA, Observatoire de Paris, CNRS, UPMC, Université Paris Diderot, 92190 Meudon, France
- ¹⁶ Argelander-Institut für Astronomie, University of Bonn, Auf dem Hügel 71, 53121 Bonn, Germany
- ¹⁷ Leibniz-Institut für Astrophysik Potsdam (AIP), An der Sternwarte 16, 14482 Potsdam, Germany
- ¹⁸ Thüringer Landessternwarte, Sternwarte 5, 07778 Tautenburg, Germany
- ¹⁹ Astronomisches Institut der Ruhr-Universität Bochum, Universitätsstrasse 150, 44780 Bochum, Germany
- ²⁰ Universität Hamburg, Hamburger Sternwarte, Gojenbergsweg 112, 21029 Hamburg, Germany
- ²¹ Jacobs University Bremen, Campus Ring 1, 28759 Bremen, Germany
- ²² Laboratoire de Physique et Chimie de l'Environnement et de l'Espace, CNRS/Université d'Orléans, LPC2E UMR 7328 CNRS, 45071 Orléans Cedex 2, France
- ²³ Center for Information Technology (CIT), University of Groningen, The Netherlands
- ²⁴ Radio Astronomy Lab, UC Berkeley, CA 94720, USA
- ²⁵ Centre de Recherche Astrophysique de Lyon, Observatoire de Lyon, 9 av. Charles André, 69561 Saint Genis Laval Cedex, France
- ²⁶ Mt. Stromlo Obs., Research School of Astronomy and Astrophysics, Australian National University, A.C.T. 2611 Weston, Australia
- ²⁷ CSIRO Australia Telescope National Facility, PO Box 76, NSW 1710 Epping, Australia
- ²⁸ National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903-2475, USA
- ²⁹ Chalmers University of Technology, 412 96 Gothenburg, Sweden
- ³⁰ Observatoire de la Côte d'Azur, Département Lagrange, Boulevard de l'Observatoire, BP 4229, 06304 Nice Cedex 4, France
- ³¹ Station de Radioastronomie de Nançay, Observatoire de Paris, CNRS/INSU, 18330 Nançay, France
- ³² Netherlands eScience Center, Science Park 140, 1098 XG Amsterdam, The Netherlands
- ³³ Centrum Wiskunde & Informatica, PO Box 94079, 1090 GB Amsterdam, The Netherlands
- ¹ Netherlands Institute for Radio Astronomy (ASTRON), Postbus 2, 7990 AA Dwingeloo, The Netherlands
- ² Astronomical Institute “Anton Pannekoek”, University of Amsterdam, Postbus 94249, 1090 GE Amsterdam, The Netherlands
- ³ Kapteyn Astronomical Institute, PO Box 800, 9700 AV Groningen, The Netherlands

Appendix A: LOFAR station field center positions**Table A.1.** LOFAR station field center positions.

Station	ETRS-X (m)	ETRS-Y (m)	ETRS-Z (m)
CS001LBA	3 826 923.942	460 915.117	5 064 643.229
CS001HBA	3 826 938.206	460 938.202	5 064 630.436
CS001HBA0	3 826 896.631	460 979.131	5 064 657.943
CS001HBA1	3 826 979.780	460 897.273	5 064 602.929
CS002LBA	3 826 577.462	461 022.624	5 064 892.526
CS002HBA0	3 826 601.357	460 953.078	5 064 880.876
CS002HBA1	3 826 565.990	460 957.786	5 064 906.998
CS003LBA	3 826 517.144	460 929.742	5 064 946.197
CS003HBA0	3 826 471.744	460 999.814	5 064 973.941
CS003HBA1	3 826 518.208	461 034.934	5 064 935.890
CS004LBA	3 826 654.593	460 939.252	5 064 842.166
CS004HBA0	3 826 586.022	460 865.520	5 064 900.301
CS004HBA1	3 826 579.882	460 917.156	5 064 900.242
CS005LBA	3 826 669.146	461 069.226	5 064 819.494
CS005HBA0	3 826 701.556	460 988.926	5 064 802.425
CS005HBA1	3 826 631.590	461 021.491	5 064 851.999
CS006LBA	3 826 597.126	461 144.854	5 064 866.718
CS006HBA0	3 826 654.179	461 136.116	5 064 824.683
CS006HBA1	3 826 612.895	461 079.974	5 064 860.746
CS007LBA	3 826 533.757	461 098.642	5 064 918.461
CS007HBA0	3 826 479.111	461 083.396	5 064 960.857
CS007HBA1	3 826 538.417	461 169.407	5 064 908.567
CS011LBA	3 826 667.465	461 285.525	5 064 801.332
CS011HBA	3 826 643.587	461 290.469	5 064 818.809
CS011HBA0	3 826 637.817	461 227.021	5 064 828.874
CS011HBA1	3 826 649.357	461 353.917	5 064 808.743
CS013LBA	3 826 346.661	460 791.787	5 065 086.876
CS013HBA	3 826 360.925	460 814.872	5 065 074.083
CS013HBA0	3 826 319.350	460 855.801	5 065 101.590
CS013HBA1	3 826 402.499	460 773.943	5 065 046.576
CS017LBA	3 826 462.450	461 501.626	5 064 935.567
CS017HBA	3 826 452.835	461 529.655	5 064 940.251
CS017HBA0	3 826 405.491	461 507.136	5 064 977.823
CS017HBA1	3 826 500.179	461 552.174	5 064 902.678
CS021LBA	3 826 406.939	460 538.280	5 065 064.610
CS021HBA	3 826 416.554	460 510.252	5 065 059.927
CS021HBA0	3 826 463.898	460 532.770	5 065 022.354
CS021HBA1	3 826 369.209	460 487.733	5 065 097.499
CS024LBA	3 827 161.630	461 409.084	5 064 420.786
CS024HBA	3 827 171.245	461 381.055	5 064 416.102
CS024HBA0	3 827 218.589	461 403.574	5 064 378.530
CS024HBA1	3 827 123.900	461 358.537	5 064 453.675
CS026LBA	3 826 391.312	461 869.528	5 064 955.653
CS026HBA	3 826 377.049	461 846.443	5 064 968.446
CS026HBA0	3 826 418.623	461 805.513	5 064 940.939
CS026HBA1	3 826 335.474	461 887.372	5 064 995.953
CS028LBA	3 825 600.841	461 260.269	5 065 604.065
CS028HBA	3 825 615.105	461 283.354	5 065 591.272
CS028HBA0	3 825 573.530	461 324.283	5 065 618.779
CS028HBA1	3 825 656.679	461 242.425	5 065 563.765
CS030LBA	3 826 014.662	460 387.065	5 065 372.068
CS030HBA	3 826 000.399	460 363.979	5 065 384.861
CS030HBA0	3 826 041.973	460 323.050	5 065 357.354
CS030HBA1	3 825 958.824	460 404.909	5 065 412.368
CS031LBA	3 826 440.392	460 273.509	5 065 063.334
CS031HBA	3 826 430.777	460 301.538	5 065 068.018
CS031HBA0	3 826 383.433	460 279.019	5 065 105.590
CS031HBA1	3 826 478.121	460 324.057	5 065 030.445
CS032LBA	3 826 891.969	460 387.586	5 064 715.032

Notes. For the stations in the Netherlands, the nomenclature CS and RS are used to refer to “core stations” and “remote stations”, respectively. See Sect. 4.1 for a description of the distinction between the two types. International LOFAR stations use a nomenclature based on the host country.

Table A.1. continued.

Station	ETRS-X (m)	ETRS-Y (m)	ETRS-Z (m)
CS032HBA	3 826 906.233	460 410.671	5 064 702.239
CS032HBA0	3 826 864.658	460 451.600	5 064 729.746
CS032HBA1	3 826 947.807	460 369.742	5 064 674.732
CS101LBA	3 825 843.362	461 704.125	5 065 381.213
CS101HBA	3 825 852.977	461 676.097	5 065 376.530
CS101HBA0	3 825 900.321	461 698.615	5 065 338.957
CS101HBA1	3 825 805.632	461 653.578	5 065 414.102
CS103LBA	3 826 304.675	462 822.765	5 064 934.074
CS103HBA	3 826 290.412	462 799.679	5 064 946.867
CS103HBA0	3 826 331.986	462 758.750	5 064 919.360
CS103HBA1	3 826 248.837	462 840.609	5 064 974.374
CS201LBA	3 826 709.325	461 913.423	5 064 713.578
CS201HBA	3 826 685.447	461 918.367	5 064 731.055
CS201HBA0	3 826 679.677	461 854.919	5 064 741.120
CS201HBA1	3 826 691.217	461 981.815	5 064 720.989
CS301LBA	3 827 413.261	460 992.019	5 064 269.684
CS301HBA	3 827 437.139	460 987.076	5 064 252.208
CS301HBA0	3 827 442.908	461 050.523	5 064 242.143
CS301HBA1	3 827 431.369	460 923.628	5 064 262.273
CS302LBA	3 827 946.312	459 792.315	5 063 989.756
CS302HBA	3 827 932.048	459 769.230	5 064 002.547
CS302HBA0	3 827 973.622	459 728.300	5 063 975.040
CS302HBA1	3 827 890.473	459 810.159	5 064 030.053
CS401LBA	3 826 766.502	460 100.064	5 064 836.210
CS401HBA	3 826 790.378	460 095.120	5 064 818.736
CS401HBA0	3 826 796.148	460 158.570	5 064 808.669
CS401HBA1	3 826 784.607	460 031.669	5 064 828.802
CS501LBA	3 825 626.175	460 641.786	5 065 640.512
CS501HBA	3 825 616.560	460 669.815	5 065 645.196
CS501HBA0	3 825 569.216	460 647.296	5 065 682.768
CS501HBA1	3 825 663.904	460 692.334	5 065 607.623
RS106LBA	3 829 261.821	469 161.961	5 062 137.050
RS106HBA	3 829 205.994	469 142.209	5 062 180.742
RS205LBA	3 831 438.959	463 435.116	5 061 025.206
RS205HBA	3 831 480.066	463 487.205	5 060 989.643
RS208LBA	3 847 810.446	466 929.381	5 048 356.961
RS208HBA	3 847 753.705	466 962.484	5 048 396.983
RS210LBA	3 877 847.841	467 456.599	5 025 437.344
RS210HBA	3 877 827.956	467 536.277	5 025 445.321
RS305LBA	3 828 721.154	454 781.087	5 063 850.822
RS305HBA	3 828 733.107	454 692.080	5 063 850.055
RS306LBA	3 829 792.203	452 829.524	5 063 221.330
RS306HBA	3 829 771.644	452 761.378	5 063 242.921
RS307LBA	3 837 941.343	449 560.431	5 057 381.027
RS307HBA	3 837 964.914	449 626.936	5 057 357.324
RS310LBA	3 845 433.443	413 580.563	5 054 755.909
RS310HBA	3 845 376.681	413 616.239	5 054 796.080
RS406LBA	3 818 468.029	451 974.278	5 071 790.337
RS406HBA	3 818 425.334	452 019.946	5 071 817.384
RS407LBA	3 811 596.257	453 444.359	5 076 770.170
RS407HBA	3 811 649.851	453 459.572	5 076 728.693
RS409LBA	3 824 756.246	426 178.523	5 069 289.608
RS409HBA	3 824 813.014	426 130.006	5 069 251.494
RS503LBA	3 824 090.848	459 437.959	5 066 897.930
RS503HBA	3 824 138.962	459 476.649	5 066 858.318
RS508LBA	3 797 202.513	463 087.188	5 086 604.779
RS508HBA	3 797 136.881	463 114.126	5 086 651.028
RS509LBA	3 783 579.528	450 178.562	5 097 830.578
RS509HBA	3 783 537.922	450 129.744	5 097 865.889

Table A.1. continued.

Station	ETRS-X (m)	ETRS-Y (m)	ETRS-Z (m)
DE601LBA	4 034 038.635	487 026.223	4 900 280.057
DE601HBA	4 034 101.901	487 012.401	4 900 230.210
DE602LBA	4 152 561.068	828 868.725	4 754 356.878
DE602HBA	4 152 568.416	828 788.802	4 754 361.926
DE603LBA	3 940 285.328	816 802.001	4 932 392.757
DE603HBA	3 940 296.126	816 722.532	4 932 394.152
DE604LBA	3 796 327.609	877 591.315	5 032 757.252
DE604HBA	3 796 380.254	877 613.809	5 032 712.272
DE605LBA	4 005 681.742	450 968.282	4 926 457.670
DE605HBA	4 005 681.407	450 968.304	4 926 457.940
FR606LBA	4 323 980.155	165 608.408	4 670 302.803
FR606HBA	4 324 017.054	165 545.160	4 670 271.072
SE607LBA	3 370 287.366	712 053.586	5 349 991.228
SE607HBA	3 370 272.092	712 125.596	5 349 990.934
UK608LBA	4 008 438.796	−100 310.064	4 943 735.554
UK608HBA	4 008 462.280	−100 376.948	4 943 716.600

Appendix B: LOFAR performance metrics

Table B.1. LOFAR primary beams.

Freq. (MHz)	λ (m)	D (m)	A_{eff} (m ²)	$FWHM$ (deg)	FOV (deg ²)	D (m)	A_{eff} (m ²)	$FWHM$ (deg)	FOV (deg ²)	D (m)	A_{eff} (m ²)	$FWHM$ (deg)	FOV (deg ²)
NL Inner			NL Outer			EU							
15	20.0	32.25	1284.0	39.08	1199.83	81.34	4488.0	15.49	188.62	65.00	3974.0	19.39	295.36
30	10.0	32.25	848.9	19.55	299.96	81.34	1559.0	7.75	47.15	65.00	2516.0	9.70	73.84
45	6.67	32.25	590.2	13.02	133.31	81.34	708.3	5.16	20.96	65.00	1378.0	6.46	32.82
60	5.00	32.25	368.5	9.77	74.99	81.34	399.9	3.88	11.78	65.00	800.0	4.85	18.46
75	4.00	32.25	243.6	7.82	47.99	81.34	256.0	3.10	7.55	65.00	512.0	3.88	11.81
NL core			NL Remote			EU							
120	2.50	30.75	600.0	4.75	17.73	41.05	1200.0	3.56	9.95	56.50	2400.0	2.59	5.25
150	2.00	30.75	512.0	3.80	11.35	41.05	1024.0	2.85	6.37	56.50	2048.0	2.07	3.36
180	1.67	30.75	355.6	3.17	7.88	41.05	711.1	2.37	4.42	56.50	1422.0	1.73	2.33
200	1.50	30.75	288.0	2.85	6.38	41.05	576.0	2.13	3.58	56.50	1152.0	1.55	1.89
210	1.43	30.75	261.2	2.71	5.79	41.05	522.5	2.03	3.25	56.50	1045.0	1.48	1.72
240	1.25	30.75	200.0	2.38	4.43	41.05	400.0	1.78	2.49	56.50	800.0	1.29	1.31

Notes. The full-width half-maximum (FWHM) in radians of a LOFAR Station beam is determined by $FWHM = \alpha\lambda/D$ where λ denotes the wavelength and D denotes the station diameter. The value of α will depend on the final tapering of the station. For these values, we have used a value of $\alpha = 1.1$ for LBA, and $\alpha = 1.02$ for HBA, as described in Sect. 12.5.

Table B.2. LOFAR angular resolution.

		Resolution			
Freq. (MHz)	λ (m)	$L = 320$ m (arcsec)	$L = 2$ km (arcsec)	$L = 100$ km (arcsec)	$L = 1000$ km (arcsec)
15	20.0	10 310.00	1650.00	33.00	3.30
30	10.0	5157.00	825.00	16.50	1.65
45	6.67	3438.00	550.00	11.00	1.10
60	5.00	2578.00	412.50	8.25	0.83
75	4.00	2063.00	330.00	6.60	0.66
120	2.50	1289.00	206.30	4.13	0.41
150	2.00	1031.00	165.00	3.30	0.33
180	1.67	859.40	137.50	2.75	0.28
200	1.50	773.50	123.80	2.48	0.25
210	1.43	736.70	117.90	2.36	0.24
240	1.25	644.60	103.10	2.06	0.21

Notes. The resolution of the LOFAR array is given by $\alpha\lambda/L$, where L denotes the longest baseline. The value of α depends on the array configuration and the weighting scheme used during imaging, i.e. natural, uniform, or robust. The values computed here assume a value of $\alpha = 0.8$ corresponding to a uniform weighting scheme.

Table B.3. LOFAR sensitivities.

Freq. (MHz)	λ (m)	Sensitivity			
		Superterp (mJy)	NL Core (mJy)	Full NL (mJy)	Full EU (mJy)
15	20.0
30	10.0	36	9.0	5.7	3.8
45	6.67	29	7.4	4.7	3.1
60	5.00	25	6.2	3.9	2.6
75	4.00	44	10.8	6.8	4.5
120	2.50	1.5	0.38	0.30	0.20
150	2.00	1.3	0.31	0.24	0.16
180	1.67	1.5	0.38	0.30	0.20
200	1.50	(2.5)	(0.62)	(0.48)	(0.32)
210	1.43	(2.5)	(0.62)	(0.48)	(0.32)
240	1.25	(5.6)	(1.4)	(1.1)	(0.73)

Notes. The quoted sensitivities are for image noise calculated assuming 8 h of integration and an effective bandwidth of 3.66 MHz (20 subbands). The sensitivities are based on the zenith SEFD's derived from 3C 295 in the Galactic halo as presented in Fig. 22. These values assume a factor of 1.3 loss in sensitivity due to time-variable station projection losses for a declination of 30 degrees, as well as a factor 1.5 to take into account losses for “robust” weighting of the visibilities, as compared to natural weighting. Values for 15 MHz have not yet been determined awaiting a good station calibration. Similarly values at 200, 210, and 240 MHz should be viewed as preliminary and are expected to improve as the station calibration is improved. The procedure for determining these values along with associated caveats are discussed in more detail in Sect. 12.6.