# Apertif: Phased array feeds for the Westerbork Synthesis Radio Telescope

## System overview and performance characteristics

W. A. van Cappellen[1], T. A. Oosterloo[1,2], M. A. W. Verheijen[2], E. A. K. Adams[1,2], B. Adebahr[3], R. Braun[4],
K. M. Hess[1,2], H. Holties[1], J. M. van der Hulst[2], B. Hut[1], E. Kooistra[1], J. van Leeuwen[1,5], G. M. Loose[1],
R. Morganti[1,2], V. A. Moss[6,7,1], E. Orrú[1], M. Ruiter[1], A. P. Schoenmakers[1], N. J. Vermaas[1], S. J. Wijnholds[1],
A. S. van Amesfoort[1], M. J. Arts[1], J. J. Attema[17], L. Bakker, C. G. Bassa[1], J. E. Bast[1], P. Benthem[1], R. Beukema[1],
R. Blaauw[1], W. J. G. de Blok[1,13,2], M. Bouwhuis[27], R. H. van den Brink [16], L. Connor[5,20], A. H. W. M. Coolen[1],
S. Damstra[1], G. N. J. van Diepen[1], R. de Goei[1], H. Dénes[1], M. Drost[1], N. Ebbendorf[1], B. S. Frank[8,12,13],
D. W. Gardenier[1,5], M. Gerbers[1], Y. G. Grange[1], T. Grit[1], A. W. Gunst[1], N. Gupta[24,1], M. V. Ivashina[11],
G. I. G. Józsa[8,9,10], G. H. Janssen[1,28], A. Koster[1], G. H. Kruithof[1], S. J. Kuindersma[1], A. Kutkin[1,18], D. M. Lucero[2,30],
Y. Maan[1], F. M. Maccagni[23], J. van der Marel[1], A. Mika[1], J. Morawietz[1], H. Mulder[1], E. Mulder[1], M. J. Norden[1],
A. R. Offringa[1,2], L. C. Oostrum[1,5,17], R. E. Overeem[1], Z. Paragi[22], H. J. Pepping[1], E. Petroff[5,14,15], D. J. Pisano[29],
A. G. Polatidis[1], P. Prasad[19], J. P. R. de Reijer[1], J. W. Romein[1], J. Schaap[1], G. W. Schoonderbeek[1], R. Schulz[1],
D. van der Schuur[1], A. Sclocco[17], J. J. Sluman[1], R. Smits[1], B. W. Stappers[21], S. M. Straal[25,26], K. J. C. Stuurwold[1],
J. Verstappen[2,1], D. Vohl[1], K. J. Wierenga, E. E. M. Woestenburg[1], A. W. Zanting, and J. Ziemke[1,31]

*(Affiliations can be found after the references)*

November 19, 2021

**ABSTRACT**

We describe the APERture Tile In Focus (Apertif) system, a phased array feed (PAF) upgrade of the Westerbork Synthesis Radio Telescope that transforms this telescope into a high-sensitivity, wide-field-of-view L-band imaging and transient survey instrument. Using novel PAF technology, up to 40 partially overlapping beams are formed on the sky simultaneously, significantly increasing the survey speed of the telescope. With this upgraded instrument, an imaging survey covering an area of 2300 deg$^2$ is being performed that will deliver both continuum and spectral line datasets, of which the first data have been publicly released. In addition, a time domain transient and pulsar survey covering 15 000 deg$^2$ is in progress. An overview of the Apertif science drivers, hardware, and software of the upgraded telescope is presented, along with its key performance characteristics.

**Key words.** telescopes – instrumentation: interferometers – surveys

## 1. Introduction

Since the start of radio astronomy, the study of the sky at radio wavelengths has provided a unique view of a diverse range of physical phenomena occurring in our own Milky Way, in external galaxies, and even out to the edge of the observable Universe. This has been done by tracing the radio continuum emission as well as the emission and absorption from radio spectral lines, in particular the 21 cm line of atomic neutral hydrogen (hereafter H i). This exploration of the radio sky could not have been done without the availability of large surveys. The legacy value of these surveys for the entire astronomical community is, therefore, huge (see e.g. Norris 2017 for an overview on continuum surveys and Giovanelli & Haynes 2015 for H i surveys).

A number of blind large-area high-resolution continuum surveys are currently available. The main ones, covering large fractions of sky, include the NRAO VLA Sky Survey (NVSS; Condon et al. 1998), the Faint Images of the Radio Sky at Twenty-Centimeters (FIRST; Becker et al. 1994), the Sydney University/Molonglo Sky Survey (SUMSS; Bock et al. 1999), the TIFR GMRT Sky Survey (TGSS; Intema, H. T. et al. 2017), the Westerbork Northern Sky Survey (WENSS; Rengelink et al. 1997), and, now ongoing, the LOFAR (van Haarlem et al. 2013) Two-metre Sky Survey (LoTSS; Shimwell et al. 2019), the LOFAR LBA Sky Survey (LoLSS; de Gasperin, F. et al. 2021), the Evolutionary Map of the Universe (EMU) survey (Norris et al. 2011), the Galactic and Extra-Galactic All-sky MWA (GLEAM) survey (Wayth et al. 2015), and the VLA Sky Survey (VLASS; Lacy et al. 2020). The list of topics they have enabled to be addressed is far too long to be described in detail, but they include star formation and nuclear activity in galaxies, magnetism across cosmic time, the life cycle of radio active galactic nuclei (AGN) and their impact on the evolution of galaxies, and detailed studies of individual galaxies and clusters of galaxies (see e.g. Norris 2017; Prandoni & Seymour 2015 and references therein).

Large-area H i surveys are more time-consuming to perform due to the low survey speed for spectral line work of existing interferometers. Therefore, the main surveys have been done with single-dish telescopes and have low spatial resolution. The largest H i surveys are the H i Parkes All Sky Survey (HIPASS; Barnes et al. 2001) and the Arecibo Legacy Fast ALFA Survey

(ALFALFA; Haynes et al. 2018). The main results from these surveys focus on the role of neutral hydrogen in galaxy evolution and the relation between gas content and star formation (e.g. Catinella et al. 2018; see Giovanelli & Haynes 2015 and Blyth et al. 2015 for overviews). These surveys have been performed with single-dish telescopes at relatively low spatial resolution. This limited the scope of the science questions they can address.

Although the above surveys led to many important results, one has to maintain the impetus of discovery by continuously improving the performance of telescopes so that new surveys with better characteristics can be performed. A number of recent technological developments opened up opportunities for major steps forward by increasing the survey speed of radio interferometers, enabling more powerful large radio surveys at high spatial resolution at frequencies around 1 GHz. Particularly relevant is the development of phased array feeds (PAFs) with low crosstalk and low noise. In addition, the increased processing power of modern computing hardware (e.g. for beam forming), together with the capability of transporting and handling very large data streams, made it possible to build low-noise PAF systems that combine broadband coverage at high frequency resolution and good sensitivity, with an increased field of view (FoV). This combination offers a significant improvement in survey speed, for both spectral line and continuum observations, providing unprecedented possibilities for carrying out a new generation of radio surveys. Such a step is also clearly a top priority on the path to the Square Kilometre Array (SKA) telescope, as can be seen from the SKA science case (Braun et al. 2015).

Taking advantage of these recent technological developments, the APERture Tile In Focus (Apertif) PAF system was developed for the Westerbork Synthesis Radio Telescope (WSRT) and completed in 2019. Apertif aims at expanding the capabilities of the telescope, notably allowing it to image large areas of the sky 39 times faster than before, with an increase in survey speed of a factor of 15 (see Sect. 12).

In this paper we provide an overview of the Apertif system, including key elements of its design and its performance characteristics. Section 2 provides a brief introduction of the WSRT. Section 3 provides a historical perspective on PAF technology and describes how the idea of Apertif was born. The science drivers and the system architecture of the upgraded WSRT are provided in Sects. 4 and 5. In Sect. 6 the radio frequency interference (RFI) environment is described, followed by a description of the PAF in Sect. 7, which includes the antenna array, receivers, and digital beam forming. The upgraded central processing (correlator and data writer) is presented in Sect. 8, and the monitoring and control (MAC) systems in Sect. 9. Sections 10 and 11 present the long-term archive and the operational model, followed by the measured system performance in Sect. 12.

## 2. The Westerbork Synthesis Radio Telescope

The WSRT was designed and built in the 1960s and was officially inaugurated in 1970. It was one of the first few telescopes that explored the use of dish-based interferometry. An overview of the rich history of discoveries and technical innovations of this radio telescope over the past 50 years can be found in Strom et al. (2018).

The WSRT consists of a linear east-west array of 14 identical steerable 25 m dishes on an equatorial mount, with a focal ratio f/D = 0.35 (Fig. 1). Instead of a solid surface, the reflectors use a metal mesh surface to reduce weight and wind load. The feed is located in a cabin at the primary focus of each reflector. Given that the WSRT is a linear array, every observation requires

a full 12-hour period in order to exploit the rotation of the Earth to image an area on the sky. The Apertif system uses 12 of the 14 dishes (see Sect. 12.8), with the remaining two dedicated to other projects, for example as part of the European VLBI Network (EVN) . The location of ten dishes is fixed. The remaining four are on a rail track and can be re-positioned to adjust the uv coverage of the interferometer, although this possibility is not used currently. The shortest baseline of the array is 36 m while the longest is 2412 m. At the observing frequency of Apertif this gives a spatial resolution of about $12 \times 12/\sin\delta$ arcsec with $\delta$ the declination of the observation. The FoV of a single WSRT dish with a single receiver has a full width at half maximum (FWHM) of about 35 arcminutes at 1.4 GHz. The installation of the Apertif system was combined with a major refurbishing (cleaning and painting) of the metal structures that support the dishes, and the electronics that drive the dish pointing motors and the angular encoders have been replaced.

## 3. Early phased array feed experiments on the WSRT

For many years, reflector antennas with multiple feeds have used electrically large feed antennas (such as horns), providing widely separated beams on the sky. Research by Ron Ekers and Rick Fisher in the late 1970s and early 1980s provided the notion that the electromagnetic fields in the focal plane of a reflector antenna carry the information to correct for distortions of the reflector surface, and/or observe a contiguous FoV that is much larger than the single beam provided by traditional feed antennas. However, the antenna technology to sample the complete electro-magnetic field in the focal plane and the processing algorithms to form multiple closely spaced, or overlapping, beams in relatively deep reflectors was not available then. In the mid-1990s, the Netherlands Institute for Radio Astronomy (ASTRON) initiated a research programme towards the SKA, which triggered research into dense aperture arrays, such as phased arrays with an electrical spacing smaller than half a wavelength. Reflector antennas with array feeds were not actively pursued in the ASTRON programme, but were identified as 'hybrid solutions'. Amongst others, Rick Fisher continued investigations into wide FoV reflector feeds, and how the output of these feeds should be processed. In 1998 ASTRON started a collaboration with Dan Schaubert (University of Massachusetts) on arrays of tapered-slot or 'Vivaldi' antennas. Originating from defense applications, such arrays demonstrated a constant effective area over a wide frequency range (Schaubert et al. 2000, 2003). Besides their application in aperture arrays, it was realised that the dense sampling enabled by the Vivaldi array would be a game-changer when placed as antenna array in the focal plane of a reflector as well (van Ardenne, 2020, priv. comm.). This application, combined with the research on the noise behaviour of such arrays (Weem 2001), formed the technological basis of the European research programme 'FARADAY' (2001–2004). The FARADAY project experimentally demonstrated that a densely packed PAF could form multiple overlapping high-efficiency beams on the sky, using a Vivaldi array and an Radio Frequency (RF) beam former (BF) at 2–5 GHz (Ivashina et al. 2002; Ivashina et al. 2004).

Triggered by the success of the FARADAY project, the PHAROS project (Simons et al. 2006) continued the WSRT PAF research with the objective of demonstrating low-noise PAF operation. For this purpose the first cryogenically cooled PAF was designed, operating in the 4–8 GHz band and combining 24 ac-

**Fig. 1.** Dishes of the WSRT.

tive elements into four high-efficiency RF beams (13 elements per beam).

The results of these early PAF experiments confirmed the potential of PAFs for the WSRT. A full-scale demonstration of PAF capabilities would require an advanced high-speed digital back-end, which led to the definition of a series of prototypes known as DIGESTIF (van Cappellen & Bakker 2010). The DIGESTIF system provided the ultimate demonstration of L-band PAF technology for Apertif on the WSRT.

## 4. Science drivers

Two major imaging surveys are performed with Apertif, a relatively shallow wide-area imaging survey (using a single 12-hour observation per pointing) covering an area up to 2300 deg$^2$, and a medium-deep survey (using $10 \times 12$ hour per pointing) covering up to 150 deg$^2$. Both these surveys will deliver continuum, polarisation, and spectral line datasets. The first data of these surveys have been publicly released[1] (see also Adams et al. in prep). In addition, a time domain transient and pulsar survey covering 15 000 deg$^2$ is performed.

The areas covered by the Apertif surveys were primarily chosen such that there is good overlap with regions that have important publicly available ancillary spectroscopic and imaging data, such as the data from the Sloan Digital Sky Survey (SDSS)[2], the Panoramic Survey Telescope and Rapid Response System (Pan-STARRS)[3] survey, the Hobby-Eberly Telescope Dark Energy Experiment (HetDex)[4], the Herschel-Atlas project[5], and the Mapping Nearby Galaxies at APO (MaNGA)[6]

---

[1] http://hdl.handle.net/21.12136/B014022C-978B-40F6-96C6-1A3B1F4A3DB0
[2] www.sdss.org
[3] www.ifa.hawaii.edu/pswww
[4] www.hetdex.org
[5] www.h-atlas.org
[6] www.sdss.org/surveys/manga

and the WHT Enhanced Area Velocity Explorer (WEAVE)[7] spectroscopic surveys. However, because the imaging surveys are focused on extra-galactic science, regions of low Galactic latitude are avoided. This, combined with the fact that in order to obtain complete *uv* coverage with Apertif, observations over the full hour angle range from –6 h to + 6 h are needed, efficient scheduling of the survey observations puts strong constraints on the location of the survey areas. As a result, for part of the area covered by the imaging surveys only limited additional data from other wavebands are available. In addition, survey observations are scheduled to be 11.5 h in duration instead of the standard 12 h, to allow for calibration and telescope movement in between observations.

Below we briefly describe the main goals of the imaging surveys. More details can be found in Hess et al. (in preparation), which also describes the observation scheduling, and the associated constraints.

### 4.1. H I and OH surveys

The two recent extragalactic H I emission surveys, HIPASS and ALFALFA, detected over 25 000 galaxies and are a major resource for statistical studies of the H I properties of galaxies. However, these surveys were carried out with single-dish telescopes and, although sensitive, they have very limited angular resolution and therefore the large majority of objects detected are spatially unresolved. The Apertif H I surveys aim to make the next step by providing the spatially resolved H I properties for a large number of galaxies for studies of galaxy structure and evolution.

Detailed spatially resolved studies of the H I emission in galaxies has been one of the key science areas of the WSRT since the 1970s and these have greatly contributed to improving our understanding of galaxies. However, although these studies

---

[7] www.ing.iac.es/weave/about.html

have been very successful, they have been limited to targeted observations of relatively small samples of galaxies in the nearby Universe, with the most recent examples being the WHISP (van der Hulst et al. 2001), the Atlas[3D] (Serra et al. 2012), and the HALOGAS (Heald, G. et al. 2011) surveys.

The goal of Apertif is to expand such imaging studies to much larger, blindly selected samples of galaxies in the nearby Universe. This is possible thanks to the improved survey speed, the high spatial resolution and the high sensitivity of Apertif. The $5\sigma$ column density detection limit reached in the data cubes of the wide-area survey at full resolution for a profile width of 20 km s$^{-1}$ is ~1–2 ×10$^{20}$ cm$^{-2}$ (depending on declination) and is just below 10$^{20}$ cm$^{-2}$ for the medium-deep survey. Because of the specific array layout of the WSRT, at somewhat lower resolution (30″) the column density sensitivities are about a factor of three better. The sensitivity and resolution of the Apertif wide-area H I survey are similar to that of the wide-area H I survey planned with the Australia Square Kilometre Array Pathfinder (ASKAP; Koribalski et al. 2020).

One of the main aims of the H I emission survey is to better understand the processes that govern the removal and accretion of gas for galaxies of all types and sizes and the role this plays in the evolution of galaxies (see e.g. Sancisi et al. 2008 for a review). This requires spatially resolving the H I properties of galaxies and the gas around them in different environments, including isolated areas such as voids, the numerous groups of galaxies that exist in the nearby Universe, and the most extreme galaxy densities in clusters. The observed properties of galaxies strongly depend on the cosmic environment in which they are located, implying that the environment plays an important role during a galaxy's formation as well as during the subsequent evolution. A variety of physical mechanisms affect the growth and shaping of galaxies such as accretion, merging, tidal interactions, and ram-pressure stripping. Resolved H I emission studies are particularly sensitive to trace these processes and an example of such a case found with Apertif is shown in Fig. 2. The availability of such information for a large number of galaxies in different environments will provide important new information on how galaxies evolve.

Another main topic for the Apertif H I survey is the search for the smallest gas-rich galaxies in the local Universe. Such objects are sensitive probes of the limits of galaxy formation. Very few such objects are known (e.g. Leo T, Ryan-Weber et al. 2008, and Leo P, Giovanelli et al. 2013). The visible mass of such systems is completely dominated by H I and only a small number of stars is present. This is possibly the result of such galaxies being almost too small to compress and cool the gas to form stars and suggests that there is a lower limit to the size of a galaxy that can form stars. How such objects formed and evolve, and whether indeed lower mass limit for star bearing galaxies exists, critically depends on the detailed conditions of normal and dark matter in the early Universe, and the predicted number of such small gas-rich galaxies in the local Universe differs greatly between different theories of galaxy formation. It is therefore important to improve our statistical knowledge on the occurrence of such systems. The large area covered by the H I surveys is particularly suited for making an inventory of these objects in the local Universe, which will provide important constraints of theories of galaxy formation.

The imaging surveys will also provide information on H I observed in absorption against relatively strong continuum sources. Given the observed frequencies (and, therefore, redshifts), Apertif is particularly relevant tracing H I in absorption in the circumnuclear regions of AGN (see Morganti & Oosterloo 2018 for a
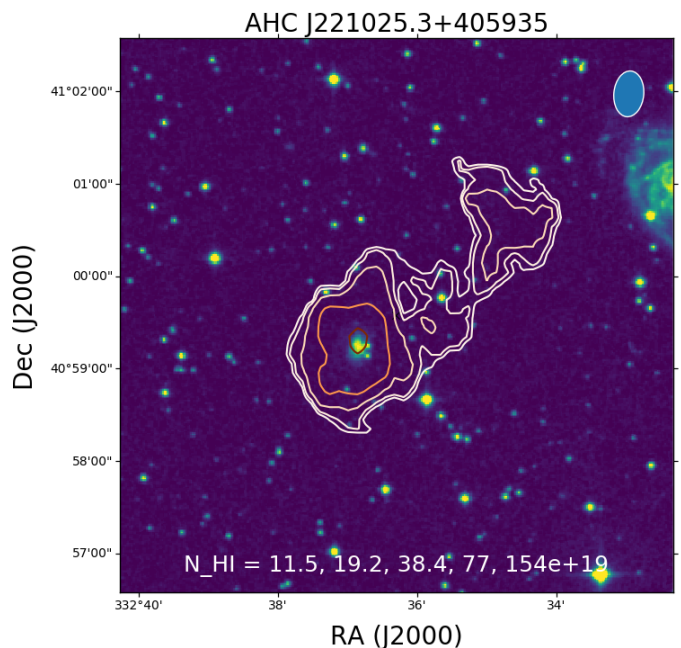
**Fig. 2.** Example of H I emission (contours) detected in a small galaxy group (colour scale) in one of the Apertif pointings of the H I wide-area survey (Hess et al. in prep). The contour levels correspond to 3, 5, 10, 20, and 40 times the noise level.

review). The presence and kinematics of the H I in these objects can trace circumnuclear disks, gas accreting and fuelling the central super massive black hole, and the presence of gas outflows that can be driven by radio jets (see e.g. Morganti et al. 2005). These gas outflows can provide information on how the AGN is impacting galaxy evolution. The WSRT has already given important contributions to such studies (e.g. Morganti et al. 2005; Geréb et al. 2015) and also provided one of the largest available surveys of H I absorption in the nearby Universe (Maccagni et al. 2017). However, these studies are based on pointed observations of radio sources for which the redshift is known. Such a pre-selection can introduce biases in the results (see e.g. Chowdhury et al. 2020). Hence, the main step forwards by Apertif will be to instead provide a blind inventory of the presence of H I in absorption without this a priori target selection. Thus, it will be possible to inspect continuum radio sources detected by Apertif for absorption regardless of other source properties.

The sensitivity for H I absorption depends on the strength of the background radio continuum and the optical depth of the H I line, which in turn depends on the spin temperature and the column density of the H I gas. The shallow H I survey will reach an optical depth of $\tau \sim 0.03$ (for 30 km s$^{-1}$ velocity resolution) against continuum sources stronger than ~100 mJy beam$^{-1}$, well suited for tracing atomic neutral hydrogen in AGN (see e.g. Morganti & Oosterloo 2018). For brighter sources, Apertif will also allow searches for broad shallow absorption, such as those indicating the presence of fast outflows (Morganti et al. 2005; Geréb et al. 2015).

The main limitation will come from the relatively hostile environment with regard to RFI at the lower frequencies covered by Apertif (i.e. below 1290 MHz; see Sect. 6). The search for H I absorption out to $z = 0.25$ is expected to deliver a few hundred new detections of gas associated with radio galaxies, greatly expanding the parts of parameter space of radio AGN where the presence of H I absorption can be studied. A comparison with
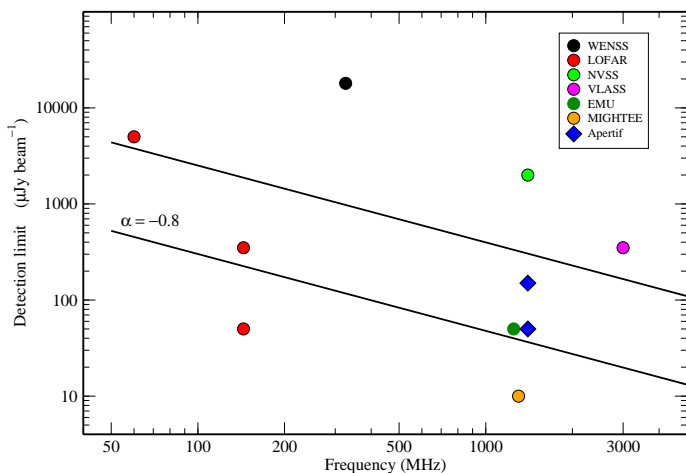
**Fig. 3.** Comparison of the $5\sigma$ detection limits of a number of radio continuum surveys and the Apertif surveys. The lines indicate a spectral index $\alpha = -0.8$, which is typical of extragalactic radio sources. The noise for the EMU survey was taken from Norris (2017). The two LOFAR HBA points refer to LoTSS (Shimwell et al. 2019) and to the LOFAR deep fields (Tasse et al. 2021). The latter cover limited areas (of a few tens of square degrees). The MIGHTEE survey, although very deep, only covers a limited number of square degrees.

other H I absorption surveys from SKA pathfinders and precursors is discussed in Maccagni et al. (2017) and Morganti & Oosterloo (2018).

Finally, Apertif also allows interesting large-area surveys to be conducted for OH megamasers, and in particular of their redshift evolution. The main spectral lines of OH megamasers are emitted at 1665 and 1667 MHz, with two satellite lines at 1612 and 1720 MHz. Therefore, given the observing band of Apertif, the main lines can be detected out to redshift $z \sim 0.3$. OH megamasers are found in luminous- and ultra-luminous infrared galaxies and are associated with galaxy mergers and their intense nuclear star formation activity. This, and the fact that the redshift of the OH megamaser comes automatically from a detection, make them a useful tracer of the merger history in the Universe. In addition, with the broad observing band of Apertif it is also possible to measure the line rations between the main and the satellite lines for a subset of the sources, which gives useful information about the physical conditions in the star forming regions. The results on the first OH megamaser found with Apertif have been published (Hess et al. 2021).

### 4.2. Continuum surveys

The images and source catalogues resulting from large-area radio continuum surveys like those listed in the introduction, are among the most important databases for many branches of astronomy. The Apertif continuum surveys bring an important addition to this list. A comparison between various surveys is shown in Fig. 3. The continuum surveys with Apertif will provide images and catalogues at 1.4 GHz with a resolution of $12''/\sin\delta$ and reaching, in the case of the wide-area survey, a noise level about 20 times lower than NVSS and with 3 times higher angular resolution (see Sect. 12.2). An example of the image obtained from a single Apertif pointing is shown in Fig. 4.

The data from the Apertif continuum surveys can be used to address a wide range of topics such as galaxy evolution, source counts, the life cycle of radio galaxies and the polarised properties of radio sources. Furthermore, for star-forming galaxies, the

combination of the radio continuum and the H I emission, which are simultaneously obtained, will allow star formation and gas properties to be connected.

A unique aspect is the synergy with the LOFAR surveys, both the LoTSS at 150 MHz (Shimwell et al. 2019) as well as the LOFAR 60 MHz survey (de Gasperin, F. et al. 2021), because both LOFAR and Apertif observe in the northern sky and there is large overlap between their survey areas. A fortunate circumstance is that the relative sensitivities of LOFAR and Apertif match the spectrum of typical radio sources and the spatial resolutions are similar (see Fig. 3). The combination Apertif-LOFAR will be particularly useful for identifying and studying radio sources in extreme phases of their evolution, such as remnants and restarted radio sources, and diffuse relic and halo emission in cluster of galaxies. For resolved sources, this will instantly yield spectral index and spectral curvature maps, which are a very rich source of information for constraining many physical parameters and processes and are powerful tools for studying radio sources with extreme radio spectra.

A proof of concept study of work based on combining Apertif data with LOFAR images is presented by Morganti et al. (2020). A significant number of radio sources in the remnant and restarted phases were found. This provided new constraints on the timescales of the duty-cycles of AGN activity and confirmed results from numerical simulations. Given that this was based on only one Apertif pointing, in the Lockman Hole region, it already shows the potential of combining radio images at 150 and 1400 MHz.

The continuum images from the medium-deep survey will reach the confusion limit and this is particularly interesting, among other topics, for investigating radio source counts. Similar studies were carried out using the WSRT (e.g. Prandoni et al. 2018; Bonato et al. 2020), but only for limited areas. Radio source counts are tightly related to the evolutionary properties of the sources (Padovani 2016; Prandoni et al. 2018; Norris 2017) and are mostly tied to the faint end of the source population. Up to now, deep observations investigating this population were limited to small fields where local density variations may influence the results. The large survey area of the Apertif surveys will allow mitigation of these local effects.

Polarised radio continuum emission contains the information about the magnetic field within the observed sources and along the line-of-sight between the observer and the sources themselves. The two components of magnetic field, those perpendicular to the line-of-sight and parallel to it, can be recovered using the observed polarisation angle and the rotation measure, respectively. These two quantities allow a reconstruction of the three-dimensional morphology of the magnetic field. The accuracy of the reconstructed magnetic field properties is dependent on the number of detected polarised sources per area on the sky. The most recent reconstruction of the Milky Way magnetic field structure by Oppermann et al. (2012) is based on the NVSS and therefore limited by the detected number density of sources ($\sim 1/\deg^2$; Taylor et al. 2009). First analysis of pointings from the Apertif wide-area survey shows number densities of 30-40 sources/$\deg^2$ allowing an improved reconstruction of the Milky Way magnetic field and of a number of large extragalactic nearby objects.

Additional information on the magnetic field along the line-of-sight and in the polarised sources themselves is their polarised spectrum. It is known that most sources become severely depolarised below 1 GHz. Most models and analytical calculations suggest a $\lambda^2$ to $\lambda^4$ dependence of the degree of polarisation on wavelength $\lambda$ (Sokoloff et al. 1998), but accurate spectra have
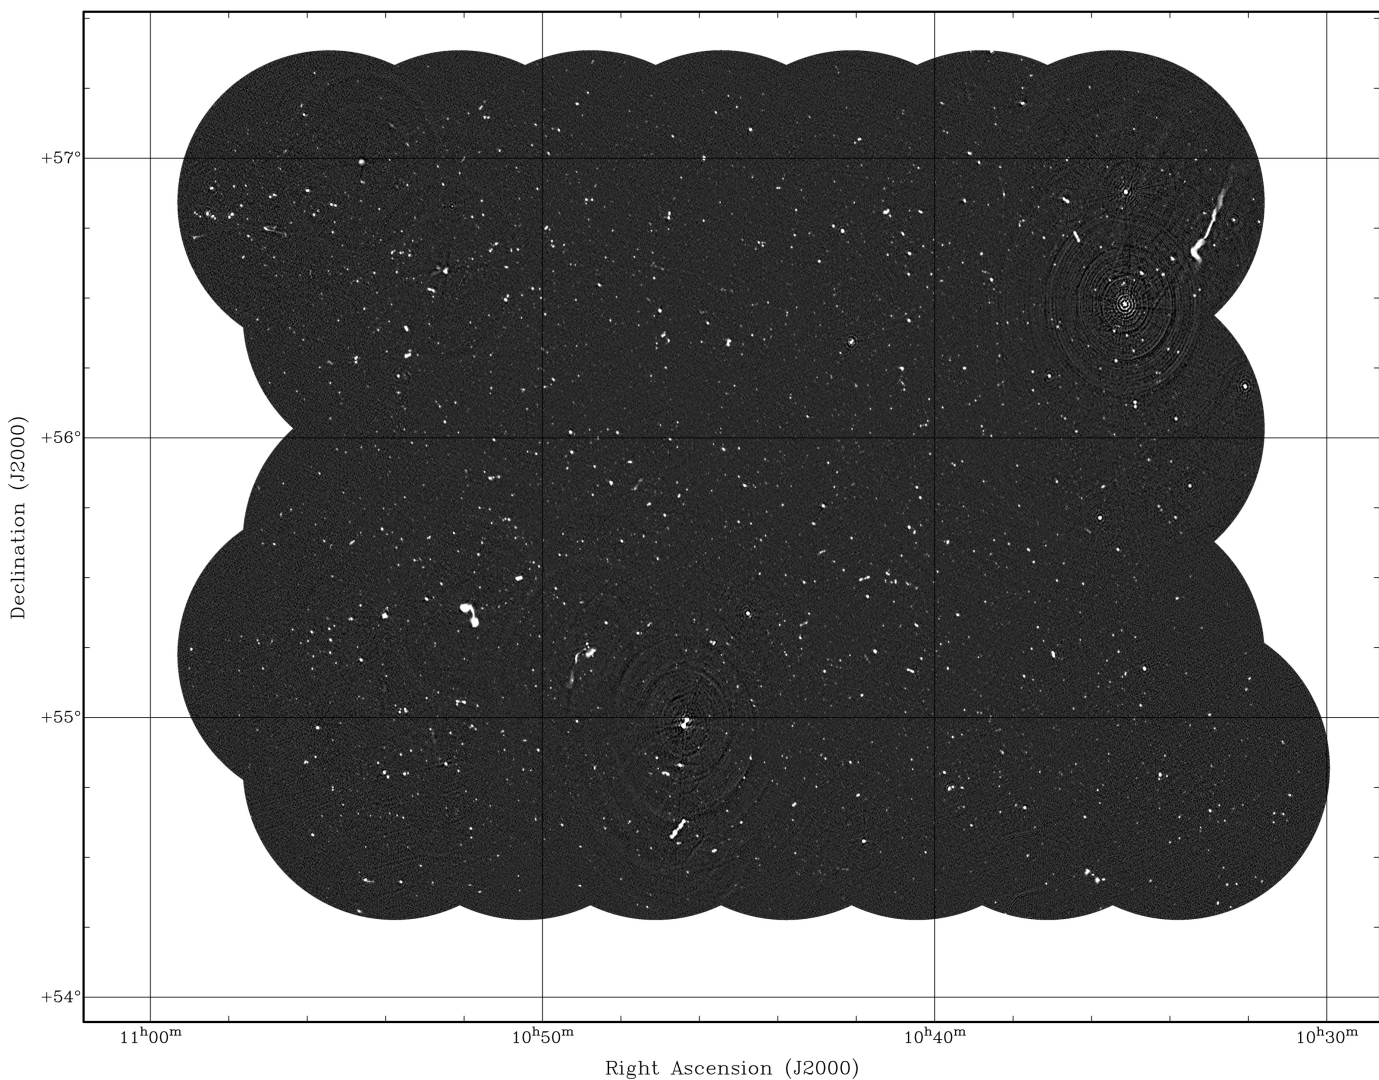
**Fig. 4.** Continuum mosaic obtained from one 11.5-hour Apertif observation, centred on the Lockman Hole region. A number of extended sources are visible (including a giant radio galaxy at the top right). The noise level is about 35 $\mu$ Jy beam$^{-1}$, except close to a few very bright sources, where calibration artefacts are still visible (see Sect. 12 for more details).

only been measured for a few sources (Schnitzeler et al. 2015). The synergy of the LoTSS survey with the Apertif wide-area survey allows a first statistical analysis of the polarised spectrum of radio sources. Due to the strong depolarisation at LOFAR wavelengths, all polarised sources detected in LoTSS will be detected by Apertif. The combination of this information will for the first time allow an investigation of the depolarisation mechanisms and the magnetic field properties causing it.

Information about the evolution of the magnetic field over cosmic timescales is contained in polarised source counts. Currently, such analyses are limited to the wide and shallow data of the NVSS (Mesa et al. 2002; Tucci et al. 2004; Stil et al. 2014) or to small fields of a few square degrees with sensitivities comparable to those obtainable with Apertif (Taylor et al. 2007; Grant et al. 2010; Subrahmanyan et al. 2010; Hales et al. 2014). All of these referenced studies show a rise in the degree of polarisation towards fainter sources, but the scatter is of the order of a few in fractional polarisation and number of detected sources. The measurements are most likely influenced by the small number of detected sources and by local variations due to the small observing area. The Apertif surveys will allow us to trace a sufficient number of polarised sources ($\sim 10^5$) over several thousand

square degrees down to flux densities of several tens of microjanskys. This allows for the first time an accurate characterisation of the polarised sky. To study the evolution of the magnetic field over cosmic timescales, information on the redshifts of sources is needed. The synergy of the WEAVE survey with the Apertif and LOFAR radio surveys, in combination with the number of sources Apertif will detect, allows us to investigate redshift-dependent polarised source counts. This is a huge step in determining the evolution of the magnetic field up to redshifts of $z \approx 2$.

### 4.3. Time domain surveys

Many transient and time variable phenomena are observable in the radio sky, including radio afterglows of supernovae, gamma-ray bursts and gravitational-wave events, on ~month timescales. Its wide FoV and high sensitivity make Apertif a powerful instrument for finding such imaging transients (Boersma et al. 2021). Faster transients include solar radio bursts lasting seconds to minutes, and radio pulses from Galactic pulsars, some as short as microseconds. However, the most recent and tantalising addition to this time-domain transient list was the discovery of bright (a

few Jy), short (of order 1 ms) radio pulses of extragalactic origin by Lorimer et al. (2007), sources now known as fast radio bursts (FRBs).

Since their discovery in 2007, the study of FRBs has become a well-developed field in its own right, with more than 100 sources now reported (see Petroff et al. 2019, for a review). The origins of these radio bursts are not yet known, but theories invoking young, highly magnetised neutron stars, black holes, and the explosions involved in their birth have gained traction (for a theory review, see Platts et al. 2019). A subset of the published FRB sources has been observed to repeat, supporting stable or non-cataclysmic progenitor theories (Spitler et al. 2016; CHIME/FRB Collaboration et al. 2019; Fonseca et al. 2020).

The first FRBs were serendipitously discovered in all-sky surveys for radio pulsars. The systems necessary to find FRBs require high time resolution ($\lesssim$ 1ms) and at least modest fractional bandwidth ($\Delta \nu / \nu \gtrsim 0.25$). The Berkeley Parkes Swinburne Recorder (BPSR) employing single pulse search software on graphics processing units (GPUs) at the Parkes radio telescope proved particularly effective in early FRB searches (Keith et al. 2010; Thornton et al. 2013; Champion et al. 2016).

Increasingly, dedicated search hardware and software have been developed with the primary goal of discovering (and in some cases precisely localising) large numbers of new FRB sources; many with the stated goal of discovering new FRB pulses in real time. The majority of these systems employ GPUs, field-programmable gate arrays (FPGAs), and other specialised hardware on dedicated compute clusters to perform the most computationally expensive tasks, such as de-dispersion (Siemion et al. 2012; Caleb et al. 2016; CHIME/FRB Collaboration et al. 2018; Law et al. 2018).

Despite the rapid advance of the field in recent years, many outstanding challenges remain in understanding the FRB phenomenon. Despite the high all-sky rate of FRBs of several thousand per day (Lawrence et al. 2017), many early instruments (such as Parkes) reported a new event only every few months due to the limited FoV inherent to single dishes or single receivers. Some of the most basic properties of the FRB population remain unknown; the underlying distributions of the population in pulse duration, dispersion measure, distance, energy, and spectral structure have all been difficult to pin down. This is partly due to a small sample size, but also in part due to insufficient instrumental resolution in the case of the pulse duration, dispersion measure, and spectral structure distributions.

Recent research was able to preserve data that capture the Stokes parameters of new FRBs and interpret the polarisation properties of individual bursts for a growing fraction of the population. Many of these bursts were highly linearly polarised (Michilli et al. 2018; Bannister et al. 2019; Fonseca et al. 2020), and analysis of the linear polarisation in some cases revealed the presence of a strong magnetic medium local to the FRB source (Michilli et al. 2018). Polarisation may be critically important for understanding the emission and local environment of FRBs. However, preserving polarisation information requires a survey to either keep full-polarisation data for all survey observations, to be analysed post-detection; or to search data in real-time while the full-polarisation data remains in a memory buffer, to be excised and saved for a burst. Each choice comes with its own set of challenges for data storage and real-time data processing power.

A subset of the FRB population has been observed to repeat, producing multiple bursts at the same (or very similar) dispersion measure, with some sources now observed to repeat over many years (Spitler et al. 2016; Oostrum et al. 2020). Sustained follow up of an FRB source is needed to eventually detect re-

peats. However, at present only 20 repeating sources are published and the overall fraction of repeating sources in the FRB population is unknown. Whether all FRBs eventually repeat is an open question, one that is being tackled currently both with observational efforts (Fonseca et al. 2020) and modelling (Gardenier et al. 2019, 2021).

In addition to the challenges of understanding the underlying population(s) of FRBs, their physical properties, and their progenitors, there are also technical challenges involved in their discovery at scale. With next generation telescopes such as the SKA it will no longer be possible to preserve the raw survey data for offline searches (Macquart et al. 2015). Instead, new sources will need to be identified in real time to capture the telescope data for later analysis. New automated FRB search techniques and pipelines taking advantage of classification and machine learning tools are being developed to prepare for this future reality (Connor & van Leeuwen 2018).

To address all these challenges, new FRB search efforts are increasingly employing interferometers to survey the sky (Caleb et al. 2017; Bannister et al. 2017; Maan & van Leeuwen 2017; Law et al. 2018; CHIME/FRB Collaboration et al. 2018). Interferometers, coherently or incoherently combining signals from many smaller elements or dishes, have the advantage of a large instantaneous FoV. Recent technological advances have resulted in new receivers such as PAFs, which place many dipoles at the focus of each dish of a telescope array.

One of the largest challenges of interferometric radio astronomy has always been computation. Beam forming within the telescope FoV requires a powerful correlator to combine the signals from all elements in phase. This is more difficult still when combining the multi-element PAF systems to form beams on the sky. Forming coherent beams and searching the time stream of each for impulsive radio signals such as FRBs provides an added technical challenge.

Faster and more agile processing units available in recent years made it possible to form more beams and search them quickly, in some cases in real time. These searches still require large compute clusters to deal with the massive amounts of data streaming from the telescope and distribute it over many processing nodes. Many FRB search efforts, including the searches described here with Apertif, now house dedicated computing clusters on-site to search the data in real time for pulses.

By combining different elements of the array and the feed, the larger FoV of an interferometric array can be sampled by many smaller beams, enabling a much more precise localisation of any new source. The instantaneous localisation ability of an interferometer depends on the length of the longest baseline, which in the case of WSRT is long in the E-W direction but very short in the N-S direction, but even more precise localisation is possible for brighter signals that appear in several beams, as demonstrated in, for example, Connor et al. (2020).

Real-time searches are also critical in searching for prompt emission from the FRB source in other wavelength regimes, but at other radio frequencies as well. Previous multi-wavelength searches for related emission following real-time FRB detections were unsuccessful. However, these early efforts triggered follow-up several hours after the initial FRB (Petroff et al. 2015).

It is still unknown how broadband FRB pulses can be, and down to what radio frequencies they are detectable. In simultaneous FRB observations between Apertif and LOFAR, FRB emission is now seen down to the bottom of the LOFAR high-band antenna (HBA) band (Pastor-Marazuela et al. 2021). Triggered searches at such low frequencies, even for one-off FRBs, are made possible by the delay that the interstellar dispersion

introduces between the 1.4 GHz and 150 MHz arrival times of the same pulse. Exploiting this delay, which is of order minutes (Maan & van Leeuwen 2017), allows for studies of, for example, the FRB emission mechanism, and the plasma environment local to the host. Such coincident searches do, however, require real-time classification and alerting.

An ideal observing setup to tackle all of these challenges would combine large FoV with high spatial resolution on the sky, for a high rate of localised FRBs. To address as many of the population unknowns as possible, such a system should also be able to resolve FRBs in time and frequency and capture polarisation information. The Apertif Radio Transient System (ARTS) is designed with these considerations in mind. Apertif provides the large instantaneous FoV, while the WSRT interferometer allows for precise spatial resolution. ARTS processes a bandwidth of 300 MHz centred on 1370 MHz.

ARTS also addresses some of the challenges above by combining a range of innovations: first, delivering high time and frequency resolution over a frequency range of 1220 MHz to 1520 MHz; second, the ability to capture full-Stokes polarisation data with a new FRB search pipeline and third, a machine learning classifier to better identify and trigger on FRB candidates.

### 4.4. Serendipitous discovery

Apart from the science planned for Apertif, its large FoV implies that the potential for detecting rare objects or making serendipitous discoveries is high. This is illustrated by the fact that in the survey data taken so far, 16 intra-hour variable radio sources were found serendipitously (Oosterloo et al. 2020). Such variable sources are among the rarest objects on the sky. They are extremely small sources that are made to 'twinkle' because their radiation passes through turbulent foreground plasma located in the solar neighbourhood and they give a unique view of the structure and conditions of the nearby ionised interstellar medium. They can change a factor of five in flux on timescales of minutes and are easily recognised in Apertif images. Before Apertif, searches with various radio telescopes revealed only a handful of such sources over the entire sky and Apertif has quadrupled the number of such rare objects known in less than a year, illustrating the power of a large FoV.

## 5. Apertif system overview

The Apertif upgrade replaces the well-known multi-frequency frontend (MFFE; Tan 1991), the Intermediate Frequency (IF) system, the frequency and clock distribution system, the correlator, most of the cabling, and all software to monitor, control, and operate the system. Figure 5 shows the top-level block diagram of the WSRT-Apertif system. Apertif samples the electromagnetic (EM) field in the focal plane of the reflectors with a densely packed array of antenna elements. The spacing between the elements is 10 cm, corresponding to $0.38\lambda$ at the lowest and $0.58\lambda$ at the highest operating frequency. After filtering and amplification, the RF signals are transported over coaxial RF cables to a telescope cabin located under every dish. In the cabin, the RF signals are converted to baseband, digitised and combined into 40 dual-polarised compound beams (CBs). To form one CB on the sky, the signals received by antenna array elements are weighted and added in a BF. Apertif uses a bi-scalar BF: The signals of all array elements of the same polarisation are used for every CB in that polarisation. Because the BF is implemented in the digital domain, multiple copies of the input signals are easily

made. Using many parallel digital BFs that each apply different weights, many CBs (each pointing in a slightly different direction) are formed simultaneously. This principle is illustrated in Fig. 6. In Apertif, every dish delivers 40 CBs. The maximum electronic scan angle is determined by the size of the PAF, the focal length and the diameter of the main reflector. The dimensions of the Apertif PAF are the largest possible in the prime focus cabin of the WSRT reflectors, and provide a good balance between the scan range (±1.5 deg in all directions), and the capital costs. The instantaneous bandwidth of Apertif, 300 MHz, is defined by the filters in the analogue electronics. The number of simultaneous beams is defined by the digital processing capacity.

The partitioning of the system was driven by the desire to build a feasible, low risk, and easy to maintain system. Considering the constraints on physical size, weight and power consumption, only the RF part of the feed is located in the feed cabin at the primary focus of the reflector. Consequently, the weight of the systems in the feed cabin is reduced from about 300 kg before the Apertif upgrade, to only 50 kg. The telescope cabin under the dish is relatively large. This lowers the required level of integration of the RF and digital backend, reducing complexity and development risks. It also allows easy access to the equipment for maintenance and repair. By using a Faraday cage inside the telescope cabin, the interference generated by the receivers and backends can be very effectively shielded from the antennas. Inside the Faraday cage, all RF parts are shielded to avoid self-generated RFI of the digital backend, and its power supply, leaking into the system. A point of attention in this architecture are the phase variations in the coaxial RF cables that connect the feed antenna to the telescope cabin.

Time and frequency references are provided by a maser and a GPS receiver. The central correlator, with an aggregated input data stream of 3.456 Tbps, correlates every beam of all dishes into visibilities. The visibilities are integrated and stored in the data writer. This is the end of the real-time system. The resulting datasets are sent to ASTRON's headquarter in Dwingeloo where an automated calibration and imaging pipeline produces calibrated image cubes. Both the calibrated cubes and the raw visibilities are stored into the Apertif long-term archive (ALTA), which is the main interface to the users. Alternatively, the signals from the dishes are tapped-off in the correlator and processed by an independent GPU-based backend for tied-array observations. (van Leeuwen 2014; van Leeuwen et al. 2021).

The beam pattern of the Apertif feed offers a great level of control, enabling aperture efficiencies of 75% and higher, and control over sidelobes and cross-polarisation levels. As described by van Cappellen & Ivashina (2011), accurate knowledge of the beam patterns on the sky is crucial for synthesis radio telescopes to produce high dynamic range, noise limited, images. Receiver gain variations in the PAF can lead to variations in the CB and consequently a calibration scheme is required. An online calibration subsystem was designed to measure the analogue gain variations during an observation and to compensate the variations in the (digital) BF to stabilise the beam on the sky. At the time of writing, the online calibration system is not fully implemented yet because the PAF system is inherently stable enough between two calibrations (i.e. a period of 2 − 3 weeks).

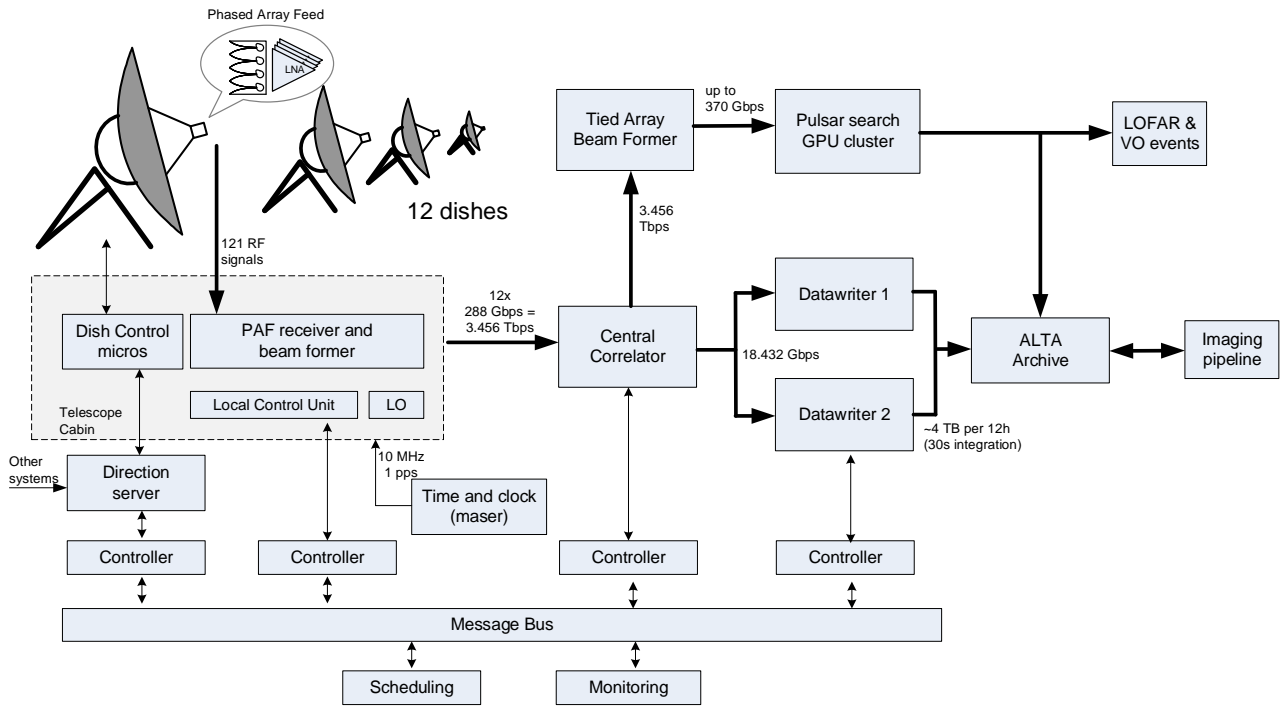Table 1 lists the main performance characteristics of the WSRT-Apertif system.

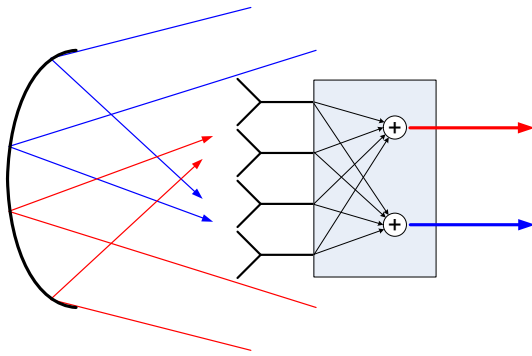**Fig. 5.** Overview of the Apertif–ARTS system.



**Fig. 6.** Concept of operation of a dense PAF.

**Table 1.** Key performance characteristics of the WSRT-Apertif.

| Characteristic | Value | Unit |
|---|---|---|
| Frequency range | 1130–1750 | MHz |
| Processed bandwidth | 300 | MHz |
| Number of beams | 40 | |
| Field of view single observation [a] | 10.5 | deg$^2$ |
| Effective field of view for surveys [a] | 5.25 | deg$^2$ |
| Sensitivity [a] | 63 | m$^2$K$^{-1}$ |
| SEFD [a] | 43.7 | Jy |
| Number of dishes | 12 | |
| Baseline lengths | 36–2700 | m |
| Spectral resolution | 12.2 | kHz |

[a] At 1400 MHz.

## 6. Impact of RFI

Figure 7 shows a measured spectrum at the WSRT site with a low-gain antenna towards the horizon. The main RFI sources are the Global System for Mobile communications (GSM) bands around 850 MHz and 1850 MHz, as well as five terrestrial TV bands below 800 MHz of the TV tower at Smilde, located 12.6 km west of the WSRT. In particular these TV bands proved to be troublesome. The TV stations use a modulation schema using a large number of sub-carriers spread over an 8 MHz band. From measurements in a WSRT dish, the expected worst case power from these signals at the input of the first low-noise amplifier (LNA) is approximately -50 dBm. Although these interferers are outside the Apertif frequency band of operation, they necessitate filters throughout the entire RF system to reject the creation of second- and third-order intermodulation products. These filters lead to an increase in the noise figure. Without these filters, the noise floor in the system would not be determined by thermal noise, but by a noise floor created by these intermodulation products. At the WSRT site, emission from high-power ADS-B aircraft secondary surveillance radar operating at 1090 MHz is also present. Its peak power is much higher than shown in Fig. 7 because it is pulsed.

The frequency range between 1300 and 1500 MHz is relatively RFI free. At frequencies below 1300 MHz significant parts of the band are badly affected by RFI, in particular for the shorter baselines. The RFI is dominated by Galileo bands E5 and E6, GPS L2, and the GLONASS global navigation satellite system, making much of this frequency range unusable. Only a few small windows are relatively free of RFI, for example the frequency range 1180 – 1195 MHz. In the frequency range 1452 – 1492 MHz RFI is present due to 5G transmitters in the area.

## 7. Phased array feed

This section describes the PAF analogue and digital hardware installed at every dish, delivering 40 dual polarised beams to the central systems. An overview of the PAF system at every dish is show in Fig. 8. The major design driver for the antenna and RF system is the wide frequency band of operation and the required noise figure.
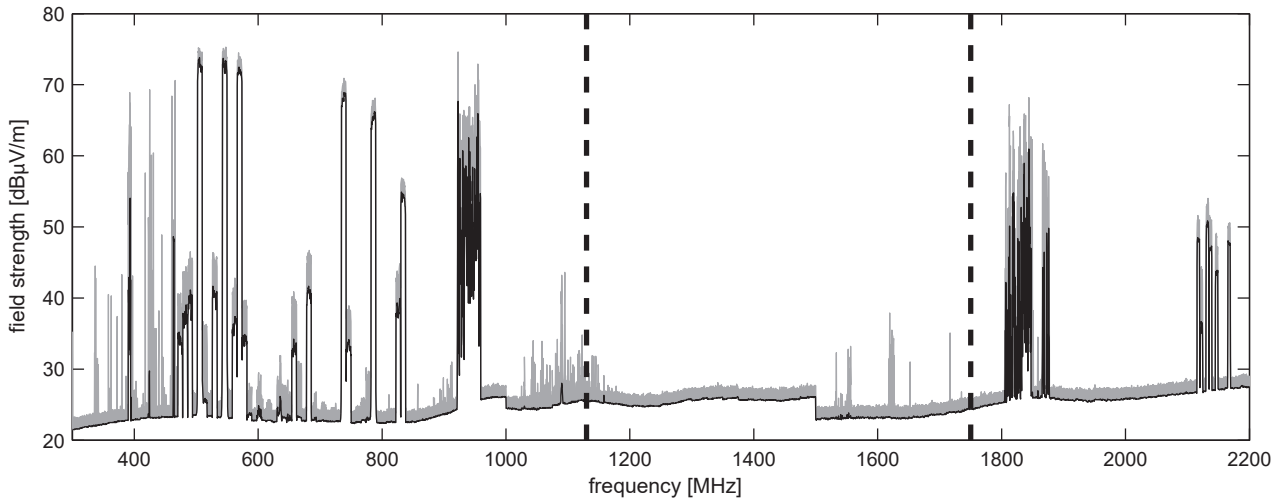
**Fig. 7.** Spectrum measured by the monitoring station at the WSRT site towards the horizon. WSRT-Apertif operates between the two vertical dashed lines. Small jumps in the noise floor are caused by varying antennas and receiver settings for each frequency band. The grey areas indicate time variability of the RFI.
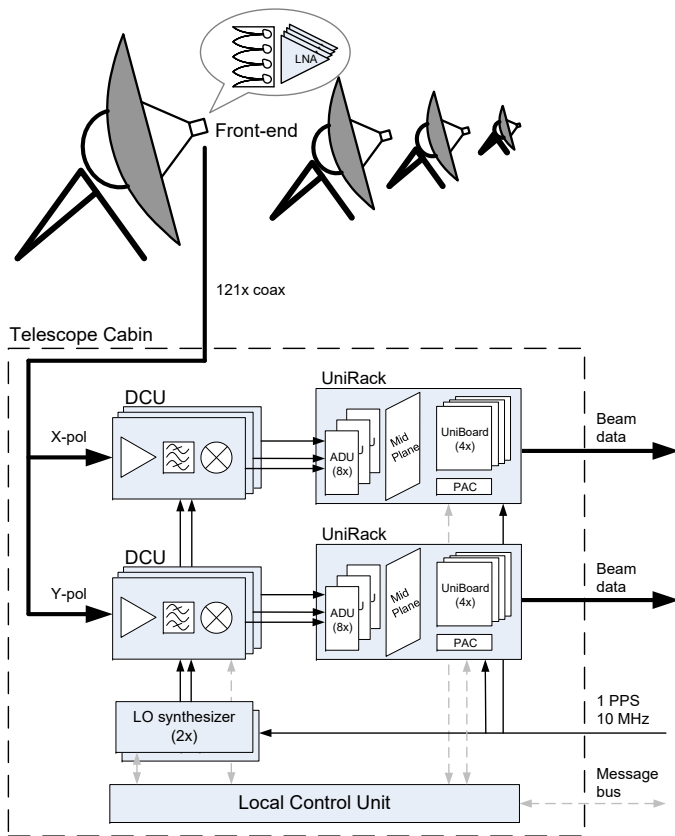


**Fig. 8.** Overview of the PAF system at every dish.

### 7.1. Antenna array

The Apertif feed arrays each consist of 121 single polarised tapered slot (or Vivaldi) antenna elements (Arts et al. 2010), 61 in one polarisation and 60 in the orthogonal polarisation (see Fig. 10). The antenna elements are configured in a rectangular grid at a pitch of 10 cm. The feed array is installed in the (primary) focus cabin of the dishes (see Fig. 9). For mechanical support, a grey polypropylene grid is attached to the front of the antenna elements. The impact of this grid on the electromag-

netic and noise performance of the antenna array is negligible. For weather protection, a foam radome is installed in front of the antenna. With a focal length of 8.75 m ($f/D = 0.35$), the WSRT reflectors are relatively deep and have an opening angle of $\pm71.1$ deg observed from the feed. A 1 m × 1 m sized PAF, the largest that fits in the focus cabin, provides a FoV of about $3 \times 3$ deg$^2$.

Figure 11 shows one of the PAF elements with integrated LNA. The entire frontend operates at ambient temperature. In the preliminary design phase, a cryogenically cooled system has been considered. Although a cryogenically cooled frontend was expected to lower the system temperature by $10 - 20$ K, it was infeasible due to its large impact on the system development, construction, and maintenance costs.

### 7.2. Frontend

Figure 12 shows the block diagram of the frontend board. The board is tightly integrated with the antenna elements to minimise losses. As opposed to horn antennas, wideband array elements like a Vivaldi do not have a sharp cut-off at the lower end of their frequency band of operation. Consequently, most of the RFI below 1130 MHz will be present at the input terminals of the LNAs. To avoid the creation of intermodulation products in the Apertif band, a very low-loss high-pass filter is integrated in the LNA. The measured $S_{21}$ of the filter is shown in Fig. 13. It suppresses the terrestrial digital video broadcasting (DVB-T) signals by more than 40 dB. The high-pass filter uses a distributed architecture with three taps. The resonance frequencies of these taps are optimised to maximise the attenuation across the stop band below 1 GHz. The losses are minimised by using a low-loss tangent board, and by finding a compromise between the conductive copper losses and the amount of attenuation. The $S_{21}$ in Fig. 13 includes mismatch losses, which do not contribute to the noise budget. The net contribution of the filter to the system temperature is $\approx 10$ K (see Table 3). In addition, out of band RF filtering is installed between analogue gain stages through the entire signal chain. The second high-pass filter is a discrete implementation of the same filter to further reduce the RFI. The third band pass filter is a ninth-order discrete filter to filter out GSM-900 and GSM-1800 above and below the frequency band of interest. The first-stage amplifier is a Skyworks 67151-396LF,
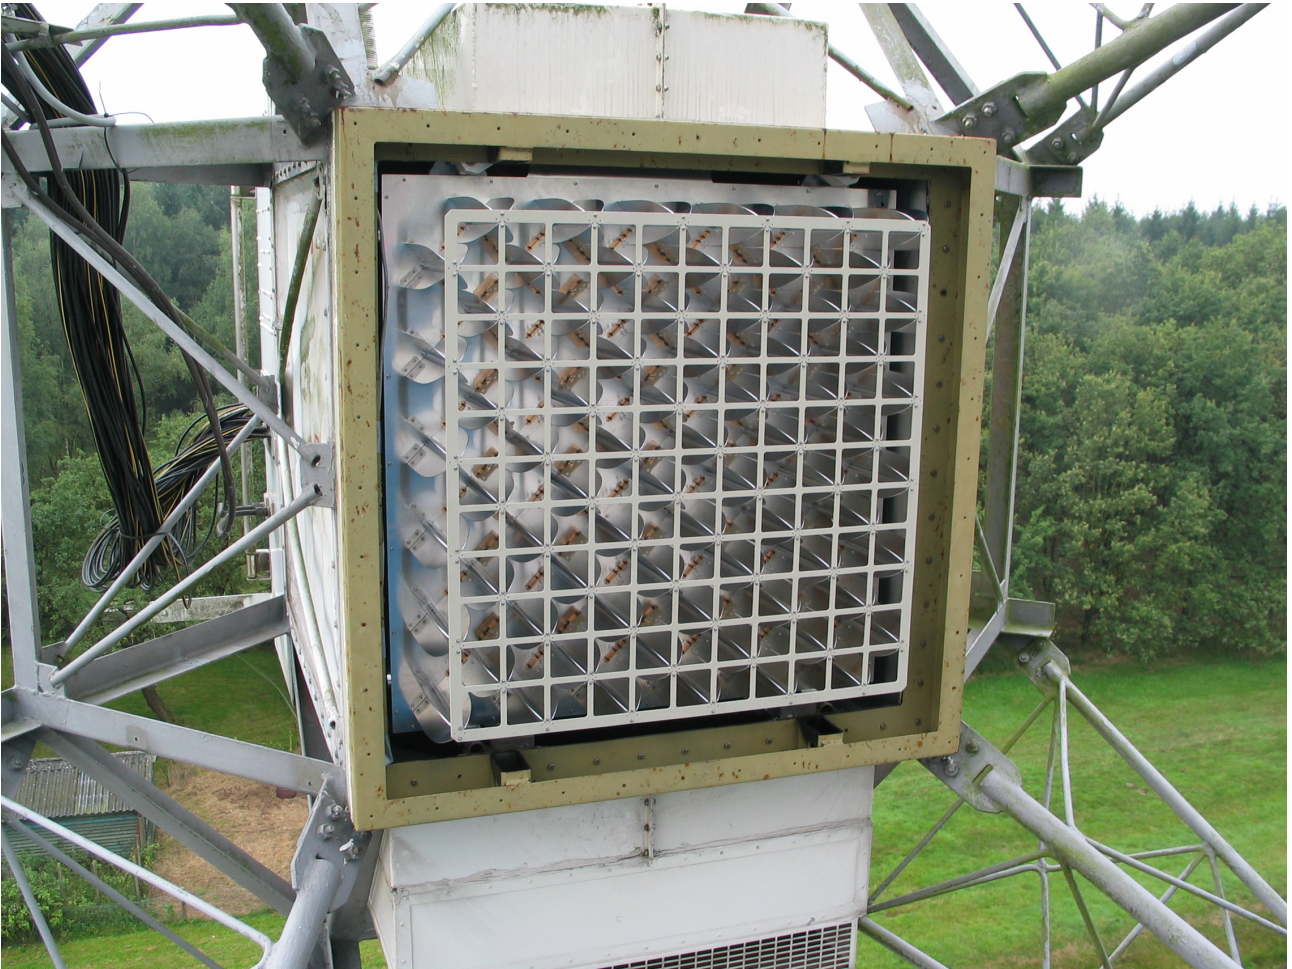
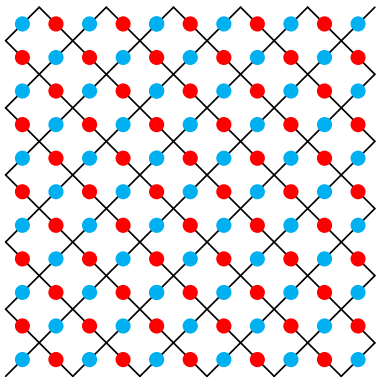**Fig. 9.** Photo of the Apertif PAF. For the photograph, the radome has been removed.



**Fig. 10.** Layout of the 121 PAF elements. The red and blue dots indicate the elements of two orthogonal polarisations.
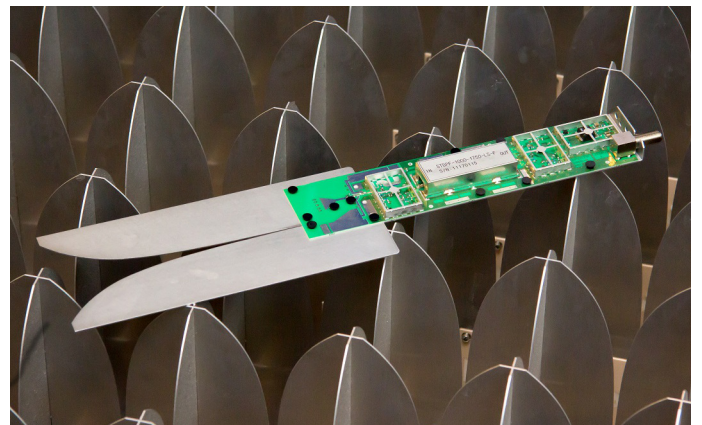


**Fig. 11.** Vivaldi antenna element with integrated LNA.

which is one of latest GaAs pHEMT transistors, and has one of the lowest-noise figures currently available using COTS devices. Its gain of 22 dB is used to overcome the largest part of the noise contribution of the other components in the RF system. The LNA receives its DC power supply over the same coaxial cable that is used to output the RF spectrum to the receiver. Protection against electrostatic discharges (ESDs) is installed at the output of the frontend board. At the input and output discharge resistors are used as well to prevent static charge buildup.

## 7.3. RF coaxial cables

The 121 RF signals from the LNAs are transported to the telescope cabin by 40 m long RF coaxial cables. Because the cables cross both the elevation and the hour-angle axis, they are subject to repetitive bending. A number of candidates for the cables were selected and subjected to accelerated lifetime tests by repetitive bending at very low temperatures. The 75 Ohm COAX9 cable performed best and is now used in Apertif.
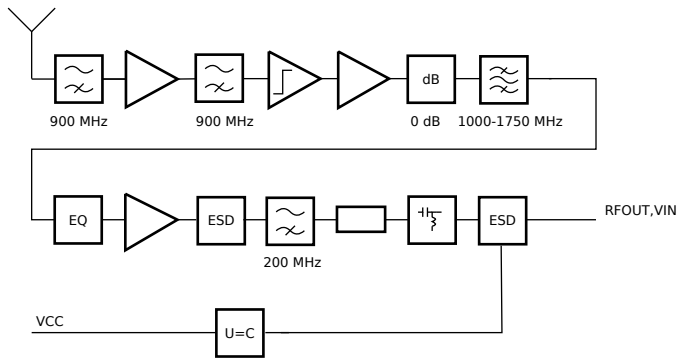
**Fig. 12.** Block diagram of the frontend board.
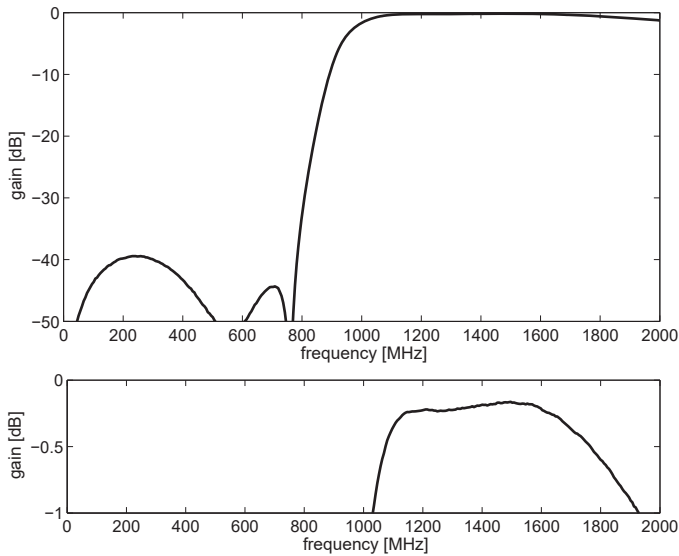


**Fig. 13.** Measured $S_{21}$ of the pre-LNA high-pass filter (top) and a close-up of the $S_{21}$ passband (bottom).

### 7.4. Telescope cabins

The output signal of the LNAs are connected to a telescope cabin (Fig. 14). The telescope cabin contains all dish-based systems (i.e. receivers and the PAF BF). The cabin is a standard 20 ft shipping container with a thermally isolated Faraday cage inside. The cabins at the four movable dishes are placed on a frame and move with the telescope structure. Every Faraday cage is equipped with an air-conditioning system for cooling (and heating) up to 6.5 kW. The shielding effectiveness of the Faraday cages was measured and is shown in Fig. 15.

### 7.5. Receiver

The down converter unit (DCU) receiver of Apertif converts the input RF band to the IF band (Fig. 16). The DCU also supplies DC power to the LNAs via the coaxial RF cables. All settings and monitoring information in the DCU subrack can be accessed through Ethernet. Four RF channels are combined on one printed circuit board. It also houses the last local oscillator (LO) distribution stages, power supply and control circuits, and the status monitoring of every channel. Eight DCU boards are combined to a DCU rack, which also houses the passive LO distribution. A separate rack is built for the filtered DC power supply. The DCU rack is shielded to suppress interfering EM sources within the

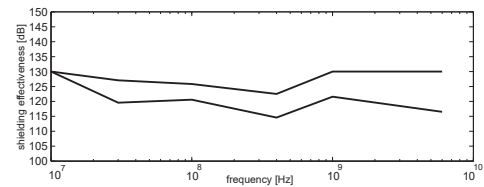**Fig. 14.** Telescope cabin, an EMC-shielded Faraday cage in a shipping container.



**Fig. 15.** Average measured shielding effectiveness according to EN 50147-1: March 1996 of the Apertif telescope cabins at two positions.

Faraday cage, like the digital backend, power supplies, control computer and the air conditioning system.

From an analysis of the RFI input spectrum and prerequisites, such as the Analog-to-Digital Convertor (ADC) capabilities, it is concluded that a double conversion scheme is most appropriate. The down-conversion scheme converts the RF frequency from 1130 – 1750 MHz to the second Nyquist zone of the ADC (450 – 750 MHz). It was very hard to find a relatively clean frequency mixing scheme without mixing products falling back into the frequency band, given the input spectrum (Fig. 7). Multiple options have been considered by modelling the RF spectrum throughout the receiver chain. The implemented mixing scheme (see Fig. 17) gives the best compromise between a clean spectrum throughout the channel, performance requirements on the filters and feasible LO frequencies using standard building blocks, such as PLLs, gain blocks, coax cables and standard printed circuit board (PCB) technology.

Starting at the top of Fig. 17, the RF input spectrum is up-converted using an adjustable LO1 frequency (4680 – 5000 MHz). This allows an intermediate frequency of 3.4 GHz. A custom bandpass IF filter rejects the image frequency and selects the 300 MHz band, while a second down-conversion mixer converts from IF to the second Nyquist zone of the ADC. The IF filter is based on ceramic coupled line technology to ensure a repetitive performance for all ≈1500 channels and a stable response across frequency. The down-conversion mixer uses an LO2 frequency of 2800 MHz, which is 3.5 times the ADC sam-
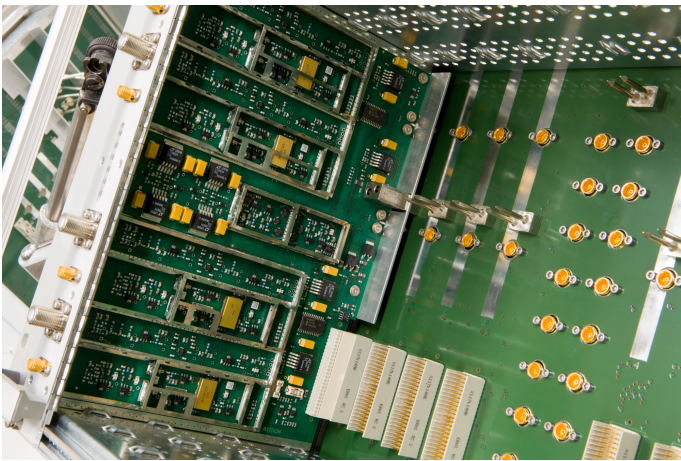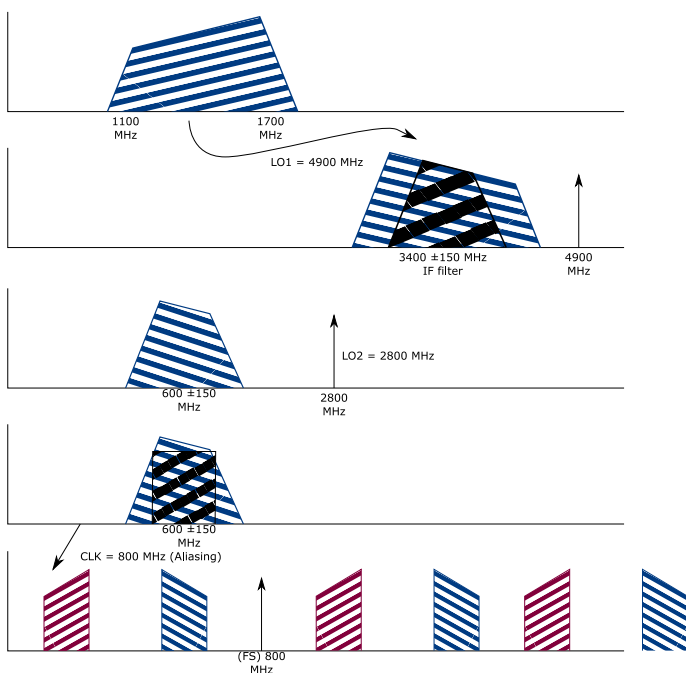
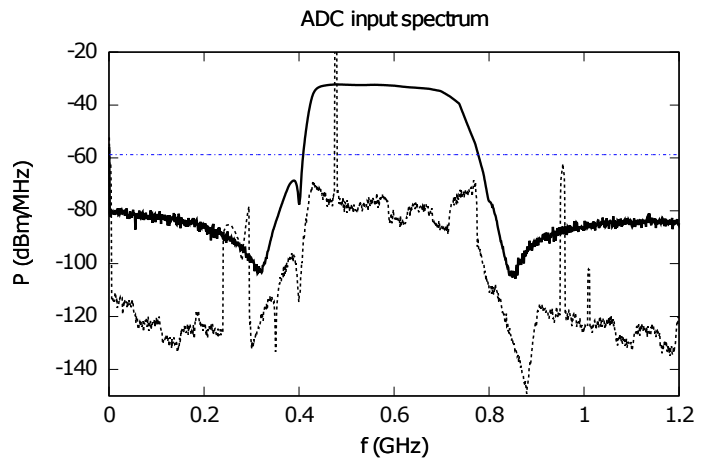**Fig. 16.** Apertif down convertor unit, covers removed.



**Fig. 18.** Simulated power spectrum at the ADC input of the noise power (solid line) and the worst-case interference (dashed line), including its intermodulation products.

### 7.6. Frequency and time references

The clock system in Apertif (Fig. 19) is used to synchronise all dishes with respect to each other, and to establish a link to the absolute time. All Apertif systems are locked to the 10 MHz clock from a hydrogen maser. A higher clock frequency would make the frequency selection grid too sparse, while lower clock frequencies would increase the phase noise contribution produced by the phased locked loop (PLL) circuits. An additional advantage of the 10 MHz reference frequency is the relative low loss at these frequencies of the coaxial cables used to distribute the clock signals. A distribution amplifier derives the 1 pps signal from the 10 MHz reference provided by the maser. The time difference between this 1 pps signal and a 1 pps signal from a GPS based clock is logged to establish an absolute time reference for the observations.

Commercially available distribution amplifiers are used to send the 10 MHz and 1 pps signals to the telescopes over existing buried coaxial cables. In the telescope cabin, two local oscillator generators (LOGs) generate the two LO signals from the 10 MHz clock reference. The LOG system is based on the Hittite HMC833 Phased Lock Loop and is built in an EM shielded 19 inch subrack. For each LO frequency range, a separate LOG system is built using different filters. The frequency can be configured remotely by an IP based Ethernet link. Furthermore, it has four isolated outputs with remotely controllable output powers. Thermal insulation of the PLL gives the unit its stability over time. To ensure phase repeatability between dishes after a reset, the output frequency is an integer multiple of the reference frequency.

The 200 MHz data processing clock that is used for the Uni-Boards in the dishes and the correlator is locked to the same 10 MHz reference as the ADC sample clock. Therefore, the processing rate in all FPGAs cannot drift between FPGAs. Zero clock drift helps to keep the internal memory usage of input FIFOs for the data transport between the FPGA small.



**Fig. 17.** Overview of the mixing schema used in the Apertif receivers.

pling clock frequency. If the clock or its harmonics would couple into signal it will be aliased back to DC or to 400 MHz. Finally, an anti-aliasing filter selects the second Nyquist zone of the ADC.

Figure 18 presents the simulated output spectrum of the RF system, when the dish is placed in the worst case position for RFI. The self-developed simulation code is loosely based on the harmonic balance method. The solid line represents thermal noise. The dashed line presents the power of the intermodulation products. The horizontal line around -60 is the maximum allowed level of the intermodulation products (i.e. 25 dB below the noise floor). It can be clearly seen that over the vast majority of bandwidth (450 – 750 MHz) the intermodulation products are well below the specification line. The line around 470 MHz is not and intermodulation product but is an in-band RFI signal from Iridium satellites.

### 7.7. Real-time digital signal processing

The analogue baseband signals from the DCUs are input to the UniRack: a shielded subrack with 64 analogue inputs that performs the analogue to digital conversion and the digital signal processing of the PAF signals, and sends the resulting 40 beams to the central correlator over optical fibres. The UniRack also
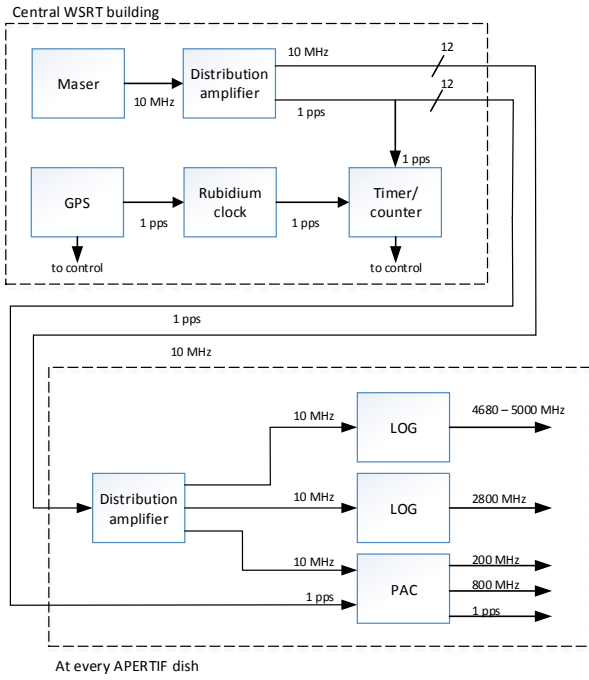
Fig. 19. Simplified diagram of the Apertif frequency and time reference distribution system.
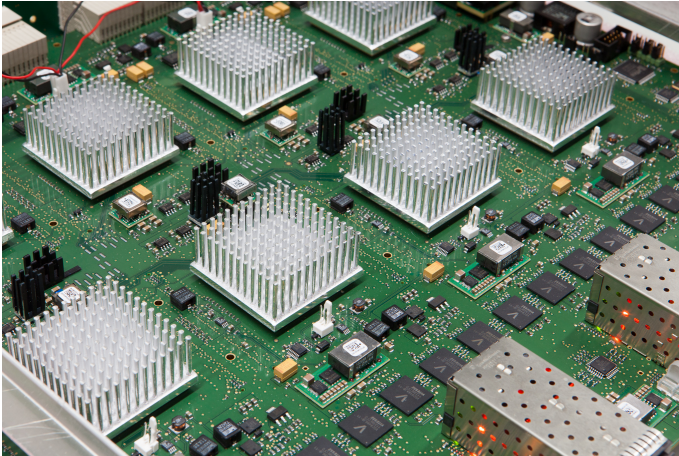


Fig. 20. UniBoard.

adds timestamps to the digital data. Apertif uses the UniBoard (Gunst et al. 2014) as the hardware platform for all digital signal processing (see Fig. 20).

One UniRack houses eight analogue to digital unit (ADU) boards, four UniBoards, and a power and clock (PAC) board. A midplane distributes the data between the ADUs on one side of the rack to the UniBoards on the other side (see Figs. 21 and 22). The PAC board provides the power and reference clock signal to all other boards via the midplane. The eight-bit ADCs operate at 800 MSPS, sampling the baseband signal in the second Nyquist zone. A variable attenuator with a resolution of 1 dB allows tuning the input signal optimally to the dynamic range of the ADC. The mixing scheme (Fig. 17) shifts the negative frequencies of the original RF band into the positive Nyquist zone of the ADC, so the band is flipped. Sampling in the second Nyquist zone flips the band again. Therefore, the net result is that

Fig. 21. Analog to digital conversion board (left), midplane (centre), and UniBoard (right), cover removed.
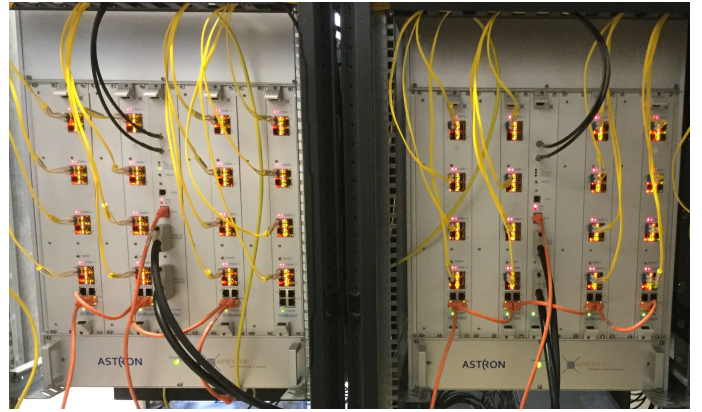


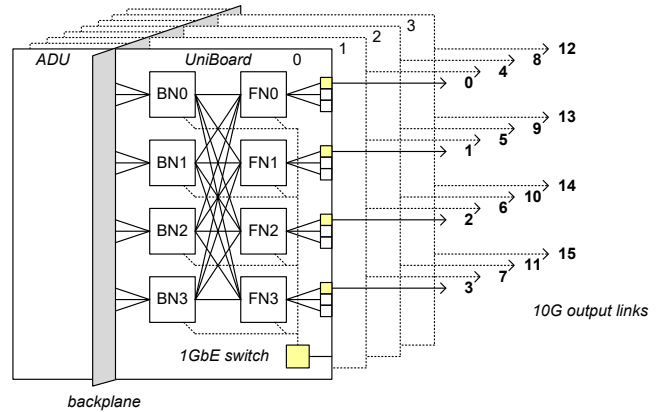Fig. 22. Photo of two UniRacks in a telescope cabin (UniBoard side of the rack).



Fig. 23. Schematic overview of datapaths in a UniRack: ADU boards on the left side, a midplane to distribute the data and FPGA processing board with 10 Gbps outputs on the right side.

after sampling, the RF frequencies are represented at baseband in the same incrementing order. The UniRack itself, including the ADU boards, is shielded to reduce the coupling of spurious signals from the digital electronics into the analogue signal chain.

Figure 24 provides an overview of the digital signal processing at the dish-level. First, the 400 MHz wide signal of every PAF element is split into $512 \times 781\,250$ Hz wide subbands (also known as coarse channels) by a critically sampled polyphase fil-
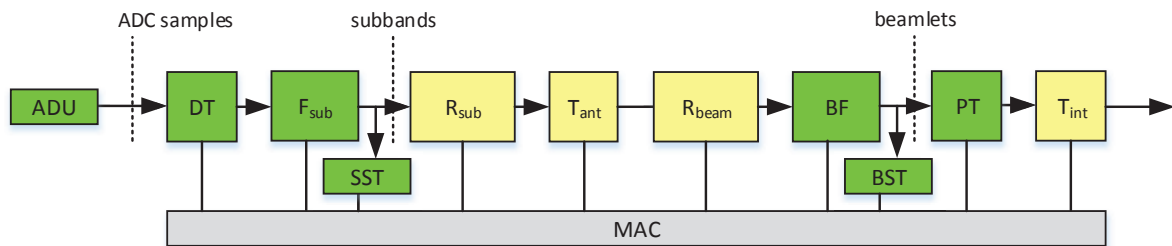
**Fig. 24.** Overview of the PAF dish-level digital processing

ter bank ($F_{sub}$). Inevitably, aliasing occurs at the band edges of the critically sampled subbands, resulting in sharp dips in the bandpass function that are hard to compensate. The $R_{sub}$ block selects the 384 central subbands covering 300 MHz bandwidth, and discards the 50 MHz band edges at each side. The BF for one subband requires the input from all PAF antenna elements, so therefore there needs to be a transpose $T_{ant}$ that groups the subbands from all antenna elements. For each of the 384 subbands the BF forms 40 beamlets from the antenna inputs. A beamlet is a quantity that represents the information of one beam spanning one subband. The BF is distributed over 64 BF units per UniRack. The $R_{beam}$ block holds and selects the subband such that the BF can create the beamlets of the CB. In each BF, the responses of all elements are multiplied with complex weights (i.e. magnitude and phase steering) and summed. The $T_{int}$ function is a corner turn that transposes and groups the data in time to provide the data in an optimal way for the correlator. The MAC function takes care of the proper operation, the subband selection and the BF weights.

To monitor and to tune up the system, the power spectrum of every PAF element is obtained by calculating the autocorrelation statistics of the subbands. Similarly, the power spectrum of every CB is calculated. These power spectra are produced by the Subband Statistics (SST) and Beamlet Statistics (BST) blocks in Fig. 24. The integration interval for these spectra is 1.024 s. These monitoring facilities run in parallel to the regular data path and are always available.

The statistics function also has a special mode in which it calculates the subband cross-correlation statistics between any two PAF elements (Schoonderbeek et al. 2020). The resulting covariance matrix of the PAF elements is used to determine the BF weights as described in Sect. 7.9. The measurement of a full $64 \times 64$ cross correlation matrix takes about 2 minutes for all subbands.

The two PAF polarisations are processed independently in two UniRacks. The beam forming for one polarisation of the PAF cannot be done on a single FPGA node for the full bandwidth, so the subband load has to be distributed across 16 FPGA nodes. Each UniBoard has four back node (BN) FPGAs and four front node (FN) FPGAs (see Fig. 23). In a subrack the BN FPGAs implement the $F_{sub}$ and $R_{sub}$ of Fig. 24, and the FN FPGAs implement $R_{beam}$ and BF. The transpose $T_{ant}$ is implemented in two hops. The first hop is via the backplane that implements a full mesh interconnect between the corresponding BN on the four UniBoards. The second hop is via the mesh interconnect between BN and FN FPGAs on each UniBoard.

In total, the Apertif digital BFs at the dishes contain 1536 ADCs on 192 ADU boards, and 768 FPGAs on 96 UniBoards. The total input data rate into the real time Digital Signal Pro-

cessing (DSP) is 9.8 Tbps. The central Apertif correlator (see Sect. 8.2) is implemented on 16 UniBoards with 128 FPGAs. Hence in total Apertif has 112 UniBoards and 896 FPGAs. Table 2 shows what fraction of the Stratix IV EP4SGX230KF40C2 FPGA resources are used by the BN filterbank image, the FN BF image at the dishes, and the central correlator image. The following resources are distinguished: The Adaptive Logic Module (ALM) that contains 2 Flip Flops + logic, Flip Flop (FF), 1 kByte block RAM (M9K), 16 kByte block RAM (M144K), and multipliers (DSP). The correlator uses more than about 85% of the FPGA resources, which is relatively high. As a consequence, it was challenging to synthesise an image that can achieve the required 200 MHz clock.

### 7.8. Geometrical delay compensation

As the direction of incidence of a celestial signal on the telescope changes during an observation, its time of incidence varies with the dish location and the angle of incidence. Delay compensation is needed to compensate for the change in geometrical delay between dishes. The geometrical delay compensation is regularly updated to follow the rotation of the Earth. The delay compensation in Apertif consists of true integer sample delay tracking (DT; see Fig. 24), followed by a residual phase tracking (PT). The DT is known as the coarse DT and the PT is known as the fine DT. The DT operates on the ADC samples at the input of the signal processor. The PT operates on the beamlets per CB and thus stops the fringe for the centre of each CB. The maximum fringe rate of the Apertif telescope is 1.28 Hz. To avoid too much de-correlation, the update rate for the PT needs to be about a factor of 100 faster than the maximum fringe rate, so a fine delay update interval of $T_{fine} = 7.8$ ms. This update rate is too high for the MAC interface and therefore the PT function uses piecewise linear interpolation, whereby the MAC only needs to provide the phase offset and phase slope coefficients for each linearised phase interval. The update interval for these linear PT coefficients is < 10.3 s. The fringe rate can also be interpreted as a differential Doppler shift between two dishes. The maximum fringe rate of 1.28 Hz is very small compared to the subband bandwidth. Therefore, the delay compensation in Apertif does not need Doppler shift correction on the sampled RF band. The residual delay that remains after the DT is 5 ns, or 4 ADC samples. To avoid significant de-correlation, the minimal subband period needs to be at least a factor of 100 larger than the delay step, so > 500 ns and thus the subband bandwidth must be < 2 MHz. The subband bandwidth is about equal to the subband rate of 781,250 Hz, so the 512 subbands that are created by $F_{sub}$ are enough.

**Table 2.** FPGA resource usage, as percentage of the maximum available number of resources.

|  | ALM | FF | M9K | M144K | DSP |
|---|---|---|---|---|---|
| Available max | 91 200 | 182 400 | 1235 | 22 | 1288 |
| BN Filterbank | 88% | 61% | 61% | 36% | 68% |
| FN Beam former | 62% | 43% | 83% | 68% | 82% |
| Correlator | 99% | 83% | 87% | 100% | 84% |

## 7.9. PAF beam forming

As introduced in Sect. 5, the PAF BF produces 40 dish-level CBs by forming weighted superpositions of the element responses. These CBs have the same role as the primary beams of dishes with a single receiver. The properties of the CBs are defined by the scheme that is used to determine the weights. Apertif implemented a beam forming scheme that maximises the sensitivity (Warnick et al. 2012) of the CBs. Although not yet implemented in Apertif, the flexibility of the PAF concept allows beam shape constraints to be added, for example to reduce the beam width variation with frequency (Ivashina et al. 2011) or mitigate RFI (Jeffs et al. 2008).

The Apertif dish-level BF implements a bi-scalar beam forming approach, in which the two sets of orthogonally placed PAF elements are processed independently. Bi-scalar beam forming incurs only a marginal loss in sensitivity (~4% for Apertif) compared to optimal full polarimetric BFs (Wijnholds et al. 2012, 2011; van Cappellen et al. 2012) and saves a significant data handling and processing load.

Apertif determines its weights by measuring array covariance matrices using the PAF subband correlator while observing a strong celestial point source (Cassiopeia A) and a relatively cold spot on the sky. Because the weights are determined with the actual system, all instrumental effects such as gain and phase differences between receiving channels, antenna mutual coupling, spill-over noise, and feed blockage, are implicitly taken into account.

In the weight calculation, two normalisation steps are implemented. The first ensures that the BF adds no gain, since that would conflict with the quantisation levels in the BF firmware. The second ensures a smooth frequency response, since the above mentioned method to determine weights is done for each subband independently, which could lead to jumps from subband to subband. To ensure a smooth frequency response, the weights are normalised such that the receiver chain that has the largest weight amplitude is real by definition. In some cases, there is no unique receiver chain with a highest weight. In that case a particular receiver chain has the largest amplitude for one subband, while in a second subband another receiver chain has the largest amplitude. To prevent the normalisation from toggling between receiver chains over the frequency band, this selection is made based on the data for a reference subband.

For all observations, Apertif uses the beam layout shown in Fig. 25. Thirty-nine beams are formed, which are laid out in a hexagonal pattern. One additional beam (beam 0) is formed at the location of the optical axis of the antenna dish, which is convenient for test purposes. The weight magnitudes of two beams are shown in Fig. 26. It can be seen that in a PAF system only a few elements per beam have high weight. All other elements have weights that are at least 10 dB lower, but are still important to optimise the details of the CB. The on-axis beam (beam 0) is dominated by the response of the central X element and the surrounding Y elements. The off-axis beam (beam 9) is dominated by two elements. The position of the highly weighted elements
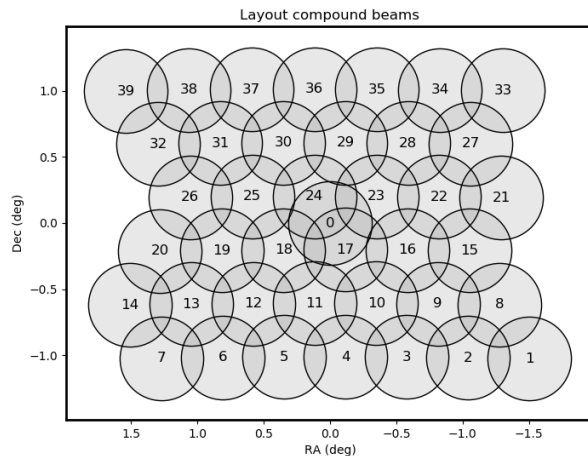


**Fig. 25.** Layout of locations of the 40 CBs over the FoV (Hess et al., in prep.).

with respect to the geometric focus point determines the pointing offset of the CB with respect to the central beam.

## 8. Central processing

The Apertif central processing combines the data from the 12 Apertif telescopes. The WSRT-Apertif system has two modes of operation: Beam-formed modes that sum the response of the telescopes, and an imaging mode in which the telescopes are correlated into visibilities.

### 8.1. Tied-array beam former

The beam-formed modes are primarily used for the ARTS, a super-computing radio-telescope instrument that performs real-time FRB detection and localisation on the WSRT, and can immediately alert LOFAR for incoming bursts. The system is extensively described in van Leeuwen et al. (2021). As shown in Fig. 27, the implementation of the BF is similar to that of the correlator, described in the next subsection. Compound beam data from the dishes are combined at the central location, using UniBoard and UniBoard[2]. On the UniBoards, the beam-forming modes are an alternative to the correlator mode described below. The system can run in a one of these two modes at a time. The subsequent transient detections functions (Sclocco et al. 2016; Sclocco et al. 2020) and post-processing software (Oostrum 2020) are implemented on a powerful GPU cluster.

### 8.2. Correlator

For imaging operations, a full Stokes correlator is implemented. Figure 28 shows a functional overview. The correlator is based on the FX architecture: First a polyphase filterbank filters the
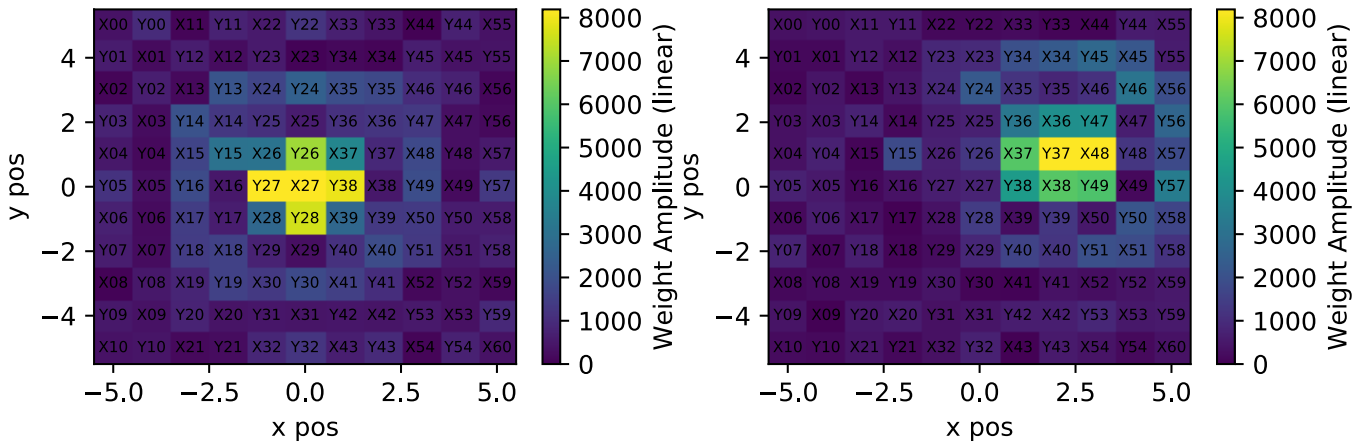
**Fig. 26.** Magnitude of the BF weights of CB 0 (left) and CB 9 (right).
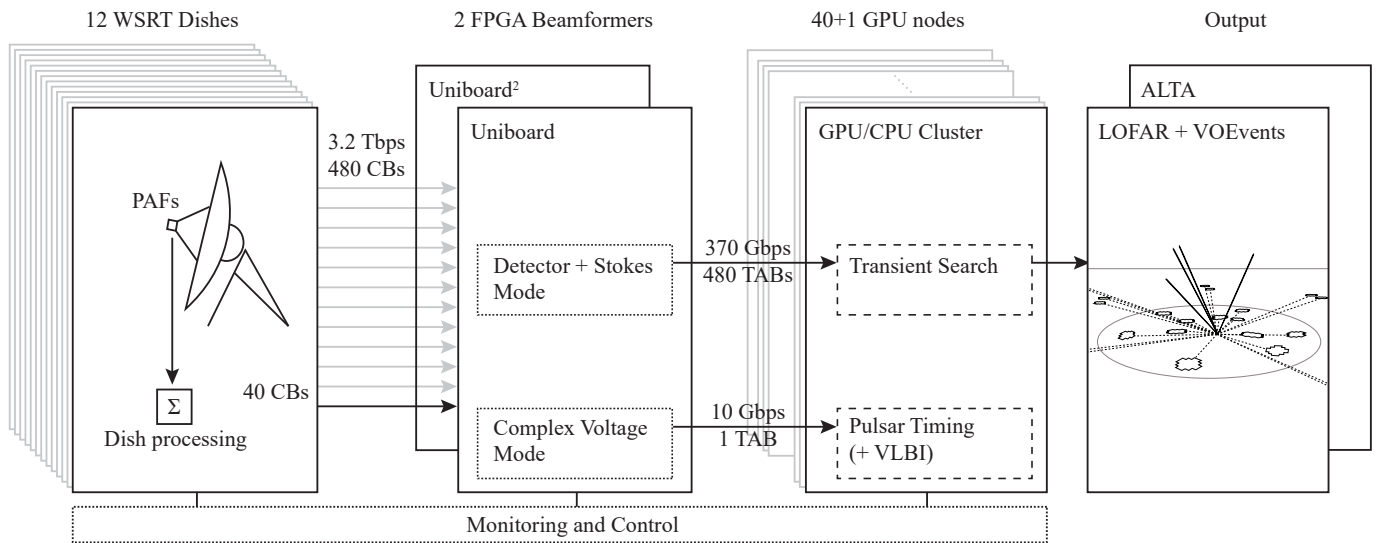


**Fig. 27.** Top-level diagram of ARTS, using the Apertif beam-formed mode (van Leeuwen et al. 2021).

781 250 Hz wide subbands into 64 channels (or 'fine channels') of 12.2 kHz each. The correlator then performs the cross-correlation between all pairs of the 24 telescope inputs (12 telescopes times two polarisations). The correlator calculates 276 cross-correlations and 24 autocorrelations, so 300 visibilities, for all 24 576 channels of all 40 CBs. The cross products are integrated for 1.024 s. After each integration interval the visibilities are offloaded to the data writer. The MAC takes care of the proper operation of the correlator, which mainly concerns monitoring the input and configuring the output, because the correlator itself has no control.

The Apertif correlator uses 16 UniBoard processing boards. To increase the number of optical inputs into the correlator, 16 optical-electrical boards (OEBs) are connected directly to the UniBoards (see Fig. 29). In the Apertif correlator application, all eight FPGAs on a UniBoard have the same function. Each FPGA has three 10G links, so in total one UniBoard can deal with 24 10GbE links from the dishes. This is just enough to connect to the 24 10GbE links that are processed by each board. Each FPGA in the Apertif correlator correlates all 300 visibilities from the 24 telescope inputs. The visibilities, a data stream of 18.4 Gbps, are output via the 1GbE interfaces on the UniBoards and
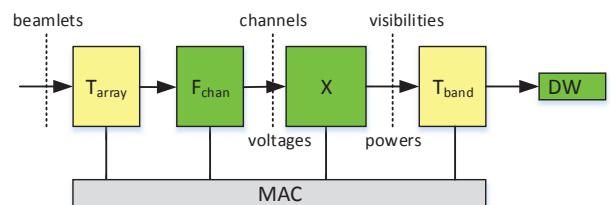


**Fig. 28.** The Apertif correlator subsystem.

arrive via four 10GbE interfaces at the data writer (see Fig. 23). The transpose function ($T_{band}$) is implemented through Ethernet switches.

### 8.3. Data writer

The main goal of the data writer is data reduction. The correlator delivers visibility data that is integrated for approximately one second, but the scientific requirements on time smearing allow for further integration of the data up to approximately 30 s, yield-
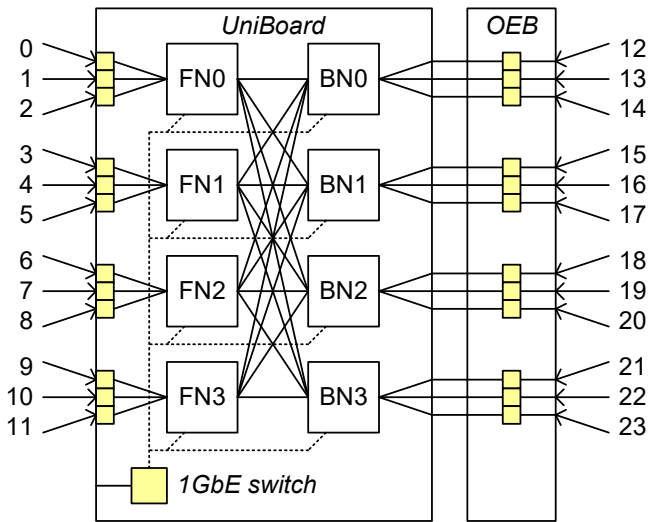
**Fig. 29.** Datapaths in the Apertif correlator.

ing a data reduction of a factor of 30. Secondly, the data writer applies phase corrections to the individual frequency channels within the subbands that were formed in the compound BF at the dishes. This correction is required to avoid phase variations of up to ±34 degrees between the lowest and highest channels in a subband for the most off-centred CBs. The phase slope across the band for off-centre CBs arises because the DT is the same for all CBs and optimal only for the central one. Lastly, the data writer writes the visibility data into separate measurement sets, one per CB. Each measurement set is augmented with all the metadata that is necessary to fully describe the data, hence ensuring that each data product is completely self-describing.

Apertif is equipped with two identical data writer systems that are used alternately to avoid one node having to write and read at a high data rate at the same time. While one node is recording a running observation, the dataset of the previous observation is offloaded to a post-processing cluster. Each data writer node is equipped with an Intel Xeon E5-2630 v3 CPU, having eight physical cores with hyper-threading; 64 GiB of RAM; eight 4 TB disks configured in RAID-6, yielding a net storage of 24 TB; and four 10GbE network interfaces. Several adjustments were made to the configuration of the system to improve the handling of large amounts of incoming User Datagram Protocol (UDP) packets.

The Input Reader is the most performance critical part of the data writer software. It needs to handle the incoming UDP data stream of approximately 19 Gbps. The header of each data packet contains a time stamp. Since data packets may not arrive in the correct time order, the Input Reader must reorder them. Packets with a time stamp after the observation's stop time are discarded immediately. The data writer uses a best-effort queue that consists of a number of so-called integration blocks. Each integration block holds data packets associated with a given integration interval, which is determined by the time stamp in the header. A new integration block is created in the best-effort queue when the time stamp in the packet belongs to the next integration interval. If the output writer cannot keep up, the best-effort queue may become full. In that case, no new integration block can be created, and the incoming data packet is dropped.

In order to reduce the size of the data products as much as possible, the visibility data are further integrated in time by the data writer. In theory, the correlator could already do so, but the UniBoards do not have enough memory available for buffering. The data writer also takes care of aggregating the different flags that can be set to indicate problems with the data. Flagged data are excised when averaging. Therefore, the time centroid of each integrated frequency channel must be recorded as metadata in the measurement set, as well as the weight that needs to be applied to normalise the data.

The BF in the PAF only performs fringe stopping per frequency subband. However, when each subband is further split up into channels at the correlator, we need to correct for the phase slope that remains between the lowest and highest frequency channels in that subband. If we do not correct for this, then the residual phase error can become as large as ±34 degrees for the most off-centred CBs.

The output writer is responsible for writing the integrated and fringe stopped visibilities of each CB to a separate measurement set, together with all the metadata that is needed in order to make the measurement set fully self-describing. All data processing takes place on data that is stored in one of the integration blocks in the best-effort queue. When all data of a given integration block have been processed, they are written to disk. Up to four CBs are handled in parallel by the output writer.

# 9. Software control

## 9.1. Monitor and control system

The Apertif MAC subsystem (Loose 2016) has four major responsibilities: (i) distribution of information; (ii) solving timing- and sequence-related issues; (iii) controlling all the involved hardware that is distributed over many locations; and (iv) providing monitoring information for state and incident tracking.

The architecture of the MAC subsystem is based on loosely coupled components that communicate via a message bus. The message bus technology assures, among others, guaranteed information delivery, persistency where needed, distribution of information over multiple machines, and solves many other communication paradigms. Complementary to this technology, many components were developed to implement the required functionality of the MAC layer of Apertif. Since information exchange between components is easy and reliable, the scope of each component can be kept small thereby making each component less complicated. Figure 30 shows the high-level design of the MAC system.

The basis for the distribution of information is the message bus technology. We chose to use the open source package QPID, which is based on the widely supported messaging standard AMQP 1.0. This package offers many features like guaranteed delivery, routing, multi-machine support, controlled lifetime of messages, persistency, broadcasting, etc. On top of these features, a thin layer was developed for Apertif to simplify communication paradigms like publish-subscribe, remote procedure call, scatter and gather. Furthermore, we defined different message categories, such as command, event, notification, etc., each having its own properties that serve their way of information exchange.

## 9.2. MAC components

To create a reliable MAC system, we also defined the behaviour of any component that is part of the MAC system. This means that any form of communication between two or more components is well defined and understood. To go one step further we
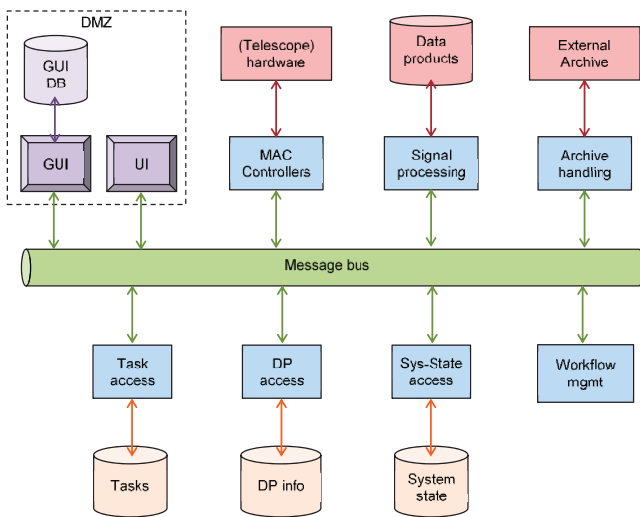
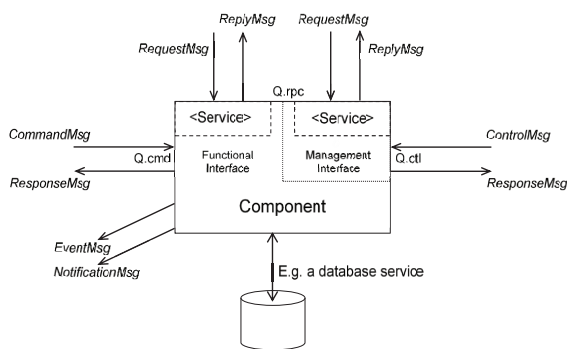**Fig. 30.** High-level design of the MAC system of Apertif.



**Fig. 31.** Typical interfaces of a component.

developed a component framework that implements this required communication behaviour. Guarded by timers and catching exceptions, the framework waits for messages on the message bus, executes the corresponding user-defined code and finally sends a response to the message bus. The framework offers many benefits: All communication paradigms are implemented only once, which, to a large extent, decreases and simplifies maintenance, and the implementation of a component by a developer is reduced to the implementation of the functionality that is needed to handle the messages. The diagram in Fig. 31 gives a schematic view on an Apertif component.

Currently the Apertif system contains three controllers for controlling all hardware: DirectionController, SignalController, and CorrelatorController. There is one instance of the Direction-Controller that is responsible for the pointing of all 12 dishes of Apertif. To control all the other hardware that is situated at each telescope, like DCUs, LOGs, PACs, ADUs, and UniBoards, each telescope has its own SignalController. The control of all these boards, more than 60 per telescope, is a complex task given the mutual dependences of the boards. To prepare the telescope for an observation many settings have to be checked or set for each of the 121 data streams that will come from the PAF. Only when all settings are correct the telescope can participate in the observation. The configuration of all the boards can take up to approximately 6 seconds when all the boards need a reset, during this period no other commands may interfere with the complex sequence of initialisation. The component framework takes care of this, so there is no extra code needed from the developer to accomplish this. CorrelatorController is responsible for configuring the 384 input streams of the telescopes in the right way and also for configuring the output streams of the correlator so that they connect to the machine where the data writer is started.

### 9.3. Monitoring

Monitoring the system can be done in two ways: actively and passively. Active monitoring can be done by sending a component a request message to report on its health, settings, or other information. The component framework will then execute the related user code to collect this information and send the answer to the requester. To collect the information of the whole system all components need to be polled.

Another way of monitoring the system is to let each component by itself report the state changes of the part of the system it is controlling. To support this, the controller framework implements a table that can contain commands that will be executed at a given interval. Depending on the implementation of the command by the developer, the command can deliver all current values as a notification message or deliver only the differences since the last time it was executed as an event message. In that case, the monitoring component receives all notifications and events, and registers the information in a database. It enables easy implementation of collecting the trends of values. The monitoring component can send out alarm messages when values or combination of values run outside their limits. In Apertif we implemented the passive mechanism but the information can still be collected in an active way also.

### 9.4. ARTAMIS

To collect and visualise monitor data for the Apertif operator and engineers, we employed Siemens' WinCC-OA Scada framework. WinCC-OA provides functionality such as a real-time event database, archiving, alarming and graphical user interface (GUI) development in a single toolkit. Being a single package, installation and implementation of the required functionality is easy and straightforward. To be able to store the collected monitor data in the database, we had to create a software component that collects data from the QPID message bus (see Sect. 9.1) and translates those messages into the proper database points of the WinCC-OA database. Once in the database, WinCC-OA provided tooling can be used for processing, alarming and visualisation of that data.

The specific tool that we developed for Apertif has been named ARTAMIS, which is an acronym for 'All-Round Telescope Array Monitoring and Information System'. ARTAMIS can be logged into via a dedicated desktop tool, or via a web based interface, which makes it possible to actively monitor the system from any location: the operator control room in the AS-TRON main building, the WSRT-based control room, the operator's home, or any other connected location. This flexibility proved to be essential in the recent COVID-19 period, where access to ASTRON's building has been severely limited by government regulations. ARTAMIS provides multiple views on the system: hierarchical views based on hardware components at different scales (array -> telescope -> BF -> UniBoards), up to the level of individual electronics boards. Furthermore, we provide a data-stream-centric view showing the network topology and sci-

**Fig. 32.** Screenshot of the ARTAMIS monitoring tool of Apertif.



**Fig. 33.** ALTA system.

ence data streams through the network from reception to storage. These views are complementary, and which one is most useful depends on the user's intent. In Fig. 32 we show an example screenshot of the graphical user interface. Apart from the hardware view, ARTAMIS provides views of the ATDB-based schedule (see Sect. 9.5), of the Archive (see Sect. 10), but also of the current state of non-APERTIF telescopes. With ARTAMIS, we thus built a single tool for the monitoring of the entire WSRT system.

### 9.5. Apertif task database

The Apertif task database (ATDB) forms the core of Apertif's operational software. Based on the survey schedule provided by astronomers, ATDB autonomously manages the full operational flow. From sending the observation parameters to the telescope system, up to data ingest into the ALTA, and removal of the raw data from the data writers. Despite being autonomous, ATDB also provides a web GUI for fine grained control by operators to troubleshoot, and to make changes in the data flow during the process. Astronomers can use ATDB to qualify observations as good or bad, after which a signal is sent to ALTA to either remove the data from the archive or transfer it to our cold storage facility at SURF.

ATDB is a web application that consists of a central database, surrounded by a pattern of micro services that all perform individual tasks based on the status of an observation in the central database. A stabilising feature of ATDB is that all the required logic is spread over small independent services that run on a myriad of different machines, but all communicate with the central database. The services are built and run in such a way that they restart whenever they fail for whatever reason, so that the mechanism can be seen as self-correcting. Services can also be taken offline and changed individually without affecting the rest of the system. The central ATDB backend can therefore remain simple and stable, limiting both maintenance and the risk of introducing bugs. In the 2.5 years of continuous operation ATDB has only failed once, due to hardware failure of an overheating computer.

## 10. Apertif long-term archive (ALTA)

The ALTA (Holties et al. 2019) provides storage and dissemination capabilities for the scientific data generated by the Apertif system. It supports telescope operations by providing continuous data intake capabilities. The scientists are supported with data management capabilities including overviews of observed and processed datasets, data selection and retrieval, quality an-
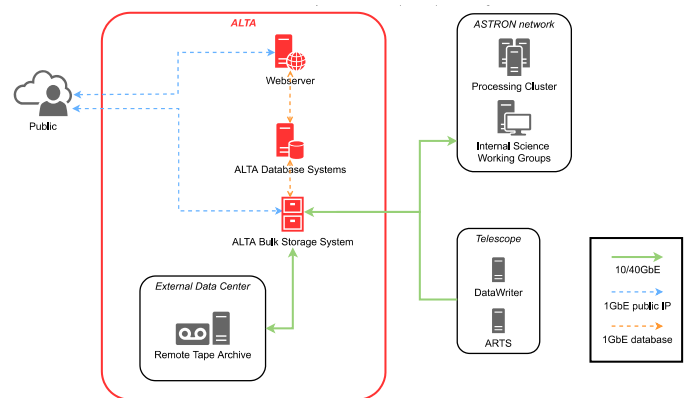
notations, and data ingest facilities for derived data products. Released data-products are published through ALTA to the public domain. A web-based user and application interface provides the required functionality for the targeted user groups to interact with ALTA. Figure 33 provides an overview of the ALTA system, its main components, and the main external systems it is connected to.

ALTA is designed to collect and manage millions of data products with a total volume of twenty petabyte over the operational lifetime of the Apertif instrument. Internal and external network connections for data transfers are based on 10 and 40 Gbps Ethernet technology. ALTA encompasses storage resources in two locations: ASTRON hosts online storage that provides 1.8 petabyte of capacity to support the operational telescope workflow and the analysis workflows for the Survey teams. The tape-based long-term storage of 'cold' data is hosted by SURF, the Dutch national high-performance computing and e-Science support centre. The backend supports an annual growth of up to four petabytes.

The basis for the data management components of ALTA is provided by the iRODS[8] framework. This framework allows the deployment of storage resources in so-called zones. Zones typically represent independently managed storage resources and can be mapped to resources hosted by different organisations. For ALTA, a single Zone has been created that spans iRODS resources at both the ASTRON (four data servers) and the SURF location (one server connected to the tape backend) as these are managed centrally as a single storage environment. Integrated functionality is implemented in Python, and executed as scheduled services in the iRODS environment. These services validate data-products during the ingest process, extract the metadata that is incorporated in the database, generate inspection plots for data quality assessment, and handle the migration between disk and tape storage based on a set of policies. The system allows for future implementation of user oriented services such as on-demand image cutout and image stitching.

The ALTA data model is based on the LOFAR data model and designed such that it can be easily mapped to Virtual Observatory (VO) compliant data structures. For the incorporation of full data provenance, inspiration was drawn from the W3C provenance, which was adapted to fulfil the constraints of the ALTA data model. In particular, the ALTA provenance model is strongly process-oriented, enabling the inspection of the full chain of processing steps and intermediate data-products, starting from the observations. A central Postgresql relational

---

[8]  https://irods.org/

database stores ALTA metadata using a schema that closely maps the structure of the ALTA data model and is fully exposed via a REST interface through the web user interface.

The GUI and the Application Programming Interface (API) components of ALTA consist of a JavaScript-based interactive user frontend and a RESTful backend. The backend builds on Django[9] and the Django REST framework[10] and handles the interaction with the database and data management components, both for the frontend and as an API for automation purposes. Data and information that is not yet released to the public domain requires an account with appropriate authorisation.

ALTA data collections are made publicly available through a release process where observations and processing pipeline runs are marked and subsequently exposed to anonymous users of the ALTA web interface. The associated data products are made available for retrieval via the Web Distributed Authoring and Versioning (WebDAV) protocol. Not all data are available for immediate access, in particular those products that are stored on tape. This applies mainly to the lower-level data products such as time-series data, uncalibrated and calibrated visibilities and the calibration tables generated by the calibration pipeline. Users are directed to a help desk service if they wish to retrieve data products not accessible immediately. Released observations, processing runs, and associated data products are registered in the AS-TRON virtual observatory server, running GAVO DaCHS[11] software, from where they become findable and accessible via standard VO aware client tools. The first scientific data release supported by ALTA has been published in November 2020[12].

## 11. Operations

The WSRT-Apertif telescope is operated as a survey instrument conducting surveys both in imaging and time-domain (tied array) mode. The operations are supported by both ASTRON observers and astronomers from the science teams. The operations are designed and run to support this dedicated survey observing mode. The cadence of observing is two to three weeks of dedicated survey observations. At the end of an observing period, system maintenance is performed. Following that, a beam weights measurement is performed. These beam weights are then used for the full period of observing that follows. Further information about how operations were designed and are implemented on a daily basis can be found below.

### 11.1. Designing Apertif operations

The focus of Apertif on conducting large surveys in a systematic way presented an opportunity to design an operational workflow that would suit this context. Like many new instruments and telescopes, the transition between commissioning and survey operations had fuzzy boundaries, which meant that we needed to make use of commissioning experience to the greatest extent possible to prepare for full survey operations. During this period, a strong emphasis was placed on automation and streamlined workflows in order to guarantee the maximum efficiency possible with limited resources. We sought to minimise the locations in the data flow that relied on humans making decisions with the goal of limiting any potential bottlenecks, which in some cases marked
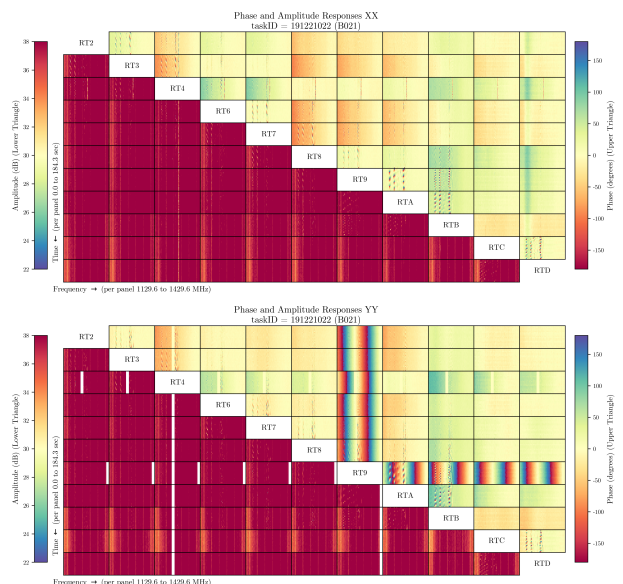


**Fig. 34.** Waterfall plots (time vs. frequency) of amplitude (lower left) and phase (upper right). Each baseline is shown for a single beam; for survey fields the central beam is shown and for calibrator observations the beam containing the calibrator is shown. A plot is made for both XX (upper image) and YY (lower image). These plots provide an initial check on data quality, especially the system tuning. RT9 has poor tuning in the YY polarisation; this is due to the missing data in the frequency range used for tuning.
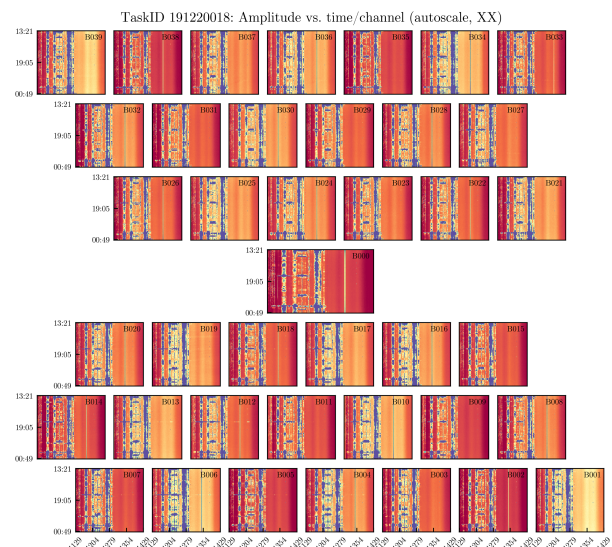


**Fig. 35.** Waterfall plots of amplitude for every beam of an observation. All baselines are averaged together for the XX polarisation, and these plots are created on a fixed colour scale. These inspection plots are particularly useful for monitoring the RFI environment.

a deviation from initial plans for the Apertif data flow. Doing this required looking critically at each stage of the observation and determining what steps could be made via automation to reduce these risks.

Decisions that we made during this time, motivated by the above, include the following eight points. The first was unifying machine-readable system specification parameters for both imaging and ARTS: We identified the key parameters needed from each team and united these into one machine-readable format, which serves as the input into the scheduling interpreter

tool (ATDBSpec, described below). This involved iteration during the commissioning phase to ensure we captured all requirements but has remained unchanged since the start of full survey operations. Schedules can be uploaded by science teams and are automatically synced to the telescope system from this uploaded location. This machine-readable format was developed so that any future interpreter could potentially replace the one currently in production (ATDBSpec).

The second was a semi-automated translation of schedule files to the system schedule: The ATDBSpec package was built in Python to sit between the machine-readable schedule files and the telescope system. ATDBSpec interfaces directly with ATDB, meaning that once it is run on a given schedule file, the observations are scheduled in the system and no further action is needed. We did not automate the process of picking up schedule files, as current procedures involve telescope operators making the final decision about whether the system is ready to observe before any observations are scheduled. ATDBSpec was later expanded to calculate schedules for drift scans, test observations and pointing calibration in addition to handling the translation of survey schedules. ATDBSpec runs entirely on the command line, developed with the goal of minimising the need to interact with a visual interface as much as possible (ATDB itself provides this interface if necessary) and with the intention of enabling future automated processes to interface with it if needed.

The third was automated system setup: The Apertif system can be readily set up in any observing mode using a Python script with minimal input parameters. This Python script wraps around existing system setup scripts and commands, bringing various aspects of setup together in one location with unified argument specification, and performs a system check at the end of the setup process to verify if it worked as expected. This Python script evolved from a shell script that was developed during commissioning, and adapted to make modifying input variables as straightforward as possible for operations. During commissioning, it was found that it was very important to ensure the right sequence of steps were followed for the various different modes of operations (e.g. imaging vs. transient surveys, or 'cold' vs. 'warm' start of firmware, which resets or preserves system status), and as such this script has been vital to setting things up with the right parameters.

The fourth was semi-automated calibration processes: We adapted Jupyter notebooks that had been used for the commissioning period (for delay and phase tuning) to be as automated as possible. A single Jupyter notebook is used to control the entire process, sending the necessary observations to the system based on which antennas are available and processing the results with visual diagnostic plots. The notebook can then be used by an operator to make a judgement call on the calibration process. By using the Jupyter notebook approach hosted on the internal server but accessed remotely via VPN, it is possible for the system to be set up from any location. Once the outcomes of calibration are approved, all relevant system files are copied to the right location and become immediately active. This process works reasonably well, but given more time and resources we would have aimed to additionally automate the validation of the calibration process and retire the notebook approach.

The fifth was automatic generation of raw data diagnostics: In order to provide a means of easily assessing data quality in the stage between observation and processing, a series of diagnostic plots were designed. These plots slice through the data across different relevant dimensions of time, frequency, amplitude and phase in order to make trends (in particular, bad data) readily visible. These plots are automatically generated once the data

are successfully uploaded to ALTA, and are directly linked in the online archive to the relevant task. A total of 18 raw data diagnostics (covering both XX/YY polarisation) are generated for each imaging observation, including the waterfall plots shown as examples in Figs. 34 and 35.

The sixth was automated diagnostic plots sent to Slack for immediate diagnostic by the first line operators: A crontab was set up to monitor the incoming observations, and automatically generate diagnostic plots if the observation is detected to be of a calibrator. For imaging, a plot of waterfall time-frequency for amplitude/phase for each cross-correlated antenna pair for the beam containing the calibrator is generated for XX/YY polarisations, and sent to the relevant operations channel (#apertif-verification) in the ASTRON Slack space. Figure 34 is an example of these particular diagnostic plots. In the case of time-domain observations, a crontab was set up on the ARTS cluster to monitor the production of local diagnostic plots from observations. This service picks up the generation of new plots and sends these to the same channel in Slack, allowing time-domain data to also be verified for quality. These approaches provide a quick and easy-to-access way to verify the results of calibration tuning and monitor calibrator or observation quality status from any location.

The seventh was automated monitoring and notifications about disk space: A crontab was set up to keep an eye on the disk space on both the data writers and the imaging processing server, with notifications above certain limits triggering warnings sent to Slack (#apertif-happili) at regular intervals. Sending these warnings to an easily accessible central location helped immensely with ensuring timely reactions to disk space problems and increasing transparency of the current disk status.

The eight was automated triggering of the imaging pipeline: We extended the automated workflow for Apertif from operations into processing via the Autocal package. This sits between ALTA and the imaging pipeline Apercal, and automatically picks up new observations, which it passes on to the pipeline. In its current form, Autocal automatically determines which observations are targets and then seeks to match the 40 flux and polarisation calibrator observations that are generally found at either side of the target. Complexity in Autocal arises from the need to sometimes schedule multiple targets in a row, the fact that not all observations are successful (meaning it needs to determine archive status for each calibrator beam and adjust accordingly), exceptions to normal scheduling and errors in scheduling either due to the system or operator input. Autocal generally performs equal to or better than a human in automatically starting up the pipeline, and operates 24/7 with systematic reproducible behaviour, thus ensuring a high efficiency for Apertif imaging processing. Upon triggering the Apercal pipeline, it sends a notification message to the relevant Slack channel (#apercal-notifications) and also reports back when the processing is complete. Autocal has also been extended to trigger the quality assessment part of Apertif imaging, and it both prepares and executes the ingest of processed data products back to ALTA. More detail about both Autocal and Apercal can be found in the Apercal paper (Adebahr et al. 2021).

Our efforts to implement as much automation as possible for Apertif survey operations were driven by the desire to ensure that human effort was invested effectively. This enabled operations to generally proceed smoothly, and most issues that surfaced tend to be due to real system problems or human error rather than issues with automated approaches. The removal of manual decision-making processes as much as possible in critical workflow paths was a key part of ensuring sustainable Aper-

tif operations within resource constraints. For example, data are automatically ingested to ALTA without waiting for a human decision on quality, ensuring the data writers are kept clear and deferring quality assessment till a later stage. On reflection, it is clear that there are more avenues in which these processes of automation could be improved and streamlined, and we hope to see continued advances in automation as a key part of observatory workflows both for Apertif and for telescopes worldwide.

### 11.2. Daily operations

The operational tools described above are a crucial part of the daily operations of observing with Apertif. The operations routine encompasses the following main aspects.

### 11.2.1. Daily monitoring

Both the state of the system and the quality of the observations are monitored on a daily basis, by both operators and astronomers. ARTAMIS (Sec 9.4) provides information about the instantaneous state of the system, over the full system chain, from the pointing position of the dishes to the state and health of the electronics and hardware and up to the data occupancy of the various on-line clusters (ALTA, ARTS cluster, data writer). Monitoring available data space and taking action to ensure ample space for continued observing are a core part of daily operations. The diagnostic plots described above provide a first view of the quality of the data being recorded and are regularly inspected to ensure ongoing observations meet the minimum scientific requirements. Overall, the monitoring of the system and data quality is used to inform the maintenance plan.

### 11.2.2. Maintenance

Every two to three weeks, survey observations are paused for system maintenance. In extraordinary circumstances, maintenance may interrupt the planned survey observations if required to return the Apertif system to an operable state that meets the minimum scientific requirements for observing. Maintenance includes both the hardware and software components of Apertif. Hardware maintenance mainly focuses on the health of the WSRT dishes (namely steering) and fixing the signal path of single antenna elements in a PAF. As Apertif is in active operational mode, software maintenance is focused on necessary bug fixes and upgrades of operating systems, with occasional deployments of small new functionalities that improve the operational flow.

### 11.2.3. System setup

Following the completion of maintenance, the system must be set up to resume survey observations. The automated setup script is used to fully reset the system with a cold start, along with setting the central frequency of the 300 MHz band to be used in the coming observations[13]. The beam weights are then measured using Cassiopeia A, as described in Sect. 7.9[14]. After the measurement, the beam weights are calculated and sensitivity plots of all CBs are produced, which are used to validate the beam weights before they are applied to the system. After the application of the beam weights, the delay and phase tuning is undertaken using the Jupyter notebook described above. The accuracy of this step in the calibration chain is crucial to ensure high signal-to-noise for tied array observations and a smooth continuum for spectral line studies.

### 11.2.4. Observing

After the PAF and array calibration are completed, the system is setup in the correct observing mode[15]. Scientists provide an observing file, and the operator on duty uses ATDBSpec to send the observations to this system. At this point, the daily monitoring resumes, and the cycle of operations continues.

## 12. System performance

Here we summarise the performance of the Apertif system. More details are given in the documentation of the first data release of the Apertif imaging surveys[16] and in a number of papers (Morganti et al. 2020, Oostrum et al. 2020, Adams et al. in prep, Hess et al. in prep, van Leeuwen et al. in prep.).

### 12.1. Sensitivity

### 12.1.1. System noise temperature

The design of the RF system aims to minimise the receiver noise temperature. Table 3 gives a breakdown of the various contributions. Starting at the antenna, the received EM wave passes through a radome, constructed from polystyrene plates. Aluminium is used for the antenna element because of its low conductivity losses, low weight and cost effectiveness. Entering the LNA board, the transition from antenna to microstrip line has been designed to be as efficient as possible. The LNA is assembled on a RO4003 substrate, which is low loss and has stable RF parameters over time. The noise contribution of the LNA includes the effects of mutual coupling of the antenna elements.

At the DCU, the mixers are the main noise contributors due to their passive topology. No other mixers were available that satisfy the high linearity, the isolation, and the large bandwidth required. The last component contributing to the noise figure is the ADC. It is used in the second Nyquist zone, having 7.4 effective number of bits and a signal to noise ratio of 47 dB. This in total gives the ADC a noise figure of 45 dB. Sufficient gain is placed in front of the ADC to overcome the huge noise figure. In addition to the components of the receiver noise listed in Table 3, other factors, such as spillover, atmospheric emission, and emission from the cosmic microwave background and the Galaxy, contribute to the overall system noise. The sum of all contributions results in a system temperature of about 70 K.

### 12.1.2. Aperture efficiency

Due to its flexibility, the Apertif PAF achieves a higher illumination efficiency than the horn feed used at the WSRT before the upgrade with Apertif. By optimising the weighting coefficients of the PAF, the illumination of the reflector is made more

---

[13] At the start of operations, we supported observing both the 1130-1430 MHz and 1220-1520 MHz bands. Currently, all survey observations are carried out in the 1220-1520 MHz band.

[14] It should be noted that while Cassiopeia A is a circumpolar source, the equatorial mounts on WSRT limit the times during which it can be observed, providing a strong constraint on the operational schedule.

[15] The delay and phase tuning are undertaken in imaging mode, so for transient surveys the system mode needs to be changed.

[16] http://hdl.handle.net/21.12136/B014022C-978B-40F6-96C6-1A3B1F4A3DB0

**Table 3.** Overview of the receiver noise contributors, all referred to the input.

| Component | $T_{rec}$ contribution [K] |
|---|---|
| Radome, antenna element | 6.0 |
| RFI filter | 10.0 |
| LNA | 32.9 |
| Coaxial cable | 0.2 |
| DCU | 3.8 |
| Coaxial cable | 0.0 |
| ADU | 2.6 |
| Total | 55.5 |

uniform, while keeping the spillover contribution low (van Cappellen & Bakker 2010). The illumination pattern of the horn feed used previously and of the on-axis Apertif PAF beam have been determined by Fourier transforming the two-dimensional complex (magnitude and phase) beam patterns (Fig. 36). The X-shaped pattern in the illumination is due to blockage by the four struts of the dish and by the focus cabin. It is clear that the PAF illuminates the reflector more uniformly than the horn feed, while the edge illumination is similarly low to avoid spillover. Measurements show that an aperture efficiency around 75% is achieved.

### 12.2. SEFD and Image noise

Figure 37 shows the measured $A_{eff}/T_{sys}$ as a function of frequency of the central CB of Apertif and compares it with that of the old WSRT. This shows that the $A_{eff}/T_{sys}$ of Apertif is about 35% lower, somewhat depending on frequency. Although the aperture efficiency of Apertif is 50% better than that of the pre-Apertif WSRT, its system temperature is about a factor of 2.3 higher. This is mainly due to the Apertif LNAs not being cooled, and that a filter had to be inserted before the LNA to suppress inter-modulation products due to strong out-of-band RFI leading to an increase in $T_{sys}$ of 10 K. Figure 37 also shows that for Apertif the behaviour of $A_{eff}/T_{sys}$ is a smooth function of frequency, while for the pre-Apertif WSRT, it varies strongly with a fairly regular pattern with bumps spaced 17 MHz. These bumps are caused by standing waves in the dish, which plagued the pre-Apertif WSRT and which are absent in Apertif (see Sect. 12.9).

The measured system equivalent flux density (SEFD) of the different CBs depend, as expected, on where they are within the overall FoV of Apertif, although the variation is very modest within the central $2° \times 2°$ of the FoV (Fig. 38). The measured SEFD of the central beam is around 560 Jy, which is ∼7% higher than nominal. The other beams within the central $2° \times 2°$ have an SEFD at most 10% higher. Outside this central region, the SEFD steeply increases with distance from the central beam, with the SEFD of the beams most distant of the field centre being around 850 Jy. This is due to the aperture efficiency of these beams being lower because only a limited number of PAF elements can be used to do the beam forming.

In practise, the noise in Apertif images is about 20% higher than expected from the theoretical performance. Apart from the SEFD being slightly higher than nominal, in every observation some data have to be flagged due to RFI, leading to higher image noise. For the first Apertif data release, the observing band used was 1130 − 1430 MHz. However, Apertif survey data products were made by only using the upper 120 MHz of the observing band due to the RFI conditions in the lower part of the full 300 MHz Apertif band, mainly due to global navigation satellites.

The best uniformly weighted continuum images have a noise level of about 30 $\mu$Jy beam$^{-1}$ while that of the line images is about 1.4 mJy beam$^{-1}$ for a velocity resolution of 7-8 km s$^{-1}$ and using Briggs weighting with robust = 0. Observations taken after January 2021 use a different frequency setup ($1230 − 1530$ MHz) resulting in a larger fraction of the observing band being free of RFI. The noise in these later continuum images is $20 − 25$ $\mu$Jy beam$^{-1}$.

### 12.3. Beam pattern and field of view

The layout of the 40 CBs is shown in Fig. 25. As described by Hess et al. (in prep.), this layout has been chosen to obtain reasonable uniformity of the noise within the FoV of a single observation after the images of the individual CBs have been mosaicked together while optimising the size of the FoV. The noise variations are below 20% in the central $2° \times 2°$ area. Thirty-nine CBs are put on a hexagonal grid with a separation of 28 arcminutes in right ascension and 24 arcminutes in declination. In addition, one CB is put on the optical axis of the telescope, which is convenient for test purposes. In a single observation, an area of roughly $3°\!.5 \times 3°$ is imaged. For survey observations, which use overlapping observations spaced $2.8°$ in right ascension and $2.3°$ in declination, there is an effective FoV of 5.25 deg$^2$. An example of a mosaic from a single Apertif observation is shown in Fig. 4.

### 12.4. Survey speed

The survey speed of a radio telescope is defined as the inverse of the time it takes to image a large area on the sky down to a certain noise limit. It therefore scales as SEFD$^{-2} \cdot \Delta\nu \cdot \Omega$ with $\Delta\nu$ the total bandwidth and $\Omega$ the FoV. This implies that the survey speed of Apertif (SEFD = 560 Jy, 39 beams, $\Delta\nu$ = 300 MHz) is about 22 times that of the pre-Apertif WSRT (SEFD = 310 Jy, 1 beam, $\Delta\nu$ = 160 MHz). In practice, however, a significant part of the observing band of Apertif is unusable due to strong RFI below 1300 MHz. Due to this loss of bandwidth due to RFI, the relative survey speed is about 10 for data taken before January 2021 and 15 for observations after that date. So despite the higher SEFD of Apertif and the impact of RFI, the gain in survey speed is substantial.

### 12.5. Beam shapes

#### 12.5.1. Element beams versus compound beams

The left panel of Fig. 39 shows the beam patterns measured on the sky, in a single polarisation, for all single elements of the PAF. Each panel covers the same $3° \times 3°$ on the sky and shows how the primary beam of a given element is positioned on the sky. It is clear that each element is sensitive to a slightly different area, but together they cover a region of about 8 deg$^2$.

One can clearly see the large optical distortions in the reception patterns of the single elements, in particular for those that are close to the edge of the PAF. However, because the PAF fully samples the radiation field in the focal plane, the signals from all elements can be combined in amplitude and phase to counteract these optical distortions and to optimise the shape of the combined beam.

The panel on the right in Fig. 39 shows the result of this optimisation for CBs formed in the direction of the element beams. We note that all images in this figure have the same colour scale. The right panel clearly shows that much of the optical distortions
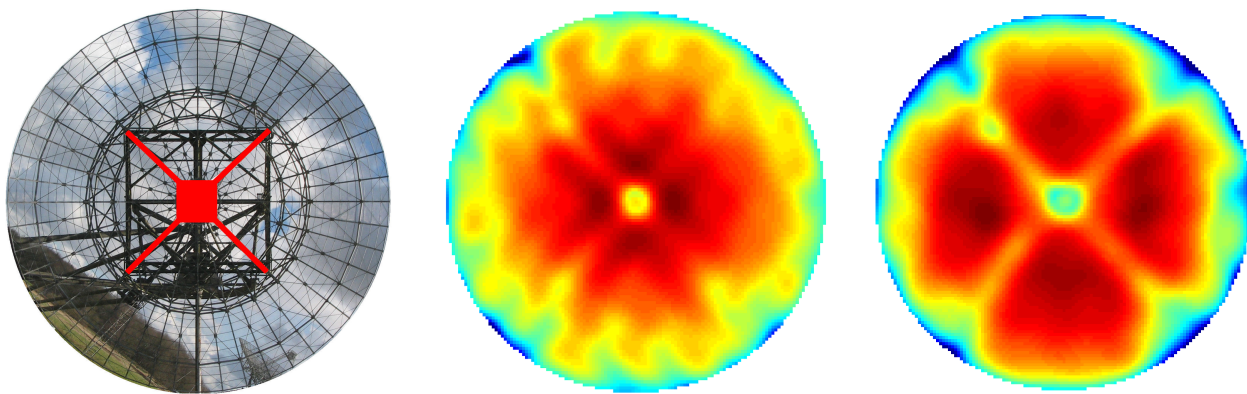
**Fig. 36.** Front view of the reflector (left) and magnitude of the aperture field of a horn feed (centre) and PAF (right) on a logarithmic colour scale (van Cappellen & Bakker 2010).
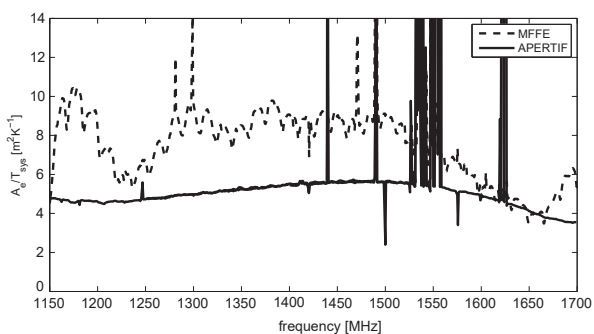


**Fig. 37.** Measured $A_{eff}/T_{sys}$ as a function of frequency of the central CB of Apertif and the pre-Apertif WSRT frontend (MFFE).
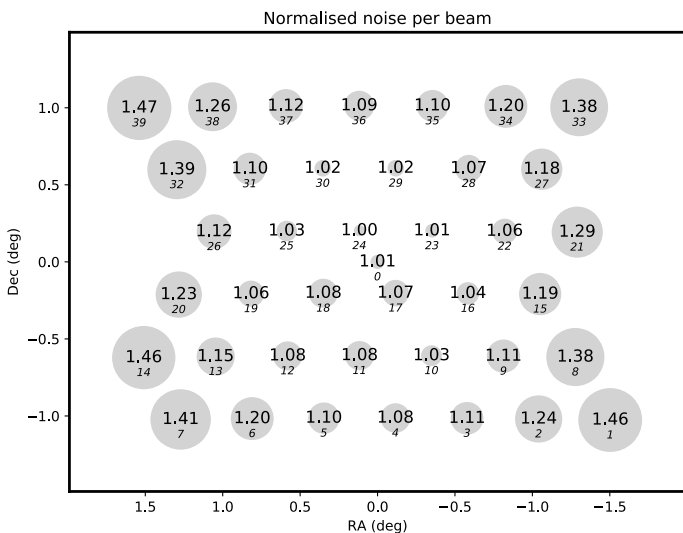


**Fig. 38.** Relative noise of the CBs over the FoV. The size of the circles scales with the relative noise levels. Beam numbers are slanted. The noise values are normalised to that of the beam with the lowest noise.

have been removed and that the shape of the beam patterns varies much less over the FoV. Some distortions remain, in particular towards the edges and corners of the FoV. To correct for these distortion one would need elements located outside the physical PAF array of Apertif.

## 12.6. Measuring compound beams

Beam weights are measured at the start of every imaging observing period (see Sect. 7.9) and are typically used for 2–3 weeks in a row. The CB shapes are empirically derived in two separate ways: drift scans across a bright source, and by comparing continuum fluxes of sources in Apertif continuum images with those from the NVSS using a Gaussian process modelling technique.

With the first method, beam maps are produced from drift scans performed periodically on Cygnus A (occasionally Cassiopeia A is used). During the drift scan measurement, the PAF is at a fixed position on the sky and Cygnus A drifts through the FoV. The change in the auto correlation amplitudes during the drift scan measurements is used to reconstruct the response of the CB response (Dénes et al., in prep.). This is then repeated 31 times to cover sufficient area to allow the measurement of all 40 Apertif beams.

In the second approach, a Gaussian process regression technique is used to reconstruct the primary beam response (Kutkin et al. in prep.). Continuum fluxes of sources in Apertif images for a given CB are compared to NVSS flux values. The ratio of the Apertif flux over the NVSS flux across the FoV corresponds to the primary beam response.

Both approaches find consistent shapes for the CBs. Figure 40 shows the sizes of the beams as measured using the drift scan method. The asymmetry of the beams located at the edge of the field is clearly seen in this plot.

The two approaches have different strengths and weaknesses. The drift scan methodology provides frequency and polarisation information about the CB shape, in addition to measuring the shape per WSRT dish. However, this approach can suffer from also detecting other sources than the main calibrator source, especially at the edges of the primary beam. It is also time-consuming observationally, leading to measurements that are performed infrequently. The Gaussian process regression approach utilises existing data, can be refined to provide shapes over a narrow range in time, and compensates any offset in flux scale between Apertif continuum images and NVSS. However, this method provides no information on the frequency or polarisation dependence of the beam, and the measurement is an average across all dishes.

For the first Apertif data release, the primary beam images based on the Gaussian process regression method, with frequency scaling information from the drift scans, were released. This choice is motivated in Adams et al. (in prep), along with a comprehensive look at the impact on the continuum fluxes.
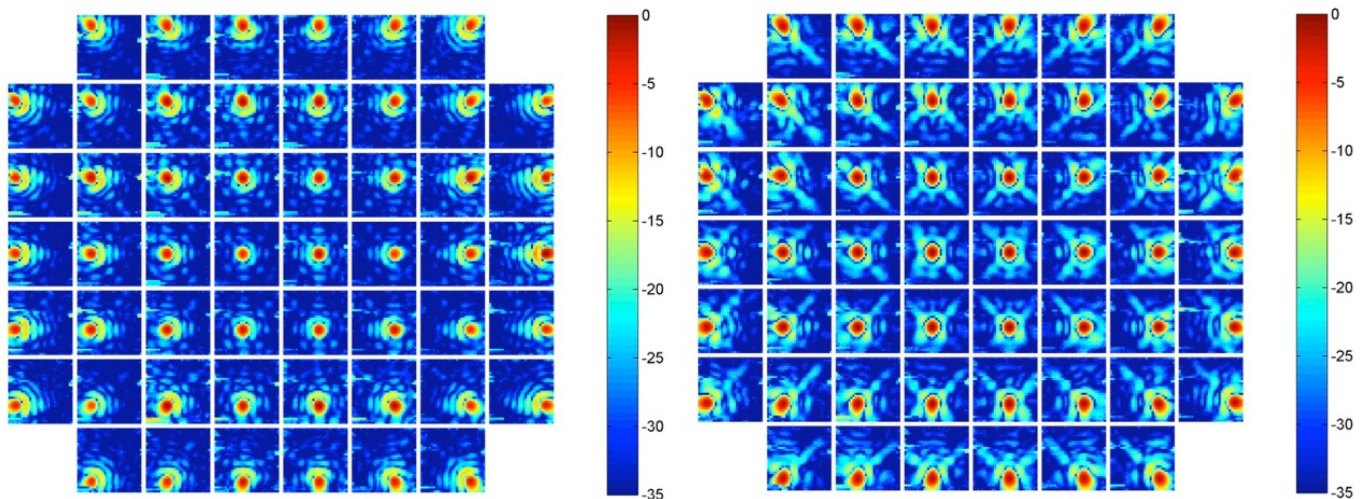
**Fig. 39.** Left: Response patterns on the sky of the individual Vivaldi elements of the PAF. Each panel covers the same $3° \times 3°$ on the sky and shows the primary beam of a given element. Right: Response of CBs formed in the same directions of the element beams. We note the significant improvement in the beam shapes compared to the element beams. The same colour scales have been used in the two figures. This makes it easy to see that the sidelobe level in the optimised CB is much lower than that of the element beams.
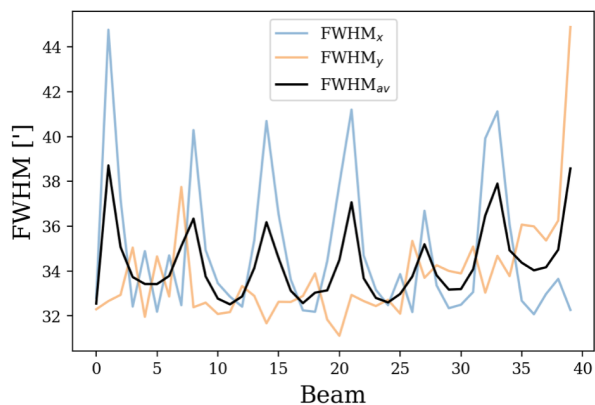


**Fig. 40.** FWHM as a function of beam number for a frequency of 1363 MHz. The black line shows the average FWHM when fitting a two-dimensional Gaussian function to the beam maps. The blue line shows the FWHM of the one-dimensional Gaussian along the major axis, and the orange line shows the FWHM along the minor axis. The shape of the CBs is not perfectly Gaussian, but a two-dimensional Gaussian function is a good approximation for the beam shapes within a few percent.

## 12.7. Polarisation

The image noise for the Stokes Q- and U-cubes with a frequency resolution of 0.78125 MHz (Adebahr et al. 2021) is ~ 120 μJy/beam per channel, leading to a final image noise after RM-Synthesis (Brentjens & de Bruyn 2005) for data with a total 150 MHz bandwidth of ~ 17μJy/beam in the polarised intensity images. This is slightly lower than the expected noise from the Stokes I continuum, which would normally be $\sqrt{2}$ higher. Since source densities in polarisation are generally lower and emission is fainter, the resulting images are less affected by sidelobe confusion and image artefacts.

Leakage from total power emission can significantly confuse emission in Stokes Q, U and V. To determine the instrumental leakage we used two different strategies. For linear feeds, instrumental leakage in Stokes Q is mostly directly correlated with Stokes I, while leakage into Stokes U is correlated with Stokes

V. Since several sources within an Apertif beam are physically linearly polarised, a separation of instrumental leakage and physical polarisation is difficult. To access the Stokes Q-leakage we used the drift scans taken for determining the shape of the CBs. Here the difference in the measurements of XX and YY corresponds directly to Stokes Q. Leakages of the order of 0.001 are found at the beam centres and 0.02 at a distance of 0.5 deg from the beam centres. For the edge beams, the leakage can be a bit higher, up to 0.1 for beams 01 and 39. For the analysis of the Stokes U-leakage, Stokes V images were analysed. Sources appearing in Stokes V are very rare and therefore any emission in this Stokes parameter can be assumed to be leakage. By analysing the fractional polarisation $FP_V = V/I$ for a large number of sources in a large number of images, we determined the leakage dependence. Values of $FP_V < 0.001$ were found for the pointing centres, increasing linearly to values of $FP_V < 0.02$ at distances of 0.5 deg, which are very similar to the values for Stokes Q.

## 12.8. uv coverage and resolution

For many years, the WSRT is participating in very-long-baseline interferometry (VLBI) observations, for example as part of the EVN, and the aim of the Westerbork Observatory is to continue to do so. Many of these VLBI observations are performed at frequencies not available to Apertif. Therefore, Apertif was installed in only twelve of the fourteen WSRT dishes, with the two remaining antennas keeping the existing MFFEs so that the WSRT can continue to perform VLBI observations at frequencies other than L band. After extensive study of the various options of which two antennas not to include in Apertif, it was decided to exclude the two most western dishes (RT0 and RT1). Although this leads to a loss of spatial resolution of about 10%, this choice creates the fewest gaps in the *uv* coverage, resulting in dirty point-spread functions with the lowest average sidelobe levels. In addition, for images tapered to 30-arcsecond resolution (which are often used for spectral line work), no gap is present in the *uv* coverage and the imaging performance is very similar to that of the pre-Apertif WSRT.
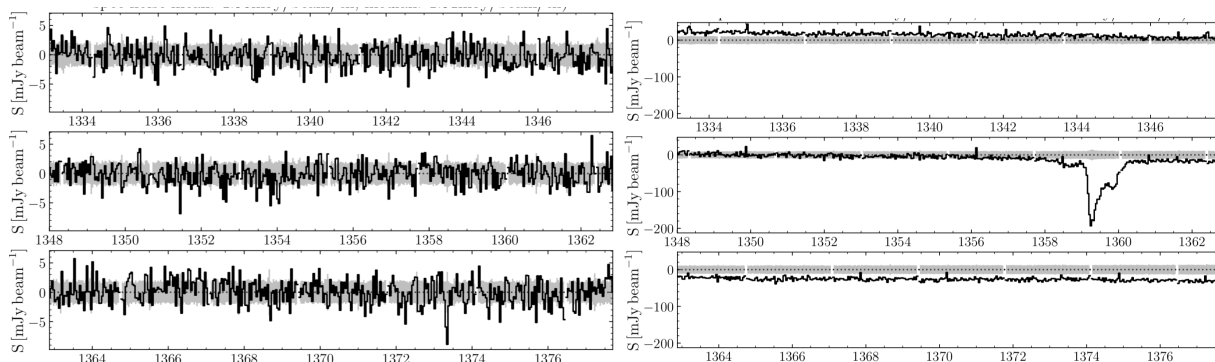
**Fig. 41.** Two continuum-subtracted spectra extracted from data cubes produced by the imaging pipeline (see Adams et al. in prep for details) against radio continuum sources of very different flux density. The source on the left has a flux density of ~20 mJy, while the flux density of the source on the right is ~2900 mJy. The horizontal axis is frequency covered by the corresponding cube. In both spectra, redshifted H I absorption is detected. Given the high flux density of the source on the right, the flatness of the spectra illustrates the good quality of the bandpass calibration that can be achieved even for bright radio sources. The offsets in the spectra relative to the zero intensity line are due to limitations in the continuum subtraction in the present version of the imaging pipeline.
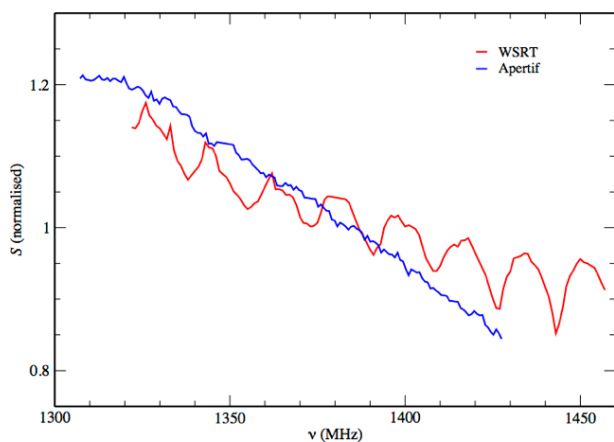


**Fig. 42.** Illumination efficiency as a function of frequency as measured for the pre-Apertif WSRT frontend (MFFE) and for the Apertif PAF. The MFFE data clearly show the periodic modulation due to standing waves, while the Apertif PAF shows no sign of such effects.

It was also decided to maintain regular coverage of the very inner uv plane, as was common practise with the old WSRT (the so-called maxi-short configuration). This means that the shortest four baselines are 36, 54, 72, and 90 m. This ensures that smooth structures with an extent of up to ~10 arcminutes can be well imaged with Apertif. Four of the twelve Apertif dishes are mounted on rails and in principle the array can be re-configured. However, in contrast to the pre-Apertif WSRT, this option is not used for Apertif. The longest baseline of the array is 2412 m. This gives a spatial resolution in uniformly weighted images of about $12''$ in right ascension and, due to the fact that the WSRT is an east-west array, $12''/\sin\delta$ in declination at declination $\delta$.

### 12.9. Standing waves

Before Apertif was installed on the WSRT, images made with the telescope were strongly affected by standing waves due to reflections between the vertex of the reflector and the frontend in the telescope focus (Popping, A. & Braun, R. 2008). Effectively, such standing waves affect the illumination of the dish as a function of frequency. In the case of the WSRT, the 8.75 m distance between the vertex and the primary focus results in variations in the illumination pattern with a period of 17 MHz. The main

impact of this on the images is that it modulates the observed spectrum of off-axis sources. Instead of being smooth, the spectrum of such sources show periodic variations (Fig. 42). In addition, these variations are polarisation dependent. This presents a huge complication for the calibration and analysis of wide-band observations since it is very difficult to take these spectral modulations into account in the calibration and data reduction, leading to very significant imaging errors for off-axis sources. Measurements show that these negative effects are absent in Apertif data (Fig. 37). Figure 42 shows the impact of this on the spectrum of a strong off-axis source as observed with the pre-Apertif WSRT and with Apertif. The pre-Apertif spectrum clearly shows the periodic modulation induced by standing waves in the dish, which are completely absent in the Apertif data. Electromagnetic simulations confirm that the absence of standing waves in Apertif is due to a much lower reflection of radiation (a factor of 1000) by the frontend back into the reflector because the Apertif PAF absorbs more energy than the WSRT horn system. In addition, the PAF acts as a diffuser for the small amount of radiation reflected back into the dish because it is reflected over a much larger solid angle.

### 12.10. Bandpass

The observing band of Apertif is the combination of 384 subbands of 0.78125 MHz each with every subband having 64 channels, each channel 12.2 kHz wide. Because the subbands are critically sampled as described in Sect. 7.7, aliasing results in a sharp drop in overall response for channels at the subband edges, namely channels 0, 1, and 63 of every subband. At the moment this issue is solved by flagging these three channels with low signal at the edges of the subbands.

The stability of the bandpass is key for a proper calibration of the data. The changes to the amplitude of the bandpass during an observation are a fraction of percent. This was derived from long (12 h), dedicated observations of strong calibrators, and from the comparison of the bandpass solutions derived from the standard calibration observations at the beginning and at the end of full synthesis observations. The well-behaved calibrated bandpasses can be seen in the two examples shown in Fig. 41, which shows two spectra taken at the location of radio sources, after subtraction of the continuum flux and showing detections of H I absorption. The two sources have very different flux density, ~ 20 mJy for the source in the left panel, while the one on

the right is a much stronger source of ~2900 mJy. The flatness of the spectra indicate good spectral stability, in particular in the case of the very strong radio source on the right.

## 13. Conclusions

We presented an overview of the Apertif system, the phased array upgrade of the WSRT: its science goals, system design, and key characteristics. The WSRT with Apertif is one of the first telescopes worldwide to successfully use PAFs for all its science operations, and it is living up to expectations. It is demonstrated that the WSRT-Apertif data can be calibrated and imaged using automated pipelines, which deliver good quality science data.

## References

Adebahr, B., Schulz, R., Dijkema, T., et al. 2021, Astronomy and Computing, 100514
Arts, M., Ivashina, M., Iupikov, O., Bakker, L., & van den Brink, R. 2010, in Proceedings of the Fourth European Conference on Antennas and Propagation, 1–5
Bannister, K. W., Deller, A. T., Phillips, C., et al. 2019, Science, 365, 565
Bannister, K. W., Shannon, R. M., Macquart, J.-P., et al. 2017, ApJ, 841, L12
Barnes, D. G., Staveley-Smith, L., de Blok, W. J. G., et al. 2001, MNRAS, 322, 486
Becker, R. H., White, R. L., & Helfand, D. J. 1994, in Astronomical Society of the Pacific Conference Series, Vol. 61, Astronomical Data Analysis Software and Systems III, ed. D. R. Crabtree, R. J. Hanisch, & J. Barnes, 165
Blyth, S., van der Hulst, T. M., Verheijen, M. A. W., et al. 2015, in Advancing Astrophysics with the Square Kilometre Array (AASKA14), 128
Bock, D. C.-J., Large, M. I., & Sadler, E. M. 1999, AJ, 117, 1578
Boersma, O. M., van Leeuwen, J., Adams, E. A. K., et al. 2021, A&A, 650, A131
Bonato, M., Prandoni, I., De Zotti, G., et al. 2020, MNRAS, staa3218
Braun, R., Bourke, T., Green, J. A., Keane, E., & Wagg, J. 2015, in Advancing Astrophysics with the Square Kilometre Array (AASKA14), 174
Brentjens, M. A. & de Bruyn, A. G. 2005, A&A, 441, 1217
Caleb, M., Flynn, C., Bailes, M., et al. 2017, MNRAS, 468, 3746
Caleb, M., Flynn, C., Bailes, M., et al. 2016, MNRAS, 458, 718
Champion, D. J., Petroff, E., Kramer, M., et al. 2016, MNRAS, 460, L30
CHIME/FRB Collaboration, Amiri, M., Bandura, K., et al. 2018, ApJ, 863, 48
CHIME/FRB Collaboration, Andersen, B. C., Bandura, K., et al. 2019, ApJ, 885, L24
Chowdhury, A., Kanekar, N., & Chengalur, J. N. 2020, ApJ, 900, L30
Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, AJ, 115, 1693
Connor, L. & van Leeuwen, J. 2018, AJ, 156, 256
Connor, L., van Leeuwen, J., Oostrum, L. C., et al. 2020, MNRAS, 499, 4716
de Gasperin, F., Williams, W. L., Best, P., et al. 2021, A&A, 648, A104

Fonseca, E., Andersen, B. C., Bhardwaj, M., et al. 2020, The Astrophysical Journal, 891, L6
Gardenier, D. W., Connor, L., van Leeuwen, J., Oostrum, L. C., & Petroff, E. 2021, A&A, 647, A30
Gardenier, D. W., van Leeuwen, J., Connor, L., & Petroff, E. 2019, A&A, 632, A125
Geréb, K., Maccagni, F. M., Morganti, R., & Oosterloo, T. A. 2015, A&A, 575, A44
Giovanelli, R. & Haynes, M. P. 2015, ARA&A, 24, 1
Giovanelli, R., Haynes, M. P., Adams, E. A. K., et al. 2013, AJ, 146, 15
Grant, J. K., Taylor, A. R., Stil, J. M., et al. 2010, ApJ, 714, 1689
Gunst, A., Szomoru, A., Schoonderbeek, G., et al. 2014, Experimental Astronomy, 37, 55
Hales, C. A., Norris, R. P., Gaensler, B. M., et al. 2014, MNRAS, 441, 2555
Haynes, M. P., Giovanelli, R., Kent, B. R., et al. 2018, ApJ, 861, 49
Heald, G., Józsa, G., Serra, P., et al. 2011, A&A, 526, A118
Hess, K. M., Roberts, H., Dénes, H., et al. 2021, A&A, 647, A193
Holties, H., Vermaas, N., & de Goei, R. 2019, in Astronomical Data Analysis Software and Systems XXIX, https://www.adass2019.nl/conference-info/posters/pdf?table=poster&id=140
Intema, H. T., Jagannathan, P., Mooley, K. P., & Frail, D. A. 2017, A&A, 598, A78
Ivashina, M., Bregman, J., & van Ardenne, A. 2002, in 12th International Conference Microwave and Telecommunication Technology, 278–281
Ivashina, M. V., bij de Vaate, J. G., Braun, R., & Bregman, J. D. 2004, Proc. SPIE, 5489, 1127
Ivashina, M. V., Iupikov, O., Maaskant, R., van Cappellen, W. A., & Oosterloo, T. 2011, IEEE Transactions on Antennas and Propagation, 59, 1864
Jeffs, B. D., Warnick, K. F., Landon, J., et al. 2008, IEEE Journal of Selected Topics in Signal Processing, 2, 635
Keith, M. J., Jameson, A., van Straten, W., et al. 2010, MNRAS, 409, 619
Koribalski, B. S., Staveley-Smith, L., Westmeier, T., et al. 2020, Ap&SS, 365, 118
Lacy, M., Baum, S. A., Chandler, C. J., et al. 2020, PASP, 132, 035001
Law, C. J., Bower, G. C., Burke-Spolaor, S., et al. 2018, ApJS, 236, 8
Lawrence, E., Vander Wiel, S., Law, C., Burke Spolaor, S., & Bower, G. C. 2017, AJ, 154, 117
Loose, M. 2016, in Astronomical Data Analysis Software and Systems XXVI, http://www.adass2016.inaf.it/index.php/participants-all/4-poster/54-loose-marcel
Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, Science, 318, 777
Maan, Y. & van Leeuwen, J. 2017, in 2017 XXXIInd General Assembly and Scientific Symposium of the International Union of Radio Science (URSI GASS), 1–4
Maccagni, F. M., Morganti, R., Oosterloo, T. A., Geréb, K., & Maddox, N. 2017, A&A, 604, A43
Macquart, J. P., Keane, E., Grainge, K., et al. 2015, in Advancing Astrophysics with the Square Kilometre Array (AASKA14), 55
Mesa, D., Baccigalupi, C., De Zotti, G., et al. 2002, A&A, 396, 463
Michilli, D., Seymour, A., Hessels, J. W. T., et al. 2018, Nature, 553, 182
Morganti, R. & Oosterloo, T. 2018, ARA&A, 26, 4
Morganti, R., Oosterloo, T. A., Brienza, M., et al. 2020, A&A, in press
Morganti, R., Tadhunter, C. N., & Oosterloo, T. A. 2005, A&A, 444, L9
Norris, R. P. 2017, Nature Astronomy, 1, 671
Norris, R. P., Hopkins, A. M., Afonso, J., et al. 2011, PASA, 28, 215
Oosterloo, T. A., Vedantham, H. K., Kutkin, A. M., et al. 2020, A&A, 641, L4
Oostrum, L. C. 2020, DARC: Data Analysis of Real-time Candidates
Oostrum, L. C., Maan, Y., van Leeuwen, J., et al. 2020, A&A, 635, A61
Oppermann, N., Junklewitz, H., Robbers, G., et al. 2012, A&A, 542, A93
Padovani, P. 2016, A&A Rev., 24, 13
Pastor-Marazuela, I., Connor, L., van Leeuwen, J., et al. 2021, Nature, 596, 505
Petroff, E., Bailes, M., Barr, E. D., et al. 2015, MNRAS, 447, 246
Petroff, E., Hessels, J. W. T., & Lorimer, D. R. 2019, A&A Rev., 27, 4
Platts, E., Weltman, A., Walters, A., et al. 2019, Phys. Rep., 821, 1
Popping, A. & Braun, R. 2008, A&A, 479, 903
Prandoni, I., Guglielmino, G., Morganti, R., et al. 2018, MNRAS, 481, 4548
Prandoni, I. & Seymour, N. 2015, in Advancing Astrophysics with the Square Kilometre Array (AASKA14), 67
Rengelink, R. B., Tang, Y., de Bruyn, A. G., et al. 1997, A&AS, 124, 259
Ryan-Weber, E. V., Begum, A., Oosterloo, T., et al. 2008, MNRAS, 384, 535
Sancisi, R., Fraternali, F., Oosterloo, T., & van der Hulst, T. 2008, A&A Rev., 15, 189
Schaubert, D. H., Boryssenko, A. O., van Ardenne, A., Bij de Vaate, J. G., & Craeye, C. 2003, in IEEE International Symposium on Phased Array Systems and Technology, 2003., 351–358
Schaubert, D. H., Chio, T. ., & Holter, H. 2000, in IEEE Antennas and Propagation Society International Symposium. Transmitting Waves of Progress to the Next Millennium. 2000 Digest. Held in conjunction with: USNC/URSI National Radio Science Meeting (C, Vol. 1, 178–181 vol.1

Schnitzeler, D. H. F. M., Banfield, J. K., & Lee, K. J. 2015, MNRAS, 450, 3579

Schoonderbeek, G., Hut, B., Kooistra, E., et al. 2020, Journal of Astronomical Instrumentation, 9, 2050007

Sclocco, A., Heldens, S., & van Werkhoven, B. 2020, SoftwareX, 12, 100549

Sclocco, A., van Leeuwen, J., Bal, H. E., & van Nieuwpoort, R. V. 2016, Astronomy and Computing, 14, 1

Serra, P., Oosterloo, T., Morganti, R., et al. 2012, MNRAS, 422, 1835

Shimwell, T. W., Tasse, C., Hardcastle, M. J., et al. 2019, A&A, 622, A1

Siemion, A. P. V., Bower, G. C., Foster, G., et al. 2012, ApJ, 744, 109

Simons, J., bij de Vaate, J., Ivashina, M., et al. 2006, in 2006 First European Conference on Antennas and Propagation, 1–6

Sokoloff, D. D., Bykov, A. A., Shukurov, A., et al. 1998, MNRAS, 299, 189

Spitler, L. G., Scholz, P., Hessels, J. W. T., et al. 2016, Nature, 531, 202

Stil, J. M., Keller, B. W., George, S. J., & Taylor, A. R. 2014, ApJ, 787, 99

Strom, R., van Ardenne, A., & Torchinsky, S., eds. 2018, 50 years Westerbork Radio Observatory, A Continuing Journey to Discoveries and Innovations

Subrahmanyan, R., Ekers, R. D., Saripalli, L., & Sadler, E. M. 2010, MNRAS, 402, 2792

Tan, G. H. 1991, in Astronomical Society of the Pacific Conference Series, Vol. 19, IAU Colloq. 131: Radio Interferometry. Theory, Techniques, and Applications, ed. T. J. Cornwell & R. A. Perley, 42–46

Tasse, C., Shimwell, T., Hardcastle, M. J., et al. 2021, A&A, 648, A1

Taylor, A. R., Stil, J. M., Grant, J. K., et al. 2007, ApJ, 666, 201

Taylor, A. R., Stil, J. M., & Sunstrum, C. 2009, ApJ, 702, 1230

Thornton, D., Stappers, B., Bailes, M., et al. 2013, Science, 341, 53

Tucci, M., Martínez-González, E., Toffolatti, L., González-Nuevo, J., & De Zotti, G. 2004, MNRAS, 349, 1267

van Cappellen, W. A. & Bakker, L. 2010, in 2010 IEEE International Symposium on Phased Array Systems and Technology, 640–647

van Cappellen, W. A. & Ivashina, M. V. 2011, in Proceedings of the 5th European Conference on Antennas and Propagation (EUCAP), 926–929

van Cappellen, W. A., Wijnholds, S. J., & Bakker, L. 2012, in 2012 6th European Conference on Antennas and Propagation (EUCAP), 634–637

van der Hulst, J. M., van Albada, T. S., & Sancisi, R. 2001, in Astronomical Society of the Pacific Conference Series, Vol. 240, GAS AND GALAXY EVOLUTION: A CONFERENCE IN HONOR OF THE 20TH ANNIVERSARY OF THE VLA, ed. J. E. Hibbard, M. P. Rupen, & J. H. van Gorkom, 451

van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, A&A, 556, A2

van Leeuwen, J. 2014, in The Third Hot-wiring the Transient Universe Workshop, ed. P. R. Wozniak, M. J. Graham, A. A. Mahabal, & R. Seaman, 79–79

van Leeuwen, J., Kooistra, E., Oostrum, L. C., et al. 2021, A&A, in prep

Warnick, K. F., Ivashina, M. V., Wijnholds, S. J., & Maaskant, R. 2012, IEEE Transactions on Antennas and Propagation, 60, 184

Wayth, R. B., Lenc, E., Bell, M. E., et al. 2015, PASA, 32, e025

Weem, J. P. 2001, PhD thesis, University of Colorado

Wijnholds, S. J., Ivashina, M. V., Maaskant, R., & Warnick, K. F. 2012, IEEE Transactions on Antennas and Propagation, 60, 4688

Wijnholds, S. J., van Cappellen, W., & Ivashina, M. V. 2011, in 2011 XXXth URSI General Assembly and Scientific Symposium, 1–4

[1] ASTRON, the Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4,7991 PD Dwingeloo, The Netherlands

[2] Kapteyn Astronomical Institute, PO Box 800, 9700 AV Groningen, The Netherlands

[3] Astronomisches Institut der Ruhr-Universität Bochum (AIRUB), Universitätsstrasse 150, 44780 Bochum, Germany

[4] SKA Organization, Jodrell Bank, Lower Withington, Macclesfield, Cheshire, SK11 9DL, United Kingdom

[5] Anton Pannekoek Institute, University of Amsterdam, Postbus 94249, 1090 GE Amsterdam, The Netherlands

[6] CSIRO Astronomy and Space Science, Australia Telescope National Facility, PO Box 76, Epping NSW 1710, Australia

[7] Sydney Institute for Astronomy, School of Physics, University of Sydney, Sydney, New South Wales 2006, Australia

[8] South African Radio Astronomy Observatory (SARAO), 2 Fir Street, Observatory, 7925, South Africa

[9] Department of Physics and Electronics, Rhodes University, PO Box 94, Makhanda, 6140, South Africa

[10] Argelander-Institut für Astronomie, Auf dem Hügel 71, D-53121 Bonn, Germany

[11] Dept. of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden

[12] The Inter-University Institute for Data Intensive Astronomy (IDIA), and University of Cape Town, Private Bag X3, Rondebosch, 7701, South Africa

[13] Department of Astronomy, University of Cape Town, Private Bag X3, Rondebosch 7701, South Africa

[14] Veni Fellow

[15] Department of Physics, McGill University, 3600 rue University, Montreal, QC H3A 2T8, Canada

[16] Tricas Industrial Design & Engineering, Zwolle, The Netherlands

[17] Netherlands eScience Center, Science Park 140, 1098 XG, Amsterdam, The Netherlands

[18] Astro Space Center of Lebedev Physical Institute, Profsoyuznaya Str. 84/32, 117997 Moscow, Russia

[19] ASML Netherlands B.V., De Run 6501, 5504 DR Veldhoven, The Netherlands

[20] Cahill Center for Astronomy, California Institute of Technology, Pasadena, CA, USA

[21] Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, University of Manchester, Manchester M13 9PL, UK

[22] Joint Institute for VLBI ERIC, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, The Netherlands

[23] INAF, Osservatorio Astronomico di Cagliari, via della Scienza 5, 09047 Selargius, CA, Italy

[24] Inter-University Centre for Astronomy and Astrophysics, Post Bag 4, Ganeshkhind, Pune 411 007, India

[25] NYU Abu Dhabi, PO Box 129188, Abu Dhabi, United Arab Emirates

[26] Center for Astro, Particle, and Planetary Physics (CAP³), NYU Abu Dhabi, PO Box 129188, Abu Dhabi, United Arab Emirates

[27] NIKHEF, National Institute for Subatomic Physics, NL-1098 XG Amsterdam, the Netherlands

[28] Department of Astrophysics/IMAPP, Radboud University, PO Box 9010, 6500 GL Nijmegen, The Netherlands

[29] West Virginia University, White Hall, Box 6315, Morgantown, WV 26506

[30] Department of Physics, Virginia Polytechnic Institute and State University, 50 West Campus Drive, Blacksburg, VA 24061, USA

[31] University of Oslo Center for Information Technology, P.O. Box 1059, 0316 Oslo, Norway