

# The Apertif Radio Transient System (ARTS): Design, commissioning, data release, and detection of the first five fast radio bursts

Joeri van Leeuwen<sup>1,2</sup>, Eric Kooistra<sup>1</sup>, Leon Oostrum<sup>1,2,3</sup>, Liam Connor<sup>1,4</sup>, Jonathan E. Hargreaves<sup>1</sup>,  
Yogesh Maan<sup>5,1</sup>, Inés Pastor-Marazuela<sup>2,1</sup>, Emily Petroff<sup>2,6</sup>, Daniel van der Schuur<sup>1</sup>,  
Alessio Sclocco<sup>3</sup>, Samayra M. Straal<sup>7,8</sup>, Dany Vohl<sup>2,1</sup>, Stefan J. Wijnholds<sup>1</sup>, Elizabeth A. K. Adams<sup>1,9</sup>,  
Björn Adebahr<sup>10</sup>, Jisk Attema<sup>3</sup>, Cees Bassa<sup>1</sup>, Jeanette E. Bast<sup>1</sup>, Anna Bilous<sup>1</sup>, Willem J. G. de Blok<sup>1,11,9</sup>,  
Oliver M. Boersma<sup>2</sup>, Wim A. van Cappellen<sup>1</sup>, Arthur H. W. M. Coolen<sup>1</sup>, Sieds Damstra<sup>1</sup>, Helga Dénes<sup>1</sup>,  
Ger N. J. van Diepen<sup>1</sup>, David W. Gardenier<sup>1,2</sup>, Yan G. Grange<sup>1</sup>, André W. Gunst<sup>1</sup>, Kelley M. Hess<sup>12,1,9</sup>,  
Hanno Holties<sup>1</sup>, Thijs van der Hulst<sup>9</sup>, Boudewijn Hut<sup>1</sup>, Alexander Kutkin<sup>1</sup>, G. Marcel Loose<sup>1</sup>,  
Danielle M. Lucero<sup>13</sup>, Ágnes Mika<sup>1</sup>, Klim Mikhailov<sup>2,1</sup>, Raffaella Morganti<sup>1,9</sup>, Vanessa A. Moss<sup>14,15,1</sup>,  
Henk Mulder<sup>1</sup>, Menno J. Norden<sup>1</sup>, Tom A. Oosterloo<sup>1,9</sup>, Emaneula Orrú<sup>1</sup>, Zsolt Paragi<sup>16</sup>,  
Jan-Pieter R. de Reijer<sup>1</sup>, Arno P. Schoenmakers<sup>1</sup>, Klaas J. C. Stuurwold<sup>1</sup>, Sander ter Veen<sup>1</sup>,  
Yu-Yang Wang<sup>2</sup>, Alwin W. Zanting<sup>1</sup>, and Jacob Ziemke<sup>1,17</sup>

(Affiliations can be found after the references)

Received 24 May 2022 / Accepted 17 January 2023

## ABSTRACT

Fast radio bursts (FRBs) must be powered by uniquely energetic emission mechanisms. This requirement has eliminated a number of possible source types, but several remain. Identifying the physical nature of FRB emitters arguably requires good localisation of more detections, as well as broad-band studies enabled by real-time alerting. In this paper, we present the Apertif Radio Transient System (ARTS), a supercomputing radio-telescope instrument that performs real-time FRB detection and localisation on the Westerbork Synthesis Radio Telescope (WSRT) interferometer. It reaches coherent-addition sensitivity over the entire field of the view of the primary-dish beam. After commissioning results verified that the system performed as planned, we initiated the Apertif FRB survey (ALERT). Over the first 5 weeks we observed at design sensitivity in 2019, we detected five new FRBs, and interferometrically localised each of them to 0.4–10 sq. arcmin. All detections are broad band, very narrow, of the order of 1 ms in duration, and unscattered. Dispersion measures are generally high. Only through the very high time and frequency resolution of ARTS are these hard-to-find FRBs detected, producing an unbiased view of the intrinsic population properties. Most localisation regions are small enough to rule out the presence of associated persistent radio sources. Three FRBs cut through the halos of M31 and M33. We demonstrate that Apertif can localise one-off FRBs with an accuracy that maps magneto-ionic material along well-defined lines of sight. The rate of one every ~7 days ensures a considerable number of new sources are detected for such a study. The combination of the detection rate and localisation accuracy exemplified by the first five ARTS FRBs thus marks a new phase in which a growing number of bursts can be used to probe our Universe.

**Key words.** radio continuum: stars – stars: magnetars – instrumentation: interferometers – methods: observational – pulsars: general

## 1. Introduction

Of all the time variable sources observable with modern radio telescopes, none have produced such a flurry of excitement in recent years as fast radio bursts (FRBs). These bursts were discovered by Lorimer et al. (2007) and have since emerged as a unique source class, characterised by bright, short radio pulses (Fig. 1) with fluences  $\mathcal{F} \sim 1$  Jy ms. Thanks to large-scale, high time resolution radio surveys, the population of known FRBs has grown rapidly, with more than 600 FRB sources now published (see Petroff et al. 2019a, 2022, for recent reviews).

While the population of FRBs has grown rapidly, and even though a ~4% subset was discovered to repeat (Spitler et al. 2016; CHIME/FRB Collaboration 2019; Fonseca et al. 2020), their progenitors remain a mystery. The detection in 2020 of a bright FRB-like burst from a magnetar in our own Galaxy (CHIME/FRB Collaboration 2020; Bochenek et al. 2020b)

favours theories that produce FRBs in highly magnetised neutron stars, perhaps in extreme environments, in distant galaxies. However, a wide range of other theories have also been proposed, from neutron stars to black holes to stellar explosions (see Platts et al. 2019, for a theory review).

Much observational progress has occurred since FRBs were first discovered; however, many fundamental questions about their nature remain. The total fraction of FRBs that repeat is unknown, with many attempts having been made to model the underlying population (Gardenier et al. 2021; James et al. 2020). Only sustained FRB follow-up can detect repeats; whether all FRBs eventually reappear drives investigation in ongoing and upcoming surveys.

New surveys also target other FRB properties. The polarisation behaviour may provide important insights into the FRB emission mechanism, but only for a subset of FRBs could polarisation information be recorded. Of those, many are highly

linearly polarised (Michilli et al. 2018; Bannister et al. 2019; Connor et al. 2020); others, however, show distinct circular polarisation (Petroff et al. 2015a; Kumar et al. 2022). Some linearly polarised FRBs reside in highly magnetised (Michilli et al. 2018) and highly variable (Hilmarsson et al. 2021) environments.

Recently, periodic behaviour on several timescales has also been observed for a number of FRBs. Two sources, FRB 20121102A and FRB 20180916B, have distinct windows of active emission, periodic on timescales of  $\sim 160$  and 16.3 days, respectively (Rajwade et al. 2020; Cruces et al. 2021; CHIME/FRB Collaboration 2020). The activity window of FRB 20180916B is frequency dependent (Pastor-Marazuela et al. 2021; Pleunis et al. 2021b), perhaps indicative of precession or, alternatively, occultation by a binary wind. Sub-burst periodicity (periodic spacing of sub-components within a single burst) has also been observed at high significance for one FRB and at lower significance for three others (CHIME/FRB Collaboration 2022; Pastor-Marazuela et al. 2022).

Despite the high all-sky FRB rate, estimated to be a few thousand FRBs  $\text{sky}^{-1} \text{day}^{-1}$  (Chawla et al. 2017), current survey instruments face many challenges. Blind searches require covering a large range of pulse duration and dispersion measure (DM) trials. This enormous search parameter space necessitates dedicated hardware and software, typically housed on massive on-site compute clusters.

Effective FRB searches further require a large instantaneous field of view (FoV) survey instrument, to maximise the probability of observing an FRB where it goes off. In the past, this has come at the expense of localisation accuracy, and many FRBs are only localised to  $10' \times 10'$ . The study of FRB host galaxies is of increasing interest, but so far only two dozen have been identified to a sufficiently precise localisation, provided by a radio interferometer, or even very long baseline interferometry (VLBI; e.g. Marcote et al. 2020).

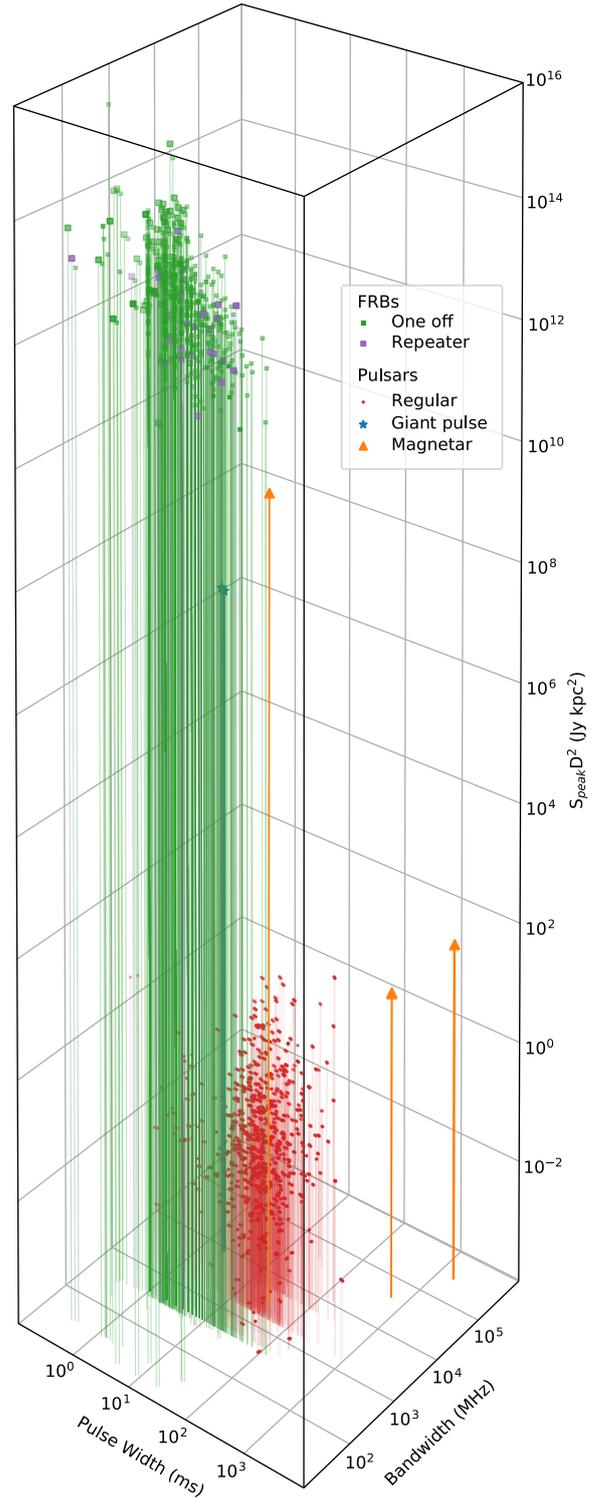
Next-generation telescopes such as the Square Kilometre Array (SKA) can no longer preserve raw survey data for offline searches (Macquart et al. 2015). New sources need to be identified in real time. This requires development of automated FRB search techniques and pipelines taking advantage of classification and machine learning tools (Connor & van Leeuwen 2018).

To address all of 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 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 further led to phased array feeds (PAFs; Sect. 2), which place additional elements at the focus of each dish in the array.

One of the largest challenges of interferometric radio astronomy is computation. Beamforming the FoV requires powerful signal processing to combine all elements in phase. This is exacerbated when combining multi-element PAF systems. Next, searching the resulting time streams for impulsive radio signals such as FRBs is an additional challenge.

Recently, faster and more agile processing has made it possible to form more beams and search them in real time. Many searches now employ large, dedicated compute clusters to deal with the massive amounts of data streaming from the telescope, and distribute it over numerous processing nodes (Sect. 3).

The raw localisation ability of an interferometer depends on its baseline length, and ability to form all resolution elements. Even more precise localisation is possible for brighter signals



**Fig. 1.** Phase space diagram of emission duration and emission bandwidth versus pseudo luminosity ( $S_{\text{pk}} D^2$ ), for different fast radio transients at peak flux  $S_{\text{pk}}$  and distance  $D$ . FRBs are similar to other fast transients in their short and relatively broadband emission; however, in pseudo luminosity, they strongly stand out from the pulsars, magnetars, and giant pulses also plotted. Catalogue data, per 2022 Feb 1, taken from FRBCAT (Petroff et al. 2016), the Transient Name Server (TNS; <http://www.wis-tns.org>), ATNF (Manchester et al. 2005), and EPN (<http://www.epta.eu.org/epndb/>) were complemented by individual data from Camilo et al. (2006), Maan et al. (2019), Dai et al. (2019), Pearlman et al. (2018), Mikami et al. (2016), Bera & Chengalur (2019), and Bochenek et al. (2020a).

that appear in several beams (cf. Sect. 5.1.4), or when refined correlation of voltage streams is possible offline.

Only interferometers can provide the arcsecond or better localisation needed to identify an FRB host galaxy unambiguously (Eftekhari et al. 2018). Using arrays such as the Australian Square Kilometre Array Pathfinder (ASKAP), the Deep Synoptic Array-10 (DSA-10), the European VLBI Network (EVN), and the Very Large Array (VLA), more than 20 FRBs were traced back to their hosts (e.g. Chatterjee et al. 2017; Bannister et al. 2019; Ravi et al. 2019; Marcote et al. 2020).

Since repeat characteristics are unknown, FRBs generally need to be localised from the discovery pulse. Identifying it in real time in the incoming stream can aid such localisation efforts by preserving data at higher resolution and/or with the full Stokes parameters (e.g. Petroff et al. 2015a; Fonseca et al. 2020).

A real-time FRB detection also enables other telescopes to be triggered to conduct multi-wavelength or multi-messenger follow-up. Several theories (Platts et al. 2019) predict emission in optical, X-ray, or gravitational waves. As of yet, no such emission has been observed for extragalactic FRBs. X-ray emission from the Galactic magnetar SGR 1935+2154 coincident with its bright radio burst (Mereghetti et al. 2020) suggests counterparts may be detectable if follow-up observations are conducted promptly. Triggers on FRBs are now possible with Virtual Observatory Events (VOEvents; Petroff et al. 2017).

The frequency range over which FRBs are detectable remains unknown. Early searches with the LOw Frequency ARray (LOFAR) and the Murchison Widefield Array (MWA) at frequencies of 100–200 MHz were unsuccessful (Coenen et al. 2014; Sokolowski et al. 2018; Chawla et al. 2020). However, studies of repeating FRBs at low frequencies, including commensal observing between WSRT, the Westerbork Synthesis Radio Telescope, and LOFAR, have proven more fruitful. LOFAR detected several bursts from repeating FRB 20180916B, leading to the discovery of its frequency-dependent activity window (Pastor-Marazuela et al. 2021; Pleunis et al. 2021b). Finding coincident emission at high and low radio frequencies generally requires real-time classification.

To investigate some of these outstanding core questions on FRBs, we designed the Apertif Radio Transient System (ARTS; van Leeuwen 2014) to harness Apertif, the APERTure Tile In Focus (van Cappellen et al. 2022) on WSRT for finding and studying FRBs. The system operates on an interferometer, the WSRT, to provide a spatial resolution of the order of  $10'' \times 10''$ . Apertif, the new PAF front-end system at Westerbork, delivers a large, instantaneous FoV of 8.2 sq. deg. Where the previous system produced a single primary beam, the new front ends generate 40. Together, these increase the FoV by a factor 30. Although the antenna elements and amplifiers are not cryogenically cooled, Apertif operates at a nominal system temperature of 70 K (Oosterloo et al. 2009). The bandwidth is 300 MHz around 1.4 GHz in two linear polarisations. Of the 12 dishes outfitted with PAFs, the eight equidistant ones that are most efficiently combined deliver a gain of  $\sim 1 \text{ K Jy}^{-1}$ . That sensitivity is comparable to the successful Parkes FRB surveys, but at a much larger FoV. A real-time search system identifies FRBs at high time and frequency resolution (82  $\mu\text{s}/195 \text{ kHz}$ ). Apertif thus has the potential to deeply survey the entire fast-transient sky in radio.

In this paper, we describe ARTS and its first discoveries. We outline its science potential (Sect. 2), describe the design and implementation (Sects. 3–5), and present the commissioning results, performance, and survey planning (Sects. 6–8). In Sect. 9 and 10, we present and discuss our first science results, which include the discovery of five new FRBs. We close with

data releases (Sect. 11), future plans (Sect. 12), and conclusions (Sect. 13).

The above sections outline the complete system, providing the reader with the overall rationale, setup, and results. More specialist readers are suggested to start with the sections relevant to their interest: Astronomers mostly focusing on the newest FRB detections could read Sects. 2, 6, 9, 10, and 13. Time-domain specialists could additionally refer to Sects. 3–5 and 7. Astronomers interested in FRB populations and surveys should also consult Sects. 7, 8, and 12. For reproducing or improving the system and results, readers can use Sect. 11 and Appendices A–D.

## 2. Apertif for the time domain: ARTS

The Apertif time-domain potential is realised through the Apertif Radio Transient System (ARTS), an instrument that can search wide fields for fast transients.

Downstream from the PAFs, high-throughput and high-performance processing hardware enables analysis of this full field and bandwidth. Processing revolves around field-programmable gate arrays (FPGAs) for channelisers, correlators, and a hierarchy of beamformers (BFs), and around CPUs and graphics processing units (GPUs) for pipelines. These are directed by modern firmware and software, powering a real-time detection system that combines the new, large FoV with the high angular resolution of the array. ARTS is directly connected to LOFAR (van Haarlem et al. 2013), also operated by ASTRON.

### 2.1. Science motivation

The system was designed around our goals of discovering and characterising FRBs, primarily; and our aims of finding and studying neutron stars, and the time-domain counterparts of slow, image-domain transients, secondarily.

#### 2.1.1. Fast radio bursts

Using FRBs as cosmological tools and understanding the underlying source population(s) requires a large number of detections – thus, a survey with good sensitivity, FoV, and time on sky. These three characteristics are provided by ARTS. As Apertif is a full-time survey machine, ARTS detects an FRB roughly every week of observing (see Sect. 12 and Maan & van Leeuwen 2017). The instrument combines high time and frequency resolution with the ability to capture full-Stokes polarisation data. It features a new GPU FRB search pipeline, and a machine learning classifier to better identify and trigger on FRB candidates.

Potentially most important in determining the formation of FRBs is localising the bursts. If FRBs are formed by young neutron stars, the galaxy in which they reside will need to have recently been forming massive stars. To identify the host galaxy with high confidence, the FRB position error box must be small enough to hold only a single candidate host. For nearby bursts, sub-arcsec localisation, using VLBI, can connect the (repeating) FRB emitters with features within the host galaxy. ARTS contains the hardware to connect the WSRT for VLBI (Sect. 4.1.3); but more importantly, the addition of Apertif onto an interferometer with a baseline of over 1 km provides good instantaneous FRB localisation (Sect. 7.4).

#### 2.1.2. Neutron stars

The surface gravity of these extremely compact stars, about  $10^9$  times the gravity on Earth, is the largest of any object visible in

the Universe. The internal densities of ten times nuclear density have not existed elsewhere since the Universe was  $\sim 1$  ms old.

That combination turns pulsars, radio-emitting neutron stars, into near-perfect cosmic time keepers (e.g. [Hulse & Taylor 1975](#)). Performing high precision pulsar timing on double neutron star (DNS) systems, informs us of the underlying binary evolution (e.g. [van Leeuwen et al. 2015](#)), and enables tests of general relativity ([Kramer et al. 2006](#); [Desvignes et al. 2019](#)).

Three groups of neutron stars are only very sporadically active in radio: rotating radio transients (RRATs), intermittent pulsars, and radio-transient magnetars. Given the odds against their detection, the number of such transient neutron stars must be comparable to that of radio pulsars ([Keane et al. 2011](#)).

ARTS performs full-FoV searches for single pulses of such pulsars, in a survey mode. It can also provide a high-time-resolution data stream, with real time coherent dedispersion, and online folding, for a timing mode. Together, these allow both searches for and studies of radio-emitting neutron stars.

### 2.1.3. Prompt and slow emission from neutron-star mergers

The orbits of DNSs decay, due to the emission of gravitational waves, until eventually the neutron stars merge. This coalescence produces radio emission, possibly on different timescales (see, e.g., [Chu et al. 2016](#)). Prompt emission could possibly form at the merger, or at the collapse of an intermediary supramassive neutron star. Incoherent radio emission will produce an image-domain, slow radio transient ([Hallinan et al. 2017](#)). Apertif allows for the study of both kinds of emission. Afterglow searches are carried out in image domain ([Boersma et al. 2021](#)) while prompt emission could be detected in time domain.

## 2.2. Key Apertif features for time domain studies

### 2.2.1. Regular interferometer

Wide-field transient surveys at both full (coherent-addition) sensitivity and high time resolution are almost always limited by the large compute demands. In the LOFAR surveys, the tied-array beams (TABs) cover about 1/9th of the potential FoV (the station beam; [Coenen et al. 2014](#); [Sanidas et al. 2019](#)). In the Square Kilometre Array (SKA), the planned 1500 SKA-Mid beams cover one-third of the primary beam ([Smits et al. 2009](#)). No sparse 2D or irregularly laid-out interferometer has the capability to find sub-ms transients at full sensitivity, over its entire FoV. One unique feature of the WSRT is its linear and regularly-spaced, E–W layout. This produces full-sensitivity TABs that are not single, small circular beams, but have a large area in the N–S direction, and multiple sensitive sidelobes ([Janssen et al. 2009](#)). This means coherent beams do not need to be computed over the entire 2D FoV, but only over the 1D distance between the main lobe and the first sidelobe (Sect. 2.3.2). Exploiting this fact allows us to overcome the compute limitations of the full-FoV surveys with for example LOFAR and SKA.

A down side of the elongated TABs is the modest N–S localisation precision for short-lived transients. The instantaneous E–W localisation is very good ( $\sim 10''$ ). Additionally, every time a source repeats the localisation ellipse has rotated on the sky, thus producing an accurate final localisation of about  $10'' \times 10''$  (see e.g. [Maan & van Leeuwen 2017](#), and Sect. 7.4).

### 2.2.2. Steerable dishes

For FRB and slow transient science, the ability of WSRT to point its dishes holds significant advantage over other wide-field surveys in the northern hemisphere. A fixed instrument such as

CHIME is able to search its full FoV in part because budget for compute was prioritised over steerable elements. This comes with the limitation that sources can only be observed for a relatively short time, at transit. Because the WSRT can point, and track much longer than fixed apertures, Apertif is able to dedicate significant follow-up time to repeating FRBs that are known to be active (cf. [Pastor-Marazuela et al. 2021](#)).

### 2.2.3. Digital backend: Commensal systems

The pulsar surveys in which the first FRBs were found (e.g. [Lorimer et al. 2007](#); [Thornton et al. 2013](#); [Spitler et al. 2014](#)) operated only during a relatively small fraction of the total available telescope time. And yet, for rare bursts such as these, the total time on sky is exceedingly important. Systems that search for FRBs in a secondary mode – commensal to other, primary projects that determine, for example, the pointing – are thus very valuable. Commensal systems such as CRAFT ([Macquart et al. 2010](#); [Bannister et al. 2017](#)) and CHIME/FRB ([CHIME/FRB Collaboration 2019](#)) have found FRBs with this approach. Exploiting the fact that the incoming Apertif PAF data and the modular hardware can be easily duplicated, ARTS was designed and built to run in parallel<sup>1</sup> to all other imaging Apertif surveys ([Hess et al. 2023](#)). Coherently combining the 8 WSRT dishes that are equidistant in the imaging “Maxi-Short” configuration, would produce a commensal transient search with the same wide field, and similar sensitivity, as the dedicated survey.

### 2.2.4. Digital backed: High resolution

The high time and frequency resolution (82  $\mu$ s, 195 kHz) of the ARTS backend offers an advantage for both the study of known pulsars and FRBs (see, e.g., [Bilous et al. 2022](#); [Oostrum et al. 2020](#), respectively) as well as the detection of new sources. Many FRBs appear to be sub-millisecond in duration, which means a large number of bursts are currently missed due to the deleterious effects of instrumental smearing ([Connor 2019](#)). ARTS searches filterbank intensity data with more favourable smearing properties than comparable instruments, boosting its detection rate and enabling the discovery of higher-DM, narrower FRBs.

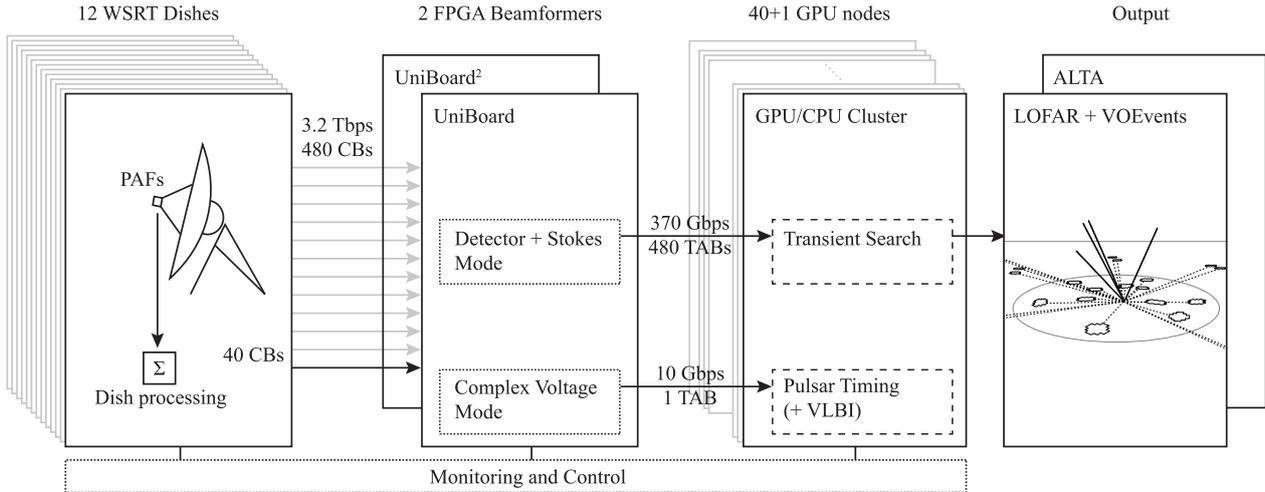
### 2.2.5. Connection to LOFAR

WSRT can localise an FRB very well in the E–W direction but only to modest N–S precision (cf. Sect. 2.3.2). But if it emits over a broad band, this same FRB will arrive later at low frequencies. Triggering the LOFAR all-sky transient buffer boards (TBBs; [Stappers et al. 2011](#); [ter Veen 2015](#); [ter Veen et al. 2019](#)) on the FRB, and detecting it at 150 MHz, would allow for localisation using the 2D array of LOFAR stations over The Netherlands and Europe, providing arcsec precision.

Previous blind LOFAR searches ([Coenen et al. 2014](#); [Karastergiou et al. 2015](#); [Sanidas et al. 2019](#); [ter Veen et al. 2019](#); [van Leeuwen et al. 2020](#)) have not found FRBs. This could partly be caused by the residual dispersion smearing that remains after trial dedispersion at such low frequencies. In the triggered case, the DM is known, allowing for more accurate dedispersion plus a significant reduction of the number of DM trials, improving detection confidence levels.

Still, also when the DM was known, low-frequency FRB observations were unsuccessful before 2021 (see e.g. [Houben et al. 2019](#)). Only with the detection of FRB 20180916B were FRBs seen for the first time below 200 MHz; and while we found

<sup>1</sup> At time of writing this mode, however, is not commissioned.



**Fig. 2.** Top-level diagram for ARTS. Different line styles indicate different types of subsystems, each described in their own section. Full lines are hardware elements (Sect. 3). Dotted boxes depict firmware and software (Sect. 4). Dashed items are science pipelines (Sect. 5). The LOFAR and ALTA output are described in Sect. 11.

**Table 1.** Observing modes, the science cases they enable, and the observing time total for 2019 (this paper).

Mode	Science case	Parallel to	Time (h)
Dedicated search	SC4	–	800
Commensal search	SC3	Imaging	0
Timing	SC1	–	50

that the all-sky low-frequency FRB rate is appreciable (Pastor-Marazuela et al. 2021), no single FRB burst was ever seen to emit at both 150 MHz and 1.4 GHz. The simultaneous ARTS-LOFAR observing that enabled these detections is detailed in Pastor-Marazuela et al. (2021).

### 2.3. High level system overview

Time-domain observations with Apertif fall under four main science cases (SCs; detailed in Straal 2018): (1) Pulsar Timing, (2) VLBI, (3) Fast-transient searching commensal to imaging, and (4) Dedicated transient searching. These were numbered to reflect increasing complexity; their priority is the inverse, the dedicated search ranking highest (Table 1). Their scientific and resulting technical requirements led to designs that were reviewed and implemented. After commissioning (see Mikhailov 2018; Straal 2018; Oostrum 2020b, and Sect. 6) the system was brought into operation. Together, a number of subsystems (Sect. 2.3.1) form and process a hierarchical series of beams (Sect. 2.3.2). Transient detections in ARTS produce triggers to allow for follow-up in close to real time, and data products stored in the Apertif long-term archive (ALTA), where they are publicly available.

#### 2.3.1. Three sets of subsystems

ARTS comprises the following major sets of subsystems. These work together as illustrated in Fig. 2.

The first subsystem set is the hardware platform (Sect. 3). It consists, first, of the WSRT dishes and the Apertif PAFs. Front-end beamformers next provide dish processing. Tied-array

beamformers built using FPGAs on high-performance processing boards (UniBoard and UniBoard<sup>2</sup>), are connected through fast networking to a GPU cluster.

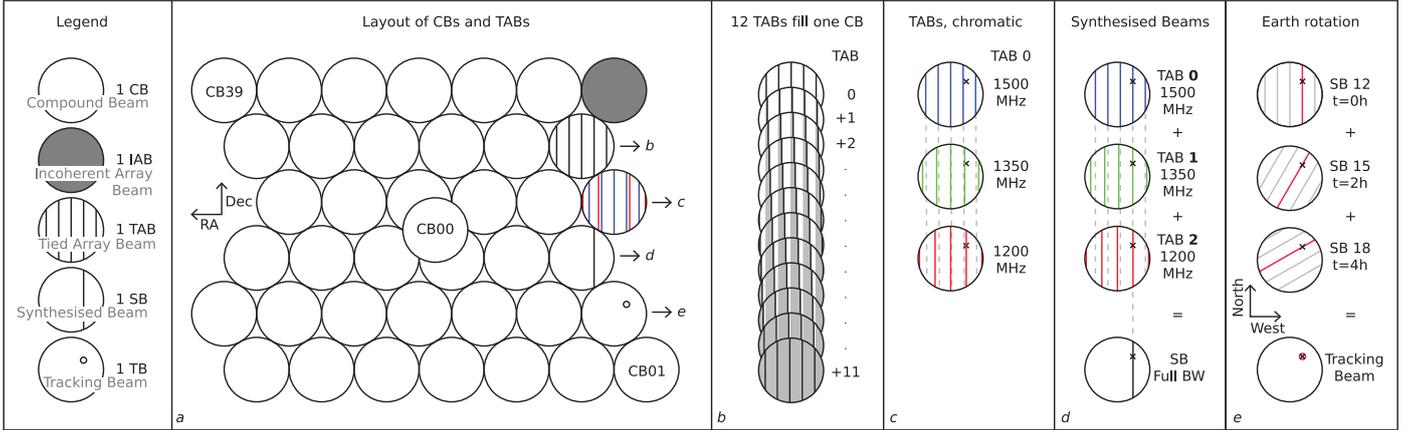
The second set is the firmware and software sub systems (Sect. 4) that control and produce one or multiple TABs, or ‘pencil beams’. These can be in Nyquist sampled, complex-voltage format for pulsar timing (and VLBI). The data in these TABs can also be ‘detected’, that is, converted to the four Stokes parameters, allowing subsequent partial integration to reduce data rates. That way, many hundreds of beams can be streamed out.

The third set of subsystems comprises the ARTS pipelines (Sect. 5). These perform transient searching and pulsar timing. For timing, the central TAB is coherently dedispersed and folded in real time, on a single multi-GPU node. For the transient search, all Stokes-I TABs are cleaned of radio frequency interference (RFI), dedispersed over a number of trial DMs, corrected for chromatic effects, and searched for transient events. Good candidates immediately trigger data dumps from a ring buffer of full Stokes-IQUV data. A deep learning implementation further classifies all candidates. Results and data are public to the outside world through VOEvents and the archive (Sect. 11).

#### 2.3.2. Hierarchical beam forming

One of the innovative aspects of ARTS is its use of hierarchical beamforming to allow for searches throughout the entire primary-beam FoV, at coherent-addition sensitivity. That is a challenge in many wide-field fast-transient instruments such as SKA (Backer 2000) and LOFAR (van Leeuwen & Stappers 2010). We give a short conceptual overview of these beam forming steps here, because they are important in all subsystem sets. Appendix D contains a quantitative description.

Figure 3 provides an overview of all constituents of the beamforming hierarchy. The base is formed by the 40 compound beams (CBs) formed by the PAF in each Apertif dish. Each CB has the same FoV as the primary dish beam, with a diameter of  $\sim 0.5^\circ$ , and produces a full-bandwidth, Nyquist sampled output data stream. Together, the CBs cover the FoV of the Apertif system as shown in Fig. 3a.



**Fig. 3.** Overview of the beam hierarchy in ARTS. Panel (a) shows the 40 CBs formed by each PAF. The output signals of these CBs can be combined into IABs or coherent TABs, whose grating responses fill the entire CB as shown in panel (b). The frequency-dependence of these grating responses shown in panel (c) can be exploited to disambiguate them by combining signals from multiple TABs to form a SB as illustrated in panel (d). Finally, a TB can be formed towards a specific locus by combining multiple SBs over time as shown in panel (e).

The data streams of corresponding CBs of multiple dishes are combined using either an incoherent beamformer or a coherent beamformer. Using the incoherent beamformer, 40 incoherent-array beams (IABs), one per CB, are formed that together cover the compound FoV of Apertif with a sensitivity scaling with the square root of the number of dishes used. When using coherent beamforming, the peak sensitivity is higher, scaling linearly with the number of dishes, but more data streams need to be analysed as the FoV per TAB is smaller. For a regularly spaced 1D array like WSRT, each TAB is not a single field: all sidelobes reach the same sensitivity as the central peak, and add to the TAB FoV. Together, 12 TABs now completely cover the FoV of their CB. This high TAB filling factor is illustrated in Fig. 3b. It was first used by Janssen et al. (2009), to find new pulsars. Such full FoV beam forming requires that we only use those dishes that are spaced equidistantly. These are the 8 dishes RT2–RT9, plus the movable dishes RTA and B if they are located on the common baseline grid.

While we exploit the TAB grating to increase the TAB FoV, it also presents a degeneracy in the location of a detected event. We may not know through which grating sidelobe the signal is coming in. Fortunately the grating responses are frequency dependent – see Fig. 3c. A source slightly further from the phase centre than the specific grating response of a certain TAB at 1500 MHz, may fall exactly in the grating response of that TAB at 1200 MHz. Thus, most grating responses will detect a source over limited frequency range. To make a detection over the full bandwidth the correct frequency-dependent TABs are combined to form a so-called synthesized beam (SB), as illustrated in Fig. 3d. In ARTS, 71 SBs are formed in each CB. As it is integrated into the subband dedispersion (Sect. 5.1.1), this frequency re-organisation itself comes with no additional computational cost; but the increased number of beams (12 TABs to 71 SBs) do require more computing downstream. The SBs provide instantaneous localisation with a resolution determined by the array size in one direction, and the size of the CB in the other.

The orientation of the grating responses on the sky changes with time due to Earth rotation. A specific locus thus moves through different SBs (Fig. 3e). By combining the best positioned SBs in time, we form the last constituent of our hierarchy of beams: the tracking beam (TB). Over a 12-h observation, ~3000 unique loci form within each CB.

### 3. Description of time domain system. I. Hardware and function

A general description of Apertif is available in van Cappellen et al. (2022). In the next three sections we describe the time-domain capabilities. A number of subsystems in Apertif are shared between imaging, time-domain and VLBI modes. Where required for context, we first provide high-level descriptions of these general subsystems. This section covers function and hardware, with additional detail available in Appendices A and B, respectively.

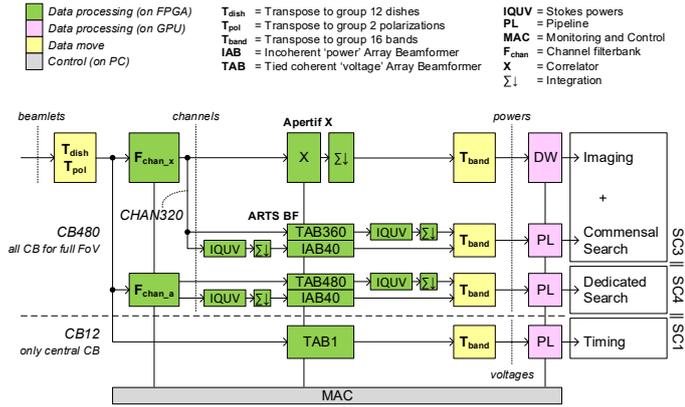
#### 3.1. Relevant general Apertif hardware and function

##### 3.1.1. Dishes in an E-W Interferometer

The WSRT consists of fourteen 25-m dishes in a linear east-west array. It started operations in 1970 (for a recent overview see Strom et al. 2018, marking its 50th anniversary). The accuracy of the steel structure, the parabolic surface of the dish, and the size of the mesh, mean WSRT can be used up to ~8 GHz. It was, however, designed for 21-cm wavelength observing. At the installation of Apertif the dishes too were refurbished and they operate as new. The mounts are equatorial. The 10 most western dishes, radio telescope (RT) 0–9, are fixed at redundant, 144 m intervals. RTA and B are immediately east of RT9, and can be moved on a rail track. RTC and D are on a track that is 1.3 km further east. Because of this dominant common baseline, the instantaneous interferometric beam has highly sensitive sidelobes.

##### 3.1.2. Phased array feed (PAF)

Twelve of the WSRT dishes, RT2–RTD, are outfitted with the Apertif PAFs. Each feed contains 121 receiver elements, as opposed to the previous generation of receivers, the multi frequency front ends (MFFEs; cf. Tan 1991; Bregman et al. 2018), that were single pixel. The Vivaldi antenna elements in the PAFs are optimised for 1500 MHz. The elements also host the low-noise amplifier (LNA), and are uncooled. The LNAs include a filter that suppresses the strong RFI present below the nominal bottom of the Apertif band at 1130 MHz. This does, however, contribute ~20 K to the total system temperature of ~70 K



**Fig. 4.** ARTS TABF subsystem and the Apertif X subsystem. The modes for the three SCs (Table 1) are mutually exclusive.

(van Cappellen et al. 2022). For the 12 dishes combined, and for the central, most sensitive CB, this produces a system-equivalent flux density (SEFD) of  $\sim 45$  Jy at 0.75 aperture efficiency. For the 8 dishes generally used in ARTS, we find a median SEFD over all CBs of 85 Jy (Sect. 7.1).

### 3.1.3. Front-end digital processing

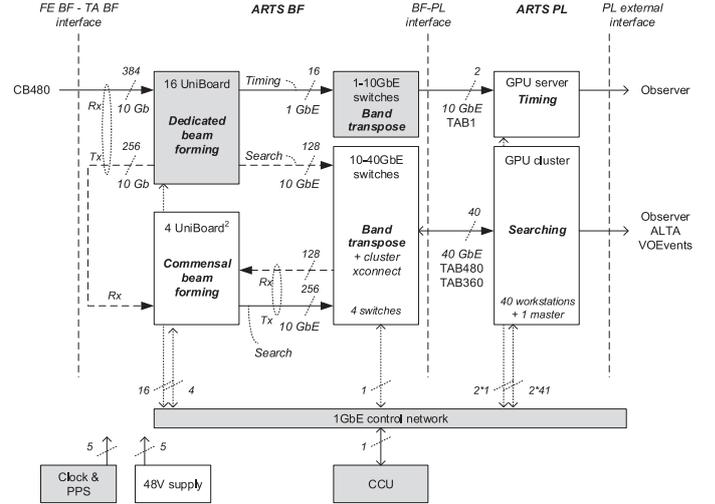
At the 12 dishes, the Apertif Front-End Beamformers (FEBFs; Appendix A.1) sample the radio frequency (RF) data from the PAF at 800 MHz. This yields a timing accuracy of 1.25 ns, at a total input data rate of 9.8 Tbps. The sampled bandwidth is a contiguous 400 MHz band, tuneable between 1130–1720 MHz. The digitised data of all antenna elements are separated into 512 subbands, each 0.781250 MHz wide. At each dish, a rack of eight UniBoards (UNBs; Szomoru 2010) beamforms the subband signals over the PAF elements, into so-called beamlets.

The CB pattern formed by these beamlets is laid out in a configuration that maximises even coverage over multiple, adjacent pointings (Hess et al. 2023). Initially Apertif was designed for 37 CBs (Oosterloo et al. 2009). Because the Apertif and ARTS hardware produce CBs in multiples of 8, the surveys now use 40.

Next, 384 out of the 512 subbands are selected for output, to achieve the final 300 MHz bandwidth. When expressed as 6-bit numbers, these fill the data transport capacity over the 384 links of 10 Gbps each, to the central building for correlation and/or tied-array beamforming. In total, the twelve Apertif-equipped dishes produce  $12 \times 40$  CBs. This data stream (called CB480) contains 3.5 Tbps of compound-beam data (Appendices A.1, B.1; and Table A.1).

### 3.2. The ARTS tied-array beamformer (TABF) and Apertif correlator (X) subsystems

The Apertif correlator ('X') and ARTS TABF both use the same Apertif FEBF CB480 data. In Fig. 4 we show how these subsystems relate. The transpose  $T_{\text{array}} = T_{\text{dish}} + T_{\text{pol}}$  groups the data from all  $N_{\text{tp}} = N_{\text{pol}} \times N_{\text{dish}} = 24$  single polarisation telescope paths (TP) in the WSRT array. The Apertif X creates visibilities by cross correlating the compound beams between all dish pairs. The ARTS TABF creates array beams by summing compound beams over the dishes. These array beams can be coherent TABs or IABs, for either voltage, power or full Stokes power data (Appendices A.3 and A.4).



**Fig. 5.** Mapping of the ARTS system on the UNBs, UNB2s, switch, and CPU/GPU cluster hardware. Only one of either the commensal or the dedicated beam forming platforms will be running at a given time. The dashed Tx links are unidirectional links. The dashed Rx links are full duplex links that are combined with the UNB2 Tx links. Parts in grey are hardware used by or shared with Apertif X.

### 3.3. The ARTS tied-array beamformer (TABF)

The function of the TABF is to create the IAB and TAB data, and to transpose the data for the ARTS pipeline processing (PL; Sect. 5). For the search data, the TABF also functions to reduce data rates, by integrating in time.

For searching, the ARTS TABF first separates the beamlets into 64 subchannels, using a similar filterbank ( $F_{\text{chan}_x}$  in Fig. 4) as Apertif X. The search modes use all 40 CBs. The TABF can combine all 12 dishes (480 CBs) but in practice only includes dishes that are equidistant in the array configuration at the time. The output of the IQUV data is integrated over 16 subchannels to achieve the required 81.92  $\mu$ s time and 195 kHz frequency resolution (Appendix A.2). The standard TABF output are integrated full Stokes IQUV array power beams, 480 for dedicated search and 360 for commensal search. Optionally, IABs can be produced. In that case, the order of beamforming and detection is switched.

For timing-mode (and VLBI) output, the FEBF transposing is bypassed (see Appendix A.1 and Fig. A.2) and the beamlets are used directly. The central CB is voltage beam-formed into a single TAB (TAB1 in Fig. 4).

The central beam-forming hardware needs to handle the large amount of input links, and the high data rate of  $\sim 3.5$  Tbps, produced by the Apertif dish processing (Sect. 3.1.3). The relatively low number of operations per bit led to the implementation of the TABFs on FPGAs. In dedicated ARTS mode, the time-domain survey determines the baselines and the pointing of the WSRT dishes. As this mode runs instead of Apertif X, the TABF can use the same 16 UniBoards used by the correlator, as detailed in the next subsection.

In commensal mode (Sect. 2.2.3), the Apertif imaging science case is in control, and the correlator runs. Firmware for both does not fit together on 16 UniBoards, but needs to run on a separate set of FPGA boards. Four UniBoard<sup>2</sup>s (UNB2s) can provide the required input/output (I/O) and compute in a setup described in Sect. 3.5.

Figure 5 shows how the ARTS FPGA beamformer maps on UniBoard and UNB2, and how the boards are interconnected.

### 3.4. Dedicated ARTS TABF on central UniBoards

The central ARTS beamformer consists of 16 UNBs with a total of 128 FPGAs (for details, see Appendix B.2). Each UNB processes 1/16th of the bandwidth (i.e. 18.75 MHz) of all dishes.

The UNBs are equipped with optical-electrical boards (OEBs; Appendix B.2) to handle the demanding I/O. All 384 10-GE ports are required for data that comes in. For X, outgoing data can fit on the 1 Gbps ports because the time integration strongly reduces the data rate. But for the TABF the complete usage of the links by the input is a challenge. The ARTS dedicated search makes 12 TABs/CB that also require 128 10-GE links, for transport and transpose of the data to the PL. For all 384 ports, the Rx and Tx are thus physically split. The UNBs receive Apertif FEBF data over Rx and send beam-formed data over Tx to the PL.

During dedicated FRB survey observations, the UniBoards generate 480 TABs with a time resolution of 81.92  $\mu$ s and 1536 frequency channels over a bandwidth of 300 MHz. For each TAB, both a Stokes I and a Stokes-IQUV data stream are created, for a total output data rate of  $\sim$ 360 Gbps.

### 3.5. Commensal ARTS TABF on Central UniBoard<sup>2</sup>s

In the central building, a set of four UNB2s (Schoonderbeek et al. 2019) is also installed, to be able to run imaging and time-domain at the same time. As that commensal search uses as input only the 8 dishes that are equidistant in the imaging Maxi Short configuration, it requires making only 9 TABs/CB. Four UNB2s each with four Arria10 FPGAs (see Appendix A.5) can create these beams, but then a channel filterbank does not also fit. Therefore the TABF for commensal search uses the output from the channel filterbank in the Apertif X, that will then be running on the UNBs. The filterbank in the Apertif X correlator already makes 64 channels, the same as the dedicated search. The UniBoard<sup>2</sup>s receive the CHAN320 data (8 dishes  $\times$  40 CBs) from the 16 UNBs, as visualised in Fig. 4 and detailed in Appendix B.3.

### 3.6. Networking and interconnect

The central UNBs receive the CB480 input from the Apertif FEBF (Fig. 5) over the 384 10G Rx links (Sect. 3.4). The 10G Tx part of these same links is used to offload the output TAB data, or to pass on data to the 4 UNB2s. For the dedicated search 360 Gbps (Table A.1) of TAB power data are carried via 128 10G links. Each UNB or UNB2 beam forms all telescopes and CBs but only for part of the bandwidth of these TAB power data. For searching, these data need to be transposed such that all bands for a single set of TABs converge in a single PL GPU node (Fig. B.4). In the data network, this band transpose function  $T_{\text{band}}$  is implemented using commercial off the shelf (COTS) 10/40 GbE switches (called the transposer in Fig. 5 and in Appendix B.5).

In commensal-search mode, the UNBs operate as correlator, and only pass on the CHAN320 data to the UNB2s. For timing, the modest data rate of 9.6 Gbps (2 pols  $\times$  300 MHz  $\times$  2 complex  $\times$  8 bit; Table A.1) is carried via 16 1GbE links.

### 3.7. ARTS GPU cluster

Data for pulsar timing are coherently dedispersed and folded by a single GPU-accelerated server (“ARTS-0”). This machine contains two 6-core CPUs, two Titan-X GPUs, two 10GE Network Interface Cards, and 16 active disks in a JBOD configuration.

**Table 2.** Hardware overview of a GPU cluster node.

CPU	2 $\times$ Intel Xeon E5-2640 v4 (40 cores total)
GPU	4 $\times$ Nvidia 1080 Ti
RAM	128 GB
Network	40 Gbps, full-duplex
Storage	32 TB in 2 independent hardware RAID0

Data for searching are much higher volume, and are processed by a GPU cluster consisting of one login head node and 40 worker nodes, all with identical hardware (Table 2). Together the latter provide a total of 160 GPUs, 1600 CPU cores, 5 TB RAM, and  $\sim$ 1.3 PB storage, for a theoretical peak performance of  $\sim$ 2 PFLOP s<sup>-1</sup>.

Each worker node processes the data from one compound beam. That way, cross detections can immediately be compared between TABs in a CB; and tracking beams made. The incoming data rates per CB are 1.8 Gbps of Stokes I data, and 7.2 Gbps of Stokes IQUV data. Half of the storage is available for the incoming Stokes I data, which when writing continuously fills up in  $\sim$ 20 h. Hence any processing that requires access to these data should take at most a few hours to run.

Transient data buffers were initially designed to reside on both the UNBs and UNB2s. The GPU cluster, however, could hold more DDR3 memory more economically, in a single location, with simpler transient-dump control functions. Thus, the ARTS BF now outputs full Stokes IQUV (for dumping), instead of only the Stokes I data (for searching). These Stokes IQUV data are only stored upon a trigger from the transient pipeline, so even though the Stokes IQUV data rate is four times that of Stokes I, its storage will fill up very slowly.

Both Stokes I and IQUV data are buffered in RAM, for 10 and 15 s, respectively, and made available to the pipelines. The system buffers about 1 TB of data for triggering overall.

### 3.8. Housing

ARTS is located in an RF-shielded high frequency cabin in the central building of WSRT. The cabin contains coolers with air-water heat exchangers. This water is pumped through an outside chiller, that operates passively when possible, and only uses compressors when required. In the cabin, three cabinets house the UNBs and switch equipment. These are air cooled through the cabin ambient cooling. The ARTS GPU servers and UNB2s are distributed over 6 racks in a single row. Three large dedicated reverse coolers create a cold and hot isle. Cold air flows through the servers to the hot isle. The CPUs and GPUs are cooled against this air. The FPGAs on UNB2 are liquid cooled, and exchange heat with the air flow outside the UNB2 box. When operational, the total ARTS system uses 50 kW of power.

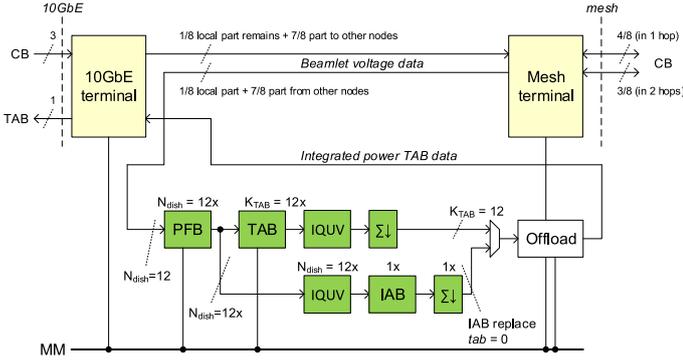
### 3.9. Control

Through the control network, ARTS and Apertif are directed by a Central Control Unit (CCU) computer that connects to the TABF and the PL via Ethernet. The CCU is synchronised with the atomic clock of the WSRT using the standard network time protocol (NTP).

## 4. Description of time domain system. II. Firmware and software

### 4.1. Firmware

In the following subsections we describe how the ARTS TABFs in Fig. 2 are implemented in hardware description language



**Fig. 6.** Top level block diagram for Dedicated Search FPGA application design on UniBoard.

(VHDL) firmware on the FPGAs on UNB and UNB2 (Sect. 3.4, 3.5). More detail is available in Appendix C.1.

Each observing mode (Table 1) has a dedicated FPGA application design that either runs on all 128 FPGAs of the TABF UNBs or on all 16 FPGAs of UNB2. The designs contain functions for the TABF itself, for monitoring during operation, and for testing during commissioning. Switching between modes occurs by changing the design image. The flash memory of UNB and UNB2 typically contains a default image and an application image. Changing observing mode requires reprogramming the flash. In practice this reprogramming takes  $\sim 5$  min for all FPGAs in parallel. This time includes steps to erase, write and read back to verify. Such reprogramming is part of the scripted mode switching that operators can initiate or schedule. Together with the required re-initialisation of the other subsystems, this switch takes  $\sim 15$  min.

There are three different FPGA images for the 16 UNBs: search, timing and correlation. The 4 UNB2 can run two images: the commensal search and the dedicated-search pass through.

The heart of the TABF system is a voltage beamformer function that weights and sums the input dish voltages. For the single voltage TAB (for timing or VLBI), the VHDL component from the subband CB beamformer in the Apertif FEBF (van Cappellen et al. 2022) was directly reused. The TABF function for surveying contains further multi-beam parameters.

The subchannel filterbank and integrator functions are weaved into one. This reduces aliasing losses. On the one hand, samples should be integrated in time to  $81.92 \mu\text{s}$  to reduce data rate. On the other hand, the fine-channel polyphase filter bank (PFB: called  $F_{\text{chan}_a}$  in Fig. 4) should separate the beamlet data into 4 channels. But as channel 0 represents information from both the low and high edge of the subband, it is best discarded; implying 25% data loss. To mitigate this, the subchannel PFB first forms 64 channels. Subchannel 0 is blanked. The Stokes power data are then not integrated in time, but over sets of 16 frequency subchannels in stead (see Fig. A.4).

Some functions in the FPGA application design require monitoring and control (MAC). These act through a memory mapped (MM) control interface as visible in Figs. 6 and A.2. Some are set only at start of measurement (e.g. PFB coefficients); others are updated regularly during the measurement (e.g. statistics, status monitoring).

#### 4.1.1. Multiple TABs on UniBoard

The dedicated-search FPGA design (Fig. 6) takes CB480 data as input. Upstream, the FEBFs already compensated for the time-varying geometrical delay due to the rotation of Earth.

That fringe stopping (FS) was applied per CB by coarse, true sample delay tracking (DT) plus fine-delay phase tracking (PT; Appendix A.1 and van Cappellen et al. 2022). The FEBFs also already transposed the CB480 data by  $T_{\text{int}_x}$  (Fig. A.2).

To form and search TABs, a series of transposes are required (Fig. B.4). The first is  $T_{\text{array}} = T_{\text{dish}} + T_{\text{pol}}$ . The wiring of the  $N_{\text{link}} = 384$  10G input links to the TABF implements the first stage of this transpose, by bringing the data from  $N_{\text{ip}} = 24$  telescope paths per band to a single UniBoard. In the second part, these paths need to come together in a single FPGA. This is achieved through the UniBoard mesh (Appendix B.2). Of incoming beamlets, 1/8th stay at the receiving FPGA, the other 7/8th are redistributed, to complete the transpose  $T_{\text{array}}$ .

The 64 subchannels are created using a PFB with an 8-tap finite impulse response (FIR) pre-filter. On these subchannels, the coherent beamformer creates multiple intermediate single-polarisation voltage TABs. The phases of the 8-bit TABF weights  $w_{\text{TAB}}$  depend on pointing nor frequency, but only on dish index  $d$  and TAB index  $n$  (see Appendix D). They are simply  $\phi_{\text{TAB}} = 2\pi d \frac{n}{n_{\text{TAB}}}$ , where  $n_{\text{TAB}}$  is the total number of TABs. For the power TAB, the intermediate voltage TABs are converted into full Stokes power data. Groups of 16 subchannels are integrated to deliver a  $81.92 \mu\text{s}$  sample time, providing more than 10 time bins over an average FRB pulse. The output data consist of separate I and IQUV data streams. The available FPGA hardware allows this design to form up to 12 TABs (see Appendix A.5 for the detailed usage).

Each FPGA in the ARTS TABF uses one 10GbE output link, so in total there are  $M_{\text{PN}} = 128$  output links (Fig. 5). Beamlets are mapped such that 5 CB directions end up at an FPGA. The GPU cluster consist of 40 workstations, one workstation per CB. Each FPGA thus only needs to output to 5 workstations in the GPU cluster. The rate per workstation is 1.8 Gbps the Stokes I data, plus 7.2 Gbps for the Stokes IQUV data (cf. Table A.1).

The design can also operate to produce incoherent-array beams (cf. Appendix C.1.4).

#### 4.1.2. Multiple TABs on UniBoard<sup>2</sup>

The CHAN320 input data for the commensal TABF have already been sub-channelised in Apertif X (Fig. 4), and offloaded by its VHDL firmware. One commensal UNB2 FPGA TABF design essentially includes 8 ported UNB FPGA designs, processing the same as one entire UNB, without requiring a mesh interconnect. The voltage TAB360 processing and other functions are similar to the dedicated search (Sect. 4.1.1). Background on the design can be found in Appendix C.1.5. When forming 9 TABs the design uses about 90% of the multiplier elements in the FPGA.

#### 4.1.3. Single TAB

The TAB1 TABF for pulsar timing or VLBI (Fig. 4) is a single coherent beamformer (Appendix C.1.3) that operates directly on the beamlet voltage samples. The fringe stopping applied in the Apertif FEBF means TAB1 weights can be held fixed. The output TAB1 is sent over 1GbE to the ARTS-0 machine. That machine can interface the voltage TAB to a VLBI system (not in place at the time of writing) or coherently fold and dedisperse it for pulsar timing (Sect. 5.2).

#### 4.2. Software: monitoring, control, operations

Control and monitoring of the overall time domain system is split over several controllers, each described in the following

subsections. In general, WSRT and Apertif are controlled by the monitoring and control system (MAC; Sect. 4.2.1). Observations are scheduled through the Apertif task database (ATDB; Sect. 4.2.2). Monitoring of the system health and component status is done through ARTAMIS (Sect. 4.2.4). Finally, the output data are moved into the Apertif long-term archive (ALTA; Sect. 4.2.5). These all interact with the ARTS-specific monitoring and control system (Sect. 4.2.3). This interaction takes place through a message bus system. The Apertif software is further detailed in van Cappellen et al. (2022) and Appendix C.2; the ARTS software in Oostrum (2020b).

#### 4.2.1. Monitoring and control (MAC)

The Apertif MAC (Loose 2016, and Appendix C.2) comprises relatively independent software components that communicate through a message bus (cf. Appendix C.2.2) and are divided over three layers: drivers, controllers and orchestration (Appendix C.2.3). The drivers handle the I/O to all hardware components. The high level controllers each cover one specific system function, such as control of a dish, the correlator, or the beamformer. Each controller typically interacts with several drivers. The orchestration layer sets different components through the entire system.

#### 4.2.2. Apertif task database (ATDB)

The Apertif task database (ATDB; Vermaas et al. 2019) workflow guides Apertif observations from specification to delivery into ALTA. ATDB is implemented as a microservices architecture. The central database is wrapped inside a Django web application and can be accessed by a range of (Python) ATDB services through a ReSTful API. Each of the ATDB services in the automated workflow has a specific task and is triggered by a specific observation status. Together, this orchestra of services forms a very flexible, adaptable and lightweight workflow.

#### 4.2.3. ARTS drivers and controllers

During system setup, a driver sets the static weights for the ARTS TABF. The rest of ARTS is configured every time an observation starts, when ATDB sends a parameter-set file (parset, as in LOFAR; Stappers et al. 2011) over the Qpid messaging bus. Two controllers direct the ARTS GPU cluster (see Fig. C.4 and Oostrum 2020b). ARTSSURVEYCONTROL, on the master node, receives the parset, which lists the CBs that need to be recorded. For each CB, pointing and processing settings are sent to the ARTSSURVEYNODECONTROL controllers on the worker nodes.

Processing on the nodes, as overseen by ARTSSURVEYNODECONTROL, revolves around the data memory buffers. Following ARTS predecessor PuMa-II (Karuppusamy et al. 2008), these buffers use PSRDADA<sup>2</sup> formats and tools. Four buffers are created: two for Stokes-I and two for Stokes-IQUV data; per set, the main buffer continuously holds all incoming data, while the trigger buffer uses `data_dbevent` to store a selected amount when an external trigger is received. Stokes-IQUV triggered data are written to disk using `dada_dbdisk`. Multiple processes read from the main Stokes-I buffer. First, ARTSSURVEYNODECONTROL starts three instances of GPU search pipeline AMBER (Sect. 5.1.1) across three GPUs. Then, it starts two data writers: `dadafilterbank` (for high-resolution filterbank files), and `dadafits` (for archive resolution FITS files). Two

instances of `fill_ringbuffer` fill the main Stokes I and IQUV buffers from the network. Finally, ARTSSURVEYNODECONTROL starts ARTS processing pipeline DARC (Sect. 5.1.2).

When an observation finishes, each instance of ARTSSURVEYNODECONTROL reports back to ATDB which FITS data products were generated for long-term archiving in ALTA.

#### 4.2.4. ARTAMIS

The dish position, subband statistics, LNA power, and operational firmware image are monitored through the all-round telescope array monitoring and information system (ARTAMIS; Schoenmakers et al. 2019). This real-time framework monitors the entire WSRT (and LOFAR), and is based on ETM/Siemens WinCC-OA. Its web-based data are available to all users. Adjustable alarms alert the telescope operators, who have an instantaneous view on the whole system, including trends.

#### 4.2.5. ALTA

The Apertif long-term archive (ALTA<sup>3</sup>; Holties et al. 2019) provides the long-term storage for raw and processed Apertif data. It covers data management, discovery, and access. The system combines 2 PB online disk-based storage, a PostgreSQL server, and a web-based user portal (all hosted at ASTRON), with tape-based scalable storage hosted at SURFsara.

Rule-based data management is provided through iRODS<sup>4</sup>. Like all raw and processed data, the ARTS filterbank files are ingested through an automated workflow. Metadata are incorporated into the database, supporting full data provenance, compliant with virtual observatory (VO) standards.

For public users, the primary entry point to access these ALTA data is through the ASTRON Virtual Observatory services (Sect. 11.2). There, data collections including the one described in this paper are published.

## 5. Description of time domain system: III. Science pipelines

### 5.1. FRB searching

Searching for FRBs is a computationally intensive process. This section presents the parts constituting the ARTS search pipeline.

#### 5.1.1. AMBER

The Apertif Monitor for Bursts Encountered in Real-Time (AMBER; Sclocco et al. 2016) is our real-time software pipeline for single pulse detection. AMBER identifies FRB candidates, that is, dedispersed signals with high signal-to-noise ratio (S/N), in the beamformed data stream (Sect. 4.1.1). AMBER is optimised for execution on highly parallel architectures such as the ARTS GPUs. AMBER is implemented in C++, with GPU kernels in OpenCL. Its design is modular: each component is separate, developed and maintained on its own. This allows developers to add or modify modules, with minimal impact on the overall code base. Moreover, modules can be reused in other projects (e.g. SKA) without a hard dependency on AMBER.

In ARTS, the pipeline sequentially applies all seven AMBER-provided processing stages to the data chunks: (1) radio frequency interference (RFI) mitigation, (2) downsampling,

<sup>3</sup> <https://alta.astron.nl>

<sup>4</sup> <https://irods.org>

<sup>2</sup> <http://psrdada.sourceforge.net/> (van Straten et al. 2021)

(3) dedispersion, (4) integration, (5) signal-to-noise evaluation, (6) candidate selection, and (7) clustering. We briefly describe each stage below.

*RFI mitigation (optional)*. Optional stages can be enabled when starting AMBER. Here, two distinct filters are applied to the input data to identify and remove – in real time – any bright, wide-band, low-DM RFI and bright narrowband RFI, respectively; this module is described in more detail in Sclocco et al. (2020a).

*Downsampling (optional)*. This module reduces the time resolution of the input data. That is particularly useful when searching for transients at high DM. There, the intra-channel smearing already exceeds the original sampling time. At high DMs, AMBER also requires longer-duration chunks of data; at high time resolution these can run into hardware memory limitations.

*Dedispersion*. This module is the base for the searching algorithm. It implements dedispersion using an algorithm derived from the brute force approach. Recognising that, in GPU dedispersion, optimising for memory limitations is more fruitful than for compute bounds, the module is highly tuned for data reuse. It can dedisperse in both a single step, or using a more efficient two-step process including subbanding. During subbanding, the module creates the SBs described in Sect. 2.3.2, as illustrated in Fig. 7. A further overview of the module design and performance is available in Sclocco et al. (2016).

*Integration (optional)*. The candidate S/N is highest when all burst emission is collected in a single time bin. We thus search over a range of sampling times. During integration, the dedispersed time series are downsampled using a user-defined set of trial pulse widths. Each integration trial thus acts as a convolution kernel, smoothing the signal and approaching the maximum intrinsic signal-to-noise ratio (S/N) at the transient pulse width.

*S/N evaluation*. This module computes the S/N of all peaks in the dedispersed, and optionally integrated, time series. Different ways of computing the S/N are implemented, and users can select the method most suitable for their data.

*Candidate selection*. Candidates whose S/N exceeds a user-defined threshold are stored and made ready for output.

*Clustering (optional)*. If requested, AMBER will cluster results in the DM and integration (pulse-width) plane. For groups of candidates with different pulse widths, but all occurring at the same time and the same DM, only the highest S/N candidate is preserved. Next, groups of candidates that occur in neighbouring DM trials are compacted, again keeping only the candidate with the highest S/N. For each resulting DM-width cluster the single candidate with the highest S/N is reported; reducing the total number of real-time output candidates. Further per-observation clustering occurs off line (see next subsection).

Eventually, all candidates found in the current input chunk are stored in a text file, and after the output is saved, AMBER continues processing the next chunk of input data.

AMBER is available on GitHub<sup>5</sup>, and portable in multiple ways. AMBER compiles and runs on different hardware platforms as it uses standardised and open languages: C++ and OpenCL. Through its combination of run-time code generation, user configurability, and auto-tuning, AMBER

**Table 3.** DARC modules and tasks, see also Fig. C.4.

Module	Task
<i>Monitoring and control:</i>	
DARCMaster	Manages all other modules
STATUSWEBSITE	Generates a web page with the status of each module
<i>Offline:</i>	
OFFLINEPROCESSING	Full offline processing pipeline
<i>Real-time:</i>	
AMBERLISTENER	Reads AMBER candidates from disk
AMBERCLUSTERING	Determines for which candidates to trigger a Stokes-IQUV dump and/or LOFAR TBB observation
DADATRIGGER	Executes Stokes-IQUV triggers through PSRDADA
LOFARTRIGGER	Executes LOFAR triggers
VOEVENTGENERATOR	Sends VOEvents to outside world

furthermore provides performance portability: it can be automatically adapted to achieve high performance for different hardware platforms (see e.g. Mikhailov & Sclocco 2018), observational parameters and searching strategies.

### 5.1.2. DARC

In real time, AMBER delivers metadata of FRB candidates. Based on these, the DARC<sup>6</sup> pipeline (the data analysis of real-time candidates; Oostrum 2021c) decides whether to store the Stokes-IQUV data buffer to disk, and/or trigger a LOFAR observation. DARC also orchestrates the automated offline processing by the deep neural network (Sect. 5.1.3) of the best FRB candidates. Oostrum (2020b) describes DARC in more detail.

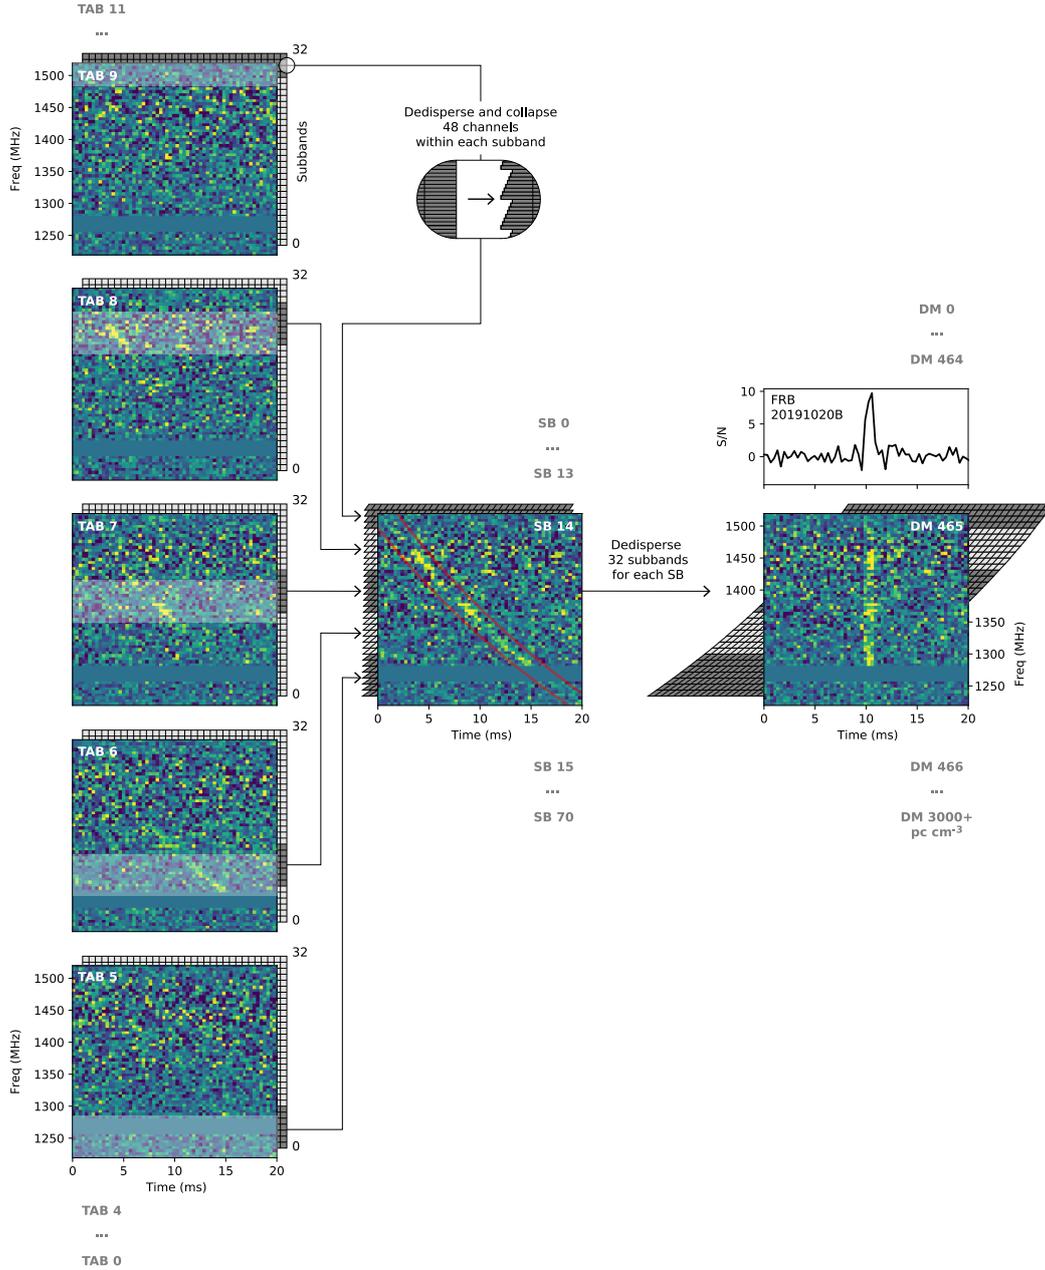
*Global design*. To ease development and improve readability, DARC, like AMBER, is split into several modules (Table 3). These fall in three categories: monitoring and control, offline processing, and real-time processing. Real-time processing is only active for the current observation, while offline processing may be running for several observations simultaneously. An overview of the DARC modules and how DARC connects to the ARTS MAC (Sect. 4.2.3) is given in Fig. C.4.

*Monitoring and control*. Each node of the ARTS GPU cluster runs an instance of DARC. The DARCMaster module receives the control signals from the ARTS MAC and manages all other DARC modules. On the master node, the STATUSWEBSITE module monitors the nodes, for a local status web page.

*Real-time analysis*. DARC runs real-time processing for every observation. Every second, AMBERLISTENER reads in the candidates from each AMBER instance. Each candidate has a beam number, downsampling factor, arrival time, width, DM, and S/N. A single transient may be detected in multiple points in this 6-dimensional parameter space. AMBER already clusters candidates that are close in time, DM, and downsampling factor. The AMBERCLUSTERING module clusters the candidates further over the DM/arrival time plane, and across all SBs of the CB.

<sup>5</sup> <https://github.com/TRASAL/AMBER>; Apache License 2.0.

<sup>6</sup> <https://github.com/loostrum/darc>; Apache License 2.0.



**Fig. 7.** Formation of the SBs during subband dedispersion demonstrated through the detection of FRB 20191020B. Combinations of 12 TABs (left column) produce 71 SBs. The TABs are combined in units of subbands, each consisting of 48 channels. Each subband is first dedispersed to a coarse DM. For SB 14 the resulting intermediate time-frequency plot is shown in the centre. We show this step here for clarity but in the production implementation it is optimised out. For ease of visibility the FRB is shown here at a DM that is reduced by a factor 20. In the second step of dedispersion, the subbands are aligned over 32 trial DMs. FRB 20191020B was thus detected as shown on the right.

If the resulting highest  $S/N$  event exceeds the known or new source threshold (Table 4), the next pipeline step, DADATRIGGER, initiates an IQUV dump and/or LOFAR TBB trigger.

The IQUV dump consists of the set of 1.024 s data chunks that encompass the dispersion delay across the band with a minimum of 2 s, plus an empty 2 s to determine noise statistics. Overlapping triggers are merged, eliminating dead time there; but disk-pool write speeds limit triggers to once per minute.

The LOFARTRIGGER module on the master node collects the worker-node inputs, representing all CBs, and selects the highest  $S/N$  trigger to be sent to LOFAR. The VOEVENTGENERATOR module can next format and send an xml VOEvent to the broker, for triggers to the outside world (Sect. 11.1).

*Offline analysis.* After recording finishes, the offline processing is automatically started. OFFLINEPROCESSING performs a deeper search for FRB candidates, and informs the team of the results via email. The module also takes care of system verification and calibration. This processing typically runs while the next observation is being analysed in real time.

The offline processing now clusters *all* candidates from the observation. Any candidates with  $DM < 20 \text{ pc cm}^{-3}$  or  $S/N < 10$  are discarded. The real-time system already compacted candidates in the DM pulse-width plane. Off line, we extend these clusters over the time dimension. Within a 0.5 s time window, and in a number of DM windows that grow as the candidate DM increases, we keep only the highest- $S/N$  candidate. For these

**Table 4.** Thresholds used during candidate clustering.

Trigger	Source type	S/N	DM range ( $\text{pc cm}^{-3}$ )	$D$
IQUV	Known pulsar	>10	$\text{DM}_{\text{src}} \pm 10$	–
	Known FRB	>10	$\text{DM}_{\text{src}} \pm 10$	–
	New source	>10	$> 1.2 \times \text{DM}_{\text{YMW16}}$	100
LOFAR	Known FRB	>12	$\text{DM}_{\text{src}} \pm 10$	–
	New source	>12	$> 2.0 \times \text{DM}_{\text{YMW16}}$	80

**Notes.**  $D$  is the maximum allowed downsampling factor for a candidate. LOFAR triggering is disabled for known pulsars.

remaining candidates, data from the TAB filterbank files are converted into the required SB data using the DARC multi-threaded SBGENERATOR tool, and cleaned of RFI. We then dedisperse to the DM found by AMBER and downsample in time to the pulse width that maximises the S/N. That refines the initial, real-time AMBER S/N estimate, which was based on only 1.024 s of signal and noise. This re-calculation is especially important for wide pulses, detected at the highest downsampling factor of 250. Offline, the method consists of a matched filter that tries many different box-car widths. For all candidates with post-processing  $S/N > 5$ , the dynamic spectra, DM-time arrays, and metadata are saved to an HDF5 file, as input for the machine learning classifier (Sect. 5.1.3).

The resulting candidate plots from this classifier, together with overall statistics, are e-mailed to the astronomer team and made available on a local website for human vetting and follow-up decision making.

In addition to FRB analysis, the OFFLINEPROCESSING module folds test pulsar data with PREPFOLD from PRESTO<sup>7</sup> (Ransom 2011). It also automatically processes drift scan data, which are used for sensitivity measurements (Sect. 7.1).

### 5.1.3. Machine learning classifier

Due to the real-time nature of the ARTS pipeline and the large number of false positives relative to true astrophysical transients, our candidate classification had to be automated. To this end, we built a binary classifier using deep neural networks to select true FRBs and discard false positives generated by RFI and noise fluctuations (Connor & van Leeuwen 2018). The publicly-available package is called `single_pulse_ml`<sup>8</sup> and uses Keras (Chollet et al. 2015) with a TensorFlow (Abadi et al. 2016) backend for the construction, training, and execution of its convolutional neural networks (Connor 2018). Our machine learning classifier was trained on tens of thousands of false positive triggers from Apertif, as well as an equal number of ‘true positives’ that were generated either by injecting simulated FRBs into real telescope data or by detecting single pulses from Galactic pulsars.

If the classifier-assigned probability that an FRB candidate is a real astrophysical transient exceeds a set threshold (currently 50%), a diagnostic plot is generated showing the frequency-time intensity array (i.e. the dynamic spectrum), a DM-time intensity array, the pulse profile, plus metadata such as the beam number, S/N, classifier probability, and width.

<sup>7</sup> <https://github.com/scottransom/presto>

<sup>8</sup> [https://github.com/liamconnor/single\\_pulse\\_ml](https://github.com/liamconnor/single_pulse_ml)

### 5.1.4. Localisation method

To localise FRBs we make use of the multi-beam information provided by our setup. We create a model of the telescope response on a grid spanned by RA and Dec, and compare this to the measured S/N.

The beam model is constructed from the hierarchical beam-forming techniques described in Sect. 2.3.2 and Appendix D, and is detailed in Oostrum (2020b). First, we create a model of the CBs, based on drift-scan data from both imaging and time-domain drift-scan data. The performance of this model is described in Sect. 7.2. For each CB, we then simulate the tied-array beamformer, followed by the construction of the TABs and SBs. This final beam model predicts the relative sensitivity of the different SBs (see Sect. 7.3).

To localise a burst, we next compare its detection footprint against this model. We determine the detection S/N of all SBs covering the CB the burst was found in, and of all neighbouring CBs. A beam model is next generated on a  $40' \times 40'$  grid with a resolution of  $2''$ , centred on the highest-S/N CB. The predicted relative S/Ns over all SBs are compared to the measured S/Ns, through a  $\chi^2$  method. The final 90% confidence localisation region is derived from the  $\Delta\chi^2$  values. As the region is generally close to elliptical, we often refer to it below as an ellipse. The most likely localisation position is generally on the semi-major axis, but often away from the centre. The localisation precision is evaluated in Sect. 7.4.

### 5.1.5. Polarisation calibration method

Although a Jones matrix was numerically determined for a model of the entire receiver chain of the Apertif prototype system (Wijnholds et al. 2012), the polarimetric quality of the final system is known from on-sky calibration.

Both in imaging and in time domain, this calibration uses an unpolarised source to determine the relative gain of each polarisation, and remove the leakage, and a known linearly polarised source to correct for the relative phase, and calibrate the polarisation angle (PA).

In the raw autocorrelation CB data, the off-axis leakage is already  $<1\%$  (Dénes et al. 2022). After calibration, the leakage is also found to be  $<1\%$  over almost the entire imaging primary beam (Adebahr et al. 2022). The on-axis leakage is even lower. Tests in time domain find this same purity, as we blindly reproduce the linear polarisation fraction of 3C286 to 1% accuracy. Such precision allows for confidence in interpreting IQUV profiles (Pastor-Marazuela et al. 2023); the rotation measures we derive are even more robust, as low leakage does not significantly impact the rotation of the PA with  $\lambda^2$  (see, e.g., Connor et al. 2020).

### 5.2. Pulsar timing

ARTS can process a single tied-array beam (Sect. 4.1.3), to perform pulsar timing and to carry out general pulsar studies. During such an observation, the complex sampled voltages are processed in real-time on a dedicated server (Sect. 3.7). This machine, ARTS-0, has two classes of operation: a recording mode and a timing mode. In recording mode, the incoming TAB1 data are written to disk in PSRDADA format (van Straten et al. 2021). For timing, ARTS-0 predicts the pulsar periods and phases for the duration of the observation using tempo2 (Hobbs et al. 2006; Hobbs & Edwards 2012). Using these predictions, 16 instances of `dspsr` fold and coherently dedisperse the data

that come in through the input ringbuffers (see Appendix C.3.1 for details). Each dspsr instance processes a 18.75 MHz sub-band containing 24 channels of 0.78125 MHz (as produced by the FEBF, Sect. 3.1.3). The 16 instances are divided over 2 GPUs, which allows the system to reduce data in real time. Subintegrations of the resulting folded profiles are stored in a PSRCHIVE (Hotan et al. 2004) format filterbank file every 10 s.

## 6. Commissioning results

### 6.1. Pulsar detections

During commissioning, we regularly observed the four pulsars B0329+54, B0531+21, B0950+08, and B1933+16. Single pulses were detectable from each. As further detailed in van Leeuwen et al. (2022b) and Oostrum (2020b) we used these detections and the radiometer equation to produce first, rough estimates of the system-equivalent flux density (SEFD) of the 8-dish system. Scintillation, intrinsic brightness variations, and the Crab Nebula background flux in the PSR B0531+21 TAB affect these measurements. The median of the four derived SEFD values is 200 Jy. PSR B1933+16 is the most stable pulsar in the set, and the SEFD its observations suggest (130 Jy) is expected to be the most accurate. Assuming an aperture efficiency of 70%, a system temperature of 70 K (Oosterloo et al. 2009; van Cappellen et al. 2022), and perfectly coherent beamforming, the theoretical SEFD of eight Apertif dishes is expected to be  $\sim 70$  Jy. The values derived here are a factor 2–3 higher. Fluctuations in pulsar brightness on that level are not unexpected. A more accurate SEFD can be derived using scans of calibrator sources, as discussed in Sect. 7.1.

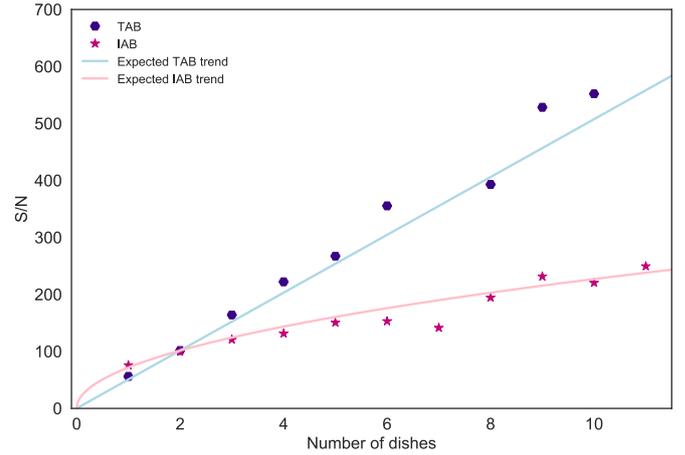
### 6.2. Coherence of the tied-array beams

Observing with an interferometric array is most profitable when full beam-forming coherence is obtained. Then, the attained TAB sensitivity evolves linearly with the number of dishes used. For incoherent addition, on the other hand, the IAB evolution follows a square-root law. We measured the efficacy of our telescope addition by observing pulsar B1933+16 over about two dozen telescope combinations (for details, see Straal 2018), in both modes (Fig. 8). The sensitivity increase is in good agreement with the expected relation  $S/N \propto \sqrt{n_{\text{dish}}}$  for the IAB, and  $S/N \propto n_{\text{dish}}$  for the TAB, demonstrating the coherence of the multidish beam-forming system.

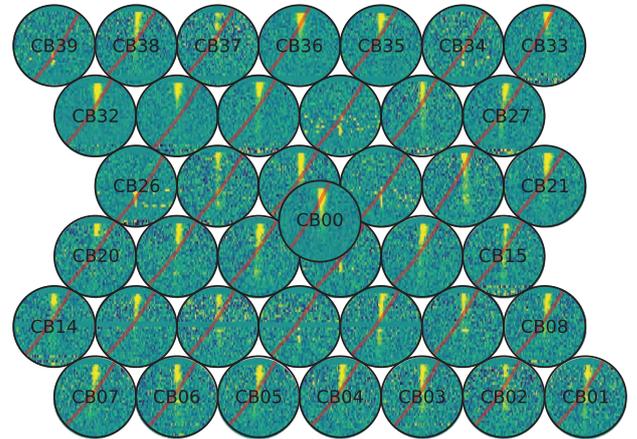
### 6.3. Timing stability

Early commissioning flagged that the synchronisation of the time stamps (the BSN, block sequence number, Appendix C.1.1) to the pulse per second (PPS) signal did not occur completely as expected when the ADCs (analogue-to-digital converters; Appendix B.1) were power cycled. Such power cycling (“cold starts”) severely reduced the effectiveness of the CB calibration solution, as the time and phase information of the CBs jumped. Firmware switches under uninterrupted power (“warm starts”) were found to keep calibration valid, and produce a system at adequate sensitivity and stability for surveys.

Such timestamp jumps, however, could impede pulsar timing with Apertif (Sect. 5.2). Especially for long-term timing, intermediate system power cycling cannot be avoided. We thus compared Apertif observations of a set of precisely timed millisecond pulsars, taken over multiple, monthly epochs with intervening power cycling, against their European Pulsar Timing



**Fig. 8.** Signal-to-noise ratio of PSR B1933+16 versus number of dishes, for coherent and incoherent addition. The two trends are scaled to intersect at the 2-dish S/N value.



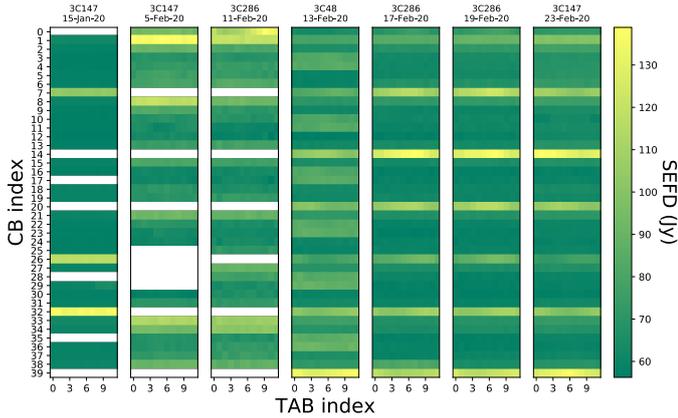
**Fig. 9.** Peryton, detected in all 40 CBs (the circles). In each, the respective time-frequency plot is displayed (similar to, e.g., the bottom panel of Fig. 27), dedispersed to  $395 \text{ pc cm}^{-3}$ . The vertical axis spans 300 MHz of bandwidth, the horizontal axis 500 ms of time. The red curve indicates the track a signal with  $\text{DM} = 0 \text{ pc cm}^{-3}$  would have followed.

Array (Kramer & Champion 2013) ephemeris. We find phase offsets that vary between epochs, of over 1 ms in magnitude. This shows the system currently does not straight-forwardly support long-term timing.

In multiple observations taken of PSR J1022+1001 without power cycling in between, the pulse phase does line up to within a  $\sim 1 \mu\text{s}$ . Short-term timing thus is phase coherent.

### 6.4. Perytons

The ARTS real-time transient detection system was tested using perytons. Perytons are a type of RFI first discovered at Parkes (Petroff et al. 2015b), generated when microwave-oven doors are opened while the cavity magnetron still operates. The frequency structure of the interference strongly resembles a high-DM source. During the Apertif Science Verification Campaign (SVC; Sect. 11.3) we emitted 8 perytons using two different ovens, from outside the WSRT control building. AMBER detected the perytons in real time at a DM of nearly  $400 \text{ pc cm}^{-3}$  (Fig. 9).



**Fig. 10.** SEFD for all TABs of all CBs on seven different days. White regions indicate absence of data. Above each figure, the calibrator source and date are indicated.

## 7. Performance

### 7.1. Compound-beam sensitivity

The ARTS sensitivity depends on the sensitivity of the individual CBs and on the calibration accuracy of the delay and phase offset between the dishes. This CB and array-delay calibration determines if the phasing is coherent, as the beam weights for TABF itself are purely geometric (Sect. 4.1.1). The system is stable for 2–3 weeks (van Cappellen et al. 2022), and is re-calibrated weekly. Following this weekly phase/delay calibration, we use quasars 3C48, 3C147, or 3C286 (Perley & Butler 2017) to determine the sensitivity of the CBs, following a drift-scan method (Oostrum 2020b) summarised below.

For each TAB of each CB, we compare the on-source and off-source values to determine the SEFD from

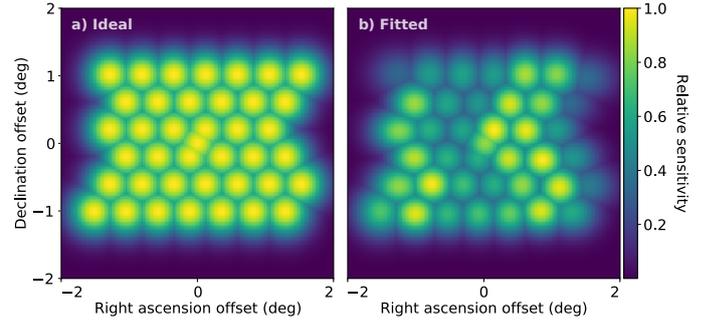
$$\frac{\text{on} - \text{off}}{\text{off}} = \frac{S_{\text{calibrator}}}{\text{SEFD}}, \quad (1)$$

where  $S_{\text{calibrator}}$  is the calibrator flux density. Figure 10 shows the derived SEFD for all 480 TABs, on seven different days. The reported SEFDs are the median values in the frequency band not affected by strong RFI. The TAB sensitivity within a CB is relatively constant. The CBs themselves vary, most likely due to variations in the weekly calibration, which can be affected by RFI or the calibrator hour angle.

The CB sensitivity is expected to depend on its position within the PAF, with central elements illuminated best. Indeed we find that beyond  $\sim 1^\circ$  from the PAF centre, the SEFD increases, by  $\sim 30\%$  for the outermost beams (Fig. 14 of van Leeuwen et al. 2022b; confirmed independently in Fig. 38 of van Cappellen et al. 2022). The median SEFD of the system is 85 Jy.

### 7.2. Field of view

Based on the beamforming models described in Sect. 2.3.2 and Appendix D, we determine the theoretical CB sensitivity pattern. As the WSRT dishes have equatorial mounts, this pattern does not rotate on the sky. As the PAF itself is planar, the CB offsets must first be transformed onto the spherical (RA, Dec) coordinates using a Gnomonic projection (see Oostrum 2020b). The resulting theoretical model at 1370 MHz of the 40 CBs is as shown in Fig. 11a. This assumes each CB reaches the same peak sensitivity. In reality, both the CB sensitivity and shape depend on their position within the PAF, and change over time as the



**Fig. 11.** Model of the 40 CBs of Apertif at 1370 MHz using (a) rotationally symmetric Gaussian beams with equal peak sensitivity and (b) 2D Gaussian fits to each beam, scaled to their measured peak sensitivity in November 2019.

system is recalibrated weekly (Sect. 7.1). Using drift scans in the image plane, the shape of the CBs was measured in June 2019 (Dénes et al. 2022; Kutkin et al. 2022). As long as none of the main PAF elements contributing to a CB fail, the shape of the CBs should remain fairly constant. A 2D Gaussian was then fit to each CB averaged over all dishes. The size of the CBs is assumed to scale linearly with wavelength.

To account for temporal variations in sensitivity (generally caused by RFI during calibration), a more accurate model is constructed for each observing session with an FRB detection. There, the peak sensitivity of each CB is determined from the nearest time-domain drift scan; leading to a model of CB sensitivity at each (RA, Dec) point on the localisation grid. The November 2019 model, for example, is shown as Fig. 11b.

From these models, we determine the Apertif FoV, out to the half-power contour, to be 8.2 sq. deg. at 1370 MHz. The FoV of a single CB is  $\sim 0.3$  sq. deg. at 1370 MHz, so Apertif increased the WSRT FoV by a factor  $\sim 30$ . At higher frequencies, the relative FoV increase is even higher.

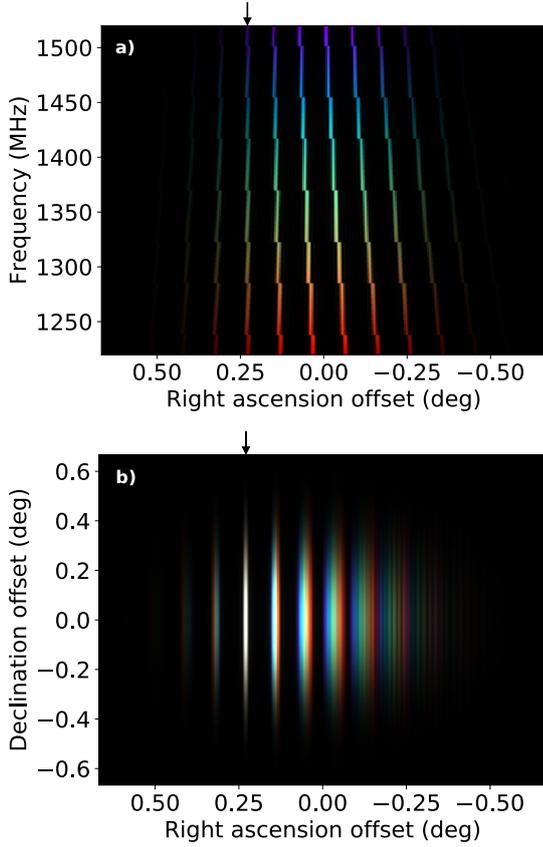
### 7.3. Tied-array beamforming performance

The FoV spanned by the CBs is tiled out with TABs and SBs (Sect. 2.3.2). To find the relative sensitivity of the different SBs, we simulate the ARTS beamformer (Sect. 3.3), recreating its pointing through the application of the geometric phase offset of the CB centre plus the additional TAB offsets (as detailed in Oostrum 2020b). We take the projected baseline length and orientation into account. The sensitive TAB sidelobes are attenuated by the CB response visible in Fig. 11. We finally simulate the mapping of the TABs to SBs, where for each of 32 subbands the most sensitive TAB is used given the required SB pointing (Appendix D). Only the main SB beam contains signal of the full frequency range, whereas the SB sidelobes do not. This is best illustrated for a beam on the edge of the CB, as shown in Fig. 12. For a main beam response, see Fig. 18.

Based on the drift scans of the calibrator quasars through the TABs of each CB (Sect. 7.1), we determine the TAB beamforming efficiency. Combining the aforementioned SEFD of 85 Jy of the CB plus TAB system with the 75% aperture efficiency and the  $\sim 70$  K system temperature leads to a beamforming efficiency  $\beta$  of 0.8. The deviation from 1.0 is likely the result of RFI, especially through its effects on the weekly CB calibration.

### 7.4. Localisation accuracy

To validate the accuracy of the localisation described in Sect. 5.1.4, we observed a field that includes both the Crab pulsar



**Fig. 12.** Model of outermost synthesized beam 70: in (a), as a function of RA and frequency. The intended main beam is indicated with an arrow. The discontinuities in the sidelobes occur when the beamforming switches to a different TAB. In (b), as function of RA and Dec. The colours indicate frequency, where white means an SB is sensitive over the full frequency range. Only the main beam in (b) is broadband, the sidelobes are not.

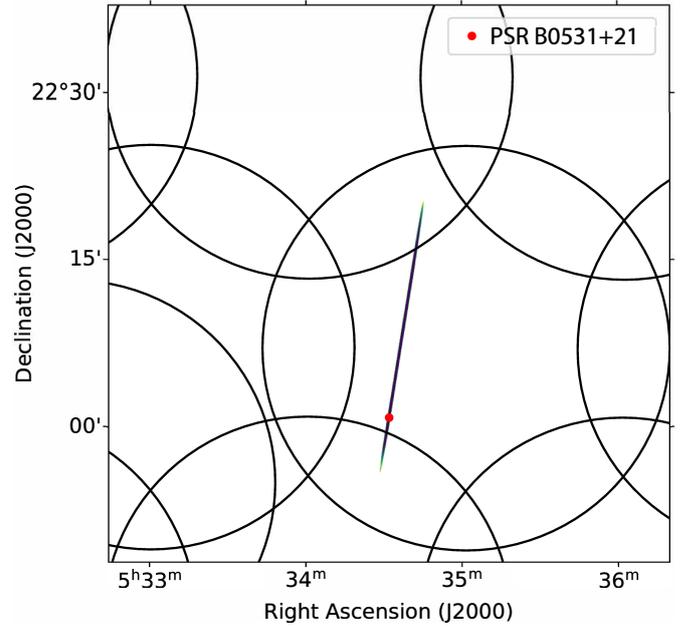
(PSR B0531+21) and PSR B0525+21. An example of the localisation of the Crab pulsar, based on the non-/detection footprint over the CBs and TABs of a single pulse, is shown in Fig. 13. The real position of the pulsar falls within the final 90% localisation region. As the likeliness is quite flat over this region, the fact that the source is away from the centre of the region does not signify any systematic offsets. Similar validation on PSR B0525+21 corroborates this outcome (Oostrum 2020b).

For sources that repeat, multiple burst at different hour angles next intersect to further pinpoint the source (Sect. 2.2.1, Fig. 3). Two or three bursts generally suffice. In Fig. 14 we demonstrate how the independent localisation of three separate bursts of FRB 20180916B (Pastor-Marazuela et al. 2021) combines to a final 90%-confidence-level localisation area of  $26'' \times 9''$ , in which the true position is contained.

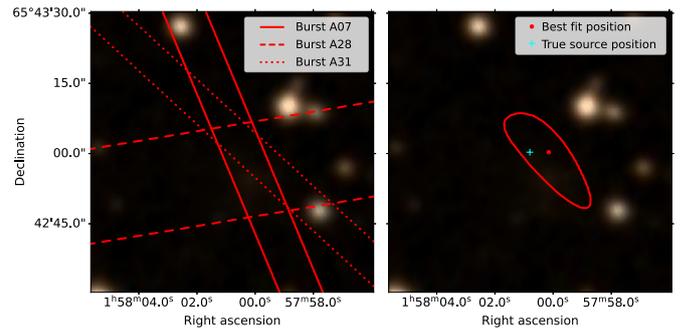
### 7.5. RFI identification and mitigation

WSRT is located in a radio-quiet zone and its RFI environment is good overall. A number of known RFI sources do, however, remain (as listed in Oostrum 2020b). These are generally stronger than astrophysical signals and can cause false-positive detections (i.e. non-astrophysical pulses erroneously classified as FRBs).

False positives hamper the search pipeline, rapidly increasing the size of the single-pulse candidates list: requiring both



**Fig. 13.** Localisation of the Crab pulsar using a single pulse. The black circles indicate the size of the CBs at 1370 MHz. The true location falls well within the 90% confidence ellipse that is displayed.

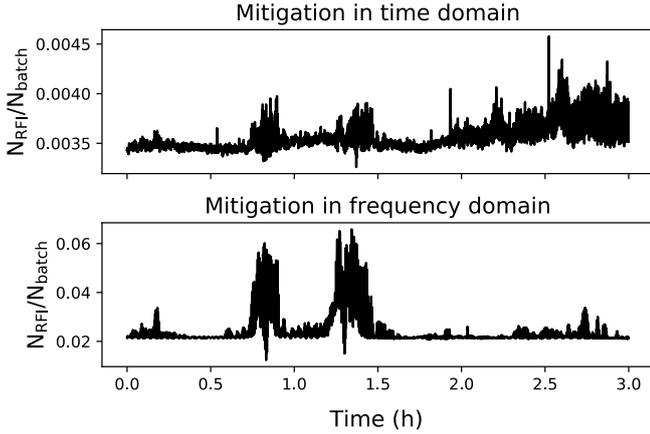


**Fig. 14.** Localisation of FRB 20180916B from three individual bursts. Left: a zoom-in on the overlap of the localisation regions for bursts A07, A28 and A31 (labels as in Pastor-Marazuela et al. 2021). Right: the combined, final localisation region.

more processing and more human verification. To reduce the impact of RFI, two mitigation strategies are currently used in ARTS: on-line RFI removal using local statistics (Sect. 5.1.1 and Sclocco et al. 2020a), and an off-line deep-learning classifier (Sect. 5.1.3).

#### 7.5.1. Direction dependence of RFI

We assessed the variability of RFI against sky direction in the AMBER results of 448 observations taken March-Sept 2019, when online RFI mitigation was not implemented yet in AMBER. Any trigger with  $DM=0$  and  $S/N > 10$  was assumed to be RFI. On a sky map of the resulting RFI trigger rates (Fig. 19 of van Leeuwen et al. 2022b), a peak at  $az=260^\circ$  and  $alt=65^\circ$  stands out from the otherwise uniform RFI distribution on the sky. The Smilde radio mast is located at this azimuth and its UHF antenna transmits in the Apertif frequency band. We likely observe a reflection on the troposphere. Mast emission that is reflected once, halfway between Smilde and WSRT, at a typical troposphere height of 17 km, is expected at an apparent altitude is  $68^\circ$ , which matches our observations.



**Fig. 15.** Fraction of samples cleaned of RFI by the two mitigation strategies, throughout a standard 3-h observing session.

### 7.5.2. Fraction of data lost to RFI

The on-line RFI mitigation in AMBER (Sect. 5.1.1) employs two methods: mitigation in time domain targets bright low-DM broadband signals, while the frequency-domain method removes spurious narrowband RFI. Both methods are applied in an iterative manner, where each consecutive step applies mitigation to an ever-cleaner set of samples. Three iterations of each method are applied. We evaluated the overall impact of RFI on our time and frequency samples, in a representative 3-h observation. We find that at most about 7% of the total bandwidth is affected by RFI for limited time periods ( $< 0.5$  h). In general, less than 3% of data is lost to RFI. Figure 15 highlights that the majority of RFI sources in our band consist of spurious narrow-band emission.

### 7.6. Observing efficiency

We express the end-to-end efficiency of the ARTS system by the fraction of on-source time for which valid data reaches the archive (Sect. 11.2). Such an outcome is prevented by, for example, dead computing nodes and pipeline failures. We compare the successful time against the total expected on-source time in dedicated mode (so, not counting against e.g. slew time, setup time, or calibration time). The ratio is prorated if a subset of beams fails. This fraction is a lower limit to our discovery efficiency though, as detections happen in real time. These can and have occurred when, for example, network issues prevented archiving, while real-time processing continued. For 2019, this efficiency was  $\sim 90\%$ , for 2020 it was  $94\%$ .

### 7.7. Injection pipeline

Throughout commissioning, each step in the ARTS pipeline (Sect. 5.1) was tested via simulated FRB injection tests and single-pulse observations of Galactic pulsars. For our simulated bursts, we developed an injection package, `injectfrb`<sup>9</sup>, as well as evaluation code for assessing completeness<sup>10</sup>. The `injectfrb` package accounts for physically realistic effects such as intra-channel dispersion smearing and sample smearing.

Synthesised FRBs were simulated with parameters drawn from a wide distribution of DMs, widths, scattering timescales, and spectral index. Fluences were drawn from a Euclidean

distribution, in line with current observations. Pulse durations were drawn from a log-Normal distribution with mean 1 ms, allowing us to test our completeness over a large range, for bursts between 0.01–100 ms. Spectral index ranged from  $-4$  to  $+4$ , and DMs were drawn from a uniform distribution between  $50$ – $3000$   $\text{pc cm}^{-3}$ . The bursts were injected off-line into filter-bank data from the ARTS observing system. They thus include the same sky noise, RFI, instrument variability, and other limitations as the live data. The simulated bursts were run through the full pipeline: first searched by AMBER, next following the post-processing steps, such as candidate clustering and machine learning classification. Injection comparison against Heimdall (Barsdell et al. 2012) helped improve, for example, the AMBER real-time S/N computation method (Sect. 5.1.1). Early on, it also helped discover significant incompleteness above  $800$   $\text{pc cm}^{-3}$ , identifying a bug in AMBER that could rapidly be fixed. Afterwards, the sensitivity of the real-life survey was found to correspond to the expected fall off in S/N with DM due to intrachannel dispersion smearing. We find that in the relevant search parameter space (Sect. 8.3), the pipeline is over 90% complete. Selection effects and biases qualitatively become more apparent at the edges of this search space: candidates at low DM are more easily misidentified as RFI while at very high DM they are smeared out. Broadband FRBs could potentially be more easily detected than strongly frequency-limited bursts.

### 7.8. Computing performance

The largest fraction of the computing requirements for ARTS is taken up by the FRB search. On each of the 40 nodes, 3 GPUs (Sect. 4.2.3) run AMBER to carry out this search. These three instances search the low, medium, and high DMs independently. Figure 16 shows the amount of time spent in each AMBER function (Sect. 5.1.1) for a 3-h observation. After the very effective optimisation of, first, the dedispersion kernel (Sclocco et al. 2016), and next, the S/N kernel, most execution time is taken up by input handling and by integration (downsampling).

Due to the low arithmetic intensity, the dedispersion kernel and most other functions are memory bound. Performance is thus best understood by comparing our throughput against the memory bandwidth of the 1080 Ti of  $484$   $\text{GB s}^{-1}$ . The high-DM instance achieves a throughput, averaged over all kernels, close to this value. Given the delays at these high DMs, most data are used only once, meaning there is little opportunity for data re-use in the dedispersion kernel. For the low and medium DM instances, however, such data re-use can be exploited, through the use of user-managed caching, allowing us to process twice as many DMs as the high-DM instance, in the same time. It is therefore by means of optimising the memory accesses that we can process observations in real time.

## 8. Survey and observation strategy

ARTS enables an all sky survey for FRBs and pulsars to pursue the science goals outlined in Sect. 2. This survey, ALERT<sup>11</sup>, the Apertif-LOFAR Exploration of the Radio Transient sky, is described below. Its first results are presented in the next section.

### 8.1. Frequency coverage

During dedicated searching ALERT uses a central frequency (cf. Appendix A.2) of  $1370$  MHz to avoid the RFI present at lower

<sup>9</sup> <https://github.com/liamconnor/injectfrb>

<sup>10</sup> <https://github.com/EYRA-Benchmark/frb-benchmark>

<sup>11</sup> <https://alert.eu/>



**Fig. 16.** Break down of functions for the three AMBER instances that run on each CB/compute node. Shown are the various total execution times per function, that together ran real-time on the 3 h observation.

frequencies. During commensal observing (SC3 in Table 1), the transient search shares the Apertif imaging survey settings. These use a lower central frequency of 1280 MHz, such that redshifted HI is observable throughout a larger cosmic volume. In the period covered here, 2019, commensal mode was not commissioned and all FRB observations used the 1370 MHz band.

### 8.2. Pointings and scheduling

In the 2019 Apertif schedule, time domain observations were allotted a 1/3rd share of the science time. These generally occurred as 2-week blocks every 6 weeks. Pointing definition followed the same field pattern used for the imaging surveys (Hess et al. 2023). To provide a relatively uniform coverage of the northern sky in the planned survey time, we limit ourselves to pointings with declination  $>10^\circ$ ; and to maximise the number of FRB detections we avoid the increased foreground dispersion in fields closer to the Galactic Plane than  $5^\circ$ .

The pointing selection was guided by our triple aims of detecting, localising, and characterising FRBs. We thus target fields that (a) produced earlier ALERT detections, such that in follow-up we can determine if or how these repeat, or (b) contain known repeater FRBs such that these can be better localised and/or studied, or (c) are blank, to search for new FRBs. In the known-source fields (a) and (b), new detections were expected at the same rate as in blank fields, which indeed we found to be the case (Sect. 9). Blank fields are scheduled using Apersched<sup>12</sup> (Hess et al. 2022). Each pointing lasts 3 h. Generally, preference is given to fields at hour angles close to zenith to minimize terrestrial RFI, which is more prevalent near the horizon.

### 8.3. Search parameter space

The survey is optimally sensitive to FRBs characterised by certain sets of parameters. Sources outside this space may still be detectable, but generally at reduced S/N. Our dedispersion plan (Table 5) starts at  $0 \text{ pc cm}^{-3}$ . Sources with DM below  $100 \text{ pc cm}^{-3}$  (similar to e.g. FRB 20200120E; Bhardwaj et al. 2021), however, may easily be misidentified as RFI in the human vetting, as they show little dispersion sweep over our 1.2–1.5 GHz band. That is especially the case for wide bursts. So for

<sup>12</sup> <https://github.com/kmhess/apersched>

**Table 5.** Dedispersion setup for ALERT.

GPU	DM range ( $\text{pc cm}^{-3}$ )	DM step ( $\text{pc cm}^{-3}$ )	Subbanding DM step ( $\text{pc cm}^{-3}$ )
1	0–409.6	0.2	6.4
2	409.6–819.2	0.2	6.4
3	819.2–3379.2	2.5	40.0

**Notes.** The three GPUs cover subsequent parts of DM-trial parameter space. Each uses two-step subband dedispersion (Sect. 5.1.1) with 32 subbands and has optimal dispersion settings.

these low-DM bursts, and for bursts above our maximum DM of  $\sim 3400 \text{ pc cm}^{-3}$ , ALERT is not complete.

AMBER searches through the incoming data at native time resolution of  $81.92 \mu\text{s}$ , and at resolutions downsampled by factors 5, 10, 25, 50, 100 and 250. Any bursts narrower than the original time samples, or wider than the last step, equalling 20 ms, will only be detected at reduced S/N.

## 9. First results

We here present the first set of ARTS results, based on observations up to 2019 Dec 31. We discovered five new FRBs, a significant addition to the  $\sim 100$  published at the time (see Petroff et al. 2016), most of which were only very roughly localised. The cut-off date is aligned with the adjustment of the Apertif time-domain survey strategy at the start of every calendar year. Results from following years will appear in the very near future (e.g., Pastor-Marazuela et al. 2023). A summary of the observed and derived properties of the 2019 FRBs is given in Table 6. It lists both the FRB\_19MMDD format names that we used in earlier publications (Connor et al. 2020; Oostrum et al. 2020), and the TNS-network designations for these bursts, in FRB\_2019MMDD[A-Z] format (Yaron et al. 2020), that we use in this paper. In the following subsections, we discuss the characteristics and localisation of each FRB separately. Where required we describe analysis methods, where they are used for the first time. The analysis of the ensemble is presented in Sect. 10.

### 9.1. FRB 20190709A

#### 9.1.1. Detection

FRB 20190709A, the first ALERT FRB, was discovered in first week of the survey. It was detected in a 3C48 calibration drift scan observation. Its fluence of  $7 \text{ Jy ms}$  puts it at about the 50th percentile of the known population at the time (cf. Fig. 29).

The FRB fluences we report are calculated from the modified radiometer equation (e.g. Cordes & McLaughlin 2003) using Eq. (2) of van Leeuwen et al. (2022b), the SEFD from the weekly calibration (per CB; Sect. 7.2) and the relative sensitivity of the SB the FRB was found in Sect. 7.3, at the best-fit position.

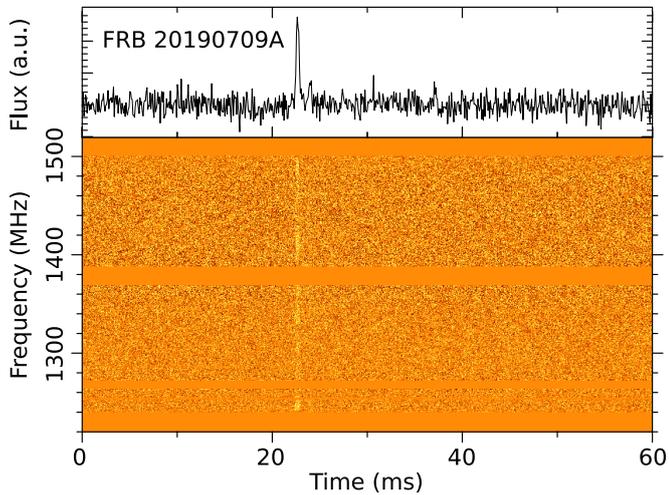
At  $663.1 \text{ pc cm}^{-3}$ , the DM of FRB 20190709A far exceeds the Galactic electron density predicted in this direction, in either NE2001 (Cordes & Lazio 2002) or YMW16 (Yao et al. 2017), as listed in Table 6. As in Connor et al. (2020) we roughly relate the  $\text{DM}_{\text{IGM}}$  contributions, for this and subsequent detections, to redshift  $z$  as

$$\frac{\text{DM}_{\text{IGM}}}{\text{pc cm}^{-3}} \approx 930 z. \quad (2)$$

**Table 6.** ALERT FRBs up to 2019 Dec 31.

FRB	Internal name	MJD	S/N	DM (pc cm <sup>-3</sup> )	DM <sub>MW</sub> (pc cm <sup>-3</sup> )	$z_{\max}$	Width (ms)	Fluence (Jy ms)	RA (J2000)	Dec (J2000)
20190709A	190709	58673.21792057	26	663.1±0.1	52 / 45	0.65	0.49±0.05	7.0±1.4	01:39:19s	+32:03:13
	190903 <sup>(*)</sup>	58729.02228880	8	664±10	53 / 46	0.65	23±2	98±20	01:32:47	+33:03:43
20190926B	190925	58752.03093855	15	957.3±0.5	51 / 44	0.97	2.2±0.2	7.6±1.5	01:42:06	+30:58:05
20191020B	191020	58776.78080366	17	465.0±0.2	102 / 101	0.38	1.04±0.10	9.1±1.8	20:30:39	+62:17:43
20191108A <sup>(**)</sup>	191108	58795.83082563	103	588.1±0.1	43 / 52	0.54	0.34±0.03	8.2±1.6	01:33:47	+31:51:30
20191109A	191109	58796.54885262	22	531.2±0.1	108 / 108	0.44	0.48±0.05	4.5±0.9	20:35:15	+61:49:02

**Notes.** First column lists TNS name, second lists previously used name (cf. Oostrum et al. 2020; Connor et al. 2020; Oostrum 2020b). MJD refers to the arrival time at the solar system barycentre at infinite frequency. DMs and widths were measured with PDMP from PSRCHIVE (Hotan et al. 2004). DM<sub>MW</sub> are the Milky Way DMs predicted using the NE2001 (Cordes & Lazio 2002) and YMW16 (Yao et al. 2017) models, respectively. We assume a 10% uncertainty on the observed pulse widths and 20% on the derived fluences. The position error regions are strongly elongated and presented individually. The best position is not always at the centre of the elliptical region. The derivation of the redshift upper limit  $z_{\max}$  is described in Sect. 10.4. <sup>(\*)</sup>Candidate <sup>(\*\*)</sup>The arrival time published in Connor et al. (2020) contained an error in the barycentric correction. In this table that is rectified.


**Fig. 17.** Dynamic spectrum (bottom) and pulse profile (top; flux density scale in relative, arbitrary units) of FRB 20190709A.

Assuming no DM contribution from the host galaxy then gives a redshift upper limit for this FRB of  $z = 0.65$ .

The FRB exhibits a strong,  $\sim 400 \mu\text{s}$  narrow component and potentially a second equally narrow component (see Fig. 17). These widths are comparable to the intra-channel dispersion smearing, suggesting the components are intrinsically extremely narrow. The separation between the components is about 1.3 ms. At  $z = 0.65$  the actual separation at the emission site would be only about 0.8 ms.

FRB 20190709A was detected before IQUV triggering (Sect. 5.1.2) was in production. In the Stokes-I data, the auto-correlation function of the burst spectrum indicates smooth structures with characteristic bandwidth of order 20–30 MHz. The Galactic scintillation bandwidth is expected to be only a few MHz. However, effects of any scintillation caused by the host galaxy medium cannot be ruled out.

### 9.1.2. Localisation

FRB 20190709A was detected in two CBs, across a total of 23 SBs. The derived localisation region is well represented by an ellipse of  $10' \times 39''$ , as shown in Fig. 18. These FRB localisation

regions are derived using statistical errors. The systematic offset that Oostrum (2020b) found for the localisations from drift scan observations (relevant for this FRB 20190709A and for candidate FRB 190903; Sect. 9.2) is solved here.

For all FRBs, we have searched for putative host galaxies using the GLADE galaxy catalogue (Dályá et al. 2018). The catalogue was created by combining previously existing galaxy and quasar catalogues (GWGC, 2MPZ, 2MASS XSC, HyperLEDA and SDSS-DR12Q). Although its primary objective is to identify potential gravitational wave host galaxies, we used this database to look for potential FRB hosts within the error region of ARTS FRBs. The catalogue is complete up to a luminosity distance  $d_L = 37^{+3}_-4$  Mpc and contains all of the brightest galaxies up to  $d_L = 91$  Mpc. While most ARTS FRBs could originate from distant galaxies beyond the completeness limit of GLADE, this catalogue does allow us to check for the presence of possible nearby host galaxies which, if they are the host galaxy, could point to a large amount of DM local to the source.

For FRB 20190709A, the GLADE catalogue contains no galaxies within the error region limits.

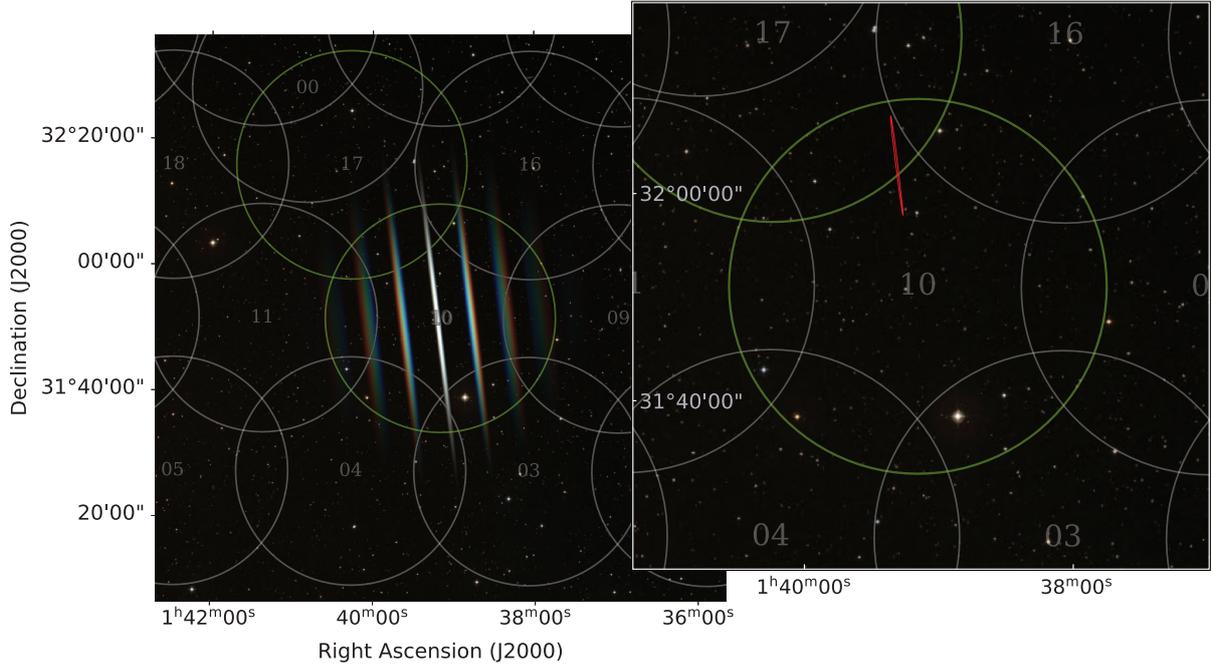
## 9.2. Candidate FRB 190903

### 9.2.1. Detection

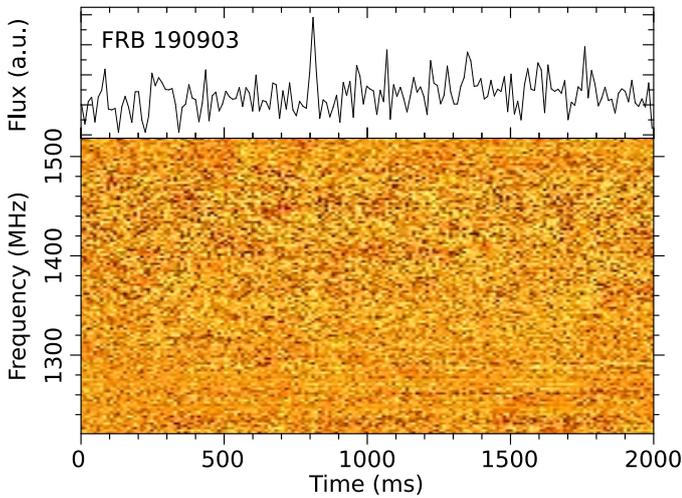
Candidate FRB 190903 (see Fig. 19) was detected just above our detection threshold. Given its low S/N it is of course quite possible that FRB 190903 is a non-astrophysical signal fluctuation. The candidate did, however, survive blind machine-learning and human vetting. Only then was it realised the FRB was from the repeat follow-up field on FRB 20190709A, at the same DM. No further repeat detections at this DM were, however, seen in 120 h of follow-up of the field (see Sect. 10.5). Its width of 23 ms is an outlier compared to the generally  $\sim$ ms widths of the unambiguous detections (Table 6).

### 9.2.2. Localisation

Because of its marginal detection in an outer beam, the localisation region of FRB 190903 is quite large: out to the SB model limit of  $40'$  from the CB centre, the region is roughly an ellipse of  $43' \times 30''$  (Fig. 20). However, we cannot exclude that

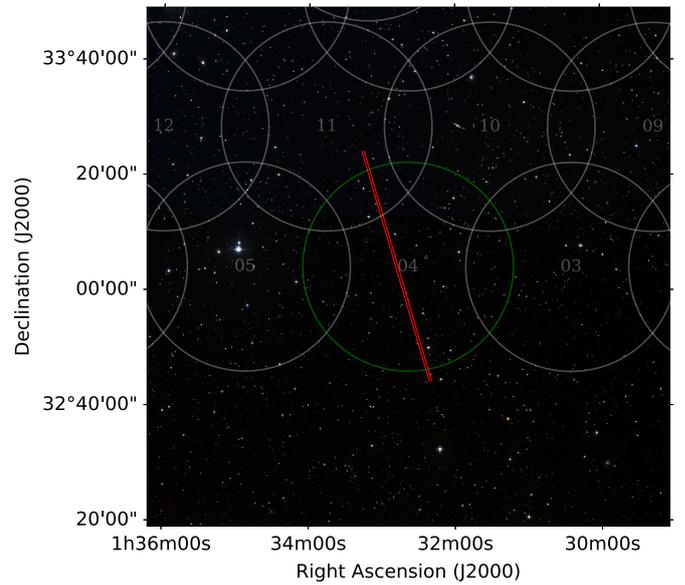


**Fig. 18.** Localisation region of FRB 20190709A. In both panels, and in the localisation region figures that follow, the CBs at 1370 MHz are shown as white (non-detection) and green (detection) circles. In the left-hand side we show in the colour scale the simulated response of the SB in which the FRB was detected most strongly (cf. Fig. 12). On the right hand side, and in the follow figures, the red, elongated and very narrow area indicates the 90% confidence level localisation area that results from combining the SB detections and upper limits of all surrounding CBs. Background image here and in the following figures are from the Sloan Digital Sky Survey (SDSS; York et al. 2000).



**Fig. 19.** Dynamic spectrum (bottom) and pulse profile (top) of FRB 190903.

FRB 190903 originated further away from main beam pattern. Given its similar DM and approximate location to FRB 20190709A, we consider that FRB 190903 might in fact be a repeat pulse of FRB 20190709A. This would place FRB 190903 far outside the main beam pattern, assuming FRB 20190709A did occur within the main beam pattern. As a precise enough model of the CB sidelobes is not currently available, we cannot further constrain the localisation of FRB 190903 within the FRB 20190709A field and thus cannot prove whether or not these two bursts do in fact originate from the same source. There are two galaxies within a  $1'$  distance from the localisation region in CB 04 at a redshift  $z < 0.65$ , one of which is located within the ellipse limits.

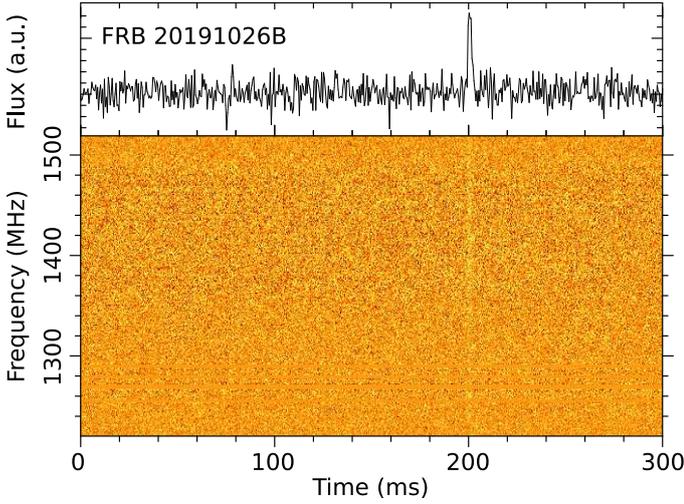


**Fig. 20.** Localisation region of FRB 190903. The circles and red area are as described for Fig. 18. The localisation area is not constrained towards lower declinations.

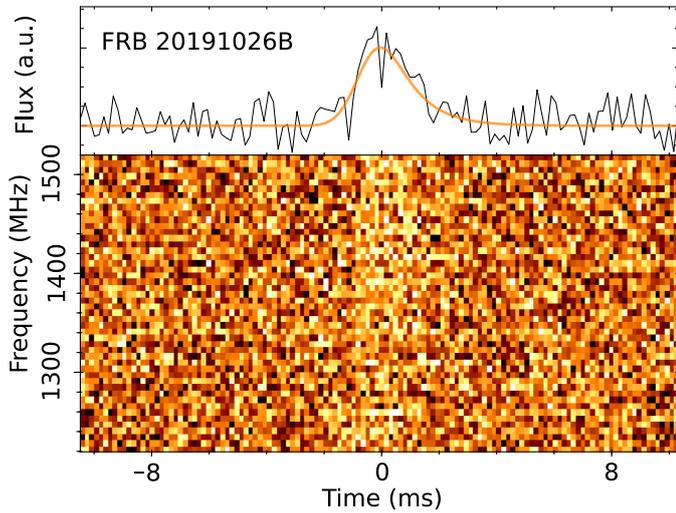
### 9.3. FRB 20190926B

#### 9.3.1. Detection

FRB 20190926B was detected with a single, 2-ms wide component at a DM of  $957.3 \text{ pc cm}^{-3}$  (Fig. 21). As the intra-channel dispersive smearing is 0.6 ms the burst is well resolved. The Galactic scintillation bandwidth on this sight line is expected to be only a few MHz, but we see no evidence of such scintillation in this burst.



**Fig. 21.** Dynamic spectrum (bottom) and pulse profile (top) of FRB 20190926B.

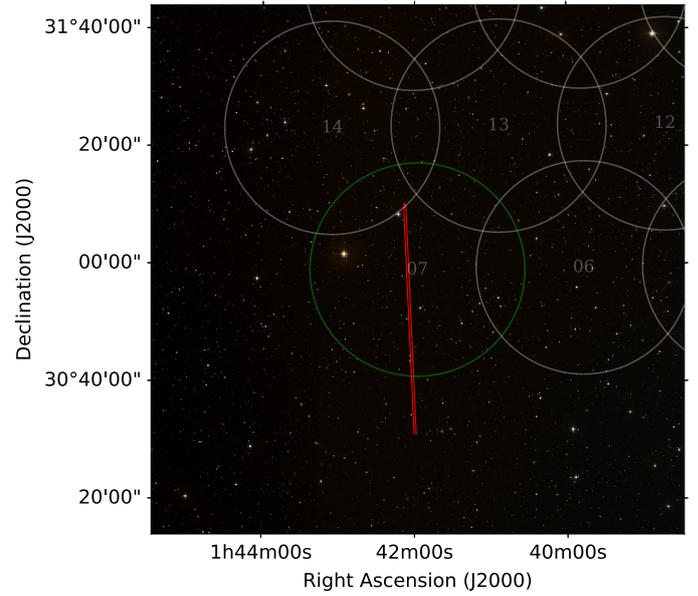


**Fig. 22.** FRB 20190926B at  $15\times$  higher time resolution than Fig. 21. The orange line in the top panel is the best fit when including a scattering tail.

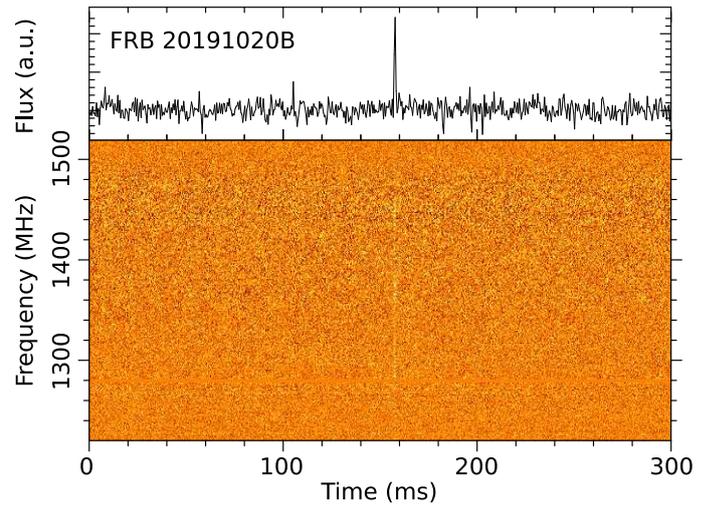
This is the only burst in the sample we present here that potentially shows exponential decay of the intensity, indicative of scattering (Fig. 22). Following the scattering characterisation described in Pastor-Marazuela et al. (2023), we fitted the dedispersed pulse profile to a Gaussian function, with an optional convolution with a decreasing exponential. The Bayesian information criterion values are the same with or without the exponential tail; we thus cannot prove its presence, or rule it out. The best resulting scattering timescale  $\tau_{sc} = 1.0 \pm 0.3$  ms is several orders of magnitude larger than the expected contribution from the Milky Way. If present, it has thus likely been produced at the host galaxy or circumburst environment.

### 9.3.2. Localisation

FRB 20190926B was detected in one CB, across a total of 7 SBs. The derived localisation region is an ellipse of  $39' \times 30''$ , as shown in Fig. 23. Because the FRB was found in a CB on the edge of the beam pattern, the localisation region is not constrained towards lower declinations. However, if the FRB was detected in a sidelobe, it should intrinsically have



**Fig. 23.** Localisation region of FRB 20190926B. The localisation is not constrained towards lower declinations.



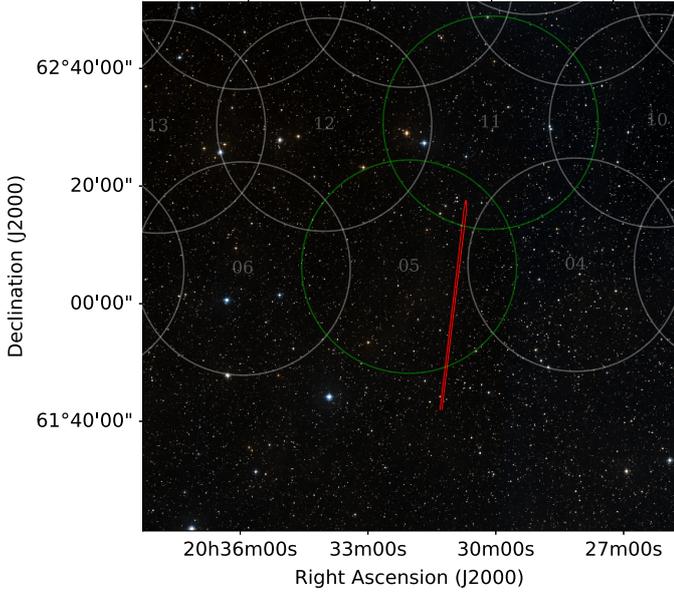
**Fig. 24.** Dynamic spectrum (bottom) and pulse profile (top) of FRB 20191020B.

been extremely bright. We conclude it more likely it occurred within the main beam of CB 07, as indicated by the green circle in the figure. The GLADE catalogue contains one galaxy, SDSS 01424.24+305214.4, redshift 0.68, located at an angular distance of  $\sim 45''$  from the localisation region, within the main beam.

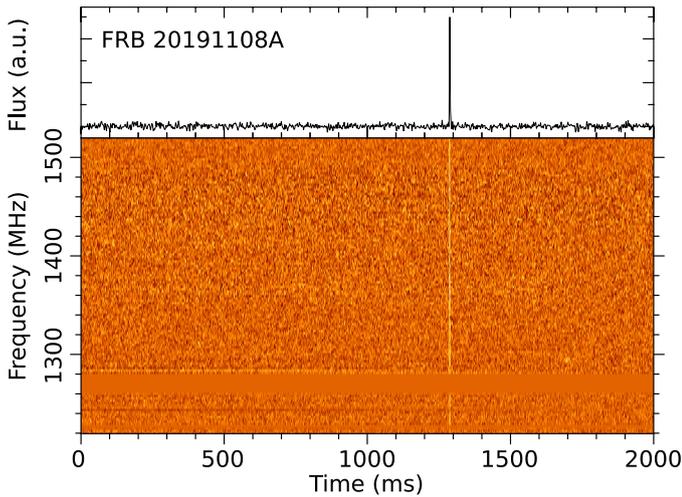
## 9.4. FRB 20191020B

### 9.4.1. Detection

With a width of about 1 ms, FRB 20191020B is also resolved (Fig. 24). The burst is consistent with a single Gaussian component. The burst spectrum is fairly uniform. This detection was the first for which IQUV data saving (Sect. 5.1.2) was successfully triggered in real time. There is, however, no quasar data available in this epoch for polarisation calibration, and the S/N is too low to perform the calibration method on the FRB data itself that was demonstrated in Connor et al. (2020).



**Fig. 25.** Localisation region of FRB 20191020B. Again, the localisation is not constrained towards lower declinations.



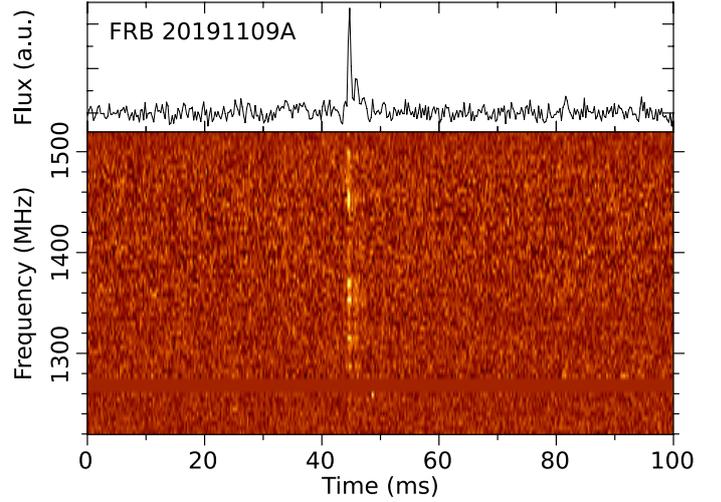
**Fig. 26.** Dynamic spectrum (bottom) and pulse profile (top) of FRB 20191108A.

#### 9.4.2. Localisation

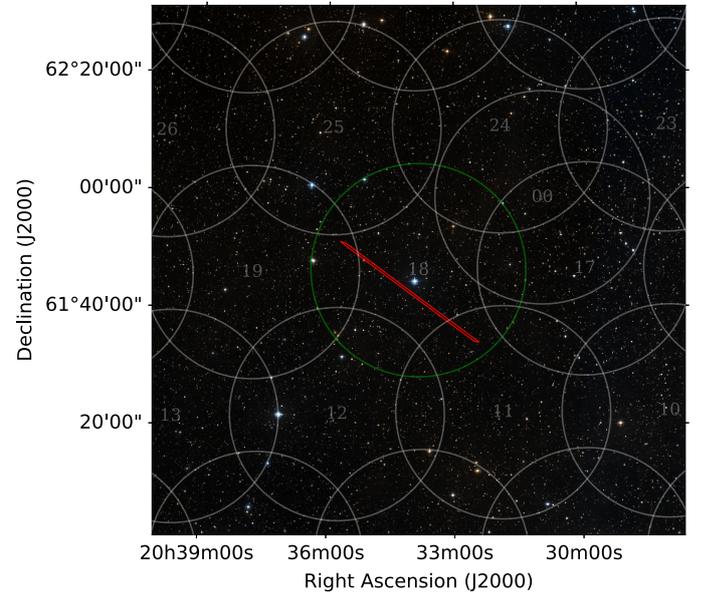
FRB 20191020B was detected in two CBs, across a total of 9 SBs. The derived localisation region is an ellipse of  $37' \times 35''$ , as shown in Fig. 25. Similar to FRB 20190926B, it was found in an outer CB and its localisation region is open-ended towards lower declinations. However, as for FRB 20190926B, we find it more likely that the FRB occurred in the primary beam of CB 05. There are no galaxies in the GLADE catalogue located at less than  $1'$  from the localisation ellipse within CB 05.

#### 9.5. FRB 20191108A

FRB 20191108A (Fig. 26) was detected in 48 SBs spread over three CBs, at a maximum S/N in the real-time pipeline of 60. The discovery DM was  $588 \text{ pc cm}^{-3}$ . Full-stokes data were captured and calibrated. The FRB exhibited a rotation measure of  $+474 \pm 3 \text{ rad m}^{-2}$ , much higher than expected from the Milky Way foreground, and much higher than seen in other sources



**Fig. 27.** Dynamic spectrum (bottom) and pulse profile (top) of FRB 20191109A.



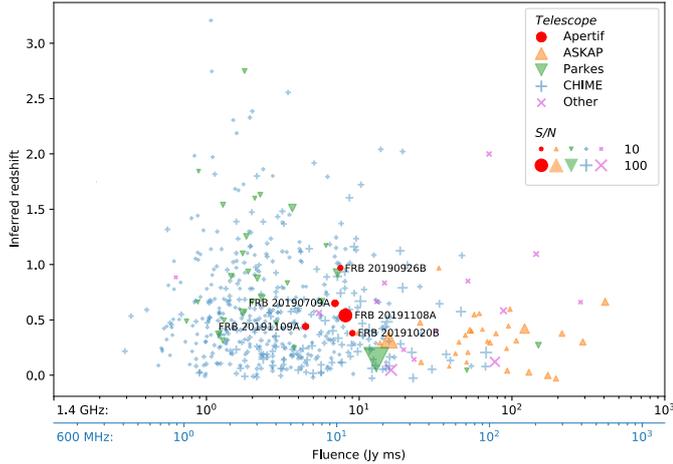
**Fig. 28.** Localisation region of FRB 20191109A. The CBs at 1370 MHz are shown in white (non-detection) and green (detection). The red, elongated and very narrow area indicates the 90% confidence level localisation area.

close to this line of sight. The FRB is localised to a small,  $5''$  by  $7'$  region that is only  $1.20 \pm 0.05^\circ$  away from the core of Local Group galaxy M33. This discovery is described in detail in Connor et al. (2020). The implications of our various FRB lines of sight through the M33 halo are discussed in Sect. 10.5.

#### 9.6. FRB 20191109A

##### 9.6.1. Detection

FRB 20191109A exhibits multiple components – 2 discernible bright components and perhaps a faint third trailing component (Fig. 27). The widths of the individual components are around  $0.6 \text{ ms}$ , and at a DM of  $531 \text{ pc cm}^{-3}$  these are resolved over the intrachannel smearing. The first two components are separated by  $1.2 \text{ ms}$ . A one-off FRB model by Falcke & Rezzolla (2014), which involves collapse of a supermassive star into a



**Fig. 29.** Redshift and fluence for the FRBs presented here compared against those in the TNS (<http://www.wis-tns.org>) also detected in or before 2019. In this semi-log plot, a spectral index of  $-0.4$  is used to compare the CHIME bursts (blue abscissa) with the other bursts (black abscissa). The Apertif FRBs are positioned around the midpoint of the 1.4 GHz FRB fluence distribution.

black hole, predicts a leading precursor, a main burst and ring-down within 1 ms or more. While there is no sign of a precursor, the temporal structure of FRB 20191109A might be reminiscent of the above predicted ringdown. This multi-component profile, with its  $\sim 1$ -ms burst separation, could also be the FRB equivalent of quasi-periodic microstructure seen in pulsars. In [Pastor-Marazuela et al. \(2022\)](#) we present additional FRBs with such structure, and provide a more detailed discussion.

FRB 20191109A also exhibits frequency structures with bandwidths of the order of 10–12 MHz. The structures are significantly correlated between the first two components, which means they are either intrinsic or are caused by scintillation affecting a broadband intrinsic pulse. The NE2001 and YMW16 electron density models both predict Galactic scintillation bandwidths to be much less than 1 MHz. Hence, the observed frequency structures are likely either intrinsic to the emission mechanism, or caused by the medium in the host galaxy.

Several repeating FRBs are known to exhibit frequency structures which drift downwards with time. By cross-correlating the spectra at the peaks of the first two components, we find an offset of about  $0.6 \pm 0.6$  MHz between the frequency structures under the corresponding components, consistent with zero. For a non-zero offset, the best-fit value would imply a subburst drift to lower frequencies with time at a rate of  $\sim 0.5$  MHz  $\text{ms}^{-1}$ , much lower than that observed for the repeating FRB 121102 ( $\sim 200$  MHz  $\text{ms}^{-1}$ ; [Hessels et al. 2019](#)).

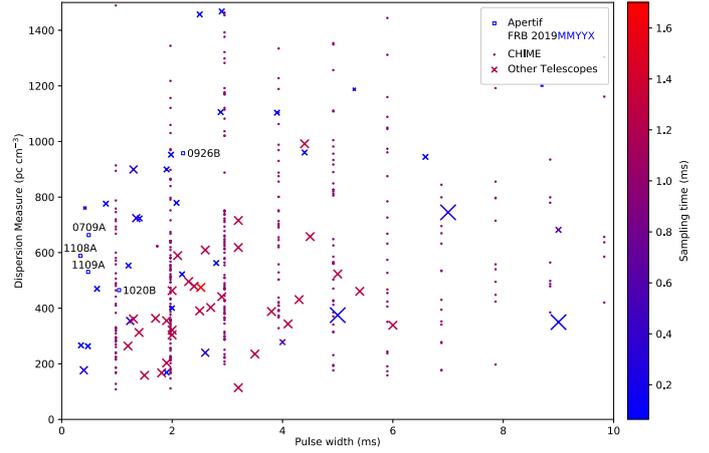
## 9.6.2. Localisation

FRB 20191109A was detected in one CB, across a total of 15 SBs. The derived localisation region is an ellipse of  $29' \times 5''$ , as shown in Fig. 28. GLADE does not contain any galaxy located at less than  $1'$  from the localisation region.

## 10. Discussion

### 10.1. Characteristics of the discovered sample

The fluences of the Apertif-discovered bursts are around the median of the known fluence distribution (Fig. 29).



**Fig. 30.** Pulse width and dispersion measure of the 5 FRBs presented in this work, contrasted against the same bursts as in Fig. 29. The Apertif FRBs are among the narrowest known, and have high dispersion measure. Markers are sized and coloured to reflect the frequency and time resolution, respectively. For Apertif, both are very high.

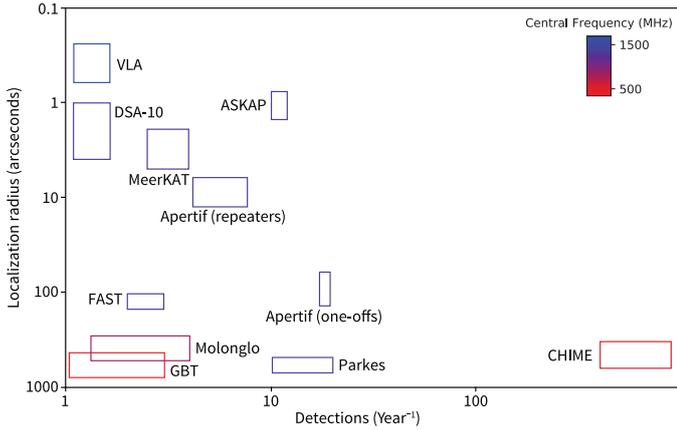
When comparing the 1.4 GHz surveys, Parkes detected bursts that were generally fainter, as its SEFD is lower (30 K; [Manchester et al. 2001](#)). In contrast, only brighter FRBs are generally detected by ASKAP, as that system searches data from incoherently added dishes ([Bannister et al. 2019](#)). As a result, following the relationship between the fluence limits and the volume up to the equivalent redshift limit, the inferred redshifts of our bursts are also higher than the ASKAP bursts, and lower than the Parkes bursts.

Given the high time and frequency resolution of ARTS (detailed below), even our distant bursts suffer less from dispersion-smearing related selection effects than the CHIME bursts (Fig. 29). Assuming a spectral index of  $-0.4$ , the best-fit value for the FRB population per [Gardenier & van Leeuwen \(2021\)](#), to compare the fluences, the CHIME bursts are generally dimmer, yet lower redshift. The Apertif FRB sample is free of selection effects out to larger distances than CHIME ([Wang et al. 2023](#)).

The stand-out characteristic of the Apertif bursts is that they are detected at significantly narrower widths than the FRB population that was known before. Figure 30 shows that 3 of the Apertif FRBs are among the 6 narrowest overall. The reason for this unique trait is two-fold. The observing system has very high native time resolution, and very good frequency resolution, such that instrument smearing is low. For comparison, this time and frequency resolution are  $10\times$  and  $5\times$  higher, respectively, than for the ASKAP incoherent-sum mode. This allows us to determine the second, more fundamentally intrinsic characteristic: the bursts display very little temporal scattering. This means we can discern individual components even if they are only a millisecond apart, as for FRB 20191109A (and as analysed for Apertif FRB 20201020A in [Pastor-Marazuela et al. 2022](#)).

Notable, too, are the generally large dispersion measures. There, the good frequency resolution limits the intra-channel dispersion smearing, in turn allowing short bursts to reach high S/Ns and be detected, whilst these same bursts might be washed out and missed by other telescopes. That brings a larger cosmic volume of bursts within our S/N reach, increasing detections at high DMs.

The 5 bursts reported here were the first discovered with the system, and IQUV polarisation recording and calibration was



**Fig. 31.** Rate per calendar year versus general localisation properties for the main FRB surveys to date. Shown is localisation radius (or equivalent for elongated regions) versus the rate per year given the observing time. Apertif localisation numbers are provided for both single-burst detections (“one-offs”), and for repeat detections at different hour angles (“repeaters”), see Sect. 7.4. Data for other surveys from [Burke Spolaor \(2018\)](#), [Gardenier et al. \(2019\)](#) and [CHIME/FRB Collaboration \(2021\)](#).

only successful for FRB 20191108A. For bursts discovered in 2020 and after, the availability of full-Stokes data is much higher ([Pastor-Marazuela et al. 2023](#)).

Finally, all 5 bursts are broadband, covering the entire 300 MHz Apertif band, even if that coverage is modulated by frequency structure. This spectral modulation and band coverage fraction are qualitatively similar to the burst sample presented for ASKAP ([Shannon et al. 2018](#)). Of the four FRB archetypes proposed in [Pleunis et al. \(2021a\)](#), only two classes are represented in our sample. Morphologically, FRBs 20190926B, 20191020B and 20191108A display straight, Gaussian profiles in time, without much hint of a scattering tail, that are broadband (class I). FRBs 20190709A and 20191109A are spectrally similar but temporarily more complex (class III). We found no bursts that are narrowband (class II) or downward drifting (class IV). In principle the broad band nature of all bursts is encouraging for simultaneous detections of the same bursts with Apertif and LOFAR (Sect. 2.2.5) although the separation of their observing frequencies is admittedly much larger than the bandwidth coverage achieved here.

### 10.2. Survey detection rate and localisation

Our 2019 results show the total number of FRB detections per calendar year is high for ALERT at Apertif. Figure 31 shows this large advance over (earlier) FRB surveys as carried out with GBT, VLA, and Molonglo. The high rate is a result of three main factors: the large time ALERT covered the sky, of about 1/3rd of the available observing hours; the large field of view enabled by the PAFs; and the high sensitivity from the full-field coherent beamforming.

As Apertif also offers 1D interferometric FRB localisation, ALERT can localise bursts significantly better than single-dish surveys do. Figure 31 shows the localisation radius (or its equivalent in case localisation regions are elongated) for the survey comparison. For Apertif we consider both the one-off case ( $\sim 100''$ , Sect. 9), and the case where multiple bursts are detected ( $\sim 10''$ , Sect. 7.4). For the rate of the latter, we assume a detected population comprised of Apertif discoveries that repeat,

plus localisations from other, known, repeater FRBs. Overall, ASKAP, CHIME and Apertif each have their own tradeoff between rate and localisation accuracy, with Apertif providing both.

### 10.3. Detection of new FRBs versus field type

The ALERT pointing strategy (Sect. 8.2) contains three types of fields. The priority order is for fields that contain a) earlier ALERT detections, b) known repeater FRBs, or c) are blank. We find the approach works well in allowing both follow-up study and new detections. The repeat visits to detection fields (a) and the new detection in there, have allowed us to chart multiple lines of sight within a field, as demonstrated in Sect. 10.5. Fields with other known FRBs (b) led to results detailed separately, for example in [Pastor-Marazuela et al. \(2021\)](#) for FRB 20180916B.

### 10.4. Prospects for counterpart identification

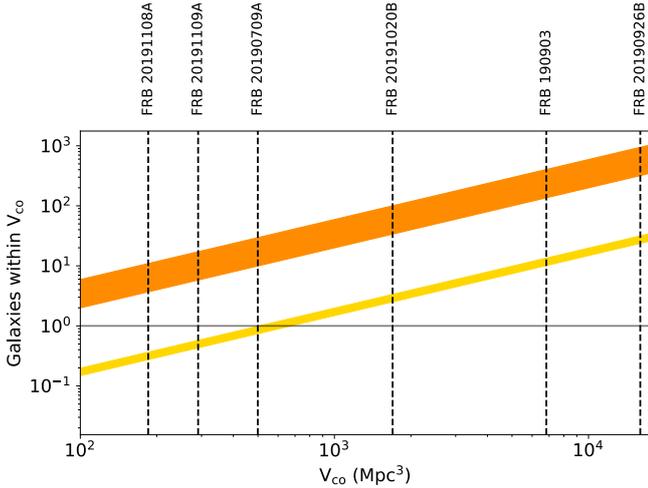
While for some FRBs we have already identified one or more galaxies in their localisation regions, there are probably more, too faint to be included in the GLADE catalogue. Here we estimate the total number of galaxies we might expect in the localisation regions of our FRBs.

The number of potential host galaxies in a localisation region depends strongly on which types of galaxies are considered to possibly host FRB progenitors: Dwarf galaxies are far more common than massive galaxies. As the first repeating source, FRB 121102, was localised to a dwarf galaxy with a high specific star formation rate ([Tendulkar et al. 2017](#)), it was thought that this type of galaxy might be related to the FRB progenitor type. However, other FRBs – both repeating and non-repeating – have now been localised to a variety of galaxies ([Chatterjee et al. 2017](#); [Bannister et al. 2019](#); [Ravi 2019](#); [Prochaska et al. 2019](#); [Marcote et al. 2020](#)). We here apply the same analysis as done for FRB 110124 in [Petroff et al. \(2019b\)](#) and [Oostrum \(2020b\)](#), and estimate the number of dwarf galaxies ( $4 \times 10^7 M_{\odot} < M_{\text{stellar}} < 10^{10} M_{\odot}$ , i.e. between the mass of the host galaxy of FRB 121102 and the maximum dwarf-galaxy mass) and massive galaxies ( $M_{\text{stellar}} > 10^{11} M_{\odot}$ ) in the FRB localisation volumes. From the stellar mass function of [Baldry et al. \(2012\)](#), and from converting the [Haynes et al. \(2011\)](#) HI mass function at ratios between 1–10, we estimate a dwarf galaxy number density of  $n = (0.02\text{--}0.06) \text{Mpc}^{-3}$ . For massive galaxies we use the luminosity function of [Faber et al. \(2007\)](#), and find  $n = (1.5\text{--}2.0) \times 10^{-3} \text{Mpc}^{-3}$ .

The expected number of galaxies in an FRB localisation region is the galaxy number density multiplied by the comoving volume out to the redshift of the FRB, assuming that the mass functions do not evolve significantly up to the maximum redshift ( $z \approx 1$ ) of our FRB sample. The redshift is estimated from the intergalactic medium (IGM) DM contribution ( $\text{DM}_{\text{IGM}}$ ) using Eq. (2).  $\text{DM}_{\text{IGM}}$  can be related to other sources of DM as

$$\text{DM}_{\text{IGM}} = \text{DM} - \text{DM}_{\text{MW}} - \text{DM}_{\text{halo}} - \frac{\text{DM}_{\text{host}}}{1+z}, \quad (3)$$

where  $\text{DM}_{\text{MW}}$ ,  $\text{DM}_{\text{halo}}$ , and  $\text{DM}_{\text{host}}$  are the DM contributions from the Milky Way, its halo, and the host galaxy (which includes the environment local to the source), respectively. For the Milky Way contribution, we take the lowest value predicted by the NE2001 and YMW16 models (see Table 6). Based on the [Yamasaki & Totani \(2020\)](#) model, we conservatively assume  $10 \text{pc cm}^{-3}$  from the Milky Way halo, and we set the host galaxy



**Fig. 32.** Expected number of galaxies in the FRB localisation areas for a range of dwarf galaxy (orange) and massive galaxy (yellow) number densities as function of comoving volume. The comoving volume depends on both the redshift of the FRB and the size of the localisation region on-sky, and should be regarded as an upper limit. The horizontal line indicates where the expected number of galaxies is one.

contribution to zero. The resulting redshift estimates thus are conservative upper limits.

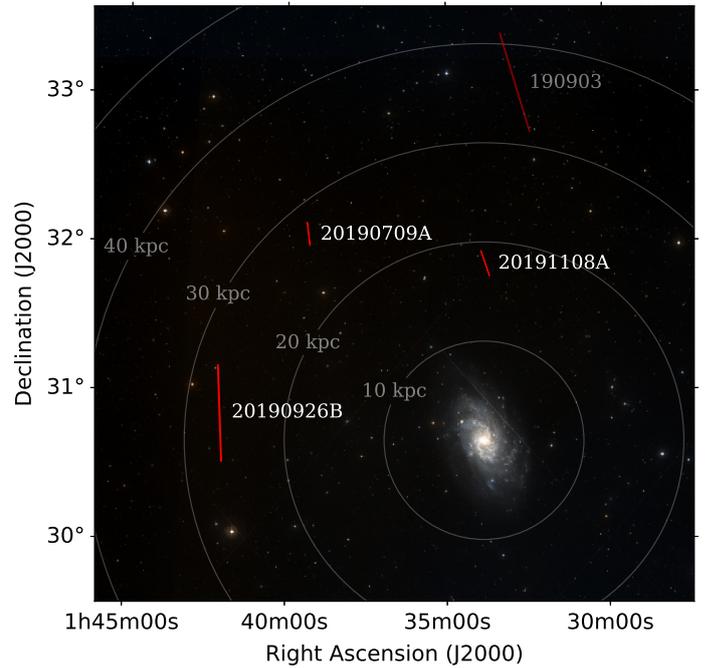
In Fig. 32 we show the resulting number of expected galaxies as a function of the comoving volume spanned by the FRB localisation region and the redshift limit. The number of expected dwarf galaxies is  $\geq 1$  for all our FRBs, up to several hundred for FRB 20190926B. If a localisation region contains a known dwarf galaxy, one thus cannot straightforwardly conclude it is the host. In contrast, the number of expected massive galaxies is less than one for three of our FRBs. For the FRB with the smallest comoving volume associated with its localisation region, FRB 20191108A, the expected number is  $\sim 0.3$ . Assuming Poissonian error bars, the probability of finding a massive galaxy in the region by chance is then 27%. It is therefore not possible to definitively associate a massive galaxy with any of our FRBs.

In order to rule out spatial coincidence at the 95% level, we find that the comoving volume should be limited to  $\sim 30 \text{ Mpc}^3$  for association with a massive galaxy, and  $\sim 2 \text{ Mpc}^3$  for association with a dwarf galaxy. Such volumes could be reached by either better FRB localisation, or by finding an FRB at a relatively small distance. For example, an FRB with the same localisation region as our best-localised burst, FRB 20191108A, but at redshift of  $z = 0.28$ , as opposed to the current upper limit of  $z = 0.54$  (see Table 6), would be localised to the required comoving volume to rule out spatial coincidence with a massive galaxy. For a dwarf galaxy association, the redshift upper limit is  $z = 0.11$ . Such redshifts are not unreasonable within the overall FRB population; about a quarter of FRBs localised to hosts are within that redshift range (we note, though, that localisable FRBs will generally be relatively close by, which biases the quoted fraction). In Pastor-Marazuela et al. (2023) we will present such bursts.

Another class of potentially interesting FRB counterparts are the persistent radio sources (PRSSs). These were discovered for the first repeater FRB 121102 (Chatterjee et al. 2017; Marcote et al. 2017), and for repeating FRB 20190520B (Niu et al. 2022). As also discussed in Connor et al. (2020), radio point sources are sparser than optical galaxies, hence ARTS localisation regions

**Table 7.** Flux densities above which no persistent radio sources are expected to be found in the FRB localisation region by chance at the 10% level, for a  $5\sigma$  sensitivity limit of  $350 \mu\text{Jy}$ .

FRB	S (mJy)
20190709A	0.21
190903	5.4
20190926B	5.3
20191020B	5.4
20191108A	0.14
20191109A	0.34



**Fig. 33.** Location of the three FRBs and one FRB candidate that fall within 50 kpc of M33.

might be small enough to identify radio sources associated with our FRBs.

For each FRB, we set a lower limit to the flux density of a persistent radio source that is ruled out to be in the localisation region by chance at the 10% level, following Eftekhari et al. (2018). The resulting flux densities are listed in Table 7. Given the  $5\sigma$  sensitivity limit for the Apertif imaging of  $350 \mu\text{Jy}$  as used in Connor et al. (2020), we note that Apertif imaging could identify persistent radio sources that are unlikely to be in the localisation region by chance for the majority of our FRB sample. In the first Apertif imaging data release (Adams et al. 2022) the sensitivity limit is even lower. Apertif thus not only has the capability to discover FRBs; it can identify the potentially associated persistent radio sources as well.

### 10.5. Probing the M33 halo

As shown in Fig. 33, three of our first four detections and one unverified FRB candidate were found in the angular vicinity of Local Group galaxy M33 (the Triangulum galaxy). This is because our first detection, FRB 20190709A, was discovered during a calibration drift scan observation of the quasar 3C48,

which is in the Triangulum constellation. FRB 20190926B and FRB 20191108A were later detected during follow up observations of our first discovery. The dispersion and scattering of this set of FRBs might help investigate the M33 halo. FRB 20191108A has the lowest angular separation from the core of M33 and has the smallest localisation region, but all three sources cut well through the M33 halo. They next also intersect the halo of the much larger, M31 (Andromeda) galaxy that is close to M33; and finally, the halo and disk of our own Milky Way. Interactions between these three Local Group galaxies produce connecting gas bridges that the FRBs also skewer. The hot gas bridge that Qu et al. (2021) identified between the Milky Way and M31, for example, is in the FRB line of sight. Combining the  $\sim 15^\circ$  angular distance between the FRBs and M31 with the bridge model of Qu et al. (2021) suggests the baryon bridge disperses the bursts by an additional  $\sim 40\text{--}200\text{ pc cm}^{-3}$ . All the FRB DMs thus ought to have components that are attributable to the plasma of M33 and M31, and to the Local Group bridges; and that amount can be no larger than the minimum extragalactic DM of the three FRBs. As the lowest extragalactic DM of the three is  $\sim 540\text{ pc cm}^{-3}$ , the electron column density in the circumgalactic medium of M33 and M31 must be less than  $\sim 500\text{--}340\text{ pc cm}^{-3}$ . A larger sample of FRBs from Apertif and other surveys may better establish this bridge floor, or even a spatial DM gradient, in the direction of M31 in the future. In that sense it is unfortunate that dedicated surveys with LOFAR did not find further FRBs in either M31 (van Leeuwen et al. 2020) or M33 (Mikhailov & van Leeuwen 2016). None of these three Apertif bursts show any evidence of temporal scattering. As described in Connor et al. (2020), FRB 20191108A shows some frequency structure, but the broad,  $\sim 40\text{ MHz}$  fluctuations are likely not due to propagation effects in the M33 halo.

### 10.6. All-sky burst rate

With ALERT we have discovered five FRBs in the 800 h observed during 2019: one FRB every  $\sim 6.9$  days. To convert this to an all-sky rate, we use the FoV derived in Sect. 7.2, of  $8.2\text{ sq. deg.}$  Using a Poissonian 90% confidence interval (Gehrels 1986), the inferred all-sky rate is  $700^{+800}_{-400}\text{ bursts sky}^{-1}\text{ day}^{-1}$ . We derive the fluence completeness threshold for this rate using the system SEFD. At the CB centres the SEFD is best, typically  $85\text{ Jy}$  (Sect. 7.1). The most conservative completeness threshold would be based on  $2 \times$  this SEFD, if we consider our FoV out to half-power. Most FRBs, however, will be found in the most sensitive part of the FoV. In 85% of the FoV, the sensitivity is at least 70% of the maximum value. We thus take that 70% as the realistic sensitivity threshold. Using the radiometer equation we then find a fluence completeness threshold of  $1.6\sqrt{w/\text{ms}}\text{ Jy ms}$  for pulses of width  $w$ . The derived burst rate is in agreement with earlier values at 1400 MHz from surveys with similar fluence thresholds (Champion et al. 2016; Bhandari et al. 2018; Rane et al. 2016).

## 11. Data and code releases

ARTS data are distributed as real-time alerts, and as a time-series archive. The ARTS codebase is publicly accessible.

### 11.1. Real time: VOEvents

ARTS sends out real-time FRB detection alerts using the International Virtual Observatory Alliance (IVOA) event messaging format, VOEvent. Detections that meet a pre-defined threshold

for viability are communicated using the FRB-specific VOEvent class, developed and standardised by Petroff et al. (2017). Events are generated locally on the ARTS cluster and broadcast via the ASTRON network. An internal link between ARTS and LOFAR allows for rapid triggering of the LOFAR system (cf. Pastor-Marazuela et al. 2021); currently this system is limited to triggers on previously discovered FRBs where the approximate DM and expected signal-to-noise ratio are known. ARTS VOEvent alerts for new detections are distributed publicly through a COMET broker (Swinbank 2014).

### 11.2. Archive: ALTA and the VO

The public access to Apertif archived data is organised through data releases in the Apertif long-term archive (ALTA) and in the virtual observatory (VO). These contain products for imaging (e.g. Adams et al. 2022) and for time domain. The ALTA web interface<sup>13</sup> provides querying, including an Aladin Lite based sky view, and download capabilities. ASTRONs VO service<sup>14</sup> is published in the VO Registry (Demleitner et al. 2015), making it accessible by VO applications by default. The data sets are accessible using the Table Access Protocol (TAP; Dowler et al. 2019).

The time-domain data are in PSRFITS format, with 1-bit sampling. The data are total intensity (Stokes I). For each pointing, all TABs are available. Data from before 2020 May 1 supply 384 channels of 0.8 MHz each, with a time resolution of 2.048 ms. Later data are archived using higher temporal and spectral resolution, and provide 768 channels of 0.4 MHz each, every 0.8 ms.

### 11.3. Data releases

Generally, data are released in batches per calendar year, together with an associated interpretation paper. Specific additional releases accompanied Oostrum et al. (2020); Connor et al. (2020) and Pastor-Marazuela et al. (2021).

Furthermore, the data specifically recorded during the Apertif Science Verification Campaign (SVC; cf. Sect. 6.4) are in a central data release on ALTA. These data include all TABs for 47 survey field from that campaign, and are publicly available<sup>15</sup>. They are in a similar PSRFITS format as the final data (Sect. 11.2), but recorded at a central frequency of 1400 MHz.

Together with this paper covering our first, 2019 discoveries, we have released the full set of 2019 survey data through ALTA<sup>16</sup> and the VO. The data releases for subsequent calendar years will follow on much shorter timescales.

The five FRBs presented in the current paper (Sect. 9) were discovered in raw, high-resolution data, but are also visible in this archived data. The resulting detection plots are displayed as Fig. 34. All archive data for these five pointings are directly accessible through the VO<sup>17</sup>. For other pointings the metadata is directly available, but the data itself first needs to be staged from tape by the ASTRON helpdesk.

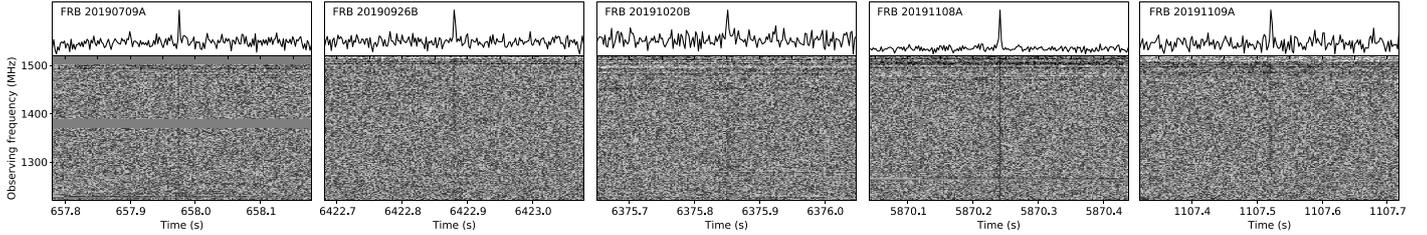
<sup>13</sup> <https://alta.astron.nl/>, includes manual

<sup>14</sup> <https://vo.astron.nl>

<sup>15</sup> [https://alta.astron.nl/science/dataproducts/release\\_release\\_id=SVC\\_2019\\_TimeDomain](https://alta.astron.nl/science/dataproducts/release_release_id=SVC_2019_TimeDomain)

<sup>16</sup> [https://alta.astron.nl/science/dataproducts/release\\_release\\_id=APERTIF\\_DR1\\_TimeDomain](https://alta.astron.nl/science/dataproducts/release_release_id=APERTIF_DR1_TimeDomain)

<sup>17</sup> <https://science.astron.nl/sdc/astron-data-explorer/data-releases/apertif-time-domain-dr1/>



**Fig. 34.** Detection of the five 2019 FRBs in the archival FITS data. Top: pulse profile, bottom: dynamic spectrum.

#### 11.4. Code releases

The codes that are used to operate ARTS and ALERT are open source. They are free for re-use and open to contributions through the TRANSients Software ALLiance (TRASAL<sup>18</sup>). The realtime search system code is released as AMBER (Sclocco et al. 2016, 2020b) and RFI<sub>m</sub> (Sclocco et al. 2020a). The accompanying real-time data handling is released as Attema et al. (2020a,b). The pipeline orchestration code suite is freely available (Oostrum 2020a, 2021a,b,c). Finally, the machine-learning post processing is released as Connor (2018).

Custom code used to generate published Apertif FRB results are released through Zenodo, for both the current paper (van Leeuwen et al. 2022a) and for previous publications (e.g. Pastor-Marazuela et al. 2020).

## 12. Future

### 12.1. Expected detection rates

Detection rate simulations help measure the operational performance of the system, help set survey strategies, and even help select the science priorities. To these ends, we use the FRB population synthesis code base `frbpoppy` (Gardenier et al. 2019) to derive the expected detection rate for ALERT. We adopt the ‘complex’ intrinsic Euclidean FRB population, for an ALERT survey with an SEFD of  $\sim 100$  Jy and the beam sensitivity pattern as presented in Sect. 7. Based on these inputs, and scaling with respect to HTRU detection rates, ALERT should detect 1 FRB every  $4.8^{+3.3}_{-2.1}$  days (where the quoted margin is the  $1\sigma$  Poissonian error). The current detection rate found in Sect. 10.6 falls within this range, implying the system to be functioning as expected.

### 12.2. Operations timescales

The Apertif surveys are operated from 1 July 2019 until 28 Feb 2022. Per that end date, the ALERT survey for pulsars and FRBs ceases. Down-sampled data from the pointings that were covered will continue to be available through the archive.

### 12.3. Possible upgrade to 2D, baseband recording system

The current system is optimised for FRB detection and localisation with a 1D array. As the TABF already coherently forms beams at array angular resolution, over the full FoV, the benefit of a complex-voltage buffer and dump system for offline analysis (such as in ASKAP; Bannister et al. 2019) is limited.

One upgrade under consideration is to extend the WSRT to a 2D array, and add such a baseband buffer system. ARTS would detect FRBs over RT2-RT9/B in real time, as now. Real-time

freeze plus off-line analysis of complex-voltage buffer data from the new and existing dishes then provides 2D localisation. A new 3 km north-south arm, of similar length as the existing 12-dish east-west arm, would provide localisations of  $\sim 1'' \times 1''$ , a good match for optical follow-up. Given the S/N produced by the highly sensitive E-W arm, the N-S arm could be more sparsely populated, using smaller dishes to match the existing FoV.

Each central UNB of the TABF (Sect. 3.4) can hold up to 4 TB of DDR3 RAM, enough to store 10 sec of CB480 voltage data (Table A.1). Such a circular buffer, covering the E-W arm, is sufficiently long for our real-time triggering. The required complex-voltage buffers for the new N-S arm could also be implemented on UNB; the absence of beam-forming requirements, however, may also allow for a simpler hardware solution. After a trigger, the voltages are combined in software correlation and imaging, providing improved FRB localisation.

## 13. Conclusions

ARTS achieves full coherent-addition sensitivity over the entire FoV of Apertif. ARTS powers the ALERT survey, and we detect one FRB every 7 days of observing. The fact that the interferometer consists of steerable dishes, allows for follow-up campaigns on known and newly detected FRBs. While none of the discoveries reported here were seen to repeat, other known repeater FRBs were detected and studied. All the while, new FRBs were found in these fields at the same rate as for blank fields. Five one-off FRBs were discovered in 2019, during the first 6 months of overall operation, and each was interferometrically localised. This combination of solid detection rates and good localisation is essential for mapping the magneto-ionic material in the Universe along different, well-defined lines of sight.

*Acknowledgements.* We thank Richard Blaauw, Raymond van den Brink, Chris Broekema, Lute van de Bult, John Bunton, Adam Deller, Paul Demorest, Roy de Goei, Peter Gruppen, Jason Hessels, Gemma Janssen, Gert Kruithof, Hans van der Marel, Rob van Nieuwpoort, Jan David Mol, Gijs Molenaar, Kaustubh Rajwade, John Romein, Gijs Schoonderbeek, Mike Sipior, Jurjen Sluman, Laura Spitler, Ben Stappers, Marc Verheijen, Nico Vermaas, and René Vermeulen for contributions to the realisation and operation of Apertif, ARTS and ALERT, and for comments and discussions. We thank Vlad Kondratiev, Joe Callingham, Sarvesh Sridhar, Emma Tiggelaar, and Matthijs van der Wiel for suggesting, providing and operating the microwave oven used in Sect. 6.4. This research was supported by the European Research Council (ERC) under the European Union’s Seventh Framework Programme (FP/2007–2013)/ERC Grant Agreement No. 617199 (‘ALERT’); by Vici research programme ‘ARGO’ with project number 639.043.815, financed by the Dutch Research Council (NWO); by the Netherlands eScience Center under the project ‘AA-ALERT’ (027.015.G09, grant ASDI.15.406); by the Netherlands Research School for Astronomy, NOVA, under ‘NOVA-NW3’, and ‘NOVA5-NW3-10.3.5.14’; and through CORTEX (NWA.1160.18.316), under the research programme NWA-ORC, financed by NWO. Instrumentation development was supported by NWO (grant 614.061.613 ‘ARTS’) and NOVA (‘NOVA4-ARTS’). PI of aforementioned grants is JvL. EP acknowledges funding from an NWO Veni Fellowship. The contributions of SMS were supported by NASA grant NNX17AL74G issued

<sup>18</sup> <https://github.com/TRASAL>

through the NNH16ZDA001N Astrophysics Data Analysis Program (ADAP). EAKA is supported by the WISE research programme, which is financed by NWO. BA acknowledges funding from the German Science Foundation DFG, within the Collaborative Research Center SFB1491 “Cosmic Interacting Matters - From Source to Signal”. KMH acknowledges financial support from the State Agency for Research of the Spanish Ministry of Science, Innovation and Universities through the “Center of Excellence Severo Ochoa” awarded to the Instituto de Astrofísica de Andalucía (SEV-2017-0709), from the coordination of the participation in SKA-SPAIN, funded by the Ministry of Science and Innovation (MCIN). KMH and TvDh acknowledge funding from the ERC under the European Union’s Seventh Framework Programme (FP/2007–2013)/ERC Grant Agreement No. 291531 (“HIStoryNU”). RM acknowledges support from the same programme under ERC Advanced Grant RADIOLIFE-320745. This work makes use of data from the Apertif system installed at the Westerbork Synthesis Radio Telescope owned by ASTRON. ASTRON, the Netherlands Institute for Radio Astronomy, is an institute of NWO.

## References

- Abadi, M., Agarwal, A., Barham, P., et al. 2016, ArXiv e-prints [arXiv:1603.04467]
- Adams, E. A. K., & van Leeuwen, J. 2019, *Nat. Astron.*, **3**, 188
- Adams, E. A. K., Adebahr, B., de Blok, W. J. G., et al. 2022, *A&A*, **667**, A38
- Adebahr, B., Berger, A., Adams, E. A. K., et al. 2022, *A&A*, **663**, A103
- Attema, J. J., Oostrum, L. C., & van Leeuwen, J. 2020a, <https://zenodo.org/record/3944623>
- Attema, J. J., Oostrum, L. C., & van Leeuwen, J. 2020b, <https://zenodo.org/record/3944631>
- Backer, D. C. 2000, in *Perspectives on Radio Astronomy: Science with Large Antenna Arrays*, ed. M. P. van Haarlem, 285
- Baldry, I. K., Driver, S. P., Loveday, J., et al. 2012, *MNRAS*, **421**, 621
- Bannister, K. W., Shannon, R. M., Macquart, J.-P., et al. 2017, *ApJ*, **841**, L12
- Bannister, K. W., Deller, A. T., Phillips, C., et al. 2019, *Science*, **365**, 565
- Barsdell, B. R., Bailes, M., Barnes, D. G., & Fluke, C. J. 2012, *MNRAS*, **422**, 379
- Bassa, C. G., Janssen, G. H., Karuppusamy, R., et al. 2016, *MNRAS*, **456**, 2196
- Bera, A., & Chengalur, J. N. 2019, *MNRAS*, **490**, L12
- Bhandari, S., Keane, E. F., Barr, E. D., et al. 2018, *MNRAS*, **475**, 1427
- Bhardwaj, M., Gaensler, B. M., Kaspi, V. M., et al. 2021, *ApJ*, **910**, L18
- Bilous, A. V., Griebmeier, J. M., Pennucci, T., et al. 2022, *A&A*, **658**, A143
- Bochenek, C., Kulkarni, S., Ravi, V., et al. 2020a, *The Astronomer’s Telegram*, **13684**, 1
- Bochenek, C. D., Ravi, V., Belov, K. V., et al. 2020b, *Nature*, **587**, 59
- Boersma, O. M., van Leeuwen, J., Adams, E. A. K., et al. 2021, *A&A*, **650**, A131
- Bregman, J. A., Schipper, B., Beereckamp, W., et al. 2018, in *Proceedings of Westerbork Telescope 50th Anniversary*, 361, 11
- Burke Spolaor, S. 2018, in *FRB2018: Finding and understanding Fast Radio Bursts*, <http://caastro.org/wp-content/uploads/2018/06/spolaor-frb2018-realfast-swinburne-1.pdf>
- Caleb, M., Flynn, C., Bailes, M., et al. 2017, *MNRAS*, **468**, 3746
- Camilo, F., Ransom, S. M., Halpern, J. P., et al. 2006, *Nature*, **442**, 892
- Champion, D. J., Petroff, E., Kramer, M., et al. 2016, *MNRAS*, **460**, L30
- Chatterjee, S., Law, C. J., Wharton, R. S., et al. 2017, *Nature*, **541**, 58
- Chawla, P., Kaspi, V. M., Josephy, A., et al. 2017, *ApJ*, **844**, 140
- Chawla, P., Andersen, B. C., Bhardwaj, M., et al. 2020, *ApJ*, **896**, L41
- CHIME/FRB Collaboration (Amiri, M., et al.) 2018, *ApJ*, **863**, 48
- CHIME/FRB Collaboration (Andersen, B. C., et al.) 2019, *ApJ*, **885**, L24
- CHIME/FRB Collaboration (Andersen, B. C., et al.) 2020, *Nature*, **587**, 54
- CHIME/FRB Collaboration (Amiri, M., et al.) 2021, *ApJS*, **257**, 59
- CHIME/FRB Collaboration (Andersen, B. C., et al.) 2022, *Nature*, **607**, 256
- Chollet, F. et al. 2015, <https://keras.io>
- Chu, Q., Howell, E. J., Rowlinson, A., et al. 2016, *MNRAS*, **459**, 121
- Coenen, T., van Leeuwen, J., Hessels, J. W. T., et al. 2014, *A&A*, **570**, A60
- Connor, L. 2018, <https://zenodo.org/record/1442657>
- Connor, L. 2019, *MNRAS*, **487**, 5753
- Connor, L., & van Leeuwen, J. 2018, *AJ*, **156**, 256
- Connor, L., van Leeuwen, J., Oostrum, L. C., et al. 2020, *MNRAS*, **499**, 4716
- Cordes, J. M., & Lazio, T. J. W. 2002, ArXiv e-prints [arXiv:astro-ph/0207156]
- Cordes, J. M., & McLaughlin, M. A. 2003, *ApJ*, **596**, 1142
- Cruces, M., Spitler, L. G., Scholz, P., et al. 2021, *MNRAS*, **500**, 448
- Dai, S., Lower, M. E., Bailes, M., et al. 2019, *ApJ*, **874**, L14
- Dálya, G., Galgóczi, G., Dobos, L., et al. 2018, *MNRAS*, **479**, 2374
- Demleitner, M., Harrison, P., Taylor, M., & Normand, J. 2015, *Astron. Comput.*, **11**, 91
- Dénes, H., Hess, K. M., Adams, E. A. K., et al. 2022, *A&A*, **667**, A40
- Desvignes, G., Kramer, M., Lee, K., et al. 2019, *Science*, **365**, 1013
- Dowler, P., Rixon, G., Tody, D., & Demleitner, M. 2019, *Table Access Protocol Version 1.1*, IVOA Recommendation, 27 September 2019
- Eftekhari, T., Berger, E., Williams, P. K. G., & Blanchard, P. K. 2018, *ApJ*, **860**, 73
- Faber, S. M., Willmer, C. N. A., Wolf, C., et al. 2007, *ApJ*, **665**, 265
- Falcke, H., & Rezzolla, L. 2014, *A&A*, **562**, A137
- Fonseca, E., Andersen, B. C., Bhardwaj, M., et al. 2020, *ApJ*, **891**, L6
- Gardenier, D. W., & van Leeuwen, J. 2021, *A&A*, **651**, A63
- Gardenier, D. W., van Leeuwen, J., Connor, L., & Petroff, E. 2019, *A&A*, **632**, A125
- Gardenier, D. W., Connor, L., van Leeuwen, J., Oostrum, L. C., & Petroff, E. 2021, *A&A*, **647**, A30
- Gehrels, N. 1986, *ApJ*, **303**, 336
- Hallinan, G., Corsi, A., Mooley, K. P., et al. 2017, *Science*, **358**, 1579
- Haynes, M. P., Giovanelli, R., Martin, A. M., et al. 2011, *AJ*, **142**, 170
- Hess, K. M., Dénes, H., Oostrum, L., & Adams, E. A. K. 2022, <https://zenodo.org/record/6988465>
- Hess, K. M., et al. 2023, *MNRAS*, submitted
- Hessels, J. W. T., Spitler, L. G., Seymour, A. D., et al. 2019, *ApJ*, **876**, L23
- Hilmarsson, G. H., Michilli, D., Spitler, L. G., et al. 2021, *ApJ*, **908**, L10
- Hobbs, G., & Edwards, R. 2012, *Astrophysics Source Code Library* [record ascl:1210.015]
- Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, *MNRAS*, **369**, 655
- Holties, H., de Goeij, R., & Vermaas, N. 2019, *ASP Conf. Ser.*, **527**, 255
- Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, *PASA*, **21**, 302
- Houben, L. J. M., Spitler, L. G., ter Veen, S., et al. 2019, *A&A*, **623**, A42
- Hulse, R. A., & Taylor, J. H. 1975, *ApJ*, **201**, L55
- James, C. W., Osłowski, S., Flynn, C., et al. 2020, *MNRAS*, **495**, 2416
- Janssen, G. H., Stappers, B. W., Braun, R., et al. 2009, *A&A*, **498**, 223
- Karastergiou, A., Chennamangalam, J., Armour, W., et al. 2015, *MNRAS*, **452**, 1254
- Karuppusamy, R., Stappers, B., & van Straten, W. 2008, *PASP*, **120**, 191
- Keane, E. F., Kramer, M., Lyne, A. G., Stappers, B. W., & McLaughlin, M. A. 2011, *MNRAS*, **415**, 3065
- Kooistra, E. 2021, <https://zenodo.org/record/4748331>
- Kramer, M., & Champion, D. J. 2013, *Classical Quant. Grav.*, **30**, 224009
- Kramer, M., Stairs, I. H., Manchester, R. N., et al. 2006, *Science*, **314**, 97
- Kumar, P., Shannon, R. M., Lower, M. E., et al. 2022, *MNRAS*, **512**, 3400
- Kutkin, A. M., Oosterloo, T. A., Morganti, R., et al. 2022, *A&A*, **667**, A39
- Law, C. J., Bower, G. C., Burke-Spolaor, S., et al. 2018, *ApJS*, **236**, 8
- Loose, M. 2016, in *Astronomical Data Analysis Software and Systems XXVI*, <http://www.adass2016.inaf.it/index.php/participants-a11/4-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, XXXIIInd General Assembly and Scientific Symposium of the International Union of Radio Science (URSI GASS), Montreal, QC, 2
- Maan, Y., Joshi, B. C., Surnis, M. P., Bagchi, M., & Manoharan, P. K. 2019, *ApJ*, **882**, L9
- Macquart, J.-P., Bailes, M., Bhat, N. D. R., et al. 2010, *PASA*, **27**, 272
- Macquart, J. P., Keane, E., Grainge, K., et al. 2015, Proc. Advancing Astrophysics with the Square Kilometre Array (AASKA14, 9-13 June, 2014, Giardini Naxos, Italy, online at <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=215>, 55
- Manchester, R. N., Lyne, A. G., Camilo, F., et al. 2001, *MNRAS*, **328**, 17
- Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *AJ*, **129**, 1993
- Marcote, B., Paragi, Z., Hessels, J. W. T., et al. 2017, *ApJ*, **834**, L8
- Marcote, B., Nimmo, K., Hessels, J. W. T., et al. 2020, *Nature*, **577**, 190
- Mereghetti, S., Savchenko, V., Ferrigno, C., et al. 2020, *ApJ*, **898**, L29
- Michilli, D., Seymour, A., Hessels, J. W. T., et al. 2018, *Nature*, **553**, 182
- Mikami, R., Asano, K., Tanaka, S. J., et al. 2016, *ApJ*, **832**, 212
- Mikhailov, K. 2018, Ph.D. Thesis, University of Amsterdam, The Netherlands
- Mikhailov, K., & van Leeuwen, J. 2016, *A&A*, **593**, A21
- Mikhailov, K., & Sclocco, A. 2018, *Astron. Comput.*, **25**, 139
- Niu, C. H., Aggarwal, K., Li, D., et al. 2022, *Nature*, **606**, 873
- Oosterloo, T., Verheijen, M. A. W., van Cappellen, W., et al. 2009, Proc. Wide Field Astronomy & Technology for the Square Kilometre Array (SKADS 2009), 4-6 November 2009, Chateau de Limelette, Belgium, online at <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=132>, 70
- Oostrum, L. C. 2020a, <https://zenodo.org/record/4292175>
- Oostrum, L. C. 2020b, Ph.D. Thesis, University of Amsterdam, The Netherlands
- Oostrum, L. C. 2021a, <https://zenodo.org/record/4560043>
- Oostrum, L. C. 2021b, <https://zenodo.org/record/4452631>
- Oostrum, L. C. 2021c, <https://zenodo.org/record/4419555>
- Oostrum, L. C., Maan, Y., van Leeuwen, J., et al. 2020, *A&A*, **635**, A61
- Pastor-Marazuela, I., Vohl, D., Oostrum, L. C., & Connor, L. 2020, <https://zenodo.org/record/4559593>

- Pastor-Marazuela, I., Connor, L., van Leeuwen, J., et al. 2021, *Nature*, **596**, 505
- Pastor-Marazuela, I., van Leeuwen, J., Bilous, A., et al. 2022, *A&A*, submitted [arXiv:2202.08002]
- Pastor-Marazuela, I., van Leeuwen, J., Bilous, A., et al. 2023, *A&A*, submitted
- Pearlman, A. B., Majid, W. A., Prince, T. A., Kocz, J., & Horiuchi, S. 2018, *ApJ*, **866**, 160
- Perley, R. A., & Butler, B. J. 2017, *ApJS*, **230**, 7
- Petroff, E., Bailes, M., Barr, E. D., et al. 2015a, *MNRAS*, **447**, 246
- Petroff, E., Keane, E. F., Barr, E. D., et al. 2015b, *MNRAS*, **451**, 3933
- Petroff, E., Barr, E. D., Jameson, A., et al. 2016, *PASA*, **33**, e045
- Petroff, E., Houben, L., Bannister, K., et al. 2017, arXiv e-prints, [arXiv:1710.08155]
- Petroff, E., Hessels, J. W. T., & Lorimer, D. R. 2019a, *A&A Rev.*, **27**, 4
- Petroff, E., Oostrum, L. C., Stappers, B. W., et al. 2019b, *MNRAS*, **482**, 3109
- Petroff, E., Hessels, J. W. T., & Lorimer, D. R. 2022, *A&A Rev.*, **30**, 2
- Platts, E., Weltman, A., Walters, A., et al. 2019, *Phys. Rep.*, **821**, 1
- Pleunis, Z., Good, D. C., Kaspi, V. M., et al. 2021a, *ApJ*, **923**, 1
- Pleunis, Z., Michilli, D., Bassa, C. G., et al. 2021b, *ApJ*, **911**, L3
- Prochaska, J. X., Macquart, J.-P., McQuinn, M., et al. 2019, *Science*, **365**, aay0073
- Qu, Z., Huang, R., Bregman, J. N., & Li, J.-T. 2021, *ApJ*, **907**, 14
- Rajwade, K. M., Mickaliger, M. B., Stappers, B. W., et al. 2020, *MNRAS*, **495**, 3551
- Rane, A., Lorimer, D. R., Bates, S. D., et al. 2016, *MNRAS*, **455**, 2207
- Ransom, S. 2011, Astrophysics Source Code Library [record ascl:1107.017]
- Ravi, V. 2019, *MNRAS*, **482**, 1966
- Ravi, V., Catha, M., D'Addario, L., et al. 2019, *Nature*, **572**, 352
- Sanidas, S., Cooper, S., Bassa, C. G., et al. 2019, *A&A*, **626**, A104
- Schoenmakers, A., Coolen, A., & Mulder, H. 2019, in *Astronomical Data Analysis Software and Systems XXIX*, <https://www.adass2019.nl/conference-info/posters/pdf?table=poster&id=116>
- Schoonderbeek, G. W., Szomoru, A., Gunst, A. W., Hiemstra, L., & Hargreaves, J. 2019, *J. Astron. Instrum.*, **8**, 1950003
- Scalco, A., van Leeuwen, J., Bal, H. E., & van Nieuwpoort, R. V. 2016, *Astron. Comput.*, **14**, 1
- Scalco, A., Vohl, D., & van Nieuwpoort, R. V. 2020a, arXiv e-prints [arXiv:2001.03389]
- Scalco, A., Heldens, S., & van Werkhoven, B. 2020b, *SoftwareX*, **12**, 100549
- Shannon, R. M., Macquart, J.-P., Bannister, K. W., et al. 2018, *Nature*, **562**, 386
- Smits, R., Kramer, M., Stappers, B., et al. 2009, *A&A*, **493**, 1161
- Sokolowski, M., Bhat, N. D. R., Macquart, J. P., et al. 2018, *ApJ*, **867**, L12
- Spitler, L. G., Cordes, J. M., Hessels, J. W. T., et al. 2014, *ApJ*, **790**, 101
- Spitler, L. G., Scholz, P., Hessels, J. W. T., et al. 2016, *Nature*, **531**, 202
- Stappers, B. W., Hessels, J. W. T., Alexov, A., et al. 2011, *A&A*, **530**, A80
- Straal, S. M. 2018, Ph.D. Thesis, University of Amsterdam, The Netherlands
- Strom, R., Van Ardenne, A., & Torchinsky, S. 2018, 50 Years Westerbork Radio Observatory: A Continuous Journey to Discoveries and Innovations, Astrophysics and space science library (ASTRON), <https://pos.sissa.it/361/>
- Swinbank, J. 2014, *Astron. Comput.*, **7**, 12
- Szomoru, A. 2010, in *10th European VLBI Network Symposium and EVN Users Meeting: VLBI and the New Generation of Radio Arrays*, 10, 98
- Tan, G. H. 1991, *ASP Conf. Ser.*, **131**, 42
- Tendulkar, S. P., Bassa, C. G., Cordes, J. M., et al. 2017, *ApJ*, **834**, L7
- ter Veen, S. 2015, Ph.D. Thesis, Radboud University Nijmegen, The Netherlands <http://hdl.handle.net/2066/147186>
- ter Veen, S., Enriquez, J. E., Falcke, H., et al. 2019, *A&A*, **621**, A57
- Thornton, D., Stappers, B., Bailes, M., et al. 2013, *Science*, **341**, 53
- van Cappellen, W. A., Oosterloo, T. A., Verheijen, M. A. W., et al. 2022, *A&A*, **658**, A146
- 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*, eds. P. R. Wozniak, M. J. Graham, A. A. Mahabal, & R. Seaman, 79
- van Leeuwen, J., & Stappers, B. W. 2010, *A&A*, **509**, 7
- van Leeuwen, J., Kasian, L., Stairs, I. H., et al. 2015, *ApJ*, **798**, 118
- van Leeuwen, J., Mikhailov, K., Keane, E., et al. 2020, *A&A*, **634**, A3
- van Leeuwen, J., Oostrum, L. C., Pastor-Marazuela, I., Vohl, D., & Wang, Y. Y. 2022a, <https://zenodo.org/record/6415175>
- van Leeuwen, J., Kooistra, E., Oostrum, L., et al. 2022b, arXiv e-prints [arXiv:2205.12362v1]
- van Straten, W., Jameson, A., & Osłowski, S. 2021, Astrophysics Source Code Library [record ascl:2110.003]
- Vermaas, N., Moss, V., & de Goeij, R. 2019, in *Astronomical Data Analysis Software and Systems XXIX*, <https://www.adass2019.nl/conference-info/posters/pdf?table=poster&id=216>
- Wang, Y. Y., van Leeuwen, J., Gardenier, D. W., et al. 2023, *A&A*, submitted
- Wijnholds, S. J., Ivashina, M. V., Maaskant, R., & Warnick, K. F. 2012, *IEEE Trans. Antennas Propag.*, **60**, 4688
- Yamasaki, S., & Totani, T. 2020, *ApJ*, **888**, 105
- Yao, J. M., Manchester, R. N., & Wang, N. 2017, *ApJ*, **835**, 29
- Yaron, O., Ofek, E., Gal-Yam, A., & Sass, A. 2020, *Transient Name Server AstroNote*, **70**, 1
- York, D. G., Adelman, J., Anderson, Jr., J. E., et al. 2000, *AJ*, **120**, 1579

- <sup>1</sup> ASTRON, the Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, The Netherlands  
e-mail: leeuwen@astron.nl
- <sup>2</sup> Anton Pannekoek Institute, University of Amsterdam, Postbus 94249, 1090 GE Amsterdam, The Netherlands
- <sup>3</sup> Netherlands eScience Center, Science Park 402, 1098 XH Amsterdam, The Netherlands
- <sup>4</sup> Cahill Center for Astronomy, California Institute of Technology, Pasadena, CA, USA
- <sup>5</sup> National Centre for Radio Astrophysics, Tata Institute of Fundamental Research, Pune 411007, Maharashtra, India
- <sup>6</sup> Department of Physics, McGill University, 3600 rue University, Montréal, QC H3A 2T8, Canada
- <sup>7</sup> NYU Abu Dhabi, PO Box 129188, Abu Dhabi, UAE
- <sup>8</sup> Center for Astro, Particle, and Planetary Physics (CAP<sup>3</sup>), NYU Abu Dhabi, PO Box 129188, Abu Dhabi, UAE
- <sup>9</sup> Kapteyn Astronomical Institute, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands
- <sup>10</sup> Astronomisches Institut der Ruhr-Universität Bochum (AIRUB), Universitätsstrasse 150, 44780 Bochum, Germany
- <sup>11</sup> Dept. of Astronomy, Univ. of Cape Town, Private Bag X3, Rondebosch 7701, South Africa
- <sup>12</sup> Instituto de Astrofísica de Andalucía (CSIC), Glorieta de la Astronomía s/n, 18008 Granada, Spain
- <sup>13</sup> Department of Physics, Virginia Polytechnic Institute and State University, 50 West Campus Drive, Blacksburg, VA 24061, USA
- <sup>14</sup> CSIRO Astronomy and Space Science, Australia Telescope National Facility, PO Box 76, Epping NSW 1710, Australia
- <sup>15</sup> Sydney Institute for Astronomy, School of Physics, University of Sydney, Sydney, New South Wales 2006, Australia
- <sup>16</sup> Joint Institute for VLBI ERIC (JIVE), Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, The Netherlands
- <sup>17</sup> University of Oslo Center for Information Technology, PO Box 1059, 0316 Oslo, Norway

## Appendix A: Detailed design I: Functions and data

Three appendices with more detail on the ARTS design follow: (A) covers function and data; (B) describes hardware; and (C) presents software and firmware.

We move through the system following the end-to-end data path shown in Fig. A.1. Functions on this path either *process* data (e.g. fringe stopping, filterbank, and beamformer), or *move* data (e.g. reordering, selection, packetising, and physical I/O).

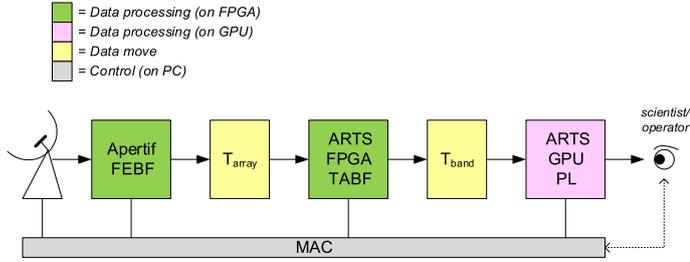


Fig. A.1: ARTS top level data path, with main interfaces. A user controls ARTS via the MAC and accesses the data products via the ARTS pipeline processing (PL).

### Appendix A.1: Relevant general subsystem: Apertif Front-End Beamformer

The Apertif Front-End Beamformers (FEBFs; van Cappellen et al. 2022) provide the CB480 data to ARTS. The FEBFs first sample the data of all antenna elements at 800 MHz. This produces  $800 \text{ MHz} \times 8 \text{ bits/sample} \times 128 \text{ ADCs/dish} \times 12 \text{ dishes} = 9.8 \text{ Tbps}$ . Of these ADCs, 121 are connected to a PAF receiver element. The two polarisations are processed independently. The sampled data are next separated into 512 subbands by means of a PFB that uses a 1024-point Fast Fourier Transform (FFT). The subband filterbank ( $F_{\text{sub}}$ ; see Fig. A.2) is critically sampled, at 781250 Hz. The PAF beamformer then forms *beamlets*, a beam formed for one subband. This requires a transpose  $T_{\text{ant}}$  that groups subbands from all 64 PAF antenna elements per polarisation. A CB is formed from 384 beamlets with identical direction, together spanning 300 MHz bandwidth (BW). With  $N_{\text{dish}} = 12 \text{ dishes} \times K_{\text{CB}} = 40 \text{ compound beams}$  the total output of the FEBF is referred to as CB480. The CB480 data flows at  $12 \text{ dishes} \times 2 \text{ polarisations} \times 40 \text{ CBs} \times 300 \text{ MHz} \times 2 \text{ complex} \times 6 \text{ bit} = 3.456 \text{ Tbps}$ , at 9 Gbps per 10G link (Table A.1).

Fringe stopping (FS) compensates for the change in geometrical delay between dishes due to Earth rotation. In the FEBF,

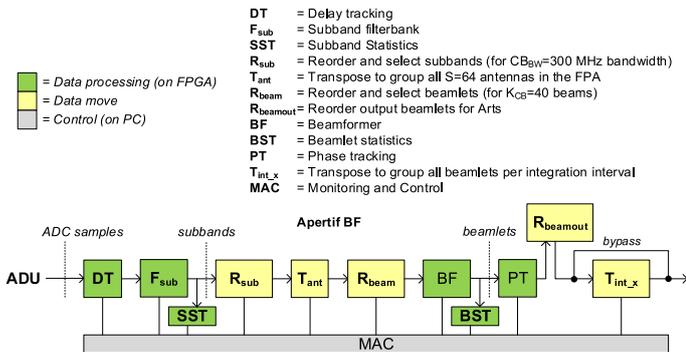


Fig. A.2: The Apertif FEBF subsystem that is input to ARTS. The  $T_{\text{int}_x}$  transpose is bypassed for ARTS Timing (TAB1).

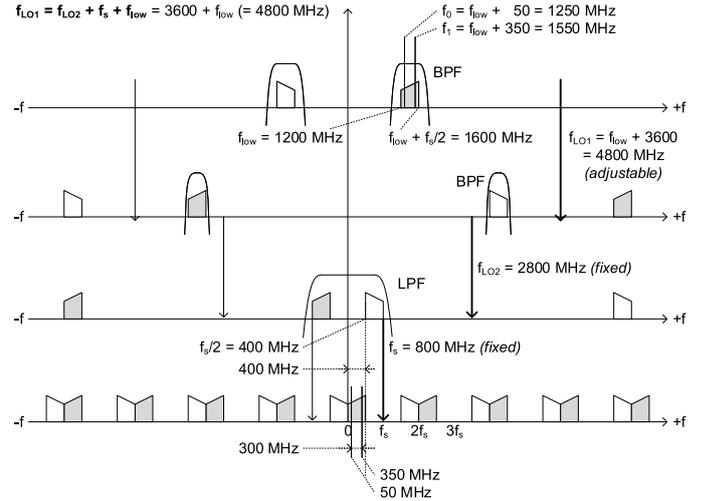


Fig. A.3: Apertif FEBF mixer and sub-sampling scheme for 300 MHz beams, from e.g.  $f_0 = 1250 \text{ MHz}$  to  $f_1 = 1550 \text{ MHz}$ . The  $\text{RF}_{\text{BW}} = 400 \text{ MHz}$ . The  $f_{\text{LO}2} = 2800 \text{ MHz}$  and is fixed. By setting  $f_{\text{LO}1}$  appropriately, the  $f_{\text{low}}$  maps to 0 Hz and  $f_0$  to 50 MHz.

it consists of true sample delay tracking (DT) operating on the ADC samples, followed by a residual phase tracking (PT) on the beamlets per compound beam. FS is thus valid for the centre of each CB.

### Appendix A.2: Frequency and bandwidth

The FEBF operates in the 1130–1720 MHz radio frequency (RF) range (van Cappellen et al. 2022). Fig. A.3 shows how 400 MHz is down converted and sub-sampled, and how the central frequency can be adjusted in steps of 10 MHz.

Typically the middle subbands 64 to 447 (50 to 350 MHz) are selected (cf. Fig. A.3 and A.4). Down stream, these subbands are split into 4 sub-channels in the TABF (Fig. A.4). By implementing this split as a 64-channel separation followed by a 16-channel integration, the impact of the channel leakage in the FFT is reduced, while effectuating the intended downsampling in time (Sect. 3.3).

The PFB for the 64 channels uses a prototype FIR filter with  $8 \text{ taps} \times 64 \text{ points} = 512 \text{ coefficients}$ . The FIR filter coefficients<sup>19</sup> are 9 bit real values. The half power bandwidth of the channels is equal to the channel frequency (12.2 kHz) and the stop band attenuation is >55 dB.

The early design choice for  $400/512 \text{ MHz} = 781250 \text{ Hz}$  channels means frequency channels need to be converted to tie Apertif into e.g. LEAP for pulsar timing (Bassa et al. 2016) and into standard VLBI. Both use (multiples of) 1 MHz channels.

### Appendix A.3: Data formats and interface data loads

The ARTS beamformers (Sect A.4) produce data in voltage and/or power data formats (Fig. A.5).

For voltage beams, such as the CB and the TAB1 for Timing, the data vector consists of the complex samples for both

<sup>19</sup> These can be recreated using `run_pfir_coeff.m` with application = 'apertif\_channel' from <https://doi.org/10.5281/zenodo.4748331> (Kooistra 2021), the code that defines the subband and channel characteristics.

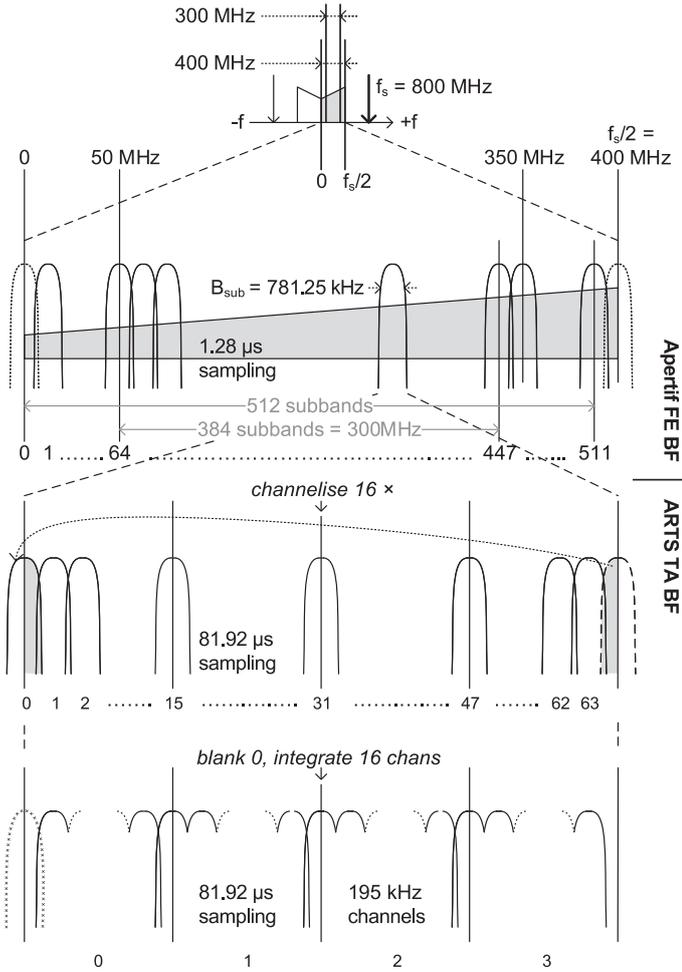


Fig. A.4: PFBs in the Apertif FEBF, with  $N_{\text{sub}} = 512$  subbands, followed by a second filterbank, plus an integration step, in the TABF.

polarisations. The search modes use integrated Stokes power data beams, such that samples can be integrated to achieve data reduction (see Sect. 4.1.1).

Table A.1 summarises the data loads on the ARTS beamformer (BF) interfaces. The default input data is CB480 for a dedicated search (Sect. A.1). For commensal searching (Sect. 3.5), the input data load in CHAN320 mode is 12 dishes  $\times$  2 polarisations  $\times$  40 CB  $\times$  300 MHz  $\times$  2 complex  $\times$  9 bit = 3.456 Tbps.

The default search output is 12 TABs per CB for TAB480. Optionally each first TAB can be replaced by an IAB for IAB40. The search mode internally first creates TAB480 voltage streams, totaling 12 dishes  $\times$  2 polarisations  $\times$  40 CB  $\times$  300 MHz  $\times$  2 complex  $\times$  8 bit = 4.608 Tbps. For each TAB the output consists of both a Stokes-I and a Stokes-IQUV data stream. As

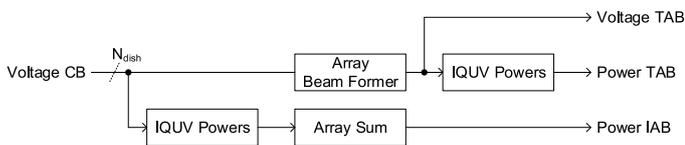


Fig. A.5: Coherent and incoherent array beamforming in ARTS.

Table A.1: Data-rate loads for ARTS TABF interfaces.

Mode	Data Type	Dish N	CBs N	TABs N	$N_{\text{int}}$	Load Gbps
CB480	CV	12	40	-	-	3456
CHAN320	CV	8	40	-	-	3456
TAB1	CV	12	1	1	-	9.6
TAB480	CV	12	40	12	-	4608
TAB480	IQUV	12	40	12	16	288
TAB480	I	12	40	12	16	72
I+IQUV/node		12	40	12	16	9
IAB40	IQUV	12	40	-	16	24

**Notes.** The number of samples per integration is  $N_{\text{int}} = 16$ , for 81.92  $\mu\text{s}$  sampling. In CV modes the data are complex voltages. I and IQUV modes are detected Stokes data. I+IQUV/node denotes the load of the combine Stokes I and IQUV data per GPU workstation, at  $(288 + 72) / 40 = 9$  Gbps.

the Stokes I data (for real time detection) and IQUV data (for data buffering) have different purpose and format, the Stokes I data are included (again) in the IQUV data stream. For downstream performance the Stokes-I streams are demultiplexed and re-ordered on the TABF such that the time index is fastest changing index. The Stokes-IQUV streams remain in the original multiplexed format.

#### Appendix A.4: Beamforming algorithms

The coherent beamformer weights and sums the input voltage data from dishes  $d$ . The weights  $w_{\text{TAB}}$  are complex values that are purely geometric and generally static (Sect. 4.1.1). The coherent beamformer operates separately per polarisation  $p$ . The single voltage output data  $v\text{TAB}$ , that directly operates on the input beamlets  $b$ , is  $v\text{TAB}(p, b) = \sum_{d=0}^{N_{\text{dish}}-1} w_{\text{TAB}}(d, b) \times \text{CB}(p, d, b)$ .

For the intermediate voltage TABs used in the survey modes, the same weight is used for all subchannels:  $v\text{TAB}(p, s) = \sum_{d=0}^{N_{\text{dish}}-1} w_{\text{TAB}}(d, b) \times \text{CB}(p, d, s)$ .

These vTABs are first converted into full Stokes power data and next integrated over  $N_{\text{int}_a} = 16$  subchannels into each channel  $c$ , resulting in the integrated power  $p\text{TAB}$  stream to the search PL:  $p\text{TAB}(c) = \sum_{s=0}^{N_{\text{int}_a}-1} \text{Stokes}(v\text{TAB}(p, s))$

In contrast to the TABs, the incoherent beamformer already operates on power data (see Fig. A.5). Weights are generally 0 or 1. It, too, next integrates over the subchannels:  $\text{IAB}(c) = \sum_{s=0}^{N_{\text{int}_a}-1} \sum_{d=0}^{N_{\text{dish}}-1} w_{\text{IAB}}(d) \times \text{Stokes}(\text{CB}(p, d, s))$ .

#### Appendix A.5: FPGA usage

Each UNB uses 8 Stratix IV FPGAs. Each UNB2 uses 4 Arria10 FPGAs. Both are clocked at 200 MHz.

The FPGA usage of the connectivity is defined by the required optical links. Each entire UNB (with OEB, see Sect. B.2) and each individual FPGA on UNB2 can connect 24 optical 10G links. The 384 10G links that come in from the FEBF thus determine that the TABF requires at least 16 UniBoards or 4 UNB2. The usage of the logic in the FPGA is split over the utilisation of the flip flops (FF) and the lookup tables (LUT). Each Stratix IV has 182,400 FF, each Arria10 has 1,708,800 FF. The memory in the FPGA is organised in Block RAMs (BRAMs), that can also be used as FIFOs. The multiplier usage follows from Sect. A.4 and the number of TABs/CB.

Table A.2: FPGA resources used versus their availability, over the ARTS TABF.

FPGA		# BRAM × 1000	# multipliers × 1000
16 UNB	128	150 / 158 M9K (94%)	122 / 165 (74%)
4 UNB2	16	41 / 43 M20K (95%)	22 / 24 (90%)

**Notes.** Listed are the UniBoard (UNB) and UniBoard<sup>2</sup> (UNB2) systems, 16 and 4 boards respectively. For the commensal search on UniBoard<sup>2</sup>, resources for the channeliser that also serves the correlator are not included. The BRAM and multiplier usage is noted as Used/Available [Type] (Percentage).

Two telescope paths are multiplexed to approach 100% digital signal processing (DSP) efficiency. The 16 UNBs can thus form up to 14 TABs/CB, where only 12 are required for the dedicated search (cf. Eq. (D.3), Sect. D.1). The 9 TABs/CB required for commensal search can exactly fit on 4 UNB2, but only if channel filterbanking is done in the Apertif correlator (Sect. 3.5). Table A.2 lists the FPGA resource usage in the TABFs.

## Appendix B: Detailed design II: Hardware

### Appendix B.1: Relevant general subsystem: UniBoards for FEBF

Each FEBF subsystem of Fig. A.2 processes one polarisation from one dish for all CBs on one subrack (Fig. B.1). There are thus 24 of these, each containing 4 UNBs. Through a backplane the UNBs connect to each other, and to 8 analogue to digital unit (ADU) boards that together host 64 ADCs, to digitise the data from the 61 antenna elements in the PAF (van Cappellen et al. 2022). On each UNB, the data are channelised on the 4 back node (BN) FPGAs and beamformed on the 4 front node (FN) FPGAs. The transpose  $T_{ant}$  is implemented through a combination of the backplane and the mesh interconnects. The 384 FEBF 10G links output the CB480 data to the TABF.

### Appendix B.2: UniBoards and OEBs for TABF

Fig. B.2 shows the UNB with the OEB. The data processing in both the Apertif FEBF and the ARTS TABF UniBoards is clocked at 200 MHz. All are locked to the same 10 MHz reference as the ADCs sample clock (Sect. B.1), ensuring the

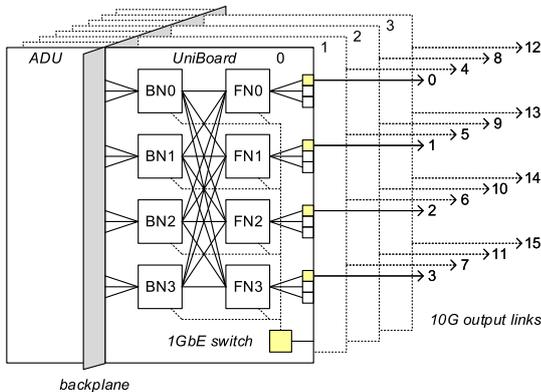


Fig. B.1: One FEBF subrack with 4 UNBs per telescope path. The FPGAs are marked BN and FN.

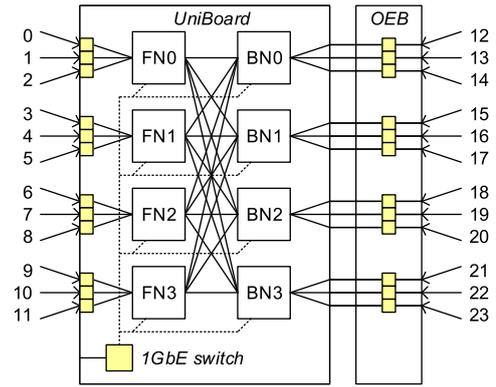


Fig. B.2: One UniBoard + optical-electrical board. This unit processes 18.75 MHz for 24 telescope paths.

processing rate in all FPGAs is held fixed relative to each other. Through this 10 MHz reference the long term stability of the ADC sample clocks is guaranteed by the WSRT maser (van Cappellen et al. 2022). The 8 FPGAs on UNB are StratixIV type EP4SGX230KF40C2. In contrast to the FEBF, there is no distinction between FN and BN in the TABF. All FPGAs have the same function. That is possible because the OEBs extends the BNs with fibre optics I/O. Each FPGA supports three 10G links. On these, each FPGA receives one polarisation from three dishes, for the same 18.75 MHz band. On the UNB, the mesh, consisting of 5 Gbps transceiver links, next interconnects the FPGAs as shown in Fig. B.2. For the dedicated FRB survey, the data are reordered on this mesh such that each FPGA processes the data from five CBs. One UNB plus OEB provides full duplex I/O for 240 Gbps. In total there are 128 FPGAs on the 16 independently-operating UNB in the TABF.

### Appendix B.3: UniBoard<sup>2</sup>s for TABF

On UNB2 all FPGAs are connected identically. There are no BNs or FNs, only PNs. There are 16 such FPGAs, type Arria10 10AX115U2F45E1SG, in the 4 UNB2s comprising the TABF, as shown in Fig. B.3. In most respects, such as connectivity, transposing, and signal processing, one PN on UNB2 replaces an entire UNB.

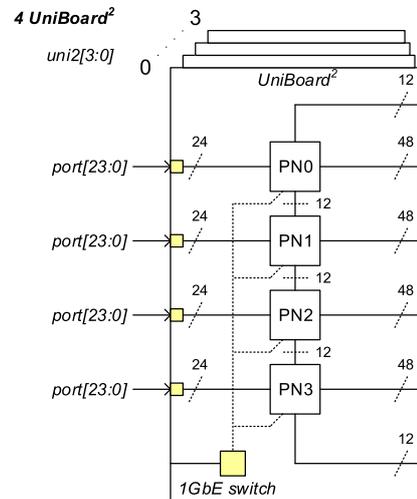


Fig. B.3: Independent PNs on the 4 UNB2s used for the commensal TABF.

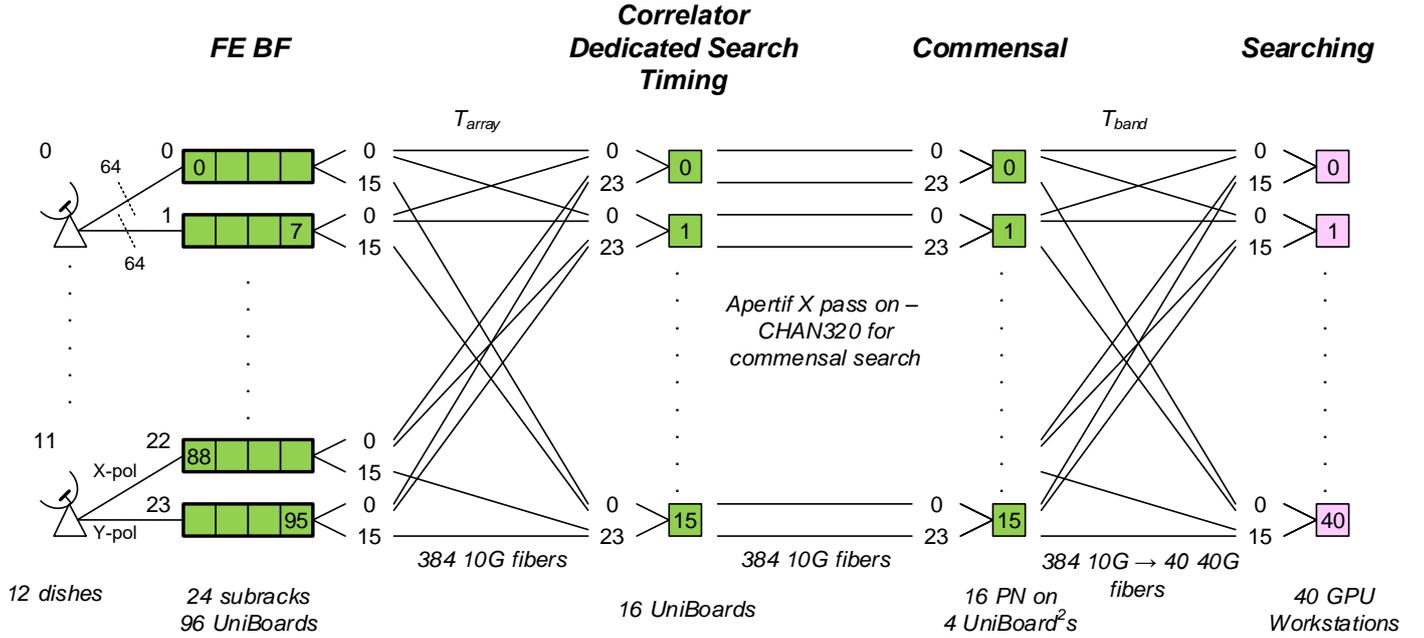


Fig. B.4: Interconnectivity in the Apertif data streams. Data move from left to right, through the Apertif Front-End Beamformer, correlator, ARTS tied-array beamformers, and pipeline processing on the GPU nodes.

#### Appendix B.4: Connecting UniBoard to UNB2

The 128 FPGAs on the 16 UNBs are each connected to the Apertif FEBF via 3 incoming Rx-only 10G links. Outgoing data are carried over the corresponding Tx-only link. The first 10G link transmits the output for the dedicated search. The other 256 ports connect directly from the UNB to the UNB2, as shown in Fig. 5.

#### Appendix B.5: Transpose using COTS 10/40GbE switches

The transpose  $T_{band}$  needs to bring together all 16 bands formed in the TABF, to have the full bandwidth per CB available in a single PL workstation (Fig. B.4). Due to the redistribution on the UNB mesh, the 40 CBs are already grouped in sets of 5. The transposer can thus operate in 8 independent parts. Fig. 5 shows a more detailed diagram for the transposer function. Each of four COTS 10/40GbE switches handles two of the 8 parallel parts. Initial designs required a custom transposer, but the switches were able to handle the uni-directional links produced by UNB and UNB2.

In commensal-search mode, the UNBs operate as correlator, and only pass on the CHAN320 data to the UNB2s. One third necessarily passes through the same 128 10G connections mentioned in Sect. 3.4, and goes through the switches, as indicated by the dashed links Fig. 5. The other two thirds is directly passed on from UniBoard via 256 10G links. By thus routing the UNB data through UNB2 the number of required ports in the transposer switches is minimised. The UNB2 output of 270 Gbps is transposed and sent to the PL. This is a factor 9/12 less than for dedicated search, because commensal mode has 9 TABs.

The Dell S6000 switches can handle the necessarily uni-directional traffic. At the output side, each switches connects to 2x5 workstations using one 40GbE port per workstation, and to the other switches.

### Appendix C: Detailed design III: Firmware and software

#### Appendix C.1: Firmware

##### Appendix C.1.1: Firmware timing

The sync interval in the CB480 data is 1.024 s, which corresponds to 800000 time samples per beamlet. An external PPS defines the start of the first block of ADC data at all FEBFs. The timing information is kept with the streaming data by means of a block sequence number (BSN) that counts subband periods. As the PPS is aligned to the top of second of Coordinated Universal Time (UTC), the BSN can be related to the wall clock time.

##### Appendix C.1.2: Requantisation in the TABF

The CB input is 6 bits wide. Internally the voltage TAB, power TAB and power IAB calculations are done using up to 18 bit. Only the channel filterbank PFB calculations are done using 9 bit FIR filter coefficients, to save multipliers. The FFT part uses 18 bit. Generally multiplications are done using the 18x18 bit multipliers, and 18-bit RAM. At the output the TABs are resized to 8 bit. If overflow occurs, the output value is clipped, preserving sign.

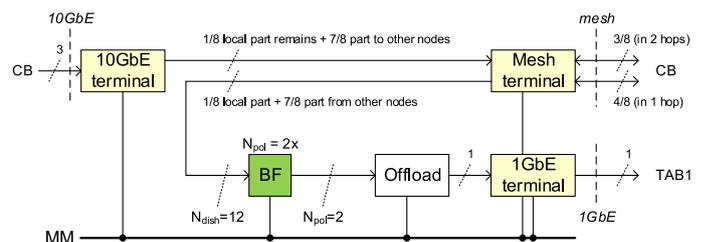


Fig. C.1: Top level block diagram for the single-TAB timing FPGA design on UniBoard.

Appendix C.1.3: TAB1 processing on UNB

The TAB1 beam former (Fig. C.1) is based on the FEBF. For its input, the transpose  $T_{int_x}$  in the FEBF (Fig. A.2) is bypassed. The UniBoard mesh re-distributes the CB480 input over the FPGAs as in Sect. 4.1.1. Through this programmable routing, the FN0 of all 16 UniBoards together receive the full 300 MHz for the central CB. Each of the 16 FN0 output nodes offloaded 600 Mbps of data via the 1GbE control interface of the UniBoards, as shown in Fig. 5.

Appendix C.1.4: IAB40 processing on UNB

The optional IAB beamformer sums the input powers from the 12 input dishes. It operates on subchannels, integrates these, and outputs them as channels. The IAB function itself does not have weights.

Appendix C.1.5: TAB360 processing on UNB2

For the commensal search, the same design is used on all UNB2 FPGAs. The block diagram for the TAB360 processing (40 CBs  $\times$  9 TABs) per 10G output is shown in Fig. C.2, and is very similar to the UNB version (Fig. 6).

Appendix C.2: Monitoring & control software

Appendix C.2.1: Architecture

The Apertif system is built as a set of loosely coupled components that communicate over the network by exchanging messages over a message bus. For this, we employ a messaging middleware layer, which ensures reliable communication.

Appendix C.2.2: Message bus

The message bus consists of queues, exchanges, and brokers. Messages are wrapped in an envelope that contains the address of the recipient, and sent to an exchange. A broker helps forward the message to a destination queue for pick up by the recipient. Messages queuing provides reliable communication, guaranteed message delivery, and is easy to use. Routing is run-time configurable.

Apertif uses the Advanced Message Queuing Protocol (AMQP), which is based on IEEE standard (ISO/IEC 19464). Apache Qpid provides an implementation of AMQP, in both C++ and in Python. The choice for Qpid was driven by experience in LOFAR. We run a broker daemon on every host. A single exchange routes messages to their destination queues. Each component has queues for difference message types.

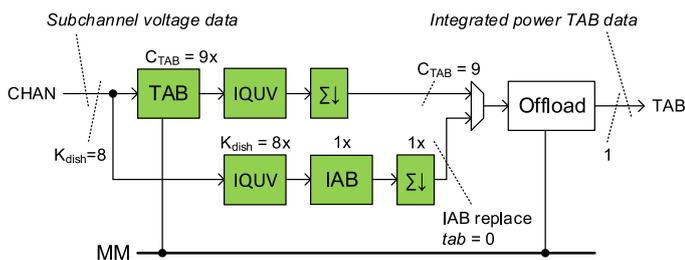


Fig. C.2: Block diagram for the commensal-search TAB360 processing.

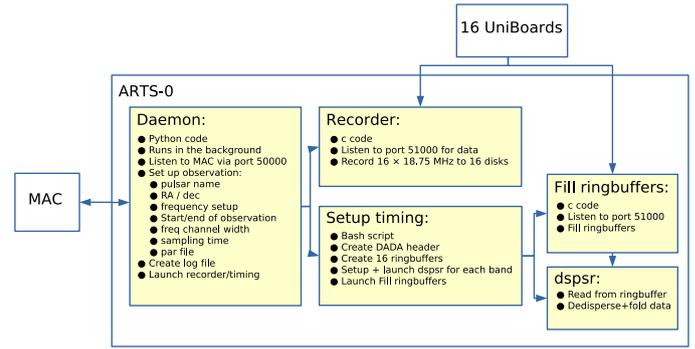


Fig. C.3: Schematic overview of the ARTS timing software pipeline

Appendix C.2.3: Layered design

The Monitoring & Control software follows a layered design, from drivers (lowest), to controllers, to orchestration (highest). The drivers interface with hardware, over Ethernet. All drivers are written in Python although the UniBoard driver is partly in C++ for performance. The controllers, next, handle Message Bus commands and serialise access to the drivers. Some high-level commands are sent to the whole system (e.g. "Start Observation"). Other, low-level commands are subsystem or driver oriented (e.g. "get\_status"). Different applications in the orchestration layer, finally, set up components throughout the system as a whole.

Appendix C.2.4: Software components

Each component has a functional and a management interface (van Cappellen et al. 2022). Both support normal-priority synchronous function calls, and high-priority asynchronous calls. Each component also generates events and notifications.

Appendix C.2.5: Controllers

All basic, shared functions of the controllers are implemented through a controller framework, that has three main tasks. A RequestHandler places new messages from the bus into the command queue. An ExecutionHandler schedules and executes these queued commands. A ResponseHandler monitors the resulting response queue.

Appendix C.3: Science-pipeline software

Appendix C.3.1: Pulsar timing

The components of the pulsar timing pipeline are detailed in Fig. C.3. The arts-0 machine contains two sets of NIC, CPU, and GPU. Each set is connected by a dedicated interconnect, and the data handling and reduction processes are pinned to the appropriate network ports, compute cores, and GPUs. This streamlining was essential for reaching real-time performance without data loss.

Appendix C.3.2: Transient searching

The detailed overview of the ARTS cluster survey controllers described in Sect. 4.2.3 and 5.1.2 are presented in Fig. C.4.

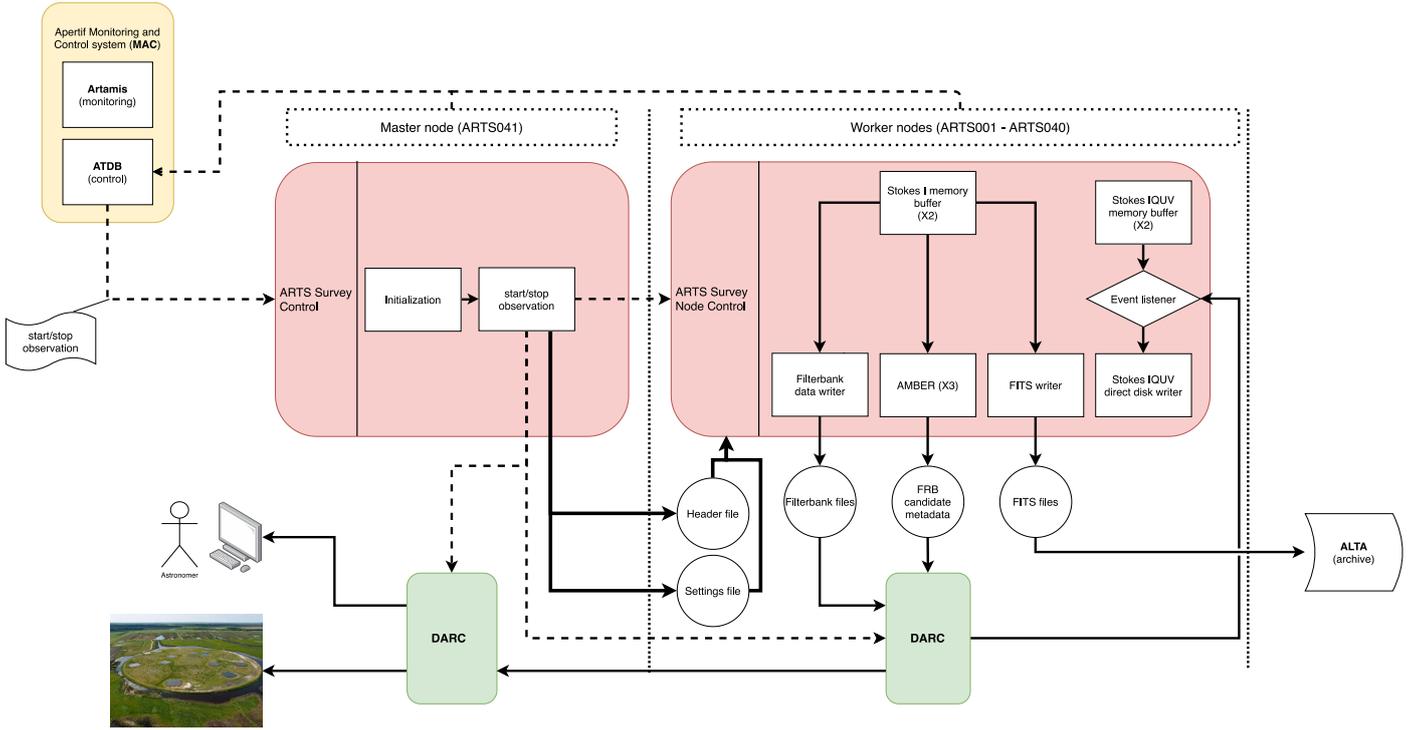


Fig. C.4: The direct interactions of the ARTS specific survey controllers. The controllers are represented by the red boxes. Dashed lines represent system control messages. Communication between the monitoring system (ARTAMIS) and the nodes, as well as any relevant communication between MAC and other components of the system (e.g. ALTA) are omitted for clarity. Since the stokes IQUV data are post processed before being used, the direct disk writer does not show any output.

## Appendix D: Tied-array beam simulations

The hierarchical beam forming scheme used for ARTS was introduced at a high level in Sect. 2.3.2. In this appendix we provide a more quantitative discussion; first on the formation of tied-array beams (TABs); next, on combining TABs to synthesized beams (SBs) and tracking beams (TBs).

### Appendix D.1: Tied array beams

To determine how many TABs are required to cover the compound beam (CB), we assume that the separation between TABs,  $\theta_{\text{sep}}$ , is equal to their half power beam width (HPBW), which depends on the observing wavelength  $\lambda$  and the projected longest baseline  $B_{\text{max}}$  as

$$\theta_{\text{sep}} = \alpha \frac{\lambda}{B_{\text{max}} \cos(\theta_{\text{proj}})}, \quad (\text{D.1})$$

where  $\theta_{\text{proj}}$  is the angle between the plane perpendicular to the linear WSRT configuration and the line of sight. We determine  $\alpha$  empirically, for  $\theta_{\text{proj}} = 0$ . We find that for the 144-m spaced arrays of 8 and 10 dishes (Apertif-8, dishes RT2–RT9,  $B_{\text{max}} = 1008$  m; and Apertif-10, comprising RT2–RTB,  $B_{\text{max}} = 1296$  m),  $\alpha = 0.78$  and  $\alpha = 0.80$  respectively.

If the TABs together cover the angular distance equal to that between a TAB main beam and its first grating response, they provide full coverage of the CB, by a combination of their main beams and first grating response. The grating distance  $\theta_{\text{grat}}$  depends on the observing wavelength and the projected common quotient baseline  $B_{\text{cq}}$  of the regular array as

$$\theta_{\text{grat}} = \arcsin\left(\frac{\lambda}{B_{\text{cq}} \cos(\theta_{\text{proj}})}\right) \approx \frac{\lambda}{B_{\text{cq}} \cos(\theta_{\text{proj}})}, \quad (\text{D.2})$$

where the approximation holds for small angles.

The number of TABs required per CB,  $N_{\text{TAB}}$ , then follows from the ratio of the grating distance and the TAB separation:

$$N_{\text{TAB}} = \frac{\theta_{\text{grat}}}{\theta_{\text{sep}}} = \frac{B_{\text{max}}}{\alpha B_{\text{cq}}}, \quad (\text{D.3})$$

where the wavelength and the projection angle cancel. This implies the number of TABs required depends on neither the projection angle nor the frequency. The number of TABs can be kept unvarying. Using Eq. (D.3), we find that 12 and 9 TABs are needed for Apertif-10 and Apertif-8 respectively.

### Appendix D.2: Synthesized beams

As shown by Eq. (D.2), the grating lobe distance is frequency-dependent. This is illustrated in Fig. D.1, which shows the grating response of the central TAB at 1500 MHz and the TAB left from the central beam at 1367 MHz for Apertif-10. The first grating response right from the centre of these two TABs coincides. This shows that a transient signal received in the first grating response of the central TAB to the right of the centre at 1500 MHz will not be detected in the corresponding grating response at 1367 MHz, but by the corresponding grating response of the first TAB left from the central TAB at 1367 MHz. For all but the main beam of the central TAB, we therefore may have to combine chunks of bandwidth from different TABs to form an SB at a given distance from the CB centre.

In principle, there is an optimal combination of chunks of bandwidth from different TABs for each position within the FoV. In practice, the number of Synthesized Beams (SBs) that we can form is limited. A practical minimum number of SBs can be

set by considering the TAB separation at the highest operating frequency, taking

$$\theta_{\text{SB}} = \theta_{\text{sep}}(f_{\text{max}}) = \alpha \frac{\lambda_{\text{min}}}{B_{\text{max}}}. \quad (\text{D.4})$$

For Apertif-10, we found  $\alpha = 0.80$  and we have  $B_{\text{max}} = 1296$  m. Assuming  $f_{\text{max}} = 1500$  MHz and a field-of-view of 30 arcmin, we find that the number of SBs required is 71. A grid where SBs are spaced at 0.5 HPBW to improve sensitivity across the field-of-view results in 151 SBs. While both modes are possible in ARTS, its standard operation uses 71 SBs, based on compute-load trade-offs throughout the entire search signal chain.

At a given frequency, the optimal contribution to the SB denoted by index  $n_{\text{SB}}$  is coming from a specific grating response denoted with index  $n_{\text{gr}}$  from TAB with index  $n_{\text{TAB}}$ . In the remainder of this section, we determine which grating response from which TAB gives the optimal contribution to a given SB. Since the position shift of the grating responses over frequency increases with distance from the main beam, it is convenient to choose the reference position to be in the centre of the CB. This intuitively leads to an indexing scheme in which 0 denotes the central position, a negative index implies a position left or westwards of the centre and a positive index denotes a position right or eastwards of the centre. For example,  $n_{\text{gr}} = -1$  denotes the first grating response left from the main beam,  $n_{\text{TAB}} = 0$  denotes the central TAB and  $n_{\text{SB}} = 5$  denotes the synthesised beam located five grid points to the right of the centre of the CB.

We assign a grating index and a TAB index at the highest frequency to each SB index. We first determine the closest grating response of the central TAB by  $n_{\text{gr}} = \lceil n_{\text{SB}}/N_{\text{TAB}} \rceil$ , where  $\lceil \cdot \rceil$  denotes rounding upwards towards the next higher integer value. Next, we determine the TAB index such that  $n_{\text{TAB}} = \{-N_{\text{TAB}}/2, -N_{\text{TAB}}/2 + 1, \dots, N_{\text{TAB}}/2 - 2, N_{\text{TAB}}/2 - 1\}$  by calculating

$$n_{\text{TAB}} = n_{\text{SB}} \bmod N_{\text{TAB}} \quad (\text{D.5})$$

and subtracting  $N_{\text{TAB}}$  if  $n_{\text{TAB}} > N_{\text{TAB}}/2 - 1$ .

Figure D.2 shows the resulting grating indices and TAB indices for all SB indices when  $N_{\text{TAB}} = 12$ . Note that the calculation above holds for even  $N_{\text{TAB}}$ . A similar convention can be

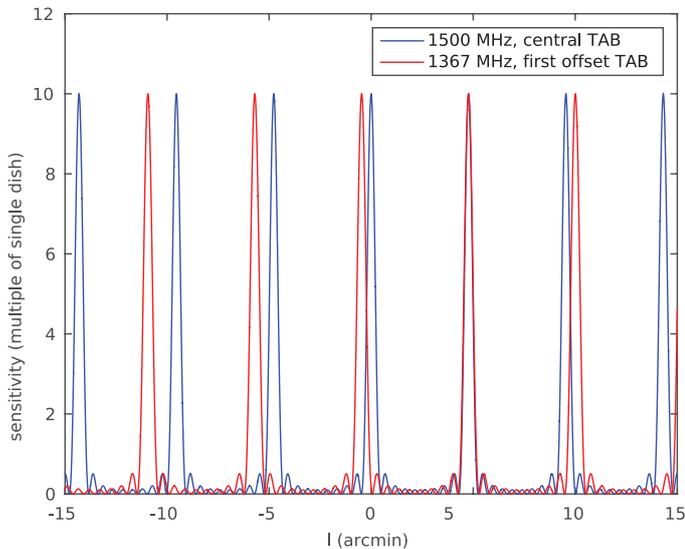


Fig. D.1: Grating response of the central TAB at 1500 MHz and the first TAB left from the centre at 1367 MHz for Apertif-10.

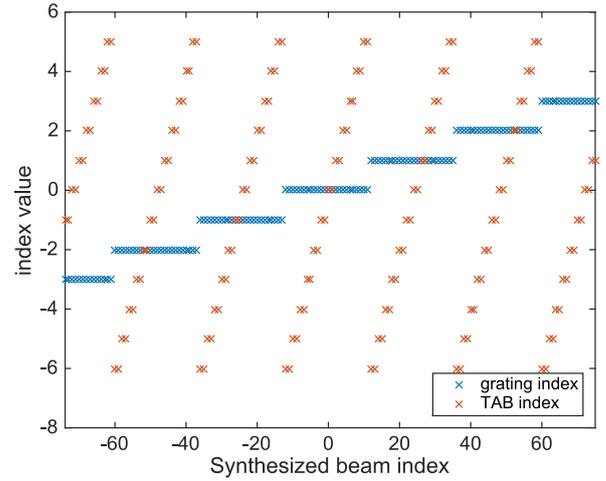


Fig. D.2: Grating index and TAB index for each SB index at the highest frequency.

defined for odd  $N_{\text{TAB}}$ . This is not spelled out here to keep this section concise.

The position of the SB with index  $n_{\text{SB}}$  is described by

$$\theta_{n_{\text{SB}}} = n_{\text{SB}}\theta_{\text{SB}} = n_{\text{gr}}\theta_{\text{grat}}(f_{\text{max}}) + n_{\text{TAB}}\theta_{\text{sep}}(f_{\text{max}}). \quad (\text{D.6})$$

Since  $\theta_{\text{grat}}$  and  $\theta_{\text{sep}}$  are frequency-dependent, at a certain (lower) frequency, the grating response of the next TAB will be at the same position and is therefore the optimal choice at that frequency. To find the frequency at which the position of the grating response of the next TAB at frequency  $f_0$  coincides with the grating response of the original TAB at  $f_{\text{max}}$ , we solve  $f_0$  from

$$n_{\text{gr}}\theta_{\text{grat}}(f_{\text{max}}) + n_{\text{TAB}}\theta_{\text{sep}}(f_{\text{max}}) = n_{\text{gr}}\theta_{\text{grat}}(f_0) + (n_{\text{TAB}} - \text{sgn}(n_{\text{SB}}))\theta_{\text{grat}}(f_0), \quad (\text{D.7})$$

where  $\text{sgn}$  denotes the signum function. Substitution of Eqs. (D.1) and (D.2) while replacing  $\lambda$  by  $c/f$  and taking  $\theta_{\text{proj}} = 0$ , we obtain

$$n_{\text{gr}} \frac{c}{f_{\text{max}} B_{\text{cq}}} + n_{\text{TAB}} \frac{\alpha c}{f_{\text{max}} B_{\text{max}}} = n_{\text{gr}} \frac{c}{f_0 B_{\text{cq}}} + (n_{\text{TAB}} - \text{sgn}(n_{\text{SB}})) \frac{\alpha c}{f_0 B_{\text{max}}}, \quad (\text{D.8})$$

which gives

$$f_0 = f_{\text{max}} \frac{n_{\text{gr}} B_{\text{max}} + (n_{\text{TAB}} - \text{sgn}(n_{\text{SB}})) \alpha B_{\text{cq}}}{n_{\text{gr}} B_{\text{max}} + n_{\text{TAB}} \alpha B_{\text{cq}}}. \quad (\text{D.9})$$

Note that a situation in which  $n_{\text{TAB}} - \text{sgn}(n_{\text{SB}}) \notin \{-N_{\text{TAB}}/2, -N_{\text{TAB}}/2 + 1, \dots, N_{\text{TAB}}/2 - 2, N_{\text{TAB}}/2 - 1\}$  may occur. Since the position of  $n_{\text{TAB}} + N_{\text{TAB}}/2$  for  $n_{\text{gr}}$  coincides with  $n_{\text{TAB}} - N_{\text{TAB}}/2$  for  $n_{\text{gr}} + 1$ , this can be solved by modifying the indices accordingly. A similar procedure can be followed for  $n_{\text{TAB}} - N_{\text{TAB}}/2 - 1$ .

From the discussion above, it is clear that for the  $n_{\text{SB}}$ th SB, for which the chunk of bandwidth close to  $f_{\text{max}}$  is coming from the  $n_{\text{gr}}$ th grating response of the  $n_{\text{TAB}}$ th TAB, the optimal chunk of bandwidth around  $f_0$  is provided by the  $n_{\text{gr}}$ th grating of the  $(n_{\text{TAB}} - \text{sgn}(n_{\text{SB}}))$ th TAB. Since the HPBW changes only slowly with frequency and the shift of the grating responses scales linearly with frequency, the optimal switching frequency will be

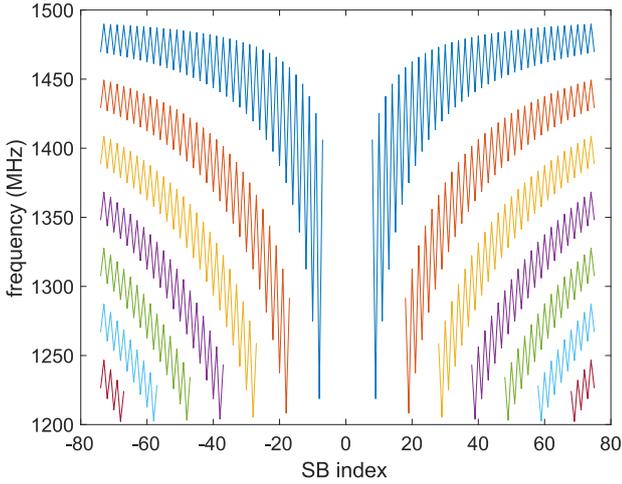


Fig. D.3: Switching frequencies for all SBs of Apertif-10 covering a FoV of 30 arcmin over the frequency range from 1200 to 1500 MHz.

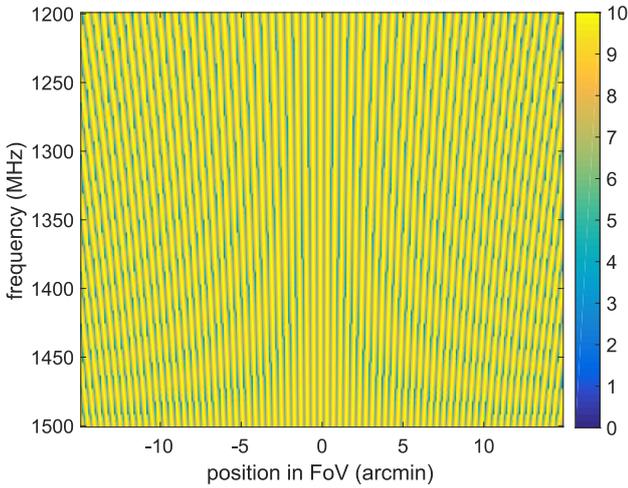


Fig. D.4: Sensitivity within the CB field-of-view as function of frequency when using the switching frequencies shown in Fig. D.3. The sensitivity is expressed in terms of the sensitivity of a single WSRT dish.

approximate halfway  $f_{\max}$  and  $f_0$  if consecutive SBs are separate by the half power beam width while the optimal switching frequency will be at approximately a quarter of this frequency interval if consecutive SBs are separated by half the HPBW. The latter applies to ARTS when using 151 SBs, so we determine the switch frequency as  $f_{\text{sw}} = 0.75f_{\max} + 0.25f_0$  for even SBs, and as  $f_{\text{sw}} = 0.25f_{\max} + 0.75f_0$  for odd SBs. The resulting switching frequencies are shown in Fig. D.3.

Fig. D.4 shows the sensitivity, expressed relative to the sensitivity of a single WSRT dish, within the CB field-of-view as function of frequency when using the switching frequencies calculated above. These sensitivity data indicate that the 151 SBs described above provide an average sensitivity across the full FoV of 81% of the maximum achievable sensitivity of the WSRT array. The latter would require a TAB phase, centred at each individual point within the CB field of view, to be formed, that is,

it would, strictly speaking, require the formation of an infinite number of TABs, which is practically infeasible.

### Appendix D.3: Tracking beams

The grating response of each TAB rotates around the centre of the FoV of the CB during an observation. This is illustrated in Fig. 3e. A given source may therefore traverse multiple SBs during an observation. To track a specific source (or position within the CB) during an observation, we may thus have to concatenate time domain data from multiple SBs. This section describes a procedure to determine which time intervals from which SBs need to be combined to track a desired position within the CB, and thus form a tracking beam (TB).

Fig. 3e illustrates how the TAB gratings rotate through the FoV during an observation. A specific point within the CB thus moves through the TAB gratings. This perspective, where the coordinate system is fixed to the TAB grating response, is shown in Fig. D.5. At a specific reference time  $t = 0$ , a specific locus can be specified by cylindrical coordinates  $(\theta_0; \phi_0)$ , where  $\theta_0$  measures the distance from the field centre and  $\phi_0$  measures the angle between the line from the field centre to the locus and the line parallel to the array, which is the line orthogonal to the grating responses. During an observation, this locus will follow a circular path through the CB with an angular velocity given by  $\omega_E$  as indicated by the red track.

The cross-over points between SBs are indicated in Fig. D.5 by vertical dashed blue lines. The area between two such lines is associated with a specific SB that can be identified by its SB index as described in Sec. D.2. If we define the central SB as index 0, we can find the SB associated with a specific locus at a specific instant  $t$  during an observation by finding out in which SB the point  $\theta_0 \cos(\phi_0 + \omega_E t)$  lies (Fig. D.5):

$$n_{\text{SB}} = \left\lceil \frac{\theta_0 \cos(\phi_0 + \omega_E t)}{\theta_{\text{SB}}} \right\rceil, \quad (\text{D.10})$$

where  $\lceil \cdot \rceil$  denotes rounding.

Each point in the compound beam is covered by the grating response of one of the SBs at  $t = 0$ . Equation (D.10) provides

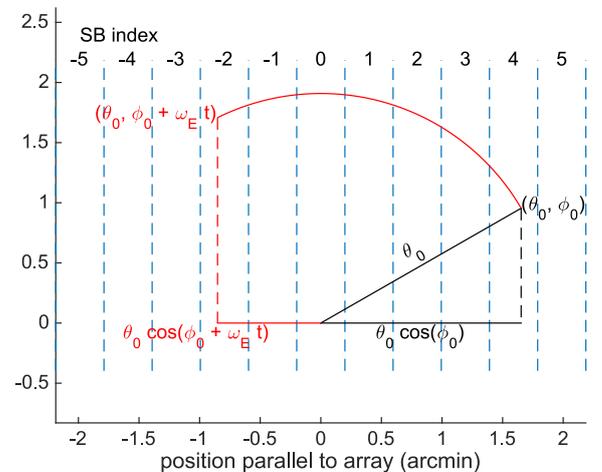


Fig. D.5: Movement of the locus at cylindrical coordinates  $(\theta_0; \phi_0)$  at the start of the observation through the synthesised beams during an observation.

the SB index at each instant  $t$  during the observation. From this, we construct a TB for the full observation for a given point in the CB. The time series associated with these TBs cover the full length of the observation and therefore allow detection of weaker sources than will be feasible to detect during the drift time for an individual SB.

To cover the full field-of-view of the CB, we can define a hexagonally close-packed grid of TBs. As the CBs themselves are arranged in a similar fashion, the ratio between the hexagonal area covered by a TB in a full synthesis observation and the hexagonal area of a CB is then equal to  $(\theta_{\text{SB}}/\theta_{\text{CB}})^2$ . The number of TBs needed to cover a CB will be approximately the inverse,  $(\theta_{\text{CB}}/\theta_{\text{SB}})^2$ . If the maximum number of TABs is used to fill the space between the grating responses of Apertif-10,  $\theta_{\text{SB}} = 0.398$  arcmin. If we assume that  $\theta_{\text{CB}} = 30$  arcmin, we would need to synthesise about 5700 TBs. For shorter observations, this number may be reduced as the TBs will have an elongated shape in the direction perpendicular to the array at the mid-point of the observation, which would allow for a larger separation between TBs along that direction.