Apertif Survey Plan

Apertif Science Team

16 March 2016

version 2.2

This document describes and motivates the surveys planned for Apertif and should be read in conjunction with the publicly available documents "Apertif Draft Surveys Program" and "Apertif Survey Team and Apertif Survey Program". Both are available at
http://www.astron.nl/radio-observatory/apertif-surveys

Executive Summary

The Survey Plan for Apertif is the result of discussions within the Apertif Science Team, in consultation with the international astronomical community. Many of the scientists planning to use SKA1 and its pathfinders have been involved in developing the Apertif Survey Plan. Therefore, the science planned for Apertif is an integral part of the science planned for SKA1 and its pathfinders. A second, important consideration for defining the Apertif surveys has been the legacy value of the Apertif data.

The Apertif surveys will consist of:

- a large-area, shallow imaging survey of HI and polarised radio continuum emission covering ~3500 deg^2;
- a medium-deep imaging survey of HI and polarised radio continuum emission covering ~ 450 deg^2;
- a time-domain survey for pulsars and fast transients over 15,000 deg^2;
- in addition to these large surveys, a small amount of observing time will be devoted to imaging a number of small survey regions (typically 10 deg^2 in size) coincident with fields from the Tier 2 of the LOFAR imaging continuum survey.

The main science topics addressed by the Apertif Surveys are:

- use resolved HI observations to investigate the connections of the gas- and total mass distribution of galaxies over a wide range of galaxy properties as obtained from other wavebands
- investigate how these relations vary with environment and what role interactions, gas accretion and removal of gas play;
- investigate the properties of the smallest gas-rich galaxies in the local Universe and confront these with theory;
- study the role of cold gas in AGN and their feedback activity;
- study the history of star formation and AGN activity of the faint radio continuum population;
- study the magnetic fields in galaxies and of the large-scale structures in which they are embedded;
- the detection, characterisation and arcsec-localisation of fast radio bursts (FRBs) over the Northern sky;
- perform a census of intermittent and normal pulsars along the Galactic Plane;
- study the (orphan) afterglows of Gamma Ray Bursts and Tidal Disruption Events.

The strengths of the Apertif Surveys, compared to other planned radio surveys will be:

- the number of *resolved* HI detections in the Apertif HI surveys will be higher than in any other existing or planned large-area HI survey in the pre-SKA1 era. Therefore the Apertif HI surveys are best suited to study the *resolved* HI properties of galaxies, and of their environment. Significant overlap exists with surveys at other wavebands, in particular SDSS, HetDex, Herschel-Atlas, ManGa and later WEAVE. Given the currently available HI survey data, studying the *resolved* HI properties will have important impact on studies of the role of gas in galaxy evolution;

- the legacy value of the Apertif surveys extends well beyond the life span of Apertif; well into the SKA1 era, no instruments will be available to perform competitive surveys in L-band of the northern hemisphere;

- the Apertif Medium-Deep HI survey is the only planned blind HI survey designed to image and resolve the faint HI emission in the outer regions of galaxies, and beyond. Therefore it is the HI survey best able to study the role of environment on galaxy evolution which is generally considered the key factor in galaxy evolution. It is also the only planned HI survey capable of detecting the smallest galaxies in the Universe, a crucial science theme within galaxy formation and evolution;

- Apertif samples a volume of the Universe ten times larger than the Parkes Multibeam, the most successful time-domain transients experiment so-far. The Apertif-LOFAR localisation of FRBs, however, is the biggest potential game changer. The 10^6-fold increase in localisation is unique to Apertif, and likely unsurpassed in the next decade;

- With its long integrations, Apertif is much more sensitive to intermittent pulsars than, e.g., the state-of-the-art GBNCC survey, with its 2-minute integrations;

- because of the match between sensitivity and resolution of LOFAR and Apertif for continuum surveys, the synergy between these two instruments is unique among the SKA pathfinders;

- given the spatial resolution of 15 arcsec of the planned Apertif continuum surveys, they will provide a very useful complement to the other main continuum survey planned for the northern sky, the VLASS, which is done at much higher resolution. Only the combination of the two surveys will give a complete picture of the radio continuum source population at GHz frequencies.

1. Introduction to the Apertif Surveys

1.1. Process

The starting point of the general science case for Apertif is described in the Apropos NWO-Groot proposal. Following the approval of this proposal, the Apertif PIs have engaged the international astronomical community in further defining the Apertif science by issuing a Call for Expressions of Interest (EoI) for Apertif surveys. The response to this call was very positive: eighteen teams submitted an EoI on a large range of science topics. The total amount of proposed Apertif observing time by the EoI teams is 20 years, the equivalent (due to increased survey speed of Apertif) of more than 700 years of observing with the 'old' WSRT. The analysis of the EOIs and how their science can be consolidated in and served by the Apertif Surveys is discussed in the *Draft Survey Plan* document.

Many of the teams that submitted an EoI for Apertif also have ambitions to participate in SKA1 and its pathfinders. Therefore, there is a very good match between the proposed Apertif science and that planned for SKA1. There has been a significant amount of coordination, through bodies such as PHISCC and SPARCS, between the various survey teams to optimise the scientific return of all surveys on all pathfinders.

In order to make the most efficient use of the limited Apertif lifetime, ASTRON management and the Apertif PIs have made the decision that the Apertif surveys should consist of a limited number of large, *commensal* surveys with high legacy value, executed by a single Apertif Science Team (AST). In this AST, all groups participating in the Apertif surveys are represented. The scientific motivation for the limited number of surveys is outlined in the Apertif *Draft Survey Plan* document which was sent for input to the community in 2015. At the same time, the community was invited to participate in the definition and the execution of the Apertif surveys and to participate in the AST. On December 7 and 8, 2015, groups that had responded to this invitation met in Dwingeloo to discuss the survey plan and the associated data products and publication policies. The first task of the AST was, based on the Draft Survey Plan sent to the community, to define the final layout of the Apertif surveys with a particular emphasis on both commensality and legacy value (i.e. also serving science themes not represented at the meeting). The result of this meeting is presented in this document.

1.2. Surveys Layout

The result of the survey definition process is that almost all science themes identified in the EoI process, and later by the AST, can be served by carrying out the surveys listed below. These have been scaled to fit in a 4-year period; after reserving observing time for VLBI, pulsar timing, and Directors Discretionary Time observations, 1600 x 12 hours of total survey time is expected. The resulting surveys are summarised here:

- a shallow, large-area imaging survey covering ~3500 deg^2, for emission and absorption HI studies as well as for continuum and polarisation work. Each pointing of this survey will be observed with a full synthesis observation of 12 h (including overhead and calibration). The spectral line noise level will be 0.65 mJy/beam over 20 km/s, giving a column density sensitivity of ~2.5×10^{20} cm^{-2} (4σ) at full resolution, or 8×10^{19} cm^{-2} at 30 arcsec resolution. The noise in the continuum images will be ~15 μJy/beam.
- a subregion of the shallow survey, covering ~450 deg^2, will be observed for 7x12 h. This medium-deep survey is a factor 2.6 deeper than the shallow imaging survey, leading to a spectral line noise of 0.25 mJy/beam over 20 km/s and column density sensitivities of ~1.0×10^{20} cm^{-2} (4σ) at full resolution and 3×10^{19} cm^{-2} at 30 arcsec resolution. The continuum noise will be ~6 μJy/beam which is very close to the confusion limit of the WSRT. This survey also serves both HI and continuum/polarisation studies. About 10% of the shallow survey area will be contributed by this medium deep survey.
- a dedicated wide-field pulsar and fast-transient survey, at both full sensitivity and high time resolution. By producing 450 tied-array "grating lobes" to cover the full field, this survey is sensitive to, e.g., all 1-

millisecond-duration fast radio bursts (FRBs) down to a flux density of 0.16 Jy. Through 3-h integrations over a 15,000 deg^2 field, dozens of FRBs will be detected. A similar number may be detected in FRB searches commensal to all Imaging Surveys. In both cases, real-time triggering of LOFAR and other telescopes then provides the localisation that has so-far been elusive.

- A small amount of observing time will be used to observe approximately 10 small survey regions (typically 10 deg^2 in size) coincident with special regions from the LOFAR imaging continuum survey (Lockman Hole, Boötes field, Elias-N1-Swire, Herschel Atlas field, NEP, EGS-Scuba2). The depth of these observations is 4x12h and intermediate between the shallow and medium-deep survey and will give continuum images which are (close to) confusion limited. The amount of observing time for this Lofar field is very limited compared to that of the other Apertif surveys.

In the following we describe these surveys in detail (Sect. 2 and 3). Additionally, we describe the data processing and pipe lines (Sect. 4), the survey operations (Sect. 5), data products and data release (Sect. 6), and conclude with the science exploitation of the surveys (Sect. 7).

2. The Apertif Imaging Surveys

All imaging surveys will be done in full polarisation, covering the frequency range 1130-1430 MHz with a spectral resolution of 12.2 kHz (corresponding to an HI velocity resolution of 2.6 km/s at z = 0 and 3.2 km/s at z = 0.256 (1130 MHz)). The highest spatial resolution images will have a resolution of 15x15/sin(δ) arcsec (but tapered data products such as 30 arcsec images, will also be made). The indicative layout of the survey regions is illustrated in Fig. 1 and is further discussed below. This layout can still be adjusted depending on the actual performance of Apertif.

The total observing time for the imaging surveys is 1100 x 12h, in accordance with the *Apertif Draft Surveys Program* and *Apertif Survey Team and Apertif Survey Program* documents.

An important additional aspect of the Apertif HI surveys is that, even well into the SKA era, no large-area high-resolution HI survey of the northern sky will be done, simply due to the fact that sufficient observing time to perform such a large-area survey will not be available on any other instrument. Therefore the Apertif surveys, in particular because of their overlap with surveys in other wavebands, will have a legacy value which extends well beyond the life span of Apertif.

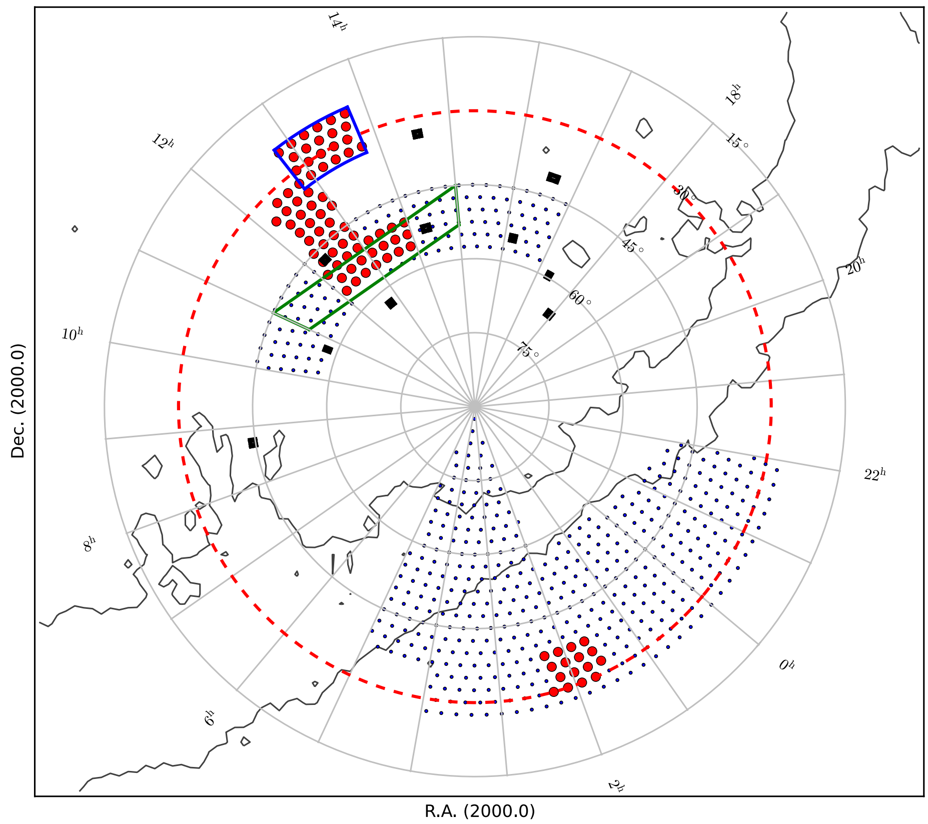

Figure 1. Layout of the Apertif Surveys. The red dots indicate pointings of the Medium-Deep Survey, while the small black dots indicate those of the Shallow Survey. The black squares give the locations of the special LOFAR fields that will also be imaged by Apertif.
The red, dashed line indicates δ = 30˚. The blue rectangle shows the location of the Herschel-Atlas field, while the green box indicated the rough outline of the HetDex survey area. The grey, curly contour indicates the locations where the Galactic foreground extinction is 1 magnitude in the *V* band.

2.1. Motivation of the selection of survey areas

The selection of the survey areas of the imaging surveys is, of course, to a large extent based on science priorities. However, an important boundary condition is also that we should be able to make an efficient observing schedule for the initially secured 4-year period of the Apertif surveys. This boundary condition implies, given that a standard Apertif observation takes 12 h, that the combined survey regions should be spread over the entire sky, in terms of Right Ascension, as much as possible. Given the fact that auxiliary data are not uniformly spread over the sky, this implies, to some extent, a compromise and is the reason that part (i.e. $22^h < \alpha < 6^h$ and $\delta > 35^\circ$) of the Shallow Survey is located outside the region covered by the SDSS. With future surveys with e.g. WEAVE, this situation can be partly remedied for spectroscopy, while Pan-STARRS will cover this part of the sky as well.

The selection of the survey area of the Shallow Survey is motivated by:

- overlap with the SDSS, ManGa, and HetDex surveys. This refers to all pointings with $10^h < \alpha < 17^h$ .We are currently coordinating with the ManGa survey team to optimise the overlap with that survey. SDSS data are also available for the strip with $22^h < \alpha < 2^h$ and $\delta < 35^\circ$ which covers part of the Pisces-Perseus supercluster. Overlap with CALIFA and S4G will also be taken into account, as well as probing galaxy populations at different environmental densities.
- observe the Zone of Avoidance for $2^h < \alpha < 5^h$. This is the region where a very large-scale structure in the galaxy distribution crosses the Galactic Plane.

The Medium-Deep survey areas are selected to cover a number of structures in the galaxy distribution in order to study the role of the environment on galaxy properties:

- cover the Herschel-Atlas area, including the Coma cluster.
- cover a large part of the HetDex survey area. Coordination in survey area and science with the HetDex and JINGLE (JCMT CO & SCUBA-2) survey teams is ongoing.
- cover the Super-Galactic Plane near the North Galactic Pole. This region contains a filament in the large-scale galaxy distribution extending from about 2 Mpc to beyond 12 Mpc.
- cover part of the Pisces-Perseus supercluster.

2.2. Comparison with other GHz surveys

2.2.1. Shallow large-area HI survey

The shallow survey will detect about 90,000 galaxies out to a redshift of z ~ 0.1, with the peak of the redshift distribution around z ~ 0.04 (v = 12,000 km/s). The noise level of the shallow survey is about a factor 2 lower than that of the ALFALFA HI survey and will cover a larger redshift range (out to z = 0.1 vs. z = 0.06). However, the more important difference is that the spatial resolution is about a factor 10 better. This implies that the next big step in understanding the HI properties of galaxies can be made by studying the *resolved* HI properties such as internal kinematics and morphology in relation to other galaxy properties as well using the properties of the fainter HI in the outer regions of galaxies, and beyond, to study the role of interaction and the environment. A number of projects have used the WSRT in recent years for doing surveys to successfully explore exactly this issue (e.g. BUDHIES, ATLAS3D, BlueDisk, Coma, Perseus-Pisces cluster), but the shallow survey will increase the number of galaxies available for study by more than two orders of magnitude. Figure 2 shows an estimate of the fraction of galaxies of which the bright, inner HI disk will be resolved by the shallow survey as function of redshift. This figure gives a lower limit, because the often irregular features in the outer regions of galaxies have not been taken into account for making this figure. An illustration, taken from the BlueDisk survey, of what a typical HI detection at the typical distance of the shallow HI survey will look like is given in Fig. 3, underlining the ability to study the resolved HI properties in galaxies.

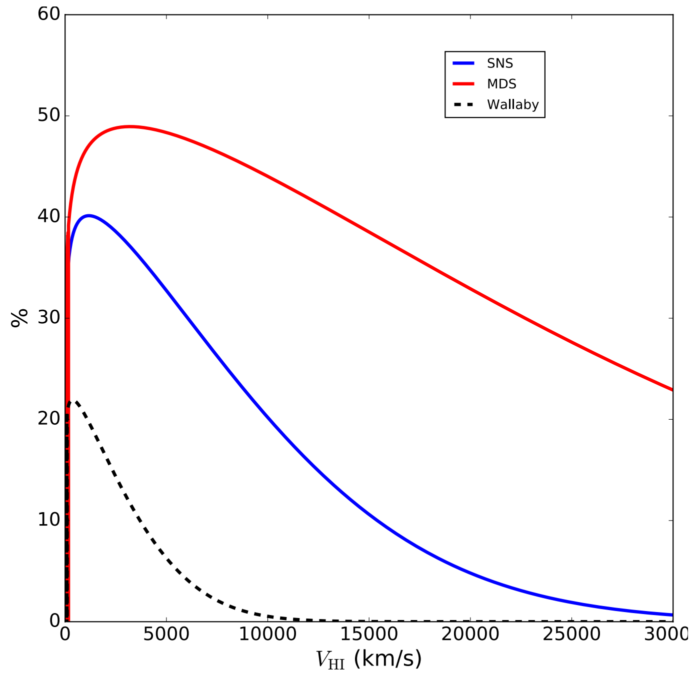

Figure 2. Fraction of galaxies of which the bright HI disk will be resolved by the shallow HI survey (blue curve) and by the Medium-Deep Survey (red curve), as function of redshift. These estimates do not take into account any fainter, extended HI beyond the regular HI disk and therefore are lower limits. For comparison, the estimates for the shallow ASKAP HI survey (Wallaby) are also given (black dashed curve).

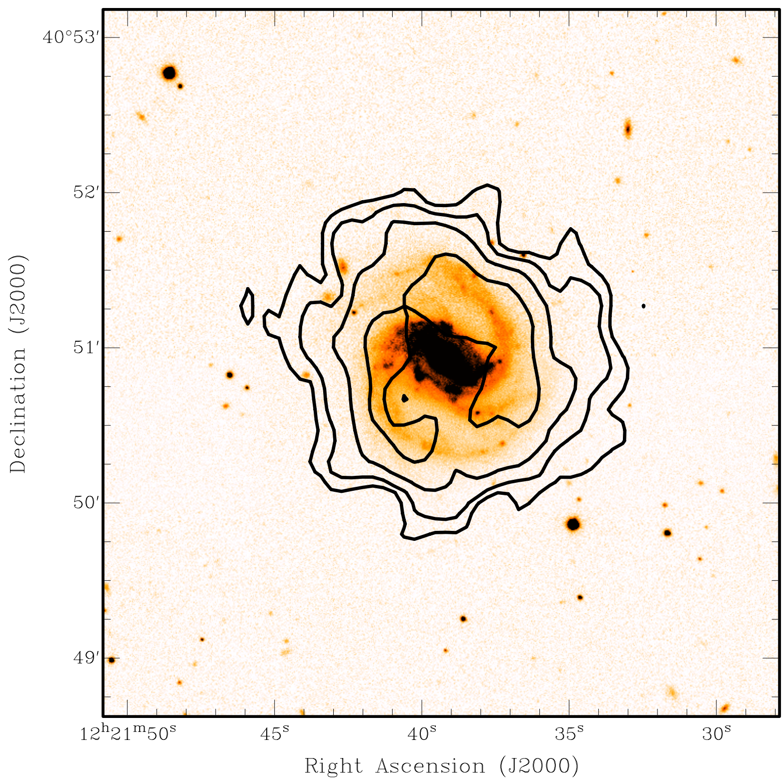

Figure 3. Illustration of what a detection will look like at the typical distance of a galaxy in the shallow HI survey on top of an optical image taken from the Sloan Digital Sky Survey. The HI data were taken with the "old" WSRT as part of a project to prepare for the Shallow Apertif HI survey (Wang et al. 2013, MNRAS, 433, 270). The angular resolution of the HI image is 30 arcsec. Contour levels are 8, 16,32, 64 10^{19} cm^{-2}.

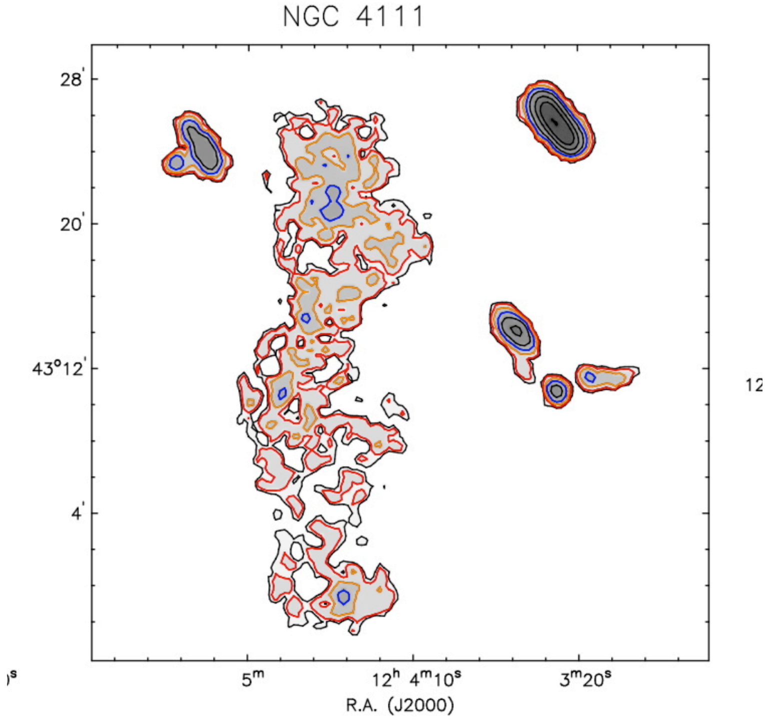

Figure 4. HI image of a large intra-galactic HI cloud in a group of galaxies, illustrating one of the type of features the Medium-Deep Survey will reveal in large numbers. This cloud shows that interactions between the galaxies have occurred which is relevant information for studying how the galaxies evolved. Contour levels are 2 (red), 5 (orange) and 10 (blue) x 10^{19} cm^{-2}. The faintest contour corresponds to the detection level of the Medium-Deep survey.

The number of detection in the ASKAP all-sky HI survey (Wallaby) will be about a factor 2 larger than of the Apertif shallow HI survey, however there are a number of important differences between the two surveys. One is a different choice of survey strategy: Wallaby aims to cover the entire observable sky while for the Apertif shallow survey we have opted to limit the survey area but have sufficient depth to detect galaxies out to larger distances. The highest redshift of detections in the ASKAP shallow survey will be around z ~ 0.05, while for the Apertif Shallow Survey it will be around z ~ 0.1.
Another important difference is that a significant resolution advantage exists over Wallaby. which will have a resolution of 30 arcsec, compared to the about 15 arcsec of the Apertif surveys. The Apertif shallow survey will be able to image the same column density at a resolution twice as high as that of Wallaby. This has the important consequence that the fraction of resolved detections of the Apertif surveys is much higher (see Fig. 2).

2.2.2. Medium-Deep HI survey

The Medium-Deep survey will detect at least 30,000 galaxies of which a large fraction will be resolved (see Fig. 2). The predicted number of detections is somewhat uncertain because it is based on the assumption of a uniform spatial distribution of galaxies while the survey volume is relatively small and contains significant cosmic structure. The main technical aim of this survey is to reach better column density sensitivities which will allow more extensive imaging of the faint HI emission in the outer regions of galaxies, and in between galaxies, of a large sample of galaxies in a range of environmental settings, over a large contiguous area allowing an unbiased view. The type of studies planned for using the Medium Deep survey are not possible with any other HI survey planned for the SKA Pathfinders. This will enable extensive studies of the role of the environment on galaxy properties. The Medium-Deep HI Survey is also the only planned HI survey which can be used to search for the smallest gas-rich companions of large galaxies in the local Universe over large volumes. Galaxies such as the Local Group member Leo T (at a distance of 0.4 Mpc and an HI mass of 3×10^5 Msun one of the smallest gas-rich galaxies known) can be detected out to distances of 5 Mpc and resolved out to 4 Mpc. The predicted number of detections for small gas-rich galaxies with HI mass below 10^7 Msun ranges from a few hundred to just below one thousand. This number is rather uncertain because the HI mass function is not well know for the relevant mass range. Furthermore, also here, due to the small survey volume, the uncertain role of cosmic structure plays a role. The Medium-Deep HI survey may possibly have the highest scientific impact of all proposed PHISCC HI surveys. Figure 4 illustrates what a typical detection of intra-galactic gas will look like and show what type of HI science the Medium-Deep survey will focus on.

As one of the outcomes of the SKA Key Science workshop held in Stockholm (August 2015), the SKA community is now considering to perform a generic commensal all-sky survey to be performed with SKA1 Mid using a total observing time of 8000 h. Figure 5 (kindly provided by Robert Braun) shows the expected HI detection rates, as function of spatial resolution of the generic SKA survey compared to what can be expected for the Apertif Medium-Deep Survey. Obviously, the total number of detection in the SKA1survey is much larger than in the Medium-Deep survey, but the parameter space covered, in terms of redshift, column density sensitivity and resolution, is very similar. Hence, the Medium-Deep survey will allow addressing SKA1-type science, well before SKA1 becomes operational.

2.2.3. Continuum surveys

An important strength of the Apertif continuum surveys compared to those planned for the southern hemisphere is the synergy with the LOFAR surveys. This will remain an important advantage until SKA1-Mid and SKA1-Low will become available. There is a very good match, both in sensitivity (see Fig. 6) and in spatial resolution, between LOFAR and Apertif. This implies that most sources will be imaged by both

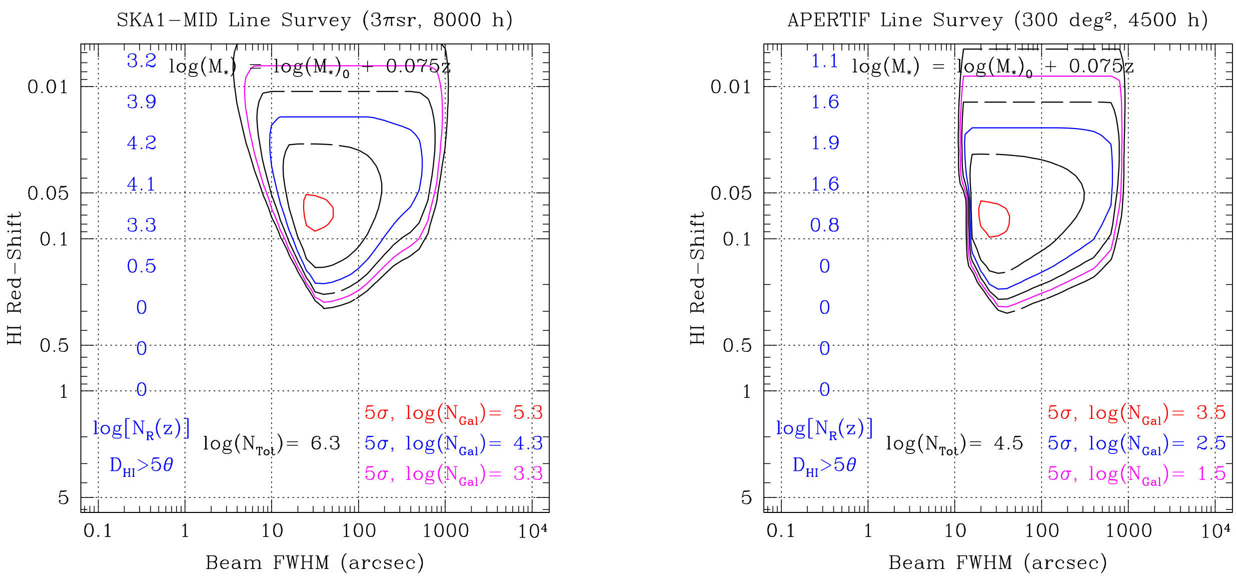

Figure 5. Predicted detection rates, as function of spatial resolution, for an all-sky survey considered for SKA1 and the Medium-Deep Apertif survey. The HI galaxy detection density (per 0.3 dex of red-shift) is plotted as function of angular resolution. The number density of well-resolved detections is indicated along the left, and along the right and the total number of detections at the bottom. Apart from the much larger number of detections expected for the SKA1 survey, because of the much larger survey area, the coverage of parameter space of the two surveys is very similar. Figures courtesy Robert Braun

instruments with similar signal-to-noise and resolution, spanning a factor 10 in frequency range. This wide frequency coverage is extremely valuable for understanding the nature of the detected sources in both the LOFAR and the Apertif surveys. The Apertif data are therefore of great added value.

All radio sources from the planned All-Sky LOFAR 120 MHz survey (depth 0.1 mJy, 5 sigma) with a spectral index flatter than -0.9 will be detected in the shallow Apertif survey. Only sources with a flux density near the LOFAR 30 MHz survey detection limit (10 mJy, 5σ) and an extremely steep radio spectrum ($\alpha < -1.5$) will remain undetected. This indicates that for > 99% of all the sources detected at 30 MHz, flux density measurements at 60, 120 and 1400 MHz will be obtained. The resulting radio colour-colour diagrams will be a powerful tool to spectrally discriminate between radio sources with extreme radio spectra such as diffuse emission from clusters, very distant radio galaxies and GPS/CSS sources. For nearby resolved sources it will instantly yield spectral index and spectral curvature maps, a very rich source of information to constrain many physical parameters and processes.

To provide similarly useful data for the deeper Tier-2 LOFAR fields, these fields will be observed with Apertif for 4x12h which will give images with a depth very close to the confusion limit. The amount of observing time for this is limited compared to other imaging surveys as well as the transient survey.

Plans are underway to undertake a large-area continuum survey, in the 2-4 GHz band, of the northern sky using the JVLA B-array (VLASS). This survey will cover all the sky observable from the JVLA, have a spatial resolution of 2.5 arcsec and a noise level of 70 µJy/beam. The Apertif continuum surveys will be a factor 5-10 more sensitive than the VLASS. In addition, it is well known that observations with the JVLA B array, such as the FIRST survey, are not well suited to image extended sources and this issue will plague the VLASS as well. For this reason, NRAO has in the past performed *two* large-area continuum surveys to complement each other: the low-resolution NVSS (45 arcsec, D-array) and the high-resolution FIRST survey (5 arcsec, B-array). No plans exist to use the JVLA to perform a low-resolution counterpart of the VLASS to remedy the issue of insufficient imaging of extended sources. However, given the properties of the Apertif

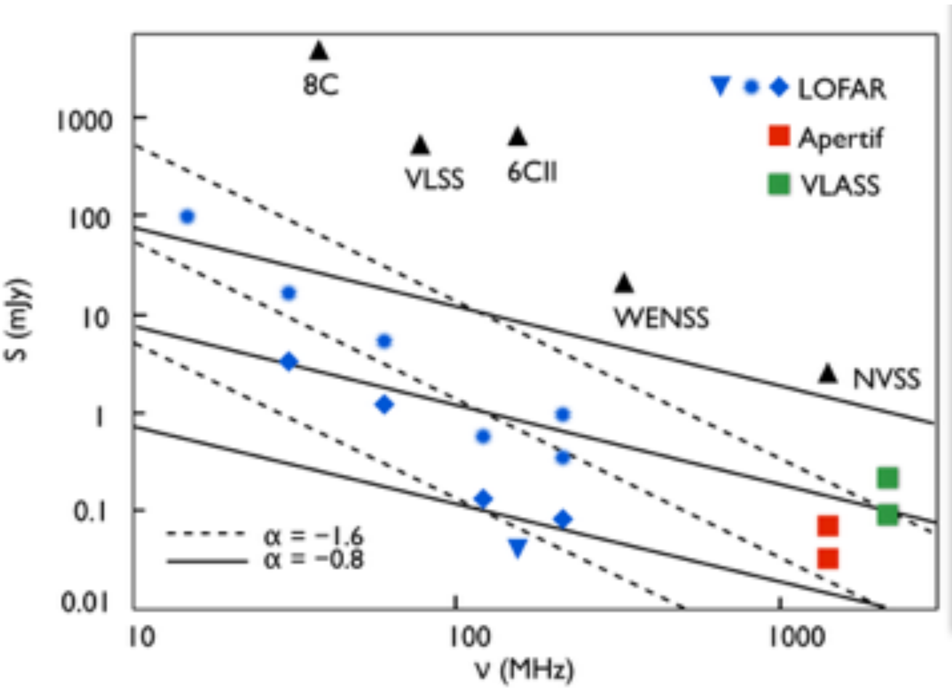

Figure 6. Comparison of the various existing and planned continuum surveys. The lines illustrate the slope of the spectrum of typical radio sources

continuum surveys, they are very well suited to be used for studying the extended emission of the sources detected at high resolution in the VLASS. Therefore the Apertif and JVLA surveys complement each other in a very useful way. The high resolution of the VLASS will also help with the identification of the continuum sources detected with Apertif. Similarly, identifications can be made with the higher resolution of LOFAR as well.

3. The Pulsar and Transient Survey (PTS "ALERT")

3.1. Aims and overview

Fast, millisecond transients appear all over the radio sky. These point to extreme energies, magnetic field strengths, gravitational fields, and densities. The physical conditions on display in these sources must far exceed any terrestrial laboratory. Some of these short-lived radio bursts can be traced to one-off flashes from pulsars but a handful of other bursts (the "FRBs"; Lorimer et al. 2007) clearly originate far outside our Galaxy, and defy any explanation. The first aim of the Apertif pulsar and transient survey is to enable major progress in understanding the astrophysics that powers these brief, extremely energetic extragalactic event. This survey, hereafter called "ALERT", the Apertif Legacy Exploration of the Radio Transient Sky, increases by a factor of 10 our exploration volume for such fast transients throughout the Universe. ALERT can detect tens of extragalactic bursts in real-time, and trigger follow-up ranging from e.g. LOFAR, for exceedingly accurate localization; up to optical and x-ray observatories. The second aim of this Apertif pulsar and transient survey is to characterize, to much lower luminosity than before, the population of intermittently active neutron stars, and better understand the Galactic neutron-star population.

All data that is recorded for imaging observations described in Ch. 2 will be searched for pulsars and FRBs in real-time. That is done in parallel to the Uniboard correlator, using a separate, smaller Uniboard^2-based beam former. This takes in the data of the 8 dishes RT2..9 that are equidistantly spaced in the MaxiShort imaging, and makes 37 x 7 = ~250 tied-array beams. This method (marked "250" in Fig. 7) provides somewhat limited sensitivity, but is less compute-demanding and thus easier to implement for commensal observing. This makes Apertif sensitive to rare, bright radio bursts.

Slow, image-plane transients, such as signs of an early afterglow, can next be identified if related to, e.g., FRBs. For FRB detections in the commensal survey, the unique simultaneous beam-formed and imaging data will allow for the immediate investigation of any slow transient. Beyond that, we will make use of the 7 repeat visits of the MDS (Sect. 2.2.2) fields to search for image-plane transients on a cadence of months to years, with emphasis on a 3-month cadence, which is the peak timescale at 1.4 GHz for the afterglows for Gamma

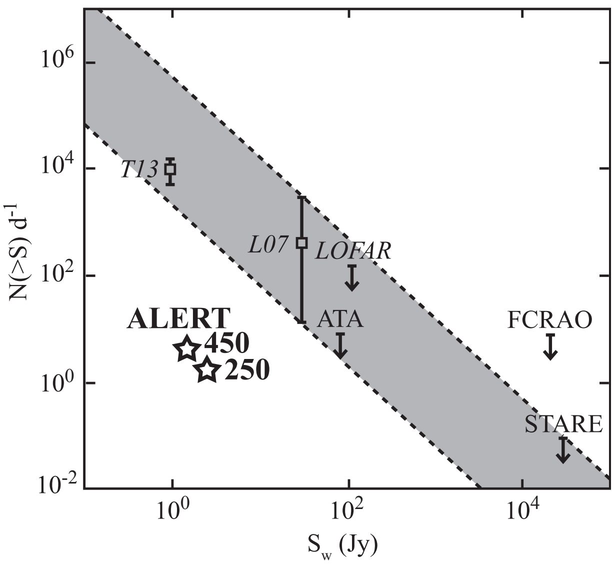

Figure 7. The limits on transient occurrence per day N versus flux S comparing previous fast transient surveys, with ALERT. The area between the dashed lines represents a homogeneous, stationary cosmological population. The dedicated (450-beam) and commensal (250-beam) Apertif surveys are sensitive to both dim and rare events, respectively. Compared to Parkes HTRU (Thornton et al. 2013, labelled "T13"), ALERT provides two orders of magnitude better sensitivity.

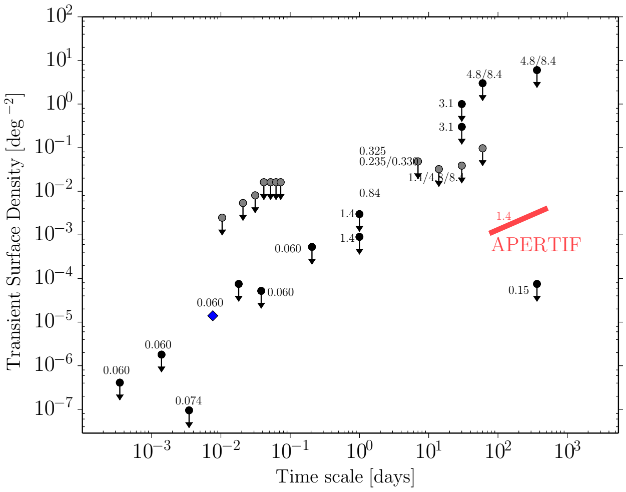

Figure 8. The slow transient surface density against time-scale for various surveys at the indicated frequencies in units of GHz. The Apertif image-transient survey (red line) goes almost a factor of 10 deeper in the exact time scale of interest at 1.4 Ghz, or order 3 months. Figure adapted from Carbone et al. 2016.

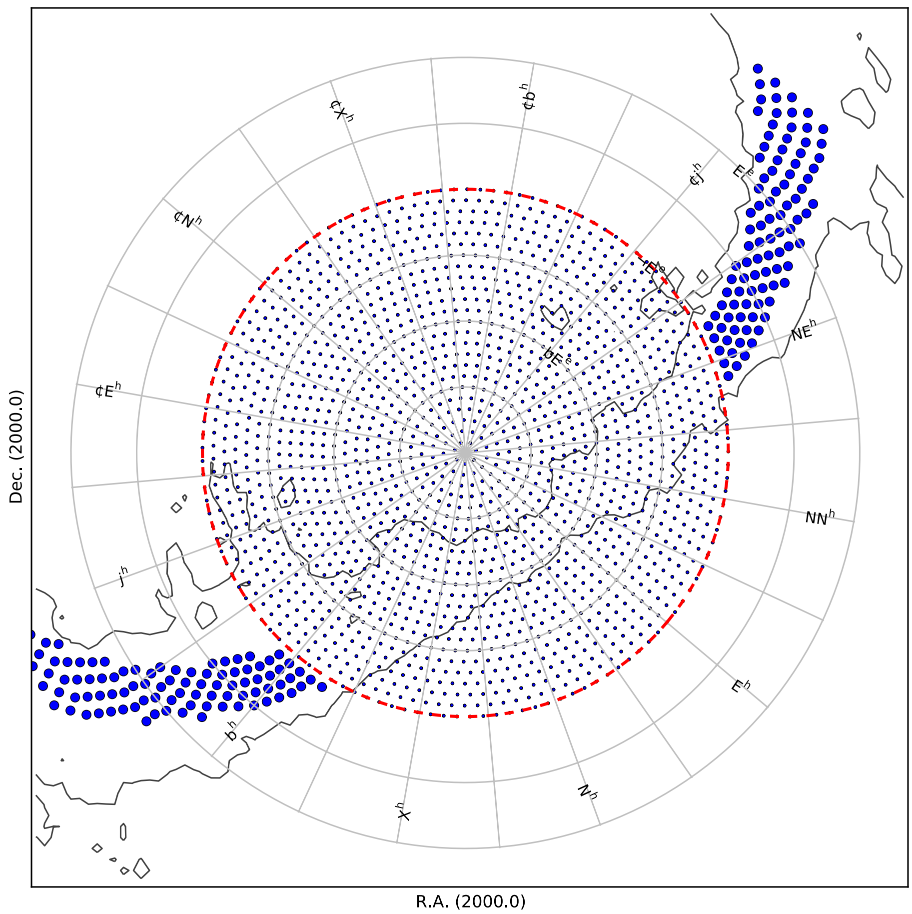

Figure 9. The pointing distribution for ALERT. There is uniform coverage of the region North of declination 30°, plus a deeper extension along the Galactic Plane.

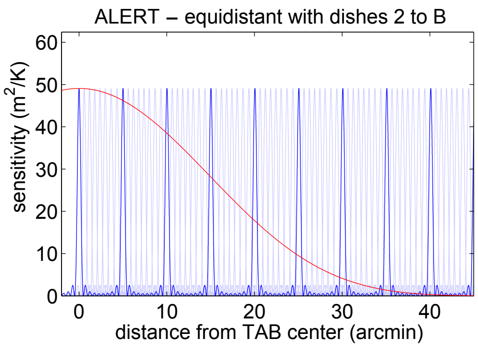

Figure 10. The sensitivity and tied-array layout for the dedicated survey. Within each of the 37 compound beams, 12 tied-array beams (TABs) are formed (blue comb-like responses). Together these fill out the primary beam (red).

Ray Bursts and Tidal Disruption Events. As illustrated Fig. 8, Apertif is uniquely sensitive to transients at that timescale and rate.

Back in the time domain, finding dim bursts from pulsars and FRBs requires combining the Apertif telescopes in the full coherent addition mode, setting the 10 telescopes RT2..B at 144-m separations. Now, the 37 x 12 = ~450 required beams necessitate the use of the 16 Uniboards that otherwise run the correlator. This, and the limited overlap in pointing region with the shallow imaging survey, require a stand-alone ALERT survey. That survey, marked "450" in Fig 7, is exquisitely sensitive, over two orders of magnitude better than the currently operating competition. We will carry out such a survey for 1600/3 x 12 hrs = 6400 hrs.

3.2. Survey Motivation and Layout

The aim of the survey is to maximise the number of detections of FRBs and pulsars and to best current knowledge FRBs are distributed isotropically on the sky. For these a uniform coverage of the Northern sky is most interesting, as the expected number of detected FRBs may show whether e.g., the current tentative preference for FRBs outside the plane continues to hold.

The density of pulsars is highest along the Galactic plane, especially toward the Galactic centre, although nearby sources, such as millisecond pulsars, appear more isotropically distributed on the sky.

Overall this leads to a combination of an all-Northern sky survey, north of declination 30°, starting roughly North of the well-surveyed Arecibo range stand-alone survey. We cover this region with 1700 pointings. With integrations of 3 hours, we are much more sensitive to intermittent pulsars than e.g., the state-of-the-art GBNCC survey, with its 2-minute integrations. We extend this area south to declination -20°, by a 10 degree strip around the Galactic plane. As this area has previously been surveyed more deeply, we integrate these 200 pointings for 6 hours. This layout is illustrated in Fig. 9.

3.3. Sensitivity

We set up and calibrate 37 compound beams. These produce a 300MHz bandwidth split in 1024 channels, at 50μs sampling. These data are searched for transient pulses in real time. An alert is sent for any strong detections, and the incoming data chunk is saved to disk.

During the dedicated survey, telescopes A and B will be moved to 144-m spacings. Compound beams are coherently added in 37 sets of 12 grating lobes (see Fig. 10 and, e.g., Janssen et al. 2009) to form 450 coherent beams. These can detect all 1-millisecond-duration FRBs of flux $S_{1400} > 0.16$ Jy and dispersion measure DM < 5000 pc/cm^3, occurring in the Apertif field of view.

Through single-pulse and/or periodicity searches, we furthermore aim to detect all pulsars with mean flux density ≳ 40 μJy and DM < 1000 pc/cm^3.

3.4. Survey strategy and prioritisation

One aspect that makes Apertif unique is its capability, in combination with LOFAR, to localise Fast Radio Bursts with unprecedented positional accuracy. The nature of FRBs has remained elusive because such localization has so-far not been possible. Host galaxies, progenitors, and independent distances or redshifts are completely undetermined. The Apertif-LOFAR localization, illustrated in Fig. 8, will create a breakthrough in understanding what causes these mysterious radio bursts.

Using the real-time FRB search in PTS, it will be possible to detect when a burst occurs anywhere in the field of view of Apertif. However, due to the linear layout of the WSRT (the white grating combs in the left-hand side of Fig. 8), it will not be possible to determine an accurate position of the burst. Here, however, nature has been kind -- the high dispersion of the radio burst means the same burst will arrive several minutes later at LOFAR. Given its two dimensional layout and long baselines, LOFAR can then localize the same FRB to several arcsec. From initial detection to localization, the steps illustrated in Fig. 11 decrease the error box by a factor of one million.

No FRBs have so far been detected at the 200 MHz frequency used in LOFAR. This can arguably be explained by the difficulty of doing an FRB search that is blind in DM and location. In the FRB localization that is part of PTS, however, both the DM and approximate position will already be known. The recent detection of FRB emission down to 700 MHz is an encouraging confirmation that some if not all FRBs are detectable below 1.4 GHz.

Survey pointings will be carried out as close to the meridian as possible, such that any localization follow-up with LOFAR will provide best sensitivity and angular resolution. Further coordination with certain LOFAR

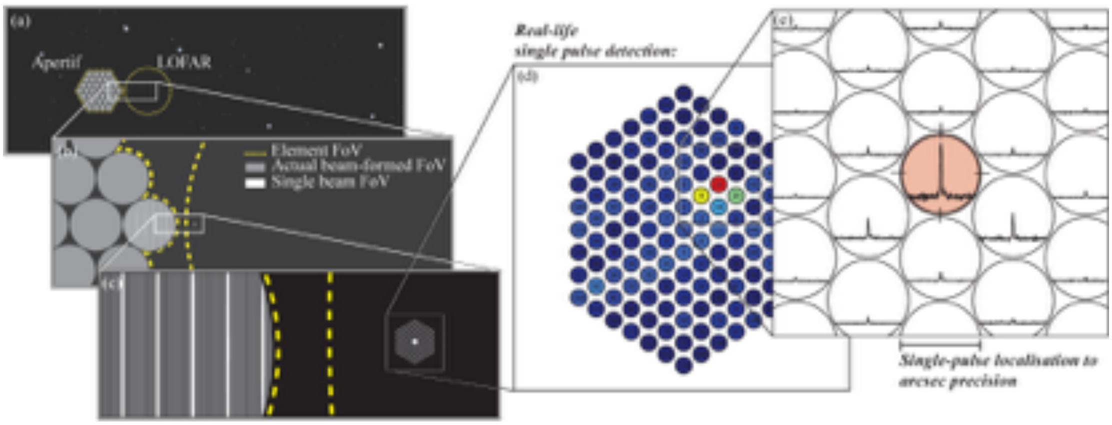

Figure 11. Left: The field of view (FoV) of Apertif and LOFAR, in successive zoom-ins. Big Dipper shown in (a), for scale. The element FoVs (dashed lines) of the Apertif dishes and LOFAR stations are similar. The gray area indicates where full sensitivity and time resolution are available. For Apertif this is visible in the top panel, for LOFAR (Dutch stations) only in the bottom panel (c). The white areas show a single beam, or resolution element; one out of 450 for Apertif and 127 for LOFAR. Clearly Apertif offers unsurpassed FoV, while LOFAR provides excellent localization. Right: LOFAR data on pulsar B0329+54 demonstrates that a single burst, with the same S/N as the first "Lorimer burst" FRB can be easily pinpointed to within one beam – offering down to arcsec localisation.

programs will be pursued, such that any PTS FRB may be found in the transient buffer boards for the directed LOFAR tile beam. While FRBs may intrinsically be 1 ms long, triggering at LOFAR will initially already record up to 5 seconds of data, and will later follow the dispersed burst down in frequency, such that the scattered pulse will be completely detected there.

Confirmation of FRB and pulsar candidates necessitates some pointings be repeated. When done at different hour angles, with rotated WSRT fan beam response, re-detections then also improve localization.

3.5. Comparison with other surveys

Compared to the best currently funded competition, the PTS field of view, and high sensitivity, mean it samples a volume of the Universe ten times larger than Parkes, the most successful experiment so-far (see Fig. 12).

The localization, however, is the biggest potential game changer. LOFAR provides every ALERT beam with over 1 million resolution elements. This crucial step is unique to ALERT, and likely unsurpassed in the next decade.

Two surveys that are currently under development and construction and may offer most competition are the upgraded Molonglo Observatory Synthesis Telescope (UTMOST) and the Canadian Hydrogen Intensity Mapping Experiment (CHIME), The FoVs of these are similar to Apertif. UTMOST has about twice the instantaneous sensitivity, but Apertif provides 10x more bandwidth. Some of the same FRB detection pipeline software is used in all three. FRB detection rates are expected to be roughly comparable (~1/week). The main difference is that UTMOST and CHIME will localize to ~1 sq. deg, while the link with LOFAR allows Apertif FRBs to be localized >10^6 times better, to 1 sq. arcsec. That is the resolution needed to e.g., detect FRB host galaxies.

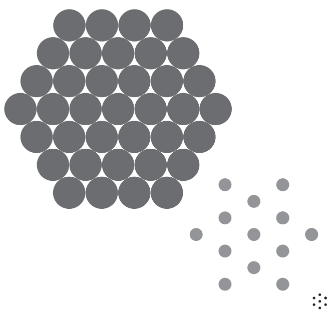

	Field deg^2	Gain K/Jy	Volume 10-3 Gpc^3	Localise arcsec^2
Apertif	8.7	1.2	**10**	**1**
Parkes	0.5	0.8	1	>100,000
Arecibo	0.01	10.5	2	10,000

Figure 12. Comparison of the probed volume, and the accomplished localization error box, for Apertif + LOFAR versus Parkes and Arecibo.

4. Data Processing

4.1. Imaging surveys

4.1.1. Overview of the Apertif Imaging Calibration Pipeline

The current implementation of the Apertif Calibration Pipeline (ACP) deals with the full calibration of Apertif stokes-I imaging data, both continuum and spectral line data. Although all Apertif observations will be done in full polarisation mode (including taking the necessary calibration data), no polarisation processing is currently foreseen in the ACP, given the current composition of the AST. However, all raw visibility data will be stored in the Apertif Long-term Archive and polarisation calibration of the observations will still be possible at a later stage, should interested groups join the AST.

The starting point of the ACP is to calibrate the observations, as much as possible, in the same manner as has been done in recent years for data from the WSRT, including the treatment of off-axis effects (although these will be at a much lower level than in data taken with the old WSRT). This implies that each Compound Beam (CB) is calibrated independently and that the combined images, made using all CBs, are constructed using a linear mosaic. The main difference with the standard calibration of WSRT data is that the calibration of the flux scale and of the bandpass does not involve a short observation of a standard calibrator at the field centre of each CB (it would take too much time to do this), but a combination of a short calibrator observation for one CB and the internal calibration of the PAF through the transmitter mounted at the apex of each dish. This allows to transfer the flux- and bandpass calibration of a single CB to all other CBs.

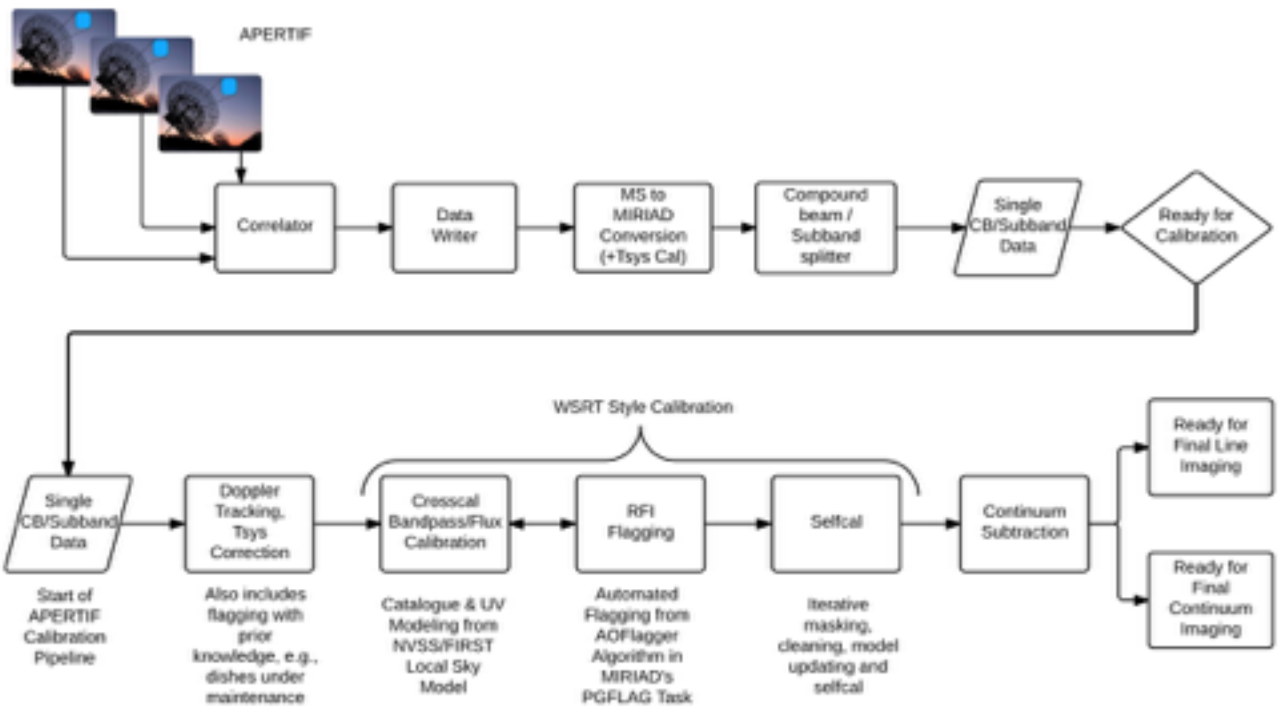

Figure 13. Schematic view of the Apertif Calibration Pipeline (ACP). The top half illustrates how Apertif data are prepared for use in the ACP, which is illustrated in the bottom half. The ACP calibrates a basic single-CB/single-subband unit of visibility data. The ACP includes Doppler correction, T_{sys} correction, flagging, bandpass/flux calibration, self-calibration and continuum subtraction. The calibrated data are then ready for final imaging of the combined spectral line and continuum data.

All calibration is based on the Miriad package. Miriad has been used extensively for calibrating WSRT data for the last 15 years and has been proven to be very efficient for this. Given that the calibration of Apertif data is very similar to that of WSRT data, using Miriad is an obvious option. However, a second, important, reason is that the AST wants to avoid a situation where one has to rely on software that is maintained and developed by teams outside the Apertif project and for which serving Apertif is not the first priority.

The general scheme of the pipeline is illustrated in Figure 13. The visibility data are prepared for processing through the ACP, which calibrates basic units of single-CB/single-subband (~15 MHz) slices of visibility data. The width of the subband is chosen such that one can still treat the data as monochromatic, which gives a huge performance benefit. Standard calibration of WSRT spectral line data does not use short observations of a secondary (phase) calibrator that are interspersed with those of the main target. Instead, it relies on self-calibration using the continuum emission in the field. Common WSRT practice is to use *clean* to generate sky models for the self calibration which implies (given the discreteness of the sky model) that the dynamic range of the images is limited to a few thousand. An extension to Miriad has been implemented by the AST which allows to use parametric sky models for self-calibration. This improves the dynamic range of the images to better than 10^4. The option of using parametric models also makes it possible to use external sky models based on the NVSS and FIRST in the initial phases of the calibration.

RFI flagging is currently done using the Miriad implementation of the AOFlagger algorithm on the polarised emission (i.e. Stokes Q, U and V). This has been extensively validated on existing WSRT data.

The ACP has been developed using a Python interface to Miriad (largely developed by the AST) and performs the necessary bookkeeping, syntax checks, exception handling and task management. The ACP can be used through a basic command-line interface (parameters are set in a parset), or through IPython/Jupyter Notebooks. The latter has the benefit of giving some level of interactiveness and to encourage

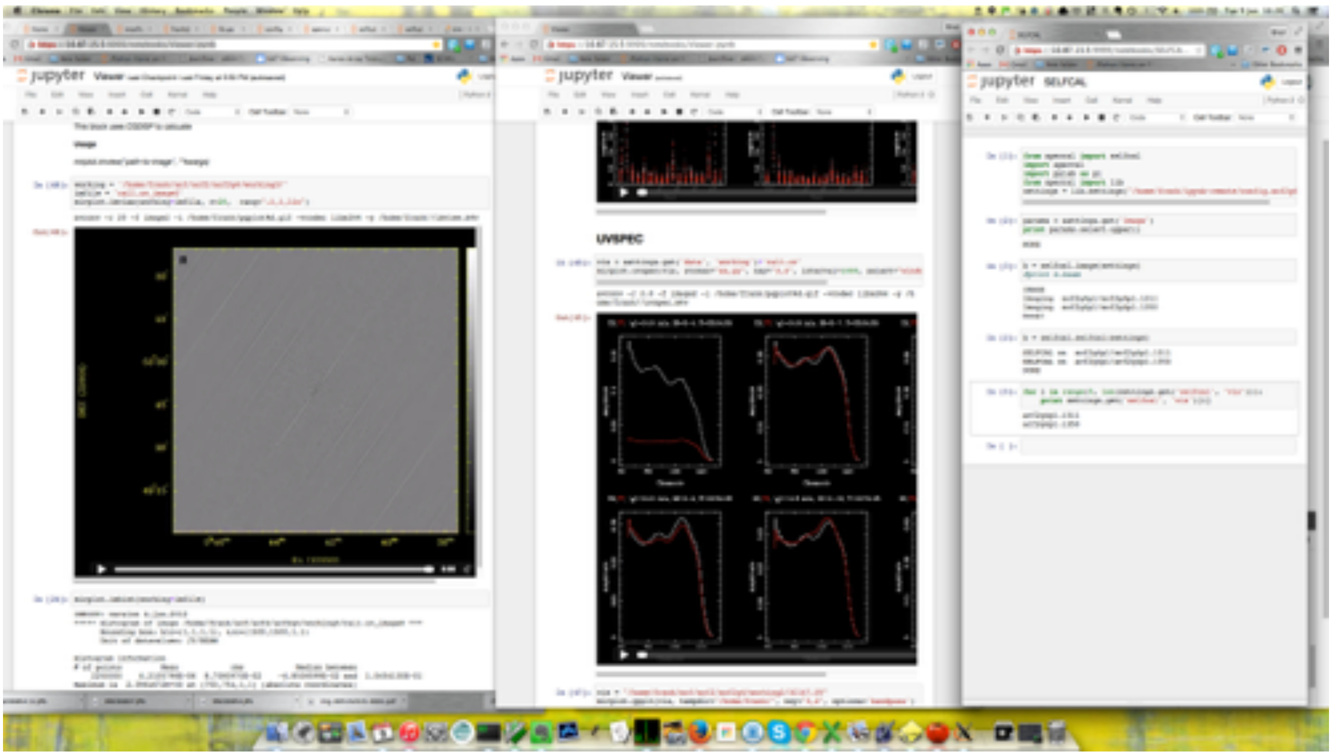

Figure 14. Impression of a number of IPython Notebooks developed for the ACP. The one on the left can be used for data inspection using images, the middle one for data inspection using spectral plots, while the one on the right is the one used for running the self-cal pipeline.

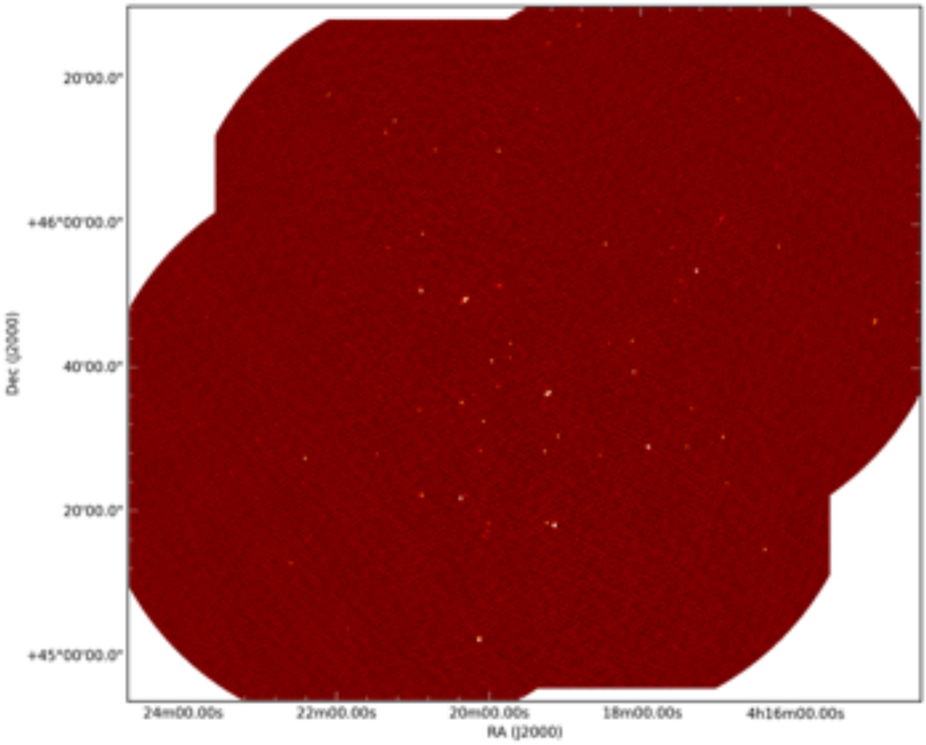

Figure 15. Image resulting from a four-point mosaic observation of one of the Apertif Reference Fields and calibrated by the ACP.

documentation of processing, which will be essential during commissioning. A number of Notebooks are available, each for a specific, limited task (e.g. flagging, data inspection, calibration). See Figure 14 for some examples.

As part of the preparation for the operation of Apertif, the AST has observed three Apertif Reference Fields with the WSRT. The purpose of these observations is to provide a reference between Apertif and the WSRT. The observations of each of these fields consists of a mosaic covering 36 pointings for which the data were obtained over nine 12-h observations, i.e. in a single day, four positions were observed. Figure 15 shows the result of such a four-point mosaic as calibrated by the ACP.

4.1.2. Data processing

The Apertif data taking is the responsibility of the Radio Observatory, based on the schedule and observation parameters that have been provided by the AST. The observations, after an initial quality check will be stored in the Long Term Archive. The software and tools required for the calibration and imaging and the scheduling and quality assessment is the responsibility of the AST. This is illustrated in Figure 16 which shows the different elements in the data path from observation to Long Term Archive. Once an observation has been taken it is stored in the archive and fed into the imaging pipeline to be processed within 24 hours in order to deliver a science ready data cube to the Apertif Long Term Archive.

4.1.3. Archive

The Apertif Long-Term Archive will contain the following data products: (i) the raw visibilities, to be stored immediately after the observations have been taken and have passed quality control; (ii) the calibration information (gain tables, flag tables etc.) and the deconvolved continuum data cubes and continuum subtracted line cubes after quality control. The calibrated data, produced by the Survey Team immediately after the observation, are subject to a proprietary period and initially only accessible to the Survey Team and their associates.

In addition, the Science Working Groups (see Sect. 7) are expected to also deliver other data products to the archive such as source lists, spectral index images, moment maps, and HI cubelets containing the object(s) of interest. The goal is to render an archive which will be the public Apertif legacy resource for many years to come.

4.2. Pulsar and Transients Survey (PTS "ALERT")

The hardware and software for ALERT passed its Critical Design Review on March 1, 2016, for implementation over Summer 2016. Full details are found in:

ASTRON-RP-1504	"Arts Architectural Design" by Eric Kooistra & Joeri van Leeuwen
ASTRON-SP-062	"Detailed Design ARTS FPGA Beamformer" by Eric Kooistra
ASTRON-RP-1506	"ARTS Design Pipeline SC3+4 Subsystem" by Alessio Sclocco, Joeri van Leeuwen & Emily Petroff
ASTRON-RP-1489	"Design Considerations for Apertif/ARTS software" by Ruud Overeem and Marcel Loose

A high-level overview is presented below.

4.2.1. Beam forming

Carrying out the ALERT pulsar/FRB search on even a state-of-the-art 1,000-CPU cluster would take over 80 years. Yet finding and following up transients demands we search all data in real time.

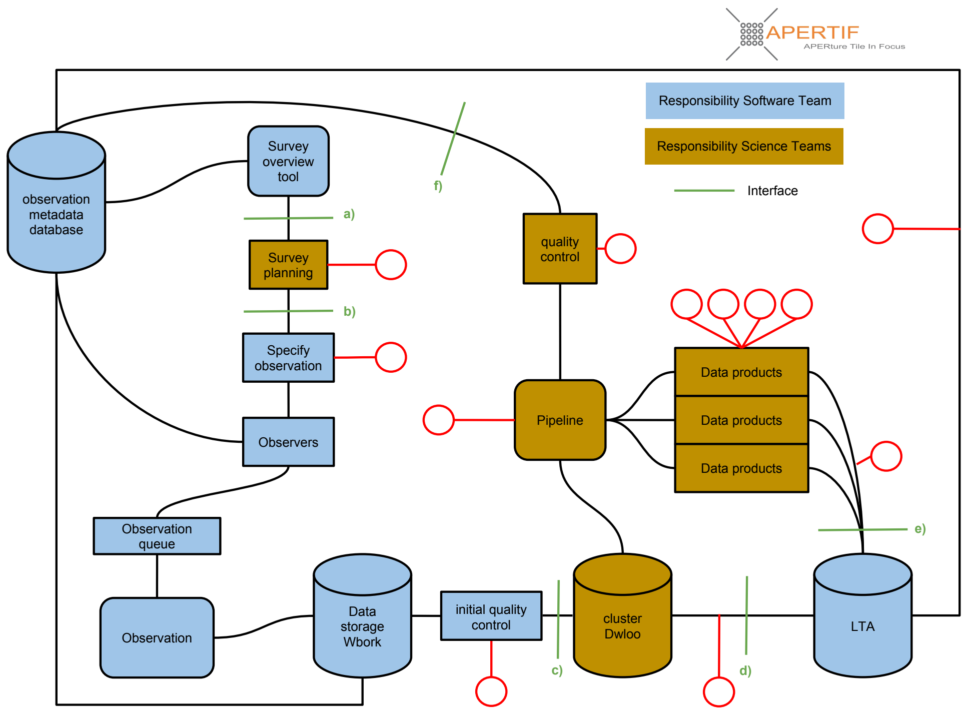

Figure 16. Overview of the data flow, from observation to science ready product in the Long Term Archive. The blue items are the responsibility of the Radio Observatory (and the necessary software is provided by ASTRON). The brown items including the necessary software are the responsibility of the Survey Team. The red circles indicate human interaction with the data.

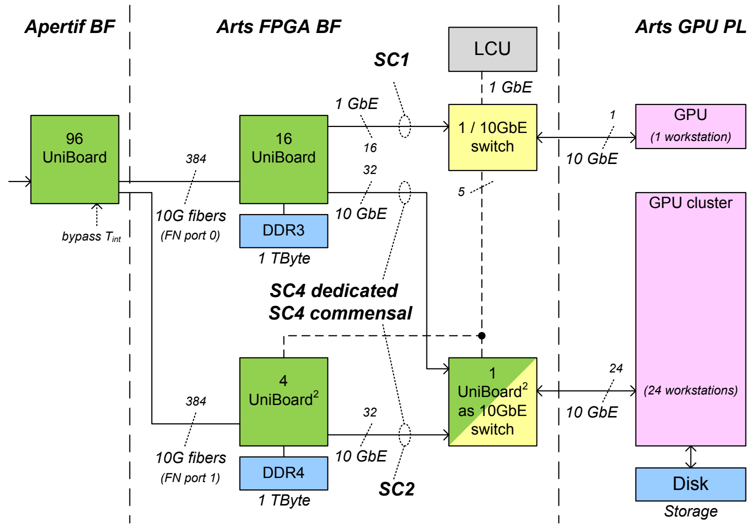

Figure 17. High level overview of the ALERT data processing hardware. Green blocks are FPGAs that carry out front-end beamforming ("Apertif BF"), or tied-array beamforming ("Arts FPGA BF"). The latter consists of two systems. One more powerful consisting of 16 Uniboards for dedicated surveys, and one more modest consisting of 4 Uniboard^2 for commensal searching. Pink blocks are the GPU Pipelines. Figure taken from ASTRON-SP-062 "Detailed Design Arts FPGA Beamformer" by Eric Kooistra.

Fast and economic alternatives for such CPU systems are however increasingly available. The "beam-forming" process, with its extreme input/output data rates, simple streaming algorithms, and low memory requirements, matches well to Field-Programmable Gate Arrays (FPGAs). The subsequent search for transients in time-distance parameter space involves moderate memory demands, high input I/O, and extremely high requirements on linear operations – thus perfectly suited for implementation on Graphics Processing Units (GPUs). Thus ALERT maps exceedingly well on a hybrid FPGA-GPU machine (see Fig. 17).

FPGAs can be customized to a high degree, to provide excellent computing efficiency. Especially when combining modular boards, containing FPGAs and fast I/O connections, specialized low-power high-performance compute systems can be rapidly built. GPUs provide excellent performance through many-core computing. This combination is technically heterogeneous for optimum scientific return.

4.2.2. Searching and archiving

Every observation is searched for FRBs and pulsars in real time. Every 4/8-hour integration is also stored on disk, and searched in-depth immediately after. For the real-time transient searching, the GPU/CPU servers buffers the beam-formed data streams and run asynchronous GPU and CPU tasks to perform different kinds of searching (4,000 parallel trial dispersion measures; thresholding for candidates above the running data mean; searching different pulse widths; periodicity searches; see Fig. 18). The first step of the real-time pipeline is dedispersion, i.e. the process of removing the effects of dispersion on the received signal. The second is time integration. The reason for this step is twofold: on one side it can be used to cancel out sub-millisecond high-intensity noise, and on the other side it can be used to increase the SNR of radio transient objects. What this step does is reducing the sampling rate of the dedispersed data by integrating and averaging over time. The third step is to compute the SNR of the dedispersed samples with the highest intensities in the processed time series, a value that will be later used to determine if a radio transient candidate has been identified. The last computational step is *detection* of values that exceed the user specified-threshold.

We have already successfully demonstrated this hardware/software system through the detection of a known FRB in a Parkes data set.

Before establishing a trigger on a candidate that has been identified in the real-time pipeline additional checks are necessary to keep the number of false triggers as low as possible. In some cases human inspection is required, but the candidate may also be given a score or rank, dependent upon parameters such as width, signal-to-noise and DM/DM Galaxy whose value dictates whether or not a trigger is issued. One possible parameter on which to score a candidate is the candidate's "broadbandness", i.e. whether the candidate has signal which is detectable across the majority of the full bandwidth (a common attribute of pulsar and FRB pulses). An example of how the candidates are presented to the user can be found in Fig. 19.

An alert is sent for any strong detections, and the incoming data chunk is saved to disk. For periodicity searches, the signals from these all fans beams are written to disk on the GPU cluster, for the duration of a survey pointing (3–6 hr). While the next observation starts, these finished data are combined, using the appropriate linear combinations, to form elliptical subbeams which have dimensions equal to the beam size that would result from an equivalent synthesis observation. These subbeams are spaced to overlap at the half-power point of both axes, and tesselate out the entire primary beam. The appropriate geometric weighting is determined by correcting for the fact that a particular position on the sky can move through different fan beams during an observation (as demonstrated in Janssen et al. 2009).

The resulting timeseries are searched for periodic signals using PRESTO. We excise radio frequency interference (RFI), and next carry out dedispersion, periodicity searches, and sifting and folding of the best candidates, as described below.

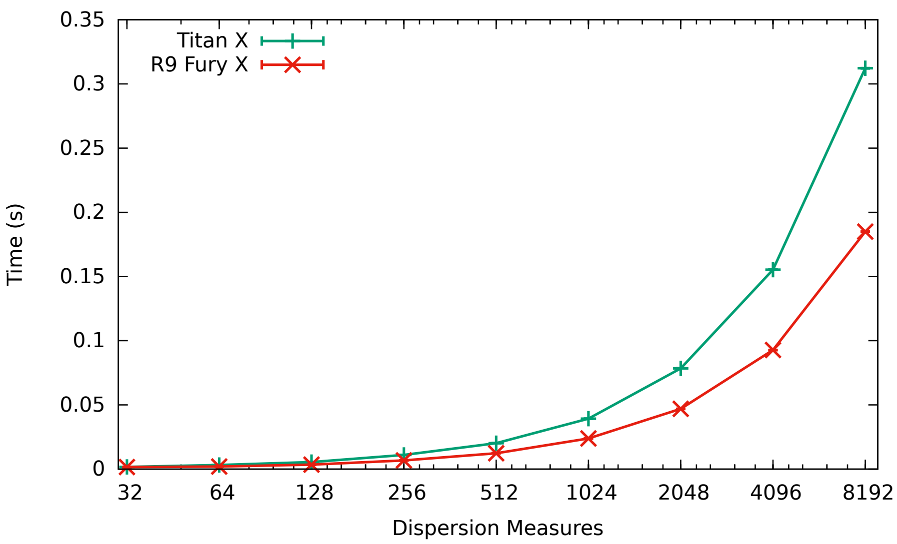

Figure 18. The performance of the demonstrator real-time system. Shown is the time taken to dedisperse 1 second of 1 beam. For the required 2,000 DM, up to 13 beams can be searched faster than real time on the default platform, the NVIDIA Titan X. That is already about 2x faster than the minimum requirement. Thus the real-time pipeline will use half of the planned 24-workstation quad GPU cluster. The other half will run the off-line periodicity search. The AMD R8 Fury X can even process 25 beams in real time, and is thus a strong contender for the final hardware implementation.

Given our success with this approach in LOFAR, we initially aim for periodicity searches assuming no acceleration. The envisioned long integrations T (3–6 hr), and the scaling of the required acceleration compute as T^4 may mean only some binary-removal searching is done, to e.g. a maximum number of bins

z_{max} of 50. PRESTO is GPU (CUDA) accelerated and half of the GPU power in ARTS is envisioned to be used for acceleration searching. After harmonic summing, acceleration candidates for all beams within a Compound Beam are sorted by significance, and duplicates in period and DM, across beams and across z_{max} values, are removed. Best candidates are next folded, optimized in period and DM, and ranked. Neural nets that are our team uses in other surveys (Coenen et al. 2015), and that will be trained on known ALERT pulsar re-detections and RFI instances, help prioritize candidates.

High time-resolution (≤ 50μs) data of every beam containing a known pulsar will be uploaded to the Apertif Long Term Archive. All other beams will be heavily down-sampled (to resolution of ~1 MHz, ~1 ms, 2-bit) and uploaded. Even at that down-sampling, the data volume already amounts to a challenging 7PB. Data will be kept at higher resolution if more storage becomes feasible. The entire survey will be kept, in this archive, publicly accessible, after the proprietary period ends.

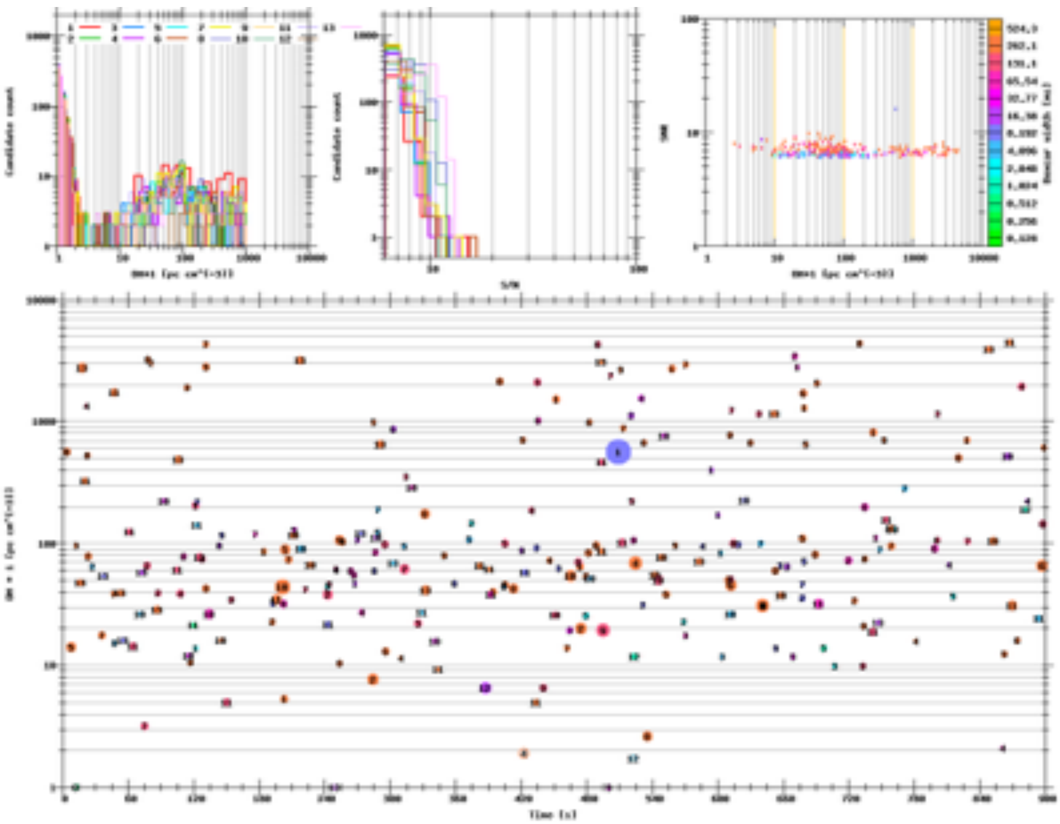

Figure 19. A sample overview plot of candidates generated in a Parkes search with the detection of FRB140514 with DM 562.7 pc cm^{-3}. The top left plot gives a histogram of candidates over the DM range searched, one for each of the 13 beams. The top centre plot gives a histogram of the S/N distribution of candidates and noise. The top right scatter plot gives the DM of each candidate and its S/N. In the bottom panel showing time and candidate DMs, the color and size of the filled circles represent the pulse width and S/N, respectively, with the primary detection beam labeled in the circle. A single bright pulse in beam 1 of width 8 ms and DM ~562 pc cm^{-3} is visible at t ~ 200 s in the DM-time plot, and in the top right corner of the DM-S/N plot; this is FRB 140515 published in Petroff et al. (2015).

5. Survey Operations

Apertif surveys will be carried out in two distinct observing modes - imaging in the maxi-short configuration, and dedicated transient searches in the Nx144m configuration (the latter being a configuration maximising the number of 144m spacings). The WSRT array will operate in one observing mode for a certain period of time (weeks or months).

In terms of support, the contribution of the Radio Observatory is limited to operation and maintenance of the telescope and instrument hardware and the production and storage of raw, uncalibrated data. The actual calibration of these data and creation of data products are solely the responsibility of the AST. This necessitates that during survey observing periods members of the Science Working Groups (see Sect. 7) will act as Survey Astronomers on Duty (SADs) and be in control of survey operations.

5.1. Daily operations of the imaging surveys

During periods of imaging, the SAD will be responsible for

1) populating and verifying the observing schedule for the coming period. Newly scheduled observations can be selected from a pre-defined pointing grid, paying attention to spatial coherence in the SNS survey area (minimize fragmentation), as well as the cadence and build-up of sensitivity across the MDS survey area.

2) ensuring that the data from the last completed observations have been copied successfully from the data writer at the telescope to the processing cluster in Dwingeloo, e.g. by verifying checksums. This will be done prior to storing the raw data in the Long Term Archive and executing the pipeline. Of prime importance is ensuring that the data writer at the WSRT maintains sufficient free storage space for upcoming observations.

3) validating the integrity of the data by checking that the observation is properly completed and all relevant metadata are in place with valid numbers. The metadata contains important information on the health of the system and is crucial for a proper execution of the pipeline.

4) performing Quality Control of the data, prior to executing the pipeline, by checking the health of the telescopes (availability, pointing, tracking), the health of the PAFs and beamformers (element performance, compound beam stability, noise source), the health of the correlator (glitches, baselines, subbands), and the environmental conditions (weather, T_{sys}, RFI).

5) starting and monitoring the calibration and imaging pipeline by checking the results of various intermediate data reduction steps such as the gain, phase and bandpass solutions for the external calibrators, the statistical results from the RFI flagger, the noise characteristics of the calibrated and continuum subtracted visibility data, the results from the self-calibration process, and the noise properties of the cleaned image cubes;

6) transferring the results from the calibration and imaging pipeline to the Long-Term Archive. These data products, as detailed in Sect. 6, should also be staged for retrieval by the various Science Working Groups.

7) preparing and submitting a standard daily report to the AST. This report should contain a concise and executive summary of the aforementioned activities. It should flag any failed observation and release that pointing for reconsideration for the schedule.

The SAD should perform these duties while physically present at ASTRON in Dwingeloo to facilitate communications with the staff of the Radio Observatory. To prevent a backlog, the SAD is expected to process the data from the previous day, weekend or national holiday. During a period of imaging observations 2 or 3 SADs should be available to ensure continuity of data processing on a daily basis. These are typically postdocs or staff members associated with Science Working Groups.

5.2. Daily operations of the transient surveys

The Pulsar and Transient Survey (PTS) is expected to run close to full time, either in dedicated (N x 144m) mode, or in reduced-sensitivity commensal mode parallel to the above imaging. Thus, every day one or more PTS SADs will be on duty. Tasks are the following.

1) Triaging the highest-significance real-time FRB detections, as these come in through SMS messages. This is a 24/7 task with an average response time of order 10 minutes. The SAD will confer with PTS team on need for multi-wavelength follow-up, then initiate such follow-up.

Tasks 2-7 fall within office hours:

2) Further investigate lower-significance FRB candidates.

3) Keep track of the pointing progress of the day. Spot-check the real-time RFI monitor. Go through the output of the "Quicklook" periodicity search for astrophysical (pulsar) signals, and an indication of the RFI environment. Mark finished observation as acceptable or failed.

4) Be available to follow up on warning and failure messages from the real-time and offline systems (including the automated archiving). In case of hardware/software down time, carry out initial trouble shooting, then organise support/maintenance.

5) If in dedicated mode, preparing observation schedule for the following day(s) by running the pointing-choice script, then moving these into the Monitoring and Control environment. Check and fill rotating schedule for SADs.

6) After observing sessions and/or telescope reconfigurations, check the tied-array beams are being properly coherently phased up, through a series of test-pulsar observations.

7) Marking standard observing results on the PTS Wiki, elevating non-standard issues with the PTS team.

For PTS, a pool of 5-8 SADs may be needed for the essential real-time responsiveness. SADs range from grad students to staff members.

6. Data Products and Data Release

6.1. Imaging Surveys

The Apertif imaging surveys are designed to be legacy projects with lasting impact on the scientific community.

HI data cubes will be produced at both 15" x 15"/sin(δ) and 30" resolution to provide both the highest possible angular resolution as well as optimal sensitivity to low column density emission. The full resolution data cubes will enable the kinematic analysis of galaxies in the MDS, but also bright inner disks of galaxies detected with the SNS. The lower resolution data cubes (with a marginal loss in point source sensitivity, but a significant increase in surface brightness sensitivity due to the unique layout of the WSRT) will serve to observe the outer gas disks of galaxies and the intergalactic medium and the diffuse continuum emission associated with galaxies, radio galaxies and clusters.

For the continuum data the complementarity with LOFAR and the LOFAR surveys is particularly relevant. Here two sets of images will be produced: images with the highest possible spatial resolution (15 arcsec; important for morphological classification of the sources, follow up identification, HI absorption studies, etc) and images at lower resolution (30 arcsec) to enhance the extended (low surface-brightness) emission (important for the study of large diffuse objects, clusters, relics etc.).

These data cubes and images, along with the associated raw and calibrated uv-data will be the main deliverables of the surveys and fall under the responsibility of the AST. These deliverables constitute the Level 0 (raw uv data), Level 1 (calibrated uv-data) and Level 2 (line and continuum cubes and images) data products. These data products will be released regularly through the long term archive (LTA) supported by the Radio Observatory (RO). The Level 2 data products consist of deconvolved and continuum subtracted line cubes and deconvolved full stokes continuum cubes from processing the entire useable band. The cubes will be produced at the highest possible resolution (15" x 15"/sin(δ)), and at 30" resolution.

Higher level data products will fall under the purview of the Science Working Groups (SWGs; see Sect. 7). The AST will however define a minimum set of Level 3 data products that should enter the archive and encourage the development of a common analysis tools among SWGs to provide these data products archive-ready. Level 3 data products will consist of the output of a standard source finding algorithm: a catalogue with positions, redshifts, flux densities, total fluxes, sizes, line widths, in addition to HI cubelets, total intensity maps, velocity fields, and continuum postage stamps. These analysis results form the Level 3 data products. As noted, these products are the responsibility of the SWGs and will depend on activities and priorities within these SWGs. A release schedule is therefore more difficult to estimate. The AST will however encourage the SWGs to release these products in a timely manner.

The AST will define a minimum set of Level 3 data products that should enter the archive and encourage the development of a common analysis tools among SWGs to provide these data products archive-ready. Level 3 data products will consist of the output of a standard source finding algorithm: a catalogue with positions, redshifts, flux densities, total fluxes, sizes, line widths, in addition to HI cubelets, total intensity maps, velocity fields, and continuum postage stamps.

In addition, it is expected that several survey publications will generate further parameterisations of sources and expanded catalogs (e.g., rotation curve fits, asymmetry parameters, optical cross-matching, etc.). These additional data products and catalogs will form the Level 4 data products, and release will follow the time scales of the publications.

The timely release of data products is a critical component for ensuring the legacy of these surveys. The 5 levels of data products discussed above will all eventually be publicly available. However, the different levels of data products represent different levels of effort necessary to create, and correspondingly, there are different release times for these different data products. We envisage the following release schedule.

Level 0, 1, and 2: raw and calibrated uv data; line/continuum cubes and continuum images
Data products associated with the shallow survey will be released no later than 6 months after the final observation of the preceding year. Assuming that regular survey operations will start in early 2017, this means that the first data release will take place the middle of 2018. Further data releases will then happen annually. These data will become available through the Apertif Long Term Archive.
The data release of Level 0, 1, 2 data products for the Medium Deep Survey follows a slightly different pattern. Here, the importance of accumulating observing time at each pointing and obtaining a uniformity in sensitivity determine the release schedule. The release dates will follow a similar schedule as the shallow survey, but data products for a given pointing will only be released when the accumulated observing time of that pointing during the preceding year has reached i) 1x12 hours (relevant for the shallow survey); ii) half of the total observing time; iii) the total observing time.

Level 3 and 4: source catalogs and published higher order data products
The level 3 products represent the basic output of a source detection scheme, i.e. catalogs containing the outputs of the source-finding and characterisation, including (but not limited to): position, velocity, size, flux and velocity width. Additional products will be cubelets, automatically generated moment maps, and spectra of detected sources. Continuum products include source lists, spectral indices, spectral index maps and rotation measure synthesis data. These data products and their public release will be the responsibility of the SWGs. The AST will monitor and encourage timely release of these data.

Several survey publications will generate further parameterisations of sources and additional catalogues (e.g., rotation curve fits, asymmetry parameters, optical cross matching, etc.). These data products and catalogs will be deposited into the Long Term Archive upon publication of the associated paper. Should publication of a paper occur before the associated data Level 0, 1 and 2 products are public, then these will also be made publicly available in the Long Term Archive, assuming there are no conflicts with other ongoing projects. The AST will be the final arbiter in this.

Limited requests (e.g., observations of single objects) by outside parties for unreleased data products will be considered on a case-by-case basis. If there are no conflicts with ongoing projects (as deemed by the AST), then these data will be made available.

6.2. Pulsar and Transients Survey

Through the PTS data products, we aim to create a digital interface to the transient sky, allowing access to both
a) the real-time data products that stream out of the transient detectors, and
b) data releases for the ~10-PB archive for legacy science.

6.2.1. Real-time products

Multi-wavelength follow-up is essential for the success of the source characterisation. PTS will have a robotic classifier for source level of confidence. This will actively broadcast VOEvent-compliant transient alerts for highly significant events, to collaborators or subscribers. PTS team members have previous experience with VOEvent triggering (e.g. COMET).

6.2.2. Archive products

For its periodicity and transient searches, PTS uses time-series at 50 μs time sampling and 1024 channels over 300 MHz. Archiving all tied-array beams at this resolution would take over 800 PB. Through downsampling in time, frequency and bit rate, this can be reduced to a more realistic 10 PB. Thus, within each PTS pointing, both dedicated or commensal, the real-time system runs a downsampling and storage mechanism. After every pointing, the resulting FITS time series (1-ms, 1-MHz, 2-bit, Stokes-I) are uploaded to the Apertif archive. Even at that reduced resolution, such an archive offers a ground breaking opportunity to future astronomers for investigating the radio transient sky (see e.g., Archibald et al., 2010, *Science*, for a high profile example using the Parkes archive). PTS members, in funded collaboration with the Netherlands eScience Center, have the general required expertise.

Furthermore, for all ~1000 known pulsars within the survey area, full-resolution time series will be stored in the archive and made public there throughout the survey.

7. Science exploitation of Apertif surveys

The scientific exploitation of the Apertif imaging and transients surveys will be project-based. Exploitation will be carried out by a variety of Science Working Groups (SWGs) coordinated by, and under auspices of, the Apertif Survey Team (AST) acting as an oversight committee. Each SWG is represented by an individual in the AST and may consist of a number of Individual Researchers (IRs, e.g., postdocs or PhD students) and/or Project Teams (PTs), each focusing on a particular scientific topic to address. Each SWG is expected to contribute to the daily operations of Apertif or to provide a comparable in-kind contribution such as (limited) access to non-public ancillary data, relevant and necessary computational infrastructure or a contribution in cash. See the document on Ground Rules for further details. Each SWG should outline which data products it will produce, on what timescale and which of those will be put in the archive.

In addition, each IR or PT within a SWG should submit a clear research plan outlining the specific scientific questions to be addressed, including a time-line for the research to be carried out as well as a publication plan, and identify the staffing to carry out this research. The proposed research should be focused and realistic in scope, thus avoiding generic and overarching themes. The representative of a SWG on the AST will introduce such a proposal for discussion among the AST members to avoid conflicts with ongoing activities of other IRs or SWGs. This is particularly relevant for protecting the research of PhD students.

SWGs may acquire new IRs or form new PTs during the survey period and hence are dynamic in their composition and research activities. Each new research initiative within a SWG should, however, be complementary to existing research efforts and explicitly not conflict with the research interests of other SWGs. In case of conflicts, proposed PTs can be asked to merge. A single PT can have members from multiple SWGs.

SWGs, PTs or IRs can not claim exclusive access to, or use of, Apertif data products but they can claim exclusive rights for a specific scientific use of these data. For instance, the integrated radio continuum flux of an extended source may be used by a PT in constructing a radio luminosity function but it should not prevent another PT from using the same data to construct a resolved SFR map for a particular galaxy in relation to the presence of OH megamasers. Hence, the same data product could serve different lines of research by multiple SWGs, PTs or IRs.

The AST will keep track of science projects, of which SWG/IR/PT is responsible for each project and the promised data products and assigned tasks. This list will be used to coordinate the division of tasks, the assignment of particular data sets, and will be the basis on which decisions about releasing pre-release data products to outside parties will be made.

New, independent SWGs may join the AST provided they make the required contributions and their research plans do not interfere with the research activities of IRs or PTs of other SWGs. The interests of new SWGs may influence the overall survey plan but adjustments will be weighed against the investments of other survey teams and the overall legacy value should not be adversely affected.

SWGs may seek collaborations with external research groups on the basis of specific scientific merit, e.g., to gain access to non-public ancillary data or relevant computer simulations. The sharing of Apertif data and external non-public ancillary data within such a collaboration will be arranged through MoUs to ensure that any data on common objects will not be distributed beyond the collaboration. To ensure limited proprietary access to Apertif data by external collaborators, any MoU should be presented for discussion within the AST prior to becoming effective. Such a discussion will pertain to the Apertif data only, as external collaborators may have their own guidelines for sharing their data or resources.

The progress of a Science Working Group will regularly be monitored and evaluated by the AST. Should significant delays in the working plan occur, in particular in the publication plan, the AST can request a SWG to restructure its activities. In the event of no improvement particular topics may then be made available again for others to work on.

7.1. Publication policy

7.1.1. Refereed journal publications

Commissioning and Early Science papers of Apertif are required to include the Apertif builders list. This team includes builders, designers and scientists that have contributed to the technical, organisational and scientific development of Apertif. The Apertif builders list will be maintained by the AST. Major survey papers will include by default the survey members that have contributed to survey planning, survey execution, software tools, data reduction, quality control and dataflow management that made the paper possible. It is expected that the membership list is a dynamic list and it is expected that individuals will join and leave. The AST is responsible for both new and expiring memberships.

The authorship list for Apertif papers should begin with the lead author of the work (as identified by the SWG PI, or, in doubt, by the AST), followed by the lead team, followed (in alphabetic order) by the other relevant contributors (e.g. suppliers of ancillary data, contributors of software specifically used for the project in question). Lead authors are normally required to be active members of Apertif SWG. However, in cases where the Apertif data is likely to be a minor part of the paper and the paper is to be led by a non-Apertif Science Team member, approval for the inclusion of Apertif data should be sought from the AST.

Draft papers should be circulated to the author list and the rest of the Apertif team, allowing sufficient time for comment. Apertif team members not listed on the paper may request co-authorship by writing to the first author (or supervisor, in case of a student first author), citing reasons for co-authorship. The modified draft paper should then be circulated to the authorship team for final review prior to submission. This stage is intended to pick up remaining minor errors, and not to raise new matters of substance. Approval by the AST before final submission is required, although such approval may be implicitly assumed unless the first author is otherwise contacted by the AST. Any disputes will be adjudicated by the AST.

7.1.2. Non-refereed papers

As it is impractical to include large numbers of co-authors, such papers only need to include key authors. Copies of submitted papers should be forwarded to the AST.

7.1.3. Other communications

Major announcements at conferences or meetings, articles for newspapers and the web, and all media releases need to be approved by the AST. Approval for normal conference presentations is not required.

7.1.4. ASTRON/Apertif Acknowledgements

The AST and the ASTRON Radio Observatory will define an "Acknowledgements" statement which should be included in internal and external papers making use of Apertif data.