

ARTS Requirements Specification

Document number: ASTRON-RS-020

Version 1.05



Document history:

Revision	Date	Author	Modification / Change
0.1-0.6	Jan-Jun, 2014	J. v. Leeuwen	L0 Science Requirements for all four Science Cases in place.
0.70	2014-Aug-29	E. Kooistra	Added document number, figure numbers, page numbers, list of references and definitions.
0.71	2014-Sep-19	E. Kooistra	<ul style="list-style-type: none"> Added feasibility study on 1 fits in with Apertif in appendix 8, 9. Updated System Requirements L1 in section 5.
0.72	2014-Sep-26	J. v. Leeuwen E. Kooistra	<ul style="list-style-type: none"> Inserted SR-0.40 (relates to SR-0.35) Clarified SR-0.6, 0.10, 0.42 Updated appendix 8 and added a section on the data storage for ARTS.
0.73	2014-Oct-17	E. Kooistra	<p>Updated review comments from Joeri van Leeuwen, Andre Gunst and Stefan Wijnholds.</p> <ul style="list-style-type: none"> Added section 8.11.1 on logging and meta data Updated section 8.2 on fringe stopping for ARTS Corrected TAB with grating lobes in Figure 8 and added Figure 9. Added Figure 13 with alternative decompositions for SC3. Added life time archive section 8.8.4.3. Added section 8.7.4 on beam data width and AGC. Added section 8.10.3.1 on using Alpha-3 of Apertif X also for early results with 3 dishes for ARTS.
0.74	2014-Dec-3	E. Kooistra	<ul style="list-style-type: none"> Corrected the description in 8.5.3 of two PFB in series. Option 3c in Figure 13 is not feasible, because the Apertif BF cannot calculate Stokes U and V.
0.75, 0.76	2014-Dec-6	J. v. Leeuwen	<ul style="list-style-type: none"> Detailed § 5.3.5, 5.3.6 & 5.3.7 Adjusted and expanded § 8.8.4.1 Improved § 8.1.2 "Flux", now consistent with Top-level SRs.
0.77	2015-Jan-13	J. v. Leeuwen	Corrected obsolete mention of 12.5 MHz ($N_{\text{band}}=24$) for 18.75 MHz ($N_{\text{band}}=16$)
0.78	2015-Jan-21	J. v. Leeuwen	Updated Fig 1 per OAR.
1.0	2015-Feb-18	J. v. Leeuwen	Updated L1s per Scientific Requirement Review, L1 are now final.

1.01	2015-Mar-17	R. Smits J. v. Leeuwen	First version of L2/L3s for SC1
1.02	2015-Mar-25	J. v. Leeuwen J.-P. de Reijer	Included final OARs (limits on relative delays between TABs). First version of space, cooling operational reqs. Updated L2/L3s for SC1
1.03	2015-Jun-03	R. Smits A. Sclocco J. v. Leeuwen	Updated L3s for SC1 First version of L3 GPU hardware estimates for SC4
1.04	2015-Jun-10	E. Kooistra	Added references [9] and [10] FR-2.3, section 8.7.3: defined TAB weights width [9] FR-2.4, section 8.2, 8.7.3: defined TAB weights update rate [9] GR-1.2, section 8.5.2: added reference to $B_{sub} = 1$ MHz change request document [10] Added ???? on $W_{tab}=4$ and $W_{power}=4$ because they may need to become e.g. 6 and 8 after analysis in [9] Section 8.8.4: UniBoard1 may be able to fit 16 GByte DDR3. Section 8.10.2: Apertif BF provides delay tracking (DT) Added Table 9 with the main MM registers Arts BF.
1.05	2015-Jun-11	E. Kooistra	Section 8.10.4: - Described 10 G Tx-only interface between UniBoard and GPU for SC3. - It may appear possible to have SC3 commensal with Apertif X, but it is not possible to say until more actual implementation resource usage figures are available. - Described option of having 16 extra UniBoards (or 4 extra UniBoard2) dedicated for Arts. - Added conclusion that all four SC of Arts will use the Uniboards to connect to the Apertif BF.

References:

1. NWOM2012-vanLeeuwen-ARTS, "ARTS – the Apertif Radio Transient System", 2012/2013, Joeri van Leeuwen
2. NOVA4-Round2-ARTS-NoLetters, "ARTS, NOVA Phase 4 Instrumentation – Round 2" , Joeri van Leeuwen
3. ASTRON-FO-017 (RS-005), "Apertif Science Requirements", CDR 2014, Marc Verheijen, Tom Oosterloo, Joeri van Leeuwen, Thijs van der Hulst
4. ASTRON-RS-006, "Apertif System Requirement Specification", CDR 2014, W.A. van Cappellen, T.A. Oosterloo
5. LOFAR-ASTRON-SRS-017, "LOFAR System Requirement Compliance", 2007, A. Gunst, M. Bentum
6. ASTRON-RP-321, "Apertif Architectural Design Document", 2009/08/11, W.A. van Cappellen
7. ASTRON-RP-1400, "Detailed Design of the Correlator System for Apertif", G. Schoonderbeek, A. Gunst, E. Kooistra, H.J. Pepping.
8. "Impact analysis of 4-bit quantization in LOFAR", March 8, 2008, S.J. Wijnholds
9. ASTRON-??-???, "Analysis of tied-array beamforming for ARTS", S.J. Wijnholds
10. ASTRON-CR-032, "Impact analysis of change request to use 1 MHz beamlets instead of 0.78125 MHz beamlets in Apertif and Arts", E. Kooistra

Contents

- 1 Conventions..... 11
 - 1.1 Requirement levels..... 11
 - 1.2 Requirement categories 11
 - 1.3 Numbering of requirements..... 12
 - 1.4 Priority specification..... 12
- 2 Scope of this document..... 13
- 3 Science and Use Cases..... 14
 - 3.1 Pulsar timing (SC1) 15
 - 3.1.1 Pulsar timing Use Case 15
 - 3.2 Streaming wide-field VLBI (SC2) 16
 - 3.2.1 Streaming wide-field VLBI Use Case..... 16
 - 3.3 Commensal Searching for Fast Radio Bursts and Pulsars (SC3) 17
 - 3.3.1 Commensal Transient Searching Use Case..... 17
 - 3.4 Dedicated Searching for Fast Radio Bursts and Pulsars (SC4)..... 18
 - 3.4.1 Dedicated Transient Searching Use Case 18
- 4 Science Requirements (L0)..... 19
 - 4.1 Science requirements for (at least) SC1 19
 - 4.2 Further science requirements for (at least) SC2 27
 - 4.3 Further science requirements for (at least) SC3 31
 - 4.4 Further science requirements for SC4..... 36
- 5 ARTS System Requirements L1..... 38
 - 5.1 Definitions 39
 - 5.2 General Requirements L1..... 40
 - 5.2.1 Apertif BF additional requirements..... 40
 - 5.3 Functional and Performance Requirements L1..... 41
 - 5.3.1 SC1, SC2, SC3 and SC4 41
 - 5.3.2 TAB for SC1, SC2 and SC4 42
 - 5.3.3 SC1 Pulsar timing..... 44
 - 5.3.4 SC2 VLBI 45
 - 5.3.5 SC3 and SC4 48
 - 5.3.6 SC3 Commensal transient search 52
 - 5.3.7 SC4 Dedicated transient search..... 54
 - 5.4 Operational Requirements L1 55
 - 5.5 Interface Requirements Hardware L1 56
 - 5.6 Interface Requirements Software L1..... 58
 - 5.7 Non-functional Requirements L1 58
- 6 ARTS System Requirements Level 2 60

6.1	General Requirements L2.....	60
6.2	Functional Requirements L2.....	60
6.2.1	SC1, SC2, SC3 and SC4	60
6.2.2	SC1 Pulsar timing.....	61
6.2.3	SC2 VLBI.....	63
6.2.4	SC3 Commensal Transient Search	63
6.2.5	SC4 Dedicated Transient Search.....	63
6.3	Interface Requirements Hardware L2	64
6.4	Interface Requirements Software L2.....	64
6.5	Operational Requirements L2	64
6.6	Non-functional Requirements L2	65
7	ARTS System Requirements Level 3	67
7.1	Functional requirements Beamforming L3.....	67
7.2	Functional requirements Pipelines L3	68
7.2.1	Functional requirements Pipelines L3, SC 1 (Timing)	68
7.2.2	Functional requirements Pipelines L3, SC 2 (VLBI).....	74
7.2.3	Functional requirements Pipelines L3, SC 3 (Commensal search)	74
7.2.4	Functional requirements Pipelines L3, SC 4 (Dedicated search)	75
7.3	Interface Requirements Software / Hardware L3	75
7.4	Functional Requirements GPU cluster L3.....	75
7.5	Operational Requirements GPU-cluster L3	76
7.6	Non-functional Requirements GPU-cluster L3	77
8	Appendix: ARTS feasibility study	78
8.1	Sensitivity	78
8.1.1	Full sensitivity	78
8.1.2	Flux	78
8.2	Delay tracking and fringe stopping.....	78
8.3	Time synchronization	79
8.4	Simultaneous operating modes	80
8.5	RF frequency grid and subband bandwidth	80
8.5.1	Remarks on tuning f_0 within 10 MHz.....	82
8.5.2	Remarks on changing B_{sub}	83
8.5.3	Channelization	84
8.6	Apertif BF output interface	84
8.6.1	Number of parallel links	84
8.6.2	Apertif BF compound beam load.....	85
8.6.3	Physical interfaces	85
8.6.4	Apertif BF output to Apertif X and ARTS	86
8.7	Beam definitions.....	86
8.7.1	Compound beams (CB).....	86
8.7.2	Incoherent Array Beams (IAB)	87

8.7.3	Tied array beams (TAB)	88
8.7.4	Beam data width	89
8.8	ARTS data interfaces.....	90
8.8.1	Data reduction.....	90
8.8.2	Data transposes.....	91
8.8.3	ARTS BF data loads	91
8.8.4	Data storage	92
8.9	Monitoring and control	93
8.10	Preliminary sub system decomposition	94
8.10.1	Data handling tasks	94
8.10.2	Apertif BF.....	95
8.10.3	Apertif X and ARTS BF.....	95
8.10.4	Preliminary mapping of ARTS on hardware	97
8.11	Logging	99
8.11.1	Meta data	99
9	Appendix: Channel bandwidth discussion for VLBI and Leap	100
10	Appendix: Power and space availability.....	101

Terminology and abbreviations:

beam	Group of beamlets that point in the same direction
beamlet	Beamformed subband, a small beam spanning one subband
channel	Unit frequency band within a beam
ephemeris	Pulsar data format file with the parameters used in the timing model
node	Processing node (PN), typically 1 FPGA chip
power beam	Full Stokes power values: I, Q, U, V
subband	Unit frequency band output of the filterbank in the Apertif BF
voltage beam	Dual polarization sample values with phase information: X_{re} , X_{im} , Y_{re} , Y_{im}
ADC	Analogue to Digital Conversion
AGC	Automatic Gain Control
Apertif	Aperture Tile in Focus
Apertif BF	The Apertif Beamformer, produces CBs
ARTS	Apertif Radio Transient System
ARTS BF	The ARTS Beamformer, produces TABs or IABs
BF	Beamformer
BSN	Block Sequence Number (timestamp)
BW	Bandwidth
CB	PAF Compound Beam, formed at dish level by Apertif BF
CPU	Central Processing Unit
CW	Carrier Wave (single frequency signal)
DM	Dispersion Measure
DT	Delay Tracking
EVN	European VLBI Network
FITS	Flexible Image Transport System
FoV	Field of View
FPA	Focal Plane Array (= PAF)
FPGA	Field Programmable Gate Array
FRB	Fast Radio Burst (Lorimer et al. 2007)
FWHM	Full Width at Half Maximum (= at -3 dB point)
GPU	Graphics Processing Unit
HPBW	Half Power Beam Width
IAB	Incoherent array beam, formed by incoherently combining dishes in ARTS BF
IH	Interface Hardware
IS	Interface Software
LEAP	Large European Array for Pulsars
MAC	Monitoring and Control
MTBF	Mean Time Between Failures
PAF	Phased Array Feed (= FPA)
PFB	Poly phase Filter Bank
PL	Pipeline processing
PN	Processing node

PPS	Pulse Per Second
PSRFITS	FITS for Pulsar data
PTA	Pulsar Timing Array
RF	Radio Frequency
RFI	Radio Frequency Interference
SC	Science Case
SEFD	System Equivalent Flux Density
SFXC	The EVN Software Correlator at JIVE
SNR	Signal to Noise Ratio
SP	Signal Path, 1 CB consists of $N_{\text{pol}} = 2$ SP
SR	Science Requirement
SRS	System Requirements Specification
TAB	Tied array beam, formed by coherently combining dishes
T_{ant}	Transpose to group data from all $N_{\text{ant}} = 64$ antenna elements in the PAF
T_{dish}	Transpose to group data from all $N_{\text{dish}} = 12$ dishes
T_{pol}	Transpose to group data from both $N_{\text{pol}} = 2$ polarizations
T_{band}	Transpose to group data from all $N_{\text{band}} = 16$ bands
T_{FoV}	Transpose to group data from all $N_{\text{CB}} = 37$ beams for the full FoV
ToA	Time of Arrival
TT	Terrestrial Time
VDIF	VLBI Data Interchange Format
VLBI	Very Large Baseline Interferometry
WSRT	Westerbork Synthesis Radio Telescope
X	Correlator

Definitions:

N_{complex}	2	Two part of a complex number, the real and imaginary part
N_{pol}	2	Number of polarizations, X and Y
N_{Stokes}	4	Number of power values in the Stokes vector [I, Q, U, V]
N_{dish}	12	Number of dishes or telescopes in Apertif
N_{sp}	24	Number of signal paths = $N_{\text{dish}} * N_{\text{pol}}$
N	1024	FFT size of the FFT in the Apertif BF subband polyphase filter
N_{sub}	512	= $N/2$, number of subbands that covers $RF_{\text{BW}}=400\text{MHz}$
N_{sel}	384	Number of selected subbands to cover $CB_{\text{BW}}=300\text{ MHz}$
N_{CB}	37	Number of compound beams
N_{gr}	12	Number of grating lobe patterns TABs to cover the full CB (SR-0.41)
N_{VLBI}	12	Number of TABs in the central CB for VLBI, choose = N_{gr} (SR-0.23)
N_{IAB}	37	= N_{CB} , number of IABs
N_{TAB}	444	Number of TABs
N_{PN}	384	= $N_{\text{sp}} * N_{\text{band}}$, total number of parallel processing nodes in the Apertif BF
N_{band}	16	Number of bands in the Apertif BF to process the full CB_{BW}
N_{link}	384	Number of physical 10G output links of the Apertif BF, = N_{PN} so 1 link per PN
f_s	800 MHz	Digitizer sample frequency
T_s	1.25 ns	= $1/f_s$, digitizer sample period
f_0		Lower edge frequency of a subband, beamlet or channel
RF_{BW}	400 MHz	= $f_s/2$, sampled RF bandwidth
CB_{BW}	300 MHz	Full bandwidth of the CB and also of the TAB and IAB (SR-0.2)
B_{sub}	781250 Hz	Subband bandwidth in Apertif BF, = beamlet bandwidth
N_{chan}	4	Number of channels per beamlet, for SC3 and SC4
B_{chan}		$B_{\text{sub}}/N_{\text{chan}}$, channel bandwidth within a beamlet, for SC3 and SC4
N_{int}	≈ 10	Number of Stokes channel power values that can be integrated
T_{Stokes}	$\approx 50\ \mu\text{s}$	Minimum required sample period for the Stokes power values
f_{Stokes}	$\approx 20\ \text{kHz}$	= $1/T_{\text{Stokes}}$, minimum required sample frequency for the Stokes power values
W_{beamlet}	6	Word width in number of bits of a beamlet voltage sample
W_{chan}		Word width in number of bits of a channel voltage sample
W_{tab}	4 → 6 ????	Word width in number of bits of a TAB voltage sample
W_{power}	4 → 8 ????	Word width in number of bits of a IAB or TAB power sample

1 Conventions

1.1 Requirement levels

Requirement levels L0, L1, and L2 are covered in this document:

L0	Science Requirements
L1	Top level system requirements
L2	Subsystem requirements
L3	Design specifications per subsystem

The L0 science requirements define what is required to be able to perform the astronomical science use cases. The L1 system requirements translate L0 into more technical system level requirements considering the entire system as a single black box with external interfaces. For L2 a first decomposition is made into subsystems and therefore it also defines the internal interfaces between these subsystems. At L3 the subsystems are further detailed to specify how they should be implemented.

1.2 Requirement categories

The L0 requirements for the Science Cases (SC) are defined by the:

SR Science Requirements

The L1 and L2 system level requirements can be separated into categories:

FR Functional Requirements
OR Operational Requirements
IH Interface Hardware Requirement
IS Interface Software Requirement
NFR Non-Functional Requirements

In the System Requirements Specifications for Apertif [4] and Lofar [5] similar requirement categories are distinguished:

- FR = Functional and performance requirements (data path, MAC, operation modes)
- OR = Operational Requirements (startup and shutdown, failure management, maintenance)
- DC = Design Constraints (environment, engineering, quality)
- IR = Interface Requirements (external, internal, power, timing, protocol, meta-data)
- SUP = Support Requirements (maintenance tests, logistics)
- VR = Verification Requirements

However in this document the categories FR, OR, IH/IS, and NFR are used. Any requirement that does not fit under FR, OR or IH/IS is put under NFR.

1.3 Numbering of requirements

Example:

FR-1.5-revB is level 1 functional requirement nr 5, second revision:

- First nr indicates the requirements level,
- second number contains the numbering within that level,
- revX indicates the revision of the requirement.

Other remarks

Changes to the revision number of a requirement will only be done after a change request in the project.

Traceability of requirements: the numbering of requirements has no meaning with respect to traceability. Their only function is to give unique identification of requirements. Dependability of requirements in terms of parent/child relationships cannot be deducted from the numbering. This must be administrated in a separate way (e.g. Excel sheet).

1.4 Priority specification

- Priority 1: Hard requirements for expert astronomers. Essential capabilities for the scientific success of this project. These capabilities are funded deliverables and will need to be accounted for to our funding agencies.
- Priority 2: The medium-hard requirements. Capabilities that increase the instrument versatility and ease of use for non-expert. These open up new scientific capabilities and increase accessibility for potential end users. These capabilities are funded deliverables and will need to be accounted for to our funding agencies.
- Priority 3: The soft requirements. Further extensions to the system that are not milestones to our funding agencies.

2 Scope of this document

This document describes pulsar, VLBI and transient processing required on Apertif [1, 4]. It assumes the input of 37 Compound Beams from each of 12 dishes. Where possible requirements already use specifications (e.g., bandwidth, field of view) currently set for the front-ends and front-end beamformer.

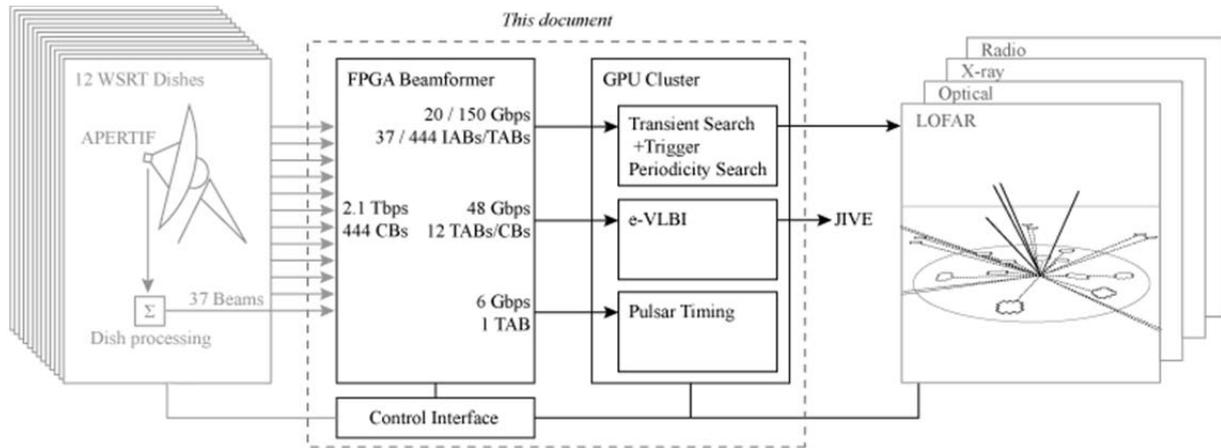


Figure 1: Top level overview of ARTS. Mentioned data rates, number of beams and hardware platforms are straw-man defaults and not requirements.

3 Science and Use Cases

The science cases are ordered to reflect increasing complexity. To build for SC4, practically all the Priority 1 requirements from SC1..3 would already be met first.

The *scientific* priority of the science cases is exactly the opposite. SC4 is the most cutting edge and new; SC1-2 only more modestly extend capabilities that the current WSRT already possesses.

Highest priority is thus SC4, then SC3, then SC1..2.

3.1 Pulsar timing (SC1)

Top-level science requirement:

“ARTS shall achieve a timing precision that has systematic instrumental errors of less than 20 ns”

Flux- and polarisation-calibrated pulsar data will be coherently dedispersed and folded at the pulsar DM and period. A standard frequency resolution of ~ 1 MHz can be adjusted down, or up to ~ 18.75 MHz band. The resulting folded profile is tagged using the observatory clock. The sensitivity is maximized (12 dishes, 300 MHz bandwidth). A single tied-array beam (TAB) in the central compound beam (CB) suffices. Additionally, baseband data can be written to disk for off-line “LEAP” combining, to detect weaker sources.

3.1.1 Pulsar timing Use Case¹

Usage narrative: An astronomer wants to measure the behaviour of a set of pulsars over a 5 year period. She submits a proposal with a source list. At the requested cadence, which is once or twice a month for each pulsar in the set, these are scheduled. At the scheduled time, the telescopes are pointed, a single TAB in the field center is formed (phased up and calibrated for polarisation and flux), over 300 MHz, at ~ 1 MHz channels (≈ 1 μ s time resolution). This TAB is coherently dedispersed and folded in real time, at the source dispersion measure and period, either for a fixed amount of time per pulsar, or until an integrated signal-to-noise (SN) of 20 is reached. Collapsed folded profiles are stored in a backed-up archive and cross-correlated with a high SN analytic template to calculate a time-of-arrival (ToA) for each observation. The ToA needs to be referred to terrestrial time (TT) by linking the measured time to the observatory clock, which is then converted to TT by using corrections to GPS and corrections to UTC provided by the Bureau International de Poids et Mesures (BIPM) in Paris. The timing models based on these ToAs produce scientific publications and contribute to the European and International efforts to directly detect gravitational waves in a pulsar timing array (PTA). The best PTA pulsars, have intrinsic timing precision of ~ 100 ns. The systematic errors introduced by ARTS should be less than 20% of that, i.e., 20 ns.

Primary Actor: Astronomer

Goal in Context: Astronomer applies for and gets pulsar timing data

Funded Extensions:

- 1: Adjustable channel widths up to 16, 18.75, and/or 20 MHz
- 2: Write baseband to disk
- 3: From expert mode to general-user mode
- 4: Off-center TABs

Open issues:

We need to make sure that the scheduling software can handle lists; i.e. when a list of start and end times for a set of pulsars is provided it needs to be possible that this list is read into the system and converted to a working schedule automatically.

¹ Use Cases based on “[Writing effective use cases](#)”, A. Cockburn, Addison-Wesley (2001),

3.2 Streaming wide-field VLBI (SC2)

Top-level science requirement:

“ARTS shall contribute to standard EVN imaging, with a field of view (FoV) of 0.000015 deg² initially, and 0.25 deg² subsequently”

3.2.1 Streaming wide-field VLBI Use Case

Usage narrative: An astronomer wants to image a field at high angular resolution with the European VLBI Network (EVN). He/she submits an observing proposal to the EVN Programme Committee which is approved. The EVN Scheduler finds a suitable date within an EVN observing session. The observation of a single epoch will typically last 8 hours. A single (phased-up) tied-array beam (TAB), covering a FoV of 0.000015 deg², is formed. The data are recorded on a Mk5 unit, or formatted to VDIF and directly streamed to the EVN correlator in subbands that are either 2, 4, 8, 16 or 32MHz wide. The full EVN array will produce an image of the target with a resolution of about 5 milliarcsecond. The on-source and off-source times are written to a log file, along with the measured SEFDs. The astronomer will need this for the proper calibration of the VLBI data. Local interferometer data are processed in parallel, with an averaging time of 1–10 seconds, and spectral channels of 16–8192 (depending on the science case, but for typical continuum projects 0.5 MHz spectral resolution is sufficient, i.e. 64 spectral channels for a 32 MHz subband). The local interferometer data will provide the astronomer with an accurate, simultaneous total flux density measurement of the target, and an image of the source neighborhood on arcminute scales, with an angular resolution of about 10 arcseconds. In order to calibrate the local interferometer data properly, the astronomer will make sure to include a WSRT primary flux density calibrator in the VLBI schedule. The results are published in major peer-reviewed journals.

Primary Actor: Astronomer

Goal in Context: Astronomer applies for and obtains high angular resolution EVN data

Funded Extensions:

- 1: Stream multiple TABs within the central-CB 0.25 deg² field of view
- 2: Stream the central CB from individual telescopes to achieve the full 0.25 deg² field of view
- 3: From expert mode to general-user mode
- 4: Automated imaging of local interferometer data; input to EVN correlator for multi-field correlation

3.3 Commensal Searching for Fast Radio Bursts and Pulsars (SC3)

Top-level science requirement:

“While Apertif is in imaging mode, ARTS will ‘commensally’ detect all 1-millisecond-duration fast radio bursts (FRBs; Lorimer et al. 2007) of flux $S_{1400} > 0.55$ Jy and dispersion measure $DM < 5000$ pc/cm³, occurring in the entire Apertif field of view”

The latency from time-on-sky to detection will be less than ~1 second, such that other telescopes (e.g. LOFAR, EVN) can still be triggered. ARTS sends out these triggers.

ARTS can self-trigger on Stokes-I detections and save buffered high-resolution polarisation data. ARTS keeps a ~12-hr ring buffer of (downsampled) Stokes-I data to search for periodic sources, off-line.

3.3.1 Commensal Transient Searching Use Case

Usage narrative: A team of astronomers want to search for *rare* radio bursts. They apply for and get allocated the right to piggy-back on imaging observations. Before each imaging session, they can use a command-line/scriptable tool to produce the commensal beam-formed schedule. This schedule sets up and calibrates ~37 compound beams (CBs) per telescope. This schedule makes sure these are incoherently added, producing 300MHz bandwidth split in ~1024 channels, at ~50 μ s sampling. The schedule tool makes sure 37 incoherent array beams (IABs) are formed, and are next searched for transient pulses over a range of ~4000 dispersion measures, in real time. An alert is sent for any strong detections, and the incoming data chunk is saved to disk. The astronomer team inspects the detections and follows up further.

Funded Extensions:

- 1: Buffer 12hr of data in transient machine
- 2: Through single-pulse and/or periodicity searches, detect all pulsars with mean flux density > 100 μ Jy and $DM < 1000$ pc/cm³.
- 3: Manually run commensal to certain imaging
- 4: Automate commensal running

3.4 Dedicated Searching for Fast Radio Bursts and Pulsars (SC4)

Top-level science requirement:

“During dedicated observing, ARTS will detect all 1-millisecond-duration FRBs of flux $S_{1400} > 0.16$ Jy and dispersion measure $DM < 5000$ pc/cm³, occurring in the Apertif field of view”

Latency, self-trigger, storage requirement as for SC3.

3.4.1 Dedicated Transient Searching Use Case

Usage narrative: A team of astronomers want to search for **dim** radio bursts. They apply for and get a block allocation of time. These are scheduled in ~12/24 hr sessions. Before each session, they can use a command-line/scriptable tool to set a telescope/beamformer/back-end settings schedule for the session. This schedule points the telescopes, and sets up and calibrates ~37 compound beams (CBs). These are coherently added in 37 sets of 12 grating lobes (“8GR8”) producing 300MHz bandwidth split in ~1024 channels, at ~50μs sampling. These 444 coherent beams (at total 2.1Tbps) are searched for transient pulses over a range of ~4000 dispersion measures, in real time. An alert is sent for any strong detections, and the incoming data chunk is saved to disk. The astronomer team inspects the detections and follows up further.

Funded Extensions:

- 1: Buffer 12hr of data in transient machine
- 2: Through single-pulse and/or periodicity searches, detect all pulsars with mean flux density > 30 μJy and DM < 1000.
- 3: Further down-sample data, upload to the Apertif/Apropos archive, where the entire survey will be kept publicly accessible.

4 Science Requirements (L0)

4.1 Science requirements for (at least) SC1

Science Requirement	SR-0.1-revB
Relates to Science Case	SC1, SC2
Description	Single central TAB: the system shall be capable of forming one full-sensitivity, full-bandwidth, linear or circular-polarization tied array beam (TAB) in the field centre.
Priority	1
Input	Up to 12x1 stream from central Apertif compound beam (CB)
Processing	Create full-sensitivity, full-bandwidth tied array beam, in the field centre, in linear polarization.
Output	Single TAB
Constraints	Beamforming is statically at the center of the telescope pointing
Verification	Detection of “standard” pulsar PSR B1919+21 at expected SNR and polarization profile. Verified by expert user.
Error Handling	
Other comments	<ul style="list-style-type: none"> • Should allow for (de)selection of telescopes that are to be included in a TAB • One central TAB is enough because a large FoV is not needed for the point sources here. • revB: added circular polarisation

Science Requirement	SR-0.2-revA
Relates to Science Case	SC1,SC2,SC3,SC4
Description	TAB bandwidth: the bandwidth shall be 300 MHz
Priority	1
Input	Single-TAB: 12x1 stream from central CB Multi-TAB: 12x37 streams for all CBs
Processing	
Output	300 MHz of bandwidth at a central frequency between 1280—1600 MHz
Constraints	As free of RFI as possible
Verification	Detection of a pulsar across entire 300MHz band
Error Handling	
Other comments	Contiguous is preferred, but RFI cluttered bands can be skipped. Applicable from single-TAB to survey multi TABs (37, 444)

Science Requirement	SR-0.3-revB
Relates to Science Case	SC1
Description	Single-TAB default spectral resolution: the default spectral resolution of the single central TAB shall be 0.5-2.0 MHz.
Priority	1
Input	12x1 stream from central CB
Processing	Channelise
Output	300 MHz of bandwidth divided over 0.5-2.0 MHz (can be 0.781 MHz) channels
Constraints	
Verification	Pulsar detected in all channels
Error Handling	
Other comments	This resolution enables initial pulsar timing

Science Requirement	SR-0.4-revA
Relates to Science Case	SC1,SC2
Description	Single-TAB adjustable further spectral resolution: the spectral resolution of the single central TAB shall also be adjustable between 1 and 2^n MHz with $n = 1, 2, \dots, 5$.
Priority	1
Input	12x1 stream from central CB
Processing	Channelize with adjustable channel widths
Output	TAB with channels adjustable in n^2 MHz widths (1, 2, 4, 8, 16, 32 MHz)
Constraints	
Verification	Bright known pulsar detected with same SNR in all widths
Error Handling	
Other comments	These wider, 2^n MHz are required for VLBI and high-precision timing In real-time folding (SR-0.14-revA) this TAB itself is not saved. In LEAP timing this TAB (SR-0.16-revA) is saved.

Science Requirement	SR-0.5-revA
Relates to Science Case	SC1,SC2
Description	Single-TAB frequency: the lower-edge frequency of the 2^n MHz subbands (SR-0.4) will be tunable to 1 MHz to coincide with LEAP and EVN observing.
Priority	1
Input	12x1 stream from central CB
Processing	Channelize with adjustable channel frequency
Output	TAB with subband edges tuneable at 1 MHz (or better yet, 250 kHz)
Constraints	
Verification	
Error Handling	
Other comments	The Aperif LO synthesizer has a frequency grid of 10 MHz, so 1 MHz tuning needs to be done in the digital domain by Arts. This fits using a 1 MHz beamlets bandwidth for the Apertif BF and avoids using a mixer and a complex subband filter in the Apertif BF (see Section 8.5.2).

Science Requirement	SR-0.6-revA
Relates to Science Case	SC1,SC2,SC3,SC4
Description	Observation Set-Up Speed: the system shall be capable of re-configuring to a new observing mode and starting within 1 minute (not accounting for slewing).
Priority	1
Input	Observing-control metadata
Processing	Initialize all components and start observation
Output	Data product
Constraints	
Verification	Test on switches between different set-ups
Error Handling	
Other comments	<ul style="list-style-type: none"> • This would be a reasonable overhead on the shortest expected observations of ~5 mins. • It also allows for reasonably fast switching from SC3 to SC4 when a Fast Radio Burst has been detected during commensal searching. • No, The triggering requirement (SR-0.40-revB) does not mention interrupting Apertif X observing. It only mentions possibly interrupting beam-forming observations.

Science Requirement	SR-0.7-revA
Relates to Science Case	SC1,SC2,SC3,SC4
Description	RFI Excision: the system shall produce final data products in which <25% of the bandwidth and <10% of the observing time is lost to RFI (either self-generated or external)
Priority	1
Input	Front-end elements and/or CBs
Processing	Clip or flag all outlier signals (like in AOFlagger) before TAB/IAB forming
Output	Cleaned TABs/IABs
Constraints	
Verification	Run PRESTO RFIfind and the offline AOFlagger on cleaned and uncleaned data to measure the difference
Error Handling	
Other comments	<p>RFI excision before beam-forming is important because then only data from the effected telescope/CB is lost.</p> <p>This is mostly important for self-generated RFI, e.g., faulty elements; with 12 x 112 > 1000 front-end elements, one will often fail. A fast turnaround (close to real-time) is then needed to prevent the entire TAB from being affected.</p> <p>This may also require real-time PAF element health monitoring.</p> <p>External RFI will more likely affect all 12 telescopes and thus does not *need* pre-beamformed excision.</p>

Science Requirement	SR-0.8-revA
Relates to Science Case	SC1,SC2
Description	Spectral leakage: the spectral leakage between adjacent channels shall be less than 50 dB
Priority	1
Input	12x1 stream from central compound beam (CB)
Processing	Channelize
Output	Spectral leakage must be less than 50 dB between adjacent channels
Constraints	
Verification	Test if highly dispersed pulsar only seen in appropriate channels
Error Handling	
Other comments	Prevents strong RFI from leaking into many more bands. The SKA1 requirement is 60dB, scaled factor 10 up from there.

Science Requirement	SR-0.9-revA
Relates to Science Case	SC1,SC2,SC4
Description	Pointing accuracy: the difference between the commanded position of the TAB and its actual position shall be less than 5% of the TAB FWHM at any time.
Priority	1
Input	12x1 stream from central CB
Processing	Steer telescopes, calculate TAB weights & phases, set these.
Output	The difference between the commanded and actual position of the TAB must be less than 5% of the TAB FWHM
Constraints	
Verification	
Error Handling	
Other comments	

Science Requirement	SR-0.10-revA
Relates to Science Case	SC1,SC2,SC3,SC4
Description	Short term output-level stability: Within a channel, the output baseline noise level will not change by more than 1% of the rms level within 1 ms.
Priority	1
Input	12 streams per CB
Processing	Digitize, set levels, downsample
Output	Stable output baseline
Constraints	
Verification	
Error Handling	
Other comments	Avoid false transient detection triggers due to e.g. switching on the noise diodes (for frontend calibration), level re-sets (e.g. for AGC), etc. during production TAB/IAB. Perhaps stagger apex noise sources between telescopes, and not cut out all 12?

Science Requirement	SR-0.11-revA
Relates to Science Case	SC1,SC2
Description	Phase stability: Over the 300 MHz bandwidth, the output signal phase will not deviate from the input phase by more than 0.1 rad.
Priority	1
Input	12 x 112 front-end elements
Processing	Form CB & TAB
Output	Stable final telescope gain
Constraints	
Verification	Correlate with EPTA telescopes, determine phase offsets
Error Handling	
Other comments	

Science Requirement	SR-0.12-revA
Relates to Science Case	SC1,SC2
Description	Hour-timescale gain stability: Over the 300 MHz bandwidth, the telescope gain toward the commanded source will not decrease by more than 10% over 1 hr.
Priority	1
Input	12 x 112 front-end elements
Processing	Form CB & TAB
Output	Stable final telescope gain
Constraints	
Verification	Test on non-scintillating set of pulsars. How often will noise apex diode really be needed for front-end element calibration?
Error Handling	
Other comments	

Science Requirement	SR-0.13-revA
Relates to Science Case	SC1,SC2
Description	Time resolution: the time resolution of the voltage samples in each spectral channel shall satisfy the Nyquist sampling rate.
Priority	1
Input	12x1 stream from central CB
Processing	
Output	Nyquist sampled voltage channels at highest time resolution of 1/channel width
Constraints	
Verification	
Error Handling	
Other comments	

Science Requirement	SR-0.14-revA
Relates to Science Case	SC1
Description	On-line dedispersion and folding: the system shall be capable of doing on-line, real-time dedispersion and folding of the TAB data given the pulsar ephemeris.
Priority	1
Input	Single channelised, voltage TAB, both polarisations
Processing	Given input parameter file, determine current-time period and dispersion measure (a “polyco”). Using the provided folding/dedispersion arguments (subchannels, sub-integration times, etc), coherently dedisperse, fold, and convert to Stokes parameters the incoming TAB data.
Output	Folded, dedispersed PSR FITS file
Constraints	Run in real time
Verification	Comparison to archive PuMa2 profiles
Error Handling	
Other comments	

Science Requirement	SR-0.15-revA
Relates to Science Case	SC1
Description	Time synchronization: the system shall be capable of providing timestamps with 10 ns accuracy, at the start of the observation and every following second.
Priority	1
Input	Observatory clock
Processing	Take observatory clock, corrected using GPS, stamp data blocks and file headers
Output	Observatory clock timestamp on FITS file and final TT time each precise to <1 ns and accurate to 10 ns
Constraints	
Verification	
Error Handling	
Other comments	Fits files just stamped with Obs clock at ns precision. TEMPO to apply corrections.

Science Requirement	SR-0.16-revA
Relates to Science Case	SC1
Description	Baseband disk recording: the system shall be capable of writing 4 x 24hr of baseband data to disk.
Priority	1
Input	Single central voltage TAB at full bandwidth
Processing	Write baseband to disk in PSRDADA or VDIF format. Hold up to 4 x 24h.
Output	
Constraints	
Verification	Fold baseband data on a pulsar, compare to archive PuMall profile and real-time folded data.
Error Handling	
Other comments	

Science Requirement	SR-0.17-revA
Relates to Science Case	SC1
Description	Batch processing: the system shall be capable of taking an input source list and producing folded FITS files.
Priority	1
Input	Source list
Processing	Determine source rise and set times Convert list to telescope pointing schedule Point telescopes Per pulsar, use appropriate parameter files to dedisperse and fold in real time (SR-0.14)
Output	Folded FITS files for list of pulsars.
Constraints	
Verification	Expert astronomer compared profiles to archive
Error Handling	
Other comments	

Science Requirement	SR-0.18-revA
Relates to Science Case	SC1
Description	Expert mode: the system shall be capable of being controlled by direct commands and scripts.
Priority	1
Input	Direct command or scripts from expert astronomer.
Processing	Scheduling, pointing, dedispersion, folding.
Output	Output data FITS file
Constraints	
Verification	Set up system, detect pulsar.
Error Handling	
Other comments	

Science Requirement	SR-0.19-revA
Relates to Science Case	SC1
Description	Non-expert mode: the system shall be capable of converting a non-expert proposal to final, retrievable FITS files.
Priority	2
Input	A non-expert astronomer submits a proposal for a pulsar timing proposal. These are carried out, and the user can retrieve their FITS files from an archive or other accessible location.
Processing	Scheduling, pointing, dedispersion, folding.
Output	FITS file accessible to non-expert astronomer
Constraints	
Verification	
Error Handling	
Other comments	

Science Requirement	SR-0.20-revA
Relates to Science Case	SC1, SC2
Description	Off-centre TAB: the system shall be capable of forming a TAB away from the centre of the central compound beam (CB)
Priority	2
Input	12x1 stream from central CB
Processing	Form TAB anywhere within the central CB
Output	Single TAB
Constraints	
Verification	Detect off-center pulsar
Error Handling	
Other comments	

4.2 Further science requirements for (at least) SC2

Science Requirement	SR-0.21-revA
Relates to Science Case	SC2
Description	Streaming e-VLBI: the system shall be capable of streaming VDIF data for a central TAB to the EVN software correlator at JIVE (SFXC).
Priority	1
Input	12 x 1 central CB
Processing	Form 300MHz TAB, convert to VDIF, stream to EVN
Output	VDIF-formatted e-VLBI stream
Constraints	
Verification	Successful correlation with other EVN dishes
Error Handling	
Other comments	

Science Requirement	SR-0.22-revA
Relates to Science Case	SC2
Description	Writing VLBI to Mk5: the system shall be capable of writing VLBI data on a Mk5 unit
Priority	1
Input	12 x 1 central CB
Processing	Form TAB, convert for Mk5
Output	To Mk5 unit
Constraints	
Verification	Successful correlation with other EVN dishes
Error Handling	
Other comments	

Science Requirement	SR-0.23-revA
Relates to Science Case	SC2
Description	Multiple TABs for eVLBI: the system shall be capable of forming up to 12 TABs anywhere within the central CB, and streaming these to JIVE.
Priority	1
Input	12 x 1 central CB
Processing	Form 12 independent TABs within CB, full bandwidth
Output	12 VDIF-formatted e-VLBI streams
Constraints	
Verification	Correlate at JIVE, compare to other EVN telescopes
Error Handling	
Other comments	Could also be 10..15 TABs but 12 is the same as SR-0.41 and thus chosen here.

Science Requirement	SR-0.24-revA
Relates to Science Case	SC2
Description	Independent repointing of TABs: the system shall be capable of individually stopping, repointing and restarting within 10 seconds each of the 12 TABs while others continue unchanged.
Priority	3
Input	12 x 1 central CB
Processing	Allow for start/stop/position of a TAB to be changeable without interrupting other TABs
Output	12 independent TABs
Constraints	
Verification	Check consistency in unchanged TABs
Error Handling	
Other comments	

Science Requirement	SR-0.25-revA
Relates to Science Case	SC2
Description	Stream individual dishes for eVLBI: the system shall be capable of formatting the central CB of up to 12 dishes as VDIF, and streaming these to JIVE.
Priority	2
Input	12 x 1 central CB
Processing	Format 12 central CB of up to 12 dishes as VDIF
Output	12 VDIF-formatted e-VLBI streams, full bandwidth
Constraints	
Verification	Correlate at JIVE, compare to other EVN telescopes
Error Handling	
Other comments	

Science Requirement	SR-0.26-revA
Relates to Science Case	SC2
Description	Expert mode: the system shall be capable of being controlled by direct commands and scripts.
Priority	1
Input	Direct command or scripts from expert astronomer.
Processing	Scheduling, pointing, formatting, streaming.
Output	VDIF stream
Constraints	
Verification	Set up system, correlate at JIVE
Error Handling	
Other comments	

Science Requirement	SR-0.27-revA
Relates to Science Case	SC2
Description	Non-expert mode: the system shall be capable of converting a non-expert EVN proposal to streaming VDIF for inclusion in the EVN
Priority	3
Input	A non-expert astronomer submits a proposal for an EVN. These on/off-source observations are carried out, and the user can retrieve the logs, as well as the final correlated EVN output from an archive or other accessible location.
Processing	Scheduling, pointing, formatting, streaming.
Output	Logs and EVN images.
Constraints	
Verification	Compare to expert mode output
Error Handling	
Other comments	

Science Requirement	SR-0.28-revA
Relates to Science Case	SC2
Description	Imaging during VLBI: the system shall be capable of automated calibration and imaging of the local interferometer data for the central CB during VLBI, with the same integration, channel and resolution settings as available for stand-alone Apertif correlation.
Priority	2
Input	12 x 1 central CB
Processing	Correlate and image as for standard Apertif imaging.
Output	VDIF-formatted e-VLBI stream
Constraints	
Verification	Field comparison to regular Apertif image of CB
Error Handling	
Other comments	<ul style="list-style-type: none"> • The implementation of this requirement is only realistic if Apertif X provides interferometry. No 2nd interferometer is envisioned. • Needed for flux calibration. Needs only correlation products for central CB. Does not need full Apertif-X <i>imaging</i> pipeline. If needed, calibration could be done <i>after</i> the target run, but that would take work and time. Current functionality at WSRT widely used.

4.3 Further science requirements for (at least) SC3

Science Requirement	SR-0.29-revA
Relates to Science Case	SC3
Description	Incoherently beam-forming 8.7 deg²: the system shall be capable of incoherently adding up to 12 telescopes, to produce an Incoherently Added Beam (IAB) for each CB. These hexagonally laid-out IABs overlap at their half-power points.
Priority	1
Input	12 x 37 CBs
Processing	Per CB, form incoherent added beam (IAB) of up to 12 dishes
Output	37 IABs
Constraints	
Verification	Compare to TAB, check if SNR is sqrt(12) lower. Simultaneously detect sources at opposite edges of the IAB.
Error Handling	
Other comments	Should allow for (de)selection of any of the 12 x 37 CBs

Science Requirement	SR-0.30-revA
Relates to Science Case	SC3,SC4
Description	Survey channelization: the system shall be capable of channelizing the 300 MHz-bandwidth TABs and IABs in spectral channels of 150-300 kHz width (preferably 300MHz/1024ch).
Priority	1
Input	12 x 37 CBs
Processing	Split bandwidth in ~1024 channels
Output	Channelized TABs and IABs
Constraints	
Verification	Input a test signal, verify it is well channelized
Error Handling	
Other comments	

Science Requirement	SR-0.31-revA
Relates to Science Case	SC3,SC4
Description	Survey time sampling: the system shall be capable of producing survey TABs and IABs at ~50μs time sampling.
Priority	1
Input	12 x 37 CBs
Processing	Sample every 50 μs
Output	Time-sampled TABs and IABs
Constraints	
Verification	Input an impulsive test signal, verify it is contained in 1 sample
Error Handling	
Other comments	

Science Requirement	SR-0.32-revA
Relates to Science Case	SC3,SC4
Description	Survey polarimetry: the system shall be capable of producing survey TABs and IABs containing polarisation information.
Priority	1
Input	12 x 37 CBs
Processing	Keep X-Y voltages or form Stokes parameters
Output	Polarimetric TABs and IABs
Constraints	
Verification	Test on input signal or known pulsar
Error Handling	
Other comments	

Science Requirement	SR-0.33-revA
Relates to Science Case	SC3,SC4
Description	Transient detection: the system shall be capable of finding all individual radio bursts up to dispersion measures of 3000 pc/cm ³
Priority	1
Input	37 IABs or 444 TABs (cf. SR-0.41)
Processing	Form Stokes-I total-intensity data. Dedisperse over DM grid, with DM steps such that DM-step smearing is less than intra-channel smearing. Find bursts of varying lengths (0.1 ms – 16 ms) and shapes (box-car, exponential tail).
Output	Trigger, burst specifications (time, DM, position)
Constraints	
Verification	Test on signal pulses from strong pulsar
Error Handling	
Other comments	

Science Requirement	SR-0.34-revA
Relates to Science Case	SC3,SC4
Description	Trigger generation: the system shall be capable of generating transient-detection trigger within ~1 second from the transient's time of arrival at the dish.
Priority	1
Input	37 IABs or 444 TABs (cf. SR-0.41)
Processing	
Output	
Constraints	
Verification	
Error Handling	
Other comments	

Science Requirement	SR-0.35-revB
Relates to Science Case	SC3,SC4
Description	Dump high-resolution, full-polarisation data: the system shall be capable of self-triggering to save polarimetric (IQUV or voltage) high-resolution data (< ~50 us) to disk for further processing. This data stream shall hold 15s of data, 5s before the burst, 5s for a burst at 5000 pc/cm ³ and 5s after.
Priority	1
Input	37 IABs or 444 TABs (cf. SR-0.41)
Processing	Determine which stream to save, write to disk
Output	
Constraints	
Verification	
Error Handling	
Other comments	At 5000 pc/cm ³ dispersion introduces a ~5 s sweep over the 300 MHz bandwidth.

Science Requirement	SR-0.36-revA
Relates to Science Case	SC3,SC4
Description	Survey data 12-hr ring buffer: the system shall be capable of recording to disk 12 hours of stokes-I, downsampled, incoming data for all IABs or TABs
Priority	1
Input	37 IABs or 444 TABs (cf. SR-0.41)
Processing	Save to disk, allow retrieval of data chunk around a specific time
Output	
Constraints	
Verification	
Error Handling	
Other comments	

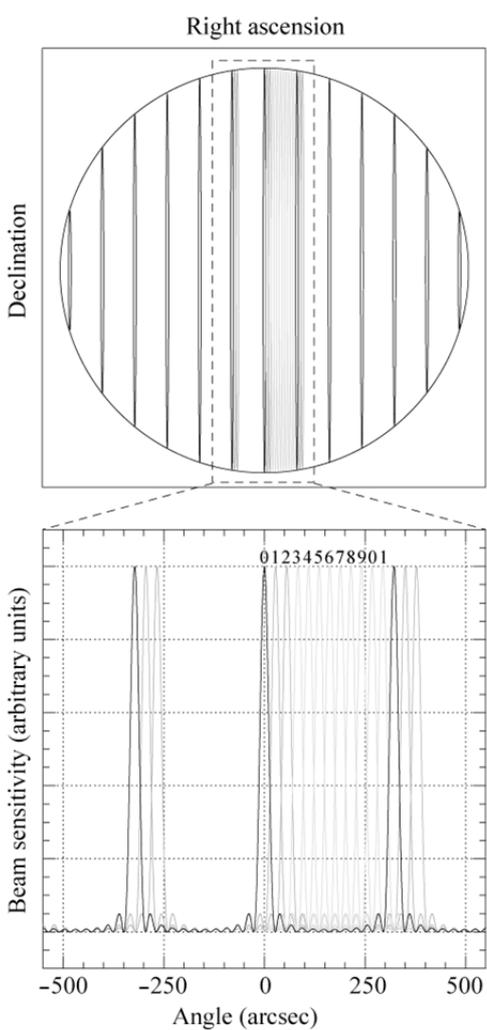
Science Requirement	SR-0.37-revA
Relates to Science Case	SC3,SC4
Description	Pulsar periodicity search: the system shall be capable of detecting all pulsars in the field with mean flux density > 100 μ Jy (SC3) and > 30 μ Jy (SC4) and DM < 1000.
Priority	1
Input	37 IABs (SC3) or 444 TABs (SC4; SR-0.41)
Processing	Use streaming data or read buffer from disk, do single-pulse and/or periodicity search, find pulsar.
Output	Pulsar position, period, dispersion measure
Constraints	
Verification	Blindly detect known pulsars down to the expected flux density limit
Error Handling	
Other comments	

Science Requirement	SR-0.38-revA
Relates to Science Case	SC3
Description	Commensal survey observing: the system shall be capable of a carrying out a transient-survey in parallel to regular imaging observations ('piggy-backing'), after set-up by an expert user.
Priority	1
Input	12 x 37 CBs
Processing	Form 37 IABs, process for transients.
Output	
Constraints	
Verification	Compare to output of dedicated IAB survey
Error Handling	
Other comments	

Science Requirement	SR-0.39-revA
Relates to Science Case	SC3
Description	Automatic commensal survey observing: the system shall be capable of automatically carrying out a transient-survey in parallel to regular imaging observations.
Priority	1
Input	12 x 37 CBs
Processing	Form 37 IABs, process for transients.
Output	
Constraints	
Verification	Compare to output of expert-user IAB survey
Error Handling	
Other comments	

Science Requirement	SR-0.40-revB
Relates to Science Case	SC3
Description	Dump buffered voltage data on trigger, and coherently beamform: the system shall be capable of buffering 5-10 seconds of voltage data from all CBs. On the self-triggering on an incoherent detection in a certain CB, that CB will be coherently beamformed as in SR-0.41, either on-line or after dumping the voltages to disk. Within a minute after the initial detection, an updated location and full-polarisation, full time-resolution (~5 us) burst profile will be produced.
Priority	1
Input	12 x 37 CBs
Processing	Buffer voltages. Freeze voltages for 1 CB on trigger. Beam-form TABs from frozen voltages (either online, or offline from disk, etc).
Output	Offline voltages or on-line TAB-formed data
Constraints	
Verification	
Error Handling	
Other comments	<p>1) Interrupting “regular” beam-forming while re-beamforming is acceptable. In this case the pipeline would need to flag out this part of the running (interrupted) observation.</p> <p>Note) This on-line re-beamforming part is <i>extremely</i> scientifically useful, but this was not realized when the funding was proposed, so it is not a milestone for funding agencies – in that sense that has priority 3.</p>

4.4 Further science requirements for SC4

Science Requirement	SR-0.41-revA
Relates to Science Case	SC4
Description	Full-sensitivity beams over 8.7 deg² : the system shall be capable of forming TABs over the full FoV
Priority	1
Input	12 x 37 CBs
Processing	For each CB, form 12 TABs at offset phases, such that the full CB is covered by the 12 grating-lobe patterns TABs
Output	12 x 37 = 444 TABs
Constraints	Should allow for (de)selection of any of the 12 x 37 CBs
Verification	
Error Handling	
Other comments	<p>In dedicated or commensal observing. Setup by expert user and automatically in survey mode.</p>  <p>Figure 2 Full CB covered by the 12 grating-lobe patterns TABs</p>

Science Requirement	SR-0.42-revA
Relates to Science Case	SC4
Description	<p>Archiving: the system shall be capable of submitting to the Apertif archive:</p> <ol style="list-style-type: none"> 1. full time-resolution data of the TABs towards the ~1000 known Northern-sky pulsars 2. down-sampled data for all 444 TABs from all ~1500 pointings.
Priority	2
Input	All TABs
Processing	<ul style="list-style-type: none"> • For #1: dedisperse, collapse to ~10MHz channels • For #2: downsample in time, frequency and bitrate to ~1ms, ~1MHz, ~2-bit stokes-I samples
Output	FITS time series for archive (stream, disk, tape, ..)
Constraints	
Verification	Detect known pulsars in archived data
Error Handling	
Other comments	All SC4 data to be archived this way, for 5-10 years after end of survey.

5 ARTS System Requirements L1

Figure 3 shows the ARTS system at L1 with the external interfaces.

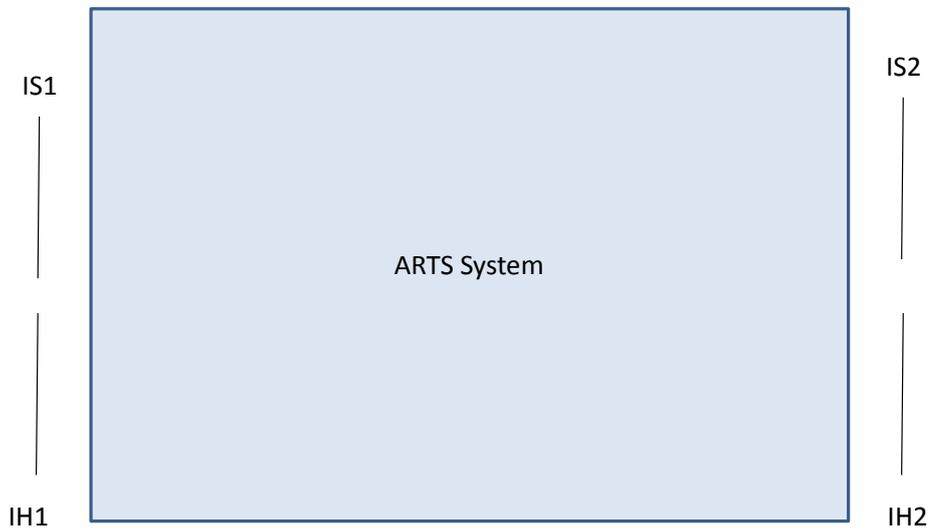


Figure 3 ARTS system at L1

The ARTS system is part of the Apertif system as shown in Figure 4. Apertif reuses the WSRT dishes but with new PAF frontends. In 2014 the Apertif BF is in the construction phase and partly exists. The Apertif X is in the detailed design phase and the ARTS is in the specification phase. Both the Apertif X and ARTS use the output of the Apertif BF.

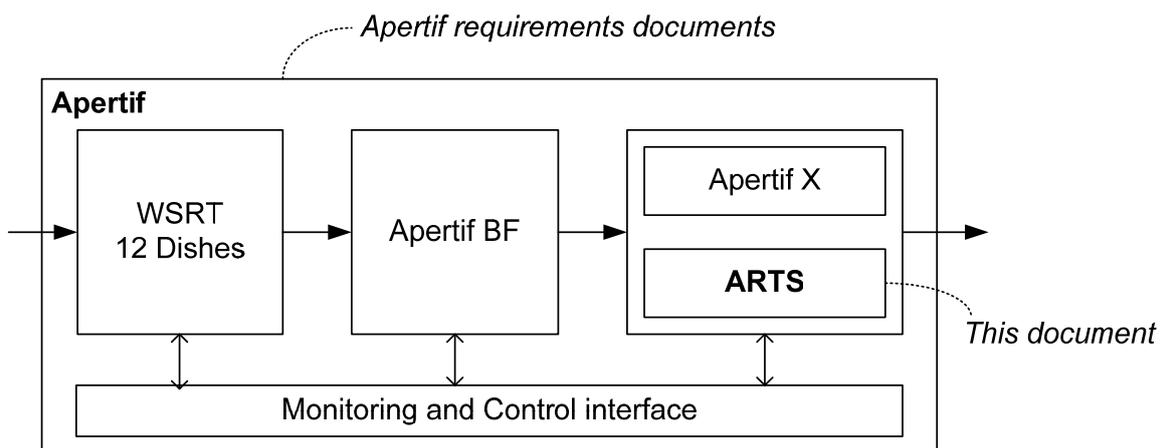


Figure 4 Top level overview of Apertif including ARTS

The pulsar surveys and transients are described in the Apertif science requirements document [1] and identified as top level operation modes in system requirements document [4]. Therefore all Apertif system requirements in [4] also apply to ARTS L1. They are not repeated here but covered by GR-1.1. This document only lists the additional system requirements that are specific to ARTS.

5.1 Definitions

Table 1 lists the various WSRT array beams that are defined within ARTS.

Beam	SC	SR	Definition
CB	-	[4]	A full bandwidth ($CB_{BW}=300\text{MHz}$), dual polarization ($N_{pol} = 2$) compound beam (CB). The CB is a 'voltage' beam. Each dish can beamform $N_{CB} = 37$ CB. In total the Apertif BF outputs $N_{dish} * N_{CB} = 444$ CBs.
IAB	3	0.29	A full sensitivity (so using one CB from $N_{dish} = 12$ dishes), full bandwidth (so same $CB_{BW}=300\text{MHz}$ bandwidth as for a CB), dual polarization ($N_{pol} = 2$) incoherent array beam (IAB). The IAB is a full Stokes 'power' beam, so the X and Y polarization need to be collected together.
TAB	1, 2, 4	0.1, 0.2	A full sensitivity (so using a CB from $N_{dish} = 12$ dishes), full bandwidth (so same $CB_{BW}=300\text{MHz}$ bandwidth as for a CB), dual polarization ($N_{pol} = 2$) coherent tied array beam (TAB). For a 'voltage' TAB the X and Y polarization can be treated separately. For a full Stokes 'power' TAB the X and Y polarization need to be collected together.
CB-12	1, 2	0.25	12 ($= N_{dish}$) central 'voltage' CBs that are input for SC1, SC2 and can be output for SC2, so one central CB per dish.
TAB-1	1, 2	0.20	1 'voltage' TAB using the central CB of each dish.
TAB-12	2	0.23, 0.24	12 ($= N_{VLBI}$) 'voltage' TABs anywhere in the central CB of each dish.
CB-444	2	0.25	444 ($= 12 * 37 = N_{dish} * N_{CB}$) 'voltage' CBs that are input for SC3 and SC4
IAB-37	3	0.29	37 ($= N_{CB}$) full Stokes 'power' IABs, covers the full FoV with IABs
TAB-444	4	0.41	444 ($= 12 * 37 = N_{gr} * N_{CB}$) full Stokes 'power' TABs, covers the full FoV with TABs

Table 1 Definition of the array beams in ARTS

Notes:

- The input to ARTS from the Apertif BF are compound beams (CB). The CB from the Apertif BF have f_0 on a 10 MHz grid and beamlet bandwidth $B_{sub} = 781250$ Hz or $B_{sub} = 1$ MHz.
- It is by sensible choice that $N_{VLBI} = N_{dish} = 12$ (SR-0.23) and that $N_{gr} = N_{dish} = 12$ (SR 0.41).
- See Table 4 and Table 8 for a list of the data loads in Gbps for the various CB, IAB and TAB beam formats.

5.2 General Requirements L1

General Requirement	GR-1.1- revA
Science Case	All
Priority	1
Relates to requirement	[4]
Description	Apertif SRS: All Apertif system requirements in [4] also apply to ARTS.

5.2.1 Apertif BF additional requirements

General Requirement	GR-1.2- revA
Science Case	All
Priority	1
Relates to requirement	SR-0.4, 0.5
Description	Apertif BF 1 MHz subbands: The default subband bandwidth of the Apertif BF beamlet output is $B_{\text{sub}} = 781250$ Hz. For combining Apertif data with data from other radio telescopes it is preferred to have $B_{\text{sub}} = 1$ MHz [10].

- SR-0.6 observation setup speed 1 min (5 min in Apertif SRS)
- SR-0.7 detect faulty PAF elements and exclude them from the CB
- SR-0.8 spectral leakage between adjacent subbands shall be less than 50 dB
- SR-0.10 short term output noise level stability
- SR-0.11 phase stability
- SR-0.12 long term CB gain stability
- SR-0.13 subband sampling rate at Nyquist rate
- SR-0.15 time synchronization between WSRT PPS and terrestrial time (TT)

5.3 Functional and Performance Requirements L1

5.3.1 SC1, SC2, SC3 and SC4

Functional Requirement	FR-1.1-revA
Science Case	All
Priority	1
Relates to requirement	SR-0.9, section 8.2
Description	Fringe stopping: The Aperitif BF performs geometrical delay tracking for the central CB at the ADC sample rate so with a delay resolution of $T_s = 1.25$ ns. If necessary any remaining fringe needs to be stopped by ARTS.

Functional Requirement	FR-1.2-revA
Science Case	All
Priority	1
Relates to requirement	SR-0.9, 0.12, 0.29
Description	Gain correction: The Aperitif BF gain is calibrated. Remaining gain differences between the CBs from different dishes or between different CB pointings need to be balanced by ARTS, if necessary to achieve the required the pointing accuracy (TAB) and expected increase in sensitivity (IAB).

Functional Requirement	FR-1.3-revA
Science Case	All
Priority	1
Relates to requirement	SR-0.15 (see section 8.3)
Description	Timestamps: The timestamp information from the Aperitif BF beamlet frames must be preserved throughout the subsequent processing in ARTS, such that at all stages the data is uniquely related to the WSRT PPS pulse that started the data processing. Lost frames or out-of-order frames must not corrupt the timestamp information of the data. The WSRT PPS can be related to terrestrial time (TT).

Functional Requirement	FR-1.4-revA
Science Case	All
Priority	1
Relates to requirement	SR-0.7
Description	Faulty element handling: In a beamformer a single faulty element can corrupt the entire beam and thus a whole observation. Therefore it must be possible to detect a faulty PAF element or a faulty CB so that it can be excluded from the measurement before the measurement is started. Faults that occur during a measurement should be flagged and logged.

Functional Requirement	FR-1.5-revA
Science Case	All
Priority	1
Relates to requirement	SR-0.7
Description	RFI removal: Outlier data caused by RFI must be removed to improve quality of the astronomical signal, and to prevent transient-detection false positives

Functional Requirement	FR-1.6-revA
Science Case	All
Priority	2
Relates to requirement	SR-0.20
Description	CB selection: The default selection for TAB-1 in SC1 and CB-12 and TAB-12 in SC2 is the central CB.
Other comments	<ul style="list-style-type: none"> • The CB selection is always the central CB, and does not need to be programmable for the CB, IAB and TAB. • Should it be possible to select the same CB for multiple IAB? No. • Should it be possible to select the same CB for multiple TAB? Yes, for TAB-12

Functional Requirement	FR-1.40 -revA
Science Case	All
Priority	1
Relates to requirement	SR-0.12-revA
Description	Hour-timescale gain stability: see SR-0.12-revA.
Other comments	<ul style="list-style-type: none"> • This is a requirement on the analog part of the APERTIF receiver chains. • This requirement is based on a L1 review update and is numbered out of order.

5.3.2 TAB for SC1, SC2 and SC4

Functional Requirement	FR-1.7-revA
Science Case	SC1, SC2, SC4
Priority	1
Relates to requirement	SR-0.6, 0.7, 0.20, 0.23, 0.24
Description	TAB weights: The TAB weights for up to the maximum of $N_{\text{TAB}} = 444$ TABs must be calculated, and the BF TAB weight should next be programmable to be able to form a TAB anywhere in the CB (SR 0.20, 0.23) and to be able to excise a CB input to by nulling it (SR 0.7, 0.41). The TAB weights must be programmed within 10 seconds (SR-0.6, 0.24). The new TAB weights may be applied when they are set, so they do not have to be applied all in parallel at the same instant.
Other comments	The TAB weights need to be programmable to keep the TAB grating lobes pointed while the earth rotates [9].

Functional Requirement	FR-1.8-revA
Science Case	SC1, SC2, SC4
Priority	1
Relates to requirement	SR-0.9
Description	TAB weights accuracy: The TAB weights must have sufficient bits to ensure accurate TAB pointing. What is the measure? The measure is that the pointing is good to 5% of the TAB FWHM (SR-0.9) [9].

Functional Requirement	FR-1.9-revA
Science Case	SC1, SC2, SC3, SC4
Priority	1
Relates to requirement	SR-0.15-revA
Description	<p>TAB/IAB/CB absolute timing: A central location for Apertif should be defined. Ideally this the same location currently used for Westerbork, in e.g. the TEMPO telescope definitions (1) and/or the VLBI definitions. For all TABs/IABs/CBs, the path length distance between this central point and the arrival locations (at the telescopes and front-ends) must be corrected.</p> <p>This should be done to absolute precision better than about 1 ns.</p> <p>For VLBI, the relative offsets in delays between TABs within one CB can next be calibrated by placing all TABs on the same calibrator for a short scan.</p> <p>This relative offset between TABs should next stay constant to about 1% of a wavelength, equaling 7 picoseconds at 21 cm, over 12 hours.</p> <p>On moving a TAB to a new pointing location with the CB, the offset should again not change by more than 1% of a wavelength, equaling 7 picoseconds.</p> <p>(1) http://sourceforge.net/p/tempo/tempo/ci/master/tree/obsys.dat</p>
Other comments	It could be that the TAB beam forming doesn't introduce <i>any</i> delay changes with time anyway; perhaps only when there is a problem with the changing weights (but then that would also affect delay vs position in the beam, and probably more substantially there).

5.3.3 SC1 Pulsar timing

Functional Requirement	FR-1.10 -revB
Science Case	SC1
Priority	1
Relates to requirement	SR 0.1, 0.2, 0.3, 0.13, 0.14 (see section 8.5)
Description	Pulsar Timing standalone: Create TAB-1 (SR-0.1, 0.2) using the default Apertif BF frequency grid for f_0 and B_{sub} (SR-0.3). The timing data of a given pulsar ephemeris is coherently dedispersed and folded (SR-0.14).
Input	$N_{dish} = 12$ central CB from the Apertif BF
Output	Folded, dedispersed PSR FITS file (SR-0.14)
Constraints	Beamforming is statically at the center of the telescope pointing
Verification	Expert-mode control (SR-0.1)
Other comments	For standalone operation a TAB channel bandwidth of $B_{sub} = 781250$ Hz or 1 MHz is acceptable (SR-0.3).

Functional Requirement	FR-1.11 -revB
Science Case	SC1
Priority	1
Relates to requirement	SR 0.1, 0.2, 0.4, 0.5, 0.9, 0.13, 0.14 (see section 8.5)
Description	Pulsar Timing with Leap: Create TAB-1 (SR-0.1, 0.2) using the Leap frequency grid for f_0 and B_{sub} (SR-0.4, 0.5). Write resulting baseband 'voltage' data to disk.
Input	$N_{dish} = 12$ central CB from the Apertif BF
Output	baseband PSR FITS / PSR DADA or VDIF file (SR-0.14)
Constraints	Beamforming is statically at the center of the telescope pointing
Verification	Expert-mode control (SR-0.1)
Other comments	The TAB center frequency f_0 of the TAB channels must fit on a 1 MHz grid relative to 0 Hz (SR-0.5, SR-0.13). The TAB channel bandwidth has to be adjustable to $B_{chan} = 1, 2, 4, 8, 16, 32$ MHz (SR-0.4) where 16 MHz is preferred (details in Section 9) . Given $B_{sub} = 781250$ Hz (or 1 MHz) from the Apertif BF the wider TAB channel bandwidths have to be synthesized by ARTS.

Functional Requirement	FR-1.12-revA
Science Case	SC1
Priority	1
Relates to requirement	SR-0.16
Description	TAB-1 disk recording: Write 4 x 24 hours of TAB-1 'voltage' data to disk in PSR DADA / PSR FITS or VDIF format .

Functional Requirement	FR-1.13-revA
Science Case	SC1
Priority	1
Relates to requirement	SR-0.18
Description	MAC expert mode: Manual control of a complete pulsar timing observation by an expert astronomer using commands and scripts.

Functional Requirement	FR-1.14-revA
Science Case	SC1
Priority	1
Relates to requirement	SR-0.17
Description	MAC batch mode: Automatically setup a series of complete pulsar timing observations from an input source list. Exact format TBD.

Functional Requirement	FR-1.15-revA
Science Case	SC1
Priority	2
Relates to requirement	SR-0.19 for SC1 SR-0.27 for SC2
Description	MAC non-expert mode: Convert an non-expert proposal into a complete pulsar timing observation. Exact format TBD.

5.3.4 SC2 VLBI

Functional Requirement	FR-1.41 -revA
Science Case	SC2
Priority	1
Relates to requirement	SR-0.1-revB
Description	Circular polarisation TABs: convert the X-Y polarisations to circular L-R polarisations
Other comments	<ul style="list-style-type: none"> See e.g., 3.3.1.5 in the document “Tied Array Adder Module, Module Specification”, “W382-02-04 dok5.doc” This requirement is based on a L1 review update and is numbered out of order.

Functional Requirement	FR-1.16-revA
Science Case	SC2
Priority	1
Relates to requirement	SR 0.4, 0.5, 0.25
Description	CB-12 output to eVLBI: Stream CB-12 data of the individual dishes in the VDIF format to the EVN software correlator (SFXC) at JIVE.
Other comments	The CB-12 input from the Apertif BF has $B_{\text{sub}}=781250$ Hz (or 1 MHz). ARTS has to reformat the frequency grid of the CB-12 to the VLBI frequency grid for f_0 and B_{chan} (SR-0.4, 0.5). The most required $f_0 = 0$ Hz and $B_{\text{chan}} = 16$ MHz?

Functional Requirement	FR-1.17-revA
Science Case	SC2
Priority	2

Relates to requirement	SR-0.28
Description	CB-12 correlation in parallel: Automated calibration and imaging of the local interferometer data for the central CB during VLBI.
Other comments	<ul style="list-style-type: none"> • This is done on the Aperif X, and only for the central CB. If the Aperif X is not running in parallel then no interferometry will be done. • Output goes to the normal Aperif imaging pipeline / archive.

Functional Requirement	FR-1.18-revA
Science Case	SC2
Priority	1
Relates to requirement	SR 0.4, 0.5, 0.21
Description	TAB-1 output to eVLBI: Stream TAB-1 data in the VDIF format to the EVN software correlator (SFXC) at JIVE.
Other comments	The CB-12 input from the Apertif BF has $B_{\text{sub}}=781250$ Hz (or 1 MHz). ARTS has to form the TAB-1 and reformat the frequency grid of the TAB-1 to the VLBI frequency grid for f_0 and B_{chan} (SR-0.4, 0.5). The most required $f_0 = 0$ Hz and $B_{\text{chan}} = 16$ MHz?

Functional Requirement	FR-1.19-revA
Science Case	SC2
Priority	1
Relates to requirement	SR 0.4, 0.5, 0.22
Description	TAB-1 output to Mk5 unit
Other comments	Removed at L1 Review.

Functional Requirement	FR-1.20-revA
Science Case	SC2
Priority	1
Relates to requirement	SR 0.4, 0.5, 0.23
Description	TAB-12 output to eVLBI: Stream TAB-12 data in the VDIF format to the EVN software correlator (SFXC) at JIVE.
Other comments	The CB-12 input from the Apertif BF has $B_{\text{sub}}=781250$ Hz (or 1 MHz). ARTS has to form the TAB-12 and reformat the frequency grid of the TAB-12 to the VLBI frequency grid for f_0 and B_{chan} (SR-0.4, 0.5). The most required $f_0 = 0$ Hz and $B_{\text{chan}} = 16$ MHz?

Functional Requirement	FR-1.21-revA
Science Case	SC2
Priority	3
Relates to requirement	SR 0.24
Description	MAC mode for TAB-12 repointing: During an VLBI observation it must be possible to start/stop/repoint a TAB without interrupting the other TABs. How should the individual TAB repointing be controlled and logged ? TBD, but similar to normal setup.

Functional Requirement	FR-1.22-revA
Science Case	SC2
Priority	1
Relates to requirement	SR-0.26
Description	MAC expert mode: Manual control of a complete VLBI observation by an expert astronomer using commands and scripts.

Functional Requirement	FR-1.23-revA
Science Case	SC2
Priority	2
Relates to requirement	SR-0.27
Description	MAC non-expert mode: Convert a non-expert VLBI proposal from VEX format into a complete observation. May need operator input.

5.3.5 SC3 and SC4

Functional Requirement	FR-1.24-revA
Science Case	SC3, SC4
Priority	1
Relates to requirement	SR-0.8, 0.30
Description	IAB and TAB channel bandwidth: The spectral channel resolution B_{chan} for the IAB in SC3 and the TAB in SC4 must be about a factor $N_{\text{chan}} = 4$ smaller than B_{sub} from the Apertif BF (SR-0.30). The spectral leakage between the channels shall be similar to the spectral leakage between the subbands (SR-0.8).

Functional Requirement	FR-1.25-revA
Science Case	SC3
Priority	1
Relates to requirement	SR-0.31
Description	IAB and TAB survey time sampling: The survey time sample period is achieved by integrating N_{int} power samples.
Other comments	<ul style="list-style-type: none"> • This downsampling may be needed for producing manageable data rates. • The straw man value for $N_{int} = 10$ to achieve $T_{Stokes} \approx 50 \mu s$ survey time sampling when $B_{chan} = 195312.5 \text{ Hz}$ (so $T_{chan} \approx 5 \mu s$). • The down sample factor N_{int} may potentially be fixed after an optimum and feasible value is determined. However being able to control N_{int} provides a means to control the IAB data rate. • The default, above, describes channelizing to B_{chan}, and then averaging N_{int} *time* samples. Decreases survey time sampling may also be achieved by first channelizing N_{int} more channels than finally needed, of width to B_{chan} / N_{int} and adding N_{int} *frequency* power samples. That has the benefit of employing a larger channelizing FFT, which better channel separation.

Functional Requirement	FR-1.26-revA
Science Case	SC3, SC4
Priority	1
Relates to requirement	SR-0.30, 0.31, 0.32
Description	<p>Form full Stokes ‘power’ beams: ARTS shall form the IAB-37 for SC3 and TAB-444 for SC4 as full Stokes (I, Q, U, V), down sampled (N_{int}) ‘power’ beams per channel (B_{chan}).</p> <ul style="list-style-type: none"> • Form the full Stokes ‘power’ beams per channel, with a spectral channel resolution B_{chan} that is about a factor $N_{chan} = 4$ smaller than B_{sub} from the Apertif BF (SR-0.30) • Form full Stokes ‘power’ beams (SR-0.32) • Integrate the IAB and TAB power data over $N_{int} \approx 10$ power samples per channel to get $T_{Stokes} \approx 50 \mu s$ survey time sampling (SR-0.31) <p>$W_{power} = 4 \text{ bit}???? [9]$.</p> <p>By contrast, for SC3/SC4, ‘voltage’ TABs are only possibly stored in a transient buffer. Otherwise they not formed, or discarded after intermediate stage.</p>
Other comments	Using the power beams allows for data reduction by integration (factor $N_{int} \approx 10$) and by only using the Stokes-I power (factor $N_{Stokes} = 4$). The voltage beams are preferred but not feasible due to limited transport and storage capacity.

Functional Requirement	FR-1.27-revA
Science Case	SC3, SC4
Priority	1
Relates to requirement	SR-0.36
Description	<p>Record to disk 12 hours of Stokes-I ‘power’ beam data: ARTS shall record to a ring buffer on disk 12 hours of IAB-37 or TAB-444 down sampled Stokes-I ‘power’ beam data. It must be possible to retrieve a data chunk around a specific time. Data format PSRFITS or comparable.</p>

Functional Requirement	FR-1.28-revA
Science Case	SC3, SC4
Priority	1
Relates to requirement	SR-0.35
Description	<p>Record to last 15 seconds of full Stokes ‘power’ beam data: ARTS shall record to a ring buffer the last 15 seconds of IAB-37 or TAB-444 full Stokes ‘power’ beam data at the full channel rate. Upon a trigger the ring buffer for that single IAB or single TAB is frozen and offloaded to disk. The ring buffer must be ready again for recording within 15 sec.</p>

Functional Requirement	FR-1.29-revA
Science Case	SC3, SC4
Priority	1
Relates to requirement	SR-0.33, 0.34
Description	<p>Search Stokes-I ‘power’ beams for transients and generate a trigger: ARTS shall immediately search the IAB-37 or TAB-444 down sampled Stokes-I ‘power’ beams for radio burst transients. Detected transients must be logged and a trigger must be generated within ~1 second from when the transient arrived at the dish.</p> <p>Incoming data is dedispersed using a dispersion plan with DM grid and steps such ensures DM-step smearing is less or equal to intra-channel smearing (following PRESTO DDplan method).</p> <p>Transient are next found as significant correlations with templates of varying lengths (0.1 ms – 16 ms) and shapes (box-car, exponential tail). This could following or using AMPP/ARTEMIS, HEIMDALL or Sclocco/vNieuwpoort codes. Method and required level of significance may need to evolve to reduce false positives.</p> <p>Triggers to LOFAR may have custom format. Generally triggers are formatted as VO Events.</p>

Functional Requirement	FR-1.30-revA
Science Case	SC3, SC4
Priority	1
Relates to requirement	SR-0.37
Description	<p>Search Stokes-I ‘power’ beams for pulsar periodicities: ARTS shall search the IAB-37 or TAB-444 down sampled Stokes-I ‘power’ beams for periodic pulsar signals. The input can be immediate streaming data (e.g., the DM trial timeseries) or read from the 12-h disk buffer after each run. Detected pulses must be logged.</p> <p>Off-line searching can be done using existing PRESTO-like CPU methods. On-line searching may need real-time folding, to be added to GPU pipeline, or stand-alone.</p> <p>PRESTO candidate plots and metrics are saved to local storage of database, for human and machine-learned inspection (TBD).</p>

5.3.6 SC3 Commensal transient search

Functional Requirement	FR-1.31-revA
Science Case	SC3
Priority	1
Relates to requirement	SR-0.29
Description	IAB-37 beamformer: ARTS shall form $N_{CB} = 37$ IAB to cover the full FoV.

Functional Requirement	FR-1.32-revA
Science Case	SC3
Priority	1
Relates to requirement	SR-0.38
Description	IAB-37 commensal with Apertif X: The IAB-37 moderate-sensitivity beams are formed in parallel to Apertif X correlation/imaging, and are searched in real time for fast transients.
Other comments	Both Apertif X and ARTS SC3 require the total CB-444 beamlet load from the Apertif BF.

Functional Requirement	FR-1.33-revA
Science Case	SC3
Priority	1
Relates to requirement	SR-0.29, 0.25, 0.12
Description	IAB programmable weights: The IAB weights can be 1 (include) or 0 (excise). The IAB weights for up to the maximum of $N_{CB} = 37$ IABs must be programmable to be able to include or excise a CB input to by nulling it (SR 0.29). Is more weight accuracy needed? Possibly, see FR-1.2

Functional Requirement	FR-1.34-revA
Science Case	SC3
Priority	1 -- 3
Relates to requirement	SR-0.40
Description	Coherently re-beamform on triggering: On detection of a transient in a certain SC3 IAB beam, the voltage buffer for the appropriate 12 CBs will be read out to: 1) write to disk for off-line coherent TAB beamforming, or 2) coherently TAB re-beamform close to real time, and search, to improve the transient localization within one minute.

Functional Requirement	FR-1.35-revA
Science Case	SC3
Priority	1
Relates to requirement	SR-0.38
Description	MAC expert mode: Manual control of a complete commensal SC3 observation by an expert astronomer using commands and scripts.

Functional Requirement	FR-1.36-revA
Science Case	SC3
Priority	1
Relates to requirement	SR-0.39
Description	MAC automated mode: Automatically control a complete commensal IAB-37 transient survey in parallel to regular Apertif X observations.

5.3.7 SC4 Dedicated transient search

Functional Requirement	FR-1.37-revA
Science Case	SC4
Priority	1
Relates to requirement	SR-0.29
Description	TAB-444 beamformer: ARTS shall form $N_{\text{TAB}} = 444$ TABs to cover the full FoV.

Functional Requirement	FR-1.38-revA
Science Case	SC4
Priority	1
Relates to requirement	SR-0.42
Description	<p>Life time archiving: ARTS will send to the Apertif archive:</p> <ol style="list-style-type: none"> 1) high time-resolution ($< \sim 50\mu\text{s}$) data of each TAB-444 beam that contain a known pulsar (~ 1000 in total). 2) down-sampled data (~ 1 MHz, ~ 1 ms, 2-bit) for all 444 beams from TAB-444, from all ~ 1500 pointings. <p>Both can be derived from the ring buffer stored under SR-0.36 .</p>

Functional Requirement	FR-1.39-revA
Science Case	SC4
Priority	1
Relates to requirement	SR-0.39
Description	<p>MAC expert and survey mode:</p> <p>Manual control of a complete SC4 observation by an expert astronomer using commands and scripts.</p> <p>Automated control over a series of SC4 observations given an input pointing grid (can be run offline before a series of observations start)</p>

5.4 Operational Requirements L1

Operational Requirement	OR-1.1
Applies to	ARTS system
Description	Maximum energy consumption/cost of energy: The ARTS system should consume less than 50 kW.
Other comments	<p>As detailed in Ch. 10, the cold water cooler at WSRT can likely cool about 50 kW worth of heat from ARTS.</p> <p>Apertif, DBBC, etc., generate about 16 kW. The cooling water machine can carry off 93 kW. Within that, 50kW for ARTS could be included.</p> <p>To rough first order: 16 Uniboards @ 400W = 6 kW 2 x 48-port (TBD) 10 GE switch = 1 kW 40 servers with RAIDS, Eth Cards @ 500 W (?) = 20 kW 80 GPUs @ 250W = 20 kW Totals 47 kW</p>

Operational Requirement	OR-1.2
Applies to	ARTS system
Description	Housing, space, location: The ARTS BF and PL equipment must fit within a space of 5-10 cabinets, to fit in the HF-cabine with the WSRT control building.
Other comments	As detailed in Ch. 10, removal of the DZB will allow for the placement of up to of order 10 racks for ARTS, if needed.

5.5 Interface Requirements Hardware L1

Interface Constraint	IH-1.1-revA
Applies to	IH1
Science Case	SC1, SC2
Relates to requirement	SR-0.1, 0.2 for SC1 SR-0.20, 0.21, 0.21, 0.22, 0.23, 0.25 for SC2
Description	CB-12 input for central CB only: The CB-12 input for ARTS originates from $N_{PN} = N_{dish} * N_{pol} * N_{band} = 384$ beamformer nodes in the Apertif BF. The load per link is $L_{BF_link1} = L_{BF_SP1_band} = 225$ Mbps (Table 4). The total load for the central CB from $N_{dish} = 12$ dishes is $L_{BF_CB12} = 86.4$ Gbps.

Interface Constraint	IH-1.2-revA
Applies to	IH1
Science Case	SC3, SC4
Relates to requirement	SR-0.29 for SC3 SR-0.41 for SC4
Description	CB-444 input for all 37 CB: The CB-444 input for ARTS originates from $N_{PN} = N_{dish} * N_{pol} * N_{band} = 384$ beamformer nodes in the Apertif BF. The load per link is $L_{BF_link37} = L_{BF_SP37_band} = 8.325$ Gbps (Table 4). The total load for $N_{CB} = 37$ CB from $N_{dish} = 12$ dishes is $L_{BF_CB444} = 3.2$ Tbps.

Interface Requirement	IH-1.3-revA
Applies to	IH1
Description	Data transport technology: The interface with the Apertif BF consists of at least $N_{sp} = 24$ 1GbE links and $N_{link} = 384$ 10Gb links.

Interface Requirement	IH-1.4-revA
Applies to	IH1
Description	Monitoring and control: The interface with Westerbork M&C is over a GbE link (TBD)

Interface Requirement	IH-1.5-revA
Applies to	IH1
Description	Expert user access: The interface with expert users is through the ASTRON LAN

Interface Requirement	IH-1.6-revA
Applies to	IH2
Description	Data transport technology: The interface with JIVE consists of $\sim 5 \times 10$ GbE links, the Apertif Archive consist of $\sim 1 \times 10$ GbE link, SURFSara/UvA/Internet consists of $\sim 1 \times 10$ GbE link to the SURF backbone.

5.6 Interface Requirements Software L1

Interface Requirement	IS-1.1-revA
Applies to	IS1
Description	Observation meta data: Essential telescope information is continuously provided to ARTS through the M&C system (similar to current set-up PuMall, TBD).

Interface Requirement	IS-1.2-revA
Applies to	IS1
Description	Meta data: Meta data in frames must indicate its origin, such that e.g. swapped cables can be diagnosed immediately independent of the beam data.

Interface Requirement	IS-1.3-revA
Applies to	IS1
Description	Data products: EK: What observation data is in de payload or file?

Interface Requirement	IS-1.4-revA
Applies to	ARTS system
Description	Data format: How is the data packed in the payload or file and what meta information is in the header ????

Interface Requirement	IS-1.5-revA
Applies to	IS2
Description	Transport protocol SC2: VDIF over TCP/IP, UDP (TBD)

Interface Requirement	IS-1.6-revA
Applies to	IS2
Description	Transport protocol SC3, 4: 1) Triggers sent out as VO Events over TCP/IP 2) Pulsar search candidate files PRESTO format rsynced to user 3) Data to archive sent as PSR FITS files over gridFTP (TBD)

5.7 Non-functional Requirements L1

Non-Functional Requirement	NFR -1.#
Applies to	ARTS system
Description	Availability: Percentage of Up-time, ARTS as a whole, > 90% MTBF, MTTR

Non-Functional Requirement	NFR -1.#
Applies to	ARTS system
Description	Required Life Time: ~ 5 years

6 ARTS System Requirements Level 2

Figure 5 shows the ARTS system at L2 with the external interfaces and the internal interfaces between the beamformer (BF) and the pipelines (PL).

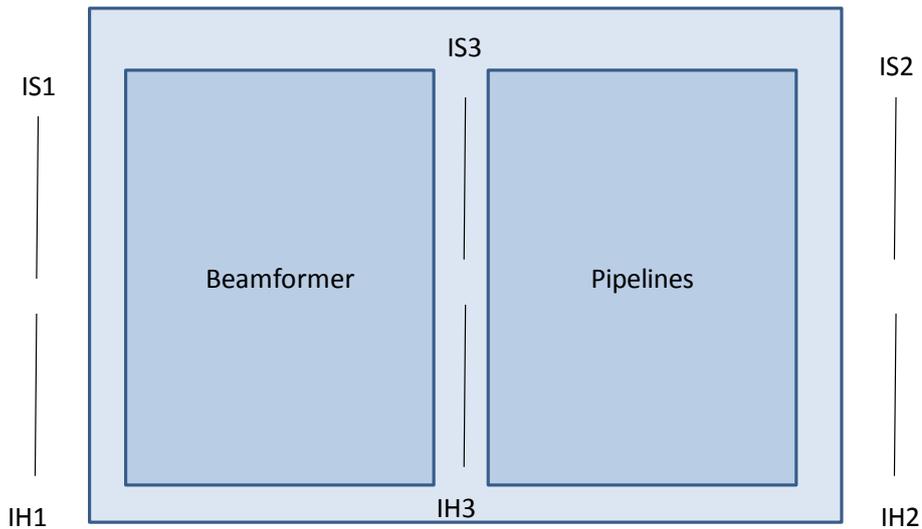


Figure 5 ARTS system at L2

6.1 General Requirements L2

6.2 Functional Requirements L2

6.2.1 SC1, SC2, SC3 and SC4

Functional Requirement	FR-2.1-revA
Science Case	All
Priority	1
Relates to requirement	FR-1.4-revA
Description	Defining a faulty element: An element is marked faulty when it produces > 5 sigma (TBD) more noise (stable or periodic) compared to its past performance or the set of other elements

Functional Requirement	FR-2.2-revA
Science Case	All
Priority	1
Relates to requirement	FR-1.7-revA
Description	Weight calculation: The BF weights needs to be calculated (potentially on a control machine) for use in the Beamformer.

Functional Requirement	FR-2.3-revB
Science Case	All
Priority	1
Relates to requirement	FR-1.8-revA
Description	Weight accuracy: The accuracy of the BF weights needs to be ≥ 5 bits [9] [Stefan Wijnholds to calculate what a pointing accuracy of 5% of the TAB FWHM implies for the TAB beamformer weights.]

Functional Requirement	FR-2.4-revB
Science Case	All
Priority	1
Relates to requirement	FR-1.8-revA
Description	Weight update rate: The BF weights needs to be updated every ≥ 1.22 seconds [9]

6.2.2 SC1 Pulsar timing

Functional Requirement	FR-2.5-revA
Science Case	SC1
Priority	1
Relates to requirement	SR 0.1, 0.2, 0.3, 0.13, 0.14
Description	Pulsar Timing standalone: As described in FR-1.9.

Functional Requirement	FR-2.6-revA
Science Case	SC1
Priority	1
Relates to requirement	SR 0.1, 0.2, 0.4, 0.5, 0.9, 0.13, 0.14
Description	Pulsar Timing with Leap: As described in FR-1.10.
Other comments	Parallel to writing baseband data to disk, the Leap data can be processed as normal pulsar timing data as described in FR-1.9. The Leap-mode can also be used for other type of observations that need baseband data.

Functional Requirement	FR-2.7-revA
------------------------	-------------

Science Case	SC1
Priority	1
Relates to requirement	SR-0.16
Description	TAB-1 disk recording: As described in FR-1.11.

Functional Requirement	FR-2.8-revA
Science Case	SC1
Priority	1
Relates to requirement	SR-0.18
Description	MAC expert mode: As described in FR-1.12.

Functional Requirement	FR-2.9-revA
Science Case	SC1
Priority	1
Relates to requirement	SR-0.17
Description	MAC batch mode: As described in FR-1.13.

Functional Requirement	FR-2.10-revA
Science Case	SC1
Priority	1
Relates to requirement	SR-0.19 for SC1 SR-0.27 for SC2
Description	MAC non-expert mode: As described in FR-1.14

Functional Requirement	FR-2.11-revA
Science Case	SC1
Priority	1
Relates to requirement	SR-0.14
Description	All folded dedispersed PSR FITS files will be stored long term (> 20 years) in an online accessible storage facility.
Other comments	Online access is restricted to members of the project until the data becomes public. Expert astronomers will have direct access to the data via a machine with full OS to allow scripted processing of the PSR FITS files.

Functional Requirement	FR-2.12-revA
Science Case	SC1
Priority	1
Relates to requirement	SR-0.14
Description	Plots of the folded profile will be generated and added to the online accessible storage facility.
Other comments	

Functional Requirement	FR-2.13-revA
Science Case	
Priority	
Relates to requirement	
Description	
Other comments	

6.2.3 SC2 VLBI

Functional Requirement	FR-2.14-revA
Science Case	
Priority	
Relates to requirement	
Description	
Other comments	

6.2.4 SC3 Commensal Transient Search

Functional Requirement	FR-2.15-revA
Science Case	
Priority	
Relates to requirement	
Description	
Other comments	

6.2.5 SC4 Dedicated Transient Search

Functional Requirement	FR-2.16-revA
Science Case	
Priority	
Relates to requirement	
Description	
Other comments	

6.3 Interface Requirements Hardware L2

Interface Requirement	IH2.1
Requirement Type	Data transport technology
Description	
Applies to	IH3
Other comments	10 GBit Ethernet, optical, copper etc

Interface Requirement	IH2.2
Requirement Type	I/O rates
Description	See Table 8
Applies to	IH3
Other comments	

6.4 Interface Requirements Software L2

Interface Requirement	IS2.#
Requirement Type	
Description	Data Products
Applies to	IS3
Other comments	

Interface Requirement	IS2.#
Requirement Type	
Description	Transport Protocol
Applies to	IS3
Other comments	

6.5 Operational Requirements L2

Operational Requirement	OR-2.1-revA
Requirement Type	
Description	Maximum energy consumption/cost of energy
Applies to	BF
Other comments	

Operational Requirement	OR-2.2-revA
Requirement Type	Housing requirement, space requirement
Description	
Applies to	BF
Other comments	

Operational Requirement	OR-2.3-revA
Requirement Type	Maintenance requirements BF
Description	
Applies to	BF
Other comments	

Operational Requirement	OR-2.4-revA
Requirement Type	Maximum energy consumption/cost of energy PL
Description	
Applies to	PL
Other comments	

Operational Requirement	OR-2.5-revA
Requirement Type	Housing requirement, space requirement PL
Description	
Applies to	PL
Other comments	

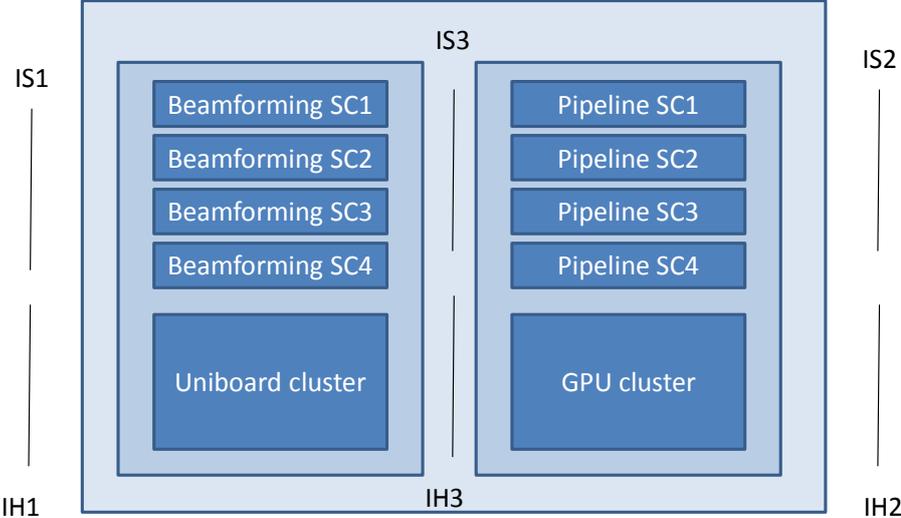
Operational Requirement	OR-2.6-revA
Requirement Type	Maintenance requirements PL
Description	
Applies to	PL
Other comments	

6.6 Non-functional Requirements L2

Non-Functional Requirement	NFR-2.1
Requirement Type	Availability
Description	
Applies to	BF
Other comments	Percentage of Up-time, MTBF, MTTR

Non-Functional Requirement	NFR-2.2
Requirement Type	Required Life Time
Description	
Applies to	BF
Other comments	

7 ARTS System Requirements Level 3



7.1 Functional requirements Beamforming L3

Functional Requirement	FR-3.1-revA
Requirement Type	
Priority	
Description	
Applies to	
Other comments	

Functional Requirement	FR-3.2-revA
Requirement Type	
Priority	
Description	Exact requirement on CB / IAB / TAB weights: TBD
Applies to	
Other comments	

7.2 Functional requirements Pipelines L3

7.2.1 Functional requirements Pipelines L3, SC 1 (Timing)

7.2.1.1 Hardware

Functional Requirement	FR-3.3-revA
Requirement Type	
Science Case	SC1
Priority	1
Relates to requirement	SR 0.14, 0.16, 0.17, 0.18, 0.19
Description	<p>GPU cluster: The GPU cluster receives the TAB-1 output from the beamformer in 16 links with each link containing the raw voltages of 2 polarizations and 19 channels of 1 MHz. The cluster performs real-time coherent dedispersion and folding.</p> <p>The system will be kept at exact UT time via an NTP server.</p> <p>A first series of benchmarks (“Results.pdf”, email Roy Smits to Joeri van Leeuwen, 2015-04-29 15:33, to be formalized) has shown that DSPSR can coherently dedispersed the full band using ~3 K40x CPUs plus 18 modern CPU cores.</p>
Applies to	
Other comments	The GPU cluster contains the par-files. Expert users have access to perform advanced data inspection and edit par-files

Functional Requirement	FR-3.4-revA
Requirement Type	
Science Case	SC1
Priority	1
Relates to requirement	SR 0.16
Description	<p>TAB-1 disk recording: The TAB-1 recorder receives the TAB-1 input from the beamformer and stores the data on disk. The disk capacity needs to be 400 TB to hold 4x24 hours of TAB-1 voltage data. The disk can be accessed by expert astronomers to move the data to another machine. The data will be kept for 6 months, after that it will be automatically deleted.</p>
Other comments	Data-loss from disk-failure needs to be kept at a risk of < 1 disk-failure per 300 hours of observation. Recommended: 16 RAID-5 virtual drives of 26 TB each. Because of the intense IO, enterprise disks or disks of similar quality are required.

Functional Requirement	FR-3.5-revA
Requirement Type	
Science Case	SC1
Priority	1
Relates to requirement	SR 0.14, 0.17, 0.18, 0.19
Description	Pulsar archive: All pulsar archive files are automatically stored on the permanent pulsar archive. The data from each project can only be downloaded by astronomers involved in that project from their own computer via a hyperlink. Each data file will have a unique hyperlink. Initially all the current archive files (40 TB) will be stored on this archive. The ongoing pulsar observations require 2 TB additional space per year.
Other comments	The storage needs to be extremely robust. No data-loss is allowed. Recommended: RAID-6 virtual drives with an exact mirror off-site + one copy to a tape-archive.

Functional Requirement	FR-3.6-revA
Requirement Type	
Science Case	
Priority	
Relates to requirement	
Description	
Other comments	

7.2.1.2 Software

Functional Requirement	FR-3.7-revA
Requirement Type	
Science Case	SC1
Priority	1
Relates to requirement	SR 0.17, 0.18 0.19
Description	Online webpage: All observation log-files as well as the plots from the observation will be automatically uploaded to an online webpage along with a hyperlink to the actual archive data. The webpage will allow astronomers to view the log-files and the plots and to download the archive files.
Other comments	Details of the individual observations and the hyperlinks to the archive data will be accessible only for the astronomers participating in the proposal related to the observation. A summary of all observations will be accessible to the general public. Data that has been on the archive for over 1.5 years will be available to (or can be requested by) the general astronomy community via the webpage.

Functional Requirement	FR-3.8-revA
Requirement Type	
Science Case	SC1
Priority	1
Relates to requirement	SR 0.1, 0.2, 0.3, 0.13, 0.14, 0.15
Description	<p>Pulsar timing standalone: The raw voltages of each of the 16 bands will go into the GPU cluster. The GPU cluster will coherently dedisperse and fold all of the 16 bands * 2 polarizations * 19 channels = 608 datastreams in real time, using a par file. The resulting archive files will be written in 10-second intervals. The header information of the archive files are written based on the meta-data of the observation, including a timestamp of the exact start of the observation. The archive files are stored on the pulsar archive. Additionally, plots are made from the resulting archive files. The cluster also continuously (?) receives meta data from the MAC control interface about the observation (similar content as with LOFAR “parset” and/or current PuMall meta data) and stores them in a log-file.</p> <p>Several default plots will be created from the archive files. These plots, along with the log files will be automatically uploaded to the online webpage.</p>
Other comments	<p>The start of the observation will be at an exact 10-second interval. No raw voltages are stored on disk. The data-streams go directly into a memory buffer for real-time processing.</p> <p>Dedispersion and folding is performed with e.g., dspsr on each band individually with:</p> <pre>dspsr -F<Nfreq>:D -E <par-file></pre> <p>where <Nfreq> is the number of frequency channels and <parfile> is the name and location of the par-file. This information is provided by the control interface.</p> <p>Data management could be done using PSRDADA, folding using DSPSR. http://sourceforge.net/projects/psrdada/ http://sourceforge.net/projects/dspsr/</p> <p>(Or using Guppi DAQ: https://github.com/demorest/guppi_daq)</p> <p>Any changes to this pipeline software, and its configuration files (clock files, par files) should be logged and under version control.</p>

Functional Requirement	FR-3.9-revA
Requirement Type	
Science Case	SC1
Priority	1
Relates to requirement	SR 0.14, 0.17, 0.18, 0.19
Description	<p>Par-files: Each pulsar observation will be dedispersed and folded using input from a small text-file containing timing information of a particular pulsar, called a par-file. Default par-files will be stored in a fixed location on the GPU cluster. Only expert astronomers will be given access to the default par-files.</p> <p>Any changes to the default par files should be logged and the original par files stored as backup.</p> <p>Non-expert astronomers can request to have their own par-file placed on the GPU cluster. He/she can specify this par-file to be used in the processing when scheduling.</p> <p>http://tempo.sourceforge.net/ref_man_sections/tz-param.txt</p>
Other comments	

Functional Requirement	FR-3.10-revA
Requirement Type	
Science Case	SC1
Priority	1
Relates to requirement	SR 0.1, 0.2, 0.4, 0.5, 0.9, 0.13, 0.14, 0.15, 0.16
Description	<p>Pulsar Timing with Leap: The raw voltages of each of the 16 bands are written directly to the TAB-1 disk recorder. All 19 channels and 2 polarizations in each band are weaved together into 1 datafile using the PSR DADA format. The header of each dada-file is written based on the meta-data of the observation, including a timestamp of the start of the observation. The voltages are written in separate directories for each band in files of 10 seconds long (760 MB per file). The directories are named according to the date in the format <code>yyyymmdd</code>. Meta data of the observation is written in the header of the data file. The files are named as: <code>yyyymmdd_hhmmss.sssss.ext</code>, where <code>yyyymmdd</code> is the year, month, day of the start of the observation, <code>hhmmss XXX sssss</code> is the number of seconds after the start of the observation and <code>ext</code> is the proper extension for the chosen dataformat (e.g. <code>dada</code>, <code>vdif</code>).</p> <p>The datafiles are OS-level accessible for expert astronomers involved in the project who will move the datafiles to an external machine.</p>
Other comments	<p>The start of the observation will be at an exact 10-second interval. The baseband data can be processed as normal pulsar timing data (as described in FR-1.9) parallel to the observation.</p> <p>Disk recording can utilize a similar protocol as used in <code>puma2cmi.py</code> for PuMa-II which runs as a daemon, listening to a fixed socket for data to record.</p>

Functional Requirement	FR-3.11-revA
Requirement Type	
Science Case	SC1
Priority	1
Relates to requirement	SR-0.18
Description	<p>MAC expert mode: The astronomer can select a default list (in straight ascii, xml, etc) of pulsar observations. Each line in this list contains: pulsar name, RA/Dec, frequency setup, length of observation, frequency channel width and sampling time. Scheduler software allows selection of pulsars from this list and allows the user to edit each of the settings. It can operate in UT-mode or ST-mode. A start-time and date of the observation needs to be given (in ST or UT, depending on the mode). The software shows when each source is visible and how often it has been observed recently, based on the information from the webpage. The selected pulsars are shown with their start and stop-times, which takes into account the telescope slew time between observations. The list of selected pulsars, along with all their settings, can be written at any time to a text-file. This text file can be edited by hand if needed and again loaded by the scheduler software.</p> <p>Once the scheduling is complete, the entire list can be saved to a controller file such that it will be observed at the given time and date. Any conflict with an existing observation will directly be reported to the user.</p>
Other comments	

Functional Requirement	FR-3.12-revA
Requirement Type	
Science Case	SC1
Priority	1
Relates to requirement	SR-0.17
Description	MAC batch mode: As described in FR-1.13.
Other comments	

Functional Requirement	FR-3.13-revA
Requirement Type	
Science Case	SC1
Priority	1
Relates to requirement	SR-0.19 for SC1 SR-0.27 for SC2
Description	MAC non-expert mode: As described in FR-1.14
Other comments	

Functional Requirement	FR-3.14-revA
Requirement Type	
Science Case	
Priority	
Relates to requirement	
Description	
Other comments	

7.2.2 Functional requirements Pipelines L3, SC 2 (VLBI)

Functional Requirement	FR-3.15-revA
Requirement Type	
Science Case	
Priority	
Relates to requirement	
Description	
Other comments	

Functional Requirement	FR-3.16-revA
Requirement Type	
Science Case	
Priority	
Relates to requirement	
Description	
Other comments	

7.2.3 Functional requirements Pipelines L3, SC 3 (Commensal search)

7.2.4 Functional requirements Pipelines L3, SC 4 (Dedicated search)

7.2.4.1 Hardware

Functional Requirement	FR-3.17-revA
Requirement Type	
Science Case	SC4
Priority	1
Relates to requirement	SR-0.33
Description	<p>GPU cluster for transient dedispersion and detection: The GPU cluster receives the TAB-444 output from the beamformer.</p> <p>A first series of benchmarks, https://app.box.com/s/2tfx3zb631epj20acfct5e97b05jpdmm to be formalized, has shown that the Sclocco OpenCL transient pipeline can dedisperse and search data with 20,000 samples/s, at 1024 channels over 300MHz, for 2000 DM steps, with 41 AMD HD7970 GPUs or 111 NVIDIA K20X GPUs .</p> <p>To be updated with Titan X values in July 2015.</p>
Applies to	
Other comments	

Functional Requirement	
Requirement Type	
Description	
Applies to	
Other comments	

7.3 Interface Requirements Software / Hardware L3

Are these needed here or does the L2 suffice?

7.4 Functional Requirements GPU cluster L3

Functional Requirement	
Requirement Type	
Description	Storage
Applies to	System/Pipeline/Node
Other comments	

Functional Requirement	
Requirement Type	
Description	Host functionality
Applies to	System/Pipeline/Node
Other comments	Other functions than receiving data from the beamformer and transport to GPUs?

Functional Requirement	
Requirement Type	
Description	Computational power
Applies to	System/Pipeline/Node
Other comments	

Functional Requirement	
Requirement Type	
Description	Memory requirements
Applies to	System/Pipeline/Node
Other comments	

Functional Requirement	
Requirement Type	
Description	Processing real time or off line.
Applies to	System/Pipeline/Node
Other comments	Control network, low latency interconnect

Functional Requirement	
Requirement Type	
Description	Communication between nodes
Applies to	System/Pipeline/Node
Other comments	Control network, low latency interconnect

7.5 Operational Requirements GPU-cluster L3

Operational Requirement	
Requirement Type	
Description	Accessibility
Applies to	System/Pipeline/Node
Verification	
Other comments	Login/head nodes

Operational Requirement	
Requirement Type	
Description	Compliancy with technology choices (nvidia, intel; openCL, Cuda)
Applies to	System/Pipeline/Node
Other comments	Login/head nodes

Operational Requirement	
Requirement Type	
Description	Maximum energy consumption/cost of energy
Applies to	System/Pipeline/Node
Other comments	Login/head nodes

Operational Requirement	
Requirement Type	
Description	Housing requirement, space requirement
Applies to	System/Pipeline/Node
Other comments	Login/head nodes

Operational Requirement	
Requirement Type	
Description	Maintenance requirements
Applies to	System/Pipeline/Node
Other comments	

7.6 Non-functional Requirements GPU-cluster L3

Non-Functional Requirement	
Requirement Type	
Description	Availability
Applies to	System/Pipeline/Node
Other comments	Percentage of Up-time, MTBF

Non-Functional Requirement	
Requirement Type	
Description	Expected life time
Applies to	System/Pipeline/Node
Other comments	

8 Appendix: ARTS feasibility study

8.1 Sensitivity

8.1.1 Full sensitivity

Both Apertif X and ARTS TAB/IAB require full sensitivity so they both use the full $CB_{BW} = 300$ MHz bandwidth (SR-0.2), all $N_{dish} = 12$ available dishes and both, $N_{pol} = 2$ polarizations (SR-0.1).

8.1.2 Flux

The limit flux for Apertif with $N_{dish} = 12$ dishes is

$$2kT_{sys} / (N_{dish} * A_{eff}) = 2 * 1.38e-23 * 70 K / ((\eta=0.7) * 12 * \pi/4 * (25 m)^2) = 4.7 * 10^{-25} W/m^2/Hz = 47 Jy$$

For SNR=6, $N_{pol}=2$, transient duration=5 ms, transient bandwidth=300 MHz, a burst of

$$47 Jy * SNR / \sqrt{N_{pol} * B * \tau} = 47 Jy * 6 / \sqrt{2 * 3E8 * 5e-3} = 0.16 Jy$$

is thus detectable (cf. Top Level SR for SC4, page 18).

For Incoherent addition of 12 dishes (SC3) this limit is $\sqrt{12}$ higher, 0.55 Jy (cf. Top Level SR for SC3, page 17).

8.2 Delay tracking and fringe stopping

Delay tracking to compensate for the geometrical delay between the dishes in the WSRT array is done directly after the digitizer at the input of the Apertif BF [7]. If necessary any remaining fringe needs to be stopped by ARTS (FR-1.1).

- The Apertif BF delay tracking has a resolution of 4 ADC samples is $4 T_s = 5$ ns, so the remaining error is ± 2.5 ns. For a $B_{sub} = 1$ MHz beamlet band width this yields a phase incoherence of $2.5 \text{ ns} * 1 \text{ MHz} * 360 \text{ degrees} \sim 1$ degree. For Apertif X a phase incoherence of 14 degrees yields a decorrelation of 1 % [7].
- The Apertif BF delay tracking is done for the central CB. The delay error for a CB at the edge is $R/c * \sin(HPBW)$, where $R=3$ km, $c=3e8$ and $HPBW=1.6$ degrees. This yields a delay error of 280 ns that needs to be tracked by ARTS [9].

8.3 Time synchronization

The relation between WSRT PPS and terrestrial time (TT) is known. The Aperitif BF uses one pulse from the WSRT PPS to start a block sequence number (BSN). The initial BSN is linked to the first digitized sample, so with ADC sample period $T_s = 1.25$ ns accuracy, and increments for every block of N samples. The BSN has sufficient range to count for many years. The BSN is a timestamp that is passed along with the data throughout the data path processing and across interfaces to preserve the synchronization to the initial PPS pulse for all subsequent data.

8.4 Simultaneous operating modes

Table 2 provides an overview of the ARTS SC in relation to Apertif X imaging mode.

Mode	Simultaneous with Apertif X	Remarks
SC1	No.	Pulsar timing requires control of dish pointing to the pulsar and control of measurement time per pulsar.
SC2	No, but requires (with priority 2) X for the central CB-12.	The VLBI needs to be the master of the dish pointing, because it needs to be set the same for all participating telescopes. The local interferometry data for the central beam provides information for a wider FoV and is used to measure the SEFD (SR-0.28).
SC3	Yes, commensal IAB-37 (with priority 1) while Apertif X does full X for CB-444	SC3 IAB-37 is explicitly intended to run simultaneously with Apertif X imaging observations (SR-0.38, 0.39).
SC4	No.	If SC4 TAB-444 and Apertif X can operate in parallel then SC3 is not needed. It depends on the observation goals whether SC4 or the Apertif X is the master of the dish pointing.

Table 2 Relation between the ARTS operation modes and Apertif imaging

Note that:

- Apertif requirement SYS-03-09 in [4] states that the Apertif operation modes in SYS-02-02 [4] are mutually exclusive. One of those modes is the “Tied Array Mode”, which is SC1/SC2 in this ARTS document. Apertif X and ARTS SC1 and stand-alone SC2 do not have to be able to run together. The ARTS SC3 requirement from this document supersedes SYS-03-09.
- The four science case of ARTS do not run simultaneously.
- If Apertif X and ARTS TAB/IAB operate at the same time then they will observe the same part of the sky, because there will be no reduced array mode using less dishes.

8.5 RF frequency grid and subband bandwidth

The Apertif BF can measure in a RF range from 1130 MHz to 1720 MHz (SYS-04-05 in [4]). In this RF range an $RF_{BW} = 400$ MHz band can be selected with a resolution of $f_{low} = n * 10$ MHz that gets digitized. The sample frequency is $f_s = 800$ MHz. Figure 6 shows how the $RF_{BW} = 400$ MHz from f_{low} is down converted and subsampled. The $f_{LO1} = 2800$ MHz and is fixed. The f_{LO2} can be adjusted in steps of 10 MHz. By setting f_{LO1} appropriately the f_{low} maps to 0 Hz and f_0 to 50 MHz.

$$f_{low} = f_{LO1} - f_{LO2} - f_s = f_{LO1} - 3600 \text{ MHz}$$

$$f_0 = f_{low} + 50 \text{ MHz}$$

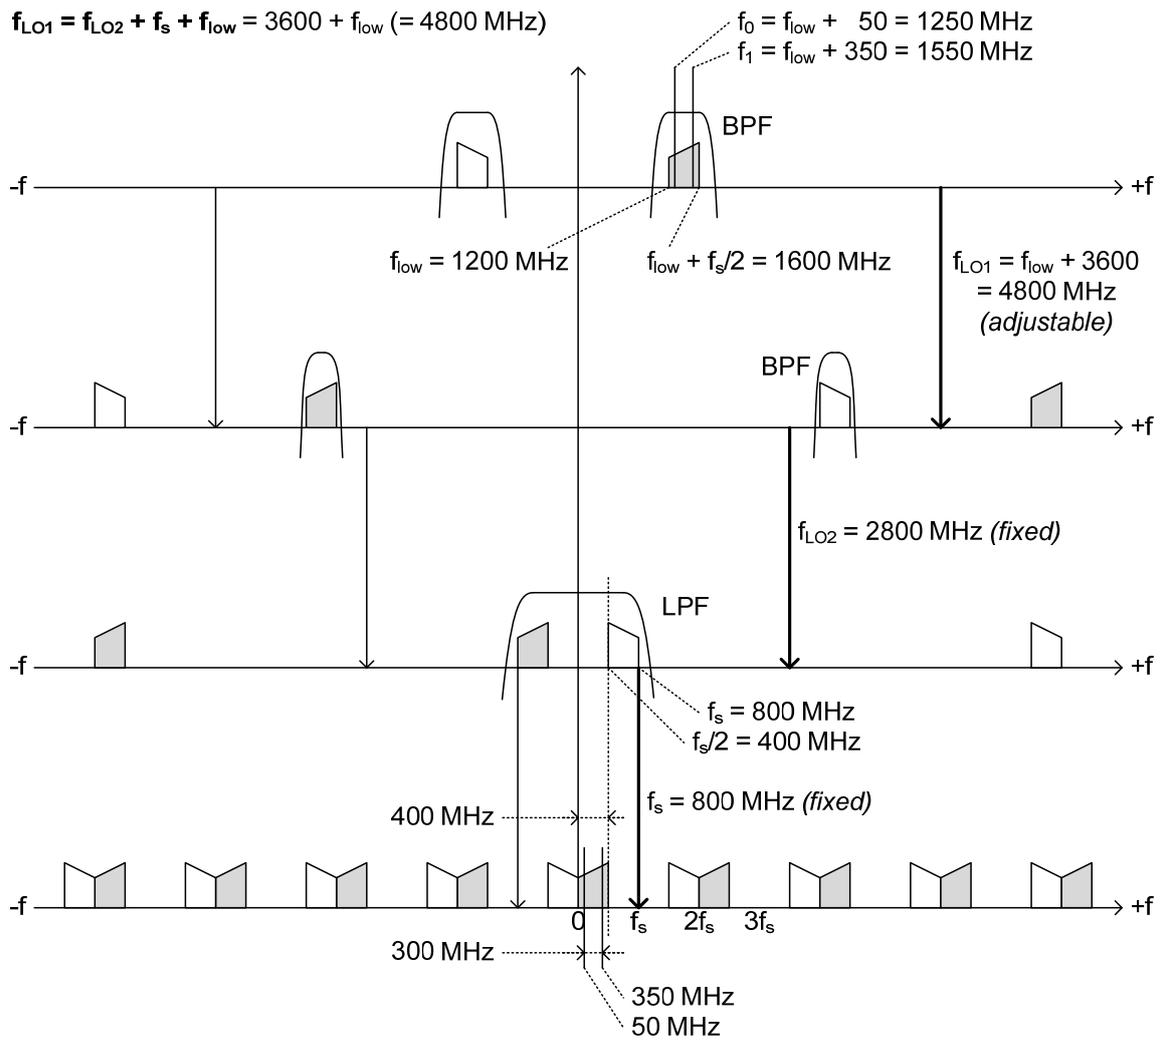


Figure 6 Apertif BF mixers and subsampling scheme for 300 MHz beams, from e.g. $f_0 = 1250 \text{ MHz}$ to $f_1 = 1550 \text{ MHz}$

A digital filterbank separates the selected $RF_{BW} = 400 \text{ MHz}$ band into $N_{sub} = N/2 = 512$ subbands using an $N = 1024$ point FFT as shown in Figure 7. Hence the subband bandwidth of the Apertif BF is $B_{sub} = RF_{BW} / N_{sub} = 400M/512 = 781250 \text{ Hz}$.

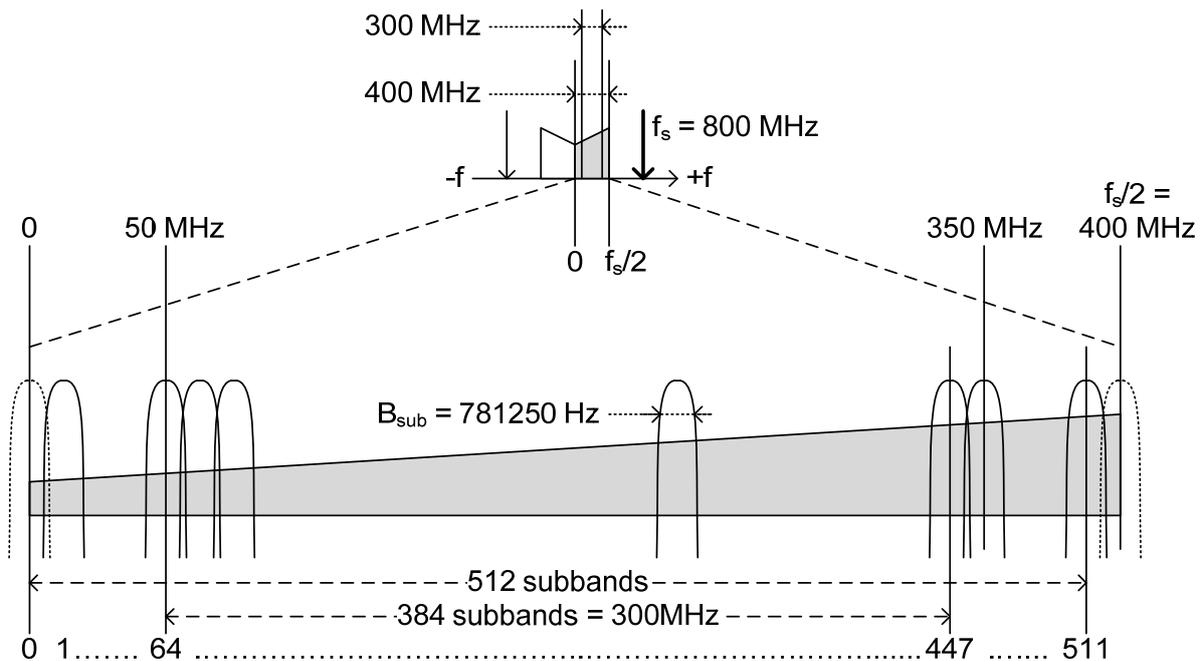


Figure 7 Poly phase filterbank in Apertif BF with $N_{\text{sub}} = 512$ subbands

In the compound beam the subbands of inputs from the PAF elements are beamformed into a beam that is called a beamlet. The beamlet therefore also has a bandwidth of $B_{\text{sub}} = 781250$ Hz. The compound beam output uses a selection of $N_{\text{sel}} = 384$ subbands to achieve the required $CB_{\text{BW}} = 300$ MHz full bandwidth of a compound beam. From Figure 6 and Figure 7 it is clear that f_0 of the CB maps to 50 MHz that falls at the centre of subband 64 ($= 50 \text{ MHz}/B_{\text{sub}}$). Hence for the CB typically subbands 64 to 447 are selected to have $CB_{\text{BW}} = 300$ MHz. As shown in Figure 7 with $B_{\text{sub}} = 781250$ Hz only subbands $0 + n \cdot 32$ fall on the 1 MHz grid (SR-0.5). For $B_{\text{sub}} = 1$ MHz all subbands would fall on the 1 MHz grid (SR-0.5).

Standalone science cases SC3 and SC4, searches for bursts and pulsars, can work with a range of grids and bandwidth of B_{sub} . Science cases SC1 (pulsar timing) and SC2 (wide field VLBI), however, need to combine the Apertif output data with data from other telescopes – then the exact grid for f_0 , and the subband bandwidth B_{sub} are important.

8.5.1 Remarks on tuning f_0 within 10 MHz

- Using a Hilbert transform it is possible to obtain the analytic signal of the digitized real input signal. The analytic signal only has the positive frequency components. Mixing the analytic signal with a complex carrier wave (CW) can be used to shift the band to adjust f_0 in the digital domain. It is unclear whether the filterbank nodes in the Apertif BF have sufficient spare processing capacity to implement the adjust of f_0 .
- Alternatively f_0 may be adjusted after the Apertif BF, but that then requires first using a synthesis filterbank to stitch adjacent beamlets back to the required beam channel bandwidth or to stitch all $N_{\text{sel}} \cdot B_{\text{sub}}$ beamlets to get back to $CB_{\text{BW}} = 300$ MHz.

8.5.2 Remarks on changing B_{sub}

- From discussions with the VLBI team and the Leap team the conclusion was that $B_{sub} = 1$ MHz is preferred over $B_{sub}=781250$ Hz. Still if the Apertif BF can only output beamlets with $B_{sub}=781250$ Hz then the approach for ARTS to achieve wider channels would be to synthesize the $N_{sel}=384 * B_{sub}=781250$ Hz beamlets back to $CB_{BW}=300$ MHz and then separate into 16MHz channels. For the full email report on the discussions see section 9.
- It is feasible to change the subband bandwidth of the Apertif BF to 1 MHz by means of using an $N = 2^5 5^2 = 800$ point FFT in the subband filterbank. Similar $B_{sub} = 2$ MHz is also feasible and may also be just allowed because Apertif requires $B_{sub} < 2$ MHz (SYS-33-01 in [4]). Changing B_{sub} to be a factor of 1 MHz does require a parameter value change of N and redesign of the filterbank in the Apertif BF to support a factor 5 in N instead of only powers of 2.
- From an implementation point of view the Apertif BF may even support wider B_{sub} as listed in Table 3, provided that the BF processing load can be divided equally over $N_{band} = 16$ processing nodes. The number of subbands per BF node N_{sel}/N_{band} has to be an integer and that causes that CB_{BW} may become too large to process are too small to be acceptable for the SC. Note that $B_{sub} = 32$ MHz is not possible because $f_s/2/B_{sub} = 800 \text{ M} / 2 / 32 \text{ M} = 12.5$ is not an integer.

B_{sub} [MHz]	$N / 2$	N_{sel}	N_{sel}/N_{band}	CB_{BW} [MHz]
0.1953125	2048	1536	96	300
0.25	1600	1200	75	300
0.781250	512	384	24	300
1	400	304 288	19 18	304 288
2	200	160 144	10 9	320 288
4	100	80 64	5 4	320 256
8	50	48 32	3 2	384 256
16	25	16	1	256
18.75	21.333 (X)	16	1	300
32	-	-	-	-

Table 3 Impact of changing B_{sub} on CB_{BW} and the number of subbands per BF node in the Apertif BF

- A typical band resolution for SC1 and SC2 is 1 MHz (SR-0.3, 0.4, 0.5). Note that $B_{sub} = 1$ MHz and $B_{sub} = 781250$ Hz share a common channel bandwidth of $B_{chan} = 15625$ Hz, if they are put through a channel filterbank that separates 64, respectively 50 channels.
- For SC1 a band of **18.75 MHz** ($= CB_{BW}/N_{band} = 300 \text{ MHz}/16$) was mentioned, with the idea that such a bandwidth (used in Apertif BF) could easily be synthesized. Such a band, however, cannot be directly formed through an integer division of $RF_{BW} = 400$ MHz and as such it is not attractive.

Summary:

The constraint on the beamlet bandwidth imposed by the Apertif BF is that $B_{sub} = 781250$ Hz. With effort it is feasible to change this to $B_{sub} = 1$ MHz in the Apertif BF [10]. Any other smaller or wider beam bandwidth will need to be derived by ARTS from a range of Apertif BF beamlets.

8.5.3 Channelization

For SC3 and SC4 respectively the IAB and TAB require a channel bandwidth of < 300 kHz (SR-0.30). These channels could be obtained with one PFB or two PFB in series:

1. Table 3 also shows that the < 300 kHz channels could already be separated at subband level by the PFB in the Apertif BF. However this does require a parameter value change of N and it is not clear whether the filterbank nodes in the Apertif BF have sufficient spare processing power and memory to fit this.
2. The alternative is that ARTS further splits the beamlet bandwidth $B_{sub} = 781250$ Hz (or 1 MHz) into 4 channels using a second PFB stage. A factor 4 is preferred over a factor 3, because $4 = 2^2$ which fits better with the FFT in a poly phase filterbank (PFB). The channels at the edge of the subband will get disturbed by the transfer of the subband filter. This could be avoided by using oversampled subbands, but that would again impact the PFB in the Apertif BF because it is critically sampled at the Nyquist rate.

Remark: The subband PFB has real input so subband 0 carries DC and $f_s/2$, but the channel PFB has complex subband data as input so from a frequency point of view channel 0 in each subband is no different than the other channels in that subband.

8.5.3.1 Spectral leakage

The spectral leakage requirement (SR-0.8) on Bchan are similar to Bsub of the Apertif BF .

Overlap (AG)?

SR-13 implies that oversampling instead of critical sampling is not required?

8.6 Apertif BF output interface

8.6.1 Number of parallel links

The Apertif BF outputs $N_{cb} = 37$ compound beams for $N_{dish} = 12$ dishes and $N_{pol} = 2$ polarization so in total for $N_{sp} = 24$ signal paths (SP). The Apertif BF processing for each SP is done by a subrack with 4 UniBoards. Per UniBoard there are 4 beamformer processing nodes (PN) and that is why $N_{band} = 16$. The whole Apertif BF contains $N_{sp} * N_{band} = 384$ beamformer processing nodes. The full bandwidth $CB_{BW} = 300$ MHz of the compound beams is processed in $N_{band} = 16$ separate frequency bands to be able to divide the processing load over multiple nodes. The compound beam data output interface of the Apertif BF is defined by:

$$CB_{pol, dish, band}[k]$$

where subscript *pol* has range $0:N_{pol}-1$, subscript *dish* has range $0:N_{dish}-1$, subscript *band* has range $0:N_{band}-1$ and array index *k* has range $0:N_{CB}-1$ where $N_{CB}=37$ is the number of compound beams in Apertif. The subscript indices indicate parallel links and the array index contains serial data on the link.

The Apertif BF output interface is fixed in the sense that the output is carried over $N_{link} = N_{PN} = N_{pol} * N_{dish} * N_{band} = 2 * 12 * 16 = 384$ parallel physical links. Therefore to create one full bandwidth CB, IAB or TAB over all SP the input information from all $N_{link} = 384$ physical links is needed. In theory the Apertif BF could regroup the $N_{band} = 16$ bands by redistributing the $CB_{pol,dish,band}[k]$ data within each subrack to get full bandwidth $CB_{pol,dish,beam}[k']$. However the UniBoards probably do not have sufficient IO and processing capacity to do this. Furthermore the required number of beams $N_{CB} = 37$ does not map evenly on $N_{band} = 16$ BF processing nodes.

8.6.2 Apertif BF compound beam load

Table 4 list loads that can be defined regarding the Apertif BF output interface assuming that the Apertif BF outputs beamlets with $W_{beamlet} = 6$ bits.

Load	Equation	Value	Description
L_{BF_SP1}	$= CB_{BW} * N_{complex} * W_{beamlet}$	3.6 Gbps	Load for 1 SP
$L_{BF_SP1_band}$	$= L_{BF_SP1} / N_{band}$	225 Mbps	Load for 1 SP per band (= per BF node)
$L_{BF_SP37_band}$	$= N_{CB} * L_{BF_SP1_band}$	8.325 Gbps	Load for $N_{CB} = 37$ SP per band (= per BF node)
L_{BF_CB1}	$= N_{pol} * L_{BF_SP1}$	7.2 Gbps	Load for 1 CB (= 2 SP, $N_{pol} = 2$)
L_{BF_CB12}	$= N_{dish} * L_{BF_CB1}$ $= N_{PN} * L_{BF_SP1_band}$	86.4 Gbps	Total load from $N_{dish} = 12$ dishes, for 1 CB
L_{BF_CB444}	$= N_{CB} * L_{BF_CB12}$ $= N_{PN} * L_{BF_SP37_band}$	3.2 Tbps	Total load from $N_{dish} = 12$ dishes, for $N_{CB} = 37$ CB
L_{BF_link1}	$= L_{BF_SP1_band}$	225 Mbps	Link load for 1 SP per band (= per BF node)
L_{BF_link37}	$= L_{BF_SP37_band}$	8.325 Gbps	Link load for $N_{CB} = 37$ SP per band (= per BF node)

Table 4 Load definitions for the Apertif BF output interface with $W_{beamlet} = 6$ bit

8.6.3 Physical interfaces

Each beamformer processing node (PN) has a one 1GbE interface and up to four 10Gb interfaces. Per dish there are 144 fibers available to transport data and control between the dishes and the central control room. Currently 32 fibers for data are planned and 1 fiber for control.

8.6.3.1 1 GbE interface

For SC1 and SC2 only one CB is needed from each dish. This single CB is also available within the full CB-444 data output via the 10Gb interface, but it could also be duplicated and offloaded via the 1GbE control interface. Per beamformer processing node the load for 1 SP is $L_{BF_link1} = 225$ Mbps. The total load of 1 CB is $L_{BF_CB12} = 86.4$ Gbps. This could be collected using 1GbE/10GbE switches.

8.6.3.2 10 Gb interface

For SC3 and SC4 all $N_{CB} = 37$ CB are needed from each dish, so the total Apertif BF output load of $L_{BF_CB444} = 3.2$ Tbps (load per physical link is $L_{BF_link37} = 8.325$ Gbps). This can only be transported via

the 10Gb interface of each BF processing node. Using one 10Gb interface is sufficient, but it is an option to use a second 10Gb interface per PN. For this the spare fibers could be used and double the number of optical modules would be required.

8.6.4 Apertif BF output to Apertif X and ARTS

Both Apertif X and ARTS SC3 and SC4 require the full CB-444 output from the Apertif BF. There are several options to transport this data:

- One 10Gb link per PN to one set of UniBoards that implement both Apertif X and ARTS
- Apertif X on one set of UniBoards may pass on the beamlet data to another set of UniBoards for ARTS
- The Apertif BF may duplicate the beamlet output and have one 10Gb to a set of UniBoards for Apertif X and another 10Gb link to another set of UniBoards for ARTS.
- A big 10GbE switch could duplicate the Apertif BF output to separate applications.

8.7 Beam definitions

Figure 8 shows a sketch that relates the field of view (FoV) of the different beams that are defined in ARTS.

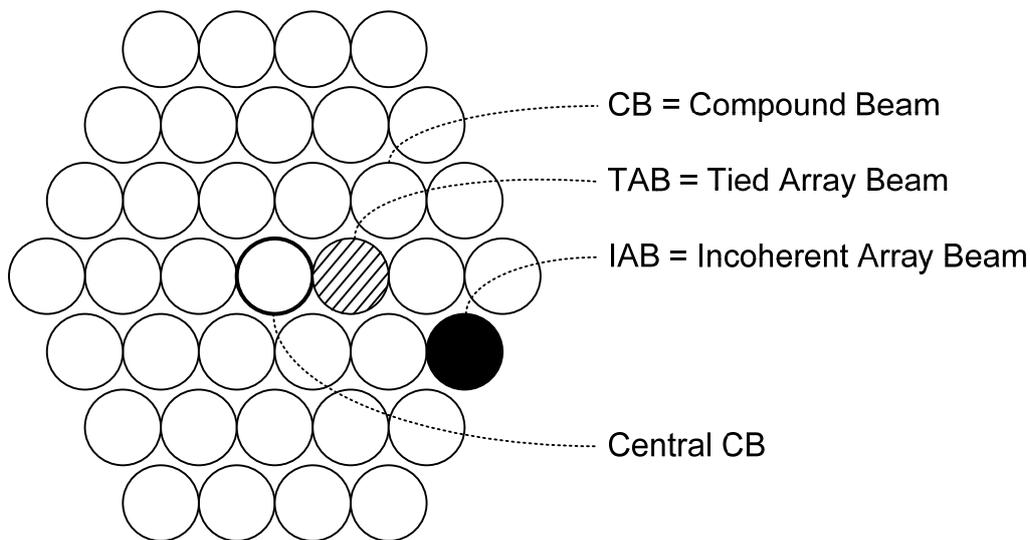


Figure 8 FoV for the CBs, the TABs and the IABs

8.7.1 Compound beams (CB)

The compound beams (CB) are formed by the Apertif BF using the antenna inputs from the PAF. The Apertif BF processes both polarizations separately. The WSRT dishes have a diameter of 25 m. At 1420 MHz the FoV of the central compound beam is 0.25 deg^2 . The compound beam is slightly narrower than the main beam of a WSRT dish. The Apertif BF increases the field of view (FoV) of the telescope by a factor 30 to 8 deg^2 by creating $N_{CB} = 37$ compound beams [5], each with a slightly different pointing offset as shown in Figure 8. For Apertif imaging the FoV could be kept constant by

letting the Aperif BF output less CB for lower frequencies and more CB for higher frequencies. For ARTS $N_{CB} = 37$ for all beamlets in the $CB_{BW} = 300$ MHz range, because only then the beamlets can be synthesized back to wider bandwidth beams up to the original $CB_{BW} = 300$ MHz.

For SC2 ARTS also needs to be able to directly output stream the Aperif BF output for the central compound beam of the individual $N_{dish} = 12$ dishes in VDIF format. As shown in Table 5 this is the only science case that requires output streaming of CB data.

Nof CBs	Format	CB	SC	SR	Priority	Remarks
12	voltage	Central	2	0.25	2	Output the central compound beam for each of the $N_{dish} = 12$ dishes in streaming VDIF format.

Table 5 Overview of the required number of CB that need to be output

If the capability of forming $N_{VLBI} = 12$ TABs (see Table 7) is already in place then requirement SR-0.25 in Table 5 to stream individual dishes for VLBI is also met by creating the TABs using unit weight for the CB from 1 dish and weight zero for the other 11 dishes.

8.7.2 Incoherent Array Beams (IAB)

The incoherent array beams (IAB) are formed by the ARTS IAB. The IAB is defined by the Stokes vector $[I, Q, U, V]$ which consists of powers that are based on both polarizations X and Y:

$$I = \sum_{dish=0}^{Ndish-1} Power(CB_{X,dish}) + Power(CB_{Y,dish})$$

$$Q = \sum_{dish=0}^{Ndish-1} Power(CB_{X,dish}) - Power(CB_{Y,dish})$$

$$U = \sum_{dish=0}^{Ndish-1} 2 Real\{CB_{X,dish} * conj(CB_{Y,dish})\}$$

$$V = \sum_{dish=0}^{Ndish-1} -2 Imag\{CB_{X,dish} * conj(CB_{Y,dish})\}$$

The incoherent beamforming makes that the IAB has the same FoV as the CB and that the IAB does not use weights. There are $N_{CB} = 37$ CB so therefore there are also $N_{CB} = 37$ IABs. The input to the IAB are dual polarization, complex channel samples (X, Y) so $N_{pol} * N_{complex} = 4$ values and the output are also $N_{stokes} = 4$ values (I, Q, U, V). The required channel data rate is $B_{chan} = 195312.5$ Hz or 0.25 MHz (SR-0.30). The Stokes vector requires a sampling period of $T_{stokes} \approx 50$ μ s or less (SR-0.31). Hence the IAB data output rate can be reduced by a factor of $N_{int} = T_{stokes} * B_{chan} \approx 10$.

Nof IABs	Format	CB	SC	SR	Priority	Remarks
37	Stokes	All	3	0.29	1	Form IABs for all $N_{CB} = 37$ compound beams.

Table 6 Overview of the required number of IABs

8.7.3 Tied array beams (TAB)

The tied array beams (TAB) are formed by the ARTS TAB. The TAB weights coherently add the input CB from the $N_{dish}=12$ telescopes to create a beam with a FoV of 0.000015 deg^2 due to the 3 km synthesized aperture of the WSRT array [2]. The coherent addition is needed to create a narrow beam and the weighting is needed to steer the TAB. The WSRT array is sparse so the grating lobes of a TAB are as strong as the central lobe. Within the FoV of the CB the TAB has many grating lobes. The WSRT array is a linear array so the grating lobes patterns of $N_{gr} = 12$ TABs are enough to cover again the full CB FoV as shown in Figure 2. By forming $N_{TAB} = N_{gr} * N_{CB} = 12 * 37 = 444$ TABs the system covers the full FoV of 8 deg^2 (SR-0.40).

The TAB is a ‘voltage’ beam (SR-0.32) implying that the phase information is used to have coherent addition and to be able to also steer the beam. For full sensitivity the TAB uses the CB input from all $N_{dish} = 12$ dishes (SR-0.1) provided their CB data is not corrupted (SR-0.7) and the full $CB_{BW} = 300 \text{ MHz}$ bandwidth (SR-0.2). Similar as for the Apertif BF the TAB can be created independently for both polarizations. A single TAB is defined by:

$$TAB_{pol, chan} = \sum_{dish=0}^{Ndish-1} w(dish, chan) * CB_{pol, dish, chan}$$

The required accuracy of the weights w can be derived from SR-0.9. The number of weights per TAB is $N_{pol} * N_{dish} * N_{sel} * N_{chan} = 2 * 12 * 384 * 4 = 36864$, so for $N_{TAB} = 444$ TABs there are 16.4M weights. The weights need to be programmable to be able to point the TAB. The TAB position in the CB will change due to the rotation of the earth as indicated in Figure 9. Therefore the weights need to be adjusted regularly to keep the pointing in the same grating lobe. This effect needs to be accounted for when a source is tracked by adjusting the TAB weights. **The required TAB weights update rate is every 1.22 s [9] ????**

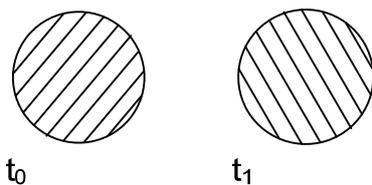


Figure 9: TAB grating lobes in CB at instant t_0 and sometime later at instant t_1 due to rotation of the Earth

Table 7 gives an overview of the number of TABs that need to be created in ARTS for the SC1, SC2, and SC4. SC3 does not use TABs.

Nof TABs	Format	CB	SC	SR	Priority	Remarks
1	voltage	Central	1	0.1	1	In the centre of the central compound beam. Stream TAB output in VDIF format. Write TAB output to Mk5 unit. Form an off-centre TAB in the central compound beam to be able to switch between both TABs, without having to move the dishes.
			2	0.21	1	
			2	0.22	1	
			1,2	0.20	2	
12	voltage	Central	2	0.23	1	Form $N_{VLBI} = 12$ TABs in the central compound beam and output stream them in VDIF format. Independent repointing of TABs within 10 s.
				0.24	3	
444	Stokes	All	4	0.41	1	Form TABs for all $N_{CB} = 37$ compound beams and for $N_{gr} = 12$ different TAB pointings per CB to cover the full FoV of 8 deg^2 .

Table 7 Overview of the required number of TABs

Similar as for IAB-37 with SC3 the TAB-444 output with SC4 is done in Stokes vector format at a sampling rate of $T_{\text{stokes}} \approx 50 \mu\text{s}$ or less (SR-0.31). Hence the data output rate can be reduced by a factor of $N_{\text{int}} = T_{\text{stokes}} * B_{\text{chan}} \approx 10$.

8.7.4 Beam data width

8.7.4.1 Voltage beam data

If the input noise 'voltage' signal has a standard deviation σ_0 then the output after quantization has a standard deviation σ given by [8]:

$$\sigma = \sqrt{\sigma_0^2 + \left(\frac{\delta}{\sqrt{12}}\right)^2} = \delta \sqrt{g^2 + \frac{1}{12}}, \quad \text{where } g = \sigma_0/\delta$$

Figure 10 shows the quantization unit δ and decision levels for 2 bit quantization. The clipping level is 2δ . With $3\sigma_0$ at the clipping level then $\sigma_0 \approx 0.6 \delta$ and $\sigma = 0.66 \delta = 1.11 \sigma_0$, so an increase in system noise power of 0.9 dB.

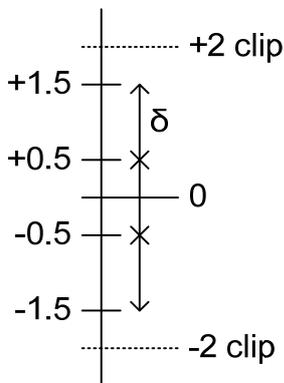


Figure 10: Decision levels and unit δ for 2 bit quantization

Using 2 bit quantization does require using automatic gain control (AGC) to keep $\sigma_0 \approx 0.6 \delta$. The AGC would have to have a time constant that does not interfere with the millisecond to second timescales seen in pulsars and transients (cf. AGC impact on PuMal). Without AGC using 4 bit quantization is sufficient to keep the input signal $3 \sigma_0$ without clipping. However in the direction of Cassiopeia-A the input signal is then too strong to fit in 4 bit. This could be compensated for using gain control via the Apertif BF weights. However instead the Apertif BF uses $W_{\text{beamlet}} = 6$ bit to provide more dynamic range which also allows some RFI.

For the ARTS TAB using $W_{\text{TAB}}=4$ bit (without AGC) or even 2 bit (with AGC) seems appropriate though.

8.7.4.2 Power beam data

Taking the power of the 'voltage' data increases the dynamic range of the data however integrating the 'power' samples reduces the dynamic range again. In fact for $N_{\text{int}} = 10$ the distribution of the 'power' samples already begins to resemble a Gaussian distribution, as can be checked with Matlab:

- Clip = 8 # Assume 4 bit quantization, so clipping range +-8
- Sigma = Clip / 3 # Put 3 sigma at the clipping level
- N=1e6
- Nint = 10
- nof_bins = 50
- X = Sigma * (randn(N,1)+i*randn(N,1)/sqrt(2))
- Hist(mean(reshape(x .* conj(x), [N/Nint, Nint])), 2, nof_bins)

For positive powers this implies that $W_{\text{power}} = W_{\text{TAB}} = 4$ bit. For $N_{\text{int}} = 10$ the 4 Lsbits can be selected. How does this change for full Stokes powers that can also become 0 and negative?

8.8 ARTS data interfaces

8.8.1 Data reduction

An important step in the data path is data reduction:

- For ARTS SC1 and SC2 the data reduction is achieved by using only the central CB-12 input and outputting 1 TAB, or CB-12, or $N_{\text{VLBI}}=12$ TABs.
- For Apertif X and ARTS SC3 and SC4 the full CB-444 input is used. For ARTS SC3 and SC4 the data reduction is achieved by integrating the signal powers over $N_{\text{int}} \approx 10$ power samples in time. A further factor $N_{\text{Stokes}} = 4$ can be achieved by only using the Stokes-I power data.

Note that the ARTS beam forming reduces the data rate by a factor N_{dish} . For SC1, SC2 with TAB-1 and SC3 with IAB-37 this indeed reduces the data rate by a factor $N_{\text{dish}} = 12$. However for SC2 with CB-12 output or TAB-12 and SC4 with TAB-444 the data rate is not reduced, because the number of beams is chosen to be equal to the number of dishes, i.e. $N_{\text{VLBI}} = N_{\text{dish}} = 12$ for SC2, and $N_{\text{gr}} = N_{\text{dish}} = 12$ for SC4. So then the data rate for the voltage CB and voltage TABs remains the same provided that $W_{\text{beamlet}} = W_{\text{chan}}$ and $W_{\text{beamlet}} = W_{\text{tab}}$.

8.8.2 Data transposes

A data transpose concerns the grouping of inputs whereby the amount of data is not reduced. This implies that another dimension of the input data needs to be distributed, because it is too much data to group and process at one node. Data can be distributed and processed separately by distributing over different frequency bands or beam directions, because these are independent.

The Apertif BF output originates from $N_{\text{link}} = 384$ links (1 link per beamformer processing node). Each link contains beamlets for all $N_{\text{CB}} = 37$ beams, but only for 1 SP (of $N_{\text{dish}} * N_{\text{pol}} = N_{\text{sp}} = 24$) and 1 band (of $N_{\text{band}} = 16$). Within ARTS the beamlets from all SP and all bands need to be collected together per CB, because for full sensitivity all SP are needed and for full $\text{CB}_{\text{BW}} = 300$ MHz bandwidth all bands are needed. This collection can be described by three transposes T_{dish} , T_{pol} and T_{band} .

- To form a 'voltage'-TAB the transpose T_{dish} is needed, but the transpose T_{pol} and T_{band} are not yet needed, because the $N_{\text{pol}} = 2$ polarizations and $N_{\text{band}} = 16$ bands can be beamformed independently.
- For a full Stokes 'power'-IAB/ TAB the transpose T_{dish} and T_{pol} are needed, but the transpose T_{band} is not yet needed, because the $N_{\text{band}} = 16$ bands can be beamformed independently.
- The transpose T_{band} is needed if the full CB_{BW} bandwidth of the CB, IAB or TAB is required.

T_{band} is mainly needed in the Science Case 3 and 4 pipelines, where the input for the incoherent dedispersion needs to be time ordered spectra over the full 300MHz band, for frequency channel c and time sample s of the Stokes I output of a certain beam:

[s1 c1] [s1 c2] [s1 c3] .. [s1 cN] [s2 c1] [s2 c2] .. [s1 cN]

. For Science Case 1 and 2, the voltage bands can be independently processed by the relevant pipelines [coherent dedispersion for SC1, re-formatting for SC2] and then combined on disk or at JIVE respectively.

At some stage (e.g. when comparing beams to see whether a signal is localized, or RFI) data for the full FoV may also need to be collected. This can be described by a transpose T_{FoV} that groups the $N_{\text{CB}} = 37$ beam directions.

8.8.3 ARTS BF data loads

Table 8 list loads that can be defined regarding the ARTS BF interface. The ARTS BF voltage beam data width is $W_{\text{tab}} = 4$ bit ????. The assumption is that without RFI (*, careful !) only in the direction of Cas-A the dynamic range of $W_{\text{beamlet}} = 6$ bits is needed, in any other direction $W_{\text{tab}} = 4$ bits are sufficient [9]. By transporting $\text{sqrt}()$ of the Stokes power values they can also be held in $W_{\text{power}} = 4$ bits ????

* Without RFI implies real-time RFI excision

Note that for voltage TAB-1 (and TAB-12/CB-12), a higher data width $W_{\text{tab}} \geq 8$ bit is preferable, as it has fewer RFI problems, while keeping data loads reasonably low.

Load	SC	Equation	Value	Description
L _{TAB1}	1	= CB _{BW} * N _{pol} * N _{complex} * W _{tab}	4.8 Gbps	Load for voltage TAB-1
L _{TAB12}	2	= N _{V_LBI} * L _{TAB1}	57.6 Gbps	Load for voltage TAB-12
L _{CB12}	2	= N _{dish} * L _{BF_CB1} * (W _{tab} /W _{beamlet})	57.6 Gbps	Load for voltage CB-12
L _{CB444_4b}	3,4	= N _{CB} * L _{CB12}	2.1312 Tbps	Load for CB-444 with W _{beamlet} = 4bit
L _{IAB1_stokes}	3	= CB _{BW} * N _{Stokes} * W _{power}	4.8 Gbps	Load for IAB-1 without integration
L _{IAB1_stokes_int}	3	= L _{IAB1_stokes} / N _{int}	0.48 Gbps	Load for IAB-1 with T _{stokes} ≈ 50 μs
L _{IAB37_stokes}	3	= N _{CB} * L _{IAB1_stokes}	177.6 Gbps	Load for IAB-37 without integration
L _{IAB37_stokes_int}	3	= N _{CB} * L _{IAB1_stokes_int}	17.76 Gbps	Load for IAB-37 with T _{stokes} ≈ 50 μs
L _{IAB37_stokes_I_int}	3	= L _{IAB37_stokes_int} / N _{Stokes}	4.44 Gbps	Load for IAB-37-I with T _{stokes} ≈ 50 μs
L _{TAB1_stokes_int}	4	= L _{IAB1_stokes_int}	0.48 Gbps	Load for TAB-1 with T _{stokes} ≈ 50 μs
L _{TAB444}	4	= N _{CB} * N _{gr} * L _{TAB1}	2.1312 Tbps	Load for voltage TAB-444
L _{TAB444_stokes}	4	= L _{TAB444} * (W _{tab} /W _{power})	2.1312 Tbps	Load for power TAB-444
L _{TAB444_stokes_int}	4	= N _{CB} * N _{gr} * L _{TAB1_stokes_int}	213.12 Gbps	Load for TAB-444 with T _{stokes} ≈ 50 μs
L _{TAB444_stokes_I_int}	4	= L _{TAB444_stokes_int} / N _{Stokes}	53.28 Gbps	Load for TAB-444-I with T _{stokes} ≈ 50 μs

Table 8 Load definitions for the ARTS BF interface (with W_{beamlet} = 6 bit, W_{tab} = 4 bit, W_{power} = 4 bit)

8.8.4 Data storage

8.8.4.1 5- 15 sec transient-capture buffer

SR-0.35 and SR-40 require a transient capture buffer:

- For SR-0.35 to store 15 sec of downsampled full Stokes IAB-37 power data (with W_{power} = 4 bit) requires (15 s) * (L_{IAB37_stokes_int} = 17.76 Gbps) / (8 bit) = 33 GByte.
- For SR-0.35 to store 15 sec of downsampled full Stokes TAB-444 power data (with W_{power} = 4 bit) requires (15 s) * (L_{TAB444_stokes_int} = 213.12 Gbps) / (8 bit) = 400 GByte.
- For SR-0.40 to store 15 sec of full rate CB-444 voltage data (with W_{beamlet} = 4 bit) or full rate TAB-444 voltage data (with W_{tab} = 4 bit) or full rate TAB-444 full-Stokes power data (with W_{power} = 4 bit) all require (15 s) * (L_{CB444_4b} = L_{TAB444} = L_{TAB444_stokes} = 2.1312 Gbps) / (8 bit) = 4 TByte.

It depends on the buffer input data rate and the buffer size where buffer can be implemented. The options for locating a transient capture buffer are:

- in ARTS PL on the GPU cluster
- in ARTS BF on UniBoard
- in Apertif BF on Uniboards (FN)

The 15 sec transient capture buffer could be implemented in the ARTS BF in DDR3 memory on the UniBoards or in the ARTS PL. This depends on whether full rate, full-Stokes data can be transported to the ARTS PL or only the down sampled Stokes-I data.

Each processing node on UniBoard has 2 DDR3 memories of each 4 GByte (tested), but could fit 8 or 16 GByte (not tested yet). If the ARTS BF for SC3 and SC4 is implemented on 16 UniBoards, then there are 128 processing nodes, so in total 1 TByte or maximum 2 or 4 TByte. The advantage of storing the beamlets on the Apertif BF front node (FN) FPGAs is that the Apertif BF has $N_{PN} = 384$ FN, so 3 time more more DDR3 memory. Apertif BF front node FPGAs beamlets can be stored. On the Apertif BF back node FPGAs ADC samples or subbands could be stored. The disadvantage of storing ADC samples, subbands or beamlets is that they are not directly available as TABs.

For SR-0.35, only the buffers that contain the IAB or TAB that had a transient trigger need to be read out.

For SR-0.40, however, the CBs are buffered. There only the CB with all the transient needs to be read out, for all 12 dishes. The CB data may then be read out directly, or after reprocessing via the IAB or TAB data path. It may potentially interrupt the ongoing IAB beamforming (with appropriate flagging for that missing real-time data).

8.8.4.2 12 hour ring buffer

To store 12 hours down sampled (integrated) Stokes-I data for TAB-444 requires $(12 \cdot 3600 \text{ s}) \cdot L_{\text{TAB444_stokes_I_int}} / (8 \text{ bit}) = 288 \text{ TByte}$ on disks.

8.8.4.3 Long-term archive

ARTS SR-0.42 requires an archive that can store TAB data for all observations done during the operational life time of ARTS. For $\sim 1\text{ms}$, $\sim 1\text{MHz}$, $\sim 2\text{-bit}$ stokes-I samples this requires $\sim 7 \text{ PB}$. The main function is to serve as long-term archive. Fast access is not essential. Data is to be downloadable for the public, through at least a simple interface, for up to ~ 10 years after the end of the surveys.

8.9 Monitoring and control

The application that runs on the FPGA can be controlled via the 1 GbE interface. Via a memory mapped master and a memory mapped bus the 1 GbE interface provides read and write access to the memory mapped (MM) registers in the application. Default every Uniboard application has some board system info. The other MM register depend on the application. Table 9 gives an overview of the main MM registers in Apertif Bf and Arts BF.

Design	MM register	Size	Rate
Apertif BF	Delay tracking coefficient		????
	Subband selection		Once per measurement
	Beamlet selection		Once per measurement
	Beamlet weights		After every noise source calibration
Arts BF	TAB weights	Section 8.7.3	Every 1.22 s [9]
	IAB on/off weights		Once per measurement
	Transient buffer control	????	

Table 9: Overview of the main MM registers in Apertif BF and Arts BF

From section 8.7.3 it follows that in total the Arts BF has 16.4 M TAB weights. The weights are complex and at least 5 bit (FR-2.3), however assume that they are packed in bytes, so 8 bit. The number of PN in the Arts BF is $N_{\text{band}}=16$ UniBoard each with 8 FPGAs so 128 PN. The TAB weights update rate needs to be at least every 1.22 s (FR-2.4), therefore assume an update rate of 1 Hz. The total bit rate for updating the TAB weights then becomes $16.4 \text{ M} * 2 \text{ (complex)} * 8 \text{ b (weight)} * 1 \text{ Hz} = 262 \text{ Mbps}$. Hence per PN the bit rate is $262 \text{ M} / 128 = 2 \text{ Mbps} = 250 \text{ kByte/s}$.

8.10 Preliminary sub system decomposition

8.10.1 Data handling tasks

The main data handling tasks for Apertif and ARTS involve:

- Control (MAC)
- Processing (DSP)
- Transport (IO)

Although initial focus often lies on the processing (e.g. algorithms) for the proper realization of the system the control and transport tasks are as important. The data handling tasks are typically mapped on the following technologies PC, GPU, FPGA and Ethernet switches. What technology to use depends on many aspects (including e.g. cost, flexibility, development time). A typical mapping is:

- PC : MAC
- GPU : streaming DSP and iterative DSP
- FPGA : streaming DSP and IO
- Switches : IO

In Figure 11, Figure 12 and Figure 13 the IO is indicated in yellow, the processing in green (FPGA) and pink (GPU) and the control in grey.

8.10.2 Apertif BF

The Apertif BF provides delay tracking (DT) on the ADC samples. The Apertif BF separates the digitized data from the PAF into subbands by means of a filterbank and then it forms beamlets for these subbands. The beamforming for one single polarization of the PAF cannot be done on a single node for the full bandwidth, so therefore the subband load has to be distributed across $N_{band} = 16$ processing nodes. The beamlet for one subband requires the input from all PAF elements, so therefore there needs to be a transpose T_{ant} that groups the subbands from all $N_{ant} = 64$ antennas. A CB is formed by a group of $N_{sel} = 384$ beamlets all with the same direction that span the $CB_{BW} = 300\text{MHz}$. Figure 11 shows the filterbank F_{sub} , the transpose T_{ant} and the beamformer (BF) that is distributed over N_{band} nodes. The MAC takes care of the proper operation, the delay tracking settings, the subband selection and the BF weights.

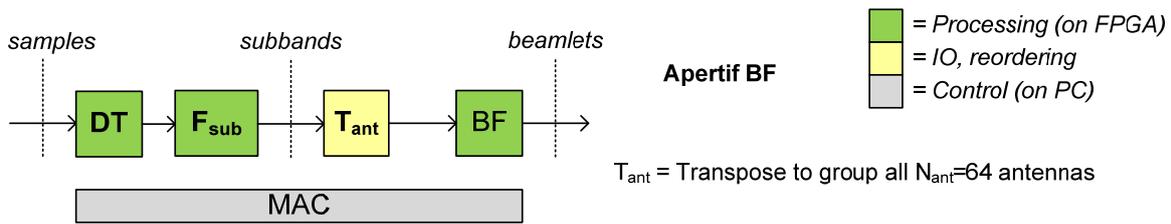


Figure 11 Apertif BF sub system

The Apertif BF for one polarization of the PAF is implemented by a subrack with 4 UniBoards, each 4 filterbank nodes and 4 beamformer nodes. In total the Apertif BF with $N_{dish} = 12$ and $N_{pol} = 2$ consists of 24 subracks with in total 96 UniBoards.

8.10.3 Apertif X and ARTS BF

Figure 12 shows a preliminary decomposition for the Apertif X and ARTS.

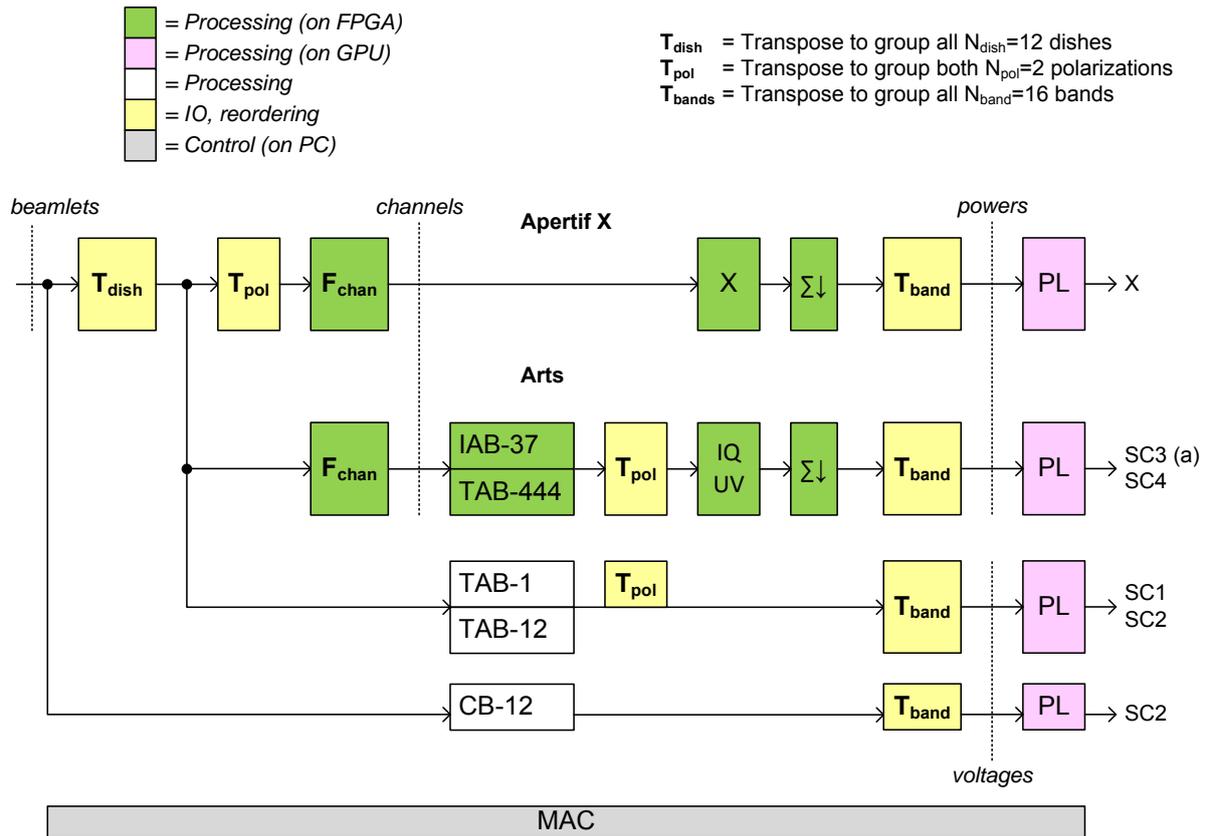


Figure 12 Preliminary decomposition of the Apertif X and ARTS sub systems

Some remarks:

- The transpose T_{dish} that collects the beamlet data from $N_{dish}=12$ dishes and the IAB / TAB processing for SC3 and SC4 operate on the maximum IO rate and are therefore planned to be implemented on FPGAs of UniBoards, rather than using a large 10GbE switch and a GPU cluster.
- The transpose T_{dish} is also needed for the Apertif X. Therefore an option is to let the Apertif X pass on the T_{dish} transposed beamlets to ARTS as shown in Figure 12. Alternatively the Apertif BF has sufficient spare IO to duplicate all output and send it via a separate set of $N_{link} = 384$ fibers to ARTS (section 8.6.3).
- The transpose T_{band} collects all $N_{band}=16$ bands to gather the full $CB_{BW}=300$ MHz. The T_{band} may also be done earlier in the data path. It probably cannot be done already in the Apertif BF (section 8.6.1).
- The full bandwidth ‘voltage’ beam data has to be transported at the full IO data rate. Reducing the IO data rate by downsampling via integration can only be done for ‘power’ beam data.

8.10.3.1 Potential for early results with Alpha-3

The quick and dirty correlator that is used for Apertif X with 1 polarization from 3 dishes can also form TABs and offload the ‘voltage’ beam data via 1GbE.

8.10.3.2 Alternative decompositions for ARTS SC1 and SC2

The TAB-1 processing for SC1 and the TAB-12 and CB-12 processing for SC2 may entirely be done on GPU hardware as indicated by the white blocks in Figure 12. The beamlet data could be provided by the Apertif BF via UDP offload over the 1GbE control network. The transpose T_{dish} is then done using 1GbE / 10 GbE switches. The total load for SC1 and SC2 is $L_{\text{BF_CB12}} = 86.4$ Gbps (see Table 4). The load per PN is $L_{\text{BF_SP1_band}} = L_{\text{BF_CB12}} / N_{\text{link}} = 225$ Mbps, so that can fit on the 1GbE control network.

8.10.3.3 Alternative decompositions for ARTS SC3

Figure 13 show two alternative decompositions for SC3. Note that the transpose T_{dish} marks the interface between the Apertif BF and ARTS. For the IAB the IQUV full Stokes powers and integration can already be done on the CB data before the beam forming. This is shown in option SC3 (b).

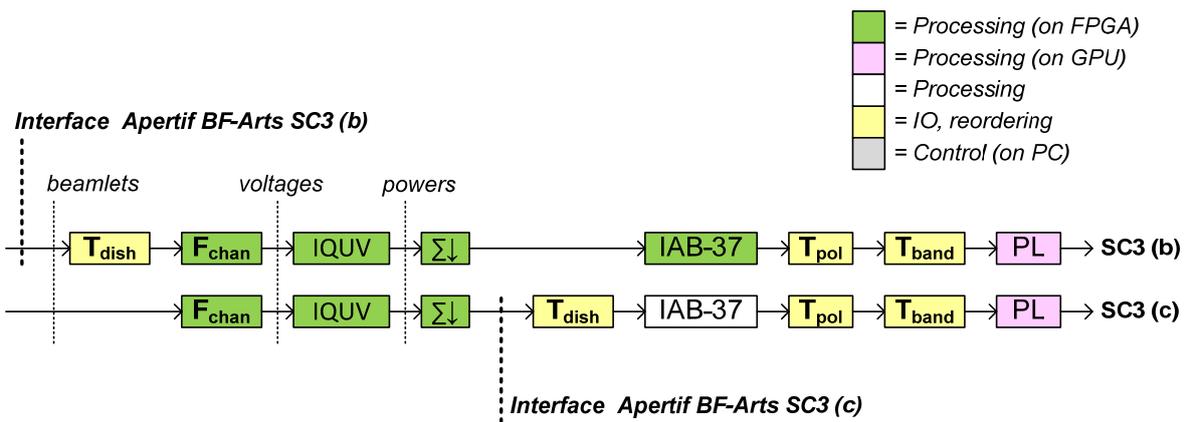


Figure 13: Preliminary alternative decompositions for ARTS SC3

If the beamlets would be separated to B_{chan} already on the Apertif BF, then only the downsampled full Stokes CB powers need to be transported to ARTS as shown in hypothetical option SC3 (c). Option SC3 (c) would also allow to directly output the CB data to the Pipelining (PL) hardware as indicated by the white blocks in Figure 13 because the total load is $L_{\text{IAB37_stokes_int}} = 17.76$ Gbps (see Table 8). The load per PN is $L_{\text{BF_SP1_band}} = L_{\text{IAB37_stokes_int}} / N_{\text{link}} = 46.25$ Mbps, so that can fit on the 1GbE control network. Unfortunately option SC3 (c) is not feasible because in the Apertif BF the processing is done per polarization. This implies that in the Apertif BF the voltages of both polarizations are not available in the same location and thus the U and V of the full Stokes values cannot be calculated in the Apertif BF. Unless per dish the X and Y polarization beamlets data would be exchanged between the two subracks using a spare 10 GbE link on the FN.

8.10.4 Preliminary mapping of ARTS on hardware

The full CB-444 beamlet load from the Apertif BF requires using the $N_{\text{link}} = 384$ 10Gb links. This large number of 10G links is probably best handled on UniBoards. The UniBoards then implement the T_{dish} transpose (rather than using a big 10GbE switch) and the IAB-37 for SC3 and TAB-444 for SC4.

The pipeline processing (PL) that includes folding, dedisperison, trigger generation, data reformatting, archiving, etc., are best done on a GPU cluster.

For SC1, 2, 3 the T_{band} transpose and links between the UniBoards and the GPUs can be done via a 1GbE network. For SC4 the T_{band} transpose needs to be done via 10G links. An issue is that the Uniboards only have the Tx link of the 10 G ports available, because the Rx link is used to receive the beamlets from the Apertif BF. If a network card or switch cannot connect to a Tx only 10G link then this could be solved using some 4 extra Uniboards to convert the single Tx link to a Tx/Rx 10 GbE link.

A large switch could remain an option for the T_{band} transpose under certain conditions. A switch that takes in 10GbE from the UniBoards but outputs 40 Gbps infiniband to the GPUs may allow for Direct Random Memory Access on the Pipeline side.

The CB-12 beamlet data is available within the CB-444 beamlet data, so UniBoard could also perform the TAB-1 for SC1 and CB-12 and TAB-12 for SC2. However any changes in f_0 and B_{sub} to interface with Leap and VLBI are better done on a GPU cluster. The CB-12 beamlet load for SC1 and SC2 could be passed on by Uniboards but could also be offloaded via the 1GbE output of the PN in the Apertif BF. The T_{dish} transpose could then be done by a 1GbE switch network that directly interfaces to a GPU cluster (rather than via UniBoards). The TAB-1 for SC1 and CB-12 and TAB-12 for SC2 can then also be done on a GPU cluster.

Apertif X, SC3 and SC4 all use the full CB-444 beamlet load from the Apertif BF. The processing load for SC3 is a factor $N_{\text{CB}} = 37$ less than for SC4. Therefore it may be possible to run SC3 commensal on the same UniBoard hardware as Apertif X. The current estimate is that the Apertif X can just fit on 16 UniBoards. Each UniBoard processes 1 of the $N_{\text{band}} = 16$ bands and gets the input from $N_{\text{SP}} = 24$ processing nodes. If more processing power is needed for the Apertif X then another set of 16 UniBoards needs to be added, so in total 32 UniBoards. The ARTS BF should also be able to fit on 16 UniBoards. **It is unclear though whether ARTS SC3 and Apertif X could fit together on 16 UniBoards.** This depends on the processing load and the amount of data transposing (storage memory, IO) that is needed towards the ARTS PL.

Alternatively there may be 16 extra UniBoards. The beamlets can be transported to these Arts UniBoards in several ways. E.g. using a second 10 GbE output port on the Apertif BF front nodes (long fibers) or by having the Apertif X UniBoards pass on the beamlet data to the Arts UniBoards (short links). With Arts on dedicated UniBoards even Arts SC4 can run in parallel to Apertif X, so whether SC3 is then still needed is TBD ????. Instead of using 16 extra UniBoards for Arts it is also feasible to use 4 extra UniBoard² to implement Arts. UniBoard² also has sufficient spare 10G ports to output the data to the GPUs. The first UniBoard² hardware is currently being tested (jun 2015).

Conclusion:

The plan is to implement the all Arts SC on the 16 UniBoards, so no direct links between Apertif BF and the Arts PL for SC1 or SC2. If possible SC3 will run commensal on the same 16 UniBoards as Apertif X.

8.11 Logging

8.11.1 Meta data

The pipeline (and derived by it, the archive; SR-042) needs access to the meta data that defines the observation data. For example the LOFAR parset has about the right amount of meta data. This involves (TBD):

- physical pointing of each dish
- pointing direction of each CB/IAB/TAB and at which time they changed
- which front-end elements are being added
- health status of the elements
- calibration output values
- which exact beamforming is used (weights, n_int, n_bits, etc)
- ...

9 Appendix: Channel bandwidth discussion for VLBI and Leap

>>> Joeri van Leeuwen <leeuwen@astron.nl> 8-7-2014 18:13 >>>

After a number of constructive and quite in-depth discussions with my instrument leads and their technical backup people, I have received feedback from the timing/LEAP (Janssen, Bassa, Smits) and the VLBI (Deller, Pogrebenko, Kettenis) teams regarding the channel bandwidth solution.

Although this is jumping ahead a little beyond my LO's I think this is a useful investigation now, as this is a potential deal breaker.

My summary: While *in principle* their software could be made to match the 781.25kHz/50 channels, it would require significant software rewriting on LEAP (~4 months) and the SFXC (~6 months), to enable the data re-routing (at the 4/8/16 MHz band overlap, you need to combine channels made out of different 781.25kHz chunks). Beyond that there is a scientific impact for LEAP long-term timing stability. Generally the conclusion was that synthesizing the 384 * 781kHz bands back to 300.0 MHz and then splitting into 16 MHz chunks would be far better.

Details -- LEAP (written by Bassa):

"In principle the LEAP project could use the 781.25kHz complex subbands produced by APERTIF. We would have to channelize these to 25 or 50 channels per subband to match 512 or 1024 channels on the 16MHz subbands that we obtain with the other telescopes.

Unfortunately this means that we lose all flexibility in choosing our channel size. Typically we prefer to have channels of 0.1MHz in size when performing the correlations and coherent addition. [JvL -- That allows for 10 us timing; while 781kHz/50 only allows for 64 us timing]

One of the major LEAP science requirements is to maintain timing stability over the monthly timescales between LEAP runs. Hence, we aim to avoid major changes to our software correlator/beamformer. Dealing with the 781.25kHz APERTIF subbands would require significant software development (several months) to incorporate APERTIF into LEAP and endanger the timing stability requirement.

For the LEAP project, the ideal case would be for APERTIF to provide 16MHz subbands to directly match the subbands obtained at the other telescopes. If that is not feasible, subbands with sizes of 8MHz, 32MHz or 64MHz are also usable."

Details -- EVN/VLBI (for SFXC, the work-horse JIVE correlator)

+ For time delay compensation the 781kHz channels are OK, not too narrow. That could still work.

+ The SFXC could then choose to split the 781kHz channels in 50 and thus match the 4 MHz bands from say Jodrell, Effelsberg, when those are split in 256.

+ This would limit the flexibility of SFXC when Wb is included.

+ These bands are incompatible with the VDIF world standard.

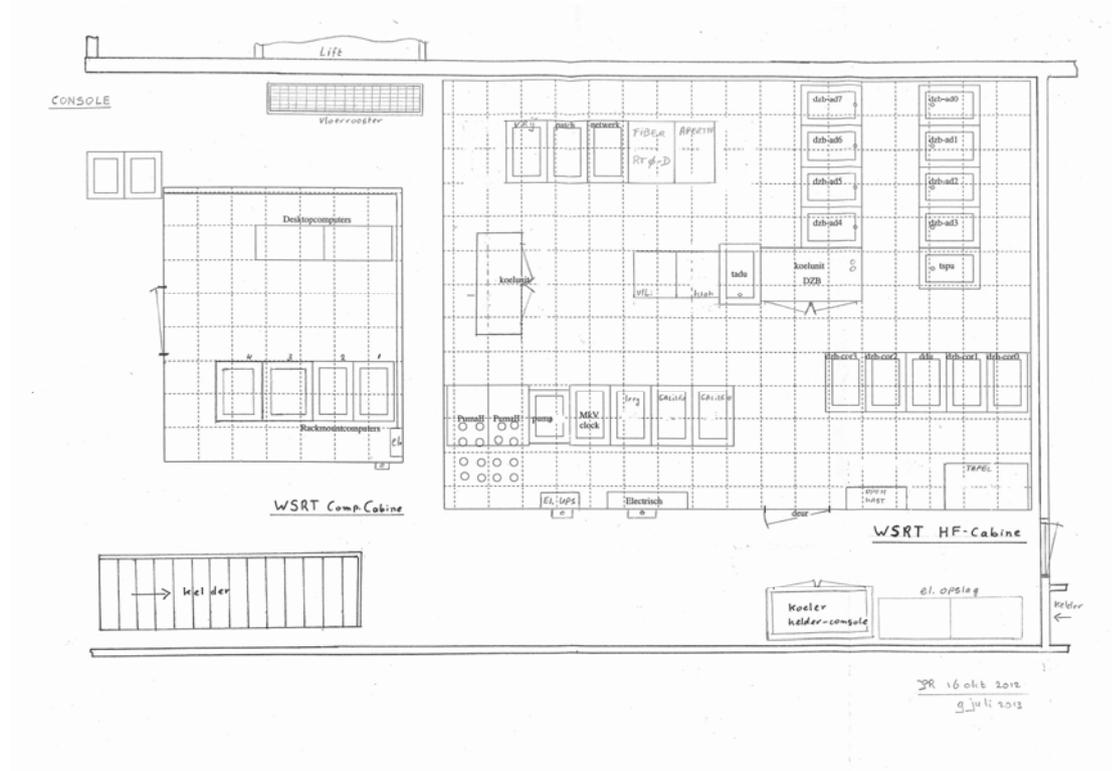
+ As there is no integer number of 781kHz Wb subbands to match to a 4MHz other-telescope subband, data rerouting and re-use would need to be implemented into SFXC. Overall this was estimated by to be ~6 months of work to test and roll out.

+ Pogrebenko and Kettenis much preferred a "stand-alone" system in which the band recombination is done before the data is streamed to SFXC and felt that could be done more quickly and independently.

10 Appendix: Power and space availability

CONCEPT

APPERTIF en ARTS backend opstelling in de WSRT hf-cabine jpr, 26mrt15



Huidige situatie WSRT hf-cabine

In de cabine wordt nu 7,7 kW ups vermogen en 31 kW huishoudnet vermogen opgenomen, en volledig in warmte omgezet. Dat is inclusief 2,4 kW voor de ventilatoren van de twee luchtkoelers. In totaal wordt er dus 39 kW, voor de hf-cabine, weggekoeld door de buiten opgestelde koudwatermachine. De cabine heeft een computer-vloer (netto 150mm) en -plafond (netto 220mm), waardoor de koude- en warme-koellucht stroomt. In deze ruimtes zijn veel obstakels, lampen, kabels, steunen en verbindinglijsten, waardoor de luchtstroom beperkt is. Hiermee is er ook een beperking van koeling.

Er staan nu 29 kasten, waarvan er 26 stuks samen 39 kW warmte produceren.

De PuMa II componenten nemen de koellucht horizontaal, daarom zitten in de vloer ervoor roosters die de koude lucht doorlaten.

Het plan is om het DZB+TaDu+PuMall te vervangen door Apertif en ARTS.

Het verwachte opgenomen vermogen bedraagt 10 en 47 kW. Samen met 6 kW vermogen van bestaande apparatuur (DBBC, klok, switch) en 15% oververmogen om te regelen (constante temperatuur regeling) is dat 73kW aan warmtelast. Daar komt het vermogen van ventilatoren nog bij.

Het werkelijke vermogen zal pas duidelijk worden als alle apparatuur functioneert, omdat deze varieert met de mate van inzet. Ik verwacht een werk-factor van 0,7, maar reken hier met 1, omdat er misschien nog apparatuur bij komt.

Het IVC (in de kelder) neemt bijna 10kW en zal ook verwijderd worden.

De netto vermogens-stijging is dan 23 kW (hf cabine) of 13 kW (gebouw)

Beperkingen:

a) lokatie

Gezien het type apparatuur, is het zeer wenselijk dat deze in de hf-cabine wordt opgesteld. De kelder, een aparte container of verspreiding over de telescoop containers is niet handig.

b) volume of vloeroppervlak

Het huidige DZB en PuMall zal verwijderd worden. APERTIF heeft 2 a 3 kasten nodig (voor 10kW) en ARTS (50kW) minimaal 5 kasten. Er is dus geen plaats probleem.

c) bekabeling

Informatie/netwerk/data kabels zijn klein en kunnen eenvoudig worden gelegd.

d) energie toevoer

De hoofd transformator is 315kVA. De WSRT gebruikt nu 200kW. Dat is zonder Embrace en zonder activiteiten op het Voormalig Kampterrein.

De eigen stroomkabels, verdelingen en stroom-filter kunnen eventueel worden vervangen door een type groter.

e) energie afvoer

De buiten opgestelde koudwatermachine kan effectief 93 kW wegkoelen. Nu doet deze ongeveer 60kW (in de winter minder, in de zomer meer ivm zon instraling).

De hoofd waterpomp, en misschien de vier kleine pompen, zullen vervangen moeten worden.

De waterleidingen naar de hf-cabine zijn voldoende groot.

De huidige watertemperatuur is 11 graden, waardoor het niet zal bevriezen en niet zal zorgen voor te droge lucht, door ontvochtiging. Dit handhaven.

De 2 luchtkasten hebben elk een fors vermogen, maar zijn beperkt door de hoeveelheid lucht die erdoor stroomt, via de computervloer.

e1) Mijn voorstel is om de computervloer geheel dicht te maken en alleen voor de voedingskabels te gebruiken.

Een of twee luchtkasten bij te plaatsen, en deze de lucht naar achteren te laten uitblazen (ipv naar beneden). Hiermee komt de koude lucht in de ruimte.

De ARTS kasten staan op een rij (of twee) en nemen zelf de benodigde koellucht.

Ik heb nog geen ideale opstelling voor ogen, maar dat gaat zeker lukken.

Misschien is het handig het verlaagde plafond te verwijderen en de data kabels in enkele goten te leggen.

Voor alle details zal een oplossing moeten komen, dat kost ontwerp- aankoop- en aanbreng- en proberen- tijd. Wie dat doet en wanneer lijkt me, dit jaar, een probleem. Maar het lijkt me wel mogelijk dit forse vermogen hier te plaatsen.

e2) Beter plan: een grote lucht-lucht koeler op het dak

Koelen met de huidige freon airco kost ongeveer 1/3 aan stroomvermogen van wat er gekoeld moet worden. Hiermee kost de aanwezige 63 kW warmte continue, 21 kW elektrisch vermogen aan freon compressor vermogen.

$21 \text{ kW} \times 24 \text{ uur/dag} \times 365 \text{ dagen/jaar} \times 0,12 \text{ euro/kWh} = 22.000 \text{ euro/jaar}$ aan stroomkosten, alleen voor de koelcompressor!

Een fraaie oplossing is om in het plafond twee maal een gat van 1x2 meter te maken (met hf-dicht rooster) en op het dak een lucht-lucht koeler te plaatsen, met daarachter een na-koeler, die op de huidige koudwatermachine werkt.

De warmte lucht van de computers is immers goed af te voeren, via een warmtewisselaar, aan de buitenlucht. (zie lijst met gemiddelde luchttemp.)

Hiermee kan bespaard worden op het elektrisch verbruik van de koelmachine.

Indien de luchtkanalen klein gekozen worden, komt daar nog het vermogen bij voor extra ventilatoren.

Een goede uitvoering van een lucht-lucht koeler is groot in omvang, en daardoor een forse investering, die alleen te realiseren is indien het gebruik een fors aantal jaren zal bedragen.

Weerstatistieken De Bilt - 2015		
Maandgemiddelde		
	Normaal	2014
Januari	3,1	5,7
Februari	3,3	6,5
Maart	6,2	8,4
April	9,2	12,1
Mei	13,1	13,2
Juni	15,6	16,2
Juli	17,9	19,8
Augustus	17,5	16,1
September	14,5	15,8
Oktober	10,7	13,4
November	6,7	8,2
December	3,7	4,8
Gemiddeld	10,13	11,71