



ASTERICS - H2020 - 653477

Demonstration of VLBI synchronization via existing SURFnet/LOFAR network

ASTERICS GA DELIVERABLE: D5.14

Document Identifier:	ASTERICS-D5.14
Date:	June 28, 2019
Work Package:	WP5 - CLEOPATRA
Lead Partner:	JIVE
Document Status:	Final
Dissemination Level:	Public
Document Link:	https://www.asterics2020.eu/documents/ASTERICS-D5.14.pdf

Abstract

A common challenge in the design of research facilities is the distribution of an accurate timing and frequency reference. In this part of the Cleopatra project, we research the suitability of the open 'White Rabbit' protocol to deliver a clock signal with sufficient stability for radio astronomical applications. In particular, we demonstrate its suitability by performing VLBI using a remote hydrogen maser reference over optical links of up to 169 km in length, without interrupting the other traffic on the link.

I COPYRIGHT NOTICE

Copyright © Members of the ASTERICS Collaboration, 2015. See www.asterics2020.eu for details of the ASTERICS project and the collaboration. ASTERICS (Astronomy ESFRI & Research Infrastructure Cluster) is a project funded by the European Commission as a Research and Innovation Actions (RIA) within the H2020 Framework Programme. ASTERICS began in May 2015 and will run for 4 years. This work is licensed under the Creative Commons Attribution- Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, and USA. The work must be attributed by attaching the following reference to the copied elements: "Copyright © Members of the ASTERICS Collaboration, 2015. See www.asterics2020.eu for details of the ASTERICS project and the collaboration". Using this document in a way and/or for purposes not foreseen in the license, requires the prior written permission of the copyright holders. The information contained in this document represents the views of the copyright holders as of the date such views are published.

II DELIVERY SLIP

	Name	Partner/WP	Date
From			
Author(s)	Paul Boven	JIVE	June 28, 2019
	Chantal van Tour	OPNT	June 28, 2019
	Rob Smets	SURFnet	June 28, 2019
Reviewed by	Giuseppe Cimó		
Approved by	AMST		

III DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	2019-04-08	First draft	Paul Boven
2	2019-04-16	Chapter on Timing Calibration	Chantal v. Tour
3	2019-04-25	Contributions to Link Setup	Rob Smets
4	2018-06-12	Update of the VLBI chapter	Paul Boven
5	2019-06-20	Fixed a Formula	Chantal v. Tour
6	2019-06-24	LOFAR chapter, VLBI update	Paul Boven
7	2019-06-28	Intro, References, layout	Paul Boven

IV APPLICATION AREA

This document is a formal deliverable for the GA of the project, applicable to all members of the ASTERICS project, beneficiaries and third parties, as well as its collaborating projects.

V TERMINOLOGY

A complete project glossary is provided at the following page:
<http://www.asterics2020.eu/about/glossary/>

CERN Conseil Européen pour la Recherche Nucléaire

ERIC European Research Infrastructure Consortium

LHC Large Hadron Collider

JIVE Joint Institute for VLBI ERIC

LOFAR Low Frequency Array

VLBI Very Long Baseline Interferometry

WR White Rabbit

WSRT Westerbork Synthesis Radio Telescope

VI PROJECT SUMMARY

ASTERICS (Astronomy ESFRI & Research Infrastructure Cluster) aims to address the cross-cutting synergies and common challenges shared by the various Astronomy ESFRI facilities (SKA, CTA, KM3Net & E-ELT). It brings together for the first time, the astronomy, astrophysics and particle astrophysics communities, in addition to other related research infrastructures. The major objectives of ASTERICS are to support and accelerate the implementation of the ESFRI telescopes, to enhance their performance beyond the current state-of-the-art, and to see them interoperate as an integrated, multi-wavelength and multi-messenger facility. An important focal point is the management, processing and scientific exploitation of the huge datasets the ESFRI facilities will generate. ASTERICS will seek solutions to these problems outside of the traditional channels by directly engaging and collaborating with industry and specialised SMEs. The various ESFRI pathfinders and precursors will present the perfect proving ground for new methodologies and prototype systems. In addition, ASTERICS will enable astronomers from across the member states to have broad access to the reduced data products of the ESFRI telescopes via a seamless interface to the Virtual Observatory framework. This will massively increase the scientific impact of the telescopes, and greatly encourage use (and re-use) of the data in new and novel ways, typically not foreseen in the original proposals. By demonstrating cross-facility synchronicity, and by harmonising various policy aspects, ASTERICS will realise a distributed and interoperable approach that ushers in a new multi-messenger era for astronomy. Through an active dissemination programme, including direct engagement with all relevant stakeholders, and via the development of citizen scientist mass participation experiments, ASTERICS has the ambition to be a flagship for the scientific, industrial and societal impact ESFRI projects can deliver.

VII EXECUTIVE SUMMARY

In this deliverable we study the distribution of stable and accurate frequency and time reference signals over optical communications fibre, using modifications of the White Rabbit standard. Link calibration and phase performance are determined, and we demonstrate the usability of the WR links by performing VLBI observations, using the optically distributed reference clock from another radio telescope.

Table of Contents

I	COPYRIGHT NOTICE	1
II	DELIVERY SLIP	1
III	DOCUMENT LOG	2
IV	APPLICATION AREA	2
V	TERMINOLOGY	2
VII	EXECUTIVE SUMMARY	3
	Table of Contents	4
1	Introduction	5
2	White Rabbit Link Setup	6
2.1	Co-existing TFT service in a carrier network	6
2.2	Network setup for both the VLBI and LOFAR TFT service	6
2.3	Network setup for the LOFAR TFT service	7
2.4	Network setup for the VLBI CAMRAS TFT service	7
2.5	Link design and operation	7
3	Timing Calibration	10
3.1	Device delay asymmetries	10
3.2	Link delay asymmetries	11
3.3	Overall timing performance	14
4	VLBI with the WR link	15
4.1	VLBI Requirements	15
4.2	Frequency Reference Stability	15
4.3	A COTS VLBI receiver and backend	16
4.3.1	Detailed flowgraph description	17
4.4	Dwingeloo Fringes, Again	18
4.5	The Last Mile	18
4.6	Parallel Link test	19
4.7	VLBI Observations	19
5	LOFAR WR Connection	22
6	Conclusions	23
	Bibliography	24

1 Introduction

In this deliverable, we aim to demonstrate the suitability of the White Rabbit protocol to deliver a very stable and accurate reference clock. In particular, we attempt to bring the signal from the hydrogen maser reference at the Westerbork Synthesis Radio Telescope (WSRT) to the Dwingeloo Telescope, and to LOFAR, co-existing on an already existing optical network with other traffic. Very Long Baseline Interferometry (VLBI) observations using this transported reference clock are then used to demonstrate the achieved stability.

The original design of White Rabbit (WR) was aimed at synchronizing the beam control electronics of the LHC at CERN, with the goal of achieving sub-ns performance on links of up to 10km distance. CERN published their design as open hardware, and it has since become a very successful project. Accurate distribution of time and frequency are important challenges for any observatory, and make this an important question for multi-messenger astronomy. The frequency stability requirements for VLBI are particularly challenging, as it requires coherence at microwave frequencies, at timescales of up to tens of minutes. Only a hydrogen maser can comfortably meet these requirements, but at a significant capital expense. The ability to share such a stable reference signal over existing fibre links makes this a very interesting solution for any kind of observatory.

The performance of White Rabbit as standardized by CERN doesn't quite meet the requirements for VLBI. The frequency noise performance of the WR switches has been improved by the design of a Low Jitter Daughter-boards (Rizzi et al. 2018), which we installed in our equipment. Additionally, a clean-up oscillator was included in the White Rabbit switch at the telescope, as described in our ASTERICS deliverable D5.4 (Koelemeij and van Tour 2017). The supported link distance was improved by the use of long haul DWDM SFPs and bi-directional optical amplifiers (BDOAs). The use of DWDM SFPs requires the use of external optical multiplexers. Choosing wavelengths that are close together reduces the asymmetry of the fiber link due to its dispersion, and reduces the sensitivity to fiber temperature changes.

The WR links that were built for this research have already been described in our ASTERICS deliverable D5.7 (van Tour et al. 2018). Two links between the Westerbork Synthesis Telescope (WSRT) and Dwingeloo Telescope (DT) were provisioned: one that utilizes the SURFnet DWDM production fibre infrastructure (link length of 169 km, via Groningen), and a direct dark fibre link (35 km). As part of the work in this deliverable, the long link was later changed to connect the LOFAR radio telescope to the WSRT, too. For reference, we include in figure 1 the link diagram from D5.7.

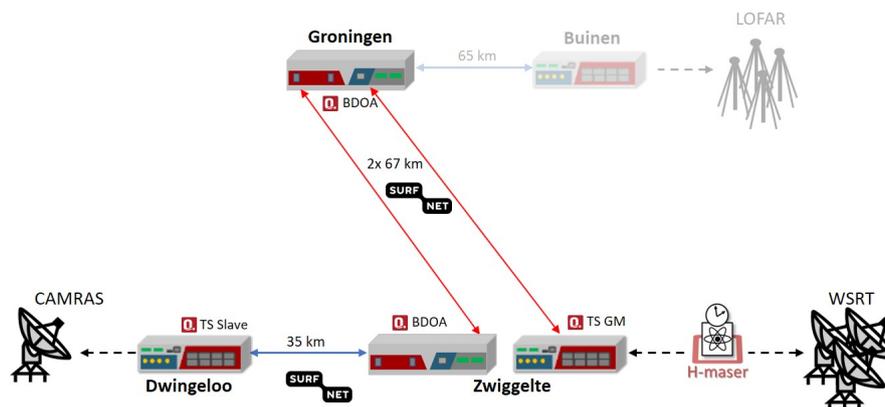


Figure 1: White Rabbit Links between the WSRT, DT and LOFAR

2 White Rabbit Link Setup

2.1 Co-existing TFT service in a carrier network

When a TFT service, such as the White-Rabbit based DWDM high precision TFT service, is to be deployed, either a dedicated fiber must be provided or the TFT service must "piggy-back" on existing services. The latter becomes an absolute requirement when fiber cost plays an important role.

In conventional DWDM transport systems, a dual fiber infrastructure is used. This means a span from location A to B contains two fibers. One fiber is used to transmit light from A to B and the other fiber from B to A. In rare occasions a single fiber is used containing counter-propagating light from A to B and B to A in this one fiber to implement bi-directional data transport. In the case of the White-Rabbit TFT service bi-directional data transport using counter-propagating light is essential to reduce uncertainty when correcting time at the slave with half the round-trip delay.

Relevant parts of the current network for the connections between ASTRON in Dwingeloo, Westerbork (a.k.a. Zwiggelte), and Groningen, are depicted in Figure 2.

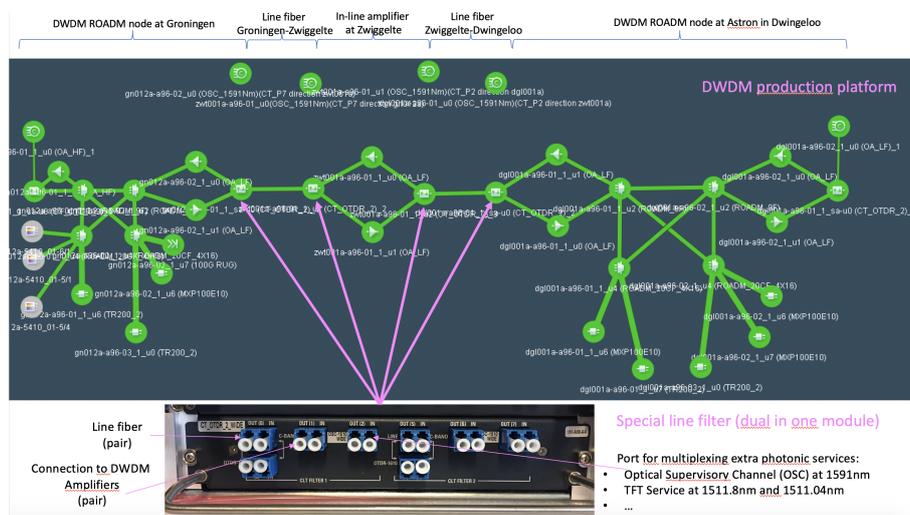


Figure 2: Topology of the SURFnet core network in the Northern branch.

The green symbols indicate components such as EDFAs, ROADMs and Add/drop structures, transponders, and light sources from old production equipment, that belong to the DWDM production network. The final component that is used to gain access to the line fibers is called the CT-filter and is used to combine and split the DWDM production uni-directional signals from other signals such as the uni-directional optical supervisory channel (OSC) and the bi-directional TFT signals.

2.2 Network setup for both the VLBI and LOFAR TFT service

In this project two Proof of Concepts (PoCs) were selected. In one case the TFT service is to synchronize the radio telescope in Dwingeloo and in the other case the LOFAR antenna field in Exloo. In both cases the reference clock is the hydrogen maser atomic clock located at the Westerbork Synthesis Radio Telescope (WSRT). This section describes the common infrastructure for these two applications while the next two sections describe the infrastructure of the LOFAR application and VLBI application, respectively. Figure 3 displays how the master White Rabbit switch is connected to the production network of SURFnet. Under normal conditions the production network uses the 1511nm window to implement an Optical Supervisory Channel (OSC) service for out-of-band management. In this scenario, the OSC has been implemented on 1591nm in order to make the 1511nm band available for the two white rabbit signals: 1511.8nm from master to slave and 1511.0nm from slave to master (return channel). In order to multiplex the White Rabbit signal with the OSC signal a 1511nm mux/demux is used. This multiplexed signal is then inserted in the line filter with the DWDM production traffic and launched into the line fiber

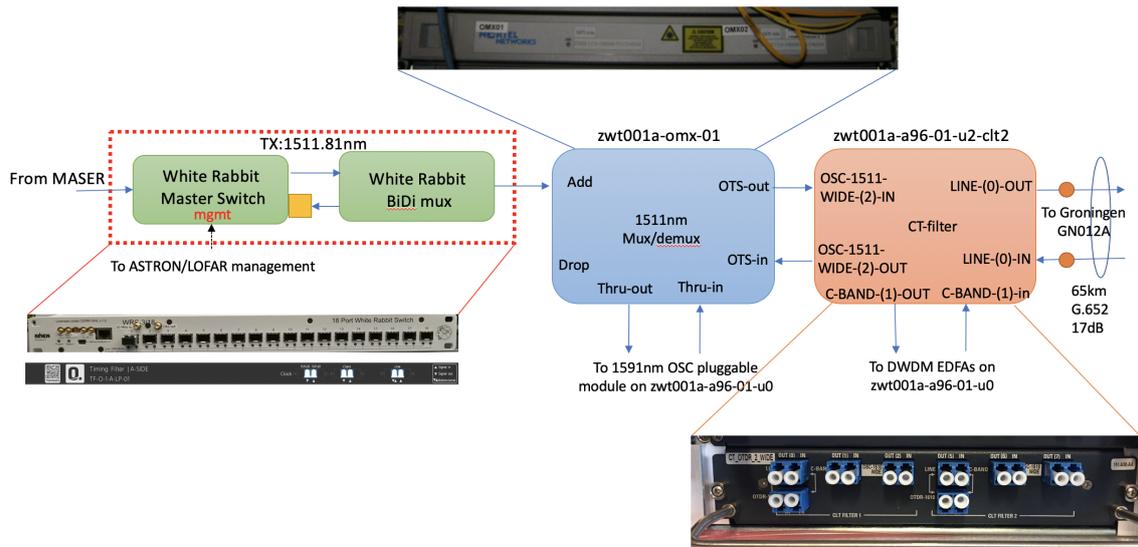


Figure 3: Master White-Rabbit switch in Zwiggelte (zwt001a) transmits light towards Groningen (gn012a) and receives light from slave switches in either Exloo for the LOFAR TFT service or Dwingeloo (dgl001a) for the VLBI CAMRAS TFT service.

going to Groningen. In Groningen the reverse operation takes place (See Figure 4) and the fiber containing both counter propagating TFT signals is connected to the BDOA amplifier. Note that the convention used here is that we the signal as if it were uni-directional in the direction from Master to Slave. In this direction the labeling on the filters, amplifiers and bidi-filters agrees with the direction of propagation.

At the output of the bi-directional optical amplifier, an ODF allows the signal to be patched to either a dark fiber going to LOFAR (Exloo) or back towards Zwiggelte and then further to Dwingeloo (using the second fiber of the fiber pair).

2.3 Network setup for the LOFAR TFT service

In the case the signal in Groningen is patched towards LOFAR in Exloo, the signal terminates at a slave White Rabbit switch in Exloo. The recovered time and frequency can then be used to synchronize the different antenna fields to the frequency and time source in Zwiggelte. From a transport perspective this is a simple termination. The schematic can be found in Figure 5

2.4 Network setup for the VLBI CAMRAS TFT service

In the case the signal is to be delivered for the VLBI CAMRAS application, the amplified master-slave signal is patched back to Zwiggelte using the second fiber of the pair between zwt001a and gn001a. The signal arriving in Zwiggelte needs to be amplified before it can be launched towards its final destination Dwingeloo. This is done as depicted in Figure 6.

Finally in Dwingeloo the slave White Rabbit switch can detect the signal from the master and transmit the return channel to the master switch.

2.5 Link design and operation

For these two PoCs three different Optical Transmission Sections (OTS), or spans, have been used. In order to facilitate link engineering using the BDOA amplifiers with equal gains, fixed-attenuators (pads) have been added to make the total loss of each OTS equal. The only exception is the span that connects to a receiving photo-diode as receiving more optical power improves performance and makes the system less vulnerable to the polarization dependent loss of the semi-conductor optical amplifier based BDOAs. With gain settings of around 20dB-24dB and input powers around -20dBm the launch power into the optical fiber varies under normal conditions between -2dBm and +2dBm, making the system insensitive to Amplified Spontaneous Emission (ASE), without invoking fiber non-linear distortion.

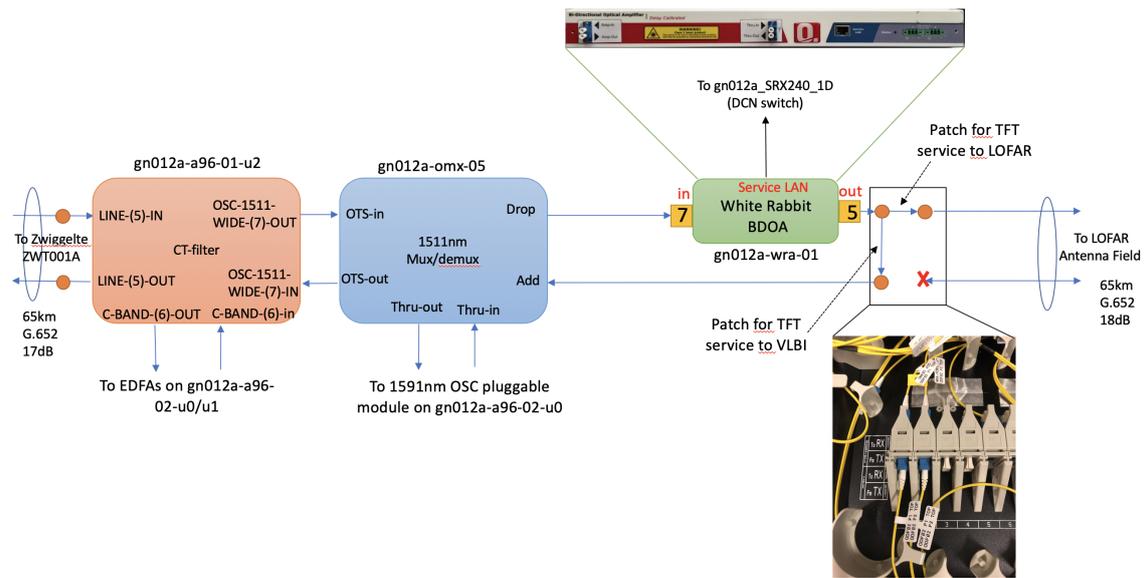


Figure 4: TFT amplifier station in Groningen (gn012a). In Groningen it is possible to patch the service to Exloo for the LOFAR application or back to Dwingelo (gdL001a) via Zwiggelte for the VLBI CAMRAS application

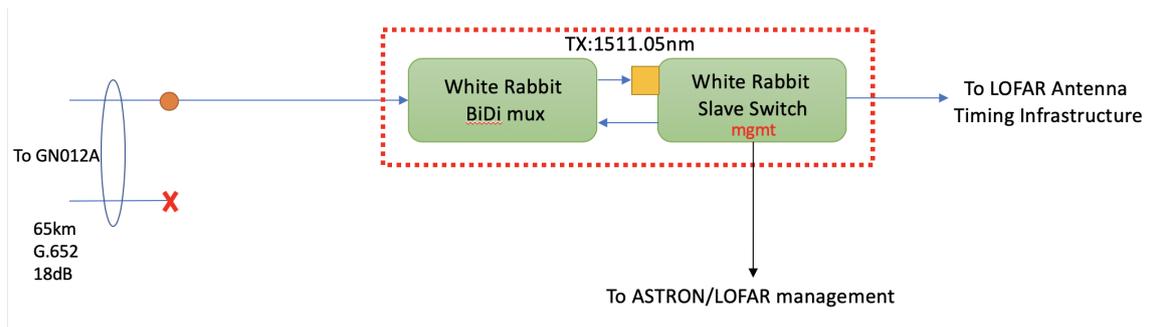


Figure 5: Slave White-Rabbit switch in Exloo receives light from the master switch in Zwiggelte and transmits light towards the master switch.

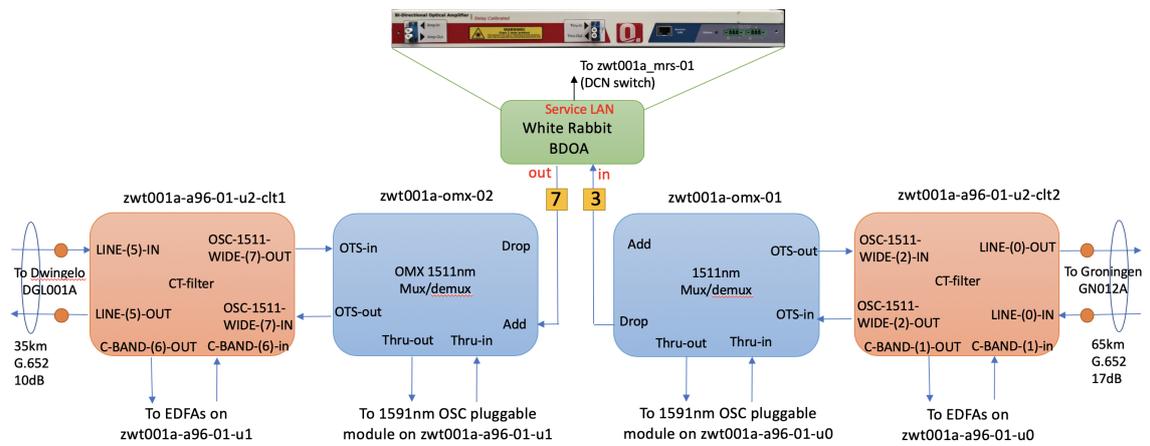


Figure 6: Amplification in Zwiggelte (Westerbork)

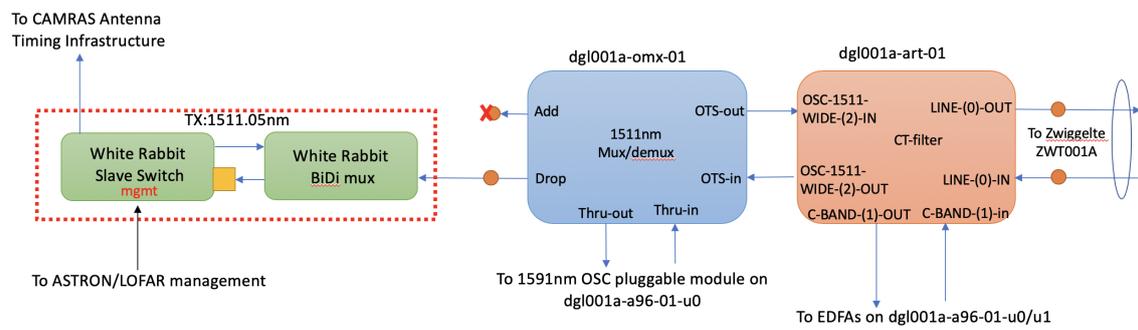


Figure 7: Slave White-Rabbit switch in Dwingeloo (dgl001a) receives light from the master switch in Zwiggelte looped back via Groningen and transmits light towards the master switch.

3 Timing Calibration

As described in ASTERICS deliverable 5.7 (van Tour et al. 2018), a WR link contains delay asymmetries, causing a phase offset between a grandmaster (GM) and slave WR switch. This phase offset, *skew*, is given by

$$skew = \frac{\delta_{ms} - \delta_{sm}}{2}, \quad (1)$$

where δ_{ms} is the upstream master-to-slave (ms) delay and δ_{sm} is the downstream slave-to-master (sm) delay. The delay asymmetries are measured and corrected for in such a way that the phase offset between the GM and slave WR switch is minimized. Figure 8 shows the lay-out of the WR links between Zwiggelte and Dwingeloo, including all devices and fiber-optic links that contribute to the phase offset between the WR switches. The next sections describe 1) the device delay asymmetries and 2) the link delay asymmetries.

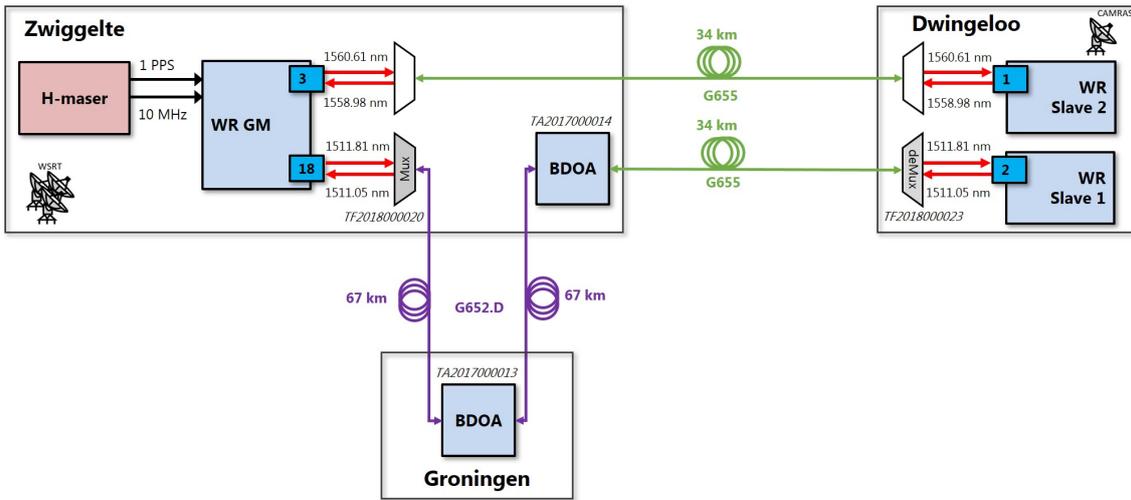


Figure 8: Link lay-out containing all relevant devices for the timing calibration.

3.1 Device delay asymmetries

The device delay asymmetries are caused by

- electrical-electrical delay asymmetries in the WR switches
- optical-electrical delay asymmetries in the SFPs
- optical-optical delay asymmetries in the optical multiplexers and BDOAs

CERN developed calibration techniques to measure and correct for these device delay asymmetries (Daniluk 2015). These techniques are used to calibrate the WR switches and optical devices, as was previously described in deliverable 5.7 (van Tour et al. 2018). The WR switches were calibrated by measuring and correcting for the phase offset between the WR GM and the WR slaves, while the WR switches were directly connected to each other using short patch-cords. The phase offset was measured with a Stanford Research SR620 time interval counter. The uncertainty in the final phase offsets between the WR GM and each WR slave is based on the 25 ps (coverage factor $k = 1$) resolution of the SR620 time interval counter (SRS) and the WR restart jitter of 12 ps ($k = 1$). The total uncertainty of 30 ps ($k = 1$) is calculated with the square root of the quadratic sum of the uncertainties,

$$\sigma_{WR\ GM - WR\ Slave} = \sqrt{\sigma_{SR620}^2 + \sigma_{WR\ restart}^2} \quad (2)$$

The optical devices that are part of the 169 km link (Zwiggelte-Groningen-Zwiggelte, Fig. 8) were delay-calibrated within OPNT premises. Table 1 lists all the measured optical device delay asymmetries. This

table also lists a total uncertainty in the phase offset between the WR GM and WR slave 1 of $\sigma_{\text{optical devices}} = 31$ ps ($k = 1$). JIVE reported that the optical multiplexers in the 34 km link (Zwiggelte-Dwingeloo) do not add significant delay asymmetries.

Table 1: overview of the optical device delay asymmetries for the 169 km link.

Device	Serial number	Location	Skew (ps)
Optical Multiplexer (Mux)	TF2018000020	Zwiggelte	-38 ± 12
Optical de-Multiplexer (deMux)	TF2018000023	Dwingeloo	37 ± 12
Bi-Directional Optical Amplifier (BDOA)	TA2017000014	Zwiggelte	-1016 ± 15
Bi-Directional Optical Amplifier (BDOA)	TA2017000013	Groningen	887 ± 21
Total			-130 ± 31

3.2 Link delay asymmetries

Apart from the delay asymmetries caused by the optical devices, long-haul fiber-optic links (> 1 km) also add significant delay asymmetries. This is caused by the chromatic dispersion (CD) of the fiber-optic link. This CD causes a temporal delay difference between traveling signals with different wavelengths. This delay difference between any two wavelengths, λ_1 and λ_2 , is given by

$$\delta_2 - \delta_1 = L \int_{\lambda_1}^{\lambda_2} D(\lambda) d\lambda, \quad (3)$$

where $D(\lambda)$ is the dispersion of the fiber-optic link, and L is the length of the fiber-optic link. Since the WR links use two different wavelengths for the upstream (λ_{ms}) and downstream (λ_{sm}) signals, the CD causes a phase offset between the GM and slave WR switch. This phase offset can be written in terms of the CD by combining Eqs. 1 and 3,

$$skew|_{CD} = \frac{L}{2} \int_{\lambda_{sm}}^{\lambda_{ms}} D(\lambda) d\lambda. \quad (4)$$

As becomes clear from this equation, the CD induced phase offset between the GM and slave WR switch is directly proportional to L , while the CD itself is not proportional to L . In order to treat the CD-induced delay asymmetry purely as a property of the fiber's material, the WR switches use the following dimensionless α parameter,

$$\alpha \equiv \frac{\delta_{ms, \text{fiber}} - \delta_{sm, \text{fiber}}}{\delta_{sm, \text{fiber}}}. \quad (5)$$

This α parameter, in combination with a round-trip delay measurement, is then used to calculate and correct for the CD-induced phase offset between the WR GM and WR Slave. Since the round-trip delay is frequently updated, any possible change in the fiber-optic link length does not require a new calibration.

To determine the α parameter of the link between Zwiggelte and Groningen (G652.D fiber), WR Slave 1 was temporarily placed in Zwiggelte. This allowed us to directly measure the phase offset between the WR GM and WR Slave 1, $skew|_{CD}$, while inserting the fiber-optic link. The exact procedure and results of this calibration are described in deliverable 5.7 (van Tour et al. 2018).

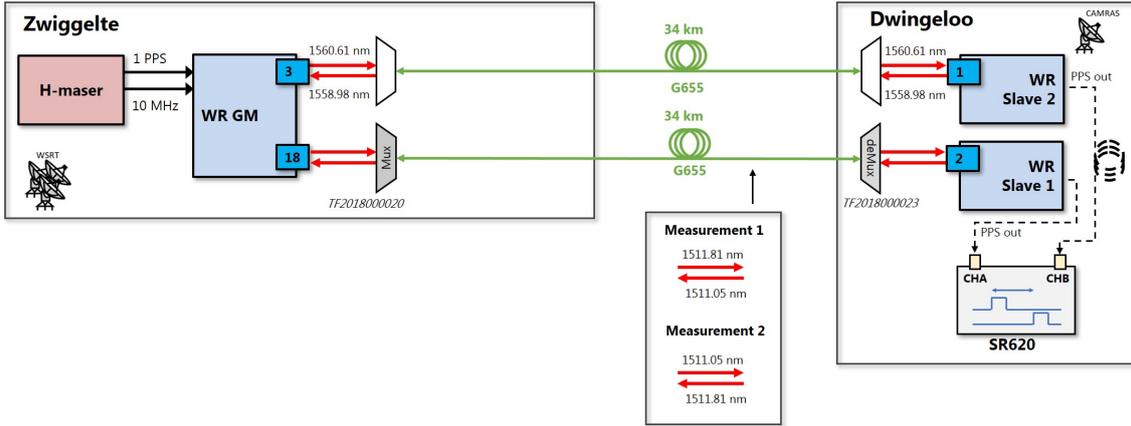


Figure 9: Schematic of the link lay-out that is used for calibration of the 1511.81/1511.05 nm Zwiggelte-Dwingeloo link. An extra 2 m of coax cable is added to the PPS output of WR Slave 2 to make sure that the time interval counter readings are always positive. The SURFnet equipment is not depicted.

The determination of the α parameters of the Zwiggelte-Dwingeloo (G655.D fiber) links required a different method, described in (Boven 2017). Figure 9 shows a schematic of the setup that was used for the calibration of the 1511.81/1511.05 nm Zwiggelte-Dwingeloo link. During the calibration, WR Slave 2 was used as a stable reference, while the upstream and downstream wavelengths of the 1511.81/1511.05 nm link were interchanged. While doing so, the round-trip delay (δ) of the 1511.81/1511.05 nm link and the phase offset between the two WR slave switches were measured with the SR620. The difference in phase offsets between measurement 1 ($skew_1$) and measurement 2 ($skew_2$), together with the measured δ , give sufficient information to calculate α . This is possible, because α (Eq. 5) can be re-written (using some basic algebra) into (Daniluk 2015):

$$\alpha = 2 \frac{skew|_{CD}}{\frac{1}{2}\delta - skew|_{CD}}, \quad (6)$$

with

$$skew|_{CD} = \frac{skew_2 - skew_1}{2}. \quad (7)$$

This equation is used to calculate the α parameter of the optical link for wavelengths 1511.81 nm and 1511.05 nm. Similarly, the calibration of the 1560.61/1558.59 nm link was performed by using WR Slave 1 as a reference, while swapping the upstream and downstream wavelengths of the 1560.61/1558.59 nm link.

The uncertainty in the calculated α parameter is given by

$$\begin{aligned} \sigma_\alpha &= \sqrt{\left| \frac{\partial \alpha}{\partial skew|_{CD}} \right|^2 \sigma_{skew|_{CD}}^2 + \left| \frac{\partial \alpha}{\partial \delta} \right|^2 \sigma_\delta^2} \\ &= \frac{4}{(\delta - 2skew|_{CD})^2} \sqrt{\delta^2 \sigma_{skew|_{CD}}^2 + skew|_{CD}^2 \sigma_\delta^2} \end{aligned} \quad (8)$$

Since $skew|_{CD}^2 \sigma_\delta^2 \ll \delta^2 \sigma_{skew|_{CD}}^2$ and $skew|_{CD} \ll \delta$, the uncertainty in α can be simplified to

$$\sigma_\alpha \approx 4 \frac{\sigma_{skew|_{CD}}}{\delta}. \quad (9)$$

So, it turns out that the uncertainty in α is determined by the uncertainty in $skew|_{CD}$. This uncertainty is calculated with (using Eq. 7):

$$\sigma_{skew|_{CD}} = \frac{1}{2} \sqrt{\sigma_{skew_1}^2 + \sigma_{skew_2}^2}. \quad (10)$$

The uncertainties in these phase offset measurements ($skew_1$ and $skew_2$) are (just as in section 3.1) based on the resolution of the SR620 time interval counter and the restart jitter of the WR switches. Therefore, Eq. 10 can be rewritten into:

$$\sigma_{skew|CD} = \frac{1}{2} \sqrt{2\sigma_{SR620}^2 + 2\sigma_{WR\ restart}^2} \quad (11)$$

Inserting the values mentioned in section 3.1 into equation 11 results in an uncertainty of 20 ps. However, the actual calibration measurements resulted in somewhat larger uncertainties (see Table 2). Those listed uncertainties are used in further calculations.

Table 2 lists a complete overview of all the results from the link delay asymmetry measurements. This table also provides a value of the estimated CD values, D . These values are estimated by assuming a constant CD over the measured wavelength range ($\Delta\lambda = \lambda_{ms} - \lambda_{sm}$). This allows us to rewrite Eq. 4 in terms of D into

$$D \approx 2 \frac{skew|CD}{L\Delta\lambda} \text{ for } \lambda = \lambda_{ms} \dots \lambda_{sm}. \quad (12)$$

The uncertainties in these estimated CD values are calculated with

$$\begin{aligned} \sigma_D &= \sqrt{\left| \frac{\partial D}{\partial skew|CD} \right|^2 \sigma_{skew|CD}^2 + \left| \frac{\partial D}{\partial L} \right|^2 \sigma_L^2 + \left| \frac{\partial D}{\partial \Delta\lambda} \right|^2 \sigma_{\Delta\lambda}^2} \\ &= |D| \sqrt{\left| \frac{\sigma_{skew|CD}}{skew|CD} \right|^2 + \left| \frac{\sigma_L}{L} \right|^2 + \left| \frac{\sigma_{\Delta\lambda}}{\Delta\lambda} \right|^2}. \end{aligned} \quad (13)$$

To calculate these uncertainties, σ_L (50 m, $k = 1$) and $\sigma_{\Delta\lambda}$ (50 pm, $k = 1$) are estimated using the specifications of the fiber-optic links and SFPs, respectively.

The estimated CD values are all in accordance with the ITU recommendations (ITU-T G.652; ITU-T G.655). To illustrate this, Fig. 10 shows the estimated CD value for the G652.D (Zwiggelte-Groningen-Zwiggelte) fiber-optic link and the estimated CD values for the G655 (Zwiggelte-Dwingeloo) fiber-optic link, together with ITU recommendations (colored areas). In case of the G652.D fiber, SURFnet also provided more accurate CD measurements (dashed line).

Table 2: overview of the properties of the fiber-optic links, and the measured link delay asymmetries. All listed uncertainties are reported with $k = 1$.

Link	Zwiggelte-Groningen-Zwiggelte (Zwt-Gn-Zwt)	Zwiggelte-Dwingeloo (Zwt-Dwg)	Zwiggelte-Dwingeloo (Zwt-Dwg)
Fiber Type	G652.D	G655	G655
λ master (nm)	1511.81	1511.81	1560.61
λ slave (nm)	1511.05	1511.05	1558.59
Length (km)	133.64 ± 0.05	34.37 ± 0.05	34.37 ± 0.05
$skew CD$ (ps)	544 ± 47	45 ± 38	116 ± 12
δ (ps)	1311048015 ± 100	34309488 ± 17	343197646 ± 56
D (ps/(nm km))	11 ± 2	4 ± 3	4 ± 1
α	$(17 \pm 2) \times 10^{-7}$	$(5 \pm 4) \times 10^{-7}$	$(14 \pm 2) \times 10^{-7}$

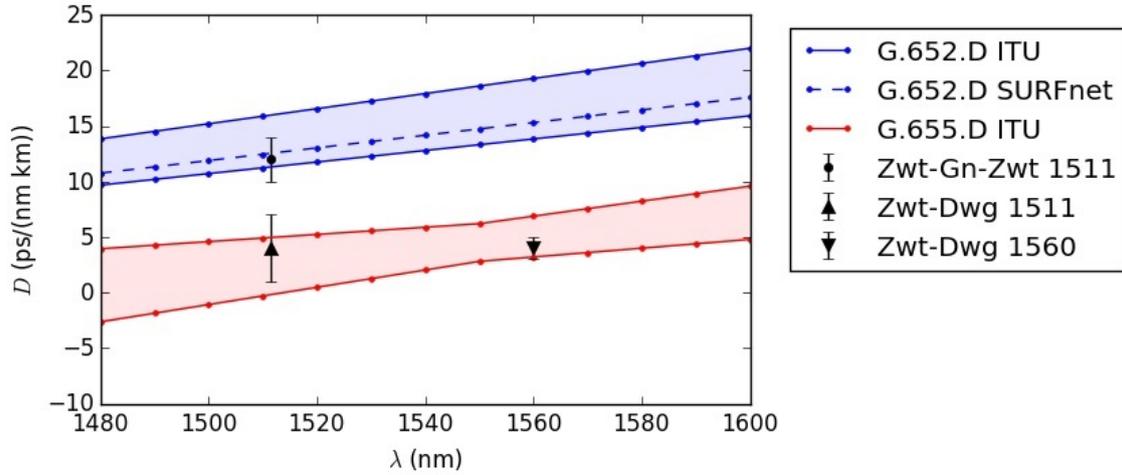


Figure 10: Estimated CD values in comparison with the ITU recommendations and SURFnet data.

3.3 Overall timing performance

To check the timing performance of the WR links, the PPS outputs of the WR slaves (see Fig. 8) are compared with each other, using the SR620. The remaining phase offset between the two PPS outputs of the WR slaves is (370 ± 30) ps with $k = 1$, where $t_{slave1} < t_{slave2}$. The total uncertainty on this phase offset, due to mentioned calibration uncertainties, can be calculated with:

$$\sigma = \sqrt{\sigma_{WR\ GM - WR\ Slave\ 1}^2 + \sigma_{WR\ GM - WR\ Slave\ 2}^2 + \sigma_{optical\ devices}^2 + \sigma_{Zwt-Gn-Zwt,\ 1511}^2 + \sigma_{Zwt-Dwg,\ 1511}^2 + \sigma_{Zwt-Dwg,\ 1560}^2}$$

$$\sigma = \sqrt{30^2 + 30^2 + 31^2 + 47^2 + 38^2 + 22^2} = 90\ ps. \quad (14)$$

This uncertainty analysis does not account for the (-370 ± 30) ps measured offset between the two WR Slaves. However, there are still multiple origins of uncertainty not taken into account. These are not included in the calculations, because more research is needed to quantify those uncertainties. One of those extra origins of uncertainty comes from the wavelength drift of the SFPs. The calibration of the Zwiggelte-Groningen-Zwiggelte link was performed in August 2018, while the calibration of the Zwiggelte-Dwingeloo was performed in March 2019. The transmitted wavelengths of the SFPs that are used in the link could have drifted in the meantime. Furthermore, the calibration performed at Zwiggelte-Dwingeloo required two pairs of SFPs, of which only one pair of SFPs is used in the final setup. Any difference in wavelength of the SFPs will contribute to the uncertainty in the link calibration. The extra uncertainty between the two WR slaves due to uncertainty in the transmitted wavelength of the SFPs is estimated at ~ 40 ps ($k = 1$), using Eq. 4 together with the estimated CD values from Table 2 and $\sigma_{\Delta\lambda}$ of 50 pm ($k = 1$). Another uncertainty that is not exactly known comes from the SR620 time interval counter. Temporal variations in the asymmetry between the two channels of the TIC would increase the uncertainty in the results. The 25 ps uncertainty that was used as an uncertainty for the SR620 does not include such type of variation. Further research should be performed to analyze the exact contribution of the SR620 to the total uncertainty budget.

In conclusion, the WR links are calibrated with better than 1 ns uncertainty with respect to each other. A real verification of the timing performance would require an extra Dwingeloo-Zwiggelte link such that a loop can be created and the WR slaves can be directly compared with the WR GM.

4 VLBI with the WR link

The main goal of this deliverable is to demonstrate the suitability of our WR system to distribute the reference clock with sufficient stability to perform VLBI observations, specifically between the Westerbork Synthesis Radio Telescope (WSRT) and the Dwingeloo Telescope (DT). The WSRT is an operational radio telescope which participates regularly in VLBI observations as part of the European VLBI Network (EVN). The DT, in contrast, has been decommissioned and is now a listed national monument. It is run by the volunteers of the CAMRAS (C.A. Muller Radio Astronomy Station) foundation. Although it is used regularly to observe e.g. pulsars and the hydrogen line, it lacks the essential equipment to perform VLBI observations.

4.1 VLBI Requirements

VLBI is one of the more challenging observational methods within radio astronomy. Part of the challenge is that the cross correlation fringe will vanish if frequency, stability, time, pointing or location are wrong, without providing a clue as to the cause of the failure. The list below summarizes what is required to participate in VLBI observations, with the major items worked out in subsequent sections of this chapter.

- Radio Telescope, with sufficient sensitivity
- Frequency Reference, with sufficient phase stability
- Time Reference, accurate to a few μs
- Receiver, tunable, with tens to hundreds of MHz of bandwidth
- Backend, to digitize and timestamp the received signal
- Location and geometry of the telescope, accurate to a fraction of the observing wavelength
- RF Noise source, for calibrating the telescope sensitivity
- Field system, to steer the telescope to a source, and control the observation

The precise location of the DT has been established a few years before in a professional GPS survey. For its geometry (distance between the azimuth and elevation axes), we use the values from the radio telescope in Defford (UK), which is an exact copy of the DT. We decided to do without the noise source, because amplitude calibration is required for imaging VLBI observations, but we are only interested in the frequency stability performance at this stage. Likewise, as only simple observations targeting a single source are foreseen in this deliverable, we skip installing a 'field system', and instead control the telescope and VLBI equipment manually.

4.2 Frequency Reference Stability

In this project, the DT will use the WSRT's hydrogen maser as its frequency and time reference, transported by WR. To demonstrate the long term phase noise performance of the link, we conducted several VLBI astronomical observations. As a criterion for success, we use the same requirements as adopted for the design of the VLBI capability of the future SKA, i.e. a coherence loss (due to the clock distribution) of less than 2%, for integration times of 1 s and 60 s. The SKA1-Mid requirements are based on a highest observing frequency of 13.8 GHz. The Dwingeloo telescope, due to its mesh dish surface, stops being useful above 8 GHz, but meeting the SKA design requirements would show the general applicability of the WR frequency transfer.

The first long distance link we had working for this project was the round-trip from the WSRT to Groningen, and back. As described in deliverable D5.7, we performed a loop-back test on this link, which showed an ADEV of approximately 10^{-12} at 1 s integration time. For reference, the figure from deliverable D5.7 is repeated below.

The coherence loss requirement can be converted into an Allan deviation (ADEV) value. This conversion can be done analytically only in the case of white phase noise (ADEV slope of -1) or white frequency noise (ADEV slope of $-\frac{1}{2}$). As already determined during D5.7 (see figure 11), we measured a slope of

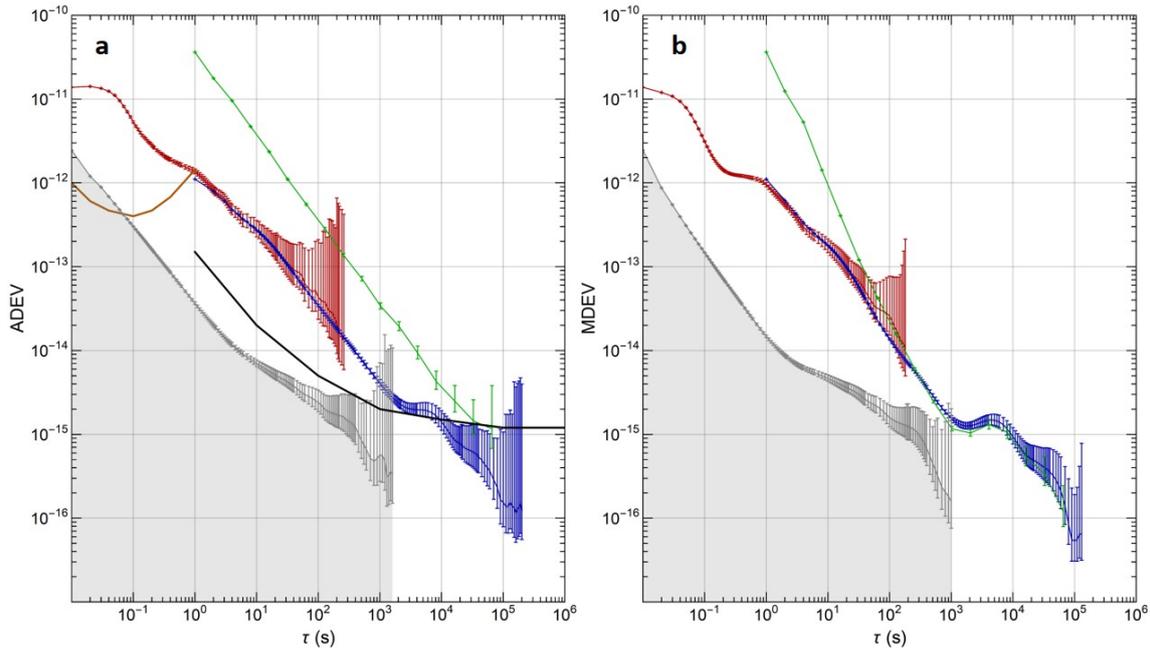


Figure 11: ADEV (left) and MDEV (right) of the 2x 67 km loopback test. Red: measured with 50 Hz bandwidth. Blue: same, 0.5 Hz bw. Black: H-maser performance specifications from sales literature. Green: PPS measurements. Brown: clean-up oscillator stand-alone performance. Gray: 3120A noise floor.

–1, which means that the WR link is in the white phase noise regime. For an observing frequency of up to 13.4 GHz, this requires an ADEV of $2.8 \cdot 10^{-12}$ at one second on the link (Alachkar et al. 2018), which we do indeed (just) meet.

4.3 A COTS VLBI receiver and backend

A VLBI backend digitizes (samples) the signals within the observed radio bandwidth and adds accurate timestamps, so the data can later be aligned with that of the other radio telescopes, and correlated. Past VLBI digitizers would use analog filters to divide the observed bandwidth into smaller sub-bands before digitizing, whereas modern backends will sample the full bandwidth, and use digital signal processing to generate the sub-bands. The bottleneck in VLBI observations is usually in storage capacity or networking bandwidth. Using fewer bits to represent the samples allows for a higher observed bandwidth, but also results in higher quantization noise. Most VLBI recordings use only 2 bits per sample to get the best sensitivity.



(a) The Ettus X310 SDR



(b) The TwinRX Receiver Module

Figure 12: COTS hardware for the VLBI Receiver/Backend (source: Ettus website)

As no VLBI backend was readily available for the DT, we chose to implement our own using a commercial off the shelf (COTS) software defined radio (SDR). Due to its high bandwidth, modularity, and inputs for PPS and 10 MHz reference signals, the Ettus Research X310 was selected. Another advantage of the X310 is the relatively large capacity of its FPGA, which enables offloading some of the VLBI

processing to it. The X310 is a modular SDR, and for its receiver, we selected a pair of ‘TwinRX’ modules, for their bandwidth and built-in filtering. These modules can be tuned to any frequency between 10 MHz and 6 GHz, and together provide up to 320 MHz of bandwidth, or 160 MHz on two polarizations. For Nyquist sampled 2 bits data, this RF bandwidth would correspond to a VLBI backend output of 1.28 Gb/s.

GNU Radio is an open source software ecosystem for software defined radio (SDR). It contains optimized implementations of many important SDR blocks for signal processing such as filtering and Fourier transforms. The GNU Radio Companion (GRC) is its graphical user interface, which makes it possible to construct complicated signal processing processes as a flow-graph of interconnected blocks. We developed a flow-graph to process the recorded samples from the backend and generate VLBI Data Interchange Format (VDIF) compliant data.

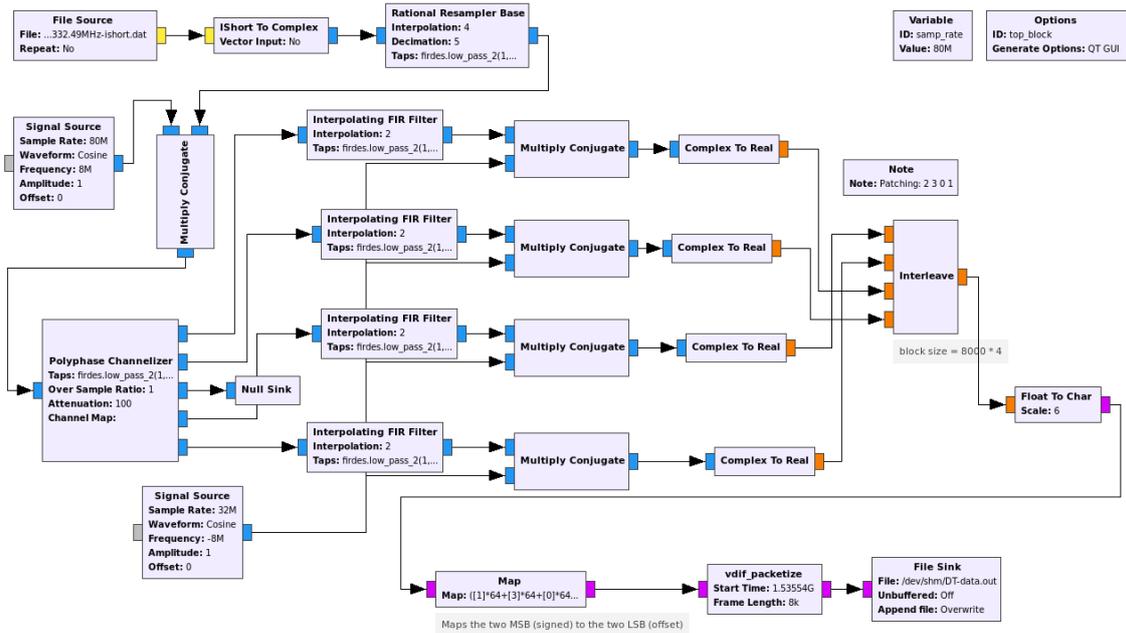


Figure 13: GNU Radio VLBI Flowchart

4.3.1 Detailed flowgraph description

Recordings from the X310 SDR are at 100 MS/s, complex (interleaved I and Q) shorts (16 bits), with an actual bandwidth of 80 MHz. In the flow-graph (figure 13), such a recording is read by the ‘File Source’ block at the top left. The interleaved shorts are subsequently converted to complex floating point numbers, and then re-sampled in the ‘Rational Resampler Base’ block to 80 MHz bandwidth. The ‘Multiply Conjugate’ and ‘Signal Source’ block together shift the spectrum up by 8 MHz. The ‘Polyphase Channelizer’ subdivides the recorded spectrum in five bands of 16 MHz each. Due to the earlier frequency shift of 8 MHz, the band edges end up together in a single sub-band, and are discarded in the ‘Null Sink’ block. The remaining four sub-bands are each 16 MHz wide, for a total observed bandwidth of 64 MHz in a single polarization.

Each sub-band is then passed through a set of three blocks, where the sampling rate is doubled (‘Interpolating FIR filter’), shifted up by 8 MHz by the ‘Multiply Conjugate’ block, and then the imaginary part of the signal is dropped. This sequence converts each sub-band from complex, via analytic, to real (upper side band) at twice the sampling rate, to be compatible with the JIVE SFXC correlator. The ‘Interleave’ block sequences groups of 32000 samples together, to get packets with 8000 bytes of payload. Re-quantizing the data to 2 bits per sample is done in two stages: first the floating point numbers are converted to 8 bit numbers (‘chars’), and then the ‘Map’ block forms a look-up table with 256 entries. This two stage approach turned out to be necessary to sidestep a small bug in the quantization of small

numbers that was present at that time in GNU Radio.

GNU Radio offers the ability to create custom blocks in C++ or Python, that can be integrated into flowcharts. The block 'vdif_packetize' is a custom block created to add VDIF headers to the VLBI data. These headers contain details such as the station code ('DW' for Dwingeloo), a timestamp, and fields describing which channel is contained in each packet. Adding the VDIF header makes the VLBI data compatible with most VLBI correlators, including the SFXC correlator at JIVE.

Given the high data throughput of VLBI observations, and all the filtering steps in the flowchart, it is not surprising that even a modern PC cannot perform this work in real-time. On a high-end workstation equipped with an Intel i9900K CPU, NVME SSD and 64 GB of RAM, it achieves a speed of about 20% of real time.

The custom module, and VLBI flowchart, have been released by us as open source software under the GPL, and are available on the ASTERICS repository, and Github.

4.4 Dwingeloo Fringes, Again

The Dwingeloo telescope participated in the first European VLBI tests in 1978 (Schilizzi et al. 1979), but the WSRT soon took over as the Dutch VLBI telescope. On August 29th 2018, 40 years after its initial VLBI observations, we obtained fringes again with the Dwingeloo telescope, to both the WSRT, and the Mark2 telescope at Jodrell Bank Observatory (JBO) in the UK. All three telescopes observed the source 4C39.25 at 1332.49 MHz. Dwingeloo only observed in left-hand circular polarization (LCP), generating four sub-bands of each 16 MHz. Due to the high output rate of the backend, the data could only be recorded to RAM, limiting recording times to 10 s. As the WR link was not in place yet, a Rubidium atomic clock (FS725) was used as a time base for the DT, while the other participating stations used their hydrogen masers. Our GNU Radio flowchart was used to convert the recorded samples for the DT into VDIF data. Correlation was performed on the SFXC correlator at JIVE. Fringes were obtained right away, in all four sub-bands.

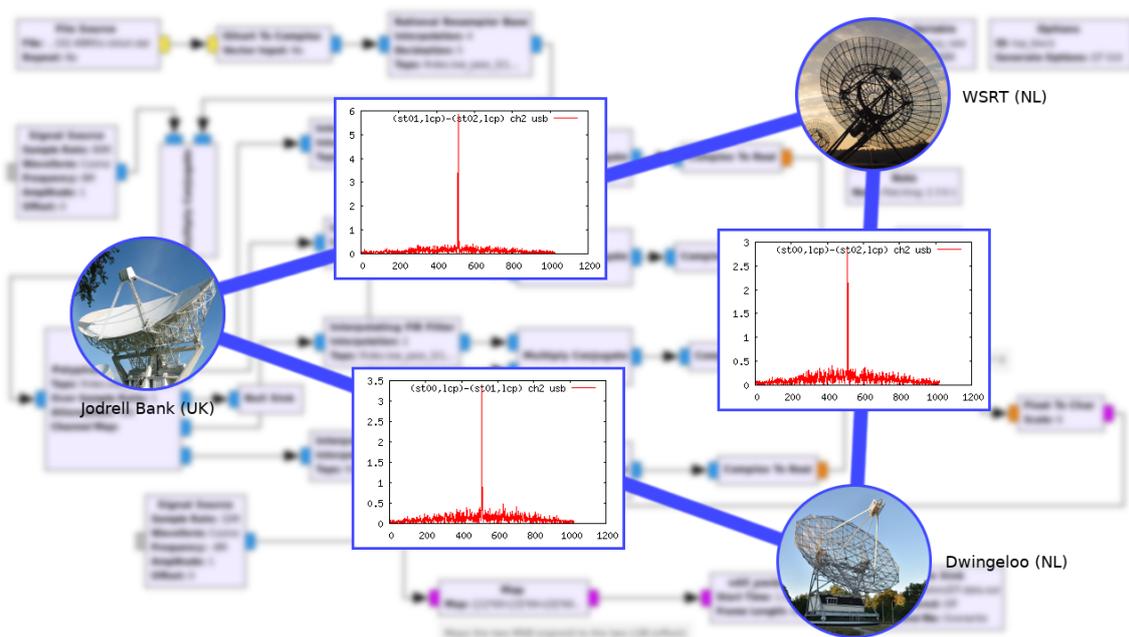


Figure 14: The Astron/JIVE daily image, showing the fringes and participating telescopes

This milestone in the project was published as a 'ASTRON/JIVE daily image', and a poster and paper for the EVN Symposium in Granada.

4.5 The Last Mile

The Dwingeloo telescope is only 60 m away from the ASTRON building, and already connected to it by multi-mode optical fibre. White Rabbit however requires single mode fibre, and an upgrade of the fibre

connection was also required to support the higher networking bandwidths generated by VLBI. The easiest route between the telescope and the SURFnet connections in the basement of the ASTRON building turned out to be 275 m long, of which about 100 m is inside the building. ASTRON was kind enough to donate 510 m of G.652.D glass-fibre to the project, left over from the recent upgrade to the WSRT. Only about 100 m of actual digging was required.



Figure 15: CAMRAS volunteers running new fiber

In December 2018, we rented a small digger, and a large group of CAMRAS volunteers joined to help with the groundwork. We installed specialized conduit for the glass fibre (as this wasn't rated for outdoor use without protection), placed warning tape above it, and closed the ground again. Two months later, the fibre itself was run through the conduit. Further work included a new fibre path inside the telescope, including a flexible 'cable wrap' around the azimuth axis of the telescope, so there would be no stresses on the fibre while turning the telescope. Finally, to complete the connection, the fibres had to be spliced together just outside the telescope, and connectorized at both ends of the run, all of which took several days to complete. In March, we had the link complete, and finally were able to connect the WR link from the WSRT maser, to the Dwingeloo telescope.

4.6 Parallel Link test

At this stage in the project, we had three operational WR links between WSRT and the DT. We had two links of 35 km dark fiber running directly between the two observatories, and the 169 km link through the SURFnet DWDM network. The shorter links had been used to calibrate the dispersion on the WSRT-DT link (see previous chapter), and will remain in place after the project. Furthermore, we had a WR switch with, and without cleanup-oscillator available in the DT. This allowed the opportunity to measure the differential performance between two links, albeit without the cleanup oscillator on one of them.

As the ADEV measurement in this case is the relative performance between two links, the required ADEV for less than 2% coherence loss becomes $\sqrt{2} \times 2.8 \cdot 10^{-12} = 4.0 \cdot 10^{-12}$, which the system meets with a small margin.

These measurements were carried out using a MicroSemi 3120a phase noise analyzer, and two WR switches, one of which was equipped with a clean-up oscillator. As can be viewed in figure 16, there is no discernible difference in performance between the 35 km (blue line) and 169 km link (purple), when compared against the other 35 km link. Not having a cleanup oscillator (green line) seems to make only a small difference, but this is because there was only one WR switch with cleanup oscillator, so all measurements (with or without cleanup oscillator) are relative to a switch without cleanup, and therefore dominated by the phase noise contributions of the WR switch without cleanup oscillator.

4.7 VLBI Observations

On March 8th 2019, the Dwingeloo telescope joined in a regular EVN network measurement experiment (NME), named N19L1. Twenty telescopes participated in this observation. With the new fibre connection, the DT VLBI backend could be connected at 10 Gb/s to a JIVE server for recording the data, removing the earlier limit of ten seconds of observation time. EVN L-band observations are mostly at

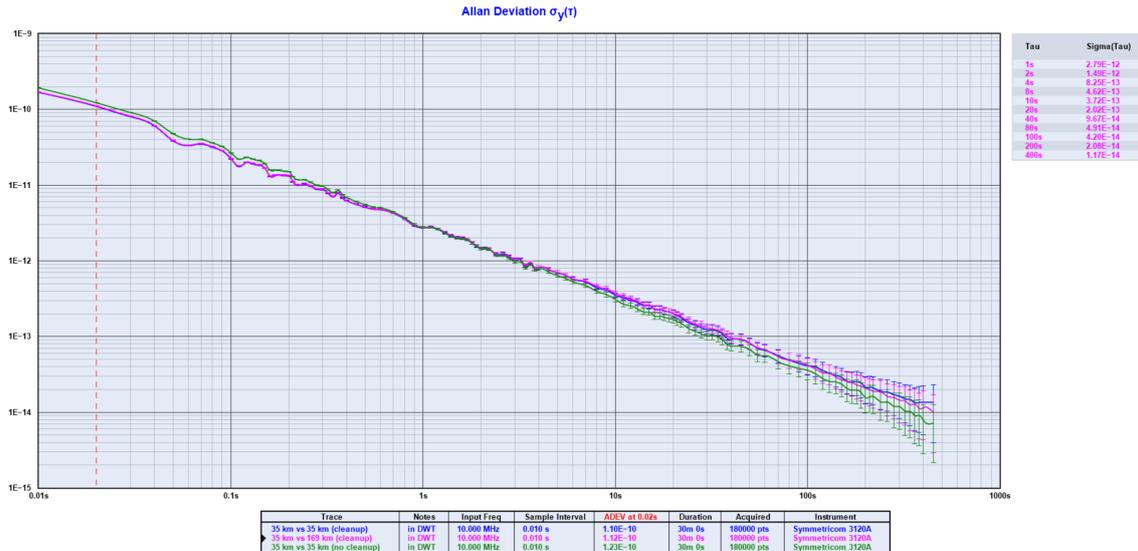


Figure 16: Relative ADEV between two parallel WR links

18 cm, whereas the Dwingeloo telescope L-band receiver is for 21 cm. Participating in this observation required that the CAMRAS volunteers installed new low-noise amplifiers for the 18 cm band, which was carried out despite the rather poor weather, with rain and high winds. Fringes between the DT and many of the participating stations were detected, in our first observation using the remote hydrogen maser. The gain of the new LNAs was a bit lower than expected, causing significant sensitivity loss, which did impact the signal to noise of the Dwingeloo fringes.

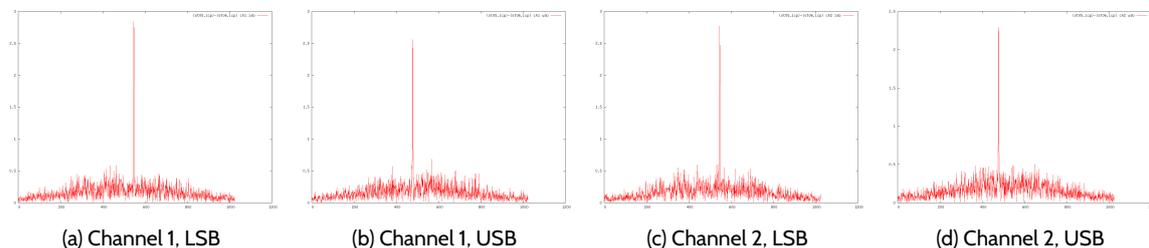


Figure 17: Dwingeloo - Effelsberg fringes, N19L1

To better study the VLBI performance of the new WR link, a number of small VLBI observations were scheduled, with the VLBI stations WSRT, Jodrell Bank and Torun participating in these tests. These were carried out at around 1330 MHz, where the DT has better sensitivity, but still away from the Galactic hydrogen line. Each observation targeted 3C84, a bright, compact source at L-band.

In order to measure the frequency stability, we use the JIVE SFXC correlator to calculate the phase on each of the baselines between the telescopes, at integration times of 1 s. These phase values are then converted into timing differences at the observing frequency, and from this, the Allan deviation is calculated. We expect the phase noise measured in such a fashion to be higher than the directly measured phase noise performance of the link itself. Otherwise, the clock distribution system itself would be the limiting factor for the phase sensitivity, and hence limit the VLBI observations. Random fluctuations in the troposphere and ionosphere will cause uncorrelated phase errors on timescales beyond a few minutes. This destroys the phase coherence between distant VLBI stations, and will cause the ADEV to increase at averaging times beyond a few minutes.

Figure 18a shows the measured ADEV of the fringe between the WSRT and the DT. This observation consisted of two scans of 90 minutes each. The included line graphing the link performance (measured directly) demonstrates that the VLBI is not limited by the clock stability. At longer timescales, the results are determined by ionospheric effects and source structure, as expected.

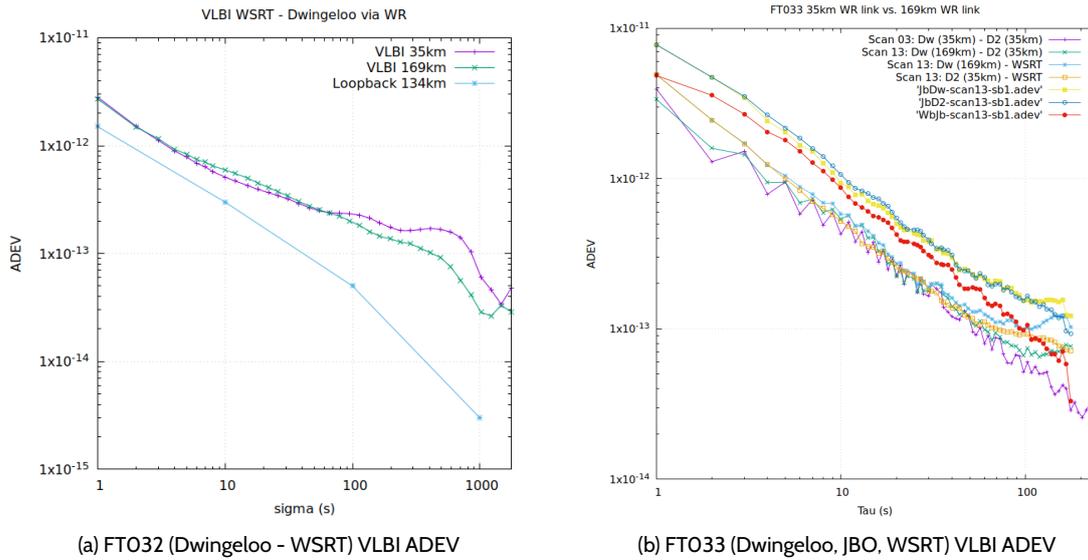


Figure 18: VLBI measurements, expressed as ADEV

During FTO33, the final VLBI observation of this project, we had access to a second USRP and WR switch, allowing a direct comparison between different WR links. The radio telescopes in Jodrell Bank and Torun were kind enough to also participate in this observation, although the Torun data turned out to have phase jumps, and no valid ADEV measurements could be extracted from it.

Figure 18b shows several measurements extracted from this observation. From lowest to highest ADEV, these are: Zero baseline (Dwingeloo receiver 1 against Dwingeloo receiver 2) carried out with a 35 km link against 35 km (purple line) and 169 km (green), showing no significant difference in performance. Fringes between the WSRT telescope and the two Dwingeloo receivers on each of the links show even less difference (blue, orange). The final three lines show the results on the baselines with Jodrell Bank, versus the long link (station Dw, yellow), the short link (station D2, green again) and the WSRT (Red). Comparing these results shows that the fringes between the WSRT and Jodrell Bank have a slightly better performance, but this is likely due to their better sensitivity compared to Dwingeloo, as at longer timescales, the results are nearly the same.

5 LOFAR WR Connection

Once the WR connection to the Dwingeloo telescope had been fully tested and characterized, it was time to connect LOFAR to the hydrogen maser at the WSRT as well. The SURFnet link was reconfigured to run from the WSRT, via Groningen (where the signal is amplified optically), to the LOFAR concentrator node in Exloo. We installed a White Rabbit switch (with low jitter daughter board) in the racks and the link was brought up.



(a) First WR light at LOFAR

(b) BDOA in Groningen

Figure 19: Connecting the WR link to LOFAR

For initial verification, the phase of the reference clock output was compared to that of a free running Rubidium clock. In the near future, the WR signal will become the timing reference of the LOFAR telescope. However, as this is an operational telescope that is constantly recording science data, such a switch-over has to be carried out with a lot of care and planning.

6 Conclusions

Our results show that, with our adaptations, White Rabbit can co-exist with other services on a production DWDM network. The timing accuracy of well below 1 ns, and a frequency stability within one order of magnitude of a hydrogen maser, make this solution a very interesting tool for radio astronomy and other sciences. Our measurements show no performance degradation between the 35 km and 169 km link. The VLBI observations have demonstrated that the remote reference clock offers sufficient phase stability to perform VLBI observations. Through these results we have shown that the existing fiber-optic network infrastructure is not only capable of transmitting high-capacity VLBI observational data, but also of distributing a common reference clock to VLBI stations at long distances. The implementation in the form of White Rabbit is especially important in view of the upcoming inclusion of White Rabbit in the next IEEE 1588 standard for time transfer.

The WR links that have been established to LOFAR and to the Dwingeloo telescope as part of this project, will remain in use after the conclusion of the ASTERICS project. This will offer improved timing to both observatories, and allow us to further study the performance of the WR links.

Noteworthy in this project was the cooperation of a volunteer organisation (CAMRAS), and their active support during the construction of the WR links to the historic Dwingeloo telescope, and during the VLBI observations. This mutually beneficial arrangement made it possible to not just measure the phase noise performance of the links, but to demonstrate that it is indeed sufficient for VLBI observations.

Bibliography

- B. Alachkar, A. Wilkinson, and K. Grainge. Frequency reference stability and coherence loss in radio astronomy interferometers, application to ska. *Journal of Astronomical Instrumentation*, 2018. URL <https://arxiv.org/abs/1801.05419v1>.
- P Boven. White Rabbit in Radio Astronomy. ICALEPCS, 2017. URL http://accelconf.web.cern.ch/AccelConf/icalepcs2017/talks/tucpl03_talk.pdf.
- G. Daniluk. White Rabbit calibration procedure. November 2015. URL <https://www.ohwr.org/project/white-rabbit/wikis/Documents/White-Rabbit-calibration-procedure>.
- ITU-T G.652. Characteristics of a single-mode optical fibre and cable. November 2016.
- ITU-T G.655. Characteristics of a non-zero dispersion-shifted single-mode optical fibre and cable. November 2009.
- J. C. J. Koelemeij and C. van Tour. Hardware for maser-level time & frequency distribution in public networks. *ASTERICS Deliverable D5.4*, 2017.
- M. Rizzi et al. White rabbit clock synchronization: Ultimate limits on close-in phase noise and short-term stability due to fpga implementation. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 65:1726–1737, 2018.
- R. T. Schilizzi, G. K. Miley, A. van Ardenne, B. Baud, L. Bååth, B. O. Rönnäng, and I. I. K. Pauliny-Toth. High resolution observation of the compact central component of the giant radio source 3c236. *Astronomy and Astrophysics*, 77, 1979.
- SRS. Measurements with the SR620. URL https://www.thinksrs.com/downloads/pdfs/applicationnotes/SR620_details.pdf.
- C. van Tour, R. Smets, P. Boven, R. Maat, P. Smets, J. C. J. Koelemeij, and A. Szomoru. Time Transfer in SURFnet/LOFAR network & general design rules for network implementation. *ASTERICS Deliverable D5.7*, 2018.