ngVLA Memo #22 Sub-Nanosecond Time Accuracy and Frequency Distribution through White Rabbit Ethernet

C. van Tour^{1,2} and J. C. J. Koelemeij^{* 1,2}

OPNT B.V., Amsterdam, the Netherlands
LaserLaB and Department of Physics & Astronomy, Vrije Universiteit Amsterdam, the Netherlands.
*Correspondence to: j.c.j.koelemeij@opnt.nl

We present a low-jitter, high-stability implementation of White Rabbit Ethernet for fiber-optic time and frequency distribution over long distances. Frequency stabilities averaging down from 7×10^{-13} at 1 s to 3×10^{-16} at 7,000 s are achieved. With the help of low-noise clean-up oscillators, the phase noise of the remote 10 MHz output frequency is strongly improved, thus reducing the rms phase jitter to 0.18 ps (integrated over the range 1 Hz – 100 kHz) while maintaining good long-term frequency stability. In addition, we study the effect of different types of off-the-shelf optical transceivers on the frequency stability when used in conjunction with a 100 km fiber-optic link. Time distribution with a residual time error on the order a few tenths of a nanosecond is routinely achieved in fiber-optic links with lengths in excess of 100 km, including optical multiplexers and bidirectional optical amplifiers. All results are obtained with commercially available, telecom-grade equipment, and as such they represent a significant step towards frequency distribution at hydrogen-maser-level stability through White Rabbit in public networks.

1. Introduction

The elements of the ngVLA require local oscillators (LOs) with low phase noise, ideally such that the total phase noise in the measurement is due to the atmosphere, not the instrument. Various studies have shown that fiber-optic links may be length-stabilized (or length-compensated) such that the frequency of hydrogen masers can be transferred over long (>100 km) distances without significant degradation of the long-term frequency stability (see for example [1,2]). The principle of length stabilization is based on the measurement of linkinduced phase variations observed after the frequency-carrying signal has made a roundtrip over the optical link [3], which may be done for optical carriers as well as radiofrequency (RF) signals transmitted over an optical carrier. The simplest implementations use amplitude modulation (AM) of an optical carrier at ~1 GHz frequencies, which is readily detected using a photodetector. More recently, a 'digital' protocol to distribute time and frequency through length-compensated fiber-optic networks was developed in the White Rabbit (WR) project, by a consortium gravitating around CERN [4]. WR is a strongly improved version of the Synchronous Ethernet and IEEE 1588 protocols, implemented in Gigabit Ethernet. Advantages of WR over previously reported methods are (i) the use of commercially available equipment and off-the-shelf high-volume components, (ii) a high level of standardization complemented by open hardware and software, and (iii) the built-in scalability of WR enabling large networks consisting of thousands of nodes, served through complex network topologies including branches and loops. Although WR is not yet stable enough to provide 'hydrogenmaser-level' frequency stability, we point out that the bitrate of WR is comparable to the AM frequencies used in previous 'analog' fiber-optic time and frequency transfer techniques [1,2]. This implies that with proper electronic and digital hardware design, WR may be upgraded so that it could transfer the time and frequency of a single hydrogen maser to many clients, and over long distances, without significant degradation of the stability, and at substantially lower cost than required for deployment and maintenance of hydrogen masers. Such cost-effective time, frequency and data transport through a single wavelength channel forms an interesting concept for current and future multi-element instruments. WR is currently under investigation for use in various antenna and astroparticle detector arrays, including the SKA, CTA, KM3NeT, LOFAR, and the European VLBI network (under the umbrella of the EU H2020 project ASTERICS [5]).

Here, we present improvements (in terms of phase noise and long-term frequency stability) over default WR for fiber-optic time and frequency distribution over long distances. This memo is organized as follows. In

Section 2, we briefly revisit our original proposal to improve the frequency jitter using a frequency lock based on measurements done with a high-precision frequency counter, and we provide arguments *a priori* showing that the improvement in jitter would go at the expense of long-term stability, leading us to abandon this idea. In Section 3 we present an alternative scheme, based on CERN's low-jitter daughterboard for the WR switch [6], used with and without a low noise, phase locked clean-up oscillator. Results are presented and discussed in Section 4, and prospects of long-haul time and frequency distribution (including a few experimental results towards its feasibility) at a stability level compatible with hydrogen masers are discussed in Section 5.

2. Originally proposed approach

In the original work plan of this ngVLA Community Studies work item, we proposed a 'digital' method to frequency lock a low-phase-noise frequency source (LO 2 in Fig. 1) to the 10 MHz output of a WR Slave device. In this case, the latter is a WR switch in Slave mode, locked to a WR switch in Grandmaster (GM) mode over a given (possibly long) length of optical fiber. The WR grandmaster device itself is locked to another low-noise frequency source (LO 1 in Fig. 1). Here, LO 1 is supposed to play the role of a hydrogen maser, while LO 2 is intended as a low-cost replacement for a second hydrogen maser in the remote (Slave) location.



Figure 1: Schematic view of the proposed enhanced WR fiber-optic low-noise frequency, time and data link. The highresolution frequency counter acts as a narrowband bandpass filter with a peak response around the inverse gate time, and should suppress residual high frequency noise of the WR slave output [7]. A slow frequency lock is implemented in software using the Raspberry Pi. VCTCXO, voltage controlled oscillator; DAC, digital-to-analog converter; GM, grandmaster; LO, local oscillator; PPS, pulse per second.

While the long-term stability of off-the-shelf WR equipment is quite acceptable for many applications (Allan deviation (ADEV) typically averaging as $1/\tau$, starting at about 1×10^{-11} at 1 s averaging), its 10 MHz output is still too noisy for frequency transfer at stability levels achieved by hydrogen masers (typically of the order of 1×10^{-5} ¹³ at 1 s averaging). Our initial plan was to clean up the high-frequency noise of off-the-shelf WR equipment by frequency-locking the clean-up oscillator LO 2. This would be achieved in a digital way, using a frequency counter (Agilent 53220A) in high-precision mode and a gate time of 1 s or longer (in order to average out the high-frequency noise), whose output would be processed by a low-bandwidth, software-controlled servo loop controlling the frequency of LO 2. In high-precision mode, the counter should be able to measure RF frequencies with close to 12 digits of resolution within one second gate time. Owing to the digital nature of the servo loop, it would be relatively easy to explore for which parameters (loop bandwidth, counter gate time) the best frequency-lock performance would be achieved. Thus, a frequency stability of about 1×10^{-12} at 1 s averaging was expected to be achieved, further averaging down as $1/\tau$. However, in earlier work we proved mathematically that the ADEV, obtained with a counter operated in the high-precision mode described above, displays at best $1/\tau^{1/2}$ behavior [7]. Thus, the gain in short-term frequency stability would go at the expense of long-term frequency stability. Because of this realization, as well as the limited improvement of the short-term stability from 1×10^{-11} to 1×10^{-12} at 1 s, we choose to follow a different approach, as detailed further below.

3. Phase noise and frequency stability improvements to WR

In order to improve the phase noise and frequency stability, we made two modifications to the WR switches used in this project. The first modification consists in the implementation of a so-called low-jitter daughterboard (LID), which was recently developed at CERN, and which contains phase locking electronics with improved jitter and thermal stability [6]. The two LIDs used here were home built, and each was integrated into a WR switch (Fig. 2). As measurements performed at CERN previously pointed out, the LID leads to an improvement of the performance of a WR link in several ways. At this point it is good to recall that WR utilizes several phase-locked loops (PLLs) to synchronize two WR devices. Any WR link will involve a Master device and a Slave device. The Master device transmits a Gigabit Ethernet (GbE) code, at a bitrate synchronous with its internal clock (which involves several PLLs). The Slave device synchronizes the frequency of its own internal clock with the incoming bit stream, and sends a second GbE code back to the Master over the same fiber (involving a few more PLLs). The Master precisely determines the round-trip phase from the returned GbE code in order to detect phase fluctuations induced by length variations of the link. The Master then relays this information to the Slave, which corrects its internal clock for the link length variations, and also corrects for the one-way propagation delay so that the Master and Slave clocks run synchronously in time, phase and frequency. This link stabilization method (which is an important part of the WR protocol) can be regarded as a 'distributed' PLL. Obviously, the performance of the entire system will depend on the quality of the local oscillators used in PLLs, as well as the phase stability of the various clock shaper circuits used in WR devices.



Figure 2: Interior of one of the modified WR switches. The LJD and PLO are indicated by the red frames.

Yet another PLL is involved if one Master device is configured as GM. In this case, the GM phase locks its internal clock to the 10 MHz and 1 PPS inputs provided by a reference clock (e.g. a hydrogen maser). The improvements realized through the LJD are twofold. First, if used in a WR GM switch, it improves the GM phase lock by use of low-noise local oscillators and clock shaper circuits with improved thermal (phase) stability. Note that the LJD electronics bypass the corresponding electronics in the original WR switch. Second, if used in a WR Slave, the LJD reduces phase noise, again through the use of low-noise local oscillators. Measurements performed at CERN (and independently reproduced by us) show that in the GM, the improvement in phase noise power spectral density (PSD) is two orders of magnitude in the range 1 - 1000 Hz, while the long term stability (ADEV) improves by one order of magnitude ($\tau = 0.01 - 1000$ s). In the Slave, the LJD improves the PSD by two orders of magnitude in the range 1 - 20 Hz, while the ADEV improves by about a factor of five [6].

The second modification made to improve the phase noise and frequency stability of the WR switches is the use of an oven controlled, low noise clean-up oscillator (OCXO) that is phase locked to the 10 MHz frequency outputs of the WR devices. To our knowledge this concept has not been reported in the context of WR before. In the WR GM switch, the 10 MHz output signal essentially is the 10 MHz frequency of an internal local oscillator (LO) that itself is phase locked to the external reference clock. In the GM Slave switch, the 10 MHz output frequency is derived from the Slave's internal clock, which (in our setup) is synchronized to the internal

clock of the WR GM through the WR protocol. The clean-up oscillators used here (further denoted as PLOs) are commercially available, packaged devices (Wenzel Associates PLO 500-30822A), which are reasonably affordable (prices well below USD 1k even for single units). Available PLL bandwidths are 1 Hz and 0.1 Hz. A schematic of the setup with the two improvements (LJD and PLO) used here is shown in Fig. 3. While ideally one would compare the WR Slave 10 MHz output with the reference clock (Rb or Cs clock in our case), we are comparing the PLO connected to the WR GM (PLO 1 in Fig. 3(b)) with the PLO connected to the WR Slave (PLO 2). The underlying reason is that both PLOs are essentially free-running on short timescales (because of their small PLL bandwidth), while having much better stability than the Rb and Cs clocks at the same timescales. A direct comparison of PLO 2 with the atomic reference clock would therefore be limited by the instability of the atomic clock at short timescales. One way to circumvent this is to use a hydrogen maser as a reference clock (which has sufficient short-term stability, but which we did not have available during the measurements).



Figure 3: (a) Setup used to characterize the performance of the default WR switches (indicated as WR GM and WR Slave), as well as the WR switches equipped with the LID. The LID is part of the WR switches and not shown here. (b) Phase locked clean-up oscillators PLO 1 and PLO 2 are used to reduce the phase noise of the WR GM and WR Slave 10 MHz outputs, respectively. The Microsemi 3120A performs phase noise and frequency stability measurements, comparing the signal provided at the device-under-test (DUT) input with that fed into the reference (REF) input.



Figure 4: (a) Long-term frequency stability (ADEV, ENBW 50 Hz) of WR over a short (2 m) fiber link, for the different configurations studied here (see also Fig. 3). Red curve, default WR; blue curve, WR equipped with LJD in both the WR GM and WR Slave devices; green curve, WR enhanced by both LJD and PLO (with 1 Hz PLL bandwidth); dashed curve, WR enhanced by both LJD and PLO (theoretical single-device stability); brown curve, two free-running PLOs versus each other; grey curve, measurement system noise floor. (b) Frequency instability measured at different ENBWs. Blue curve, ADEV (ENBW 50 Hz) obtained with WR switches equipped with LJD; purple curve, same as blue curve, but ADEV obtained with 0.5 Hz ENBW; black curve, active hydrogen maser ADEV obtained with 0.5 Hz ENBW (manufacturer specification); grey curve, noise floor ADEV (ENBW 0.5 Hz).

4. Results

As we are primarily interested in using WR as a method to distribute the signals of a reference clock (*e.g.* a hydrogen maser) to remote locations, we will focus on the performance of the WR Slave switch. Figure 4(a) shows both the short-term and long-term stability (ADEV, measured at 50 Hz equivalent noise bandwidth (ENBW)) obtained with default WR switches, with WR switches equipped with the LJD (see also Fig. 3(a)), and WR switches equipped with LJDs and PLOs (Fig. 3(b)), respectively. The LJD improves the stability by about a factor of five, in good agreement with results earlier obtained by Rizzi *et al.* [6]. The LJD implementation improves the root mean-square (rms) phase jitter, integrated over the range 1 Hz – 100 kHz, from 13 ps to 2.9 ps. It should be noted that this is the combined jitter contributed by both the GM and Slave 10 MHz outputs. Therefore, the performance of an individual output may be better, and given the similarity of the circuits used for both 10 MHz outputs, the rms jitter of the individual outputs may be a factor $\sqrt{2}$ smaller.

Adding the PLOs has a significant impact on the frequency stability and phase noise. With the setup of Fig. 3(b), using PLOs with 1 Hz PLL bandwidth, the ADEV is reduced to the low 10⁻¹² range for averaging times between 0.01 and 0.1 s, during which the PLO is essentially in a free-running state. Around 1 s a local maximum is visible, which we ascribe to the action of the PLO's internal servo loop. From 2 s onwards, the servo loops takes over, forcing the PLO phase to follow that of the 10 MHz output of the WR Slave's LJD. On longer timescales, the ADEV with the PLO remains about a factor two larger than that of the LJD only. In Fig. 4(a) we also plot the ADEV of two free-running PLOs. Comparing this with the ADEV of the LJD-PLO setup suggests that the PLO's servo loop starts having an effect on timescales as short as 0.1 s, leading to a degradation of the stability. It would be interesting to compare this situation with a setup equipped with PLOs with 0.1 Hz servo bandwidth, as this might lead to (somewhat) improved stability. Unfortunately, however, one of the two 0.1 Hz devices we had available turned out to be malfunctioning during the measurement campaign.

In terms of phase noise, the PLOs lead to a huge improvement. In Fig. 5 phase noise power spectral densities are shown, obtained for the various configurations studied in this work. While the LID implementation already significantly improves the default WR Slave phase noise in the low-frequency range, the PLO leads to a further suppression of noise by 30 to 40 dB over the frequency range 10 Hz - 100 kHz, with a high frequency noise floor at -168 dBc/Hz. Several spurs are visible in the range 50 Hz - 1 kHz, but none of them exceeds -122 dBc/Hz. As a result, the rms jitter of the 10 MHz output is as low as 0.18 ps over the range 1 Hz - 100 kHz, an improvement by nearly two orders of magnitude over default WR.

Given the widespread use of hydrogen masers in multi-element radio astronomy, it is interesting to compare the results obtained here with that of a commercial active hydrogen maser (Microsemi MHM 2010, standard option). When comparing ADEV stability measures, one must ensure that ADEV values are obtained using the same ENBW. The results of Fig. 4(a) were obtained with an ENBW of 50 Hz, whereas the MHM 2010 ADEV is specified for an ENBW of 0.5 Hz. To make a fair comparison, we show in Fig. 4(b) ADEV traces for the WR LDD setup measured both with 50 Hz and 0.5 Hz ENBW. The latter may be directly compared with the ADEV of the hydrogen maser, also shown in Fig. 4(b). Another comparison with hydrogen masers can be made in terms of the phase noise. Figure 5 displays the specified upper limit of the power spectral density of the MHM 2010 maser's 10 MHz output, which is comparable to the phase noise of the WR Slave with LD and PLO.

The results presented here suggest that the improved WR system may in certain cases provide a good alternative for hydrogen masers. This may be true in particular for situations in which the observational phase coherence is limited by atmospheric conditions, rather than by local-oscillator stability. Another advantage of a WR-based frequency distribution system is that it also distributes phase. In other words, all nodes in the WR network will have the same phase, which may allow reducing the 'search window' in initial fringe searches. Furthermore, the WR PLL electronics may be enhanced by higher quality LOs (which would be an extension of the work by Rizzi *et al.* [6]), so that the intrinsic phase noise and stability are improved further, and additional clean-up oscillators are not necessary.



Figure 5: Phase noise power spectral density of the 10 MHz WR Slave output for the various configurations studied here. Red trace, default WR; blue trace, WR equipped with LJD both in the WR GM and WR Slave device; green trace, WR enhanced by both LJD and PLO (with 1 Hz PLL bandwidth); brown trace, two identical free-running PLOs; grey trace, measurement system noise floor; black curve, hydrogen maser phase noise (10 MHz output, manufacturer specification). Rms phase jitter values, integrated over the range shown, are 13 ps for default WR, 2.9 ps for WR with LJD, and 0.18 ps for WR with LJD and PLO).

5. Prospects for long haul hydrogen maser time and frequency distribution

The results presented in Sec. 4 were obtained using a short (few meters) fiber-optic link. However, for largescale multi-element facilities such as the ngVLA, time and frequency distribution would have to be realized over distances in excess of 100 km. Over such distances, the damping of real-life fiber-optic links may readily exceed 30 dB. Given the 35-40 dB maximum sensitivity of GbE optical receivers, this implies that some form of regeneration may be required if links substantially longer than 100 km are considered. However, any form of regeneration will add noise to the optical signal (and therefore degrade the frequency transfer stability). Of all available methods, regeneration by optical amplification offers the best performance, and optical amplifiers are widely used in long-haul fiber-optic telecommunications systems. Although designed as a GbE system for communication over relatively short range (<10 km), one may ask if WR too could be used in conjunction with optical amplifiers to extend its range. A known issue here is that WR differs from typical long-haul telecommunications systems in that it requires *bidirectional* optical transmission, which is incompatible with the unidirectional nature of the optical amplifiers used in long-haul telecommunication systems. In recent years, we have developed bidirectional amplifiers and demonstrated their suitability for WR [8]. We have furthermore shown that WR can be used to transfer time over 320 km in live networks with a residual time error of less than 1 ns [9]. In our laboratory, we routinely perform link and equipment calibrations such that residual time errors below 100 ps are achieved over a distance of 100 km.

Running WR on longer links introduces two potential additional issues which may affect both the short-term and long-term stability of the system. First, the long fiber link leads to an additional delay in the WR 'distributed' phase measurement system that detects and corrects for path length changes, thus reducing the bandwidth of the control loop. Also, the random variations of the link will generally be larger, which (especially in combination with the lower control bandwidth) will lead to increased noise. An excellent theoretical treatment of these effects can be found in [3]. Notwithstanding these issues, link stabilization based on roundtrip measurements of RF-modulated optical signals has produced impressive results, demonstrating that hydrogen-maser-level frequency transfer may be done over distances well beyond 100 km [1]. Second, as WR communicates using different upstream and downstream wavelengths, a differential delay due to chromatic dispersion (CD) exists, which leads to a phase offset. While the fixed part of this phase offset can be calibrated and digitally removed, this is less easy to accomplish for the time-varying part of the CD differential delay. One potential source of such varying delays is wavelength drift of the optical transmitter [10]. For example, a drift over 0.01 nm over the course of 10,000 s (induced by *e.g.* a slow temperature change of the equipment) of an optical wavelength, transmitted over a 100 km link with a 'typical' CD coefficient of order 10 ps/(nm km), would cause a relative frequency error of order 10^{-15} . One remedy consists in using temperature stabilized optical transmitters, such as those used in Dense Wavelength Division Multiplexing (DWDM) networks, for which stable wavelengths are required.



Figure 6: WR link configurations used during the comparison of various types of SFP transceivers. (a) Back-to-back (connected through a few-meter-long fiber patch cord). (b) 100 km of spooled fiber. (c) 100 km of spooled fiber with a bidirectional optical amplifier (OPNT, delay calibrated) located halfway. All equipment was kept within the same room.



Figure 7: (a) Frequency instability (ADEV, ENBW 0.5 Hz) between the 10 MHz outputs of the WR Slave and reference clock for the 100 km WR link with default switches, CWDM SFPs, and one bidirectional amplifier (blue curve). The red curve was obtained with the same configuration, except for the switches now carrying DWDM SFPs. (b) ADEV curves (0.5 Hz ENBW) measured for a WR link (default WR switches, CWDM SFPs, no optical amplifier) running over 2 m of fiber (red curve) and 100 km of fiber (blue curve). (c) Same as (b), but now showing the phase noise power spectral density for the 2 m fiber link (red trace) and 100 km fiber link (blue trace). The vertical dashed line indicates the servo spike located at ~510 Hz which one would expect to observe in the case of noise dominated by link delay variations and a stabilization system with infinite control bandwidth, cut off by due to the finite roundtrip time through the 100 km of optical fiber [3].

Due to partially unavailable equipment near the end of this project, we were only able to test the WR switches equipped with LID and PLO in combination with a short fiber link. In the future we plan to conduct tests of the LJD and PLO WR implementations in our test link, which consists of 100 km of spooled fiber, and which can be equipped with a bidirectional optical amplifier halfway. Nevertheless we have been able though to perform measurements on default (unmodified) WR switches connected through the 100 km fiber link with and without bidirectional optical amplifier (Fig. 7). We have compared the long-term stability using commercially available Coarse Wavelength Division Multiplexing (CWDM), small form-factor pluggable (SFP) transceivers, with the stability achieved with temperature-stabilized DWDM SFPs. The results are shown in Fig. 7(a) and Fig. 8, and indicate that in the case of CWDM SFPs there is a significant correlation between phase error and (diurnal) ambient temperature changes, which appears to be absent for DWDM SFPs. The ADEV curves in Fig. 7(a) indicate how these presumed wavelength variations impact the frequency stability. As long as the link stabilization works properly (*i.e.* the round-trip phase being stabilized to nominally the same, fixed reference value), the ADEV decreases as $1/\tau$. However, if the phase reference value drifts in an uncontrolled way (for example due to a time-varying CD delay asymmetry), the additional phase noise tends to reduce the ADEV slope. From Fig. 7(a) it can be seen that for our 100 km link, the ADEV starts levelling off around τ =25 s if CWDM SFPs are used, whereas with DWDM SFPs the ADEV starts to deviate from $1/\tau$ behavior only at τ =250 s. We attribute this behavior to the better wavelength stability of DWDM SFPs (whose use seems mandatory if hydrogen-maser-level stabilities of order 10⁻¹⁵ are required). A more thorough test of this hypothesis would involve online measurement of the actual wavelengths of the SFPs versus ambient temperature and the observed phase error.



Figure 8: (a) Phase difference between the 10 MHz outputs of the reference clock and WR Slave (red curve) and ambient temperature (blue curve) versus time for the 100 km WR link (default switches with CWDM SFPs and one bidirectional amplifier). (b) Same as (a), but for WR switches equipped with DWDM SFPs. The data shown in (a) and (b) were taken sequentially.

Finally, we performed a measurement (with default WR switches equipped with CWDM SFPs) to verify that even for 100 km link length, the reduced link stabilization bandwidth (due to fiber roundtrip propagation delay) and additional fiber noise do not contribute to the phase noise and frequency instability (at least at the level of stability achievable with default WR switches). Figure 7(b) shows the short-term frequency stability of the WR Slave relative to the reference clock, both for a short (2 m) link and a 100 km fiber link. The ADEV curves do not show any significant difference for the two configurations. The power spectral density, measured at the WR Slave 10 MHz output, also shows no difference between the two configurations (Fig. 7(c)). This indicates that in default WR switches, the Slave phase noise is not dominated by the limitations associated with long fiber links as described above. In conclusion, we have shown that commercially available WR equipment can be readily upgraded such that RF frequencies may be transferred with low phase noise and stability approaching that of a hydrogen maser of fiber-optic links. We have also addressed a few potential issues associated with fiber links of 100 km length or more, and we routinely achieve time offsets of 100 ps and below over links involving 100 km of spooled fiber and bidirectional optical amplifiers. While earlier work on length-compensated fiber links suggests that hydrogen-maser-level frequency transfer over long-haul optical links is well feasible, it will take further research to find out whether upgraded WR equipment as described here is in fact capable of doing so. To this end, the WR hardware might also receive a second upgrade with improved LOs and PLLs, thus extending the recent work of Rizzi *et al.* [6], to improve the intrinsic phase noise and stability further.

Acknowledgements

This work was supported in part by the ASTERICS project, funded through the European Union's Horizon 2020 research and innovation program under grant agreement No 653477, and in part by the Netherlands Organization for Scientific Research (NWO-TTW) through projects 12346 and 13970. We thank Lex van der Gracht, Adrian Savencu, Konstadinos Mourtziapis, Gertjan de Graaff, and Rob Smets for fruitful discussions and technical assistance during the experiments.

References

- 1. O. Lopez et al., Eur. Phys. J. D 48, 35-41 (2008).
- 2. M. Fujieda et al., J. Nat. Inst. Info. Comm. Tech. 57, 209-218 (2010).
- 3. P. A. Williams, W. C. Swann, and N. R. Newbury, J. Opt. Soc. Am. B 25, 1284 (2008).
- 4. https://www.ohwr.org/projects/white-rabbit/wiki
- 5. https://www.asterics2020.eu/cleopatra
- 6. <u>https://www.ohwr.org/projects/wr-low-jitter</u>
- C. van Tour, "Ultrastable optical frequency transfer over actively compensated amplified fiber links", M.Sc. thesis, 2016, Delft University of Technology (unpublished).
- 8. E. Dierikx et al., IEEE Trans. Ultrason., Ferroelect. Freq. Control 63, 945-952 (2016).
- 9. <u>http://www.opnt.nl/news/accurate-time-synchronization-in-live-network-vodafone/</u>
- 10. H. Li et al., IEEE Trans. Nucl. Sci. 62, 1021-1026 (2015).