White Rabbit clock characteristics

Mattia Rizzi
CERN, Geneva, Switzerland
University of Brescia, Italy

Maciej Lipiński, Tomasz Wlostowski
Javier Serrano, Grzegorz Daniluk
CERN, Geneva, Switzerland

Paolo Ferrari
Stefano Rinaldi
University of Brescia, Italy

***Abstract*—White Rabbit (WR) extends the Precision Time Protocol (PTP) to provide synchronisation with sub-nanosecond accuracy and sub-50 picoseconds precision. The protocol aspects of the WR extension are currently studied and integrated into the upcoming revision of PTP. In the context of this PTP revision, mechanisms are added to allow the control of the Layer 1 (L1) syntonisation by the PTP protocol. This article focuses on the frequency transfer characteristics of the L1 syntonisation in WR.**

We first explain the interaction between L1 syntonisation and PTP synchronisation in a WR device and describe the architecture of its Phase-Locked Loop (PLL). We then characterize the frequency transfer through a WR network in two ways: measuring the characteristics of the WR switch according to the Synchronous Ethernet (SyncE) metrics defined in ITU-T G.8262, and performing phase noise analysis. The results of the measurements allow us to propose improvements that might be useful for different types of WR applications. Metrology laboratories might be interested in the optimisations made to significantly reduce the phase noise. On the other hand, the telecom industry might be interested in the modifications that make the WR switch SyncE-compliant but deteriorate its performance. Notably, the latter was achieved merely by modifying the software that implements the WR PLL.

I. Introduction

The White Rabbit (WR) project is a multilaboratory, multicompany, and multinational collaboration to develop a versatile solution for control and data acquisition systems where sub-nanosecond synchronisation accuracy is required. The project was started within an effort to renovate the CERN control and timing system, and initially all WR applications concerned accelerators. The open nature of the project and the fact that WR is based on widely-used standards, made it a preferred solution in a much wider range of applications. All of these applications benefit from the high synchronisation accuracy that is provided by the WR extensions to the Precision Time Protocol (WR PTP) [1] and its implementation.

The protocol aspects of WR are studied by the P1588 Working Group [2] and constitute the basis for features and a profile that are likely to be included in the new revision of the IEEE1588 standard. The protocol aspects are intended to support implementation-specific mechanisms that provide the high accuracy.

There are two key elements that distinguish the implementation of WR PTP from other PTP implementations:

- tight cooperation between the PTP synchronisation and the Layer 1 (L1) syntonisation, and
- enhanced timestamping precision using phase detection.

This article analyses the implementation and the performance of the L1 syntonisation in WR and then studies its possible further improvements. It first explains the architecture and implementation of the Phase-Locked Loop (PLL) used in WR and the cooperation between PTP synchronisation and L1 syntonisation. The current implementation of the L1 syntonisation in WR is characterised using Synchronous Ethernet (SyncE) metrics and phase noise transfer analysis. These measurements allow to suggest a number of possible modifications to optimize phase-noise and to achieve compliance with SyncE. In both cases, the impact on the L1 syntonisation and PTP synchronisation performance is evaluated.

II. PTP synchronisation and L1 syntonisation in White Rabbit

The relation between the PTP synchronisation and the L1 syntonisation is described in this article as the relation between the *local PTP clock*, the *L1 tx clock signal* and the *L1 rx clock signal*, all depicted in Fig. 1. The *local PTP clock* of a WR node provides the node's local estimate of the time of the Grandmaster (GM) to which it is synchronised. This clock has a *time counter* that is a digital time representation incremented at each rising edge of the *local PTP clock signal*. The *L1 rx clock signal* is recovered from the reception of data from the medium while the *L1 tx clock signal* is used in the transmission of data over the medium.

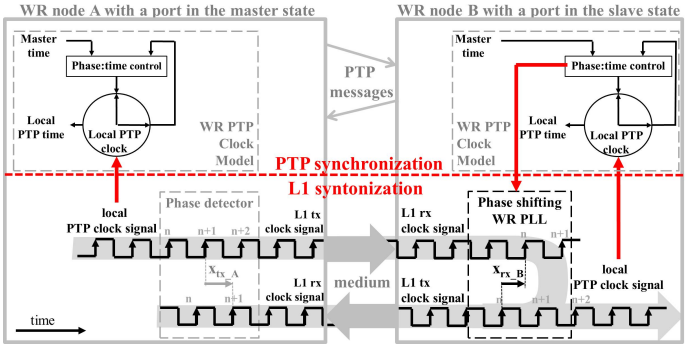

Fig. 1. L1 and PTP clock signals.

Unlike in many PTP implementations, in WR the PTP synchronisation and the L1 syntonisation are made to tightly cooperate. In particular, the local PTP time and the *L1 tx/rx clock signals* are congruent and coherent in the WR network [3]. This is because each of the WR nodes implements the *WR PTP clock model* depicted in Fig. 1. Thus, on a link directly connecting two WR nodes A and B, the *L1 tx clock signal* transmitted by node A at its port in the master state is syntonised to its *local PTP clock*. WR node B has its port in the slave state and its *local PTP clock* is syntonised to the *L1 rx clock signal* received from node A. As a result, the *local PTP clocks* of both WR nodes are frequency-traceable to the same source, their GM, without intervention of the PTP protocol. The PTP measurement of link-delay and offset from

the master is used only to adjust the value of the *time counter* and the phase of the *local PTP clock* of WR node B.

The synchronisation in WR is maintained by adjusting the phase (phase-steering) rather than manipulating the *time counter* value. The WR PLL not only syntonises the *local PTP clock* to the recovered *L1 rx clock signal* but also maintains the desired phase offset between these two clock signals, a value so-called *setpoint*. The architecture of the WR PLL is explained in the next section.

III. WR Phase-Locked Loop

The WR PLL, detailed in [4], is a phase-shifting digital PLL that uses Digital Dual Mixer Time Difference (DDMTD) [5] to obtain the phase error between the input clock signals.

A. Digital Dual Mixer Time Difference (DDMTD)

The DDMTD uses digital mixing to produce output clock signals of lower frequency than that of the input clock signals. The operation of DDMTD is explained in Fig. 2. The input clock signals are sampled with D-type flip-flops that are clocked with an offset clock signal, clk_{DDMTD}, generated from one of the inputs. Its frequency, f_{DDMTD}, is very close

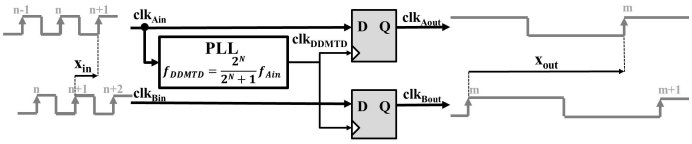

Fig. 2. Digital Dual Mixer Time Difference phase detector.

to that of the input clock signal, f_{in} and is specified as follows:

$$f_{DDMTD} = \frac{2^N}{1+2^N} \cdot f_{in} \quad (1)$$

where N is an implementation-specific value that is 14 in WR. The sampling operation performed by the flip-flops is similar to analog mixing and low-pass filtering. Thus, the output clock signals, clk_{Aout} and clk_{Bout}, are of a frequency that is proportional to the frequency of the input clock signals. The phase, expressed in radians, between the input signals is equal to that between the output signals. Therefore, the time-difference between the edges of the input and output clock signals is proportional and can be expressed as follows:

$$x_{in}[ns] = \frac{1}{1+2^N} \cdot x_{out}[ns] \quad (2)$$

The output time-difference is significantly larger and it can be measured with much greater precision than the input phase. It is used to calculate the time-difference between the input clocks. This technique is used in the WR PLL described in the next section.

B. DDMTD-based software WR PLL (SoftPLL)

The architecture of the SoftPLL is depicted in Fig. 3. It uses DDMTD implemented in a Field Programmable Gate Array (FPGA) to compare the *local PTP clock signal* to an input *clock signal* that can be either:

- the *L1 rx clock signal* recovered at the slave port, or
- the *clock signal* coming from an external reference.

The outputs of the DDMTD are lower-frequency clock signals, as explained in III-A. The rising edges of these signals are timestamped using a *time counter* incremented by the *DDMTD clock signal*. These timestamps, called phase-tags, are fed into the software implementation of a Proportional-Integral (PI) controller that runs in an embedded CPU [6] inside the FPGA. The controller steers two Voltage-Controlled Crystal Oscillators (VCXO); FRETHE025 generates the *DDMTD clock signal*, VM53S3 generates *local PTP clock signal*. Fig. 3 depicts a simplified block diagram of the WR PLL, which actually consists of two PLLs: Helper and Main.

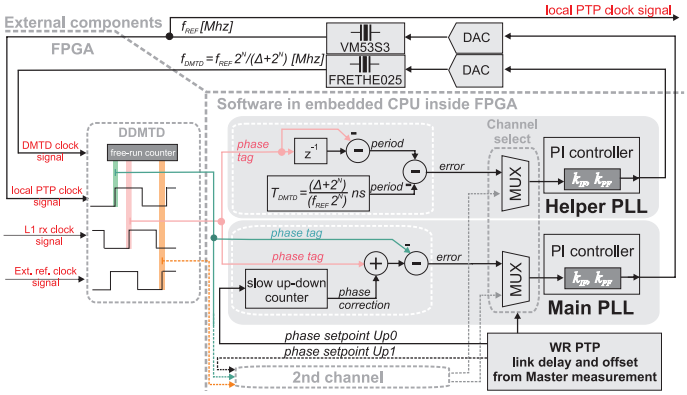

Fig. 3. Overview of the White Rabbit Phase-Locked Loop design.

The *Helper PLL* controls the *DDMTD clock signal*. This PLL works by comparing the difference between subsequent phase-tags to the "ideal" period of the *DDMTD clock signal*.

The *Main PLL* controls the *local PTP clock signal* that is a copy of the *L1 rx clock signal*, phase shifted by a programmable *setpoint* that is provided by the WR PTP. The PLL works by comparing the phase-tags of the *L1 rx clock signal* to the phase-tags of the *local PTP clock signal*, corrected for the *setpoint*. Any change of the *setpoint* value is applied with an LSB-step increment.

The WR switches use a 62.5MHz clock signal which determines a number of characteristics of the SoftPLL. The SoftPLL that receives a 62.5MHz input signal produces a *DDMTD clock signal* of 62.496185 MHz. The down-converted clock signals produced by the DDMTD have frequency of a 3.814kHz. Since each rising edge of these clock signals is timestamped, the phase-tags are provided to the SoftPLL at the DDMTD frequency. This is indeed the sample rate of the SoftPLL and so its Nyquist frequency is 1.9kHz. With the current parameters of the SoftPLL, the resolution of its phase-tags is 0.977ps and its bandwidth is 30Hz.

The SoftPLL determines the characteristics of the frequency transfer through a WR switch. These are characterised in the next section according to the ITU-T G.8262 guidelines.

IV. SyncE characteristics of L1 syntonisation

The characteristics of the frequency transfer through a WR switch are measured according to ITU-T G.8262 recommendation for a synchronous Ethernet equipment slave clock. Such a measurement allows to compare the "WR clock" with the "SyncE clock". Table I summarizes the results of tests that are

Test name	G.8262 section	Test result	Measured values	Setup
Frequency offset	6	Passed	4.256ppm	Fig 4-1
Pull-in range	7.1	Passed	8ppm	Fig 4-1
Hold-in range	7.2	Passed		Fig 4-1
Pull-out range	7.3	Passed	8.8ppm	Fig 4-1
Wander generation	8.1	Passed	Fig. 5	Fig 4-2 Fig 4-3
Jitter generation	8.3	Passed	pk-pk: 0.01UI RMS:$<$ 0.01UI	Fig 4-2
Wander tolerance (op-1 only)	9.1.1	Failed	SoftPLL unlocks above f=1.0 [Hz], A=0.25 [μs]	Fig 4-4
Jitter tolerance	9.2	Failed	SoftPLL unlocks	Fig 4-2
Wander transfer	10	Failed	Fig. 6	Fig 4-4

TABLE I
WR CLOCK CHARACTERISTICS ACCORDING TO ITU-T G.8262 METRICS.

described in a number of documents available on the WR web-pages dedicated to tests [7]. The table does not include tests of transient response and holdover (section 11 of ITU-T G.8262). These features are not supported by the current release of the WR switch. Tests of solutions under development can be found in [8].

The results in Table I were obtained using a dedicated SyncE tester, Calnex Paragon-X, and general-purpose measurement equipment. The measurement setups are depicted in Fig. 4 and described below:

1) A CS4000 Cesium Frequency Standard (Cs) is the external reference for the Paragon-X that is syntonised to the free-running device under test (DUT) over a 1GbE fiber link.
2) The CS4000 is the external reference for the Paragon-X. The DUT is syntonised to Paragon-X over 1GbE fiber link connected to port 2 (P2). The Paragon-X is syntonised to the DUT over a 1GbE fiber link on port 1 (P1).
3) The CS4000 is the external reference of the GM. The 10MHz output of the GM is used as the external reference for the CNT-91 time interval counter (TIC). The DUT is syntonised to the GM switch over a fiber link. The CNT-91 TIC is configured to take Time Interval Error (TIE) measurements every 1ms. The acquired data is filtered above 10Hz.
4) The CS4000 is the external reference for an Agilent 33250A function generator and two CNT-91 TICs. A customized GM WR Switch is fed with a phase-modulated 10MHz clock signal generated by the function generator. The GM is modified such that it does not use the SoftPLL to syntonise to the input 10MHz clock signal. So, the GM switch uses the input signal directly to generate the *L1 tx clock signal*. The DUT is syntonised to the *L1 tx clock signal* of the GM switch over a 1GbE fiber link. The two CNT-91 TICs are configured to make Time Interval Error (TIE) measurements of the 10MHz outputs of the GM and the DUT at a 1kHz sampling rate.

The tests results in Table I show clearly that the currently available WR switches are not compliant with SyncE. Although the wander and jitter generation of the WR switch are orders of magnitude better than required by ITU-T G.8262, the WR switch fails the tests of wander transfer as well as wander and jitter tolerance. Fig. 6 shows that the transfer function of the WR switch has a bandwidth of 30Hz and a phase gain of 3.3dB at 16Hz while ITU-T G.8262 requires a gain smaller than 0.2dB and a bandwidth between 1Hz and 10Hz for EEC-Option 1, and 0.1Hz for EEC-Option 2.

The reasons for the failures of the tests are investigated in section VII, which proposes modifications to the SoftPLL that make the WR switch compliant with SyncE. Section VII provides also measurements of WR performance with the proposed changes. The section that follows characterizes L1 syntonisation using phase noise analysis.

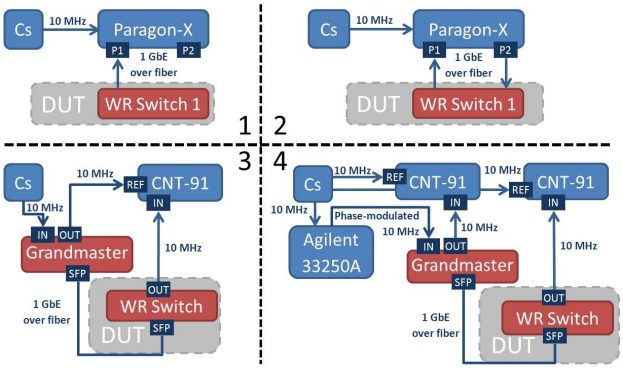

Fig. 4. Set-ups used to obtain WR clock characteristics.

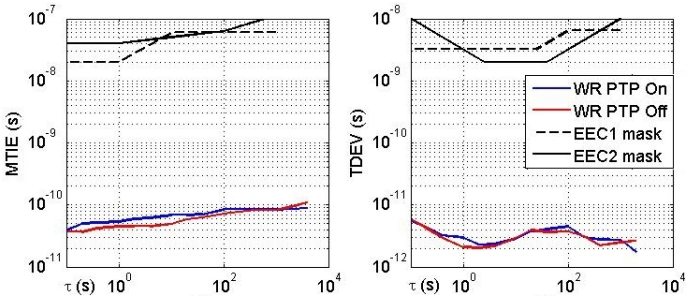

Fig. 5. Wander generation of the WR switch.

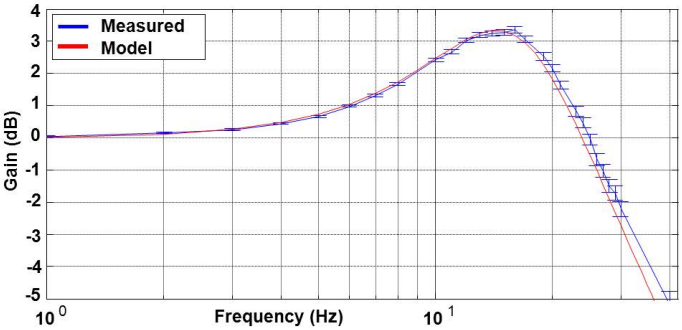

Fig. 6. Transfer function of the WR switch.

V. PHASE NOISE MEASUREMENT OF L1 SYNTONISATION

The phase noise of frequency transfer through a WR network is measured to evaluate the current performance and identify potential improvements. The measurement is done at each state of a linear daisy chain of 3 WR switches in a setup depicted in Fig. 7.

The measurement setup includes a CS4000 Cesium Frequency Standard, Microsemi 3120A High-Performance Phase Noise Test Probe, and 3 WR switches. The CS4000 is the external reference for the GM WR switch and the Microsemi Test Probe. The Test Probe is connected to the output of each of the WR switches. This output provides a 10MHz clock signal derived from the 62.5MHz *local PTP clock signal* using a AD9516 PLL to minimize additional phase noise. This setup is used to measure the phase noise of the of *local PTP clock signal* at the GM WR switch (GM), WR switch 1 (SW1) and WR switch 2 (SW2).

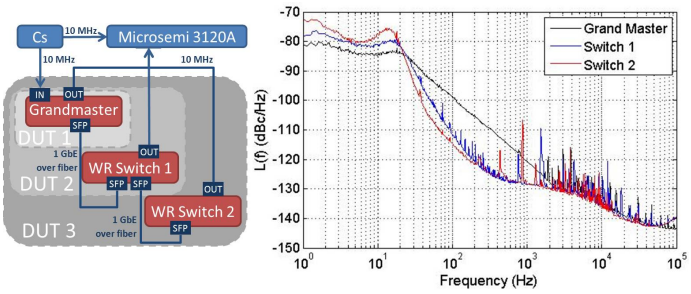

Fig. 7. Phase noise measurement setup and phase noise plot.

Fig. 7 shows the results of these three measurements. The effect of gain peaking is clearly visible in the graph. The increase of phase noise in the 1Hz-10Hz region suggests possible phase noise leaking from the voltage controlled oscillator (VM53S3). Table II provides the integrated RMS jitter in different regions of the spectrum, i.e. 1Hz to 10Hz, 1Hz to 2kHz and 1Hz to 100kHz. The values of jitter in the first region allow to evaluate the jitter in the bandwidth of the SoftPLL with respect to the jitter over the entire measurement bandwidth.

GM lock to ext. ref.	SoftPLL BW	Meas. at	RMS jitter 1Hz-10Hz	1Hz-2kHz	1Hz-100kHz
SoftPLL (current)		GM	4.7ps	9.0ps	9.1ps
	30Hz (current)	SW 1	7.1ps	11.0ps	11.0ps
		SW 2	8.8ps	14.0ps	14.0ps
	200Hz (modified)	SW 1	5.0ps	10.0ps	10.0ps
		SW 2	5.1ps	10.0ps	10.0ps
ext. PLL (modified)		GM	<0.1ps	1.0ps	5.9ps
	30Hz (current)	SW 1	4.4ps	4.8ps	4.9ps
		SW 2	4.8ps	6.0ps	6.1ps
	200 (modified)	SW 1	1.2ps	3.2ps	3.3ps
		SW 2	1.5ps	4.4ps	4.5ps

TABLE II
INTEGRATED RMS JITTER IN DIFFERENT REGIONS OF THE SPECTRUM.

Additionally to frequency-domain analysis, Table III provides time-domain analysis. Allan Deviation is measured by the Microsemi 3120A Test Probe at each of the three switches for different values of integration time with an equivalent noise bandwidth (ENBW) of 50Hz. Both, Allan Deviation and phase noise analysis confirm an accumulation of phase noise in the lower frequencies of the spectrum.

GM lock to external reference	Meas. at	Allan Deviation (ADEV) τ=0.01 s [s]	τ=0.1 s [s]	τ=1 s [s]	τ=10 s [s]	τ=100 s [s]
SoftPLL (current)	GM	9.2e-10	1.3e-10	1.3e-11	1.3e-12	1.3e-13
	Current release of SoftPLL (BW: 30Hz)					
	SW 1	7.4e-10	1.6e-10	1.9e-11	1.9e-12	1.9e-13
	SW 2	6.9e-10	2.1e-10	2.7e-11	2.6e-12	2.6e-13
	Modified SoftPLL (BW: 200Hz)					
	SW 1	1.1e-9	1.4e-10	1.4e-11	1.4e-12	1.4e-13
	SW 2	1.1e-9	1.4e-10	1.4e-11	1.4e-12	1.5e-13
ext. PLL (modified)	GM	1.2e-11	1.7e-12	4.1e-13	7.7e-14	-
	Current release of SoftPLL (BW: 30Hz)					
	SW 1	2.1e-10	6.4e-11	1.3e-11	1.1e-12	-
	SW 2	2.9e-10	8.3e-11	1.3e-11	1.1e-12	-
	Modified SoftPLL (BW: 200Hz)					
	SW 1	1.9e-10	2.2e-11	3.3e-12	3.2e-13	-
	SW 2	3.6e-10	3.8e-11	5.1e-12	5.1e-13	-

TABLE III
ALLAN DEVIATION, EQUIVALENT NOISE BANDWIDTH OF 50HZ.

The source of the phase noise in the lower frequencies of the spectrum is analysed in the next section and methods to improve the frequency transfer over WR network are proposed.

VI. IMPROVEMENT OF THE PHASE NOISE

The measurements presented in the previous section indicate that the phase noise performance of the frequency transfer in the WR network can be improved at 1) the GM syntonizing to the external reference, and 2) the WR switches syntonizing to the GM. These improvements are described in the following subsection.

A. VCO noise leaking

The phase noise spectrum of the WR Switch between 1Hz-10Hz suggest an undesired phase noise accumulation. We use our simulation model to tune SoftPLL such that the phase noise accumulation is minimized, which is shown in the measurement results.

The measurements discussed in section IV and depicted in Fig. 6 agree closely with the transfer function estimated by our model of the SoftPLL. We therefore consider the model valid and use it to predict the impact of leaking of the phase noise from the VCO (VM53S3). The effect of the VCO phase noise, as estimated from the model, is depicted in Fig. 8-a (red line). At frequencies between 1Hz and 5Hz, the modelled VCO's leaking phase noise is comparable with the phase noise of the GM reference (black line). The modelled noise increases to the phase noise floor in the spectrum of the WR Switch 1 phase noise (blue line).

In order to prevent the phase noise leaking, the SoftPLL was modified to provide stronger rejection of the VCO phase noise. The modified SoftPLL has a bandwidth of 200Hz with the VCO rejection characteristics of -58dB@1Hz, -34dB@5Hz and -24db@10Hz (compared with the characteristics of the current SoftPLL: -48dB@1Hz, -20dB@5Hz and -7db@10Hz). The phase noise of the WR Switch 1 with modified SoftPLL is depicted in Fig. 8-a (green line). The phase noise of WR Switch 1 with modified SoftPLL does not exhibit any phase noise accumulation in the 1Hz-10Hz range. This improves the Allan Deviation. The new ADEV measurements are provided in Table III. Theoretically, the larger bandwidth of the modified SoftPLL could lead to increased jitter due to a less aggressive filtering of the phase noise above 30Hz. However, measurement with a cascade of two WR Switches running the modified SoftPLL and connected to the GM show a decrease of jitter compared to the non-modified SoftPLL (current), as presented in Table II.

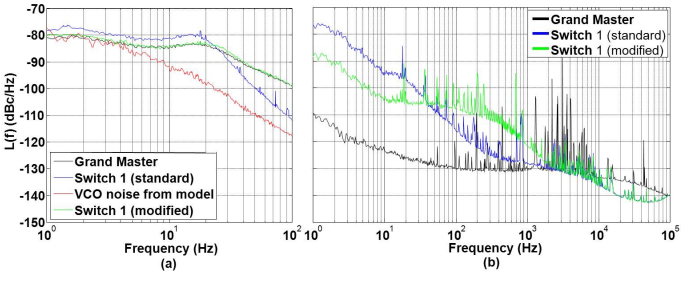

Fig. 8. Phase noise plots: (a) with modified SoftPLL; (b) with modified GM.

B. Syntonisation of GM to the external 10MHz reference

The GM WR Switch locks to the external 10MHz reference using a SoftPLL that requires 62.5MHz input signal, as depicted in Fig. 3. It is the clock conversion that introduces undesired phase noise and needs to be optimized to improve performance of the frequency transfer in the WR network.

The SoftPLL operates with a 62.5MHz frequency because its usual input clock signal is a divided 125MHz clock from a Gigabit Transceiver. In the GM, however, the SoftPLL is used to lock to the external 10MHz reference. To provide the SoftPLL with the required clock signal, the 10MHz input is multiplied using an internal PLL of the FPGA, called a Mixed Mode Clock Manager (MMCM). The phase noise spectrum of the 62.5MHz output of the MMCM is measured with an Agilent E5052B Phase Noise Analyser. It shows a large phase noise power located between 10kHz and 2MHz, with an RMS jitter (integrated from 10kHz to 2MHz) of 155ps. This high frequency phase noise would be filtered out by an analogue PLL. However, the SoftPLL cannot filter such noise because of its digital nature (Nyquist bandwidth of 1.9kHz) and the digital nature of the DDMTD. The analogue equivalent of the DDMTD has a low pass filter after the mixer to limit the bandwidth of the measurement system. The DDMTD has no analogue components. Instead the input phase noise is filtered by the bandwidth of the D-type flip-flops (hundreds of MHz) and by the filtering action of the deglitching algorithm [4]. Consequently, the remaining phase noise is aliased over the Nyquist bandwidth of the SoftPLL acting as white noise. This effect can be seen in Fig. 7. The Allan Deviation slope of the GM, see Table III, is the inverse of the integration time which indicates White PM or Flicker PM noise. This is a confirmation of the aliased white noise in the SoftPLL bandwidth.

In order to verify that the MMCM is indeed the problem and to measure the improvement of using an external PLL, the available on board AD9516 PLL is used to directly multiply the external 10MHz input reference. The SoftPLL is bypassed and the multiplied clock signal is used as the *local PTP clock signal*. This solution allows to verify our suspicions with minimal changes to the hardware.

Fig. 8-b provides the results of phase noise measured at the modified GM and at WR Switch 1 using either the current release of SoftPLL (BW: 30Hz) or its modified version (BW: 200Hz). The results confirm that the modified GM has a better phase noise profile in the lower frequencies. This is also confirmed by the values of RMS jitter in Table II. The downside of the GM modification are the spurs that are especially visible above 1kHz. They degrade the integrated jitter over the entire bandwidth to 5.9ps RMS. The spurs that occur on the GM are partially filtered by the SoftPLL noise transfer function of WR Switch 1. The origin of the spurs is likely to be related to suboptimal power supply decoupling and noise coupling to the controlled oscillators.

Table III provides measurement values of the Allan Deviation with the modified GM and with two WR Switches using either the current or the modified SoftPLL. The modification of the GM does not improve significantly the Allan Deviation results for the current SoftPLL, as expected. However, the modified SoftPLL benefits from the lower phase noise of the GM with a significant reduction of ADEV.

The drawback of the modified SoftPLL is a worse phase noise in the 100Hz-1kHz region due to phase noise accumulation. A controlled oscillator with at least 10dB better phase noise in the 1-10Hz spectrum would help to pick the best noise profile since less bandwidth would be required to reject the noise coming from the local oscillator.

The modification of the GM does not allow the GM to phase-align its *local PTP clock signal* with the PPS input. Therefore, it has limited usefulness. However, the ongoing design of new WR switch hardware will likely include a dedicated external PLL to lock to the external reference.

VII. SyncE-compliance

Compliance of L1 syntonisation in WR with SyncE specifications (ITU-T G.8262) might be helpful in a future adaptation of the WR-based features integrated into the new revision of PTP. This section proves that WR can be made SyncE-compliant only through software changes in the SoftPLL.

The SoftPLL fails to pass the wander tolerance tests by unlocking when the modulating frequency reaches 1Hz and the time amplitude reaches 250ns. Increasing the de-locking threshold is not a solution since the wander frequency might drive the local oscillator (VM53S3) out of its range.

In order to meet the SyncE characteristics without changing the hardware, the SoftPLL software was modified as follows:

1) The bandwidth was reduced to 5Hz.
2) Multiple Unit Interval tracking of error was implemented.
3) Locking logic was modified to allow tracking of large zero-mean wander.

The three tests from subsection IV were repeated using the modified SoftPLL: wander tolerance (op-1 only), jitter tolerance, and wander transfer. All tests were successfully passed. The transfer function of the modified SoftPLL is depicted in Fig. 9-1. It meets the noise transfer (sec. 10

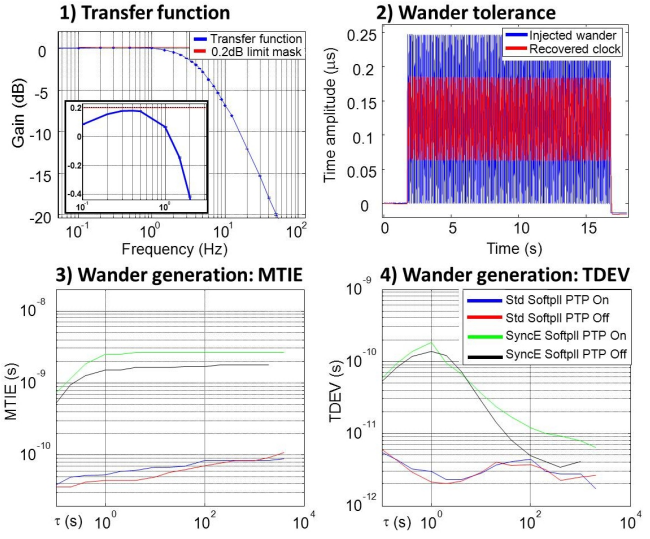

Fig. 9. Characteristics of WR switch with SyncE-compliant SoftPLL.

of G.8262) requirement of below 0.2dB gain peaking and confirms that the bandwidth is within the range specified for the EEC-option-1. Fig. 9-2 shows that the WR switch with the modified SoftPLL correctly transferred the wander noise defined in section 9.1.1 of G.8262. The TIE was measured with respect to a CS4000 and its plot shows the injected wander reference signal (blue) and the signal recovered by the DUT (red). The DUT tracked the reference signal without losing the lock and quickly followed the reference signal phase when the wander noise injection finished (at second 17). The switch was able to lock even after the injection of wander was initiated. The MTIE and TDEV of the switch running the modified SoftPLL are depicted in Fig. 9-3 and Fig. 9-4 for the case when the WR PTP is enabled and when it is disabled. In the latter case, the switch is only syntonised.

In order to compare the frequency transfer using the currently available SoftPLL and the SoftPLL modified to be SyncE compliant, phase noise was measured at WR Switch 1 using the setup depicted in Fig. 7 (Microsemi 3120A) where the GM is modified as explained in VI-B. Fig. 10 shows that

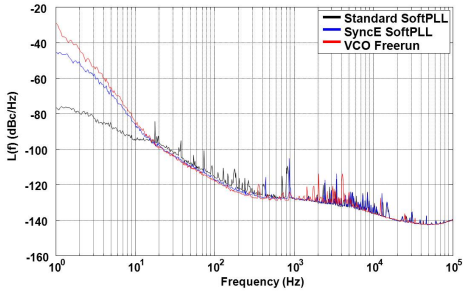

Fig. 10. Phase noise plot.

the unmodified SoftPLL (black) has a very low integrated jitter of 5ps RMS (from 1Hz to 100kHz). The SyncE-compliant SoftPLL (blue) has a much higher jitter in the 1-10Hz bandwidth that results in a total integrated jitter of 100ps RMS. This is attributed to the the VCO (VM53S3) that exhibits high phase noise in the 1-10Hz region when not controlled (red trace). An oscillator with a better phase noise profile in that region (e.g. -70 dBc/Hz at 1Hz) can lower the disparity of performance between the two versions of the SoftPLL.

VIII. PTP Synchronisation with modified L1 syntonization

The performance of the WR PTP synchronisation with the described modifications to the L1 syntonisation is evaluated in a cascade of 10 WR switches, as depicted in Fig. 11. The time error between the 10MHz output of the first switch, the Grandmaster, and that of the other switches is measured over 30min using a 10GS/s oscilloscope. For each version of the Grandmaster, i.e. the release and the modified version, three measurements are performed, one for each SoftPLL versions:
M1 – the SoftPLL in the currently available WR switches;
M2 – the SoftPLL with 200Hz bandwidth described in VI-A;
M3 – the SoftPLL compliant with SyncE described in VII.
The same version of SoftPLL is used in all the 9 switches.

The modifications affect mostly jitter that is calculated as standard deviation of the measured time error.The modification

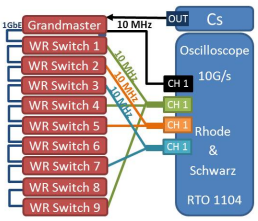

WR SW	Standard deviation [ps]						Mean offset [ps]					
	Release GM			Modified GM			Release GM			Modified GM		
N°	M1	M2	M3	M1	M2	M3	M1	M2	M3	M1	M2	M3
1	9.9	6.3	241	8.2	7.7	245	-18	<1	132	-3	11	60
2	14.8	8.1	310	10.5	8.3	345	53	46	391	60	174	262
3	19.7	9.2	411	12.1	8.7	470	16	16	263	39	46	-9
4	25.2	9.8	600	15.5	9.2	680	-17	2	205	21	19	113
5	31.0	11.2	NL	18.6	10.1	NL	-63	-48	NL	-37	-23	NL
7	53.7	13.2	NL	28.4	10.8	NL	46	40	NL	81	98	NL
9	98.0	14.9	NL	37.6	11.6	NL	-126	140	NL	-89	-86	NL

NL: SoftPLL unlocked during test, unable to keep the synchronisation.

Fig. 11. Characteristics of WR switch with SyncE-compliant SoftPLL.

of the Grandmaster alone improves jitter as it minimizes the initial noise that is amplified by the cascaded SoftPLLs. Similarly, increased bandwidth of the SoftPLL (M2) alone improves jitter as it reduces the amplification of the initial noise. The best results are achieved when both modifications are applied, in such case the precision after 9 hops is 11.6ps and the accuracy is below 100ps.

On the other hand, the SyncE-compliant version of the SoftPLL shows increased jitter and is unable to keep synchronisation after 4 switches. This is due to the poor phase noise performance below 10Hz of the VM53S3 oscillator, a better oscillator would help to keep the synchronisation. Still, the measurement shows that the WR PTP synchronisation with the SyncE-compliant SoftPLL allows sub-ns synchronisation in a chain of 2 switches.

IX. Conclusions

This article provides comprehensive information about frequency transfer and its characteristics in a WR network. This characterisation is done using ITU-T G.8262 guidelines and phase noise analysis to allow easy comparison with currently available frequency transfer techniques in a wide range of applications. It is shown that further adaptation for different applications is achievable, sometimes only by software modification. Indeed, the compliance with timing characteristics of a synchronous Ethernet equipment slave clock was achieved this way. On the other hand, hardware modification is desired to improve the phase noise profile of WR.

The increasing number of WR applications pushes the limits of its performance while the standardisation makes it likely to be used in less stringent applications that require compatibility with legacy equipment. We have shown that adaptation in different directions is feasible, straightforward and often requires only software modifications.

References

[1] E. Cota, M. Lipiński, T. Włostowski, E. Bij, and J. Serrano, "White Rabbit Specification: Draft for Comments," www.ohwr.org/documents/21.
[2] IEEE P1588 Working Group. ieee-sa.centraldesktop.com/1588public.
[3] O. Ronen and M. Lipinski, "Enhanced synchronization accuracy in ieee1588," in *2015 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS)*.
[4] T. Włostowski, "Precise time and frequency transfer in a White Rabbit network," Master's thesis, Warsaw University of Technology, may 2011.
[5] P. Moreira, P. Alvarez, J. Serrano, I. Darwezeh, and T. Wlostowski, "Digital Dual Mixer Time Difference for Sub-Nanosecond Time Synchronization in Ethernet," *FCS 2010 IEEE International*, 2010.
[6] LatticeMico32 Open. http://www.ohwr.org/projects/lm32.
[7] WR Test Reports. www.ohwr.org/projects/wr-switch-testing/wiki.
[8] M.Lipinski. SyncE Characteristics of a White Rabbit Switch. http://www.ohwr.org/attachments/3833/WR-SyncE-characteristics-report.v3-2015-02-27.pdf.