

Rathera

Sub-nanosecond Network Synchronisation An Introduction to White Rabbit

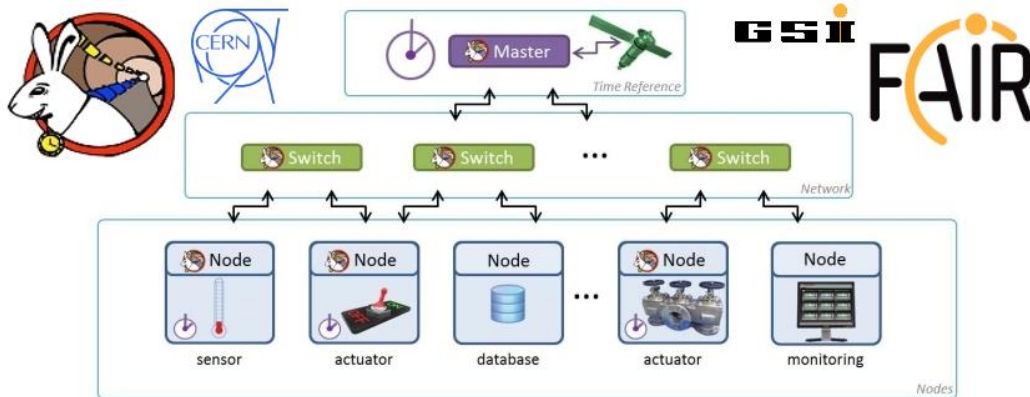
Prepared by Alex Scarbro

Issue: 2

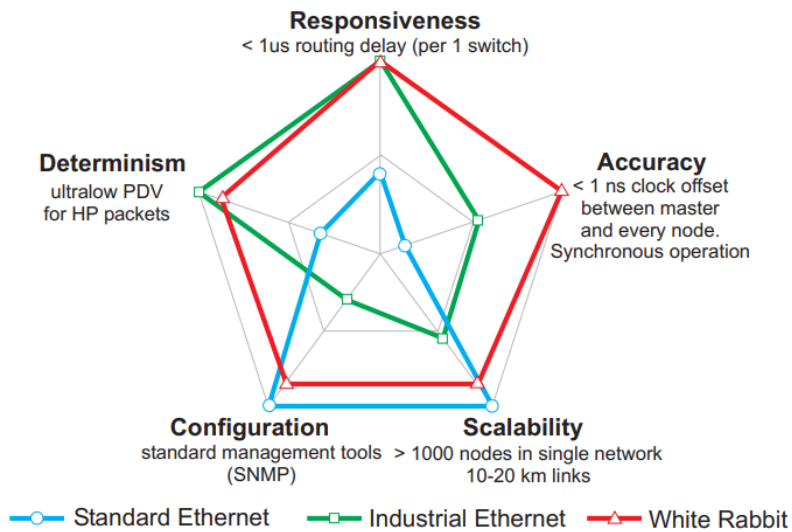
Issue	Date	Changes	Signed by
DRAFT	1/4/22	Initial draft	-
1	19/4/22	First release	A. Scarbro
2	20/4/22	Updated PTP standardisation. Replaced commercial footer from Rathera template. Amended applications.	A. Scarbro

Introduction

White Rabbit is a network timing technology originally developed in partnership between CERN (Geneva) and GSI (Darmstadt) to provide their Large Experimental Physics programmes (LHC & FAIR) with deterministic data transfer and state of the art synchronisation accuracy.



White Rabbit enables a large-scale network based on Ethernet (IEEE802.3), that supports more than a thousand “nodes” spread over many tens or even hundreds of kilometres. All of the nodes are synchronised together with sub-nanosecond accuracy. For optimal performance, White Rabbit utilises a dedicated fiber optic network, however it is fully interoperable with traditional Ethernet network traffic.

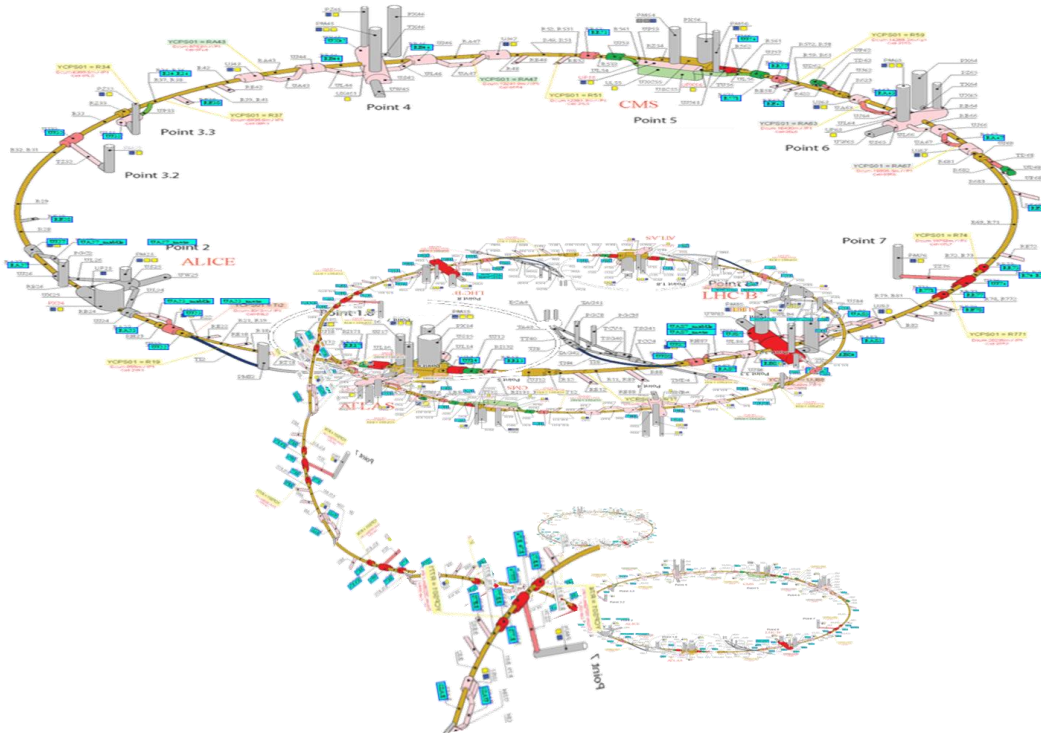


Rather aptly, “White Rabbit” is a reference to the character in Alice’s Adventures in Wonderland.




History

White Rabbit was conceived in 2006 as collaboration between the European Organization for Nuclear Research (CERN) and GSI Helmholtz Centre for Heavy Ion Research. For CERN, it aimed to ultimately replace the existing General Machine Timing (GMT) system which is responsible for synchronising all the beam control and measurement events that take place around the 27km long Large Hadron Collider (LHC). Arguably the largest and most complex machine mankind has ever built.



For GSI, White Rabbit was to be used as the primary timing system for the planned International Accelerator Facility (FAIR).



Historically, large experimental physics programmes would develop their own bespoke timing solution since commercially available equipment lacked the required performance and flexibility.

This would inevitably require custom hardware and dedicated cabling that was generally separate from the control and data networks.

With this in mind, the White Rabbit team set out to break this cycle of reinvention and so they developed a scalable and modular platform that required very little application-specific customisation. The solution was designed to utilise the many thousands of kilometres of single mode fiber optic cable already installed around the LHC.

The key requirements for White Rabbit were:

- **Determinism** – It must limit the worst-case message delivery time to a fixed upper bound.
- **Accuracy** – It must achieve sub-nanosecond synchronisation of all the slave nodes with respect to the network “grand master”.
- **Reliability** – It must guarantee the robustness of delivery of timing and data across the whole of the network.
- **Openness** – It must be completely open (in terms of hardware, software and gateway) and not tied to any particular manufacturer (thus avoiding vendor lock-in).


CERN Open Hardware License

The White Rabbit PTP Core technology (WRPC) and many derivative products, such as Switches and Nodes, are released under mix of licenses including CERN’s Open Hardware License (OHL). The OHL was developed as a legal tool to promote collaboration among hardware designers and support the freedom to use, study, modify, share and distribute hardware designs, and products based on those designs.

CERN is the custodian of the OHL and releases new versions and variants from time to time. The current release is version 2 and it comes in three variants:

- CERN-OHL-S (strongly reciprocal, [txt](#), [pdf](#))
- CERN-OHL-W (weakly reciprocal, [txt](#), [pdf](#))
- CERN-OHL-P (permissive, [txt](#), [pdf](#))

The [CERN OHL v2 home](#) can be visited to access more information about the variants, including a rationale document explaining drafting decisions and a list of frequently asked questions.



The [Open Hardware Repository](#) hosts all elements of the White Rabbit technology and many of its derivative products (three of which are pictured above). It provides a place for electronics designers across the world to collaborate on open hardware designs, with a similar philosophy to the free software movement.

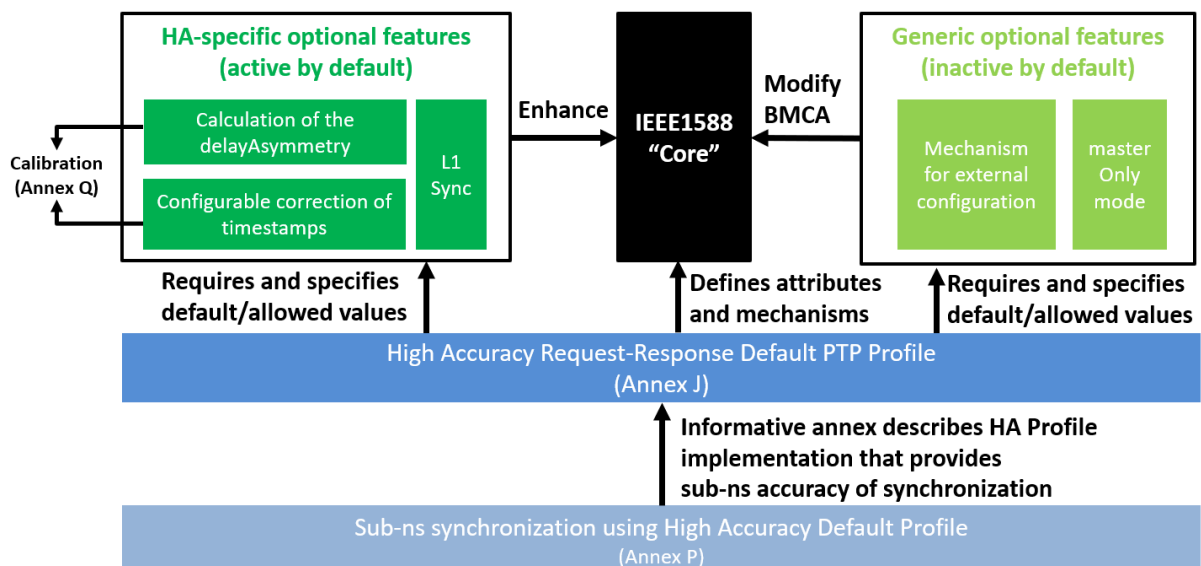
Vendors

As a result of being released under an Open Hardware License, standard White Rabbit hardware is available from many vendors around Europe and further afield. Where required, this Open Hardware is complemented with customised Open or Closed License solutions.

Vendor	Country	Hardware (Open & Closed Licensed)
Seven Solutions	Granada, Spain	White Rabbit Switch & Low Jitter Variant Simple PCIe FMC carrier – SPEC WR-LEN WR-ZEN
Creotech	Piaseczno, Poland	White Rabbit Switch Simple PCIe FMC carrier - SPEC
INCAA Computers	Apeldoorn, Netherlands	CompactRIO White Rabbit - CRIO-WR Simple PCIe FMC carrier - SPEC Simple PXI express FMC Carrier - SPEXI
Janz Tec	Paderborn, Germany	Simple VME FMC Carrier - SVEC
OPNT	Amsterdam, Netherlands	White Rabbit Switch
SyncTechnology	Beijing, China	FMC WR mezzanine - Cute-WR-DP Mini-WR White Rabbit Switch
Rathera	Shipleigh, United Kingdom	Custom integrated WR timing solutions

IEEE Standardisation

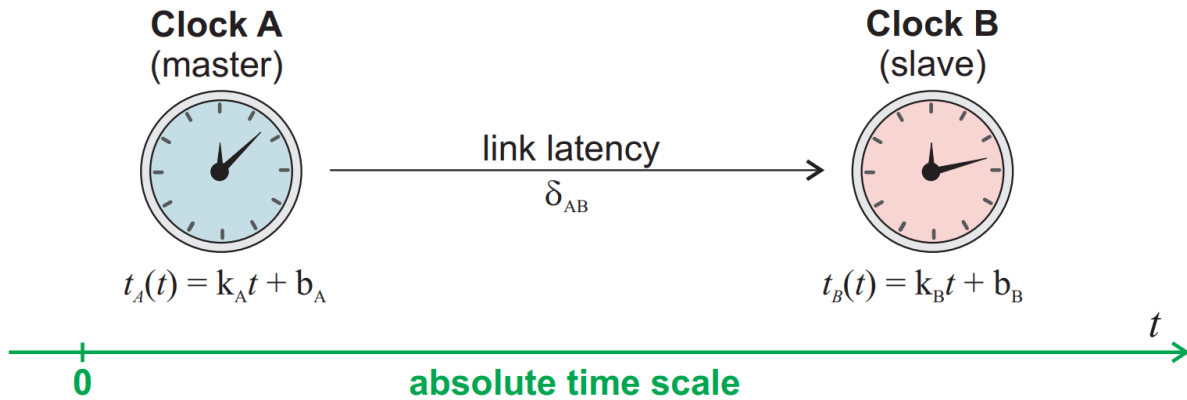
Thanks to the incredibly dedicated work of Maciej Lipinski and the wider team at CERN, their White Rabbit technology has been formally adopted in to the IEEE1588-2019 Precision Time Protocol (PTP) High Accuracy (HA) profile.



Their efforts will hopefully enable its wider adoption and allow for interoperability of equipment from many different vendors.

Challenge of Accurate Time Synchronisation

The principle of synchronisation of the slave to the master clock is described below. Here the two clocks are physically separated and so there exists some minimum latency in the communication between them (δ_{AB}).



The two clocks are ticking at different rates (denoted by k_A & k_B) and with different offsets (denoted by b_A & b_B) with respect to the absolute time scale.

To synchronise these clocks, two steps are required:

1. The slave clock must be made to tick at the same rate as the master clock ($k_B = k_A$). This process is known as Syntonisation and effectively describes two oscillators being made to have equal frequency, but not necessarily equal phase.
2. The slave clock offset (b_B) must be made to match the master offset (b_A), known as Offset Adjustment.

Having syntonised the two clocks, the instantaneous time of the master clock (t_A) can be transferred to the slave, with this communication taking time δ_{AB} . The clock offset (Δ) can then be calculated and used to advance the slave clock by the necessary amount.

$$\Delta = t_A - t_B + \delta_{AB}$$

At first inspection, the problem of time transfer might appear straightforward. However, when we try to achieve accuracies measured in tens of picoseconds, we must account for phenomena such as varying link latency, link asymmetry, limitations of the physical media interfaces and many different sources of noise and jitter.

Foundation Standards

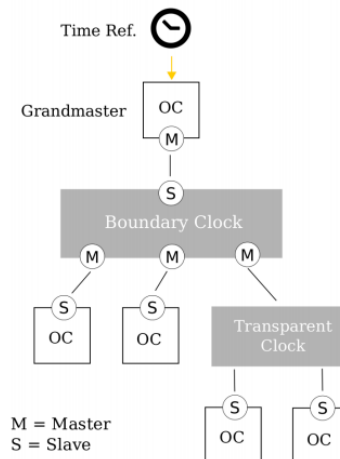
White Rabbit is built upon a foundation of two established technologies, namely the Precision Time Protocol (IEEE1588 PTP v2) and Synchronous Ethernet.

IEEE1588-2008 Precision Time Protocol

The IEEE1588 Precision Time Protocol (PTP) is used to distribute a common notion of time across a standard computer network. With the help of supporting hardware, it can achieve clock accuracies in the sub-microsecond range making it particularly relevant for real-time measurement and control applications.

PTP sits in the niche between the Network Time Protocol (a purely software-based application layer solution) and a GPS timing receiver. The former achieving little better than millisecond accuracy and the latter as low as a few nanoseconds. However, A GPS-based solution is dependent upon external infrastructure (i.e., an antenna & coaxial cable) and requires an unobstructed view of the sky to maintain performance.

PTP has a hierarchical structure with a system time reference feeding a “grand master” at the top of the tree. Nodes with only a single master or slave port (denoted by M or S) are called Ordinary Clocks (OC).



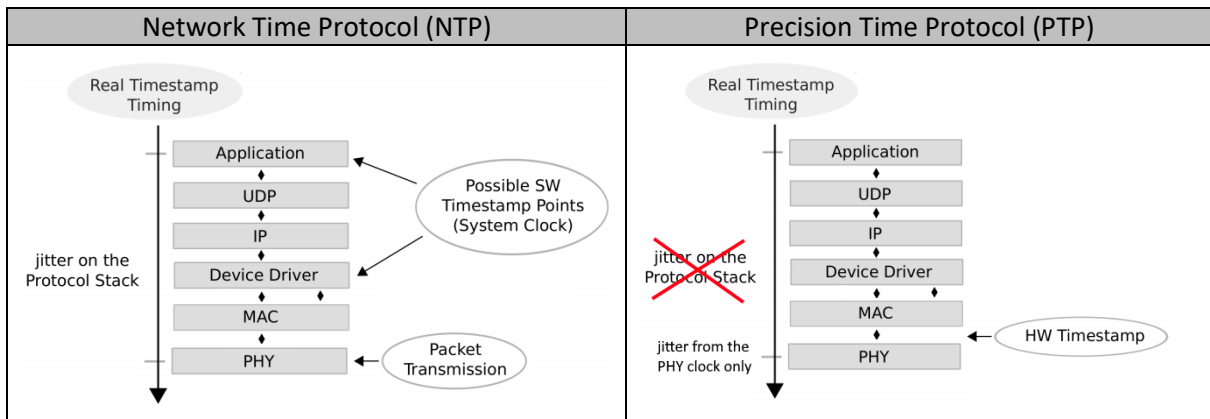
A PTP network of nodes can either be connected point to point or with Transparent Clocks (TC) and/or Boundary Clocks (BC) in between.

A Transparent Clock measures the time taken for a PTP timestamp packet to traverse its network ports and adds this delay in to the packet. Therefore, its effect on timing is effectively transparent to the nodes connected to it.

A Boundary Clock uses the Best Master Clock algorithm to select which one of the master clocks connected to its slave port(s) it must synchronise to. The remaining ports on the BC are therefore configured as master ports for the OC (slave) nodes below. This allows other master clocks to assume the “grand master” role in the event of a failure.

PTP nodes can be interconnected using standard Ethernet switches, but the indeterminate “Store & Forward” delay between switch ports will degrade timing performance compared to that which can be achieved with PTP-aware equipment (e.g., BC or TC switches).


PTP is an evolution of the NTP concept that instead uses a Delay Request-Response algorithm to estimate the total round-trip delay. It can leverage dedicated hardware to perform timestamping of the PTP packets at the point of transmission and reception on the network medium.



The hardware timestamping support in PTP eliminates the timing jitter present in the software protocol stack of the NTP implementation. However, accuracy is still degraded by any asymmetry in the master-slave network links.

Early implementations of PTP required specialised ASICs to perform time stamping, nowadays this functionality is commonplace in consumer-grade copper and fiber optic Physical Interfaces (PHYs).

The Delay Request-Response algorithm is shown below:

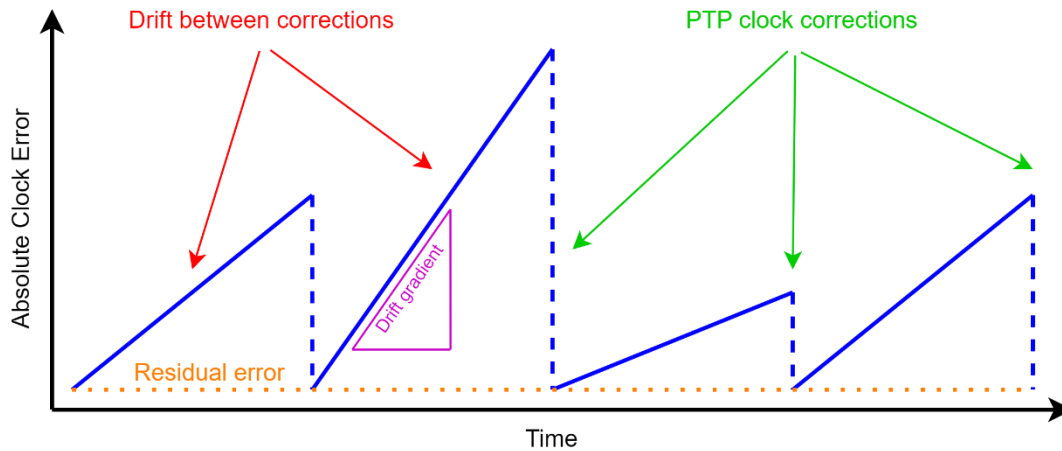


The message exchange begins with the master sending a Sync message to the slave noting its transmission time (t_1). The slave receives the message and notes its arrival time (t_2). The master then sends a Follow_Up message containing the t_1 timestamp value. Next the slave sends the Delay_Req request to the master noting its transmission time (t_3). Again, the master receives the request and notes its arrival time (t_4), sending this back to the slave in the Delay_Resp response message. The round-trip time (δ) can then be estimated:

$$\delta = (t_4 - t_1) - (t_3 - t_2)$$

Note that there is no means of determining the exact one-way delay with this algorithm (master to slave or slave to master). PTP assumes the delay is equal in each direction.

This PTP delay measurement process must be run periodically to maintain accuracy as the master and slave clock frequencies are not synchronised with each other. Therefore, it is necessary to make a tradeoff between instantaneous timing accuracy and the consumption of network data bandwidth.



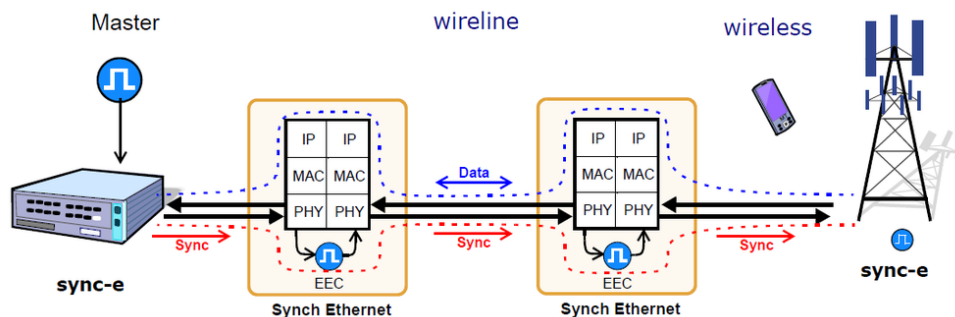
The diagram above demonstrates the exaggerated slave clock drift between corrections. The rate of drift is proportional to the difference in master & slave clock frequency. A residual absolute error remains due to any uncompensated link asymmetry.

PTP uses an epoch of 1st January 1970 00:00:00 and corrects the slave clock with respect to International Atomic Time (TAI). TAI is a monotonic timescale and so does not require a mechanism to deal with leap seconds. UTC can still be derived from TAI through a simple calculation using an offset value transmitted in the PTP data structures.


PTP can provide a One Pulse Per Second (1PPS) clock output and, by disciplining a local oscillator, a fixed reference frequency output (10MHz being most common). This allows PTP to be used instead of GPS & GNSS timing receivers in locations where visibility of the satellite constellation is intermittent, impossible or intentionally contested (jammed).

Synchronous Ethernet

Abbreviated to SyncE, this telecommunications technology allows the distribution of a common frequency standard throughout a network. SyncE utilises a dedicated Phase Locked Loop (PLL) in the receiving PHY to extract a clock that is synchronous to the data stream.



The recovered clock accuracy is not affected by the volume of network traffic. However, the jitter performance is highly dependent upon the clock recovery process and the number of cascaded devices.



As shown in the above Bode Plot, the clock recovery PLL can have “peaking” in its transfer function just before roll-off occurs. When many devices are cascaded (all with similar transfer functions), this peaking can be multiplied having a degrading effect on the loop phase and gain margin. Ultimately this results in the potential for system instability which is highly undesirable. It is therefore necessary, in cascaded applications, to adjust the PLL parameters to be slightly over-damped.

How does White Rabbit work?

As previously outlined, IEEE1588 PTP v2 can achieve sub-microsecond accuracies (under good conditions). However, to achieve sub-nanosecond performance, White Rabbit must address a number of shortcomings with the standard PTP (v2) implementation.

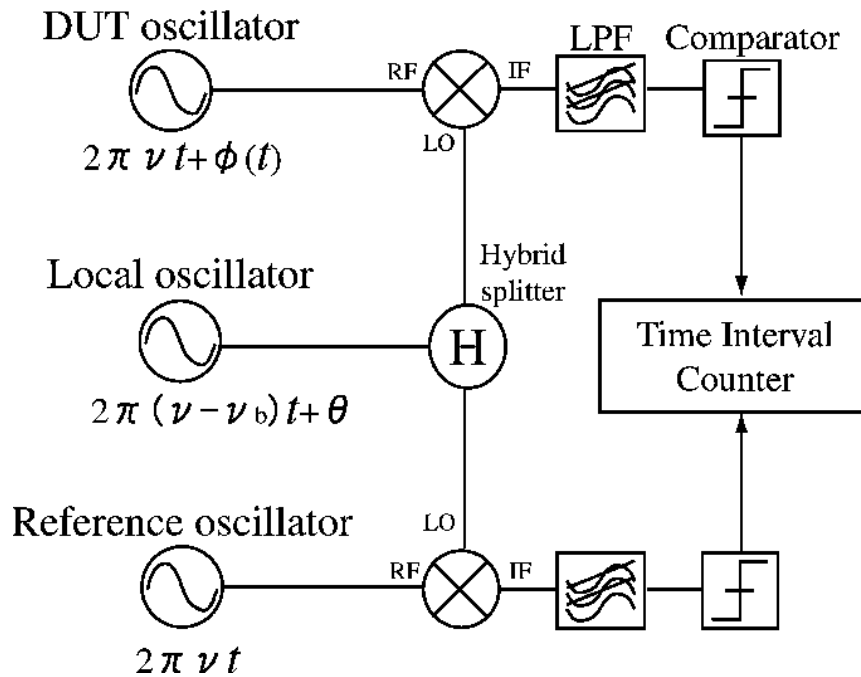
1. PTP makes no attempt to syntonise the frequency of the slave clock to the master clock.
2. PTP requires a relatively high message exchange rate to maintain good accuracy.
3. PTP has limited measurement precision and resolution of timestamps.
4. PTP does not autonomously measure and subsequently correct for link asymmetry.

To address points 1 and 2, White Rabbit uses Synchronous Ethernet (SyncE) to distribute a common clock frequency to all nodes in the network. This removes the problem of the slave clock drifting (relative to the master) between PTP clock corrections. As a result, the timing accuracy can be significantly improved without increasing the PTP message exchange rate. In fact, the rate can usually be decreased making more network bandwidth available for data & control.

To address points 3 and 4, White Rabbit increases the timestamp precision considerably by casting the fine delay measurement as a phase detection problem (using the Digital Dual Mixer Time Difference technique). These measurements are utilised both for the initial fixed delay & physical link asymmetry calibration and thereafter during normal PTP operation.


Digital Dual Mixer Time Difference

The Digital Dual Mixer Time Difference (DDMTD) phase detector is a corner stone for the high performance and yet low-cost implementation of White Rabbit. The DDMTD is a digital equivalent of the established analogue method (shown below) that allows the phase difference of high frequency signals to be measured with picosecond resolution whilst only processing at a fraction of the input signal frequency.




Here the DUT and Reference Oscillators are at the same frequency (ν), but have a difference in phase between them of $\phi(t)$. Both oscillators are mixed with a common “local oscillator” (LO) that has a small offset in frequency (known as the beat frequency, ν_b). The IF output spectra contain the (high) sum and (low) difference frequency mixing products. The IF output signals are passed through low pass image-rejection filters, leaving only the difference product at the beat frequency. Whilst the DUT & Reference oscillator frequencies are translated down to a low beat frequency, their phase relationship remains unchanged. Comparators and a simple digital counter are then used to measure the time interval between rising edges of the beat frequencies.

To use a mechanical analogy, the DMTD operates in much the same way as a set of callipers where our offset local oscillator acts as our vernier scale to substantially improve our resolution without overcomplicating the measurement hardware.



Whilst the analogue approach yields excellent results, a White Rabbit implementation requires several phase detectors, some of which may have many inputs (for example in the White Rabbit 18-port Network Switch). A digital implementation of the DMTD is therefore be desirable.



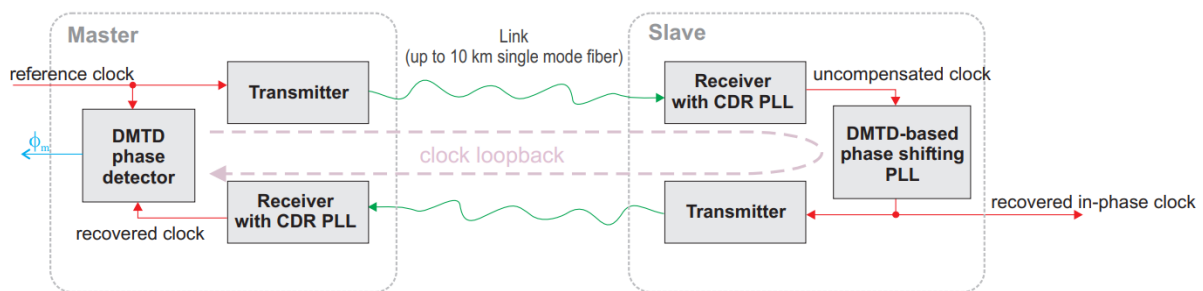
The diagram above shows a simplistic Digital-DMTD structure. Here the “local oscillator” (denoted at clk_{DDMTD}) is generated using a fractional PLL derived from one of the two input signals and the mixers are replaced with D-type flip-flops. The clk_{Aout} & clk_{Bout} signals are de-glitched and fed to an interval counter as previously described (not shown here).

For example, the White Rabbit Switch operates with an N value of 14 and so with an input clock signal of 62.5MHz (typical for a 16-bit wide 1.25Gbps PHY with 8b/10b encoding), the clk_{DDMTD} frequency is 62.496186MHz and the beat frequency (clk_{Aout} & clk_{Bout}) is 3.814kHz.

$$X_{in}[ns] = \frac{1}{1 + 2^N} \cdot X_{out}[ns]$$

The difference in time between the two input signals (X_{in}) is magnified by the ratio of the input frequency to the beat frequency (a factor of 16,385, with $N=14$) to generate a difference in time between the two output signals (X_{out}). With the interval counter clocked at 62.5MHz, the resolution of X_{out} is 16ns and thus the measurement resolution of X_{in} is 0.977ps.

This all-digital approach scales much more efficiently where many DMTD measurement channels are required and could be entirely implemented within the logic of a Field Programmable Gate Array (FPGA). However, the PLL blocks present in FPGAs have notoriously poor jitter performance and so the clk_{DDMTD} “Helper” synthesizer is generally realised partially as a PLL onboard the FPGA with an external voltage-controlled crystal oscillator (VCXO). Thankfully, the same clk_{DDMTD} signal can be shared across all DDMTD instantiations.




The DDMTD module is used for several functions in a White Rabbit master-slave link:

- In the master to measure the time difference between its reference clock and the recovered clock looped-back by the slave.
- In the slave (in conjunction with a Software-PLL) to measure the time difference between the recovered clock and its local PTP clock along with performing phase shifting.
- In both the master and slave to perform measurement of the fixed transmit and receive latencies.

White Rabbit Link Delay Model

The PTP algorithm along with the significantly enhanced timestamp precision (provided by the DDMDT technique) delivers a very accurate measurement of round-trip delay from master to slave and slave back to master. To be able to absolutely align the slave clock with the master we must determine the link delay & link asymmetry. This is achieved by understanding and accounting for all of the fixed & variable delays that contribute to the overall Link Delay Model shown below.



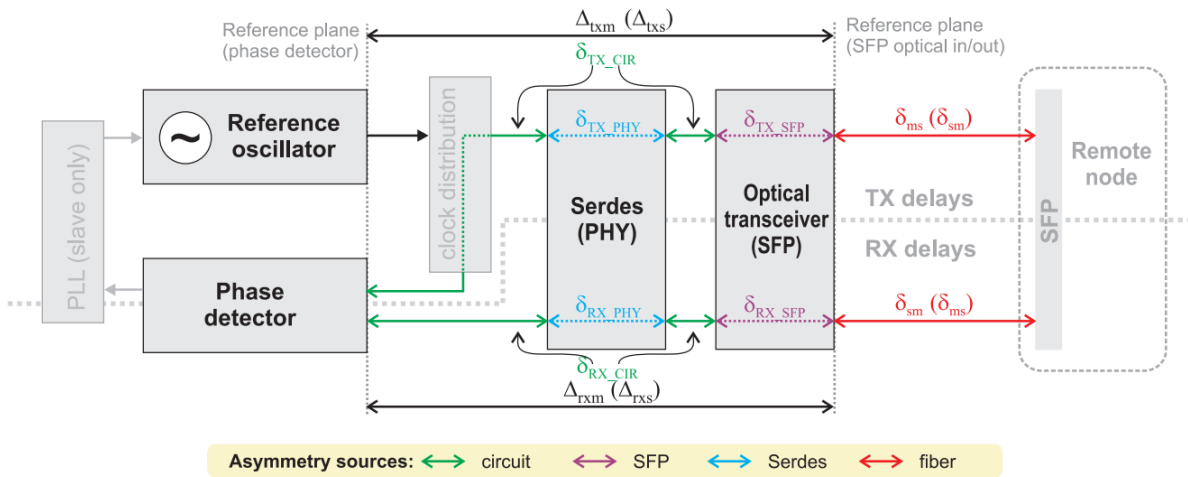
The White Rabbit master and slave each have a number of fixed delays (Δ_{TXM} , Δ_{RXM} , Δ_{TXS} , Δ_{RXS}) in their transmission and reception paths. These include delays through the high-speed transceivers, the routing from the high-speed transceivers to the package pins, the PCB traces between the package pins and the SFP cage and delays internal to the SFP optical transceivers.

There are delays internal to the high-speed receivers that vary between connection cycles as the deserializer synchronises to the inter-symbol data boundaries - the so called “bitslide” delays (ϵ_M and ϵ_S). These values are made available to the WR algorithm by the PHY logic and so are automatically accounted for in the correction calculations.



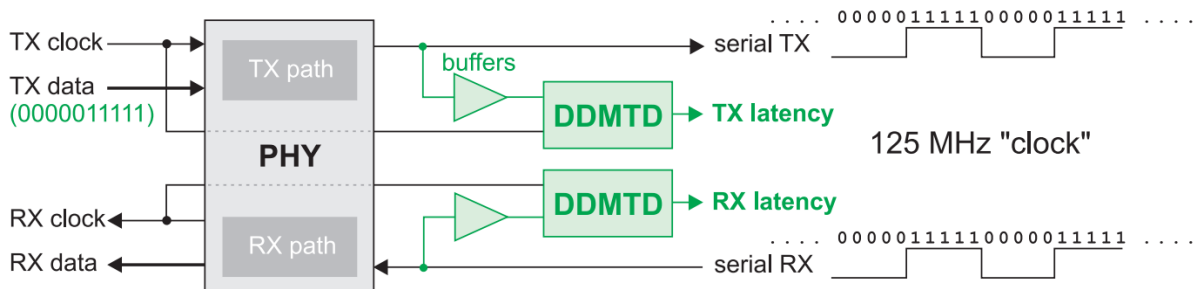
There are variable delays of transmission and reception through the medium (δ_{MS} and δ_{SM}). These delays primarily change due to environmental factors such as ambient temperature. In the case of bidirectional fiber optic (where a single fiber is used with a different wavelength travelling in each direction) there will be a difference in the propagation velocities in each direction resulting in link asymmetry.

A detailed breakdown of the sources of delay are shown in the below diagram.



The fixed delays caused by routing within the FPGA fabric, package & PCB traces ($\delta_{TX_CIR}/\delta_{RX_CIR}$) and the SFP optical transceivers ($\delta_{TX_SFP}/\delta_{RX_SFP}$) are determined during the formal module calibration procedure and are constant as long as the FPGA “gateway” (bitstream) is not changed and the ambient temperature is maintained.

The TX & RX latencies within the serializer/deserializer ($\delta_{TX_PHY}/\delta_{RX_PHY}$) are measured during the White Rabbit Link Setup “CALIBRATION” step upon physical connection to a White Rabbit network.



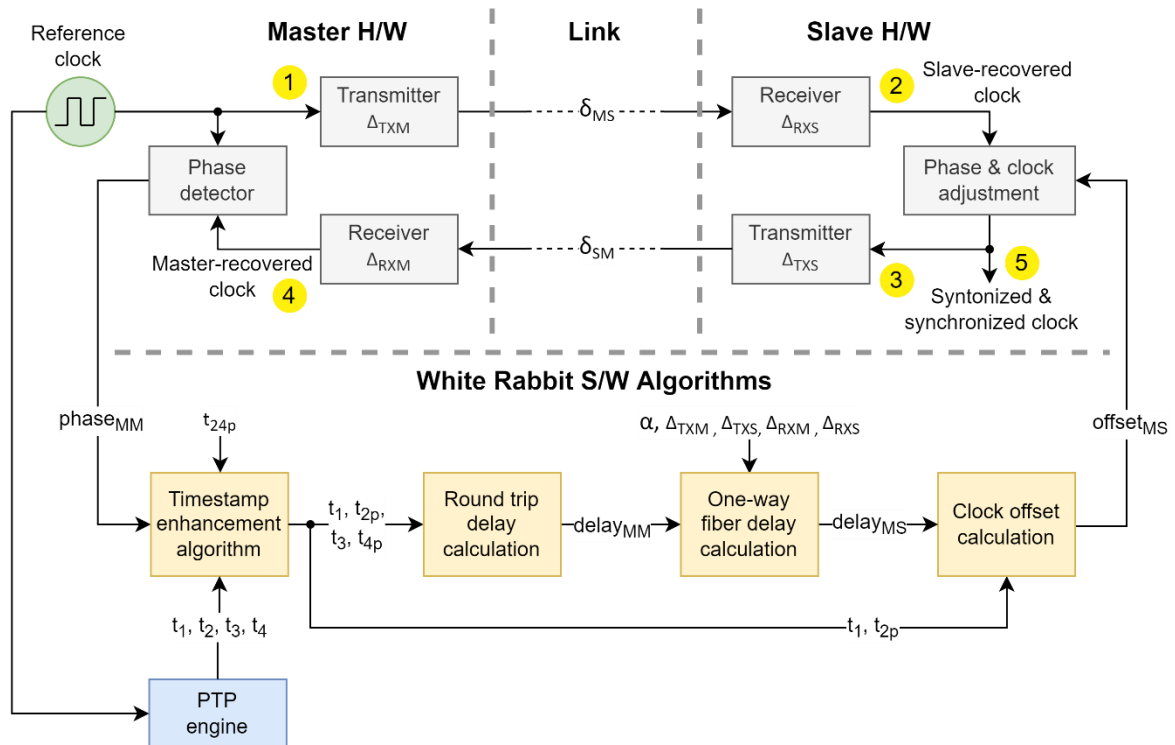
The transmit latency is determined by sending a sequence of special 10-bit “K28.5” characters, where half of the bits are logic-1 and half are logic-0. When transmitted at 1.25Gbps, this effectively results in a 125MHz square wave being generated on the serializer output. A DDMTD is used to measure the delay between this square wave and the 125MHz transmit reference clock.

The receive latency is measured in much the same way, but with the assistance of the node at the other end of the link which is instructed to transmit a sequence of “K28.5” characters.

The fiber asymmetry (α) is measured as part of the formal calibration procedure and is constant for a given combination of fiber type (e.g., G.652) and operating wavelength.

Synchronisation Process

With knowledge of PTP, the Link Delay Model and the DDMTD technique we can now describe how White Rabbit utilises these elements to accurately synchronise the slave clock.



The process begins with the encoding of the master's transmitted data using the master's reference clock (1). The reference clock is effectively transferred over the link to the slave and recovered along with the data in the receiver (2).

The slave-recovered clock is phase shifted using a DDMTD-enabled PLL and used to encode the slave's transmitted data (3). Again, this clock is effectively transferred over the link to the master and recovered along with the data in the receiver (4).

Another DDMTD is used as a phase detector in the master to measure the difference in phase between its reference clock and the looped-back slave clock ($phase_{MM}$).

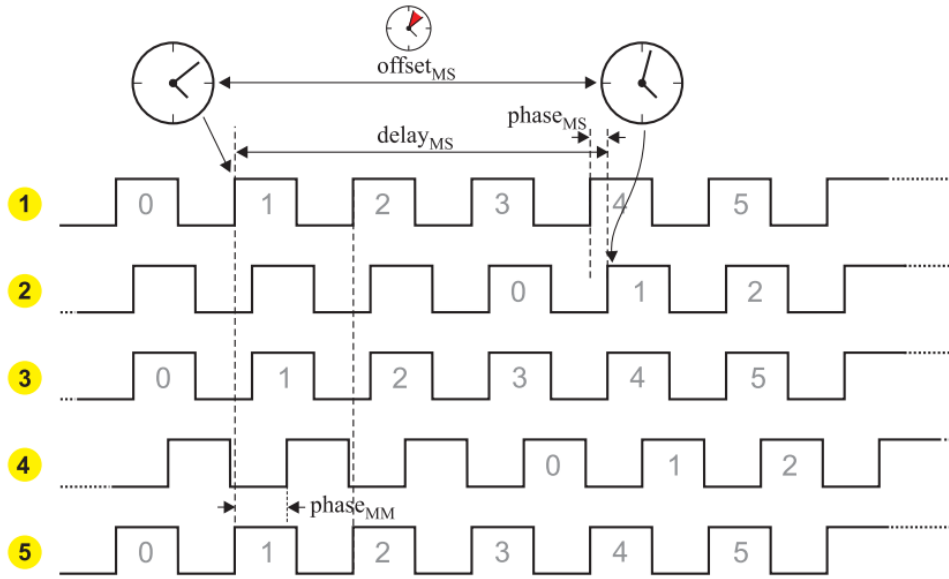
The value of $phase_{MM}$ is described as:

$$phase_{MM} = \{\Delta + \delta_{MS} + \delta_{SM} + phase_s\} \text{ modulus } T_{ref}$$

Where:

- Δ is the sum of all fixed delays ($\Delta_{TXM} + \Delta_{RXS} + \Delta_{TXS} + \Delta_{RXM}$).
- T_{ref} is the period of the 125MHz Gigabit Ethernet reference clock (8 ns).
- $phase_s$ is the programmable phase shift performed in the slave.

The White Rabbit algorithms use the PTP timestamps and $phase_{MM}$ measurements to servo the slave to produce the desired syntonized and synchronized clock (5). A more detailed treatment of those algorithms follows.



The receiver PTP timestamps (t_2 and t_4) can contain glitches and so are measured on both the rising and falling edges. The timestamp enhancement algorithm uses the value of $phase_{MM}$ relative to the expected phase transition value (t_{24p}) to select whether to use the rising or falling edge to generate the enhanced precision t_{2p} and t_{4p} timestamps.

The phase transition value (t_{24p}) is device specific and is measured once during factory calibration and programmed in to non-volatile storage. It can also be calibrated on-demand with the assistance of a connected master.

It is not necessary to enhance the precision of the transmit timestamps since these are generated from the local clock domain (i.e., their timestamps are always an integer number of cycles).

The round-trip delay ($delay_{MM}$) is calculated from the enhanced PTP timestamps:

$$delay_{MM} = (t_{4p} - t_1) - (t_3 - t_{2p}) = \Delta + \delta_{MS} + \delta_{SM}$$

With prior knowledge of the fiber asymmetry (α), the one-way fiber delay (δ_{MS}) is calculated:

$$\delta_{MS} = \frac{1 + \alpha}{2 + \alpha} (delay_{MM} - \Delta)$$

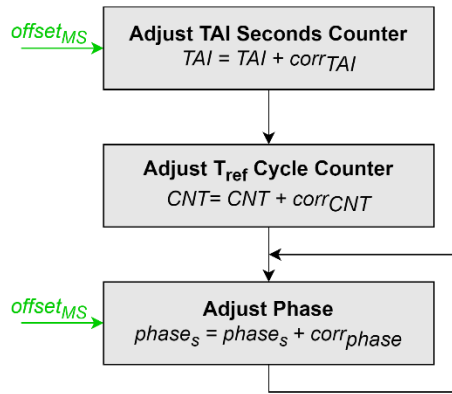
By adding the fixed delays in the master to slave path the master-slave delay ($delay_{MS}$) is calculated:

$$delay_{MS} = \frac{1 + \alpha}{2 + \alpha} (delay_{MM} - \Delta) + \Delta_{TXM} + \Delta_{RXS}$$

Lastly the current offset of the slave clock with respect to the master clock is calculated ($offset_{MS}$):

$$offset_{MS} = t_1 - t_{2p} - delay_{MS}$$

The flow diagram below describes the process in which the slave's counters & phase is adjusted using the calculated $offset_{MS}$ value:



The whole seconds value of $offset_{MS}$ is first used to correct the slave's TAI Full Seconds Counter:

$$corr_{TAI} = \left\lfloor \frac{offset_{MS}}{1s} \right\rfloor$$

The reference clock T_{ref} Cycle Counter (used to generate the 1PPS signal) is then adjusted by the whole number of T_{ref} cycles in the remainder:

$$corr_{CNT} = \left\lfloor \frac{offset_{MS} - corr_{TAI}}{T_{ref}} \right\rfloor$$

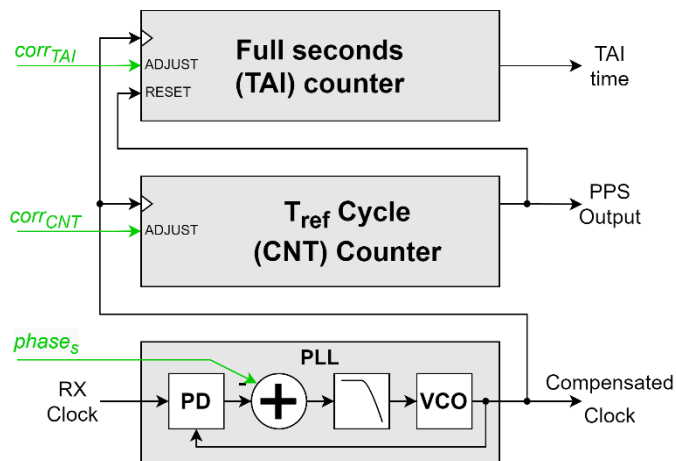
Lastly, the fractional T_{ref} remainder of $offset_{MS}$ is used to correct the phase of the slave clock.

$$corr_{phase} = offset_{MS} - corr_{TAI} - (corr_{CNT} \cdot T_{ref})$$

This process is periodically repeated to correct for drift in the link delay:

$$corr_{phase} = offset_{MS} - offset_{MS_previous}$$

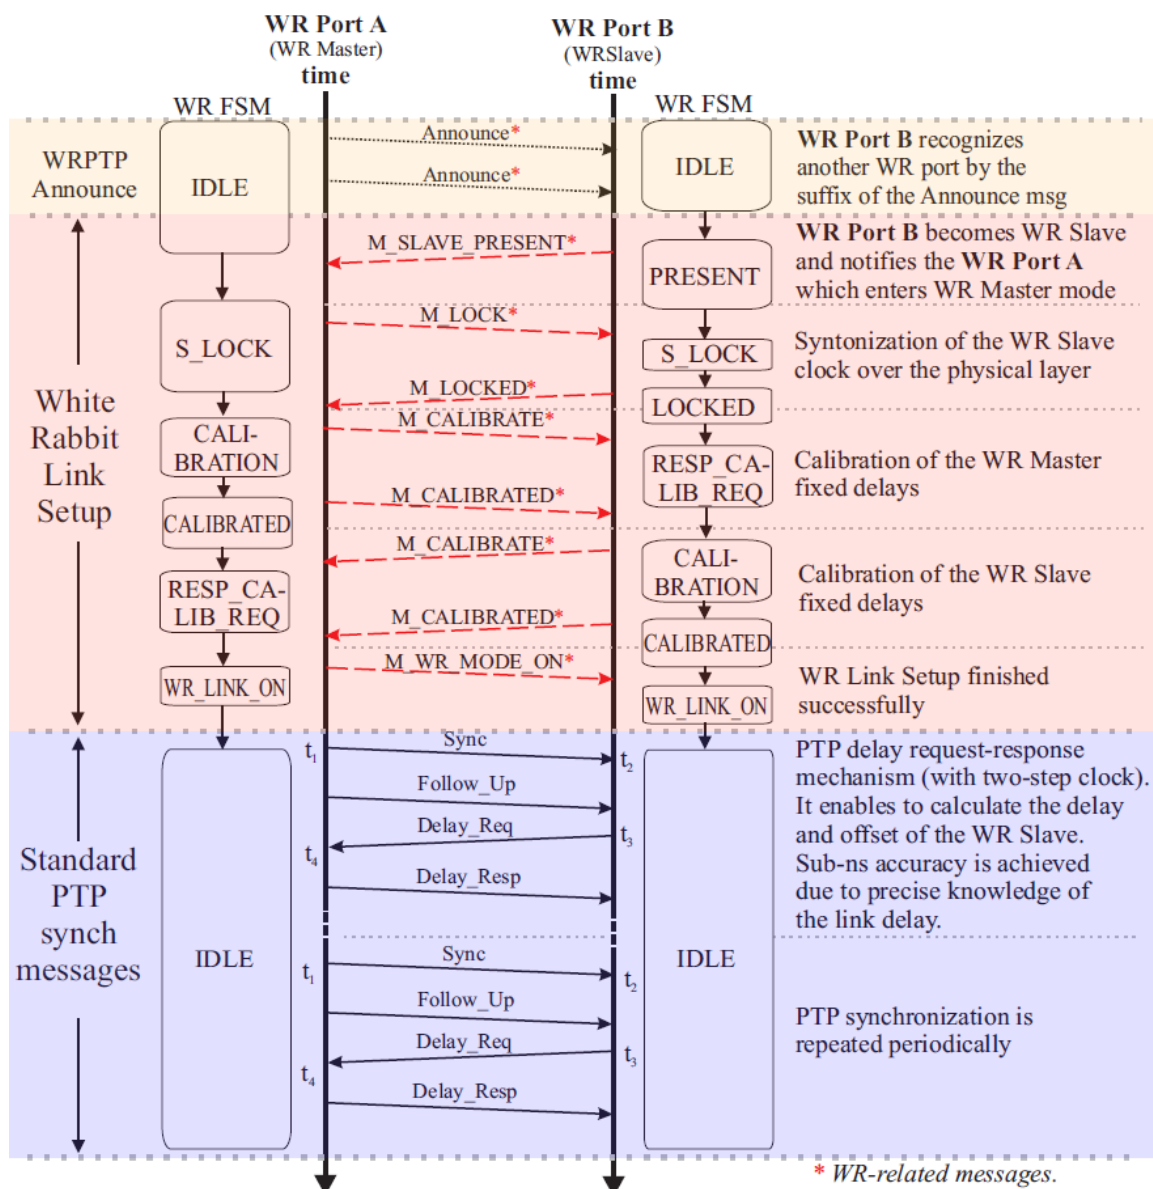
A possible implementation of the slave's clock servo is shown below. The TAI and CNT counters are derived from the compensated clock oscillator and so all three remain synchronous throughout the correction process.



Message Exchange Pattern

The sequence diagram below documents the White Rabbit & PTP message exchange process between a master and slave node. The three main phases are:

1. **Announce Phase** – The master makes any attached nodes aware that they are connected to a White Rabbit capable network.
2. **WR Link Setup Phase** – The slave clock is syntonised to the master clock and the fixed high-speed transceiver delays are measured both ends of the link.
3. **PTP Phase** – The PTP protocol executes to perform coarse calculation of the round-trip propagation delay of the link. The fine delay measurement is performed periodically by the White Rabbit core to extend the PTP timestamp precision.



A non-White Rabbit capable PTP v2 device (connected to “WR Port B”) can still participate in the PTP message exchanges and achieve synchronisation, however it would not benefit from the syntonisation, fine delay measurement and asymmetry correction provided by White Rabbit.

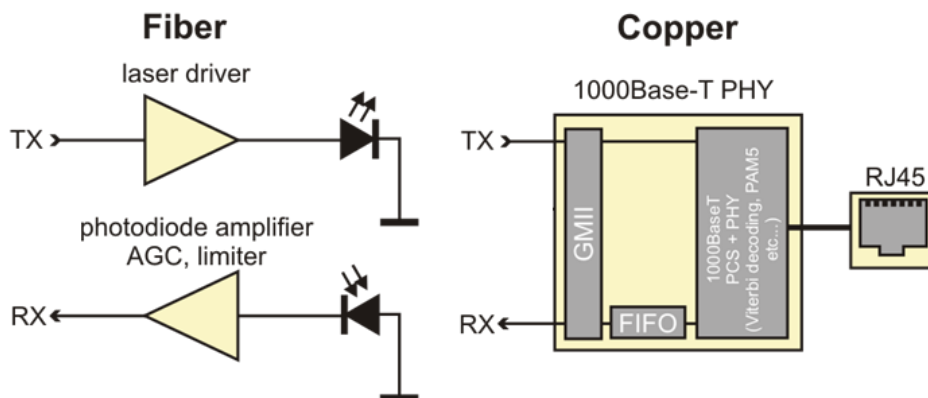
Ethernet Compatibility

White Rabbit is backwards compatible with IEEE1588 PTP v2, allowing it to be integrated with heterogenous timing networks, albeit with reduced accuracy and precision.

White Rabbit implements IEEE802.3 Gigabit Ethernet over a Single Mode fiber optic with 100BASE-BX10 being the preferred electro-optic standard. A commodity Small Form-factor Pluggable (SFP) optical transceiver module is used for medium and long-distance links. For particularly short links, optical attenuators may be required. Direct Attach Copper (DAC) modules can be used for very short hops if they incorporate an MSA compliant identification EEPROM.

White Rabbit does not yet formally support 10Gbit Ethernet, however this is being undertaken as part of the White Rabbit Network Switch v4 development (WRS4). The first revision of this new 24-port switch will initially operate at one gigabit, however future releases of the gateway/firmware should add support for 10Gbps on a subset of the ports.

The optical SFP modules used contain only a driver & laser diode in the transmit path and a photodiode and amplifier/AGC in the receive path. Essentially a bit state in the electrical domain is replicated in the optical domain and vis versa (albeit with a fixed delay). These SFP modules provide a direct interface between the high-speed serial transceiver on the FPGA and the optical medium.



White Rabbit cannot achieve the stated sub-nanosecond performance over 1000Base-T Ethernet (copper twisted pair). This is because FPGAs do not include the necessary electrical interface for a 1000Base-T (copper) physical medium and so an external PHY is required. These PHY devices contains logic, a FIFO and possibly other clocks domains that introduce an unknown variable delay between the data interface and the copper medium. 1000BaseT SFP transceiver modules also incorrectly report the link and auto negotiation states.

Component & Resource Requirements

As we will show, the bare minimum of additional components & resources is required to upgrade a system to support White Rabbit. For an application already using a compatible FPGA and SFP optical transceiver, the increase in Bill of Material (BoM) cost is less than £35 for single unit quantities.

Supported FPGA Families

White Rabbit is (currently) only realisable on a Field Programmable Gate Array (FPGA) from one of the supported families. Wrappers have been written for the device-specific elements (e.g., the High-Speed Serial Transceivers) to make them operate deterministically in the White Rabbit regime.

The list of supported device families is constantly growing. The table below provides a snapshot at the time of authoring this document.

Manufacturer	Family	Transceiver	Toolchain
AMD (Xilinx)	Virtex-5	GTP-5	ISE
	Spartan-6	GTP-6	
	Virtex-6	GTX-6	
	Artix-7	GTP-7	Vivado
	Kintex-7	GTX-7	
	Virtex-7	GTX-7	
	Virtex-7	GTH-7	
	Kintex Ultrascale	GTHE3/GTH-US/GTY-US	
	Kintex Ultrascale+	GTHE4	
	Zynq-7 (Kintex-7 + ARM uP)	GTX-7	
Intel (Altera)	Arria II	Various	Quartus
	Arria V	Various	

FPGA Resource Utilisation

The resource utilisation of a typical instantiation of the White Rabbit PTP Core on a Xilinx Spartan-6, XC6SLX45T FPGA is shown in the table below. This utilisation includes a 32-bit soft-core processor and its associated instruction and data memory.

Slice Logic Utilization	Used	Available	Utilization
Number of Slice Registers	6,791	54,576	12%
Number of Slice LUTs	8,956	27,288	32%
Number of occupied Slices	3,345	6,822	49%
Number of MUXCYs used	1,532	13,644	11%
Number of bonded IOBs	26	296	9%
Number of RAMB16BWERS	56	116	48%
Number of RAMB8BWERS	3	232	1%
Number of BUFIO2/BUFIO2_2CLKs	1	32	3%
Number of BUFG/BUFGMUXs	7	16	43%
Number of BSCANs	1	4	25%
Number of DSP48A1s	3	58	5%
Number of GTPA1_DUALs	1	2	50%
Number of PLL_ADVs	2	4	50%

Required Hardware

The White Rabbit PTP Core requires three digitally tuneable clock signals:

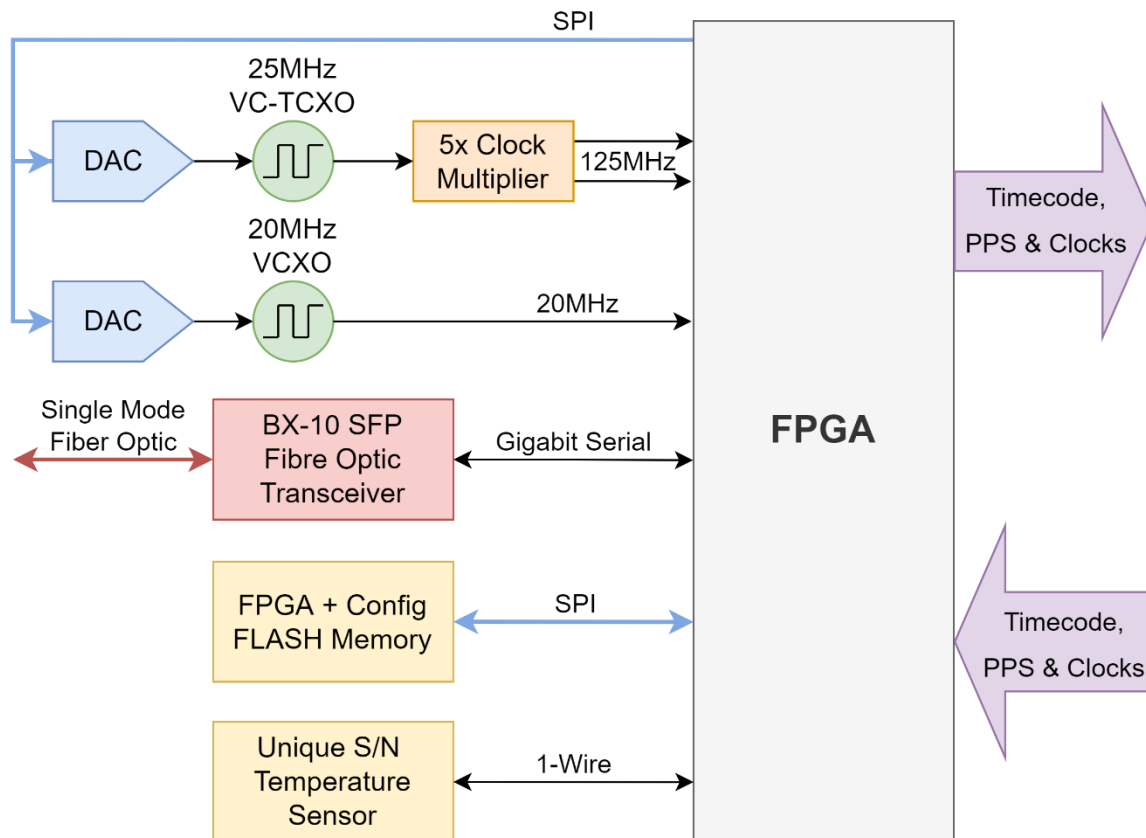
1. 125MHz/62.5MHz clock for the 8-bit/16-bit PHY respectively.
2. 125MHz clock for the system reference, PPS generation and soft-core clock.
3. 20MHz clock for the DDMTD PLL.

The first two clocks (125MHz/62.5MHz) are derived from a single 25MHz Voltage Controlled Temperature Compensated Crystal Oscillator (VC-TCXO) that drives a clock multiplier. This is typically implemented as a single-chip frequency synthesiser with integrated fan-out buffer. Because of the need to maintain a very clean, low-jitter performance, these clocks are generally distributed as LVDS

signals. The VC-TCXO will typically need to have a temperature stability of 2.5ppm and a tuning range of more than 10ppm.

The third clock (20MHz) can be a lower cost, less stable (e.g., 50ppm) single-ended Voltage Controlled Crystal Oscillator (VCXO).

The two oscillators are (usually) tuned by a voltage-output Digital to Analogue Converter (DAC) with a common voltage reference. A commodity low speed monotonic 16-bit DAC is adequate for most applications.



It is not necessary to use a high-performance ovenised 50ppb oscillator. In fact, doing so places extreme stability requirements on the upstream grand master clock as a result of the more limited PLL “capture & lock” range. The tuning DAC & voltage reference combination could also need to have a resolution approaching 20-bits (1ppm) which would be considered Instrumentation Grade. Where a very low phase noise clock required, it may be best to lock a high-performance oscillator to a spare output on the fan-out buffer with a narrow loop bandwidth. This provides the long-term stability of White Rabbit master clock (inside the loop bandwidth) whilst providing lower phase noise (outside the loop bandwidth).

Most FPGAs are accompanied with an SPI or Quad-SPI FLASH memory to store the bitstream. This bitstream defines the configuration & connectivity of all logic elements, peripherals and memory in an FPGA. Once the FPGA has been initialised, White Rabbit takes control of the FLASH memory interface and utilises unused space towards the top of this non-volatile memory to store its settings. Therefore, no additional EEPROM/FLASH memory devices are required.

White Rabbit interfaces with the FLASH device via single bit wide SPI, rather than Quad-SPI, and so the bit-width interface flags in the memory need to be programmed accordingly.

An optional 1-Wire temperature sensor provides environmental monitoring and a unique serial number from which a pseudo-globally unique Ethernet MAC address can be derived (if not defined in the FLASH configuration memory).

Power Consumption

A White Rabbit PTP Core implemented on a 7-Series Xilinx FPGA along with the VC-TCXO, VCXO, DACs, voltage reference, clock generator and SFP transceiver will consume circa 2W (assuming >95% efficient regulators are used for the FPGA core, I/O and MGT rails).

White Rabbit PTP Core (v4.2)

The White Rabbit PTP Core (WRPC) is a single FPGA Hardware Description Language (HDL) module that can either be instantiated as a standalone timing module (i.e., the only HDL in an FPGA) or integrated into a much larger System-on-Chip design requiring enhanced synchronisation & determinism.

The WRPC makes building White Rabbit in to a system much easier by shielding the engineer from having to deal with all of the low-level precision timing complexity and only presenting the application with the useful data, control and monitoring interfaces.

PTP Modes

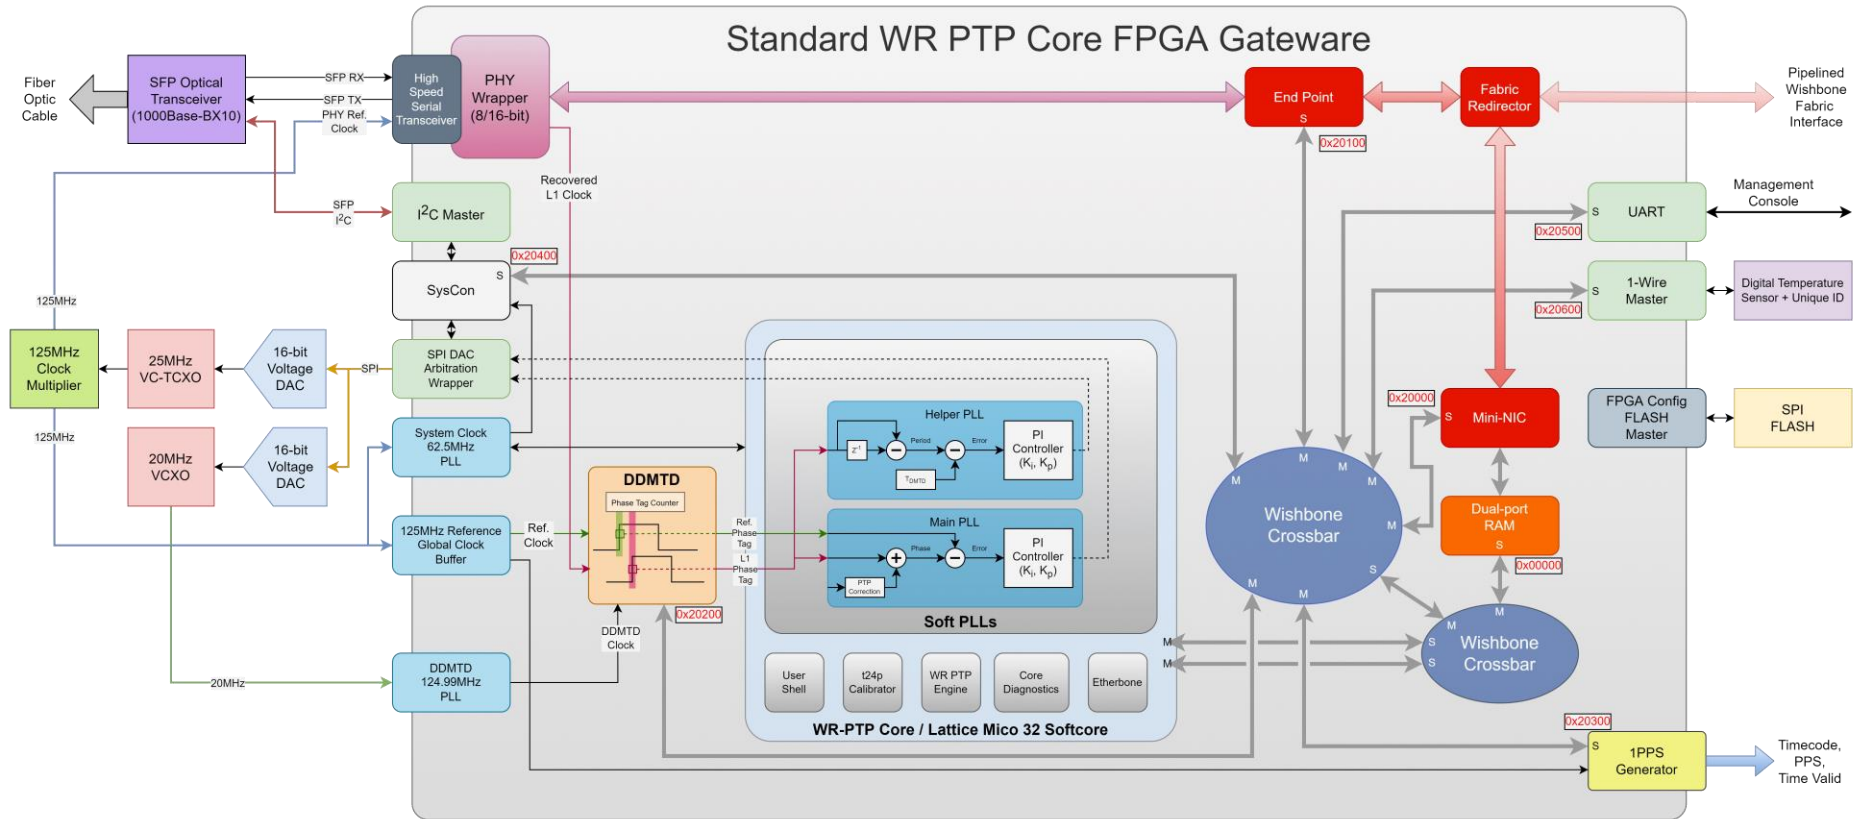
The WRPC implements an IEEE1588 PTP Ordinary Clock and can operate in one of three modes:

- **Grand Master:** In this mode, the WRPC is fed with an external timecode, 10 MHz clock and 1PPS reference signal from an Atomic Clock or GPS Disciplined Oscillator. As a Grand Master, the WRPC provides the source of synchronisation for all other White Rabbit compliant devices in the network.
- **Master:** In this mode, WRPC operates with its internal clock “free-running” (i.e., not disciplined to an external clock/timing reference). Again, the WRPC provides the source of synchronisation for all other White Rabbit compliant devices in the network (albeit only with the stability performance of its internal oscillator).
- **Slave:** In this mode, WRPC synchronizes and synchronises its internal clock to the White Rabbit network Grand Master/Master. Here the WRPC provides a timecode, 1PPS and reference clock output for application logic or attached equipment.

The PTP mode can be changed at any time through the management interface.

WRPC Architecture

A typical implementation of the WRPC is outlined in the diagram below. It shows the interaction between the external components and the internal gateway/firmware functions that form the WRPC.



Network traffic from the SFP module is received by the High-Speed Serial Transceiver and passed to the End Point module for timestamping. Regular Ethernet frames are routed in the Fabric Redirector to the Pipelined Wishbone Fabric Interface for the application logic to process. Out of Band (OOB) White Rabbit/PTP related data is redirected to the Mini-NIC and stored in the Dual Port Memory block for processing by the LM32 soft-core processor.

Likewise, regular Ethernet frames from the application logic (connected to the Pipelined Wishbone Fabric Interface) or OOB frames from the LM32 soft-core, can be transmitted via the Fabric Redirector, the End Point and lastly the High-Speed Serial Transceiver. This approach allows a networked HDL design to be connected to the WRPC, whilst the White Rabbit traffic and internal behaviour remains transparent to the application logic.

Wishbone Interconnect Architecture Overview

The Wishbone is an Interconnect Architecture developed by [OpenCores.org](https://opencores.org). This interconnect is intended as a general-purpose interface and, as such, it defines the data exchange mechanism between connected cores.

Wishbone is analogous to the flexible, variable-width, vendor-independent busses found in computers. Where Wishbone differs is that it is targeted at SoC applications that are not encumbered by the need to drive backplanes/long multi-drop busses.

Features of the Wishbone Interconnect Architecture include:

- Simple, compact and with very few logic resource requirements.
- Supports Read/Write, Block Transfer and Read-Modify-Write cycles
- Data and Operand bus widths up to 64-bits.
- Supports both Big-Endian and Little-Endian data ordering.
- Several core interconnection methods including:
 - Point-to-Point
 - Shared Bus
 - Crossbar Switch
 - Switched Fabric
- Handshaking protocol allows each IP core to throttle its data transfer speed.
- Supports single clock data transfers.
- Uses a Master/Slave architecture for more flexible system designs.
- Multi-Master capable.
- Supports several arbitration methodologies including:
 - Priority arbiter
 - Round-robin arbiter

The WRPC utilises Wishbone to interconnect the LM32 soft-core, via Crossbar Switches, to all of the peripheral HDL modules that implement a Wishbone Slave Interface.

LM32 Processor

At the heart of the WRPC is a Lattice Mico32 (LM32) soft-core processor. The LM32 executes the open-source White Rabbit PTP Core Software ([wrpc-sw](#)) which controls all the HDL modules and performs the following functions:

- WR PTP Engine – A minimal implementation of the PTP daemon with profile extensions for White Rabbit.
- Software PLLs – Phase Locked Loops controlling the DDMTD, Main & Auxiliary oscillators.
- User Shell – A command-line interface for configuration and monitoring.
- t24p calibrator – An automated procedure for measuring t2/t4 phase transition value.
- IPv4, ICMP, BOOTP, ARP – Optional protocols for use when Etherbone is enabled.
- Diagnostics
- HDL low-level device drivers (for EEPROM, SysCon, PPS-gen, Packet filter, UART etc).

The WRPC-SW is written in the C language and is built using the LM32 Cross-Compiler. It can be extended for application-specific requirements; however, program storage space is limited to the FPGA's maximum Block-RAM size.

The LM32 core is clocked from a 62.5MHz PLL (derived from the 125MHz global reference clock) and is connected to all the HDL modules within the WRPC gateway via a Wishbone interface.

HDL Modules

In addition to the LM32 soft-core, the WRPC includes several HDL modules that each perform an interface or timing function:

- **Mini-NIC** – A compact Network Interface Card implementation
- **Endpoint** – An Ethernet Media Access Controller (MAC) with timestamping.
- **Soft-PLL** – An interface between the Software-PLLs and the DDMTDs.
- **PPS Generator** – A module responsible for generating the 1PPS signal and providing counters for the TAI and nanoseconds CNT values.
- **Wishbone SysCon & GPIO** – An important module that coordinates all activities of the internal logic within the WISHBONE interconnect and, in this instance, it also provides a number of register-mapped GPIO lines for the status indicators and I2C/SPI busses.
- **UART** – A Universal Asynchronous Receiver-Transmitter module through which the user interacts with WR PTP Core command shell.
- **1-Wire Master** – An interface to the external temperature sensor/unique ID chip.

Each HDL module has a unique Base Address allocated within the Wishbone memory space (Also indicated by red text in the WR PTP Core diagram).

Module Name	Base Address (Bytes)
WRPC Internal Memory	0x00000
Mini-NIC	0x20000
Endpoint	0x20100
Soft-PLL	0x20200
PPS generator	0x20300
Wishbone SysCon & GPIO	0x20400
UART	0x20500
1-Wire Master	0x20600
Aux WB Master	0x20700
WR Core Diagnostics	0x20800

Registers for the control and monitoring of the HDL modules exist at specified offsets from the base address. An example for the WR Core Diagnostics module is shown below.

F.1 WR Core Diagnostics

[version 0x00000001]

Diagnostics information accessible via WR

F.1.1 Memory map summary

SW Offset	Type	Name	HW prefix	C prefix
0x0	REG	Version register	wrc_diags_ver	VER
0x4	REG	Ctrl	wrc_diags_ctrl	CTRL
0x8	REG	WRPC Diag: servo status	wrc_diags_wdiag_sstat	WDIAG_SSTAT
0xc	REG	WRPC Diag: Port status	wrc_diags_wdiag_pstat	WDIAG_PSTAT
0x10	REG	WRPC Diag: PTP state	wrc_diags_wdiag_ptpstat	WDIAG_PTPSTAT
0x14	REG	WRPC Diag: AUX state	wrc_diags_wdiag_astat	WDIAG_ASTAT
0x18	REG	WRPC Diag: Tx PTP Frame cnts	wrc_diags_wdiag_txfcnt	WDIAG_TXFCNT
0x1c	REG	WRPC Diag: Rx PTP Frame cnts	wrc_diags_wdiag_rxfcnt	WDIAG_RXFCNT
0x20	REG	WRPC Diag: local time [msb of s]	wrc_diags_wdiag_sec_msb	WDIAG_SEC_MSB
0x24	REG	WRPC Diag: local time [lsb of s]	wrc_diags_wdiag_sec_lsb	WDIAG_SEC_LSB
0x28	REG	WRPC Diag: local time [ns]	wrc_diags_wdiag_ns	WDIAG_NS

Here, the register for the monitoring PTP state is located at an offset of 0x10 (i.e., an absolute address of 0x20810) and the lower 8-bit define the current PTP state.

WRPC Diag: PTP state

HW prefix: wrc_diags_wdiag_ptpstat
 HW address: 0x4
 SW prefix: WDIAG_PTPSTAT
 SW offset: 0x10

31	30	29	28	27	26	25	24
-	-	-	-	-	-	-	-
23	22	21	20	19	18	17	16
-	-	-	-	-	-	-	-
15	14	13	12	11	10	9	8
-	-	-	-	-	-	-	-
7	6	5	4	3	2	1	0
PTPSTATE[7:0]							

- **PTPSTATE** [*read-only*]: PTP State
 - 0: NONE
 - 1: PPS_INITIALIZING
 - 2: PPS_FAULTY
 - 3: disabled
 - 4: PPS_LISTENING
 - 5: PPS_PRE_MASTER
 - 6: PPS_MASTER
 - 7: PPS_PASSIVE

For applications where an engineer wishes to extend the WRPC, the “Wishbone Slave Generator” utility ([wbgen2](#)) can be used to generate custom VHDL/Verilog cores which implement a Wishbone Bus Slave. This tool has support for registers, memory blocks, FIFOs and interrupts. The utility accepts a C-like file syntax with an abstract description of what you would like the slave to present on its Wishbone interface.

This utility generates:

- An automatically allocated memory layout
- VHDL/Verilog code for the slave module
- C header files for driver development
- HTML documentation

A utility called [Cheby](#) has been recently developed that will eventually replace wbggen as the preferred Wishbone interface generation tool.

White Rabbit PTP Core Interface Signals

A brief summary of the key WRPC signals is provided below. Not all signals are required for a standalone slave implementation.

- **PHY Interface**
 - The WRPC HDL presents a generic interface for connecting to the PHY since the capabilities and performance of the High-Speed Serial Transceivers vary across different FPGA families. A platform-specific wrapper makes the White Rabbit PTP Core FPGA-agnostic whilst adding custom logic to improve the determinism of the PHY.
- **Timecode Outputs**
 - TAI (40 bit wide) - Full seconds count in the International Atomic Time scale.
 - CNT (28 bit wide) - Fractional part of TAI (0 – 124999999).
 - One count per Tref period (e.g., 8ns).
 - 1PPS - One Pulse Per Second. Pulse generated when the CNT = 0. Programmable width.
- **Reference Clock Output**
 - WR Clock - White Rabbit synchronised & synchronised 125MHz reference clock
- **External Reference Inputs**
 - Ext 1PPS – Needed in Grand Master mode to qualify next 10MHz clock edge.
 - Ext. 10MHz – Needed in Grand Master mode for precise timing of slaves.
- **Status Outputs**
 - Timecode Valid – Asserted once the slave has fully synchronised with the master.
 - Ethernet Link – Indicates a link is present.
 - Ethernet Activity – Indicates network activity.
 - PPS – Pulse-stretched version of the 1PPS signal for driving an LED.
- **Management**
 - UART TX & RX – Provides the serial command shell (115,200 baud, 8N1).
 - Reset WRPC Input – Active low reset for LM32 only.
- **Other Interfaces**
 - SFP I²C – Connects to the MSA identification EEPROM in the SFP optical transceiver. Also enables digital diagnostic and wavelength selection in supported modules.
 - DAC – Data words & latches for DDMTD, Main and Auxiliary oscillator tuning DACs. An external module must deal with arbitrating the bus access between DACs.
 - FLASH SPI – Used by the WRPC to read/write from the upper portion of the FLASH configuration memory after FPGA initialisation.

- 1-Wire – Connects to DS18B20 temperature sensor with unique serial number.
- Pipelined Wishbone Fabric Interface – Provides a means of the application logic to send and receive regular Ethernet frames.

WRPC Monitoring & Control

The WRPC provides a serial management console with a rich feature set and so commands are provided for:

- SFP identification
- Database manipulation
- Calibration
- Debugging
- Boot script editing
- MAC/IP address query & override
- PLL configuration
- PTP mode
- VLAN membership
- Status monitoring

The initialisation of the WRPC is logged to the serial console during start-up and includes:

- Reading of MAC from FLASH memory or generating from 1-Wire serial number.
- Reading or calibrating t24p phase transition value from memory.
- Calibrating the LM32 timing loops.
- Initially locking the PLL.
- Detecting the SFP type and matching to the internal database.
- Starting the PTP Engine in the selected mode (Grand Master, Master, Slave).

Issuing the “gui” command (shown below) provides a real-time display of the nodes’ operating mode, synchronisation status and key timing parameters.

```

WR PTP Core Sync Monitor: PPSI - LEN board
Esc = exit

TAI Time:                Sat, Mar 6, 2021, 04:28:15

WR-LEN mode : WRC_SLAVE_WRO
-----
Link status:

wr0 : Link up   (RX: 2040, TX: 614), mode: WR Slave   Locked   Calibrated
IPv4: BOOTP running

wr1 : Link down
-----

Servo state:            TRACK_PHASE
Phase tracking:         ON
Synchronization source: wr0

Timing parameters:

Round-trip time (mu):   64211797 ps
Master-slave delay:     32092282 ps
Master PHY delays:     TX: 234636 ps, RX: 283095 ps
Slave PHY delays:      TX: 205320 ps, RX: 218812 ps
Total link asymmetry:   27233 ps
Cable rtt delay:       63269934 ps
Clock offset:          1 ps
Phase setpoint:        11082 ps
Skew:                  1 ps
Manual phase adjustment: 0 ps
Update counter:        324
--

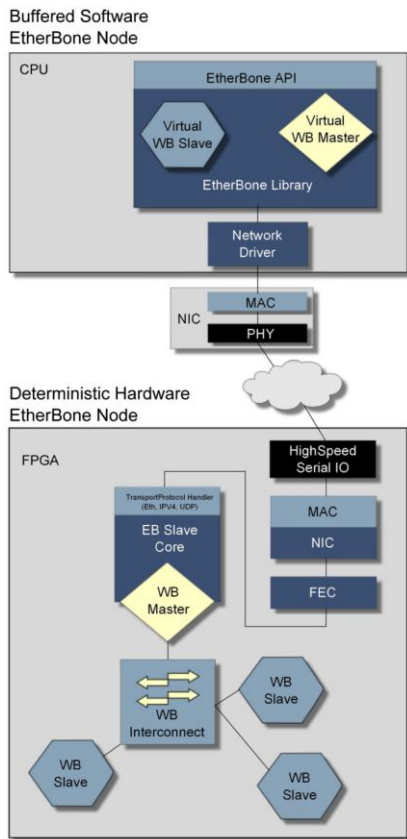
```

By building the WRPC-SW with the “CONFIG_SNMP” flag, an SNMP Agent is enabled allowing SET & GET requests to be issued to the core from an SNMP Manager.

By building the WRPC-SW with the “CONFIG_CMD_LL” flag, it is possible to enable a low-level command (“devmem”) that allows the entire Wishbone address space to be “peeked & poked” from the console.

It is also possible to access the Wishbone register address space via the optical Ethernet network using the EtherBone protocol. This requires an additional (open) FPGA core to be built that connects to the Wishbone Fabric Interface.

EtherBone encapsulates read and write memory transactions in to standard UDP/IP frames. Software implementations of a Wishbone Master are available for several programming languages. Wishbone’s operation is transparent, deterministic and low-latency.



0	4	8	12	16	20	24	28	31
Version	IHL		Type of Service		Total Length			
Identification				D F	M F	Fragment Offset		
Time To Live		Protocol		Header Checksum				
Source IP								
Destination IP								
Options & Padding								
Source Port				Destination Port				
Length				Checksum				
Data								
Magic (0x4E6F)				Version	P R	P R	AddrSz	PortSz
Probe ID / Potential padding to 64-bit alignment								
B C A	R C A	R F A	C Y C	W C A	W C F	WCount		RCount
Potential padding to 64-bit alignment								
BaseWriteAddr								
WriteVal 1								
⋮								
WriteVal M								
BaseRetAddr								
ReadAddr 1								
⋮								
ReadAddr N								

Lastly, a Wishbone-compliant master can be implemented between the WRPC Crosspoint Switch and application logic. This allows control and monitoring of White Rabbit from an embedded soft/hard-core microcontroller.

SFP Identification

As outlined in the Link Delay Model, the Δ_{TX} and Δ_{RX} Fixed Delay values include both the circuit (δ_{CIR}) and SFP (δ_{SFP}) delays. The WRPC therefore needs a mechanism to identify the connected SFP type so it can recall the correct calibration parameters. This is achieved by interrogating the ID EEPROM fitted in Multiple Source Agreement (MSA) compliant SFP transceivers.


At start-up, or when instructed from the command shell, the WRPC reads the ID and performs a look-up in its internal database to recall the correct Δ_{TX} , Δ_{RX} and fiber optic asymmetry (α) values. The database allows parameters to be stored for several uniquely identifiable SFPs.

The “sfp” shell command provides a means of listing, adding and removing SFP entries in the database.

Fiber Optic & Wavelength Selection

To achieve the very best performance with White Rabbit it is necessary for the optical fiber path length in each direction to be as closely matched as possible whilst ensuring those paths both experience the same environmental conditions (i.e., temperature). This is most easily achieved by using a single fiber with a different wavelength travelling in each direction.

White Rabbit aims to adhere to the Gigabit Ethernet standard and so the bidirectional IEEE 802.3ah 1000Base-BX10 SFP transceivers are the most suitable for this application. These transceivers are inexpensive and capable of operating up to 10km over one strand of single mode fiber. The standard specifies that the Downstream module transmits at 1490nm, whilst the Upstream module transmits at 1310nm.



The SFP transceivers are easily identified with a **violet** marker on the Downstream release latch and a **blue** marker on the Upstream.



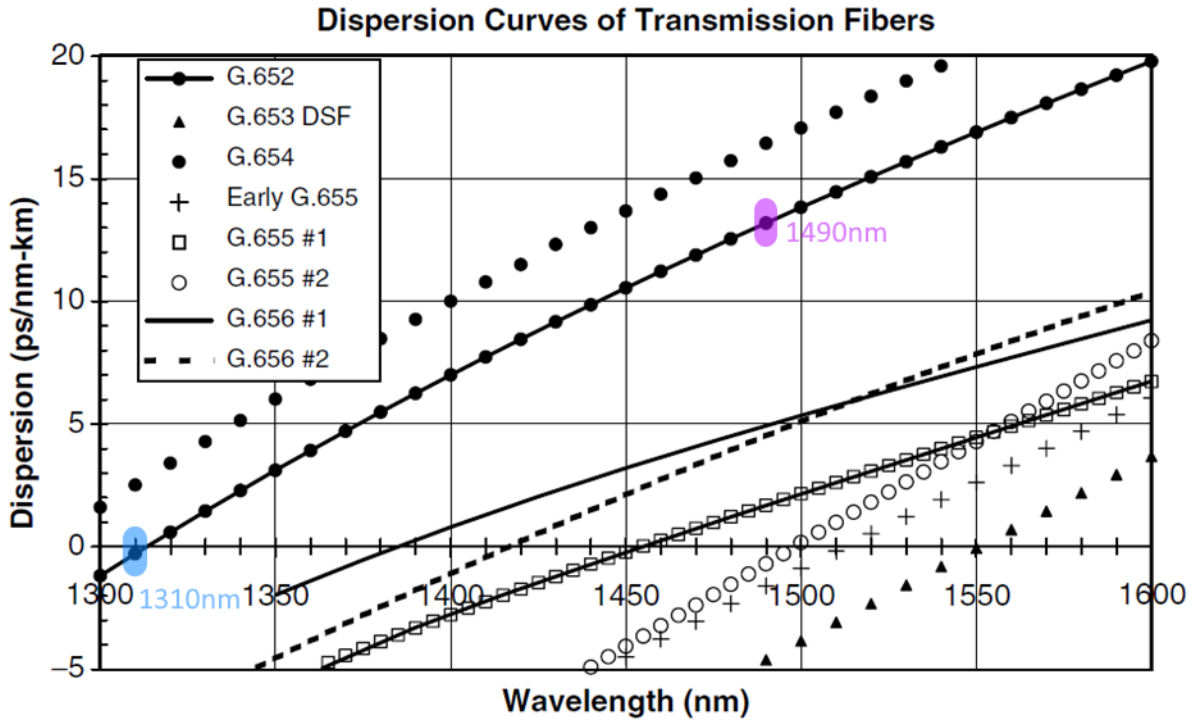
Non-standard SFP modules are available that operate at different wavelengths (e.g., 1490nm & 1550nm) and these can have a much longer operating range of up to 80km.

Having selected the operating wavelength, consideration should next be given to the type of optical fiber to install. The International Telecommunications Union (ITU) has developed many formal standards for the performance of optical fibers targeted at different applications. The standards define many parameters including the mechanical configuration, materials, cut-off wavelength, attenuation and chromatic dispersion.

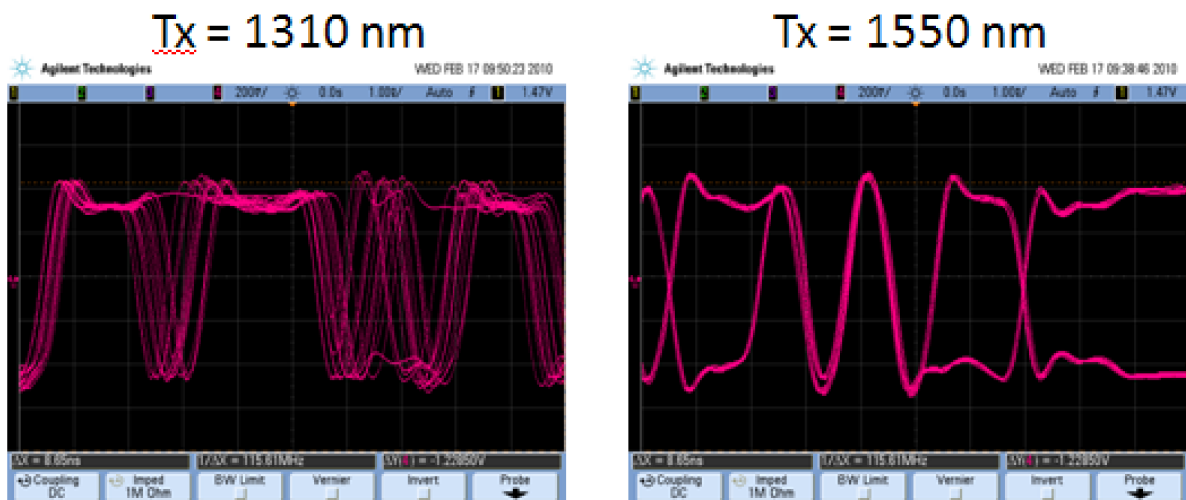
Common ITU optical fiber standards include:

- [ITU G.652](#) - Chromatic dispersion optimised for 1310nm.
- [ITU G.655](#) - Chromatic dispersion optimised for 1550nm.
- [ITU G.657](#) – Improved Bending-Loss performance compared to G.652.

The Chromatic Dispersion performance defines the delay per km at a given wavelength and this parameter is of particular importance in White Rabbit applications. The plot below shows the dispersion performance for several common fiber types with points highlighted at 1310nm and 1490nm for G.652.



The impact of chromatic dispersion is shown in the oscilloscope captures below. These plots show a 1310nm and 1550nm wavelength signal after travelling down 10km of G.655 fiber optic.




The 1310nm plot clearly shows significant Inter-Symbol Interference (ISI). This will manifest itself as jitter in the recovered clock and, in extreme cases, bit errors in the data stream.


As already stated, the G.655 is optimised for the lowest chromatic dispersion at 1550nm, rather than 1310nm, however this only serves to exacerbate an underlying performance issue related to how the laser beam is produced.

Optical transceivers often use different semiconductor technologies to generate the beam at different wavelengths. Two common types are the Fabry-Perot laser (FP, with a typical spectral width $\Delta\lambda = 1.6\text{nm}$ at 1310nm) and Distributed Feed-Back laser (DFB, with a typical spectral width $\Delta\lambda = 0.09\text{nm}$ at 1550nm). The optical power spectrum of the FP (left) and DFB (right) are shown below.

Tx = 1310 nm (LD Transmitter)




Tx = 1550 nm (DFB LD Transmitter)




In the above example, the optical fiber has the highest dispersion at the wavelength where the FP laser diode has the widest spectral width. Swapping the optical fiber for 1310nm optimised G.652 we see much less degradation (due to ISI) at 1310nm, whilst still achieving acceptable performance at 1550nm.

Tx = 1310 nm



Tx = 1550 nm



The resulting “rule of thumb” is that the fiber optic should be chosen to have the lowest chromatic dispersion at the wavelength of the laser with the widest optical spectra.

Calibration

Calibration of White Rabbit network components is required to achieve the very highest synchronisation performance. Careful calibration is particularly important where many White Rabbit devices are to be cascaded.

The White Rabbit Link Delay Model describes the two categories of delay in a White Rabbit master-slave connection:

- Fixed transmission and reception delays (Δ_{TXM} , Δ_{RXM} , Δ_{TXS} , Δ_{RXS}) which are internal to the nodes.
- Variable physical medium delays (δ_{MS} and δ_{SM}).

For a given build (“Place and Route”) of the FPGA design, PCB tracking and SFP transceiver combination, measurements are required to determine the unknown $\Delta_{TXM}/\Delta_{RXM}/\Delta_{TXS}/\Delta_{RXS}$ parameters. This step must be performed for every device in a White Rabbit Network and this includes all the (used) ports of any White Rabbit Switches.

An absolute calibration of the timing reference plane in White Rabbit nodes is complicated by the electro-optic conversion that takes place in each direction of a link. Whilst new techniques are being developed, the preferred methodology is still to perform a relative alignment against a “golden” unit. The full calibration procedure and derivations of the formulae are available on the Open Hardware Repository. Here we will provide a very basic overview.

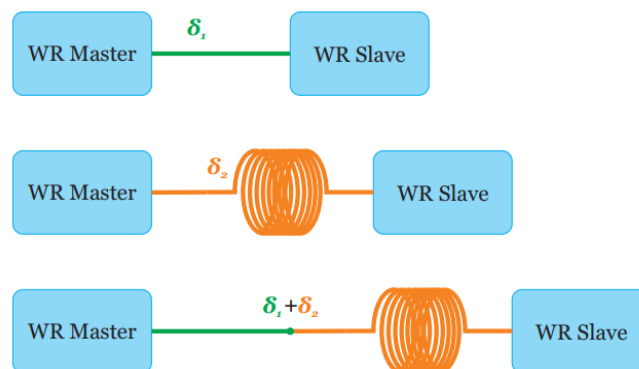
Reference Fiber Length Measurement

The first step of the calibration process is to determine the propagation delay of two fiber optic cables:

1. A “Helper” fiber (f_1 , a few metres long, and green in the diagram).
2. A “Long” fiber (f_2 , ideally a few kilometres long, and orange in the diagram).

Measurements are performed of the round-trip delay with fibers f_1 , f_2 and $f_1 + f_2$ installed between two uncalibrated WR nodes (master & slave fixed delays set to zero)

This step uses the substitution method to deduce the round-trip time of each fiber.

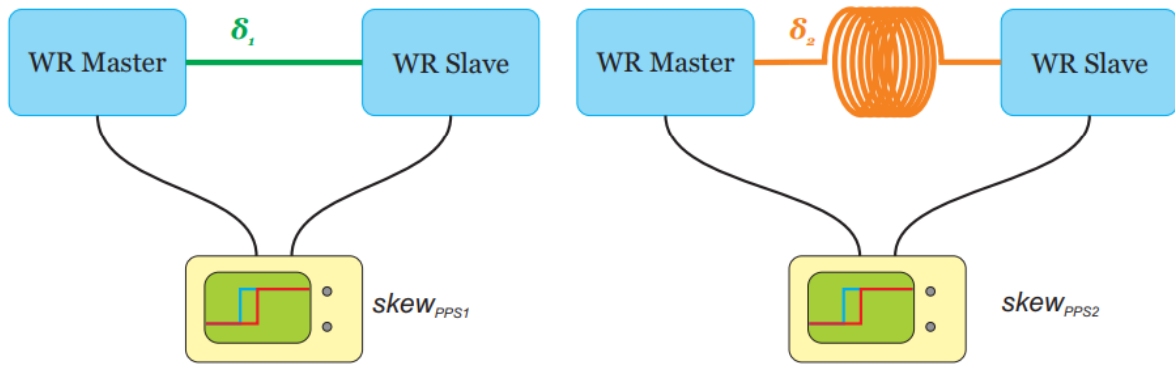


Fiber Optic Relative Delay Coefficient Measurement

When a single fiber optic is used with a different wavelength travelling in either direction, the propagation velocities will not be equal which gives rise to asymmetry (δ_{MS} is not equal to δ_{SM}).

This asymmetry is denoted by the relative delay coefficient (α) and its value is required by the White Rabbit algorithms to calculate δ_{MS} or δ_{SM} from the round-trip delay.

Here skew measurements are performed on the 1PPS outputs with fiber f_1 and then fiber f_2 fitted between the two uncalibrated WR nodes.




Calculations are performed using the skew readings and fiber delay to determine the relative delay coefficient (α):

$$\alpha = \frac{2(skew_{PPS2} - skew_{PPS1})}{\frac{1}{2}\delta_2 - (skew_{PPS2} - skew_{PPS1})}$$

Golden Calibrator

To complete this step, two nodes of the same PCB design, FPGA bitstream and complementary (Upstream & Downstream) SFPs are required. One node is selected as the “golden” calibrator node.



The round-trip delay is again measured and the sum of the transmit & receive delays calculated.

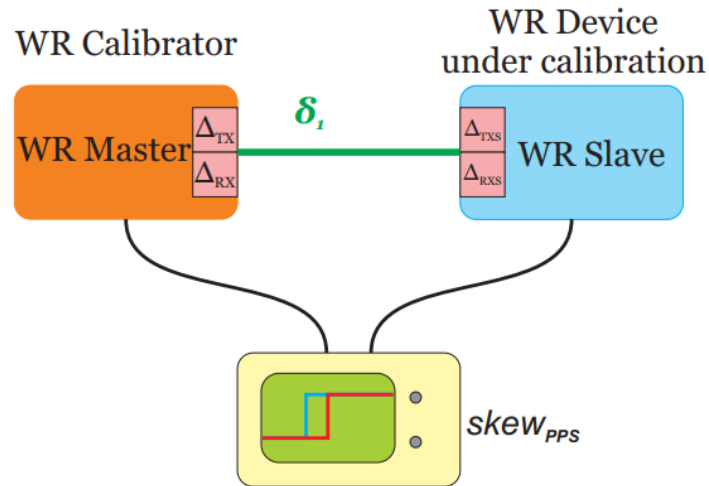
$$\Delta_{TX} + \Delta_{RX} = (delay'_{MM1} - \delta_1)/2$$

A convention is asserted that the transmit and receive delays are equal as any asymmetry is accounted for later in the actual node calibration. The calibrator node is programmed with the calculated values.

Node Fixed Delay Alignment

We can at last now align our White Rabbit network nodes using our f_1 “helper” fiber of known round-trip time and the defined “golden” calibrator.

A measurement is then performed of the total round-trip time ($delay_{MM}$), the “golden” calibrator transmit/ receive fixed delay ($\Delta_{TXM}/\Delta_{RXM}$) and the Device Under Calibration “bitslide” value (ϵ_S).



The coarse transmission and reception delays are then calculated, programmed in to the DUC's SFP database and the PTP daemon restarted:

$$\Delta'_{TXS} = \Delta'_{RXS} = \frac{1}{2}\Delta_S = \frac{1}{2}(delay_{MM} - \Delta_{TXM} - \Delta_{RXM} - \epsilon_S - \delta_1)$$

Next, the skew is measured between the "golden" calibrator and the DUC. The skew value is used to provide the final asymmetry correction to the DUC's fixed transmit and receiver delay parameters:

$$\begin{aligned}\Delta_{TXS} &= \frac{1}{2}\Delta_S - skew_{PPS} \\ \Delta_{RXS} &= \frac{1}{2}\Delta_S + skew_{PPS}\end{aligned}$$

The DUC's PTP daemon is again restarted and should synchronise with a sub-nanosecond offset.

Temperature Sensitivity

It should be noted that the fixed delays stored in the White Rabbit PTP Core (v4.2) database do not have a temperature coefficient and so cannot currently account for variations in the node temperature. Whereas variation of propagation delay in the fiber optic (for example, due to temperature changes) are automatically corrected for by the White Rabbit algorithms.

Temperature compensation has been investigated as part of the Large High Altitude Air Shower Observatory (LHAASO) project. Measurements have been performed over temperature on a "CuteWR" node (a White Rabbit node in an FMC form-factor) and this was found to exhibit a 10ps/°C deviation.


An experimental build of the White Rabbit PTP Core was subsequently created that added linear temperature coefficients to the fixed delay parameters. When re-measured, this change resulted in a reduction in deviation to 1-2ps/°C.

Performance Testing

To demonstrate the performance of White Rabbit, a test setup was devised with two WR-LEN modules, configured as slave nodes, connected to a White Rabbit Switch v3 – Low Jitter (WRS3-LJ), configured as a Free-Running Master.

One WR-LEN module was connected to the WRS3-LJ with a 0.5m fiber optic cable. The second WR-LEN module was connected to the WRS3-LJ with either a 0.5m or 6.4km fiber optic cable.

The 1PPS outputs of the two WR-LEN modules was fed to the A & B inputs of a CNT-90 Timer/Counter/Analyser. This was set to measure the time difference between the rising edges of the two 1PPS signals (“Time A-B”). Since this relative measurement is between nodes, it is not necessary to discipline the (WRS3-LJ) to an external timing reference.




Results

Measurements were taken every second for 1 hour with the 0.5m fiber fitted and then repeated again with the 6.4km fiber fitted. No reconfiguration/restarting of the master or slaves was performed - the fiber optic cable for one WR-LEN was simply swapped.

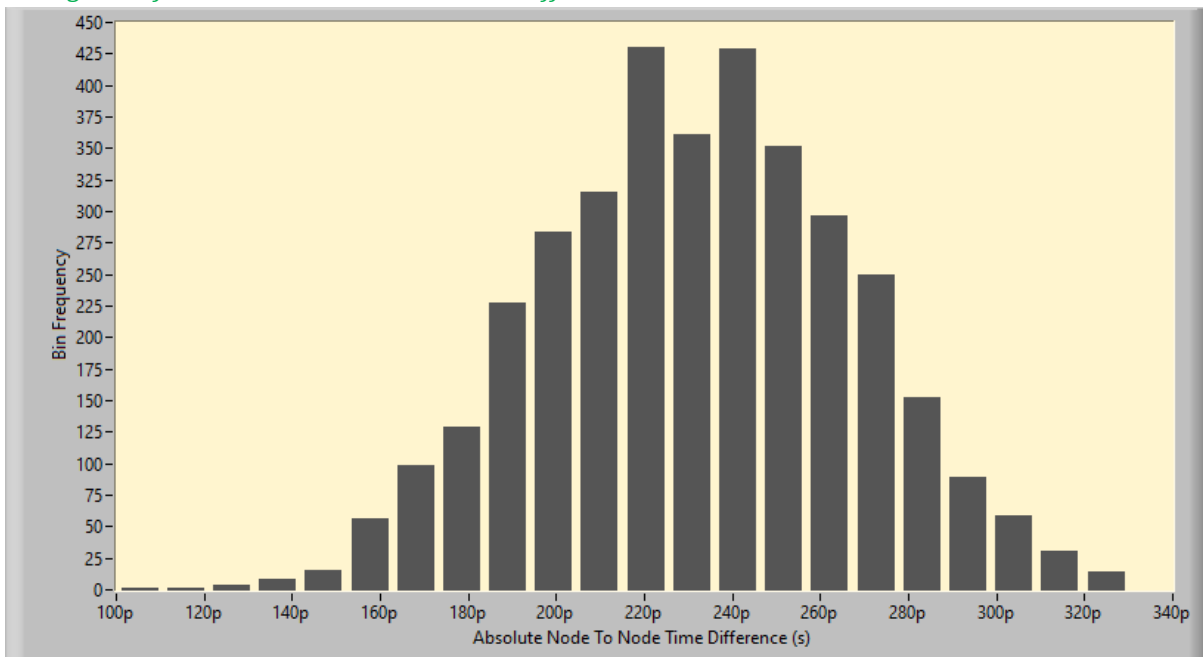
0.5m + 0.5m Fiber Optic

- Mean: 231.1ps
- Min/Max: 100ps/330ps
- Standard Deviation: 33.8ps

Absolute Note to Node Time Difference vs Sample




Histogram of Absolute Note to Node Time Difference



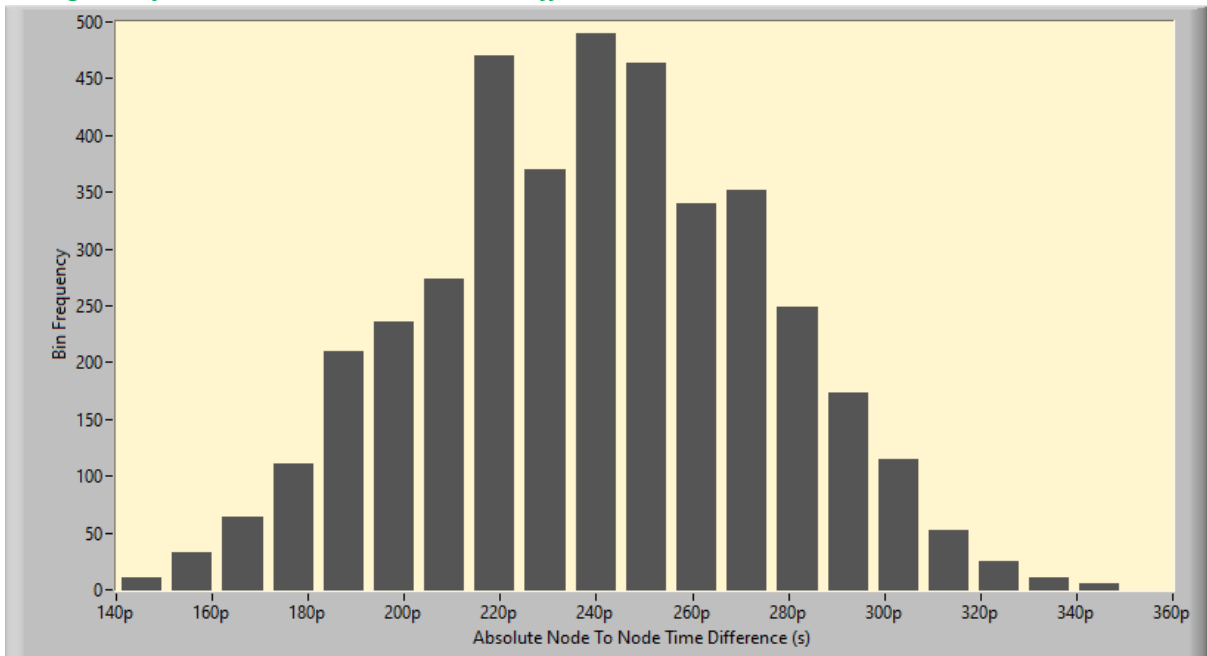
0.5m + 6400m Fiber Optic

- Mean: 239.7ps
- Min/Max: 140ps/350ps
- Standard Deviation: 34ps

Absolute Note to Node Time Difference vs Sample



Histogram of Absolute Note to Node Time Difference




Where is White Rabbit currently used?

The following is a brief outline of just a small number of current programmes & systems utilising White Rabbit technology.

CERN

White Rabbit was tailored specifically to counter the limitations of the existing CERN GMT controls infrastructure. In addition to the inherent benefits outlined above, White Rabbit also provided improved bandwidth, a more comprehensive set of diagnostic tools and integration of data, timing and control networks.

The Large Hadron Collider (LHC) is a truly enormous system-of-systems and so White Rabbit is being very gradually phased-in over a number of years to replace the GMT whilst minimising disruption.



In the meantime, White Rabbit has also found applications at CERN in:


- The [Real-time distribution of magnetic field data in the synchrotrons \(BTrain\)](#)
- The [LHC Instabilities trigger distribution system](#)
- The [Open Analogue Signal Information System \(OASIS\) Trigger Distribution](#)
- The [Low-Level RF System in the Super Proton Synchrotron](#)

Since White Rabbit was first used within CERN, the technology has also expanded its application outside the field of particle physics and is now deployed in many different industries, including telecommunications, financial markets, smart grids, space and quantum computing.

Below is a brief summary of a subset of the current deployments.

GSI – FAIR


GSI hosts the International Accelerator Facility (FAIR) where matter that usually only exists in the depths of space is produced for scientific research. The FAIR 1.1km circumference ring accelerator generates particle beams of both natural elements and antiprotons of unparalleled intensity and quality.



Again, the FAIR has a GMT system to trigger & synchronize equipment and software actions, timed according to the accelerator cycles. This has been designed from the ground up with White Rabbit technology to synchronize 2000+ nodes with up to 2km of fiber between end points.

KM3NeT

KM3NeT is a research infrastructure housing the next generation neutrino telescope. Located in the deepest seas of the Mediterranean, KM3NeT will open a new window in to the properties of the elusive neutrino particles. Once completed, the telescopes will have detector volumes of several cubic kilometres of clear sea water. Sites in France and Silly will together study astrophysical neutrino sources from GeV to PeV energies.



The arrays of sensors indirectly detect extra-terrestrial neutrinos by means of the Cherenkov radiation, which is induced during the path of charged relativistic particles through the sea water.

When a neutrino interacts with the telescope, surrounding matter produces a muon, which travels through the water at a higher speed than the light. Such particles generate a blue light pulse known as Cherenkov radiation. The arrival time of the light is detected and used for reconstruction of the muon trajectory, and, therefore, that of the original neutrino.

White Rabbit provides sub-nanosecond synchronization between the Digital Optical Modules - the main detection components of the system.

EISCAT

EISCAT is an international effort to conduct ionospheric and atmospheric measurements using the most sophisticated imaging radar ever built. This includes observing the effects of the aurora borealis. The system is distributed over three countries: Finland, Norway and Sweden, and all the facilities are located north of the Arctic circle.



The EISCAT3D phased array radar system is heavily reliant on a very accurate control of the time offset and phase synchronization across the array. The required accuracy is in the sub-nanosecond range within each site, and within microseconds between sites. The synchronization of clocks throughout the EISCAT3D network will be based on White Rabbit technology.

SKA

The Square Kilometre Array (SKA) project is an international effort to build the world's largest radio telescope, with eventually over a square kilometre (one million square metres) of signal collection area. The SKA project has a strict sub-2ns timing accuracy requirement between receiver systems that is not realisable using current synchronisation technologies such as NTP, PTP or per-node GPS. White Rabbit was therefore chosen to provide very accurate distribution of 1PPS across the array.

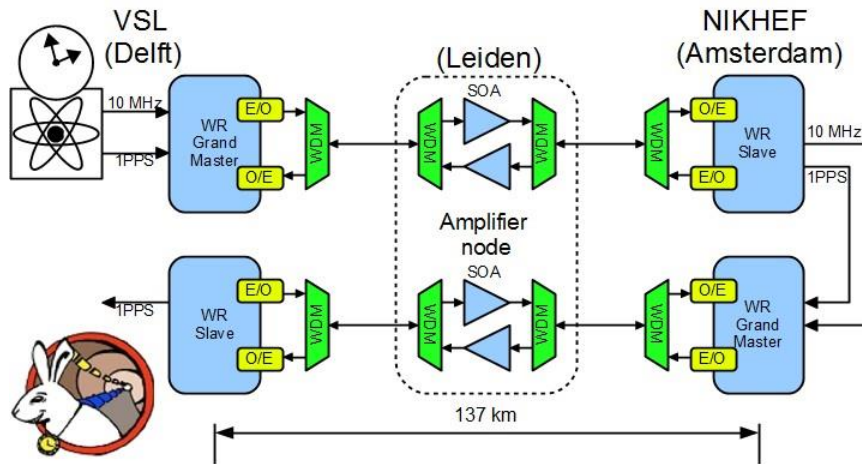


VSL

VSL is the Dutch National Metrology Institute and one of their many responsibilities is the dissemination of time (UTC) and frequency across the Netherlands.



For this purpose, White Rabbit links have been installed between the atomic clocks at VSL in Delft and the National Institute for Subatomic Physics (Nikhef) in Amsterdam using two 2x137km fibers. Their system strives for an uncertainty of well below 1 ns. This is a significant improvement over the previous GPS-based techniques which limited their accuracy to 5-8ns.




VTT MIKES

The Finland Centre of Metrology VTT Technical Research Centre were merged in 2015 to form VTT MIKES and operates as a national metrology institute.



VTT MIKES's customers include both Finnish and international companies as well as the public sector. VTT MIKES currently operates a number of nodes connected by 10 White Rabbit Switches for the dissemination of UTC and comparison of reference clocks. The long-term goal is to interconnect an ensemble of 10 atomic clocks distributed across the country.



A 950-km WR-PTP link has been realized between Espoo in the south and Kajaani in the middle of Finland. After calibration, the time transfer accuracy of this link (when compared against a GPS Precise Point Positioning reference) was found to be hold within $\pm 2\text{ns}$ over a three-month period.

Deutsche Börse

In a world far away from Large Experimental Physics programmes and national laboratories, White Rabbit has found a home at the German Stock Exchange (Deutsche Börse). Here co-located stock trading companies pay a handsome fee to receive the latest pricing information before their competitors. Putting this in to context, the fastest time from a new price being announced to the buy/sell requests arriving by return is 82ns (a 25m distance at the speed of light).



White Rabbit technology is used to provide tight synchronisation of customers' system clocks with the T7 trading system and precision time stamping of the arrival time of new pricing data for customer auditing.

White Rabbit is now seeing world-wide adoption within the financial trading sector.

Prospective Applications of White Rabbit

Here we discuss just a few examples of possible applications of White Rabbit.

Time & Frequency-as-a-Service

With the unquenchable thirst for faster and faster internet connection speeds, fiber optic is now being deployed all the way in to homes (FTTH) and business premises (FTTP).

Whilst White Rabbit is currently only targeted at FPGA-based applications, Application Specific ICs (ASICs) are being developed. Support for IEEE1588-2008 PTP has been commonplace for several years in Ethernet network switches, routers and adapters. With White Rabbit achieving standardisation through the IEEE1588-2019 PTP High Accuracy Profile, it is now hopefully just a matter of time before it becomes ubiquitous in core network infrastructure, perhaps heralding the arrival of Time & Frequency-as-a-Service available to all homes, businesses. This would allow critical infrastructure to move away from GNSS/GPS disciplined clocks that might otherwise be spoofed or jammed.

Distributed Test Equipment

White Rabbit can be utilised as a generic trigger, control and data transfer resource for Test & Measurement Equipment. By extending the LAN Extensions for Instrumentation (LXI) standard with precision timing functionality it can provide capabilities not otherwise possible with COTS equipment alone.

One example is a Distributed Oscilloscope where thousands of digitizers are spread over a large White Rabbit network, all synchronised to a common clock. A trigger event on one node can be timestamped and distributed all other nodes. The nodes would utilise their common notion of time to determine which sample in their pre-trigger buffer occurred at the same time as the broadcasted trigger event. They would then send their trigger-aligned sample buffer to a central processor for display and analysis. This is all possible thanks to White Rabbit's fixed upper bound on packet delivery.



CERN's OASIS is an example of this application of White Rabbit. The screenshot above shows aligned captures from across the Large Hadron Collider (LHC). The OASIS installation, as of 2021, consists of over 500 multiplexed digitisers, capable of digitising more than 5000 analogue channels with more than 250 acquisition triggers sources.

Summary

White Rabbit is set to transform timing and synchronisation over commodity optical networks. Through its standardisation we expect it to become ubiquitous in a vast number of markets across the globe.

Key Points

- It provides time & frequency with sub-nanosecond accuracy & picoseconds of precision.
- It is capable of delivering Cesium Atomic Clock levels of long-term stability to a system.
- It can be realised in a volume comparable to a traditional Double-Ovenised Crystal Oscillator.
- It can consume as little as 2W on the latest generation of FPGA devices.
- It requires only a modest increase in Bill-of-Material cost and FPGA resources.
- It is Open Hardware we can all utilise thanks to CERN, GSI, NikHEF et al.

Future Outlook

A White Rabbit Collaboration Organisation has been proposed as a method of steering future development, providing product certification and marketing White Rabbit to a much wider audience.

White Rabbit Development Update 2022/2023

- The White Rabbit Switch v4 continues development with schematics recently reviewed. New 1Gb/s model expected late 2022 with 10Gb/s gateway upgrades expected 2023/2024.
- The White Rabbit PTP Core v5.0 update is currently underway.
- WRPC Software is being reorganised.
- A new UART-based bootloader is being developed.
- The PTP Ported to Silicon (PPSi) is being updated to improve compliance with IEEE1588.
- New FPGAs & multi-gigabit transceivers are being added to the supported families list.
- Migration of the WRPC soft-core from Lattice Mico32 to a RISC-V based solution has begun. This has a more modern GCC-based toolchain, should make debugging easier and could yield a smaller compiled binary size.
- Support is being added for programmable VCXOs.
- A new Absolute Calibration (AbsCal) technique has been developed providing an alternative to the current “golden calibrator” method.
- The Low Phase Drift Calibration (LPDC) has been developed to mitigate the imprecision of the “bit-slide” measurement ($\sim\pm 35\text{ps}$ for GTX transceivers) which limited the accuracy to no better than 100ps. LPDC allows the offset of the master and slave to be reduced to $\pm 10\text{ps}$ after re-establishing a link either due to re-connection or power cycling.

Who do we have to thank for developing White Rabbit?

The prominent members of CERN's White Rabbit development team at (in 2017) are pictured below:



Left to Right: Grzegorz Daniluk, Javier Serrano, Tomasz Wlostowski, Maciej Lipinski, Adam Wujek, Dimitris Lampridis

Significant contributions continue to be made by other groups including commercial partners, GSI (Darmstadt) and NikHEF (Amsterdam) both in development of hardware & software and in the support of vendors & users around the world.

Where to go for more information?

The following links will provide you a starting point to understanding the White Rabbit technology, its implementation and its application both inside and outside Large Experimental Physics programmes:

- M. Lipinski, T. Włostowski, J. Serrano, P. Alvarez- [White Rabbit: A PTP Application for Robust Sub-nanosecond Synchronization](#)
- P.P.M. Jansweijer, H.Z. Peek - [Measuring propagation delay over a 1.25 Gbps bidirectional data link](#)
- Tomasz Włostowski MSc Thesis - [Precise time and frequency transfer in a White Rabbit network](#)
- Grzegorz Daniluk MSc Thesis - [White Rabbit PTP Core the sub-nanosecond time synchronization over Ethernet](#)
- M. Lipinski - [White Rabbit: a next generation synchronization and control network for large distributed systems](#)
- Grzegorz Daniluk - [White Rabbit PTP Core Overview](#)
- OHWR - White Rabbit Project - <https://ohwr.org/project/white-rabbit/wikis/home>
- OHWR - White Rabbit PTP Core - <https://ohwr.org/project/wr-cores/wikis/wrpc-core>
- OHWR - Software for White Rabbit PTP Core - <https://ohwr.org/project/wrpc-sw/wikis/home>
- OHWR - White Rabbit Calibration - <https://ohwr.org/projects/white-rabbit/wiki/Calibration>
- OHWR - White Rabbit IEEE1588-2019 - <https://ohwr.org/projects/wr-std/wiki/wrin1588>

- White Rabbit Development on Discourse - <https://forums.ohwr.org/c/white-rabbit-dev>
- M. Rizzi, M. Lipinski, T. Wlostowski, J. Serrano, G. Daniluk, P. Ferrari, S. Rinaldi - [White Rabbit Clock Characteristics](#)
- H. Li, G. Gong, W. Pan, Q. Du, J. Li - [Temperature Effect and Correction Method of White Rabbit Timing Link](#)
- M. Kreider, R. Bar, D. Beck, W. Terpstra, J. Davies, V. Grout, J. Lewis, J. Serrano, and T. Wlostowski – [Etherbone: Open borders for system-on-a-chip buses: A wire format for connecting large physics controls](#)
- M. Kreider – [Etherbone: A Network Layer for the Wishbone SoC Bus](#)
- D. Lampridis, T. Gingold, D. Michalik, T. Pereira da Silva, A. Poscia, M. H. Serans, M. R. Shukla - [Renovation of the Trigger Distribution in CERN's Open Analogue Signal Information System \(OASIS\) using White Rabbit](#)
- The Open Hardware Repository - <https://ohwr.org/explore/projects>
- Open Hardware at CERN - <http://cds.cern.ch/record/2109248/files/CERN-Brochure-2015-002-Eng.pdf>
- ITU Standards - <https://www.itu.int/rec/T-REC-G/en>

Credits

Thank you to the following authors & organisations for sharing their images, papers & presentations. This document would be awfully bland without your generosity!

- T. Włostowski - Precise time and frequency transfer in a White Rabbit network.
- P.P.M. Jansweijer, H.Z. Peek - Measuring propagation delay over a 1.25 Gbps bidirectional data link.
- E.Gousiou - The White Rabbit Network.
- M. Kreider, R. Bar, D. Beck, W. Terpstra, J. Davies, V. Grout, J. Lewis, J. Serrano, and T. Wlostowski - Open borders for system-on-a-chip buses: A wire format for connecting large physics controls.
- Media Centre - CERN (<https://home.web.cern.ch/resources>)
- Media Centre - GSI (https://www.gsi.de/en/press/media_center)
- Media Centre - KM3NeT (<https://www.km3net.org/pictures-and-videos>)
- Media Centre - EISCAT (<https://eiscat.se/press>)
- Media Centre - SKA (<https://www.skatelescope.org/quickimages>)
- Media Centre - VSL (<https://www.vsl.nl/en>)
- Media Centre – VTT MIKES (<https://www.vttresearch.com/en/research-expertise/metrology-vtt-mikes>)
- Media Centre - Deutsche Börse (<https://www.deutsche-boerse.com/dbg-en/media/media-database/photo-database>)