UNIVERSIDAD
DE GRANADA

Programa de Doctorado
en Tecnologías de la Información y la Comunicación

DISTRIBUCIÓN ROBUSTA DE TIEMPO
MEDIANTE REDES DE FIBRA ÓPTICA DE LARGA DISTANCIA

Robust Time Distribution
Through Long-Distance Optical Fiber Networks

Memoria de Tesis doctoral

Autor
José López Jiménez

Directores
Manuel Rodríguez Álvarez
Antonio Javier Díaz Alonso

Departamento de Ingeniería de Computadores, Automática y Robótica
Escuela Técnica Superior de Ing. Informática y de Telecomunicación

—

Granada, 2023

Editor: Universidad de Granada. Tesis Doctorales
Autor: José López Jiménez
ISBN: 978-84-1195-032-9
URI: https://hdl.handle.net/10481/84683

Declaración

El doctorando / *The doctoral candidate*
José López Jiménez

y los directores de la tesis / *and the thesis supervisors*
Manuel Rodríguez Álvarez / Antonio Javier Díaz Alonso

garantizamos, al firmar esta tesis doctoral, que el trabajo ha sido realizado por el doctorando bajo la dirección de los directores de la tesis y hasta donde nuestro conocimiento alcanza, en la realización del trabajo, se han respetado los derechos de otros autores a ser citados, cuando se han utilizado sus resultados o publicaciones.

guarantee, by signing this doctoral thesis, that the work has been done by the doctoral candidate under the direction of the thesis supervisors and, as far as our knowledge reaches, in the performance of the work, the rights of other authors to be cited (when their results or publications have been used) have been respected.

Granada, a 20 de junio de 2023.

Directores de la tesis *Thesis Supervisors*	Doctorando *Doctoral Candidate*
Manuel Rodríguez Álvarez Antonio Javier Díaz Alonso	José López Jiménez

A mi madre,
a mi familia, por todo su apoyo
y a los que querrían estar conmigo hoy.

Distribución robusta de tiempo mediante redes de fibra óptica de larga distancia

José López Jiménez

Palabras clave: Sistemas de sincronización de tiempo, White Rabbit, Resiliencia, Holdover, Error de tiempo, Fiabilidad, Redundancia, Sincronización de tiempo, IEEE 1588, Diseminación de tiempo, Redes de larga distancia, Time as a Service.

Resumen

La creciente dependencia de diversos sectores, desde las telecomunicaciones hasta los mercados financieros, en la sincronización precisa resalta la necesidad de sistemas de transferencia de tiempo robustos. Los sistemas distribuidos, que constituyen la espina dorsal de nuestro mundo digital, requieren una base de tiempo común para hacer posibles la correlación de datos, el monitoreo del estado del sistema en su totalidad y la ejecución de procesos coordinados dentro de intervalos de tiempo acotados. Si bien los Sistemas Globales de Navegación por Satélite (GNSS) han proporcionado referencias de tiempo válidas a un coste asequible, su vulnerabilidad hace necesaria la integración de alternativas terrestres más resilientes.

Este trabajo se centra en el uso del protocolo White Rabbit (WR), en el que se basa el perfil de IEEE 1588-2019 *High Accuracy*, para ofrecer capacidades de sincronización en enlaces de larga distancia. Inicialmente diseñado para su uso en redes privadas con modelos de retardo de enlace conocidos y capaz de alcanzar una precisión por debajo de un nanosegundo, existen desafíos significativos que deben abordarse con respecto a la extensión de WR en redes públicas que abarcan distancias más largas, condiciones de red dinámicas y equipamiento específico para larga distancia. Para garantizar la fiabilidad del sistema de transferencia de tiempo, el uso del protocolo WR en estas circunstancias debe ir acompañado de capacidades de fiabilidad mejoradas que permitan a los dispositivos de la red funcionar de forma autónoma o cambiar rápidamente de una referencia de tiempo a otra.

Esta tesis está estructurada en cuatro partes. La primera parte está dedicada a la revisión del estado del arte, donde examinamos las fuerzas impulsoras que llevan a muchas áreas de la industria a adoptar sistemas de transferencia de tiempo precisos y resilientes, tanto regulatorias como competitivas. Luego, revisamos los conceptos fundamentales de las señales de tiempo y la sincronización, junto a los métodos más utilizados y los más avanzados para transferencia de tiempo entre ubicaciones lejanas,

y nos enfocamos en los elementos que afectan la robustez y resiliencia de WR en redes de larga distancia, es decir, la transmisión de señales de datos a través del medio óptico, y los mecanismos que pueden implementarse para mejorar la redundancia y resiliencia en un sistema de transferencia de tiempo. Esto es seguido por la introducción de las metodologías experimentales, herramientas y plataformas de desarrollo que se han empleado.

En la segunda parte de la tesis abordamos los aspectos de resiliencia de la solución de transferencia de tiempo tanto a nivel de nodo como de red. Primero, describimos y abordamos las limitaciones introducidas por el ruido de señal, analizamos su impacto en el mantenimiento, e implementamos una solución experimental de *holdover* que extiende el tiempo de *holdover* de un nodo WR garantizando una exactitud por debajo de un nanosegundo durante el tiempo suficiente para realizar el cambio a una referencia de tiempo alternativa. A esto le sigue la descripción de los métodos desarrollados para proporcionar a los nodos WR mecanismos de tolerancia a fallos en topologías en anillo, demostrando la viabilidad del enfoque tanto para datos como para sincronización.

La tercera parte consiste en un análisis extensivo de los dispositivos de red específicos y fenómenos físicos que causan un impacto en el rendimiento de la sincronización vinculada a enlaces de larga distancia. Los elementos identificados se implementan en un banco de pruebas de laboratorio, y se evalúa la influencia de los elementos identificados en términos de asimetría y retardo frente a condiciones ambientales variables. Este análisis resalta la importancia de evaluar los componentes de la red en la ruta óptica y modela su impacto.

Y la cuarta y última parte corresponde a la prueba y verificación de los mecanismos de resiliencia en un escenario del mundo real. Después de una visión general de las necesidades de sincronización en la industria de tecnología financiera y la descripción de los requisitos técnicos y regulatorios, se delinea un enlace de *Time as a Service* (Taas, o "tiempo como servicio") que abarca 44 km de longitud en el área metropolitana de Madrid, y se destacan y verifican sus mecanismos de autonomía y resiliencia. Los resultados de este escenario, donde se garantiza la trazabilidad de la señal de tiempo al Tiempo Universal Coordinado (UTC) y se pone a prueba la robustez del sistema, sirven como validación de los modelos y mecanismos descritos en las partes anteriores.

Robust time distribution through long-distance optical fiber networks

José López Jiménez

Keywords: Time Transfer Systems, White Rabbit, Resiliency, Holdover, Time Error, Reliability, Redundancy Time Sychronization, IEEE 1588, Timekeeping, Time Dissemination, Long-distance networks, Time as a Service.

Abstract

The increasing reliance of various sectors, from telecommunications to financial markets, on precise synchronization highlights the need of robust time transfer systems. Distributed systems, the backbone of our digital world, require a shared common time base to enable data correlation, system-wide state monitoring, and executing coordinated processes within bounded time intervals. While Global Navigation Satellite Systems (GNSS) has provided valid and cost-effective time references, their vulnerability necessitates the integration of more resilient terrestrial alternatives. This work centers around the use of the White Rabbit (WR) protocol, on which the IEEE 1588-2019 High Accuracy profile is based, to offer synchronization capabilities over long distance links. Initially developed for use in private networks with well-known link models and capable of sub-nanosecond accuracy, there are significant challenges that need to be addressed regarding the extension of WR over public networks spanning longer distances, dynamic network conditions and specific long-distance equipment. To ensure the reliability of the time transfer system, the use of the WR protocol in these circumstances must be followed by enhanced reliability capabilities that allow the devices in the network to function autonomously or to swiftly shift from one time reference to another.

This thesis is structured in four parts. The first part is devoted to the review of the state of the art, where we examine the driving forces that push many areas of the industry to adopt precise and resilient time transfer systems, both regulatory and competitive. Then, we review the fundamental concepts of time signals, the most widely used and pioneering methods to transfer a time reference between locations, and we focus on the elements that impact the robustness and resiliency of WR over long distance networks, that is, the transmission of data signals over the optical medium, and the mechanisms that can be implemented to enhance redundancy and resiliency in a time transfer system. This is followed by the introduction of the experimental methodologies, tools and development platforms that have been employed.

In the second part of the thesis we address the resiliency aspects of the time transfer solution at both the node and network levels. First, we describe and address the limitations introduced by signal noise, analyze its impact on holdover, and we implement an experimental holdover solution that extends the holdover time of a WR node ensuring sub-nanosecond performance for long enough to perform switchover to an alternate time reference. This is followed by the description of the methods developed to provide WR nodes with fault-tolerant mechanisms in ring topologies, demonstrating the viability of the approach for both data and timing.

The third part consists of an extensive analysis of the specific network equipment and processes that cause an impact on the performance of synchronization that is bound to long distance links. The identified elements are deployed in a lab testbench, and the influence of the identified elements in terms of asymmetry and delay is tested against variable environmental conditions. This analysis highlights the importance of evaluating the network components in the optical path, and models the impact thereof.

And the fourth and last part corresponds to the test and verification of the resiliency mechanisms in a real world scenario. After an overview of timing in the fintech industry and the outline of the technical requirements, a Time as a Service (TaaS) link spanning 44 km in the metropolitan area of Madrid is described, and its autonomy and resiliency mechanisms are highlighted and verified. The results of this scenario, where the traceability of the time signal to Coordinated Universal Time (UTC) is guaranteed and the robustness of the system is tested out, serve as validation of the models and mechanisms described in the previous parts.

Agradecimientos

Todos los que me conozcan sabrán que estoy pasando un mal rato enfrentándome a escribir estos párrafos. Así que mis disculpas si he pasado por alto algún nombre. Aunque no aparezca aquí, os llevo en el corazoncito.

En primer lugar, quiero expresar mi gratitud hacia mis directores, Javier y Manolo por su valiosa orientación y por todas las ganas que me han insuflado a lo largo de estos años. Su experiencia han sido de inmenso valor para este trabajo.

También me gustaría agradecer a todos aquellos con los que he compartido no solo espacio sino muchos momentos en el CITIC. Aprendí mucho durante los años que estuve allí. Sois muchos para nombraros uno por uno, así que aprovecharé este espacio exclusivamente para pedir disculpas si el ruido de los ventiladores de los equipos ha causado pérdidas de audición en alguno de vosotros. Son daños colaterales de *hacer la ciencia*. Gracias específicamente a Fran, por custodiar mi taza durante mi ausencia y que algún día recogeré. Gracias por las *masterclass*, y por las feroces críticas a los Peugeot 308 de varias generaciones, al Jeep Renegade y al Astra GSi.

Por supuesto, en el lado de Seven/Orolia/Safran también tengo muchísimos motivos por los que agradecer a muchas personas. Desde Trini, que me acogió en los primeros meses de hacer pruebas con el *horno chino*, a toda la gente con la que trabajo actualmente y que hacen más agradables las horas que echamos trabajando. Se difuminan aquí un poco los agradecimientos del CITIC con los de Seven, pero debo mencionar, por supuesto, a José Luis, Miguel, Emilio, Rafa, Jorge Machado, Mario, Paco... con los que también llevo años aprendiendo (y sufriendo) y por supuesto a todos aquellos con los que he compartido algunos pacharanes en la Posada.

En el lado más personal, gracias nuevamente a mi madre y mi hermana no solo por haberme apoyado en todo, sino por haberme aguantado en algunas ocasiones. Y gracias (quizás disculpas sea un término más cercano a la realidad) especialmente a

Isaac por estos últimos meses de reclusión.

Y volviendo a un plano más general, me gustaría agradecer a toda la gente más lejana que he encontrado por el camino con muchas ganas de compartir conocimientos de forma entusiasta. Me refiero, sobre todo, a personas con las que he compartido tiempo en proyectos (de Nikhef, NPL, OP, ...) que claramente disfrutan con lo que hacen y lo contagian a los demás. Igualmente, gracias a la comunidad del OHWR por todo el trabajo del que yo me he beneficiado como punto de partida para este trabajo.

En conclusión, estoy seguro de que sin todos los que he mencionado en estas líneas yo no estaría ahora mismo finalizando este trabajo. Gracias por todo.

Jose.

Contents

1. Introduction **1**

1.1. Motivation 1
1.2. Objectives 5
1.3. Framework 7
1.3.1. CLONETS and CLONETS-DS 7
1.3.2. TOWR 8
1.3.3. AMIGA 8
1.3.4. WRITE 10
1.3.5. Doctorado Industrial 10
1.4. Chapter Structure 11

2. State of the Art **13**

2.1. Introduction to time transfer systems 14
2.2. Driving forces for reliability and redundancy in time transfer 16
2.2.1. Requirements for reliability and redundancy in the industry 17
2.2.1.1. Telecommunication networks 17
2.2.1.2. Finance 19
2.2.1.3. Power industry 20
2.3. Common ground: definitions and terminology 21
2.3.1. Noise model of a time signal 22
2.3.2. General definitions 22
2.3.3. Noise characterization 24
2.3.4. Types of oscillators 29
2.4. Time transfer: principles, techniques and protocols 31
2.4.1. One-way time transfer 31

2.4.2. Two-way time transfer 32
2.4.3. Techniques and protocols 33
2.4.3.1. Satellite-based techniques 34
2.4.3.2. Analog signals 38
2.4.3.3. Optical fiber links 39
2.4.3.4. Network protocols 39
2.4.3.5. Introduction to the White Rabbit protocol 44
2.4.4. White Rabbit over long distance 48
2.5. Introduction to fiber optic transmission in long-haul networks . . . 49
2.5.1. Limiting factors in fiber optic transmission 53
2.5.1.1. Attenuation 53
2.5.1.2. Chromatic dispersion 54
2.5.1.3. Polarization Mode Dispersion 56
2.5.1.4. Nonlinear effects 57
2.5.1.5. Sagnac effect 58
2.6. Mechanisms and methods involved in redundancy and resiliency . . 58
2.6.1. Redundant networks 58
2.6.2. Holdover 61
2.6.3. Switchover and failover 61

3. Materials, methods and tools 63
3.1. Hardware platform: the White Rabbit Switch 63
3.1.1. Software 64
3.1.2. Hardware and FPGA architecture 66
3.1.3. White Rabbit Switch - Low Jitter & Expansion Board 67
3.2. Development environment 68
3.3. Instrumentation and other materials 71
3.3.1. Long distance testbench 71
3.3.2. Laboratory instrumentation 73
3.3.3. Clock references 74
3.3.4. Software tools 74
3.4. Experiment design and methods 74
3.4.1. Experimental design Overview 76
3.5. Conclusion 81

4. Mechanisms for Improved Resiliency in the White Rabbit Switch: stability, holdover, and redundant capabilities 83
4.1. Introduction . . . 83
4.2. Background . . . 84
4.2.1. Synchronization and Syntonisation in White Rabbit . . . 84
4.2.2. White Rabbit syntonization (DDMTD and SoftPLL) . . . 85
4.3. Noise improvements in the WRS-LJ . . . 88
4.4. WRITE expansion holdover board . . . 90
4.5. The development of a holdover mechanism . . . 92
4.5.1. Factors affecting holdover performance . . . 95
4.5.2. Types of mechanisms to improve holdover . . . 96
4.5.3. WRS-LJ Experimental Holdover Implementation . . . 99
4.5.4. Results . . . 104
4.5.4.1. Abracon AOCJY-10.000 . . . 104
4.5.4.2. Morion MV341 . . . 112
4.6. Redundant timing system with low-latency data capabilities . . . 118
4.6.1. Redundancy protocols and White Rabbit . . . 119
4.6.1.1. PTP clock types in WR-HSR . . . 120
4.6.1.2. Implementation of WR for redundant ring networks 121
4.6.1.3. Switchover in the WR ring network . . . 122
4.6.2. Data implementation of HSR . . . 123
4.6.2.1. HSR Overview . . . 124
4.6.2.2. Link Redundancy Entity (LRE) . . . 125
4.6.2.3. Fast Switchover Unit (FSU) . . . 128
4.6.3. Results . . . 128
4.6.3.1. Switchover timing results . . . 129
4.6.3.2. Data resiliency results . . . 131
4.7. Conclusion . . . 132

5. Analysis and modeling of existing error and asymmetries in long distance network equipment 135
5.1. Optical network elements for long-distance links . . . 137
5.1.1. Equipment transceivers . . . 138
5.1.2. Optical transponder . . . 141
5.1.3. Optical routing and switching . . . 141
5.1.4. Amplification and regeneration . . . 144

5.1.5. Dispersion compensation 146
5.2. Temperature dependence and time error in unidirectional long distance WR links 147
5.2.1. Precision of CWDM/DWDM wavelength reproducibility . . 149
5.2.2. Reference setup: 50 km unidirectional link 150
5.2.3. Impact of the WR devices 153
5.2.4. Impact of DWDM MUX/DEMUX modules 153
5.2.5. Impact of EDFAs 154
5.2.6. Impact of DCMs 155
5.2.7. Summary 158
5.3. Conclusion 158

6. Time as a Service for finance based on WR **161**
6.1. Introduction 161
6.1.1. Overview of timing in the fintech industry 162
6.1.2. Technical requirements of a financial timing service 165
6.1.3. Time as a Service 167
6.2. TOWR project 168
6.2.1. Link overview 169
6.2.2. Time server hardware 169
6.2.3. Network time dissemination 170
6.2.4. Customer side and Calibration 173
6.2.5. Resiliency features 174
6.2.6. Service monitoring 178
6.3. Conclusion 179

7. Conclusion **181**
7.1. Summary of findings 181
7.2. Main contributions 185
7.3. Publications 187
7.3.1. Journal contributions 187
7.3.2. Conferences 188
7.4. Future work 188

A. Introducción **193**
A.1. Motivación 193
A.2. Objetivos 197

A.3. Marco de trabajo 199
A.3.1. CLONETS y CLONETS-DS 200
A.3.2. TOWR 200
A.3.3. AMIGA 201
A.3.4. WRITE 202
A.3.5. Doctorado Industrial 203
A.4. Estructura de los capítulos 204

B. Conclusiones **207**
B.1. Resumen de los resultados 207
B.2. Contribuciones principales 211
B.3. Publicaciones 213
B.3.1. Publicaciones en revistas 213
B.3.2. Publicaciones en conferencias 214
B.4. Trabajo futuro 214

References **217**

List of Figures

1.1. Graphic summary of the contents of this thesis. 6

2.1. Schematic view of an example time transfer system. Timescales are generated by NMIs and other research institutions, which then are disseminated using methods like GNSS. A time server provides a network grandmaster with this timescale reference, which is distributed to other devices in the network. 15
2.2. Relationship between accuracy and stability. From [35] 23
2.3. Examples of different power-law noises contained in a signal. From [34]. 25
2.4. Different noise processes are characterized by their slopes in the ADEV plot. A modified version of the ADEV is necessary to discriminate the white PM noise from flicker PM noise. 27
2.5. Visual representation of the MTIE of a time signal. 28
2.6. Stability performance of the classical microwave atomic frequency standards compared with temperature-compensated (TCXOs) and oven-controlled (OCXOs) crystal oscillators. Reprinted with permission from [41] 30
2.7. One way time transfer. 32
2.8. Two way time transfer. 33
2.9. Principle of GNSS common view time transfer. When the baseline between the stations A and B is short, the ionospheric and tropospheric errors are highly correlated. 35
2.10. TWSTFT using a geostationary communications satellite. From [47] 36
2.11. PTP message exchange. 43

2.12. Schematic depiction of a WR network showcasing the multiple roles and connections of the same device in different locations of the topology. The red and blue arrows represent the duality of the control data (in red) and the WR time network (in blue). From [67]. 46

2.13. White Rabbit Link Model. From [71]. 47

2.14. Fixed DWDM spacing grid vs. flexible DWDM grid. The dynamic nature of the flexible grid allows the network management to optimize the system capacity, allocating no more than the necessary bandwith for each application. 52

2.15. Relation between transmission wavelength and fiber attenuation. The three transmission widnwos refer to the wavelengths at which data can be transmitted more efficiently. The 850 nm window is typically used for short-distance communications. The second window allows for longer distance transmission, but the 1550 nm window is the typical choice for long-haul systems due to its lower attenuation and compatibility with common amplification technologies. From [85]. 54

2.16. Depiction of pulse broadening. The energy of the pulse is spread over time with longer distances, impacting on the capability of the receiver to tell symbols apart. From [86]. 56

2.17. Relation of chromatic dispersion with wavelength for the most common fiber types. From [86]. 57

3.1. The White Rabbit Switch. 64

3.2. High level diagram of the WRS. 65

3.3. Software architecture of the WRS. From [98] 66

3.4. Detail of the SCB and holdover expansion board of a WRS-LJ with its top cover open and front casing removed. The main features are highlighted. 69

3.5. Development platform and stages of the development methodology. The core components of the WR firmware are highlighted (Xilinx Virtex 6 FPGA, an LM32 softCPU and a Buildroot embedded Linux environment). The figure also depicts an schematic view of the iterative development process from system design to functional verification. 72

3.6. Representation of elements of the long distance testbench, laboratory instrumentation and software tools used in this study. In the top half, the long distance testbench, whose components are detailed in Chapter 5, is complemented by the use of the PHM frequency reference, as well as the climate chamber that allows the modification of environmental conditions. In the bottom half, the lab instrumentation is represented, along with screenshots and logos of the main software tools used to process the data obtained from instrumentation. 75
3.7. This diagram shows a schematic view of a setup in which time synchronization data can be recorded using a time and frequency counter. 78
3.8. Schematic diagram of a experimental setup evaluating the performance of a long distance WR-based time service, assisted by a GNSS time reference. 80

4.1. Overview of the WR phase-locked loop design. From [69] 86
4.2. Experimental setup for determining the additive phase noise introduced by the WRS and the WRS-LJ 91
4.3. Comparison of phase noise levels for a Grandmaster and a Slave in both standard WRS and WRS-LJ. 91
4.4. The WRITE Holdover expansion board with a surface-soldered Abracon AOCJY-10.000 OCXO. 93
4.5. Illustration of time error evolution in different scenarios. 98
4.6. Schematic view of the circular buffers that are used to store the latest n measurements and outputs of SoftPLL history. 101
4.7. Diagram showing the most relevant functions of the main SoftPLL control loop, including the holdover implementation modifications. 102
4.8. Experimental setup for the evaluation of holdover performance. The device under test is the slave WRS-LJ, that includes the expansion holdover board. 105
4.9. Time error in holdover mode using the Abracon OCXO during the first 60 seconds. Rates 1:N can be observed in the legend of the figure. rxxxxx represents rate 1:x, while r00000 represents an additional run of rate 1:1 106
4.10. Time error in holdover mode using the Abracon OCXO during the first 60 seconds. Data shown represents the medium sized time windows (from 1:64 to 1:2048) 107

4.11. Time error in holdover mode using the Abracon OCXO during the first 10 minutes. The run with rate 1:4096 shows a sudden phase shift after 9.5 minutes that has been likely introduced during the postprocessing stage. For the sake of transparency, this run is included as observed. 107
4.12. Time error in holdover mode using the Abracon OCXO during the first 10 minutes. The only data represented are the medium-sized history windows. 108
4.13. Time error in holdover mode using the Abracon OCXO up to 90 minutes. 109
4.14. MTIE figures for all the Abracon AOCJY-10.000 holdover runs. . . 110
4.15. Evolution of the TE for the MV341 in holdover mode during the first 60 seconds since launch. 114
4.16. Evolution of the TE for the MV341 in holdover mode during the first 10 minutes since launch. 114
4.17. Closeup of the TE evolution during the first minutes of runs for the Morion MV341 OCXO. Although almost unnoticeable in the previous figure, there are random frequency fluctuations that change the instant rate of frequency offset. 116
4.18. Evolution of three runs of the experiment that were run for longer periods (between 4.5 and 12 hours). 116
4.19. Calculated MTIE for all the holdover runs of the Morion MV341 OCXO. 117
4.20. HSR network operation with multicast traffic. From Wikipedia, uploaded by H. Kirrmann, 2017, licensed GPL. 124
4.21. HSR frame format including the 6-byte HSR tag. From Wikipedia, uploaded by H. Kirrmann, 2017, licensed GPL. 125
4.22. Diagram depicting the Link Redundancy Entity and its relation with its interfaces in the WRS FPGA project. From [121]. 126
4.23. Experimental setup for the redundant topologies experiment. The ring is comprised of one grandmaster and six doubly-attached slaves. From [121]. 129
4.24. PPS offset between Grandmaster and each of the slaves during a worst case switchover scenario. From [121] 130

5.1. WR links over long distance without need for amplification. Courtesy of Safran Electronics and Defense Spain. 139

5.2. Manufacturer-provided pictures of some of the common elements in DWDM systems for long-haul operation. From Fiberstore. 140
5.3. Artist's impression of the 5 km diameter central core of SKA antennas. From SKA Project Development Office and Swinburne Astronomy Productions, CC-BY 3.0. 142
5.4. Schematic view of a DWDM bidirectional link. The output of the DWDM is designed for unidirectional fibers but the optical filters recombine and segregate the desired wavelengths 142
5.5. Principle of operation of diffraction gratings applied to signal routing. From Optically Multiplexed Systems: Wavelength Division Multiplexing. Creative Commons attribution 143
5.6. Bypassing unidirectional amplifiers using OADMs to redirect the frequency signal over custom bidirectional amplifiers in a metrological experiment. From [9]. 144
5.7. Schematic view of a WR link between two points in the metropolitan area of New York City and Chicago, using WR devices as intermediate regeneration stations. Courtesy of Safran Electronics and Defense Spain. 147
5.8. Long distance testbench. Three WRS are shown on top of the rack, and MUX/DEMUX devices, OADM, EDFAs and DCMs follow from top to bottom. Spools of fiber are contained in the grey metal cases, and the device in the right side corner is a climate chamber that allows us to establish a precise temperature of operation. 148
5.9. Schematic view of the reference setup of the long distance testbench. 151
5.10. Variation of the measured RTT and the 1-PPS output of the reference setup. Data spans roughly 48 hours. 151
5.11. Schematic view of the thermal analysis of the WRS. 153
5.12. Schematic view of the thermal analysis of the DWDM MUX/DEMUX. 155
5.13. Schematic view of the thermal analysis of the EDFA. 156
5.14. Results of the thermal analysis of the EDFA in terms of RTT. The changes of temperature in the FPGA are correlated to the observed changes in the RTT. 156
5.15. Detailed zoom-in of results of the thermal analysis of the EDFA in terms of RTT. 157
5.16. Schematic view of the thermal analysis of the DCM. 157
5.17. Results of the thermal analysis of the DCM in terms of time error. . 158

6.1. Order book depth chart from a currency exchange. The x-axis reflects the unit price, while the y-axis depicts the cumulative order depth. Bids are shown in green, while asks are shown in red. From "Order book depth chart", Wikimedia Commons, CC-SA-AT 164
6.2. Map showing the time source generation site, the customer datacenter site, and the link path. 170
6.3. Schematic view of the time server hardware, demonstrating the redundant elements. Courtesy of Safran Electronics and Defense Spain. 170
6.4. Overview of the TOWR link architecture. Both time servers provide a WR grandmaster with an independently generated time reference. A WR-Z16 acts as boundary clock, and selects the time reference that is transmitted through the network. The client side is continuously monitored and compared to a DOWR GNSS receiver. Courtesy of Safran Electronics and Defense Spain. 172
6.5. Offset between the customer node and the monitoring DOWR over 25 summer days. The 24-hour periodicity is due to the variable ionospheric delays introduced by the GNSS receiver. Courtesy of Safran Electronics and Defense Spain. 175
6.6. Time output at the customer site during a period of succesive time server failover events. Spikes in the 1-PPS signal represent lack of 1-PPS signal during reconfiguration. Holdover capabilities are disabled. Courtesy of Safran Electronics and Defense Spain. 176
6.7. Comparison of five holdover executions of TOWR holdover in different weekends. The maximum expected error of 1.5 microseconds after 24 hours is highlighted in red. Courtesy of Safran Electronics and Defense Spain. 177
6.8. Daily time error reported for the month of December. Every Saturday the link is automatically disabled to evaluate the performance of the holdover solution. The limits of the time error established in the SLA (a maximum of 15 seconds error from initial calibration) are shown as horizontal red bars. Holidays and weekends when the market is closed are shown in grey. Courtesy of Safran Electronics and Defense Spain. 177

A.1. Resumen gráfico del contenido de este trabajo. 199

List of Tables

2.1. Characterization of noise processes according to their power-law components [38]. 25
2.2. PTP parameters in use by the BMCA. The specification of values for Clock Class and Clock Accuracy can be found in Table 4 and Table 5 of the PTP standard. 44

4.1. RMS Jitter results for WRS and WRS-LJ 90
4.2. Compatible oscillators with the WRITE Holdover expansion board and main stability characteristics 94
4.3. Examples of total buffer length with several undersampling parameters. 101
4.4. Summary of time elapsed to reach 1, 10, and 100 ns TE thresholds for all the analyzed runs. The color gradient of each cell is relative to the best and worst case in each column. The 100 ns threshold in the 1:2048 run was never exceeded during the length of the measurement. 111
4.5. Summary of the statistic parameters of the Morion MV341 holdover executions. In this analysis, the observation time always starts at $t = 0$ for any tau. 114
4.6. Change of offset before and after switchover 130
4.7. Latency results for HSR WR Switches 132
4.8. Bandwidth results for a HSR WRS ring with six devices 132

5.1. Reproducibility test of CWDM SFPs 150
5.2. Additive RTT and time error analysis with MUX/DEMUX devices . 154
5.3. Summary of RTT data and thermal coefficients. 159
5.4. Summary of time error data and thermal coefficients. 159

Acronyms

4G Fourth generation of mobile telecommunication technology

5G Fifth generation of mobile telecommunication technology

ADEV Allan Deviation

AMIGA Analysis of the Interstellar Medium of Isolated Galaxies

ARM Advanced RISC Machine

ARP Address Resolution Protocol

AVAR Allan Variance

BC Boundary Clock

BMCA Best Master Clock Algorithm

BME Bolsas y Mercados Españoles

CAM Content Addressable Memory

CD Chromatic Dispersion

CDMA Code-Division Multiple Access

CDR Clock and Data Recovery

CERN European Center for Nuclear Research

CLONETS CLOck NETwork Services

CLONETS-DS CLOck NETwork Services - Design Study

CPU Central Processing Unit

CSP Central Signal Processor

CV Common view

CWDM Coarse WDM

DAC Digital to Analog Converter

DAN Double attached node

DCF Dispersion Compensation Fiber

DCM Dispersion Compensation Module

DDMTD Digital DMTD

DMTD Dual Mixer Time Difference

DSP Digital Signal Processing

DUT Device Under Test

DWDM Dense WDM

E2E End to end

EBI External Bus Interface

EDFA Erbium-Doped Fiber Amplifier

EMPIR European Metrology Programme for Innovation and Research

ESA European Space Agency

ESMA European Securities and Markets Authority

FBG Fiber Bragg Grating

FFU Fast Forwarding Unit

FIFO First in, first out

FM Frequency Modulation

FPGA Field-Programmable Gate Array

FSU Fast Switchover Unit

GLONASS Gloval'naya Navigatsionnaya Sputnikovaya Sistema

GM Grandmaster

GMT Greenwich Mean Time

GNSS Global Navigation Satellite System

GPS Global Positioning System

HAL Hardware Abstraction Layer

HDL Hardware Description Language

HFT High Frequency Trading

HSR High-availability Seamless Redundancy

HVAC Heating, Ventilation and Air Conditioning

HY Hybrid Clock

IAA Andalusian Institute of Astrophysics

ICMP Internet Control Message Protocol

IDE Integrated Development Environment

IEC International Electrotechnical Commission

IEEE Institute of Electrical and Electronics Engineers

ILA Integrated Logic Analyzer

IP Intellectual Property

IP Internet Protocol

IRIG Inter-range instrumentation group

ISP Internet Service Provider

ITU-T International Telecommunication Union - Telecommunications Sector

IX Internet Exchange

LACP Link Aggregation Control Protocol

LAN Local Area Network

LEO Low Earth Orbit

LRE Link Redundancy Entity

LSB Least Signifficant Bit

LTE-TDD Long Term Evolution - Time Division Duplexing

LVDS Low Voltage Differential Signaling

MAC Medium Access Control

MEMS Microelectromechanical System

MEO Medium Earth Orbit

MiFID Markets in Financial Instruments Directive

MiFIR Markets in Financial Instruments Regulation

MMCM Mixed Mode Clock Manager

MSTP Multiple Spanning Tree Protocol

MTIE Maximum Time Interval Error

MUX/DMUX Multiplexer/demultiplexer

NF Noise figure

NIC Network Interface Controller

NIST National Institute of Standards and Technology

NMI National Metrology Institute

NPL National Physics Laboratory

NREN National Research and Education Network

NTP Network Time Protocol

NZDSF Non-zero Dispersion Shifted Fiber

OADM Optical Add/Drop Multiplexer

OCXO Oven-controlled Crystal Oscillator

OEO Optical-electrical-optical

OHWR Open Hardware Repository

OS Operating System

OTN Optical Transport Network

OXC Optical Cross-connect

P2P Peer to peer

PCB Printed Circuit Board

PCS Physical Coding Sublayer

PHM Passive Hydrogen Maser

PI Proportional Integral

PLL Phase-locked loop

PM Phase Modulation

PMD Polarization Mode Dispersion

PMU Phasore Measurement Unit

ppb parts per billion

ppm parts per million

PPP Precise Point Positioning

PPP-RTK Precise Point Positioning - Real Time Kinematic

PPS Pulse Per Second

PPSG PPS Generator

PPSi PTP Ported to Silicon

PROFINET Process Field Network

PRP Parallel Redundancy Protocol

PRTC Primary Reference Time Clock

PSD Power Spectral Density

PTB Physikalische Technische Bundesanstalt

PTP Precision Time Protocol

RMS Root Mean Square

ROA Real Observatorio de la Armada

ROADM Reconfigurable OADM

RSTP Rapid Spanning Tree Protocol

RT Real Time

RTS Regulatory Technical Standard

RTT Round-trip time

RTU Routing Table Unit

SCB Switch Core Board

SDP Science Data Processor

SEC Securities and Exchange Commission

SFP Small form-factor pluggable

SKA Square Kilometer Array

SLA Service Level Agreement

SMD Surface-mount Device

SMF Single Mode Fiber

SNR Signal-to-noise ratio

SPB Shortest Path Bridging

SPOF Single Point of Failure

STP Spanning Tree Protocol

SyncE Synchronous Ethernet

TaaS Time as a Service

TAI *Temps Atomique International*

TC Transparent Clock

TCP Transmission Control Protocol

TCXO Temperature-compensated Crystal Oscillator

TE Time Error

TIC Time Interval Counter

ToD Time of Day

TRILL Transparent Interconnection of Lots of Links

TSU Timestamping Unit

TWSTFT Two-way Satellite Time and Frequency Transfer

UART Universal Asynchronous Receiver-Transmitter

UDP User Datagram Protocol

UTC Coordinated Universal Time

VCTCXO Voltage-controlled Temperature-compensated Crystal Oscillator

VHDL VHSIC Hardware Description Language

WDM Wavelength-Division Multiplexing

WR White Rabbit

WRITE White Rabbit for Industrial Timing Enhancement

WRS White Rabbit Switch

WRS-LJ White Rabbit Switch – Low Jitter

WSS Wavelength Selective Switch

Chapter 1

Introduction

The purpose of this first chapter is to present an overview of the context in which the study presented in this thesis is developed. Initially, the background and the context in which the thesis is framed are briefly introduced in the section 1.1, high-level notions are provided to facilitate both the understanding of the difficulties existing today in the competent area of work and the proposed and implemented solutions/improvements. The section 1.2 formally states the scientific objectives to be achieved, and then, in the section 1.3, this introduction is complemented with a summary of the national and international projects in which we have collaborated during the development of this work. Finally, this introduction closes with a schematic description of the content that can be found in the rest of the chapters.

1.1 Motivation

Accurate time synchronization has become a relevant necessity with ever broadening areas of application. The levels of precision and accuracy that just a few years ago were requirements of the most demanding scientific infrastructures in the world have progressively made their way into new applications. As of today, the need to achieve synchronization traceable to a few nanoseconds can be found in sectors of the economy as diverse as telecommunications networks, financial markets, or other critical infrastructures such as Smart Grids [1], [2].

Within this diversity of application fields, one can appreciate a common point among all of them: the extension of these systems over geographically wide areas, or alternatively, the need to interoperate with other distant systems. Often, systems distributed over large areas need their components to operate in a coordinated manner,

for which they need to have a common notion of time.

Historically, it is often said that the idea of "modern" time, that is, that there is a single master time reference in a geographic area to which all clocks adjust (as opposed to the local solar time that previously existed) is a creation of the industrial era that emerged in response to the need for trains within the same railway network to share a common time. Thus, in the 1840s the concept of Railway Time emerged in England [3]. Over time, the use of each locality's solar time fell into disuse due to the advantages of using standardized time. In 1884, Greenwich Mean Time (GMT) was chosen as the universal time standard at the International Meridian Conference and remained so until 1972, when the Coordinated Universal Time standard (known by its pseudoacronym UTC) took over and the world moved away from the Earth's rotation period as a reference to using atomic clocks.

Just as Physics and Metrology have made enormous strides over the past two centuries in pursuit of ever more precise and stable time references, the work of numerous scientists and engineers, including this present study, has consisted in developing methods and techniques to disseminate the commonly used time standard of each era to the applications that require it, maximizing the technically possible metrics according to each historical period and solution (precision, extension, reliability, etc.).

Turning back to the present, over the past few decades, communication network infrastructures have proliferated. Their capacity in terms of bandwidth, number of nodes served, and ubiquity has continued to grow to this day. These networks are instrumental technology for the development of numerous intelligent distributed systems. The growth in quantity and complexity of these systems results in the magnitude of the data that traverses these networks where, by design and for the sake of backwards compatibility, the importance attributed to synchronization or reliability is quite low. In fact, in the Ethernet networks commonly in use today, there is no guarantee that the data sent through the network will reach its recipient in the same order in which they were sent, nor is a maximum latency guaranteed. Even more, the delivery of data to the recipient is not even guaranteed [4]. This is what is known as a best effort delivery mechanism.

This is an apparent contradiction since the normal operation of distributed systems often implies correlating data from physically distant sensor platforms, launching a process start within a bounded time interval, or simply precisely monitoring the overall state of the system to be able to perform a forensic analysis that reconstructs the order and intervals between all the captured steps. To carry out such tasks, it is essential to meet the requirement of sharing a common time base. Otherwise, system

components would not be able to temporally tag events, nor determine the time it takes for a message to reach its recipient. Without these premises, it is impossible to guarantee the causality of the events that occur in the system.

Therefore, architects of such systems must include a specific solution to maintain synchronization throughout the network. The two most common methods in use today to maintain equipment synchronization are techniques based on Global Navigation Satellite Systems (GNSS) or techniques based on packet exchange in the aforementioned best effort networks.

The first of these technologies is, in itself, a distributed system: global navigation satellite systems such as GPS (Global Positioning System) or GLONASS (from the Russian transliteration of Global'naya Navigatsionnaya Sputnikovaya Sistema or global satellite navigation system). These systems are used to distribute not only the end user's location but also time. This time source offers significant advantages: its coverage is global, the degree of precision is relatively high for the required investment and the maintenance of the source and its traceability are guaranteed by strong public institutions. However, GNSS signals, which are received by the end user as a weak radio signal that was transmitted from a constellation of satellites tens of thousands of kilometers away, are vulnerable to a whole range of possible issues or deliberate attacks. The signal is susceptible to atmospheric turbulence, must be obtained through an outdoor antenna, and is vulnerable to interference or spoofing attacks [5].

From this weakness of GNSS-distributed time, coupled with the criticality of many of the infrastructures that rely on it for their operation, comes the need for managers of these systems to provision or complement their time sources with others distributed through different means, and/or to provide equipment that is fault-tolerant and can manage multiple sources simultaneously.

Among the alternative sources and distribution methods to GNSS, those that transmit information through communication networks stand out. Within this family, there is a variety of techniques and protocols that, depending on their complexity and implementation requirements, can provide a variable level of accuracy and precision. The error range that can be achieved with these protocols can vary from a few milliseconds of error (as achieved by the software implementations that exist in many consumer electronic devices) to the few tens of picoseconds that can be achieved with other packet exchange protocols assisted with specific hardware. White Rabbit (WR), an extension of the Precision Time Protocol version 2 (PTPv2) [6] born to meet the needs of a top-tier scientific infrastructure (the European Center for Nuclear Research or CERN), is an example of the latter [7].

However, there are difficulties in achieving the same level of precision when time must be distributed over long distances for two fundamental reasons: first, the nature of the transmission medium makes it impossible to reach great distances using the same link model as in a local scale network due to physical phenomena such as signal attenuation and dispersion. Second, when using other techniques and devices to enable signal transmission to more distant distances, new elements are introduced into the link model that must be taken into account. There is an added difficulty to the previous ones if the transmission is carried out through a third party managed transport network (i.e. a internet service provider or ISP) which can alter the network at any time. While these modifications are usually transparent in terms of data service, which is not affected, the impact on a time distribution system can introduce faults in the synchronization signals and/or interruption of the service during a determined time interval.

Regarding the robustness of the synchronization system, it is worth considering that a possible failure in the time distribution system of the infrastructures and the use cases mentioned above can lead to malfunctions or complete shutdowns of the systems and, ultimately, to large economic losses and legal liabilities. To mitigate this risk, there are various approaches that increase the resilience of the system to failures. In the face of the possibility that a single time source is compromised and becomes a single point of failure, the implementation of redundancy techniques is suggested to allow switching between different time sources (switchover and failover). Conversely, in case of sudden failure of the device or system that is used as a time reference, specific hardware techniques can be implemented to ensure that the error of a local oscillator of a device does not exceed a certain threshold (holdover). This is necessary to maintain the integrity of the time signal during the switchover interval to an alternative source or, in case there is none, to prolong the operating status autonomously, without receiving external time corrections, for as long as possible until the failure has been resolved.

A time system that takes into account all the concepts and difficulties described as part of this general overview can be provided as a service. The companies that require it have a need for a complex synchronization system, but they may not possess the necessary infrastructure or expertise to ensure a signal that meets all the requirements of precision, accuracy, traceability, and high availability. These types of solutions are known as Time as a Service or "TaaS". Under this concept, time is a product supplied by a third party that takes care of demonstrating the integrity of the time signal at all times. TaaS solutions are common in the field of financial services. Within the frame-

work of this work, we illustrate the opportunity to collaborate in the deployment and maintenance of a pilot project of a time distribution system as a "TaaS" service based on time generation at a point geographically distant from the end customer, the distribution of a time reference through the public network of an ISP maintaining traceability to UTC, and redundancy and resilience mechanisms both at the source and at the end nodes.

To conclude this introduction, it is worth emphasizing that the focus of this work is on the extension of WR technology over public data networks, a challenging task given the variability and lack of inherent control in these infrastructures. Unlike other works in the literature, which are based on having complete control of the network, this study delves into the complexity of working with public networks and how they can be used to provide robust, high-precision time synchronization. Through this approach, we aim to provide a unique and valuable contribution to this field of study, exploring new frontiers of what is possible in terms of time synchronization across large-scale, public networks. The formal objectives of this study will be outlined in the following section.

1.2 Objectives

This thesis addresses the challenge of high-accuracy long distance time dissemination over optical fiber links, a service traditionally reliant on satellite methods or network protocols like PTP or Network Time Protocol (NTP) that cannot achieve the same level of performance. **We hypothesize that deployments of WR links can be made fault-tolerant, with resiliency features in both the node and the network, achieving an accuracy and precision equal to or better than GNSS systems, using infrastructure from existing telecomunication fiber networks**. This approach offers a cost-efficient solution that does not rely on the deployment of dedicated network infrastructure, is easy to deploy and reaches the high level of capillarity that the telecom fiber networks can provide, and fulfills the demanding requirements of multiple critical infrastructure industries like smart grid or fintech.

Our aim is to develop methods for deploying highly accurate WR links through public access networks, treated as regular traffic, while maintaining optical signal integrity. Given the dynamic nature of public networks, enhanced resilience and redundancy features are necessary for WR equipment to withstand changing conditions and maintain signal integrity and service availability. The lines of work followed to achieve these objectives are graphically represented in Figure 1.1.

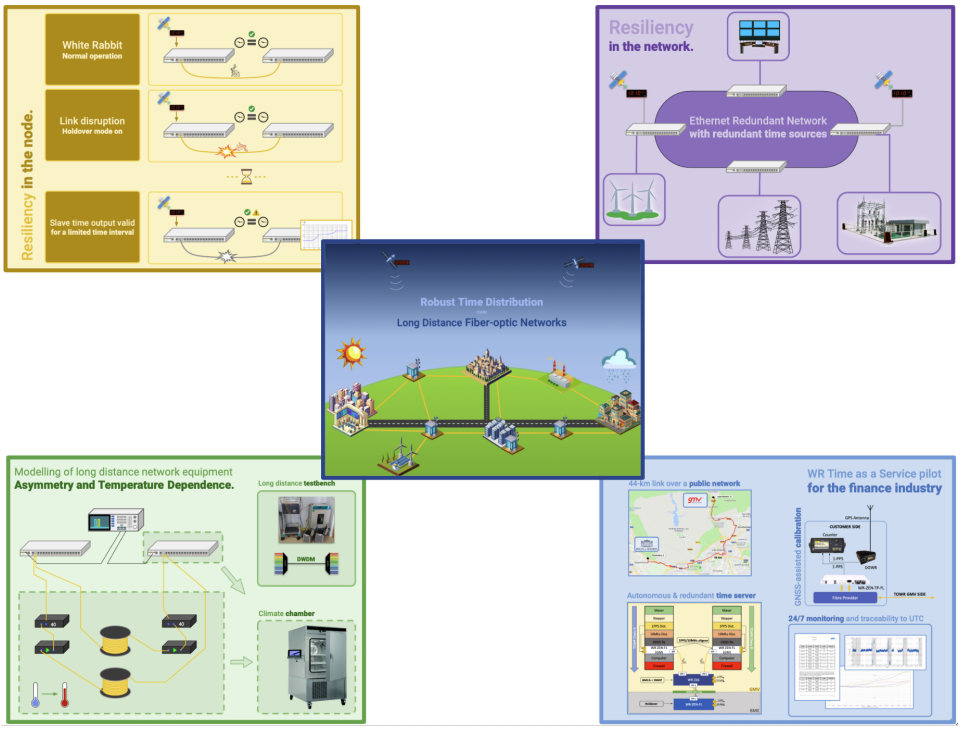

Figure 1.1: Graphic summary of the contents of this thesis.

The scientific objectives are specified in the following list, which reproduces the same points that were defined in the corresponding Research Plan when this project was started:

1. Study of the optical elements used in long-distance networks, as well as the problems of distribution of the synchronization signal in the optical medium. This will include investigating attenuation effects, dispersive effects, and any other factors that may affect time synchronization in optical Ethernet networks.

2. Development of component and network models that allow the deployment of synchronization service on long-distance optical fiber networks. This model will integrate information from components and equipment provided or obtained in the laboratory to determine their effect on the final system.

3. Analysis of network topologies and evaluation of reliability mechanisms in case of partial or total network failure. This analysis will be conducted for various applications, with particular emphasis on telecommunications and critical infrastructures, and risk mitigation strategies will be designed.

4. Design and implementation of topology selection mechanisms and switchover / holdover functionalities for long-distance networks. This objective includes the development of software and the description of Hardware Description Language (HDL) hardware for Field-Programmable Gate Array (FPGA), so that these functionalities applicable to long-distance networks are incorporated into WR systems.

5. Application and use cases. Experimental validation of the developed solutions and empirical verification with other time and frequency distribution methods. This will include the definition of success metrics, the performance of field tests, and the collection and analysis of data to demonstrate the effectiveness of the proposed solutions.

1.3 Framework

This research work has been carried out within the framework of several national and European research projects. This section provides a brief description of these projects: CLONETS, AMIGA, TOWR, and WRITE. Additionally, this thesis has been partially funded by the Ministry of Science and Innovation of the Government of Spain (MICIN) as part of the *Programa de Formación de Doctores en Empresas (Doctorado Industrial)* in its 2017 call within the company Safran Electronics and Defense Spain (then known as Seven Solutions).

1.3.1 CLONETS and CLONETS-DS

The CLONETS (CLOck NETwork Services, grant ID 731107 from the call H2020-INFRAINNOV-2016-2017 of the European Commission) project aims to prepare the transfer of new-generation technology to industry and strengthen coordination between research infrastructures and research and education telecommunications networks, with the goal of preparing the deployment of a sustainable and pan-European network that provides high-performance time services to European scientific infrastructures [8]. This project involves a consortium formed between European metrological institutes, academic institutions, and industry players involved in the topic. This network will be designed to be compatible with the distribution of time and frequency to a multitude of agents that require lower performance services than those of the metrological infrastructures, responding to needs created by developments such as cloud computing, the internet of things, and industry 4.0.

The CLONETS project has a series of high-level objectives through which information is obtained from national research infrastructures and their needs for time and frequency technologies and possible applications of a time and frequency distribution network for other areas of industry are analyzed, focusing on distribution through optical fiber networks. It also aims to define roadmaps and strategies that lead to the definition and construction of the pan-European time and frequency distribution network.

There is a continuation of this project, called CLONETS-DS [9] (grant ID 951886 from the call H2020-INFRADEV-2018-2020 of the European Commission), focused on the development of a design study that collects the needs of the scientific community and defines the architecture of the network."

1.3.2 TOWR

TOWR is an industrial project co-funded by the NAVISP (Navigation Innovation and Support Programme) of the European Space Agency along with the companies GMV Aerospace & Defence and Seven Solutions [10] identified by the activity code NAVISP-El2-033. The goal of the project is the development of a time distribution service using optical fiber in the Metropolitan Area of Madrid. This is achieved by generating its own time base from very stable oscillators, transferring GNSS time to align the reference oscillators, and using WR technology to distribute the time reference to the final customer through a public communications network. The project focuses on resilience mechanisms and signal traceability to meet the requirements of sectors such as finance. The pilot deployment of the project was carried out with the Spanish stock market operator, Bolsas y Mercados Españoles (BME).

Reviewing section 1.2 it can be seen that the objectives of this project are largely aligned with those of this thesis. Many of the results obtained were put into practice in the pilot deployment and will be described in detail in chapter 6.

1.3.3 AMIGA

AMIGA6

AMIGA6 (grant AYA2015-65973-C3-1-R MINECO/FEDER, EU) is an AMIGA project (Analysis of the Interstellar Medium of Isolated Galaxies) from the Ministry of Science and Innovation led by the Andalusian Institute of Astrophysics [11].

AMIGA6 is based on the results of previous AMIGA projects and is part of the preparatory work that AMIGA is doing for the scientific exploitation of the Square

Kilometer Array (SKA) radiotelescope [12]. The life cycles of HI (neutral atomic hydrogen clouds), both in isolated galaxies and in dense groups, remain poorly known since their low-density columns are only reachable by SKA. The main objectives of AMIGA6 are: to refine the models of cold gas accretion using isolated galaxies, and to analyze the role of HI tidal removal in the suppressed star formation in Hickson's compact groups. In preparation for the challenge of exploiting SKA data, AMIGA6 complements fundamental science with applied research in 3 work packages, thus contributing to SKA's Big-Data consortia. AMIGA6 is led by the coordinator of Spanish participation in the SKA, and brings together all the Spanish groups involved in the SKA data flow: Signal and Data Transport (SaDT), Central Signal Processor (CSP), and Science Data Processor (SDP). The relationship with this thesis falls within the group of Data and Signal Transport. In collaboration with *Universidad de Granada*, an assessment of time transfer technologies and the difficulties inherent in the network architecture of the SKA radiotelescope project was carried out.

AMIGA7

The AMIGA 7 Data Processing in Hardware subproject (grant number RTI2018-096228-B-C32 from the Spanish Ministry of Science and Innovation) [13]) investigates optimization techniques and synthesis of circuits that implement both signal processing algorithms of the CSP and scientific algorithms of the SDP and the Regional Centers of the SKA radio interferometer in reconfigurable hardware. AMIGA7 covers a key period for SKA, following the end of the consortium designs and the success of their Critical Design Reviews, until the beginning of construction.

Reconfigurable devices are firmly established in the architecture of the SKA's CSP modules, where they are widely used. DSP circuits usually allow the use of fixed-point formats that offer higher performance and lower cost and consumption than their floating-point equivalents. The project proposes the application of signal sensitivity analysis techniques based on their dependencies to optimize automatic methods. The application of these methods will enable new levels of optimization of the CSP modules in their evolution towards the construction of SKA.

On the other hand, the potential use of high-energy efficiency accelerators based on FPGAs for scientific computing in the SKA Regional Centers is of great interest to reduce their high energy cost. In this context, the project continues the development of the Witelo environment for the synthesis and automatic generation of VHDL (VHSIC Hardware Description Language) for high-performance implementations in FPGA and thus circumvent the limitations of existing tools. The goal is to complete

the environment to allow iterative 'stream' processing of unstructured data meshes, optimizing access to external memories and implementing local reduction operations.

On a second front, the project evaluates the feasibility and efficiency of FPGA designs of SKA scientific algorithms. Specifically, the primary objective is the artificial intelligence algorithms used in the analysis of the images obtained with SKA given their possible high level of parallelism.

1.3.4 WRITE

WRITE (White Rabbit for Industrial Timing Enhancement) is the European metrology project within the EMPIR (European Metrology Programme for Innovation and Research) initiative with code 17IND14-WRITE that involved metrology institutes from six European states (Italy, United Kingdom, France, Sweden, Netherlands, and Finland) along with companies involved in the development of time distribution technologies [14].

The overall goal of the project is to demonstrate all the metrological steps necessary for the industrial adoption of WR, including the improvement of devices and the study of an effective implementation in ordinary industrial computing infrastructures without degradation of the technique's performance compared to the results obtained in controlled research laboratories or using dedicated fiber infrastructures.

There are two specific objectives of the WRITE project that are particularly closely related to this work. The first of these is the development of validated techniques for redundant and resilient time transfer to industry end users through multiple time links from the source to the user, and providing alternative clock sources; the second, the development of new WR devices with improved performance that better interact with existing industrial protocols and standards.

1.3.5 Doctorado Industrial

This thesis project has been developed within the framework of the *Programa para la formación de doctores en empresas*, popularly known as the *Doctorado Industrial*, and co-financed between the Ministry of Science and Technology of the Government of Spain (at the beginning of the call, the Ministry of Economy, Industry and Competitiveness) and the company Safran Electronics and Defense (at the beginning of the call, Seven Solutions) with code DI-17-09646.

The research plan presented to *Universidad de Granada* and containing the objectives described in section 1.2 was used as a basis for the development of a scientific-

technical memory written in collaboration with the company Seven Solutions and evaluated by independent experts from the *Agencia Estatal de Investigación* (State Research Agency).

This program, whose access is controlled by a rigorous evaluation of proposals by the State Research Agency, allows the participant to develop a set of unique skills and competences as a result of the combination between research and academia with the experience and practical approach typical of the industrial world. This memory aims to highlight both aspects to demonstrate both the academic skills developed and the practical implications of the research.

It is noteworthy that this dual, academic and practical approach, opens the opportunity to participate in European research initiatives, like the previously mentioned projects, from a unique perspective in which the researcher is doubly involved, both in development from a theoretical point of view, as well as closer to industry and project management.

Also, these types of projects require that the results of the associated research not only be disseminated in the present doctoral thesis memory or in scientific journal publications, but these results must have an industrial component that allows their exploitation through the development of products and services. In this context, the researcher is not only at the forefront of his own research work, but also integrated into a product development team with results currently available on the market.

The program of the *Doctorado Industrial* project was elaborated with the understanding that it should align with the timeline of *Universidad de Granada*'s doctoral program as closely as possible. Thus, a temporal planning of the research project was elaborated, described as lines of work that culminate in justifiable milestones with associated deliverable documents. This scheme remains aligned with the Research Plan presented to the Postgraduate School of *Universidad de Granada*.

1.4 Chapter Structure

After providing an overview of the field of application and the motivation for this work, the organization of this document continues with the structure detailed below:

- In Chapter 2, the state of the art is presented to the reader. It consists of a review of literature focused on pre-existing work and an analysis of the technologies intended to be used in the developments described in subsequent chapters.
- Chapter 3 aims to present the *working environment*: the methodology used, the

platforms that serve as a starting point, and the universe of tools employed.

- Chapter 4 presents the techniques and methods developed to increase the reliability of a node, focused on extending its ability to maintain synchronization during time intervals when visibility of the reference has been lost, and on the development of redundant topologies that allow switching without reconfiguration time.

- Chapter 5 is focused on analyzing the specifics introduced by the optical medium itself for long-distance transmissions, and the problems that the techniques used by data network service providers can introduce to the time signal.

- Chapter 6 describes the pilot deployment of a time *TaaS* service in which participation was involved in its development and maintenance, validating in the real world the concepts developed in previous chapters. The work presented in this chapter was originally published in a peer-reviewed article [15]. Here, it is presented with further elaboration in the wider context of the overall thesis.

- Chapter 7 gathers the main contributions of this work and points out some possible lines of future work.

Chapter 2

State of the Art

> One very clichéd and overused way of starting an academic or technical is by defining the term that one is going to discuss. Experts argue that this can lead to a lack of engagement from readers out of boredom and disgust. It is a good piece of advice because…What is time anyway?
> Oh no, I tricked myself. In fact, the definitions of "time" that are commonly found in dictionaries possess a certain philosophical quality, which some might find amusing. *Time is the continued sequence of existence and events that occurs in an apparently irreversible succession from the past, through the present, into the future.*
> Wow. *Apparently* irreversible. I am not ready for this discussion.
> However, other remarkably more scientific sources define time like *what a clock reads* [16]. OK, now we are talking. Let's do this.
>
> The author

This chapter provides a comprehensive review of the state of the art in time and frequency transfer methods, with a particular emphasis on their resilience and redundancy features. We begin by examining the competitive and regulatory motivations that drive various industry sectors managing critical infrastructure to enhance their time transfer systems for improved resilience and reliability. In addition, we highlight the accompanying specifications and standardization efforts.

To establish a common foundation, we present an introductory section on timing concepts and terminology before more deeply discussing the technical details of time transfer mechanisms, techniques, and protocols. The discussion on noise theory and its associated statistics is essential for the developments presented in Chapter 4, as

well as the review of the basic principles and protocols of redundant networks.

Following the review of commonly used time transfer mechanisms and protocols, we culminate with a detailed examination of the White Rabbit (WR) protocol as the preferred choice for the advancements made in this work. We then procede with a literature overview of previous efforts in the field of long-distance WR protocols, which is complemented by the content in Chapter 5 as well as the subsequent section on transmission in long-distance fiber networks, before culminating the state of the art study with an evaluation of resiliency and redundancy techniques, with the focus applied on Ethernet networks and time transfer systems. It is worth noting that, with very few exceptions, the knowledge and ideas presented in this chapter are applied in practice in the pilot deployment showcased in Chapter 6.

2.1 Introduction to time transfer systems

Time transfer is a central process to numerous industries and technologies as it allows accurate synchornization of clocks over large distances. The need for precise, reliable, and redundant time transfer has only grown with the advent of newer technologies and the increasing reliance on automation in multiple industries, as this study will reflect. This section offers a brief overview of the crucial components in time transfer systems, seggregating the technologies and components according to their role in the system, and setting up the scope of this study.

In this introduction to time transfer systems, we will consider a model that comprises three principal components: time generators, time servers, and time distribution (Figure 2.1). At the end of that model, with the highest level of capillarity, the devices and mechanisms that consume time in different systems are located.

In the most basic terms, time is tracked by using a time generation device or process that uses a physical phenomenon to generate a periodical, distinct event. Throughout history, mankind has based its time tracking on multiple periodic natural and artificial phenomena (e.g. the rotation of the Earth, the phases of the Moon, or the oscillations of a pendulum). For modern industrial and scientific applications, time generating devices are artificial gadgets that exhibit excellent long-term stability, low phase noise and minimal frequency drift. Since one can start tracking time using any generator from an arbitrary point in time, it is needed to agree on a common reference starting point. Following this, we can consider that a clock is the combination of the time generating device along with a counter that tracks the number of periodic events.

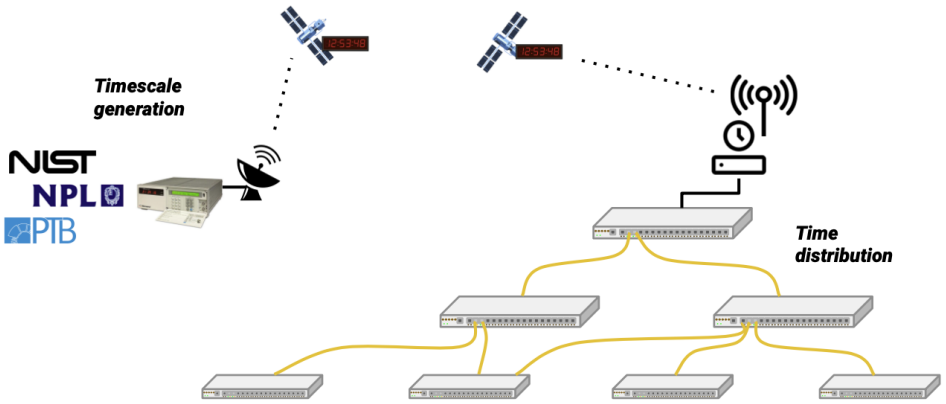

Figure 2.1: Schematic view of an example time transfer system. Timescales are generated by NMIs and other research institutions, which then are disseminated using methods like GNSS. A time server provides a network grandmaster with this timescale reference, which is distributed to other devices in the network.

The best performing time generators to date are atomic clocks based on the resonant frequencies intrinsic to certain elements like Cesium, Rubidium or Hydrogen. These atomic standards exploit the consistent frequency of the atoms electromagnetic emissions when they transition between energy states. Additionally, there exists a subclass of atomic clocks known as optical clocks. These leverage the frequency of light, which is higher than the microwave frequencies used in traditional atomic clocks, to enable even more precise measurements of time.

Although extremely stable and of common use in the metrological and scientific spheres, atomic clocks are too costly and heavy for the requirements of most industry applications. It is common use for National Metrology Institutes (NMIs) to realize and maintain timescales (i.e. the frequency standard along with the previously defined *counter*) based on these very stable time generation references, and then disseminate them by using the rest of the elements of a time transfer system. The combination of a frequency standard and its counter of ticks is referred to as *timescale*. Some of the most widespread timescales are Coordinated Universal Time (UTC) and *Temps Atomique International* (TAI). In applications where time traceability is important, time measurements are always related to an official, reputable scale through a series of comparisons with limited uncertainty.

Time servers are devices that play an essential role in the dissemination of time in a time transfer system. They function as intermediaries, receiving the precise time information and frequency signal generated by atomic clocks or any other precise

time generator, and it prepares it for distribution to various devices and systems that require synchronization. Many of these time servers obtain their time reference from a Global Navigation Satellite System (GNSS) that derives its timescale from the atomic clocks present in the satellite segments of the systems. Then, the time server consumes the time reference that it receives, and it implements mechanisms that allows the transmission of the time information utilizing different time protocols like Network Time Protocol (NTP), Precision Time Protocol (PTP), WR, or other mechanisms not necessarily based on packet-based networks.

Prior to any further analysis, it is important to acknowledge the inherent heterogeneity in network system topologies. The variety of diverse network technologies, protocols, and even transmission media requires a meticulous system design to ensure compatibility and interoperability. In the context of this work, given the widespread adoption of Ethernet optical networks for data transmission across physically extense networks, our investigation will center specifically on packet-based netowkr protocols over Ethernet optical networks.

GNSS techniques are also discussed due to their relevance as time sources in these systems, but there is a growing concern regarding the vulnerability of GNSS to external attacks [17], and public efforts are being promoted towards the development of alternative ways to disseminate timing without relying on the GNSS layer for dissemination [18]. Another possible response, as exemplified by the CLONETS initiative, is based on the use of public fiber optic networks to disseminate the timescales and frequency standards from NMIs to the places of interest to the industry. The distribution of time is completed in several tiers of service depending on the requirements of the time consumers and the capillarity of the network [8], [9].

2.2 Driving forces for reliability and redundancy in time transfer

Although the origin is unclear, there seems to be a consensus in literature regarding what *reliability* entails as a concept. Reliability is a probability. It is the term used to express a certain degree of assurance that a device or a system will operate successfully in a specified environment during a certain time period [19]. In the context of a timing system, it is the certainness that the system will meet the time accuracy and availability requirements within the established requirements, even in the case of failure within the system [20].

In contrast, to mitigate the risk of system collapse caused by Single Points of Fail-

ure (SPOFs), redundancy is employed in a system by incorporating duplicate or alternative components, processes or pathways, ensuring system resiliency and continuity in the face of potential failures. In the case of a timing system, redundancy can be achieved through the use of multiple time sources, communication links, and power sources. For example, a redundant timing system may use multiple physical communication links to ensure that the time information can be transmitted even if one of the links fails. This provides an additional layer of assurance for critical systems.

The term "resiliency" has different meanings across various disciplines. Crawford S. Holling, a Canadian ecologist, is credited with proposing the term "resilience" as a "system parameter" in his paper *Resilience and Stability of Ecological Systems*. In this document, he defined resilience as "a measure of the persistence of systems and of their ability to absorb change and disturbance and still maintain the same relationships between populations or state variables". While definitions of resilience vary across disciplines, they generally refer to the ability of a system to withstand change by adapting to a new situation.

In the context of this work, resilience refers to the combination of system improvement derived from both reliability and redundancy. A resilient timing system is prepared for sudden changes and can quickly recover from disruptions.

2.2.1 Requirements for reliability and redundancy in the industry

Reliability and redundancy are critical features in timing networks used in sectors such as Finance, Telecommunications, and Power. These features are necessary to ensure the delivery of accurate and reliable time even in the face of intentional attacks, component failures, and natural threats. Achieving resilience in these networks requires the use of alternative timing sources that devices can *failover* to in case of primary source failure, as well as the implementation of mechanisms that allow the devices to operate autonomously in the face of a disruption of normal operation.

2.2.1.1 Telecommunication networks

As new networks and communication technologies continue to evolve, the demands for synchronization also evolve rapidly. The telecommunications industry is one of the main drivers of technological advancements in synchronization. Although the critical infrastructure supporting this industry is well-established, there is a growing need to implement backup timing solutions for GNSS time sources. GNSS is often used as the primary source for timing, but it is vulnerable to interference and external

attacks. In addition to GNSS, atomic clocks and frequency standards are commonly selected sources in certain locations to provide highly accurate and stable timing [21].

To provide some historical context, early mobile communication networks like CDMA-2000 (from Code-Division Multiple Access) [22] required base stations to synchronize with system time with an accuracy of $\pm$ 3 microseconds, while allowing for a holdover of up to $\pm$ 10 microseconds in case of failure for at least 8 hours. With the advent of next-generation mobile communications like 4G LTE-TDD (Long Term Evolution - Time Division Duplexing), the synchronization requirement has become even more stringent, with small cells now requiring an accuracy of $\pm$ 1.5 microseconds. However, holdover requirements for 4G depend on the size of the cell. While larger cells with a radius of 3 km need to maintain holdover accuracy for up to 18 hours, small cells typically rely on low-cost oscillators that can only maintain phase accuracy for a few minutes [23].

As 5G continues to expand and new applications emerge, synchronization requirements are becoming increasingly specialized to specific services. Unlike previous generations, there is no single standard synchronization requirement that must be met across all 5G technologies. One of the most challenging requirements for 5G is carrier aggregation time alignment error, which measures the accuracy of time synchronization across different cell sites. The required time alignment error for carrier aggregation is stringent, with a maximum limit of 260 nanoseconds for inter-site synchronization and 130 nanoseconds for intra-site synchronization [24].

While GNSS-based timing distribution has been the go-to solution for many industries, it may not be accurate enough for the most demanding services in 5G, necessitating the deployment of more advanced techniques like common view (CV) time transfer or other alternative approaches. However, implementing these solutions will imply on-site calibration and performance monitoring to ensure that the quality of the solution is met in the field.

As of today, the most accurate standardized packet-based time transfer protocol over optical fiber (IEEE 1588-2019 High Accuracy, based on WR) meets the most stringent requirements of the industry in terms of accuracy. However, its original implementation lacked interoperability and reliability features that have hindered its adoption in the telecommunications sector for many years. The inclusion of the methods that help WR achieve its performance as a High Accuracy profile in the latest IEEE 1588 PTP standard is an important leap towards improving interoperability. Additionally, the development of redundancy and resiliency features in industrial WR devices with a suitable technology readiness level will help overcome the cur-

rent limitations and enable the integration of WR technology into current and future infrastructure.

2.2.1.2 Finance

The rise of modern information technology and communication networks has revolutionized the financial industry, making almost instant trading possible on a global scale. However, this disruptive change has also led to new challenges, including an arms race among certain sectors to achieve the lowest possible latency and fastest response times to changing market conditions [25].

Fintechs, particularly those involved in High Frequency Trading (HFT), rely on minimizing their latency with respect to the market and optimizing the reaction times of their decision-making devices because the slightest additional delay in their market orders can have significant impacts on their transactions. HFT relies on the ability to execute trades in milliseconds or even microseconds, and state-of-the-art technology is critical for achieving an advantage with respect to the rest of the market participants.

Regulatory efforts such as the Markets in Financial instruments Directive (MiFID) [26] have been put in place by the European Securities and Markets Authority (ESMA) to ensure fairness, transparency, and stability in financial markets, partly as a response to the practices mentioned in this section. MiFID has been instrumental in ensuring post-trade transparency by requiring the recording and timestamping of all events, cross-venue order monitoring, and the detection of market abuse. Entities subject to Regulatory Technical Standard 25 [27] must store records of their transactions, systems and algorithms for a minimum of five years. Some other requirements of the directive are relevant for the content of this work:

- Business clocks that timestamp any reportable event must be synchronized with UTC or the timing signals of a satellite system, provided that any offset from UTC is accounted for and removed from the timestamp.
- Operators, members and participants of trading venues must comply with a maximum divergence from UTC of 100 microseconds, with a timestamp granularity of 1 microsecond or better.
- A system of traceability to UTC must be established by documenting the system design, functioning, and specifications.

It is important to note that the regulations require traceability of the source of time to UTC. Any measurement of time must be related to a UTC realization through

an unbroken chain of measurements, each with its associated uncertainty. This demand for traceability at all times emphasizes the need for reliable timing solutions in the financial industry. To ensure this reliability, multiple time sources are preferred, and the installation of robust oscillators with a substantial guaranteed holdover time can benefit this sector, allowing it to operate even during outages of their primary time references.

2.2.1.3 Power industry

The Smart Grid is a representation of an advanced electrical grid that uses modern technology, such as digital communications and automation, to improve the efficiency, reliability and sustainability of electricity production and distribution [28]. The integration of information and communication technologies enables comprehensive control and monitoring systems, leading to more efficient energy generation, particularly when a larger portion of energy is generated from renewable sources that cannot be adjusted to accommodate instantaneous demand. Real-time visibility of the grid's state facilitates the management of load and generation fluctuations. Additionally, the impact of outages can be mitigated by enhancing the situational awareness of grid managers, who can then isolate faulty areas and prevent cascading failures. All of these capabilities rely on a common time notion distributed across the system.

Not all Smart Grid applications impose stringent timing requirements. For example, customer-side intelligent metering can tolerate a significant offset without any quantifiable effect. However, other applications in the generation and distribution sectors are much more sensitive. Synchrophasor measurements supply the system with estimates of electrical signal waveforms in magnitude and phase. Grid operators install Phasor Measurement Units (PMUs) at dispersed points within the electrical system to provide the monitoring layer with real-time data on voltage and current waveforms relative to UTC [29], and accurate time synchronization is crucial for PMUs taking samples of waveforms throughout the power system.

According to the IEC/IEEE 60255-118-1:2018 (for International Electrotechnical Commission / Institute of Electrical and Electronics Engineers) standard [30], "a time source that reliably provides time, frequency, and frequency stability at least 10 times better than values corresponding to 1% Total Vector Error is highly recommended". In numerical terms, the required accuracy is at least 2.6 microseconds. Additionally, the standard requires a granularity of at least 1 microsecond and traceability to UTC. Other capabilities, such as line differential protection and sampled analog values, also necessitate comparable accuracy levels. However, these capabilities only require rel-

ative accuracy to other points within the network and do not impose traceability to UTC [31].

For the use case of the electrical grid, deploying a time synchronization system over a fiber optic network is an efficient solution that leverages the existing network infrastructure that is already deployed for other necessities. A resilient time distribution network is also a cost-effective solution considering the cost of deployment and maintenance of GNSS-based receivers in a large number of remote and uninhabited locations, which can be an easy target for hostile agents. Considering this, the IEC has developed several technical documents that support the deployment of network time protocols on power infrastructure. For example, the IEC 61850 [32] standard for communication networks and systems in substations recommends using precise time synchronization protocols like PTP, and defines in its clause 61850-9-3 the PTP Power Utility profile, which aims at an accuracy better than one microsecond after crossing 15 network bridges with transparent clocks [20]. This profile supports precise timestaming of voltage and current measurement for differential protection, wide area monitoring and protection, busbar protection and event recording [20].

As for redundancy, the development of Smart Grid techniques has driven the advancement of Ethernet-based high-availability protocols. The IEC has also developed standards such as IEC 62439-3 [33], which defines the High-availability Seamless Redundancy (HSR) and Parallel Redundancy Protocol (PRP) to enhance the reliability and fault tolerance of communication networks used in power systems. Their principle of operation and capabilities are addressed in section 4.6.

2.3 Common ground: definitions and terminology

Before exploring in more depth the intricacies of the state-of-the-art advancements in time synchronization, it is necessary to establish a common foundation of understanding that will allow us to ensure clarity and consistency in terminology. This introductory section aims to provide concise definitions and explanations of key technical terms and concepts, such as oscillators, noise, and frequency stability, which are fundamental to our discussion. To the greatest extent feasible, the concepts and mechanisms presented in this section are based on those found in existing technical documents and standards. This approach offers the benefit of increased consistency and clarity when compared to other publications, as well as building upon established knowledge and practices in the field.

2.3.1 Noise model of a time signal

A signal source has a sine wave output that is described by the expression [34]:

$$V(t) = [V_0 + \epsilon(t)] \sin[2\pi\nu_0 t + \phi(t)] \tag{2.1}$$

where V_0 stands for the nominal peak amplitude of the signal, $\epsilon(t)$ represents the amplitude deviations of the signal, ν_0 is the nominal frequency and $\phi(t)$ encompasses all the phase deviation processes.

In the field of time synchronization and frequency stability, the focus is put on the $\phi(t)$ term, from which we can derive the instantaneous frequency [34].

$$\nu(t) = \nu_0 + \frac{1}{2\pi}\frac{d\phi}{dt} \tag{2.2}$$

And the fractional frequency deviation can be derived from the instantaneous frequency as in:

$$y(t) = \frac{\Delta f}{f} = \frac{\nu(t) - \nu_0}{\nu_0} = \frac{1}{2\pi\nu_0}\frac{d\phi}{dt} = \frac{dx}{dt} \tag{2.3}$$

where $x(t) = \phi(t)/2\pi\nu_0$ are the time fluctuations derived from the phase fluctuations [34]. This $x(t)$ term is, in essence, the magnitude that is measured by lab instrumentation and represents an array of equally spaced samples of the time error measured.

2.3.2 General definitions

- Accuracy: a sense of closeness of a measurement to the true value of the quantity being measured. In the case of an oscillator, the accuracy refers to the frequency deviation with respect its nominal value [35]. Since phase fluctuations are inherent to frequency sources, the accuracy must be evaluated over a period of time and not based on its instantaneous value. Also, note that high accuracy does not imply high precision.

- Precision: a sense of reproducibility and consistency of a set of measurements (note: high precision does not imply high accuracy). In the case of the frequency of an oscillator, a precise result does not refer to whether the output has a larger or a smaller offset with respect to an ideal reference, as long as it stays the same over time [35]. Note that high precision does not imply high accuracy.

- Stability: related to the previous term, the term stability refers to the capability of an oscillator to produce the same time or frequency offset over time, regardless of its accuracy. This concept can be illustrated by thinking of a stable oscillator that, unless adjusted, produces a frequency with a large offset. On the other hand, an unstable oscillator that is periodically adjusted might be able to keep a high accuracy for a short period of time. Figure 2.2 helps illustrate this concept.

- Synchronization: the process of aligning the clock of one device to that of another device with high accuracy. In other words, synchronization ensures that the time kept by two clocks is the same or very close, regardless of their underlying timebase or frequency [36].

- Syntonization: the process of aligning the frequency of two clocks. This is particularly important in systems that require high frequency stability and accuracy [36]. Note that the syntonization process does not require the time base to be aligned at any moment.

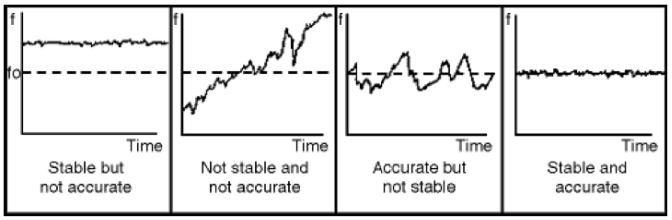

Figure 2.2: Relationship between accuracy and stability. From [35]

The following definitions are a free interpretation of terms derived from the ITU-T (International Telecommunication Union - Telecommunications Sector) recommendation G.810 *Definitions and terminology for synchronization networks* [37]:

- Jitter: short-term variations of the defining events in a timing signal from their ideal positions in time.

- Wander: same definition as jitter but applied to long-term variations.

- Phase noise: an inherent phenomenon in clock and frequency systems that refers to the random fluctuations in the phase of an oscillator or signal.

- Synchronous network: a network in which all clocks have the same long-term accuracy.

- Free running clock: one that is not being controlled by any syntonization process and has no access to stored data that could be derived from a previously connected external reference.

- Holdover mode: an operating condition of a clock which is no longer being synchronized to its reference and is using stored data that was acquired while it was locked, to minimize its phase and frequency variations.

- Locked clock: one that is being actively controlled so that the time error to an external reference is bounded, and the long term average frequency equals the one of the reference.

- Frequency stability: "the spontaneous and/or environmentally caused frequency change within a given time interval".

- Time error: the difference between the time of one clock and a reference clock, considered ideal. This concept can also be generically referred to as *offset* in this document. Time error is represented as $x(t)$ in section 2.3.1.

2.3.3 Noise characterization

The noise sources that degrade the stability of an oscillator can be modeled as a combination of processes that can be differentiated as a sum of power-law noises that are characterized by the spectral density of the fluctuations of their fractional frequency $y(t)$ as in:

$$S_y(f) = h_{-2}f^{-2} + h_{-1}f^{-1} + h_0 + h_1 f + h_2 f^2 + \cdots \tag{2.4}$$

where $S_y(f)$ is the power spectral density (PSD) of the fractional frequency fluctuations at Fourier frequency f, and h_α are the coefficients of the power law processes. The PSD of a time series $x(t)$ describes the distribution of its power into the frequency components that conform that signal:

$$S_{xx}(f) = \lim_{T\to\infty} \frac{1}{T} \left| \int_{-T/2}^{T/2} x(t) e^{-i2\pi f t} dt \right|^2 \tag{2.5}$$

The noise type has a different effect and receives a different name according to its exponent (Table 2.1. PM stands for phase modulation, while FM stands for frequency modulation). An example of the most common of these noises can be seen in Figure 2.3,

which can help understand how averaging contributes to remove noise in a time data set.

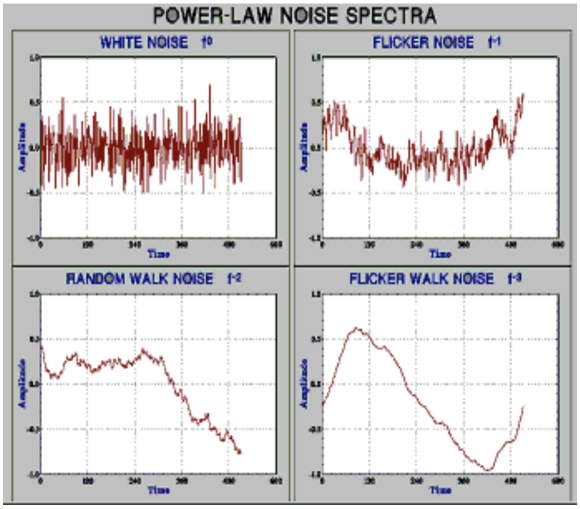

Figure 2.3: Examples of different power-law noises contained in a signal. From [34].

Table 2.1: Characterization of noise processes according to their power-law components [38].

Noise Type	Exponent
Random Run FM	-4
Flicker Walk FM	-3
Random Walk FM	-2
Flicker FM	-1
White FM	0
Flicker FM	1
White PM	2

From the power spectral density of the fractional frequency fluctuation, the PSD of the phase fluctuations and ultimately the single side band phase noise can be obtained [39]. This can also be used to connect the power-law noise coefficients with a time-domain stability statistic such as the Allan Variance (AVAR) [39]. While in some fields the stability of the data can be estimated from the standard variance of the data, this solution can be only applied with data whose results are time independent and under certain assumptions of the noise characteristics. In oscillator data, which contains time-dependent noise, the mean and the standard deviation will not converge over time to any particular value.

For that reason, non-classical statistics such as the AVAR and the Maximum Time Interval Error (MTIE) serve as a way to provide a quantitative measure of a clock stability over different intervals of time. These metrics, among others, are used to characterize the stability of time signals because using other traditional statistical tools such as the standard deviation do not converge under the presence of certain types of noise.

Allan Variance

The Allan Variance $\sigma_y^2(\tau)$ of a time series y (also commonly found as its square root, the Allan Deviation or ADEV) is a statistical measure defined as the two-sample variance of the difference in frequency between two equally spaced time intervals [34]:

$$\sigma_y^2(\tau) = \frac{1}{2(N-1)} \sum_{i=1}^{N-1} (y_{i+1} - y_i)^2 , \tag{2.6}$$

where τ is the sampling interval, and N is the number of samples in the time series. Or, in the case of a dataset of phase measurements:

$$\sigma_y^2(\tau) = \frac{1}{2(M-2)\tau^2} \sum_{i=1}^{M-2} (x_{i+2} - 2x_{i+1} + x_i)^2 \tag{2.7}$$

The ADEV is a useful tool to identify different kinds of phase noise. There are several types of phase noise that produce a unique distribution in their power spectrum and have different effects on the ADEV. The slope of the ADEV can help determine the interval of averaging needed to remove the noise types defined in Table 2.1 (Figure 2.4). In other words, by measuring the ADEV over different averaging times, it is possible to identify the type of phase noise that is present in a signal, reveal long-term effects that may not be apparent in short-term measurements, determine its possible sources and characterize its impact on the frequency stability. Beyond the stochastic processes that cause noise, the ADEV can also reflect systematic effects that would be hard to tell otherwise. For example, a heating system cycling every few hours will be reflected in the ADEV plot as a bump in the periodicity of the heating cycles.

Maximum Time Interval Error

While the AVAR is a metric of the statistical stability of an oscillator, the time error (TE) of a clock, with respect to a frequency standard, is defined as *the difference between the time of that clock and the frequency standard one* [37]. The TE $x(t)$ of a clock

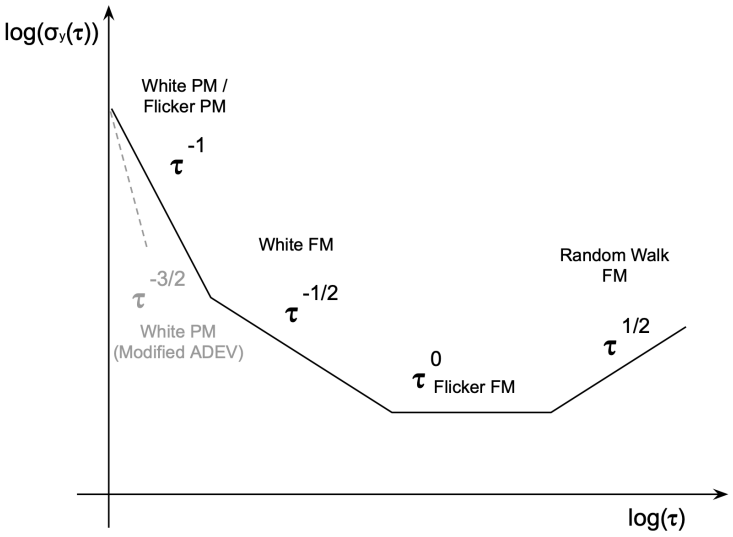

Figure 2.4: Different noise processes are characterized by their slopes in the ADEV plot. A modified version of the ADEV is necessary to discriminate the white PM noise from flicker PM noise.

that generates $T(t)$ in relation to a reference $T_{ref}(t)$ is defined as:

$$x(t) = T(t) - T_{ref}(t) \tag{2.8}$$

Based on the model of the signal of an oscillator that was provided in section 2.3.1, we can conceptualize a model for the time error that includes systematic errors in the initial synchronization, the frequency offset and the frequency drift [40]:

$$x(t) = x_0 + (y_0 - y_{0,ref})t + \frac{D - D_{ref}}{2}t^2 + \frac{\phi(t) - \phi_{ref}(t)}{2\pi\nu_{nom}} \tag{2.9}$$

and if we assume the reference to be ideal (i.e., that $T_{ref}(t) - t = 0$):

$$x(t) = x_0 + y_0 t + \frac{D}{2}t^2 + \frac{\phi(t)}{2\pi\nu_{nom}} \tag{2.10}$$

where x_0 is the initial time error, y_0 represents the initial fractional frequency offset, and D is known as the frequency drift (a linear estimation of the rate of change of the fractional frequency deviation, essentially attributed to aging and environmental conditions [37]) from a specified initial value. The TE is the initial function from

which many other stability parameters can be calculated, and it is a useful metric to characterize the error that is expected over an observation time. The Maximum Time Interval Error (MTIE) is a function that obtains, as its name implies, *the maximum peak-to-peak delay variation of a given timing signal with respect to an ideal timing signal within an observation time for all observation times of that length within the measurement period* [37]. The expression of the MTIE in all the possible observation intervals τ within a measurement period T is defined as:

$$\mathrm{MTIE}(\tau, T) = \max_{0 \leq t_0 \leq T-\tau} \left\{ \max_{t_0 \leq t \leq t_o+\tau} [x(t)] - \min_{t_0 \leq t \leq t_o+\tau} [x(t)] \right\} \tag{2.11}$$

This figure is commonly used in the industry to characterize the time stability properties of time sources, specially in the field of telecommunications networks. By quantifying the maximum error over a specified interval, the MTIE provides an expected worst-case scenario for the performance of a time or frequency source over a time interval (Figure 2.5).

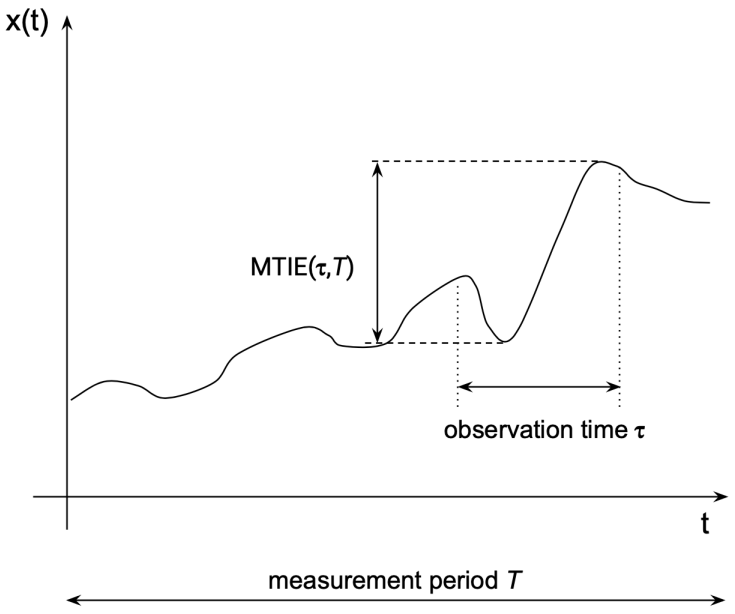

Figure 2.5: Visual representation of the MTIE of a time signal.

2.3.4 Types of oscillators

After discussing the fundamental concepts of noise characterization and the related stability metrics, it becomes necessary to consider these principles in the context of the devices that show these behaviors: oscillators. Different classes of oscillators have unique properties that give rise to specific noise characteristics and stability behaviors. Understanding these variations is essential for proper oscillator selection and subsequent system design, mainly determined by performance requirements and budget constraints.

Oscillators are devices that generate a periodic signal of a given frequency, built around a physical phenomenon [41]. Imperfections and random processes cause their output to always include some level of noise. The type of oscillator, its design, and the underlying physical principles all contribute to define the noise profile, thereby impacting on the oscillator stability. From the ubiquitous quartz crystals to the latest research on optical clocks, every clock source has properties that makes it suitable for a range of applications.

Quartz oscillators function based on the piezoelectric properties of quartz crystals. An electric field causes the quartz to generate a steadily oscillating signal. These oscillators are compact, economical, and reliable, which makes them commonplace in consumer electronics. However, they are prone to aging and are sensitive to environmental factors including temperature, humidity, and pressure [41]. A more advanced variant, oven-controlled crystal oscillators (OCXOs), mitigate some stability issues of quartz oscillators by housing the quartz crystal in a controlled temperature casing. The stable temperature minimizes frequency variations caused by environmental temperature change. While OCXOs offer better short-term stability, they also have larger size, higher power consumption, and cost more than their simpler counterparts [41].

The atomic clocks mark a significant leap in timekeeping technology. They use microwave emissions from specific atoms to create an extremely precise frequency signal [42]. Rubidium clocks, for example, use the hyperfine transition of electrons in rubidium-87 atoms. Cesium atomic clocks, on the other hand, use the frequency of microwave spectral lines emitted by cesium atoms, and hydrogen maser clocks operate similarly, relying on the emission frequency of hydrogen atoms [42]. These clocks offer exceptional long-term stability, though they are often impractical for many applications due to their size, power draw, and cost.

The vanguard of the timekeeping technology is currently being pushed by optical clocks. These devices also leverage the oscillations of atoms or ions transitioning

between energy states, but at much higher frequencies in the optical range of the spectrum. These higher frequencies imply that the *slices* of time that they measure with their ticks is much finer, resulting in greater precision and accuracy for timekeeping. Compared to cesium and rubidium cloks, optical clocks offer significantly improved stability [43]. While cesium clocks have been the primary standard for defining the second for several decades, there is a high degree of probability that the second will be defined in terms of optical clocks in the future [44].

It is noteworthy to mention that every type of oscillator is characterized by different stability behavior over short and long timescales. Short-term stability is often affected by random noise and environmental factors, while long-term stability can be compromised by systematic errors and the inherent limitations of the oscillator's physical principles. Figure 2.6 gives an overview of the relative stability levels of different types of clocks.

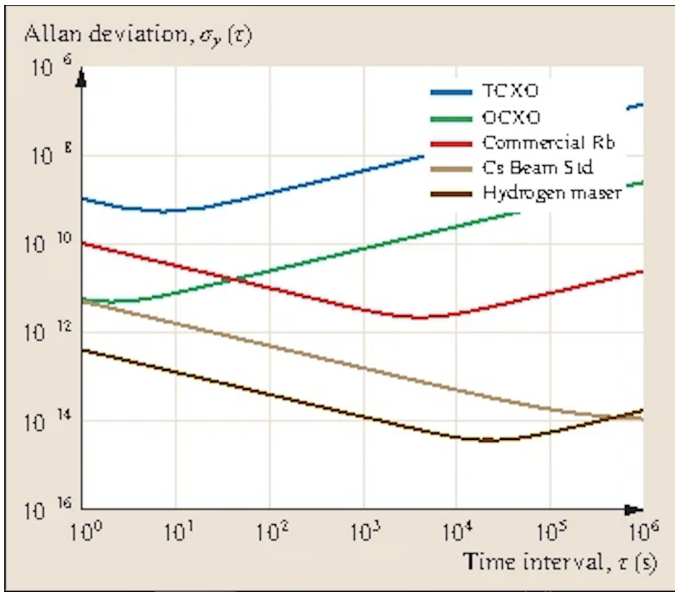

Figure 2.6: Stability performance of the classical microwave atomic frequency standards compared with temperature-compensated (TCXOs) and oven-controlled (OCXOs) crystal oscillators. Reprinted with permission from [41]

2.4 Time transfer: principles, techniques and protocols

The previous section has made clear that in any system, clocks are subject to various sources of errors, instabilities and other factors that cause them to drift apart, leading to discrepancies in timekeeping between different systems. To maintain synchronization, it is necessary to continuously monitor and correct these discrepancies.

Furthermore, when time information is transmitted between two clocks, it must travel through a medium, be it a network cable, a fiber-optic link, or even the atmosphere. The propagation delay of the medium is the time it takes for the time information to travel between the source and the destination. This delay can be influenced by factors like the physical distance, the medium's characteristics, and the transmission speed of the signal. Accurately accounting for these propagation delays is essential to achieve accurate time synchronization.

Besides the performance of the oscillators involved and the delay introduced by the transmission medium, there are other factos that may impact the accuracy and precision of a time transfer system. Environmental factors can introduce instabilities in the delay of the transmission medium, and electronics and software can introduce jitter and noise that limit the accuracy of the measurements done. In some cases like GNSS time dissemination, even relativistic effects play a role, and the motion of the satellites and the gravitational potential difference between the satellites and the receivers must be taken into account.

Time transfer techniques are designed to overcome these challenges by providing methods to measure and compensate for the errors and delays associated with timekeeping and signal transmission. These techniques can be broadly classified into one-way and two-way time transfer methods, each with its own set of advantages and limitations In the following sections, we will explore the underlying principles of the methods and their most common implementations.

2.4.1 One-way time transfer

Consider the most simple scenario in which clock A is a reference clock and clock B is physically distant from A, and not synchronized to it, as depicted in Figure 2.7. In the case of one-way time transfer, Clock A sends a time signal to clock B through a medium that introduces a delay δ_{AB}. If a high degree of accuracy is needed, clock B must consider the propagation delay between them. In this scenario, clock A sends its time signal as measured by its timebase, and clock B receives it as measured by its own timebase. Under this circumstances, we can name t_1 the time of departure of the time

signal from clock A, t_2 the time of arrival according to clock B, and offset_{AB} as the time offset between both clocks. Thus the offset can be estimated and used to adjust clock B:

$$t_2 = t_1 + \delta_{AB} + \text{offset}_{AB} \tag{2.12}$$

$$\text{offset}_{AB} = t_2 - t_1 - \delta_{AB} \tag{2.13}$$

The one-way time transfer delays are limited by the granularity of the clocks involved, and by the estimation of the delays involved. The designer of such a system must consider that not only the propagation through the communication channel must be considered, but also the internal delays of the transmiter and the receiver.

Also, this scenario is a simplified case that would only perform adequately if the delay between the clocks is static. Otherwise, more information needs to be added to solve the equations for time. Still, this principle of operation is what allows GNSS receivers to acquire time from the satellite constellations. In that case, either the location of the receiver is accurately known, or the visibility of several satellites is needed simultaneously to solve the equations for both position and time.

2.4.2 Two-way time transfer

In the case that both clocks can send their time signals to one another, each of the clocks can measure the time of arrival of the time signal according to their own timebase. Under those circumstances, if we keep the notation of t_1, t_2, delay_{AB} and offset_{AB} from the previous section, and we add t_3 and t_4 for the time of departure and time of arrival of the time signal sent from clock B to clock A, and considering that the propagation delays might be asymmetrical, therefore considering delay_{BA} as a distinct variable. This scenario is shown in Figure 2.8 and leads to the equations:

$$t_2 = t_1 + \delta_{AB} + \text{offset}_{AB} \tag{2.14}$$

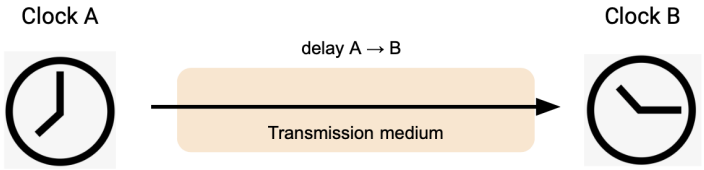

Figure 2.7: One way time transfer.

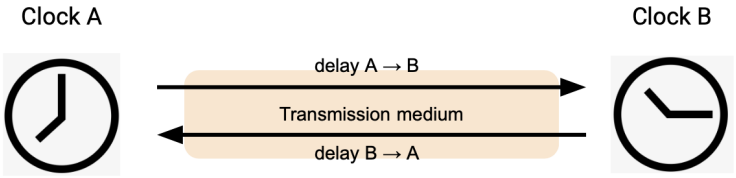

Figure 2.8: Two way time transfer.

$$t_4 = t_3 + \delta_{BA} - \text{offset}_{AB} \tag{2.15}$$

and the offset_{AB} can be calculated as

$$\text{offset}_{AB} = \frac{1}{2}\Big((t_2 - t_1) - (t_4 - t_3) - (\delta_{AB} - \delta_{BA})\Big) \tag{2.16}$$

which only holds true if the delays in both directions are equal. The time transfer mechanims that implement two-way exchanges can determine both the offset and the round-trip time ($RTT = (t_2 - t_1) + (t_4 - t_3)$) of the propagation medium, but they do not have a mechanism to calculate the asymmetry in each of the directions of the delay. Also, the transmission and reception delays at both ends of the exchange have not been considered in this analysis.

2.4.3 Techniques and protocols

The fundamental principles of one-way and two-way time transfer have been implemented into a diverse ecosystem of techniques and protocols. These methods vary wildly in terms of their propagation medium, accuracy and precision and ease of deployment, each of them catered for different use cases and applications. First, we will describe the fundamental characteristics of the most common GNSS and satellite-based time transfer techniques. After that, we will enumerate other techniques based on analog signals and optical fibers. Lastly, we will expand on protocols designed for deployment in packet-switched telecommunications networks, placing particular emphasis on WR, as it constitutes the basis of the present work.

2.4.3.1 Satellite-based techniques

One way GNSS time transfer

GNSS time transfer, a widely utilized method for time dissemination, relies on satellite navigation systems such as GPS, GLONASS, and Galileo. These systems transmit RF signals containing essential timing, navigation, and ephemeris data that enable GNSS receivers to accurately determine their location and time.

Under typical circumstances, a GNSS receiver requires data from a minimum of four satellites to solve the pseudorange equation, thus obtaining a precise measurement of location and time. However, in the context of GNSS time dissemination, the receiver's position can be predetermined and accurately established. Consequently, one-way data from a single satellite suffices to acquire the time reference from the highly precise atomic clocks incorporated within the satellites. Still, the ionosphere and troposphere introduce variable delays depending on day/night time and elevation angle. The navigation messages broadcast by the satellites contain a model to compensate for the ionospheric delays, but instabilities as large as tens of nanoseconds may remain [45].

The accuracy of the GNSS time transfer varies greatly depending on the specific hardware and techniques used. While geodetic techniques using multi-frequency antennas and receivers (which improves the model of the propagation delays) can achieve precisions of near 100 ps and accuracy of about 3 ns, the standard GPS service guarantees an accuracy below 40 ns for 95% of the time [46].

Common-view GNSS time transfer

This approach is an improvement on top of the one-way technique to compare two clocks at remote locations [45]. Adhering to the naming conventions of previous sections, Clock A and Clock B measure their time difference with respect to a GNSS satellite visible from both locations simultaneously. Subsequently, the clocks exchange their measurements to compute the difference. Ideally, GNSS time cancels out, as delay fluctuations between the satellite transmitter and the two receivers are highly correlated.

It is important that the measurements be made at the same time, and using the same satellite as "pivot", to prevent the stability of the satellite clock from becoming a factor. The common-view technique surpasses one-way GNSS time transfer by reducing common-mode errors. This method performs optimally when the distance between receiver stations is small relative to the distance to the satellite, as the

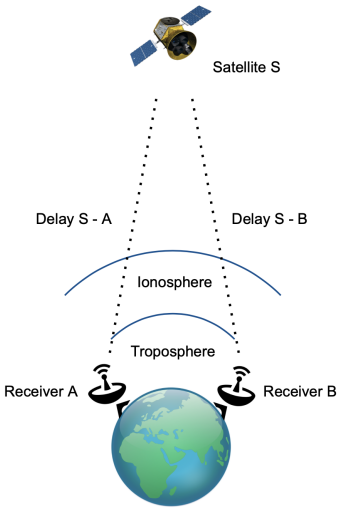

Figure 2.9: Principle of GNSS common view time transfer. When the baseline between the stations A and B is short, the ionospheric and tropospheric errors are highly correlated.

propagation paths exhibit greater similarity when the baselines are shorter.

Two way Satellite Time and Frequency Transfer

Two-way Satellite Time and Frequency Transfer (TWSTFT) is a technique used to compare time and frequency references between different locations. The methodology is centered around the exchange of signals between two ground stations via a geostationary satellite. Typically, the first ground station transmits a time-coded signal to the satellite, which receives and retransmits it to the second ground station. Upon receiving the signal, the second ground station compares the received time signal with its local atomic clock. This process is simultaneously performed in the opposite direction, which allows the stations to calculate the round-trip time and the offset between the stations following the principles of two-way time transfer. Assuming that the signals traverse symmetrical paths, many of the errors can be considered negligible because they are in the reign of common-mode [47].

Other GNSS techniques

While common-view GNSS Time Transfer and TWSTFT are techniques that have proven highly effective, they are not without disadvantages. CV GNSS time transfer requires the simultaneous observation of a same GPS satellite by two clocks at dif-

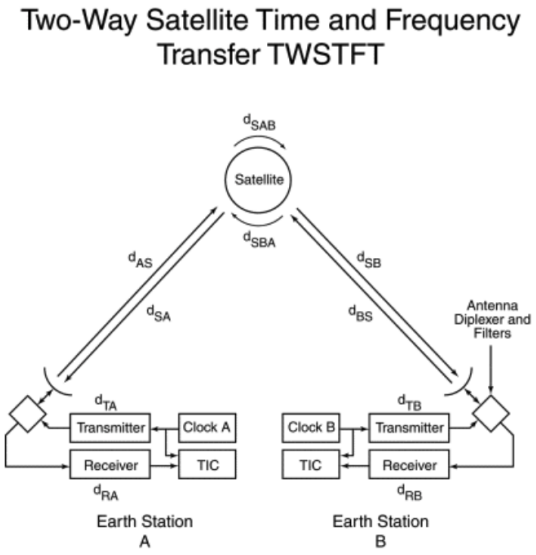

Figure 2.10: TWSTFT using a geostationary communications satellite. From [47]

ferent locations, which restricts its operation to certain time windows and makes it susceptible to errors from the satellite's clock. On the other hand, TWSTFT, although it offers continuous measurements and eliminates the satellite error clock, it involves a more complex setup and higher costs, but these are not the only available techniques that allows users to increase the performance of GNSS systems regarding the generation of a more precise time source or comparing remote timescales.

PPP (Precise Point Positioning) [48] is a technique used to improve the positioning and time accuracy of the data provided by a GNSS by using precise satellite clock corrections and orbital data using a single receiver. This technique uses a global network of reference stations to calculate errors in the GPS satellite's clock and the satellite's actual position compared to its prediction. The corrections are then applied to the measurements made by a single GPS receiver to improve its position estimate. PPP is a method that calculates the clock and position information in a post-processing approach, which makes it unsuitable for real-time conditions.

In contrast, PPP-RTK (PPP - Real Time Kinematic) [49] is a closely related technique that aims to equally improve the positioning and timing accuracy of the GNSS receiver, but the PPP-RTK technique relies on real-time data from the reference stations, in addition to precise satellite orbits and clock data, information about the satellite phase biases. Thus, instead of requiring a convergence time of hours, the assistance of an instantaneous stream of data allows PPP-RTK to operate in real time

conditions.

These are just two of the advanced techniques to reduce the errors of the computed position and time of a GNSS receiver. Several others also involve the use of differential GPS, essentially improving the performance of a receiver by mitigating the common errors shared between a stationary receiver at a known location and a mobile receiver that receives the corrections.

Other satellite constellations

All GNSS constellations have installed their space segment in Medium Earth Orbit (MEO) with an altitude between 19 000 and 24 000 kilometers. The satellites that conform these GNSS constellations feature atomic clocks with very stable long-term behavior, and their orbit is distant enough that their orbit can be predictably modeled and corrected. At a lower altitude (below 2 000 km), many Low Earth Orbit (LEO) formed by a larger number of satellites provide communications, positioning, navigation and timing capabilities for terrestrial receivers akin to GNSS. Given their lower altitude, the signal of these satellites reaches the surface of the earth with a much higher power level, often up to 1 000 times higher, which makes them more robust than GNSS signals. However, this does not comes without advantages. The lower orbit causes the satellites to suffer a larger atmospheric drag, which makes the orbits more unstable and requiring corrections more often. Also, these constellations do not have time sources with the stability of GNSS, and rely on other time sources (like GNSS, among others) for their operation. A profound analysis on the LEO satellite clocks is available in [50].

In addition, efforts like the European Horizon 2020 DEMETRA project proposed the use of geosynchronous satellites as one of several timing services that address specific needs of potential users in the industry [51].

Vulnerabilities of satellite-based time transfer

The widespread use of GNSS time dissemination has led to an increasing reliance on its accuracy and availability in a variety of sectors, including critical infrastructure such as the electrical grid, telecommunication networks and defense. However, the broad availability of the signals originated in satellites thousands of kilometers away from the surface of the Earth comes with the problems derived from the weakness of its signals. These signals are vulnerable to jamming, spoofing and unwanted interference [17], which poses a significant risk to both safety and security.

The term *jamming* refers to intentional interference caused by powerful radio frequency tramsmitters tuned at the same frequency as the GNSS signals. The consequence of such attacks can be severe, potentially disrupting the critical infrastructure [52] and causing system failure. Moreover, these attacks are difficult to mitigate and they can be executed using low-cost, commercially available equipment, which make them an effective weapon in areas involved in geopolitical tensions [53].

Another more sophisticated approach to disrupting the GNSS signal data is *spoofing*, which refers to the transmission of counterfeit GNSS signals designed to deceived the receiver into accepting false time and position information. This can be achieved by simply retransmitting the authentic time signals after a certain delay (*meaconing*) [54], or using more complex techniques that may lead to misrouting or causing a time error in the victim of the attack [55].

Other effects that may impact the normal operation of GNSS systems include, but are not limited to, natural interference (scintillations) and multipath effects, specially in dense urban environments that favor the appearance of so-called *urban canyons*. More details on GNSS threats and the proposed solutions can be found in literature [5].

2.4.3.2 Analog signals

One prevalent approach for implementing one-way time transfer involves the utilization of a network of timing cables to disseminate time signals (e.g., single or multiple pulses per second) from a master clock to the connected nodes. The electrical signal serves as a reference for the slave node's timebase. Since the analog signal lacks Time of Day (ToD) information, these systems often rely on an alternative parallel method to acquire the ToD, or the master may encode the information by modulating the analog signal, as exemplified by the case of Inter-range instrumentation group (IRIG) timecodes [56].

The use of analog signals, owing to their simplicity, has become a standard across numerous industries. However, the transmission delay associated with this approach depends on the calibration of the cabling, which is both time-consuming and prone to errors, making it unsuitable for large-scale systems with hundreds or thousands of nodes.

2.4.3.3 Optical fiber links

In the recent years, optical fiber links have been used for frequency dissemination between metrological institutions, which allows them to compare optical signals with breaking-edge accuracy and stability. These applications are the vanguard of research on diverse fields of physics such as possible violations of the equivalence principle of General Relativity [57], or even possible time variations of the mass ratio between electron and proton [58], and are made possible by disseminating the frequencies generated by the latest generation of optical clocks over optical fibers, achieving uncertainties below 10^{-19} over hundreds of kilometers [59].

Significant steps have been taken to practically demonstrate the capability of long-distance frequency comparisons over fiber links. Notably, the *Observatoire de Paris* (OP, SYRTE department) and *Physikalisch-Technische Bundesanstalt* (PTB) in Germany have been key players in these developments. As part of the European Metrology Research Program (EMRP) project "International Timescales with Optical Clocks" (ITOC), a long distance, stable optical fiber link of 1415 km was established between the PTB in Braunschweig and OP. This link enabled the frequency comparison of a strontium optical lattice clock at the PTB with a similar clock at OP, illustrating the promising capabilities of these links in precision timekeeping [60].

Beyond metrological comparison of frequency standards, there are other optical solutions based on active compensation of fiber delay fluctuation that allows fiber optic time and frequency dissemination over hundreds of kilometers that achieves a time stability below 30 ps over 30 days of operation, and an accuracy calibration below 50 ps [61]. Still, these solutions are based on custom and costly hardware implementations that hinders their adoption in a commercial scale. As an alternative, time transfer based on network protocols can often be implemented in software or make use of off-the-shelf network components that ease its deployment and limits costs.

2.4.3.4 Network protocols

Typically, time protocols are designed to be network-independent and can be encapsulated across a diverse range of both residential and industrial network standards. For example, both NTP and PTP can be utilized over any network that employs the Transmission Control Protocol/Internet Protocol (TCP/IP) suite. This suite serves as a global benchmark for data transmission across networks. Moreover, PTP provides explicit support for fieldbuses such as DeviceNet, ControlNet, and PROFINET [62]. However, it's crucial to note that the precise characteristics of the physical and link

layers can influence time performance. For this reason, this analysis presumes the use of an Ethernet network with an optic fiber as the transmission medium.

Network Time Protocol

Network Time Protocol (NTP) [63] is, most likely, the most widely deployed synchronization protocol in the world. It follows the principles of two-way time transfer for its operation, where a client exchanges packets with a known server and the result of the transaction is, at its most essential form, the offset between clocks and the round-trip time of the exchange. Its use since the beginning of generalized computer networking in the 1980s has lead to a large variety of implementations, usually completely hardware-agnostic, that is focused on providing service to a massive amount of users. Depending on the specifics of the architecture, NTP can typically maintain time to within tens of milliseconds over the public Internet, to a few hundred of nanoseconds under controlled environments [64].

NTP employs a hierarchical, distributed architecture to maintain accurate timekeeping across a network. This tiered structure consists of stratum levels, ranging from stratum 0 to stratum 15. Lower stratum levels indicate higher accuracy and reliability. The hierarchical structure also assists in load distribution and redundancy, ensuring that the failure of a single server does not disrupt time synchronization across the entire network. Primary servers, also known as stratum 1, synchronize directly with reference clocks (stratum 0), which include atomic cocks or GPS receivers. Secondary servers, or stratum 2, synchronize with stratum 1 servers, inheriting their time informaiton. The hierarchy continues downward to lower stratum levels.

This protocol implements a combination of algorithms that help improve the accuracy and stability of the time synchronization. These algorithms including Clock Selection, to assist the client in choosing the best possible server according to network conditions and the information provided by the clock server, Clock Filter, that processes multiple time samples from a server to estimate the most accurate offset an delay, and Clock Discipline, to adjust the local clock based on the calculated offset and delay variation employing a phase-locked loop (PLL) or frequency-locked loop (FLL) mechanism [65]. Still, the accuracy of NTP is limited unless dedicated hardware is employed, and the assumption of equal path delays for both directions also introduces a systematic offset.

Precision Time Protocol

Precision Time Protocol (PTP), also known as IEEE 1588 [62], is another protocol used to synchronize clocks over computer networks that is considered to address a segment of the market where neither NTP nor GNSS solutions provide an adequate solution. This protocol is used mainly in industrial environments where NTP cannot achieve the required accuracy, and GNSS-derived time may not be a sensible solution due to the necessary capillarity of the application.

Like NTP, the principle of operation of PTP is based on the two-way time transfer achieved by exchanging network packets containing specific information about the protocol, the capabilities of each device and, most importantly, the timestamps that determine the time of departure and arrival of the frames. Although it is not mandated by the standard, most PTP implementations base their accuracy on the generation of hardware-assisted timestamps. This removes a large source of uncertainty caused by the lack of determinism in the processing time introduced by the software implementation. PTP implementations using this approach typically reach sub-microsecond accuracies, while some of them, under tightly controlled conditions, can achieve accuracies below 100 ns.

The first version of the protocol was published in 2002 (IEEE 1588-2002) and is not compatible with the newer ones: IEEE 1588-2008 (PTPv2) is the most widely supported version to date, and IEEE 1588-2019 (PTPv2.1) is the most recent version and maintains backwards compatibility with the 2008 version. PTPv2 introduced the concept of *PTP profiles*, predefined sets of parameters and configurations that improve interoperability between vendors by ensuring that they can comply with specific protocol configurations (e.g. limited message rates, mandated network topology or transport protocol, requiring some features and prohibiting others…).

This protocol is based on a master/slave hierarchy, where the master serves as the primary reference source for slave clocks, and there are multiple clock types defined within PTP: Ordinary Clocks (OC) are end devices with a single network connection, Boundary Clocks (BC) act as both masters and slaves with multiple network connections, and Transparent Clocks (TC) are intermediate devices that measure and compensate for network delays.

PTP uses a set of message types to facilitate clock synchronization, including *Sync*, *Follow_Up*, *Delay_Request*, *Delay_Response*, and *Announce* messages. *Sync* and *Follow_Up* messages are used for clock synchronization, while *Delay_Request* and *Delay_Response* messages are employed to measure the propagation delay. *Announce* messages help in

the selection of the best master clock based on factors such as clock accuracy and user-defined priority. Notice that PTP supports both one-step and two-step operation: one-step operation includes the timestamps of transmitted frames on the fly, while two-step uses *Follow_Up* messages to include the timestamp of the previous message. Figure 2.11 shows the PTP standard message exchange for one-step and two-step operation. The naming of the events is coherent with Section 2.4.2 for the sake of simplicity.

The selection of the best master clock in PTP is performed by the Best Master Clock Algorithm (BMCA). Its principle of operation is discussed in the following section.

PTP grandmaster selection: the Best Master Clock Algorithm

The basis of PTP time transfer is the master-slave message exchange (as described in the previous section). It follows that the time dissemination in a PTP network is a hierarchical process in which a device is synchronized to a master and, if suitable, provides time to its slave nodes. Ultimately, this scheme results in a logical tree topology in which all nodes refer their time from the same, single device. The PTP master that serves as ultimate reference of time in a PTP network is called *grandmaster*. As the name implies, a grandmaster is the highest-ranking clock within its network, and typically, it obtains its time from an external source.

A network with such features becomes subject to a large problem: the timing of the whole network is reliant on a single device. If the grandmaster fails or there is a network outage in a vital link, the whole network will lose its reference. To mitigate this situation, PTP specifies that supporting devices must implement the BMCA. PTP relies on the BMCA to determine the best reference clock in a PTP domain. It is a distributed algorithm that continuously evaluates the quality of available time sources to elect the most accurate and stable timekeeping device.

Each PTP device in the network has a set of properties associated with its clock, which determine the outcome of the BMCA decision-making process. These key properties include the Clock Identity, which is a unique identifier; the Clock Class, defining the clock's accuracy and type; the Clock Accuracy, providing an estimation of the clock's timekeeping error; and the Clock Variance, indicating the clock's stability. Also, each device maintains its own Priority1 and Priority2 values, which are user-configurable and help determine the grandmaster in case of a tie.

The BMCA operates in a distributed manner, with each PTP device in the network executing the algorithm independently. Devices exchange Announce messages

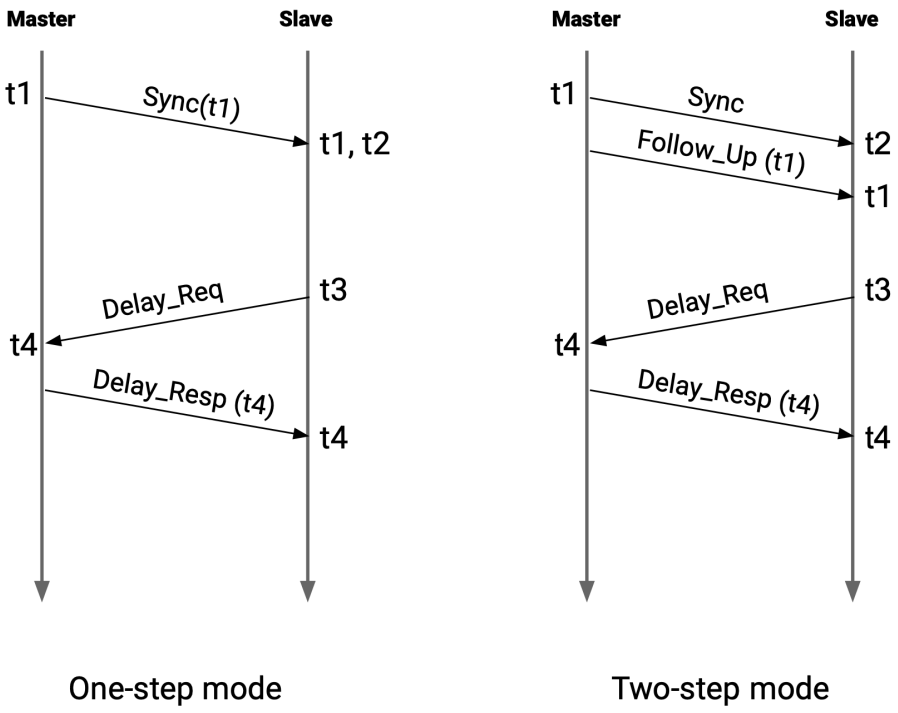

Figure 2.11: PTP message exchange.

at regular intervals, containing their clock properties and information about the current grandmaster. Upon receiving an Announce message, a device compares its clock properties with those of the sender and the current grandmaster. The comparison follows a predefined hierarchy, prioritizing Clock Class, then Clock Accuracy, followed by Offset Scaled Log Variance and finally considering the user-defined priorities.

The BMCA's distributed nature provides resilience against network failures and adaptability to changing conditions. If a grandmaster fails or its clock properties degrade, the BMCA enables the network to quickly elect a new grandmaster, minimizing disruption to the time synchronization process. Additionally, the BMCA allows for seamless integration of new devices with superior clock properties, ensuring that the PTP network consistently maintains the highest possible time accuracy and stability.

However, the reconfiguration time can pose challenges for critical applications that demand continuous service. During this period, clock synchronization across the network may be disrupted or degraded, leading to potential inconsistencies. Critical applications such as those found in telecommunications, industrial automation on financial trading systems rely on consistent timekeeping to ensure proper operation,

Table 2.2: PTP parameters in use by the BMCA. The specification of values for Clock Class and Clock Accuracy can be found in Table 4 and Table 5 of the PTP standard.

Parameter	Value	Description
Clock Class	0-255	Clock quality indicator
Clock Accuracy	From <1 ps to >10 s	Estimated clock accuracy
Offset Scaled Log Variance	log2(s)	Estimated clock stability
Priority1	0-255	Clock priority, user-specified, pt. 1
Priority2	0-255	Clock priority, user-specified, pt. 2

safety, and regulatory compliance. These use cases can benefit from the use of seamless redundancy protocols that provide zero time recovery. These protocols enable devices to receive time synchronization messages from multiple sources simultaneously, ensuring continuous operation in the event of a device or network failure.

2.4.3.5 Introduction to the White Rabbit protocol

Many new features have been made part of PTP in the latest 2019 version. Among others, the inclusion of performance statistics monitoring as part of the protocol, or the support for multiple simultaneous domains (i.e. instances of the PTP protocol). But the most relevant newly supported feature of IEEE 1588-2019 in relation to this work is the inclusion of a High-Accuracy Delay Request-Response Default PTP profile, which formalizes WR as part of the IEEE 1588 standard. In the following section we describe the characteristics of WR in more depth.

WR was born as part of an open initiative led by CERN in which other scientific institutions, metrological institutions, universities and private companies also took part to develop a time synchronization protocol that could comply with the stringent requirements needed for the timing system of the CERN accelerator complex [7]. The specifications of the products, the software and hardware design of WR devices are open source, all compiled in the Open Hardware Repository (OHWR)[1].

Leaving the recent standardization as the High Accuracy profile of PTPv2.1 aside [66], WR was implemented as an extension on top of PTPv2 aiming for time accuracy below 1 ns and a precision of a few dozen picoseconds for distances of up to 10 kilometers. The results shown in the first published experiments [7] demonstrate the achieved performance over three cascaded 5-km links.

As an extension to PTPv2, WR operates on the same principle as PTP, which in-

[1] https://www.ohwr.org

volves essentially exchanging frames containing information about timestamps. However, there is an unique set of closely intertwined techniques that allow WR to achieve a significantly improved performance with respect to the protocol in which it is based.

In this section, we will discuss of the key factors that contribute to WR's superior accuracy and precision. These include the use of syntonization of the physical layer clocks, an enhanced timestamping mechanism that can measure the phase relationship of the signals, and an improved link delay model that leads to more sophisticated calibration capabilities.

Syntonization of physical layer

In section 2.3.2, the difference between the concepts of syntonization and synchronization was already established. The first step in the synchronization process of WR is making sure that the clock of a slave is frequency-locked to the clock of its master, in a similar manner to how it is done in Synchronous Ethernet (*SyncE*). In fact, the WR clock does not comply with the SyncE specification because the latter is burdened with lots of legacy requirements that would hinder its performance.

If a WR slave receives a message that identifies its master as a WR-capable device, the syntonization process is started. The clock and recovery circuitry (CDR) of the slave generates a copy of the clock that the master device used to encode its Ethernet stream, and the slave will use a software-implemented PLL (SoftPLL) to lock the local oscillator clock to the one recovered from the master. The data used to steer the local frequency is a discrete stream of phase measurements performed by means of a series of DDMTDs (Digital Dual-Mixer Time Difference). Thus, in a WR network, all the nodes share copies from the clock of the Grandmaster device. The syntonization of the physical layer is extensively discussed in section 4.2.2.

It is worth noting that the recovery clock circuit is also put in use in the master side. The slave uses its local copy of the master clock to encode its Ethernet stream. At the master side, the "loopback" clock is recovered back by the master, and the phase relationship between the master clock and the "loopback" clock is used to enhance the timestamp resolution and ultimately align the phase of the slave clock.

Enhanced resolution of timestamps

One of the shortcomings of PTPv2 is the limited resolution of the hardware-generated timestamps. Considering the equations 2.14 to 2.16, one finds that the determination of the offset and the round-trip time of a PTP link is limited to the

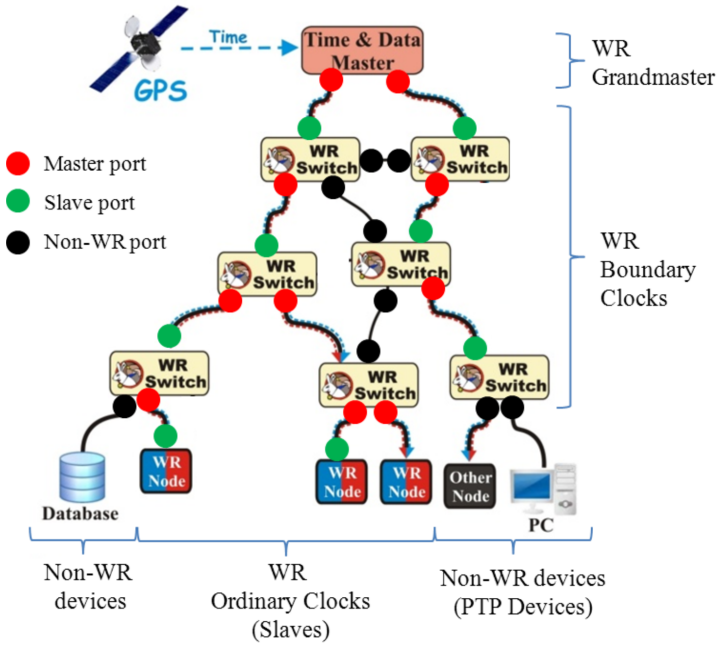

Figure 2.12: Schematic depiction of a WR network showcasing the multiple roles and connections of the same device in different locations of the topology. The red and blue arrows represent the duality of the control data (in red) and the WR time network (in blue). From [67].

resolution of the clock with which the timestamps are generated. In the case of Gigabit Ethernet, this would mean that the theoretical precision would be limited to 8 ns. WR takes advantage of the previously mentioned syntonization of the physical layer to measure the phase relationship between the master and the slave clocks.

Once the clock of the slave is frequency-locked, the WR protocol will proceed with the exchange of PTP messages that are timestamped for their transmission and reception time. While the transmission time is synchronous to the clock that governs the transmitting side, the reception timestamp is composed of a number of clock cycles plus a phase remainder. The DDMTDs [68] that continuously monitor the syntonization of the clocks can be used to enhance the resolution of the timestamps with the phase component. This information is shared between both parts of a WR link to calculate the phase relationship between clocks that was mentioned in the previous section.

The low-level syntonization and synchronization mechanisms in WR are studied in more depth in 4.2.1, where we identify the factors that limit the resiliency capabil-

ities of WR devices extending the pathway found in literature [69], [70].

Enhanced link model

Another limitation of the PTPv2 protocol (or any packet-exchange based protocol for the matter) is that asymmetries cannot be calculated. The endpoints of a PTP link do not have any mechanism to determine whether the time of flight of the master-to-slave path is the same as the time of flight of the slave-to-master path. For that reason, asymmetry is neglected, and both paths are assumed to be symmetrical. This is a limitation in the achievable accuracy that WR solves by having a more detailed and controlled link model.

The link model of WR is comprised of two types of delays (see Figure 2.13): fixed delays and dynamic delays. The fixed delays are attributed to the electrical paths inside the FPGA, the electronic board in which the components are mounted and the small form-factor pluggable (SFP)transceiver. Therefore, for a given link, these components can be separated into four: Δ_{txm} for the fixed transmission delay of the master side, Δ_{rxs} for the fixed reception delays of the slave side, Δ_{txs} for the fixed transmission delays of the slave side, and Δ_{rxm} for the fixed reception delays at the master side. The values of these components are considered a constant and depend on calibration. On the other hand, the transmission medium delays δ_{ms} and δ_{sm} are variable, and they are evaluated with every packet exchange.

However, as in the case of PTP, WR cannot tell if δ_{ms} and δ_{sm} have different values. To compensate this asymmetry, the WR link delay model assumes that the communication is done using a single optical fiber for both directions, and that each direction will employ a different wavelength. Under this assumption, the refractive index depends on the wavelength of propagation, which results in different propagation delays δ_{ms} and δ_{sm}. This difference is predetermined in the asymmetry coefficient α:

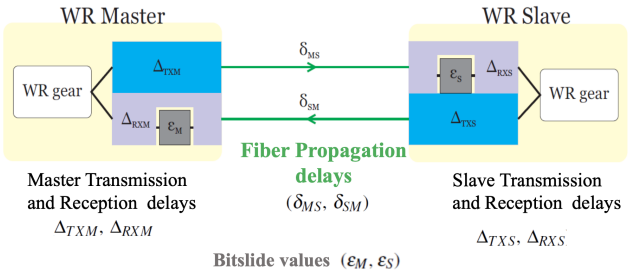

Figure 2.13: White Rabbit Link Model. From [71].

$$\alpha = \frac{\delta_{ms}}{\delta_{sm}} - 1 = \frac{n_{ms}}{n_{sm}} - 1 \tag{2.17}$$

where n is the refractive index associated to the propagation wavelength in each direction. The α parameters are calibrated in laboratory conditions for the selected wavelengths of transmission and type of fiber used.

2.4.4 White Rabbit over long distance

The WR design was initially conceived for links spanning up to 10 km. However, the utilization of off-the-shelf equipment allowed users to substitute components limiting the transmission to 10 km with those rated for greater distances. Nonetheless, extending the link length introduces a myriad of challenges.

Numerous experiments have been conducted to assess WR links over long distances in controlled environments. The Finnish metrology laboratory VTT deployed a dual fiber link across 950 km of the Finnish research network FUNET, achieving an accuracy of less than 2 ns compared to their GPS PPP reference [72]. Additionally, within an active telecommunications network in the Netherlands, a 2x137 km link was established between VSL and Nikhef sites, resulting in an uncertainty budget of 8 ns, while similar experiments have also been performed by the German ISP Deutsche Telekom [73] and the Swedish Netnod [74]. OP has also investigated and validated a series of optimizations for WR devices geared towards enhancing long-range link performance, with tests conducted on laboratory setups ranging from 25 to 500 km [75]. Notably, the longest WR link to date was deployed between Chicago and New Jersey (1,350 km) through a collaboration between Seven Solutions and Optiver US [76]. This link achieved a mean offset of 2.98 ns and a standard deviation of 10.4 ns when measured against a GNSS reference with a 15 ns root mean square (RMS) jitter.

In the aforementioned long-range deployments by the scientific community, maintaining accuracy across a public network is accomplished by developing a highly precise model of each intermediate component, calibrating the delay introduced by every part of it, and occasionally employing components specifically customized for this purpose.

These intermediate components are necessary for two reasons: one of them refers to the transmission of the signal towards longer distances, requiring amplification and the associated specific hardware. The second is related to managing the transport of the signal through network equipment that not only introduces asymmetries, but is also capable of dynamically switching the optical path. The focus of the current WR

link model is on addressing the linear effect introduced by the use of different wavelengths and the fixed delays at the end of the link. Although this approach is suitable for simple point-to-point connections, it falls short when dealing with deployments that involve amplification, optical switching, and other networks components. These additional modules can significantly impact the link overall latency and asymmetry, making the current link model insufficient for accurately characterizing and correcting these effects.

Besides the difficulties of accomodating the new components into the link model, calibrating every component is a practice suitable only for scientific or experimental deployments. Such methods are generally incompatible with the policies of network operators, who are not only exceedingly hesitant to grant access to their network architecture and equipment, but they may also modify the optical path without prior notification, deeming the effort useless. The unique challenges associated with WR long-range time transfer over typical network equipment found in an ISP's core network are extensively discussed in the following section.

2.5 Introduction to fiber optic transmission in long-haul networks

The overview of the WR protocol in the previous section stated the initial requirements for the design of the project. Regarding the reach of working links, the only original requirements were related to ensure the use of off-the-shelf ethernet components and a reach of 10 kilometers maximum. The final design is able to comply with those requirements, but is also easily extended to longer distances if the adequate hardware is selected. The study also reviewed some of the experimental deployments that have been done by metrological institutions, and stakeholders in the industry such as telecommunication networks operators.

The results of that analysis led us to the following conclusion: the deployment of WR links over long distances is possible, and the accuracy and stability of the time service can be maintained only if there is a close control of the network architecture over which the WR link is to be deployed. However, this control involves overcoming new challenges that the original WR link model did not consider. Mainly, we can discern two different main challenges: there is the need to have a precise model of every network component in between, which may be difficult if the network in which the link is to be deployed is owned by a external party that cannot disrupt the service, and also this external party will have deployed mechanisms that can dynamically change

the optical path for the purposes of ensuring the operation of the network in case of maintenance or unexpected failure.

For better understanding of the transmission options that we have, let us define define some terms related to the transmission options in fiber. One of the main characteristics of the core networks is that it is a shared infrastructure that agglutinates the signals from other regions of the network with a very high level of capillarity. Therefore, the network will have multiple points in which those signals are inserted and extracted from the network.

The name given to the method by which multiple signals are combined into the same shared medium is called multiplexing. This method is used by network operators to increase the capacity and efficiency of the fibers. Hence, the data streams from the final users are logically aggregated into different physical channels, and those channels are aggregated into a common physical medium (in this case, an optical fiber). The most commonly used multiplexing technique in the industry for the current generation of networks is Wavelength Division Multiplexing (WDM), which utilizes multiple wavelengths (or *colors*) of light to transmit multiple signals simultaneously. Inside this realm, we can distinguish between Coarse and Dense Wavelength-Division Multiplexing (CWDM and DWDM):

- CWDM systems use a wider channel spacing, typically around 20 nm. The ITU recommendation G.694.2 [77] defines 18 wavelength channels between 1271 nm and 1611 nm. The larger channel spacing allows for the use of less sophisticated and less expensive optical components. However, the CWDM systems allow for a lower total capacity because of their less efficient use of the spectrum. The initial WR specifications indicate the use of the CWDM channels corresponding to wavelengths 1310 nm and 1490 nm, although other channels can be used if the necessary considerations regarding network design and link asymmetry are met.

- DWDM systems, on the other hand, use a much narrower channel spacing, typically between 0.8 and 1.6 nm (100 GHz and 50 GHz channel spacing) [78], enabling a significantly higher number of channels. One drawback of the dense channel spacing is the need for more precise optical components such as the lasers of the transmitters. A DWDM laser transmiter implements a control system and a cooling mechanism to precisely stabilize the wavelength of the signal. As the discussion unfolds, we will discover that the more demanding requirements of these systems are an advantage for the purposes of maintaining

accuracy and precision in a time transfer system.

The CWDM and DWDM grid systems serve as an stable and reliable means of transmitting multiple data channels over the same fiber. Nevertheless, with the advent of increased data rates, the 50 GHz grid became a limiting for newer, more complex modulation schemes for rates beyond 100 Gbps. To efficiently accomodate the newer systems and signals, the ITU-T proposed a flexible, finer grid where bandwidth resources can be dynamically associated to optical signals depending on its requirements. Given the finer bandwidth slices of the flexible grid, the attainable efficiency and total capacity of a network is higher than that of a static DWDM system, as is depicted in 2.14. The flexible DWDM grid is also defined in the latest version of ITU-T G.694.1 [78].

However, since the WR equipment is not compatible with the newer, higher-rate transmission speeds and modulation schemes, the scope of the study will be limited to the static DWDM grid and its associated components. Transmission systems that support flexible grid have different requirement and optical components, an overview thereof can be found in [79].

Another technology that impedes the dissemination of physical layer clock featured by WR is the Optical Transport Network (OTN) [80]. This is a ITU-T recommendation that defines a series of network elements connected by optical fiber links to provide functionality of transport, multiplexing, routing and monitoring of optical channels carrying client signals. One of the functionalities provided by OTN is to aggregate several lower-rate signals into a single higher-rate stream, encapsulating the lower-rate signals frames into a digital wrapper. Since this application modifies the physical layer of the WR Ethernet stream, the correct operation of WR cannot be guaranteed.

Beyond wavelength selection, there is another basic configuration choice to make when designing an optical fiber network or when designing an application on top of it of the like of a WR time service: the use of one fiber per direction of propagation, or the use of one fiber for both directions:

- Unidirectional transmission, which uses one fiber for each propagation direction, is common in public networks such as the ones owned by ISPs or educational and research institutions. For time transfer purposes, it has the main drawback of difficulting the estimation of link asymmetry. The length of the master-to-slave and the slave-to-master paths can be considerably different.

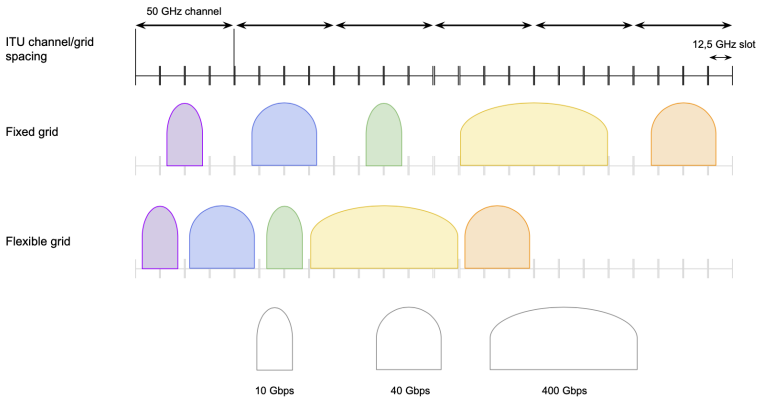

Figure 2.14: Fixed DWDM spacing grid vs. flexible DWDM grid. The dynamic nature of the flexible grid allows the network management to optimize the system capacity, allocating no more than the necessary bandwith for each application.

- Bidirectional transmission, on the other hand, uses the same fiber for both directions. While less common in long-distance deployments over public networks, this configuration is more advantageous for the time transfer applications: the bidirectional transmission ensures that both directios have the same optical length, and any transient fluctuations in delay will impact both directions equally. For these reasons, the usual WR links are deployed using two different wavelengths over bidirectional transmission.

Following these definitions, the scientific community employs two common terms to refer to long distance time or frequency transfer link: *dark fiber* and *dark channel* [81]. Dark fiber refers to unused (therefore *unlit*) optical fibers within a fiber-optic cable containing several fibers. These fibers are considered "dark" because they are not transmitting any data. Dark fiber can be leased to organizations that need to expand their network capacity or deploy new applications, such as the case of a time or frequency transfer system. The dark fiber approach gives the users the possibility of more control of the link because the fiber is not shared with any active application. On the other hand, dark channel refers to the lease of a dedicated channel within the DWDM grid of a fiber that is used to transmit data. This is a much cost-effective solution, although the restrictions put in place by the manager of the network are higher.

Once these concepts have been introduced, we can continue our discussion with the general factors that limit the transmission of data over fiber.

2.5.1 Limiting factors in fiber optic transmission

Fiber optic technology has revolutionized the field of data communication, providing unparalleled speed and bandwidth for transmitting vast amounts of information across large distances, and it is still a very active research field where records are periodically broken in terms of capacity (319 Tb/s over 3000 km in June 2021 [82]) and transmission speed (99.8% of the speed of light with hollow-core fibers, [83]). However, despite its many advantages, the fiber optic medium is not without its challenges. The design of a reliable fiber optic network requires a deep understanding of the various effects that can impact the performance and integrity of the transmitted signal.

Some of these effects, such as attenuation and chromatic dispersion, are well-known and widely studied. However, there are numerous other phenomena that warrant consideration when designing a fiber optic network. These include polarization mode dispersion (PMD), scattering, nonlinear effects, bending loss, connector losses, absortion and mode coupling, among others. More importantly, some of these effects may have negligible effects for data transmission, while having a larger impact on the design of a time transfer link. In the following sections we are going to review and discuss the most noticeable of these effects from the point of view of the design of a WR time transfer link, and the strategies to mitigate them.

2.5.1.1 Attenuation

Light traveling in a fiber loses power over distance due to causes such as intrinsic absorption (interaction of the light with the silica molecules), extrinsic absorption (interaction of the light with impurities introduced in the fiber during the manufacturing process or by contaminants such as water) and scattering (caused by fluctuations in the density and composition of the material and by the presence of impurities) [84]. Attenuation is one of the key reasons that define the wavelengths used in fiber optic communication, because it depends strongly on wavelength and, consequently, only the lowest attenuation parts of the spectrum are used for transmission. Figure 2.15 depicts the dependence of attenuation with wavelength and the historical three windows for which optical components are manufactured.

As a reference value, the SMF-28 Ultra fiber from Corning, one of the main manufacturers of optical components for scientific and industrial applications, specifies a maximum loss of 0.32 dB/km for 1310 nm, 0.21 dB/km for 1490 nm, and 0.18 dB/km for 1550 nm. Together with the distortion mechanisms, attenuation is one of the two

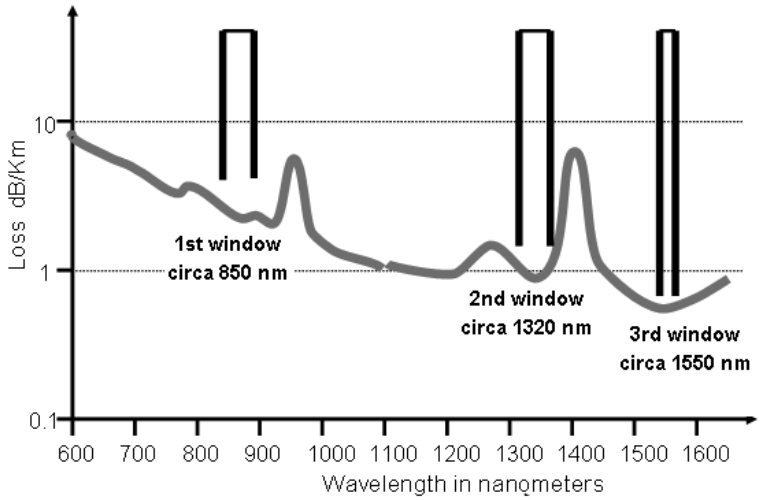

Figure 2.15: Relation between transmission wavelength and fiber attenuation. The three transmission widnwos refer to the wavelengths at which data can be transmitted more efficiently. The 850 nm window is typically used for short-distance communications. The second window allows for longer distance transmission, but the 1550 nm window is the typical choice for long-haul systems due to its lower attenuation and compatibility with common amplification technologies. From [85].

main effects that limits the maximum possible distance of transmission before the signal needs to be amplified or regenerated back to an appropiate level.

2.5.1.2 Chromatic dispersion

Chromatic dispersion is an inherent limitation of fiber optic transmission systems and, in fact, is a limitation that will be found in every transmission medium besides the vacuum. Chromatic dispersion is an effect of an intrinsic property of the silica glass of which fiber is manufactured: the refractive index of the medium is a function of the wavelength of the propagating light [86]. The refractive index n is, by definition, the relationship between the speed of light in a vacuum c, and the speed of light in a transparent material v:

$$n = \frac{c}{v} \tag{2.18}$$

The relationship between the refractive index n and the wavelength of propagation λ can be described by the Sellmeier equation, which models the dispersion characteristics of the material:

$$n^2(\lambda) = A + \frac{B}{1 - C\lambda^{-2}} + \frac{D}{1 - E\lambda^{-2}} \tag{2.19}$$

where A, B, C, D and E are the Sellmeier coefficients for the material. [87]. Using the same reference fiber as in the previous section, for Corning SMF-28 Ultra the effective index of refraction is 1.4676 for 1310 nm and 1.4682 for 1550 nm. Note that the Sellmeier coefficients show a slight dependence with temperature, which then make the refractive index vary slightly with temperature changes.

Optical sources do not generate an ideally a pure light that is composed of an exact wavelength. Instead, a modulated light signal is composed of a certain slice of wavelength spectrum. The consequence of this is that each of the wavelength components of the optical signal will find a slightly difference refractive index, and therefore they will travel at slightly different speeds. Ultimately this effect leads to *pulse broadening* [84] and is the other main limiting factor of transmission in fiber systems. Given a sufficiently long link, the energy of one symbol and the next one will be spread over time and the receiver is not able to tell symbols apart, as exemplified in Figure 2.16. This effect is called *material dispersion*.

Besides material dispersion, there is another dispersion factor known as *waveguide dispersion*. This rises from the fact that the energy that traverses the core and cladding of the fiber find slight differences in the refractive index of the material, and therefore travel at different speeds [84].

The sum of these two terms is known as chromatic dispersion (CD), and is characterized as a function of the wavelength, measured in ps/(km nm). The most commonly deployed fiber in network, compliant with ITU-T G.652, has a small chromatic dispersion in the wavelengths around 1310 nm, but it has a larger CD in the 1550 nm window. Dispersion in these fibers limits the maximum transmission length without compensation techniques in DWDM networks. Consequently, other fibers were designed with different dispersion profiles: G.653 is designed for the lowest dispersion in the 1550 nm window, but it presents a higher occurrence of other nonliner effects that render them useless for the deployment of DWDM. Another design, specified in G.555 for Non-Zero Dispersion Shifted Fibers (NZDSF), were designed to suppress the nonlinear effects and have reduced dispersion in the 1550 nm window, although they are not suitable for use in the region of 1310 nm [86]. See Figure 2.17 for an estimation of dispersion values over different wavelengths for the most common kinds of fiber.

Beyond pulse broadening that limits the distance of operation, chromatic disper-

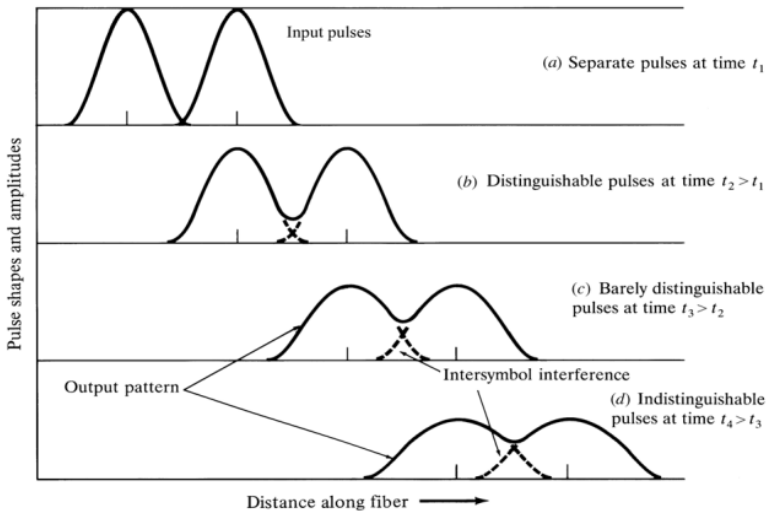

Figure 2.16: Depiction of pulse broadening. The energy of the pulse is spread over time with longer distances, impacting on the capability of the receiver to tell symbols apart. From [86].

sion has a direct impact on the accuracy of a WR link that goes unnoticed if we only focus on data transmission. If we recall the link model of a standard, bidirectional WR link, we will find the α parameter that is in charge of correcting the asymmetry introduced by the propagation of different wavelengths in each direction. This parameter is an estimation of the chromatic dispersion for a given fiber. The calibration procedure of WR enables the user to estimate the asymmetry parameter in a lab environment, but this is not suitable for long distance links, while definitely a considerable source of error. Assuming a 80-km link over a G.652 fiber with $D_{CD} = 17\text{ps}/(\text{km nm})$, the introduced error over the 160-km round trip is of 2.72 nm per every nanometer of separation between wavelengths in the DWDM grid. This can be as small as 2 ns error using adjacent channels, or 154 ns in the worse case scenario (71 channels apart) using dense, stable DWDM wavelengths.

2.5.1.3 Polarization Mode Dispersion

A light signal at a given wavelength in a single-mode fiber actually is transmitted through two orthogonal polarization states. While they are initially aligned at the beginning of their path over the optical fiber, these two modes will eventually be misaligned due to having travelled at different transmission speeds. This phenomenon is called polarization mode dispersion and is caused by imperfections in the manufacturing process of the fiber (slightly elliptical cores) or by the birefringence properties

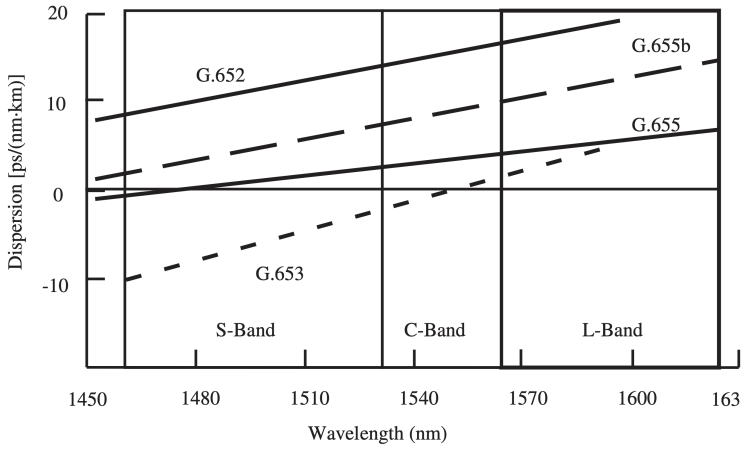

Figure 2.17: Relation of chromatic dispersion with wavelength for the most common fiber types. From [86].

of the material, that is, changes in the refractive index depending on the polarization and propagation direction of light. PMD increases with the square root of the distance, and for Corning SMF-28 Ultra the PMD is limited to 0.1 ps/$\sqrt{\text{km}}$. The effect of PMD drifts slowly with time in a random way, and it is difficult to compensate.

We assume that PMD is unlikely to significantly affect the performance of a WR link, based on the premise that the anticipated value of PMD in the fiber is inherently low. Following the analysis in [88], we expect the random fluctuations in delay due to PMD to be obscured amidst other, more substantial sources of instability in the system.

2.5.1.4 Nonlinear effects

Many other nonlinear processes impose limitations on the performance of an optical transmission system, but their impact is already limited in a properly designed optical network and therefore its effect on timing performance is considered negligible [89]. Some of these effects include nonlinear inelastic scattering processes like stimulated Raman scattering and stimulated Brillouin scattering. Others are the consequence from interactions between a high power optical signal and the glass fiber material, like self-phase modulation, cross-phase modulation and four-wave mixing.

2.5.1.5 Sagnac effect

The Sagnac effect, named after the French physicist Georges Sagnac [90], is an interferometric effect that is observed in a system rotating relative to an inertial frame of reference. Essentially, it demonstrates the effect of rotation on the speed of light. If we imagine a stationary circular path and we send two beams of light from a given point in opposite directions, they will meet at the same point (assuming that both beams have identical optical properties and the medium is ideal). However, if the circular path is rotating, the light beam traveling in the same direction as the rotation has to cover a longer path than the light traveling in the opposite direction, causing a phase shift when the two beams are combined again.

Although this effect is not relevant for data transmission purposes, it becomes of importance for long-distance time transfer methods due to the rotation of the Earth. Under the scenario here described, an East-West link can be seen as a fragment of a rotating fiber ring along the Earth surface, and the beam of light traveling in the East-West direction will take slightly less time than the West-East counterpart. The impact depends on the topology of the link and the latitude, but the contribution to error is estimated as 500 ps for every 100 km in worst case conditions [91].

2.6 Mechanisms and methods involved in redundancy and resiliency

This section presents a discussion of the methods and mechanisms that can be used in packet-based networks to increase the reliability and resiliency of a time transfer service deployed on top of it. The analysis is based on two principles: ensuring that a single failure of a link or device does not entail that the whole system is vulnerable to downtime or general failure, and ensuring that, in case of a link failure, the nodes of a hierarchical timing network can continue operating autonomously for a limited period of time. These two are necessary but not sufficient requirements in order to switch from one primary time source to another in the case of failure or disruption, be it manually or automatically.

2.6.1 Redundant networks

This section introduces the principles, methods, protocols and terminology towards the end of achieving reliable and resilient time transfer. There are several redundant network topologies that offer more than one physical path to ensure network

availability. These are briefly described here:

Mesh topology

Mesh topologies are described in various ways, ranging from extreme definitions where each device is interconnected to every other device, to more broad definitions that encompass any network that allows for some level of link redundancy [92], [93]. The former rarely applies in modern networks, while the latter can arguably envelop other topologies. Yet for our purposes, we adhere to the broad definition, focusing on networks that feature at least a single redundant path between nodes, as such networks must accommodate the redundancy for seamless operation.

These mesh networks employ efficient load balancing and routing capabilities to optimize the use of their redundant pathways, managing the excess of potential communication routes. Without proper routing, data packets could congest the network via suboptimal paths, leading to latency issues and potentially causing system failure in adverse scenarios like broadcast storms.

Mesh network redundancy is typically managed using various iterations of the Spanning Tree Protocol (STP), or leveraging shortest-path algorithms via protocols like Shortest Path Bridging (SPB) or Transparent Interconnection of Lots of Links (TRILL).

The STP family, including its Rapid Spanning Tree Protocol (RSTP, IEEE 802.1D-2004) and Multiple Spanning Tree Protocol (MSTP, IEEE 802.1q) variants [94], [95], are designed to prevent loop formations in Ethernet networks. They operate by blocking redundant links to establish a single active route between network devices. If a link fails, these protocols recalibrate the active routes to sustain traffic flow without loop creation.

Conversely, SPB employs a link state protocol enabling network devices to exchange information about the network's topology. Each device builds a full network map, and the protocol identifies the shortest route between any two devices [96]. SPB uniquely allows all links to remain active for traffic forwarding, ensuring more efficient resource utilization.

Link aggregation

This topology involves combining multiple physical links between network devices to form a single logical link with higher bandwith and redundancy capabilities. Usually, these kind of topologies are not only used for the purposes of redundancy but

also because they provide load balancing capabilities. Switch-to-switch aggregation is a specific type of link aggregation that is used to connect two switches together to create a higher-speed connection. The standard protocol for link aggregations is the Link Aggregation Control Protocol (LACP), which was last reviewed for IEEE 802.1AX-2020 [97], and was initially derived from the proprietary implementation by Cisco, EtherChannel. Other proprietary implementations of link aggregations are tied to the network equipment vendor, as in the case of Split Multi-Link Trunking Technology by Avaya, or Juniper's Aggregated Ethernet.

Ring network

A ring is a network topology in which each device is connected to exactly two other devices, forming a closed loop. This disposition ensures that all the devices are connected even after the failure of one of the links, while reducing the cost of cabling. However, the reduced cost in cabling and network equipment may be counterbalanced with the increased complexity of the end nodes, that need to integrate switching abilities.

A ring network can be seen as a special case of mesh network and, therefore, can operate on the basis of the spanning tree protocol by blocking a single link of the network. However, there are protocols specifically designed to leverage the ring topology to the point of achieving seamless redundancy. High-availability Seamless Redundancy (HSR), as mentioned in section 2.2.1.3, is a protocol specifically designed to ensure high-availability and real-time performance in industrial automation and control systems. The protocol achieves fault tolerance by relying on the parallel transmission of data over the two redundant paths of the ring topology. It was standardized under the IEC 62439-3 [33] and aims to achieve zero recovery time in case of failure. In an HSR network, every frame transmitted by a node is duplicated, one for each path of the ring. The receiver checks for duplicates and discards the second frame. As a result, HSR guarantees that at least one copy of the frame will reach the destination node even in case of a single link failure.

Parallel redundancy

The standard IEC 62439-3 does not only define HSR but another protocol to achieve seamless redundancy by means of duplicated network infrastructure. By connecting nodes through two disjoint parallel networks, the Parallel Redundancy Protocol (PRP) can be deployed with the same goal of zero recovery time in the event of

network failure.

In a PRP network, each node is connected to two different local area networks (LANs). When a source transmits data, it sends duplicate frames simultaneously over both networks. The receiving node will identify and eliminate the duplicate frame while ensuring that one copy of the frame is received, even in the event of failure.

The main disadvantage of this approach is the higher cost of deploying and maintaining two physically disjoint networks, although it provides more flexibility than the ring option. To maximize the reliability, other aspects of the network such as power sources, manufacturer of equipment and configuration can be managed independently to avoid any common vulnerability to affect both networks simultaneously.

Detailed operation of the PRP protocol can be found in [33].

2.6.2 Holdover

The ability to transfer frequency and time from a reference device to a *slave* device hinges on the slave's capacity to periodically measure its accuracy with respect to the reference and adjust its local clock and timebase accordingly. However, when the communication link between master and slave is disrupted, this comparison cannot be performed. In such cases, holdover serves as a mechanism ensuring continuous operation and maintaining a stable time and frequency reference when the reference clock is unavailable or compromised. The primary goal of holdover mode is to maintain the time and frequency output of the signals within acceptable accuracy and stability limits for a defined period until the reference signal is restored, or, in worst-case scenarios, provide sufficient time for an orderly downtime preparation that minimizes service impact.

Holdover relies on a local oscillator that has been previously synchronized to the primary reference. When the device enters holdover mode, the local oscillator maintains its frequency and time outputs based on the historical data acquired during the time of normal operation, as well as on information provided by local sensors.

Based on this definition and the concepts defined in section 2.3, Chapter 4 expands on the analysis of challenges and limitations that a holdover implementation must consider, and presents an experimental holdover solution for a time transfer network.

2.6.3 Switchover and failover

In the context of redundant networks, especially when these are a vital part of ensuring the availability of the network in case of disruption, switchover and failover are

the mechanisms that allow for seamless transition between redundant components, while maintaining the stability of the network and the services deployed atop.

The difference between these two concepts resides on the event that triggers the transition. A switchover is a planned, controlled transition of network traffic from one component or path to another. Switchover situations are typically executed to facilitate maintenance, upgrades, or optimizations in the network infrastructure while minimizing the disruption of the services that the network provides. In the specific case of time networks, switchover can happen in terms of a different time source. In contrast, failover is an automatic reactive process that is triggered upon the detection of a failure or performance degradation in the system, aiming to ensure service continuity or minimize the impact of the disruption.

Successful failover or switchover operations rely on the presence of well implemented redundancy mechanisms and failover. In the case of failover, the device that detects the performance degradation or failure needs the redundancy mechanism to promptly (or even seamlessly) provide access to the backup resources. Throughout the reconfiguration time, an adequately designed holdover mechanism complements the effectiveness of failover and switchover operations by mitigating the degradation of the time signals by maintaining a stable and accurate clock signal until the visibility of a reference time source is restored.

Chapter 3

Materials, methods and tools

This chapter begins by introducing the materials used in the analysis and development of the thesis project. As the project's objectives are centered around achieving long distance synchronization over an optical medium using the White Rabbit (WR) protocol, we present the specific materials for these applications as well. Next, we describe the development framework and tools utilized to achieve the scientific and technical objectives outlined in section 1.2, along with the rest of the equipment and measurement instrumentation employed. In addition to describing the materials (both hardware and software) involved, we provide a general overview of the commonalities shared by all experimental work in subsequent chapters, as well as defining the indicators used to quantify the performance of timing signals generated by the developed solutions, and discussing the metrics and statistical figures that are relevant for analyzing the resulting data.

3.1 Hardware platform: the White Rabbit Switch

The deployment of the WR protocol requires the use of dedicated hardware. For the purposes of implementation and development in this project, the White Rabbit Switch (WRS, Figure 3.1) has been selected as the reference hardware platform.

The WRS follows the principles of openness that are one of the main features of the WR project. The design of the WRS was commissioned by CERN to Seven Solutions, and its hardware design and firmware are openly available[1]. The WRS, in addition to its data switching capabilities, includes the necessary hardware elements

[1] `https://ohwr.org/project/wr-switch-hw/wikis/home`

Figure 3.1: The White Rabbit Switch.

to support WR operation in all 18 of its ports. With its ability to disseminate its time signal using off-the-shelf networking gear and standard protocols, with selected components the device can achieve a distance of more than 100 km in a single hop while keeping the time accuracy below one nanosecond. The internal architecture of the WRS is comprised of an Advanced RISC Machines (ARM) processor that operates on an embedded Linux OS, and a Xilinx Virtex-6 FPGA. An embedded Linux OS hosts the necessary drivers, configuration tools and general purpose services for the WR implementation. The programmable hardware contains, among other elements, the Dual-Mixer Time Difference (DMTD) phase measurement modules that provide enhanced timing capabilities, Ethernet endpoints, switching core logic, and the Real Time (RT) subsystem responsible for time-sensitive operations. The RT subsystem includes the logic that generates the local time signals of the device, and a LatticeMico32 (LM32) softCPU that handles real-time phase measurement data and the syntonization mechanism. The next sections expand on this brief overview.

3.1.1 Software

The WRS software architecture comprises several daemons running on the embedded Linux system, operating on the ARM processor. These daemons interact with the FPGA-implemented devices via an External Bus Interface (EBI) bus, with the FPGA serving as a software accelerator for the software in terms of time syn-

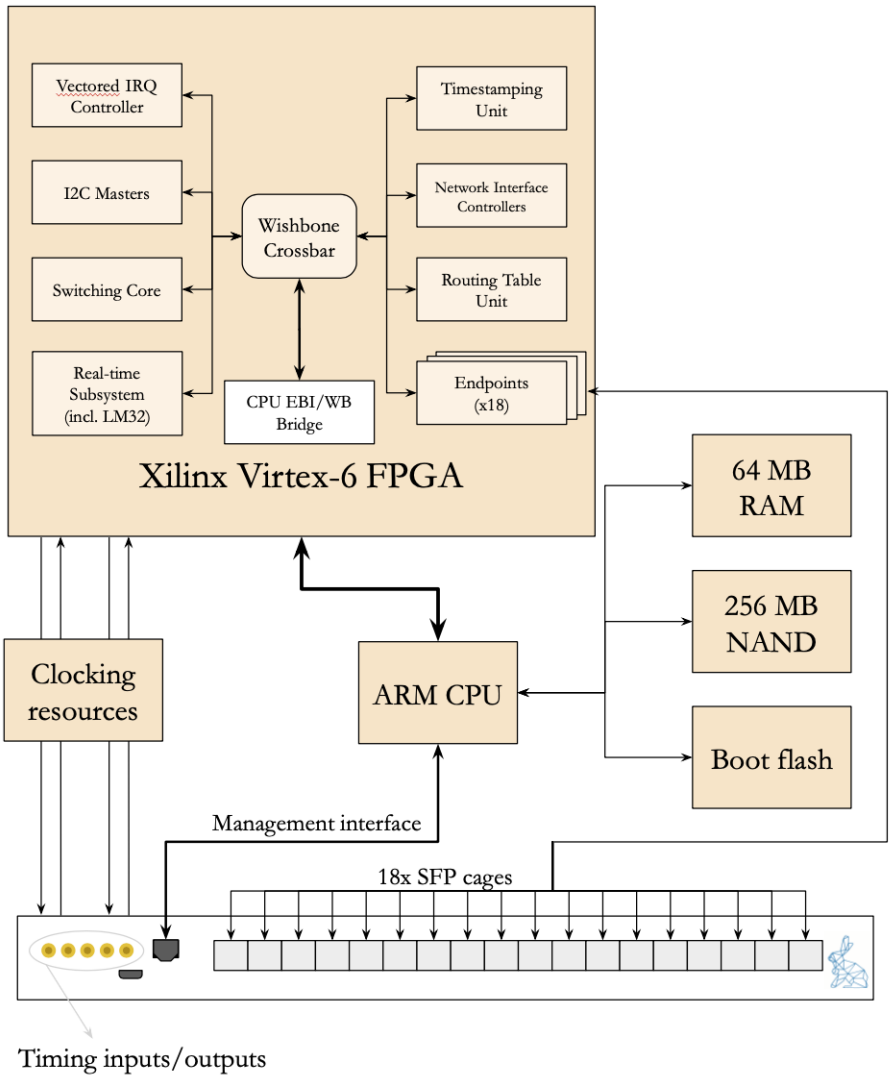

Figure 3.2: High level diagram of the WRS.

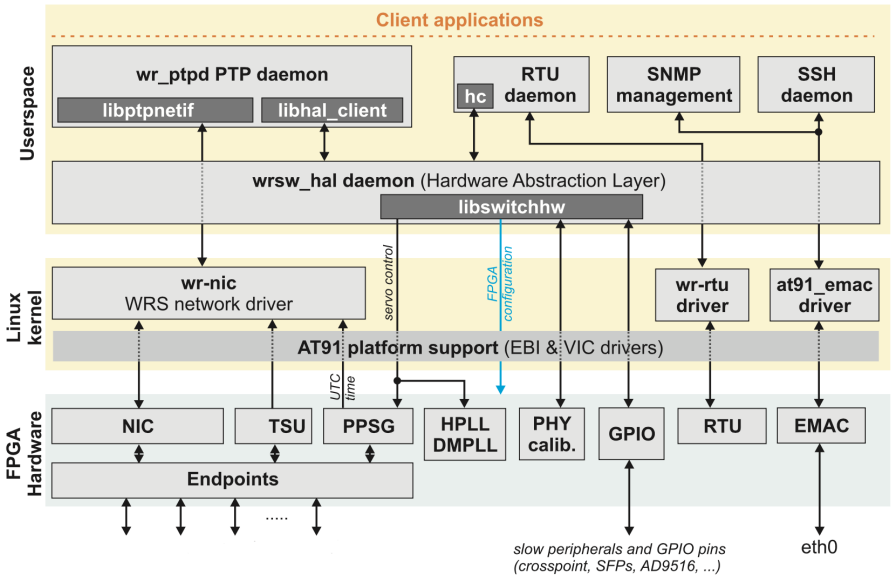

Figure 3.3: Software architecture of the WRS. From [98]

chronization, data switching and management of hardware components. A Hardware Abstraction Layer (HAL) manages the communication between the various software components and the FPGA.

The ARM CPU runs several key daemons that are relevant to this work. The PTP daemon (PTP Ported to Silicon, PPSi) retrieves hardware-generated timestamps from the timestamping unit (TSU) and implements the WR-PTP protocol, while the routing table unit (RTU) daemon manages the routing table governing the Switching Core. Figure 3.3 shows the general architecture of the WRS software and its relationship with the FPGA-implemented modules and the rest of the hardware. The pulse-per-second generator (PPSGen) module implements the timebase of the device, serving as the *clock* of the device, and the low-level phase syntonization is complemented by the round-trip phase measurements computed by the PTP software daemon and communicated to the software-implemented PLL via the HAL mechanisms (represented in the figure as HPLL and DMPLL for *helper* and *main*).

3.1.2 Hardware and FPGA architecture

The WRS consists of two boards: the switch core board (SCB) and the Mini Backplane. The SCB is the main printed circuit board (PCB) that contains the FPGA and CPU components, as well as the clocking resources specific to WR. The Mini Back-

plane is an extension of the SCB board, designed to host 18 SFP cages in the final device. Figure 3.2 provides a high-level diagram of the switch's architecture.

The WRS FPGA is responsible for managing the critical components of the WR protocol that are closely linked to the physical layer, such as phase measurement, timestamping, and syntonization. Additionally, it handles the switching logic that is essential to the operation of an Ethernet switch. The FPGA is comprised of several submodules that are briefly explained as follows.

- The RT subsystem is responsible for managing the time synchronization components. It includes the LM32 SoftCPU that implements the SoftPLL software, the hardware side of the SoftPLL, and the PPSGen module that manages the device's time counter registers and generates output time signals.
- The Endpoints module includes 18 ports, each of which is dedicated to implementing physical coding sublayer (PCS) and Medium Access Control (MAC) functionality for the Ethernet ports. It interfaces with the Gigabit Transceivers and generates precise timestamps for incoming and outgoing frames.
- The network interface controller (NIC) serves as an interface between the ARM Linux kernel and the FPGA-implemented endpoints.
- The RTU learns the routes of through traffic based on the operation of the address resolution protocol (ARP) and can also be configured manually.
- The Switching core stores and forwards frames according to the rules defined by the RTU.

For more in-depth information on the WRS architecture, readers can refer to [98].

3.1.3 White Rabbit Switch - Low Jitter & Expansion Board

One advantage of the open hardware initiative is that users are able to fork, adapt, and improve existing platforms to better suit their needs. In a previous study, Rizzi et al. [69] analyzed the limiting factors of the WRS in terms of noise, stability, and frequency transfer characteristics. The study concluded that several feasible and straightforward adaptations could be made to improve the WRS's performance. As a result, certain experimental developments were made to enhance the short-term stability of frequency transfer, which were later published in another paper by Rizzi et al. [70] Many of the improvements proposed in these publications, among others, have been

incorporated into a new revision of the WRS main board design. This new product is called the White Rabbit Switch - Low Jitter (WRS-LJ)[2]. Throughout this thesis, the relevant limitations of the WRS and the improvements implemented in the new WRS-LJ design will be reviewed in greater depth. Notably, the WRS-LJ includes a new connector that accommodates an expansion board, enabling the device to use different experimental oscillators as the reference clock for the device, which can enable new functionalities regarding the resilience of the time service. The novelties of the WRS-LJ and its associated expansion board will be explored more deeply in Chapter 4. Figure 3.4 highlights the novelties of the new hardware design.

Beyond the WRS, Seven Solutions offers a broader family of devices under the WR umbrella. The WR-Z16, WR-ZEN, and WR-LEN are among these other members, expanding the WR device landscape for diverse needs and use cases. While our primary focus in this book is the WRS, we will occasionally touch upon these other devices. Whenever such references occur, we will clearly specify the particular device in question to maintain clarity and precision.

3.2 Development environment

The development environment for working with the WRS is complex due to the need to manage the FPGA-implemented modules, ARM, and the SoftCPU inside the FPGA simultaneously. Communication between these subsystems is critical for real-time processes such as syntonization and phase measurement. This communication occurs through a Wishbone bus structure to which multiple registers and devices are connected.

In addition to the internal communication challenges, each subsystem must be programmed using specific toolchains and Integrated Development Environments (IDEs). The Virtex 6 series of FPGAs requires the use of Xilinx ISE[3], a software tool for the synthesis and analysis of Hardware Description Language (HDL) designs. Although Xilinx ISE has been discontinued since 2013, its use is necessary when targeting FPGAs from the 6-series, such as the Virtex inside the system being developed. Alternatively, the Xilinx Vivado design suite can be used for more recent FPGA series. The design flow for Xilinx ISE is similar to Vivado and involves initial synthesis of a solution from source code and constraint files, followed by optimization, mapping, placement, and routing of the synthesized design in a model-specific implementa-

[2] `https://ohwr.org/project/wrs-lj-hw/wikis/home`
[3] `https://www.xilinx.com/products/design-tools/ise-design-suite.html`

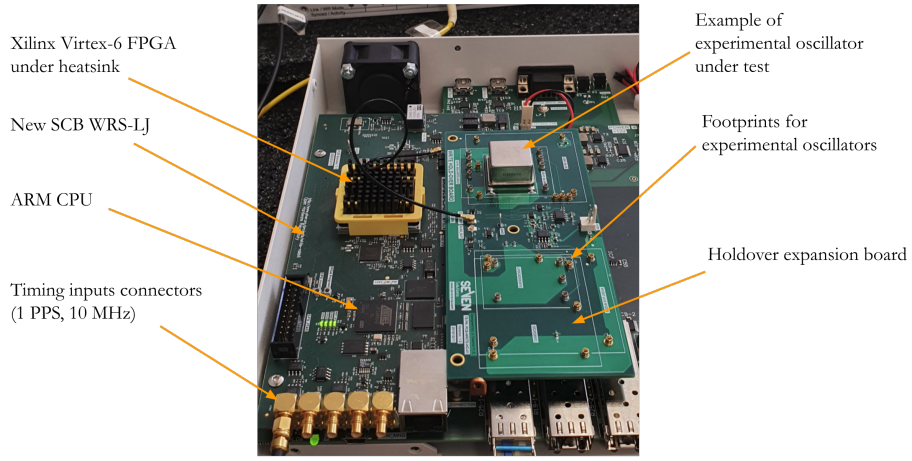

Figure 3.4: Detail of the SCB and holdover expansion board of a WRS-LJ with its top cover open and front casing removed. The main features are highlighted.

tion phase. The final step involves generating a binary file for programming into the FPGA. The ARM Linux system is generated using Buildroot[4], a tool that automates the process of building and customizing embedded systems by providing configuration tools, handling dependencies, generating a bootloader, and providing a cross-compiled toolchain. The LM32[5] SoftCPU, on the other hand, does not contain any OS and the software that runs is completely standalone. Therefore, the LM32 toolchain is the only tool needed to obtain and deploy the software.

In terms of debugging tools, there are two areas of focus: the SoftPLL code and FPGA design modifications. The SoftPLL code is a bare metal application that runs in the HDL-described CPU and does not have the same debugging tools as those available for other options more integrated with the hardware vendor such as Microblaze CPUs using the Xilinx XSCT tool. Instead, the debugging process has relied on informal methods such as printing values and monitoring the thousands of phase measurements that are generated every second through a custom visualization tool that captures all SoftPLL data piped out via a User Datagram Protocol (UDP) data stream.

For FPGA design modifications, ISE provides simulation capabilities whereby the user provides the source code of the module and a suitable testbench to serve as stimulus for the module under test. Xilinx also provides debugging *Intellectual Property* (IP) cores, known as ChipScope ILAs (Integrated Logic Analyzers)[99], to probe certain

[4] `https://buildroot.org`

[5] `https://www.latticesemi.com/en/Products/DesignSoftwareAndIP/IntellectualProperty/IPCore/IPCores02/LatticeMico32.aspx`

signals and sample them with a set periodicity. However, informal methods were used to test the modifications and developed modules due to two limitations. The first limitation concerns the challenge of considering the vast amount and complexity of the stimulus required to simulate network transactions involving PTP and enhanced phase measurement of the WR protocol. Additionally, accurately reproducing the random noise conditions present in real-world scenarios poses an even more significant challenge. Even if the complexity of developing the testbenches were something that could be faced, the computing power that would be needed to effectively obtain useful results would be prohibitive for currently available computing resources. As a result, validation of FPGA module operation was conducted using a real prototype whose correct behavior was verified by probing the status of internal signals using ChipScope debugging tools and producing stimulus from real hardware.

The design and development process of the software and FPGA firmware in this study can be modeled as a multifaceted iterative workflow as depicted in Figure 3.5. This process begins with the system design phase, where the solution is translated into an initial code draft. During this stage, high-level designs and requirements are transformed into an actual software and firmware code iteration that can be implemented on the LM32 SoftCPU and Virtex-6 FPGA respectively. Architectural concepts, data flow diagrams and functional specifications help envision the structure of the system before addressing the building phase.

The initial coding framework is then developed, elaborated and refined. This is where the IDE of choice, the Xilinx ISE tools and the LM32 toolchain come into play (*System Design* in Figure 3.5). ISE provides a comprehensive environment to implement and analyze the FPGA-based designs. It allows the translation from the hardware description using languages such as VHDL or Verilog into the components and connections that map the mentioned design into the FPGA components. At the same time, the LM32 toolchain provides a suite of tools for programming the LM32 softCPU that is part of the FPGA project, including compilers, assemblers, and linkers.

Following the System Design phase, the system transitions into the *Implementation & analysis stage*, which involves a close evaluation of the designs as implemented on the hardware. The designs are synthesized and mapped into the FPGA, followed by an elaborate timing and methodological analysis to ensure that the design meets the system requirements. At this point, any issues or bugs that arise are addressed, resulting in a series of iterations that incrementally improve the design's performance and functionality.

The final stage is the deployment of the software and firmware into the actual

hardware. After successful implementation and analysis, the system is tested on the WRS to ensure it operates as expected. Here, our debugging tools like ChipScope Pro are critical in examining internal signals in the FPGA, but the performance of the timing solution is also measured using real world equipment and laboratory instrumentation. The cycle then repeats starting with design changes, implementation and analysis, and deployment until the system's performance meets the defined requirements.

3.3 Instrumentation and other materials

To ensure that the developed solution functions as intended and monitor its performance, external laboratory equipment validation was used. Such equipment is expected to be calibrated, functional, and unbiased, making it a reliable method of validating the obtained solutions. To simulate realistic scenarios where long distance WR time dissemination is required, typical network elements found in production networks and long fiber spools can be utilized to establish Ethernet links spanning hundreds of kilometers while staying in the lab environment and having both ends of the link available for the measuring instruments. While the long distance equipment is only briefly described here, its physical significance will be extensively discussed in Chapter 5.

3.3.1 Long distance testbench

The network equipment used in the Ethernet links forming the backbone Network or the Core Network differs from that commonly used in consumer and enterprise grade LANs. In this work, long distance experiments were conducted using CWDM and DWDM SFPs, which convert the electrical Ethernet signal into the optic domain for transportation. CDWM is typically used in WR links for distances up to ten kilometers, whereas DWDM is prevalent in the Backbone network for effectively multiplexing large amounts of Ethernet data streams into a single-mode fiber.

The use of DWDM equipment in conjunction with access to a series of G.652 D. fiber spools of lengths ranging from 5 to 50 kilometers enabled the simulation of links of up to 150 kilometers in the laboratory. These fiber spools are accompanied by DWDM add/drop filters, dispersion compensation modules (DCMs), erbium-doped fiber amplifiers (EDFAs), optical power meters, and other auxiliary elements. Depictions of these elements can be found in the corresponding section of Figure 3.6.

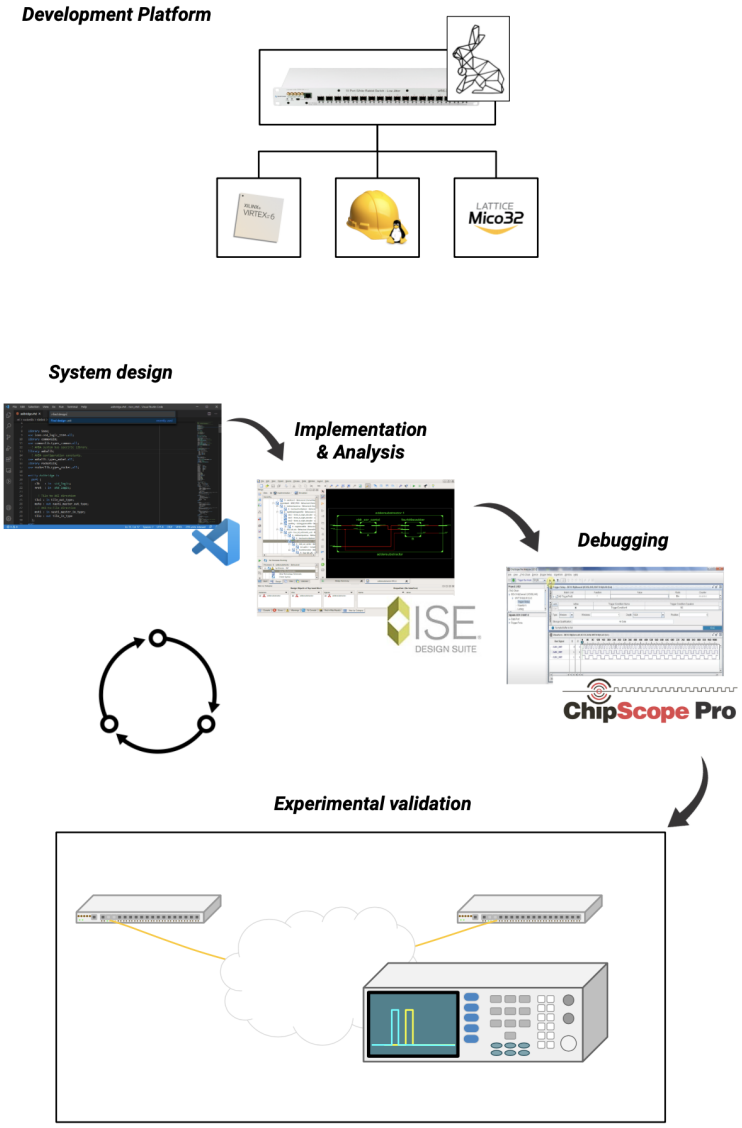

Figure 3.5: Development platform and stages of the development methodology. The core components of the WR firmware are highlighted (Xilinx Virtex 6 FPGA, an LM32 softCPU and a Buildroot embedded Linux environment). The figure also depicts an schematic view of the iterative development process from system design to functional verification.

Please note that this brief metion of the long-distance testbench is just a part of the methodology description. However, the comprehensive details and in-depth explanation of this testbench will be provided in Chapter 5. The mention here is solely for context, to give a more complete picture of our methodology.

3.3.2 Laboratory instrumentation

Throughout the development and validation phases of this project, we primarily utilize instrumentation to measure time differences and other significant statistics in analog electrical signals. This purpose requires the use of laboratory equipment ranging from the most general and versatile, to some more specialized devices for this field (bottom part of Figure 3.6). To this end, we employed digital oscilloscopes as the *staple*, versatile tool for determining not only time properties of signals but also as a debugging tool that confirms the presence and levels of the analog signal.

Additionally, we utilized time and frequency counters, such as the Keysight 53230A[6], to acquire data from periodic signals like frequency, time offset, and their associated statistics. The two main benefits of using this equipment are the effective absence of dead time between triggering events, and the ability to program data acquisition for periods ranging from seconds to days. This allows the user to extract relevant events in a comma separated values (CSV) file that can be further processed in software.

Finally, we also utilized a Microsemi (now Microchip) 3120A Phase Noise probe[7]. This device can measure the amplitude, phase and frequency stability of a clock signal and generates complex statistics that are useful for the analysis of the obtained solutions. The significance of the parameters here described is discussed in section 2.3.3.

In addition to measurement instrumentation, we employ a climate chamber to simulate varying environmental conditions and assess their impact on the accuracy of the time distribution solution. Specifically, we expose parts of the network deployments to changing environmental conditions, as modifications in these conditions can significantly affect the solution's accuracy.

[6] `https://www.keysight.com/us/en/product/53230A/350-mhz-universal-frequency-counter-timer-12-digits-s-20-ps.html`

[7] `https://www.microsemi.com/product-directory/phase-noise-and-allen-deviation-testers/4131-3120a`

3.3.3 Clock references

In a time synchronization system, reference oscillators play a crucial role, not only in devices that disseminate time signals but also in the verification instrumentation used to assess accuracy and stability. However, these elements have inherent limitations, since they are electronic components subject to imperfections during fabrication and various random processes that generate different types of noise during normal operation. Noise in oscillators is a complex field beyond the scope of this work, although some notions are briefly provided to help frame its relevance in the experimental methodology section.

Our work employed multiple oscillators from different sources, including temperature-compensated crystal oscillators (TCXO), oven-controlled crystal oscillators (OCXO), a passive hydrogen maser (PHM) from the Time and Frequency Transfer Laboratory at *Universidad de Granada*, and other oscillators such as rubidium standards and miniature cesium clocks in collaboration with partners of the WRITE project.

Having oscillators based on different physical principles is significant because no oscillator is perfect. Referring back to the observations made in section 2.3.4, each type of oscillator suffer from noise, albeit at different levels within different periods. For instance, while the signal from an OCXO is more stable than that from a cesium standard in the short term (i.e., seconds or minutes), it will drift further apart in the medium or long term, being outperformed by the cesium standard.

3.3.4 Software tools

This section applies mainly to further processing the data obtained from the lab instrumentation. This was performed using general purpose software and scripting using, among others, MATLAB, GNU/Octave, bash scripts, and Python. Two pieces of software that must also be mentioned due to the provided service are the AllanTools[8] (a set of Python scripts to process and plot stability statistics of time-related signals) and the instrument software of the Microsemi 3120A Phase Noise probe.

3.4 Experiment design and methods

The successful design and execution of experiments in any scientific field is highly dependent on the thoroughness and quality of the preliminary research conducted. In the context of studying synchronization solutions over distant locations, a critical

[8] `https://allantools.readthedocs.io/en/latest/index.html`

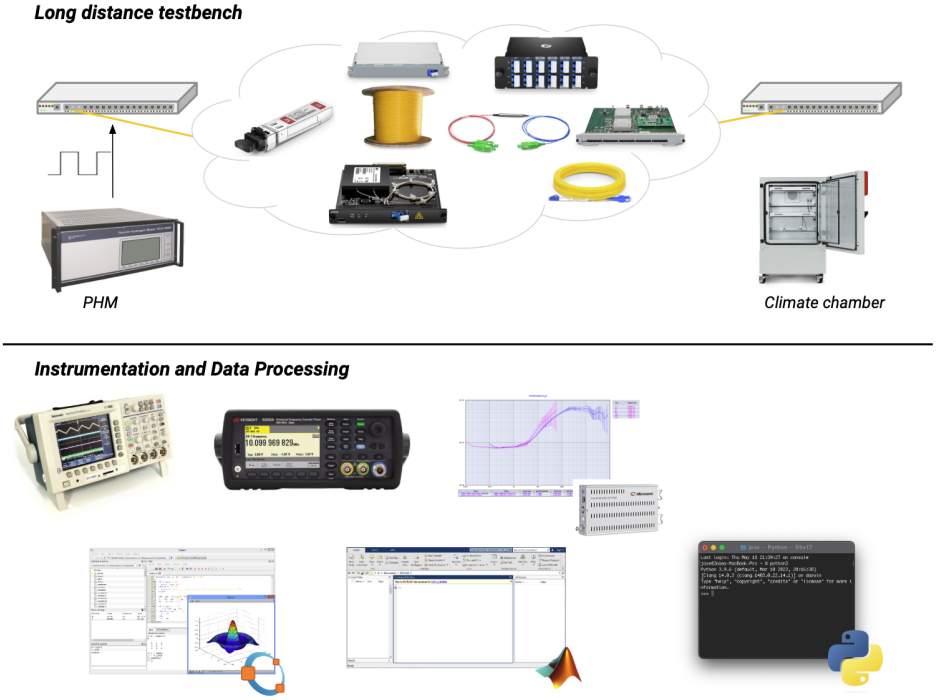

Figure 3.6: Representation of elements of the long distance testbench, laboratory instrumentation and software tools used in this study. In the top half, the long distance testbench, whose components are detailed in Chapter 5, is complemented by the use of the PHM frequency reference, as well as the climate chamber that allows the modification of environmental conditions. In the bottom half, the lab instrumentation is represented, along with screenshots and logos of the main software tools used to process the data obtained from instrumentation.

component of this preliminary research is the state-of-the-art review. However, this analysis is not merely a static summary of existing research findings at a given point in time. Rather, it is a neverending and iterative process that can greatly inform and enhance the experimental and methodological design.

The field of application of this work is constantly evolving, and new findings and innovations are regularly emerging. Therefore, it has been essential to revisit and update the state-of-the-art review regularly to ensure that the research remains relevant. Moreover, by keeping up with the latest developments in the field, the researcher can identify any knowledge gaps and potential areas for further investigation. These insights can be used to address these gaps and ultimately lead to more impactful and significant contributions.

The next section will integrate the most up-to-date information gathered from the state-of-the-art review with a meticulous experimental design, to tackle the significant challenges and meet the scientific objectives established earlier. This will be preceeded by a section containing basic definitions that will help set common ground between the reader and the author.

It is important to note that this section will provide a broad overview of the experimental design, while the specifics will be discussed in the corresponding sections. Before moving on to the overview of our experimental setup, the reader is encouraged to review the time concepts, naming conventions and metrics discussed in Chapter 2.

3.4.1 Experimental design Overview

Although all experiments and tests conducted in the context of this project are related to evaluating the performance of a time transfer setup, there are distinct differences among the experiments in each chapter or topic.

- The first part of Chapter 4 experiments focus on the performance of the slave node of a link when the connection to the master is lost. In this context, the most relevant metric is the accumulated time error since the link was lost, and the stability of the time signal.

- In contrast, results involving the redundancy of the data paths not only involve the measurement of the time transfer accuracy, but also the reaction time in the case of a link disruption and other network health and performance parameters.

- Experiments of Chapter 5 are centered around the influence of environmental conditions of a functioning long distance WR link. Therefore, the experiments

involving this chapter will be done over a working WR link in which the components are put under certain stress. Under these premises, other variables beyond time transfer accuracy such as link delay, determinism, indirect measurements of dispersion and temperature coefficients are significant.

- In Chapter 6, an integral Time as a Service (TaaS) deployment is evaluated under real world conditions, therefore the usual access to both ends of the link do not apply and other approaches (i.e. auxiliary GNSS-based time transfer) that posed their own challenges had to be followed.

Time transfer performance

The essential metric that is evaluated in virtually every work in this domain is the performance of the time transfer, which is determined based on the results of an experiment as depicted in Figure 3.7. The broadest description of the setup configuration encompasses a master device that serves as time reference, and a slave device that aims to synchronize its clock to the master. Depending on the specific configuration, there might be more than one node between the master and the slave, or there may be singular features in the devices or the network in between that make the setup worth of research. With this setup the goal is to evaluate the time error between the devices, and to obtain a large data series that facilitates the characterization of the time signal, and the identification of possible anomalous phenomena including glitches, drifts or any other.

The time error between the devices is computed externally by means of the one-pulse-per-second (1-PPS) analog output. This is an electrical signal that rises abruptly once every second, and the precise timestamping of the rising edge can be determined through employment of a time and frequency counter (or an oscilloscope). It is crucial to consider the length of the cables utilized, as even a 20 cm segment of coaxial cable results in an additional delay of one nanosecond, which is significant when dealing with the levels of accuracy that are dealt with in this context.

The electrical properties of the analog signals are equally important. The slew rate (change of voltage per unit of time) is limited by the properties of the signal amplifiers in use. To ensure the slew rate does not introduce a significant source of error, it is preferred to use identical hardware in both the master and slave components when possible. Moreover, the instrument's acquisition setup must be configured with care to guarantee precise timing. Specifically, the trigger should be set to occur when the slope of the 1-PPS signal is at its maximum, which minimizes uncertainty.

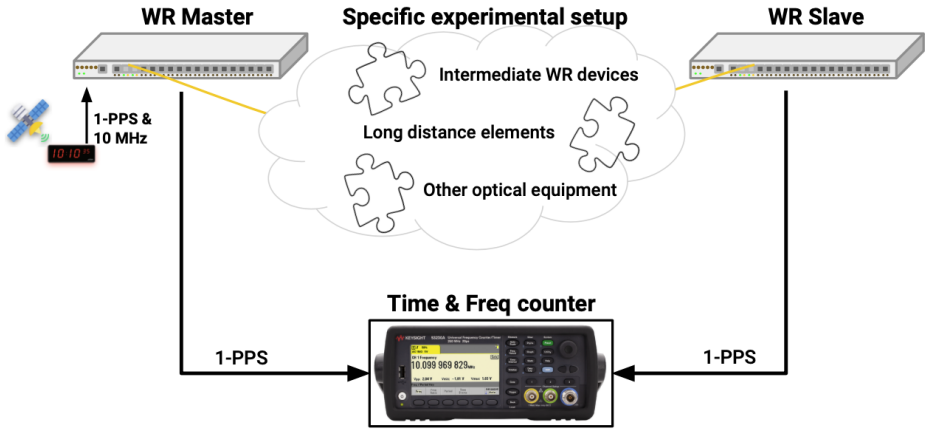

Figure 3.7: This diagram shows a schematic view of a setup in which time synchronization data can be recorded using a time and frequency counter.

Taking into account all these factors, the time and frequency counter can serve as a useful tool for acquiring a data set of the 1-PPS time error between the master and slave devices, evaluated for every occurrence of the 1-PPS signal. This time series allows for a quick assessment of whether the system is synchronized, especially if the dispersion between occurrences is low or at least constant in the case of uncalibrated devices. However, this data can also be analyzed using more advanced statistical tools such as the ADEV and the MTIE to gain a deeper understanding of the properties of the time transfer setup regarding noise and stability.

Holdover evaluation

Starting from the most basic setup described in the previous figure (single , direct link between two devices), the communication between the master and the slave devices will be interrupted to force the slave into holdover mode. This means that the slave device will attempt to maintain its output frequency and stability without new data from the reference signal. The data collected by the time and frequency counter will have the same nature as in the previous setup, measuring the time error between the master and the slave. However, there are two important considerations to address in this test.

Firstly, the measurement of the device under holdover mode is strongly influenced by stochastic processes that alter the instantaneous frequency of its system oscillator. This means that replicating a given test under the same environmental conditions can result in different outcomes. To characterize the typical and worst cases, it is necessary

to repeat experiments multiple times.

Secondly, the stability of the master reference is a limiting factor in this setup. The oscillator of the master device is also subject to various types of noise and environmental factors that impact its frequency stability. This creates a significant issue when trying to measure the time error between the slave and the master. To address this problem, it is very important to have access to a reference source that is at least one order of magnitude more stable than the device under test. Otherwise, instead of characterizing the stability of the slave, the instabilities of the clocks at both ends of the link will be combined and added up in the data acquired.

Fault-tolerant network topologies

The last part of the chapter dedicated to resiliency mechanisms, that is, section 4.6 discusses high-availability and fault-tolerant network topologies and protocols. In addition to timing considerations, network data performance is also a relevant factor to evaluate in these topologies. Specifically, in a high-availability redundant ring topology, nodes must be able to forward frames quickly to minimize the delay introduced by each node of the ring, and the redundancy data protocol mechanisms must be processed efficiently to reduce resource consumption destined to storing frames so to deliver frames to the destination application as soon as possible. The three metrics considered for these tests are data latency, frame loss and data bandwidth.

To test the performance of this topology, a network comprising six WRS devices, each with custom HSR-capable firmware is set up. One of these devices has the role of grandmaster, while the remaining devices are doubly synchronized slaves. Data latency is characterized by using the internal timestamping capabilities of WR devices, while the bandwidth and frame loss tests is determined by a using general-purpose network benchmarking tools between two PCs attached to different nodes of the ring. One PC injects data through a UDP stream, while the other acts as receiving end and evaluates the integrity of the received data.

The timing analysis focuses on assessing the scalability of time transfer accuracy across all the nodes of the ring, as well as the switchover capabilities in the event of a network disruption on one link.

GNSS-assisted measurements

In the context of Chapter 6, ensuring traceability to UTC and determining the time error between the location of the time service customer and the time source are

challenging due to the distance between them. Although the viability of the deployed service has been verified in lab environments that can recreate an optical link spanning hundreds of kilometers, there is still uncertainty due to the independent ISP that owns the medium and can change the apparent length of the link. This can be disruptive for the timing service, even though it is harmless for the ISP's typical service.

One possible approach to overcome this limitation is deploying another WR link back to the point where the time source is located, via different locations. This would allow us to verify that the time signal can be deployed through different points of the network and travel back to the point of origin, where we could measure the time error between the time source and a slave device colocated where the time master rests. However, the cost of renting exclusive DWDM channels from an ISP makes this approach impractical.

Therefore, we opt to check the WR time service against a GNSS time source located at the premises of the customer. Using calibrated components that are verified by a metrology institution, we can guarantee traceability to UTC within a certain margin. However, this approach suffers from limitations typical of GNSS clock sources, such as the delay of the GNSS signals traveling through the ionosphere that cause a time error with a daily periodicity (bound to $\pm$ 10 nanoseconds). In the deployed setup, the 1-PPS time signal coming from the slave is compared to the 1-PPS generated by the GNSS time source using a time and frequency counter (Figure 3.8). It is important to observe that the GNSS timing source is a much larger source of instability

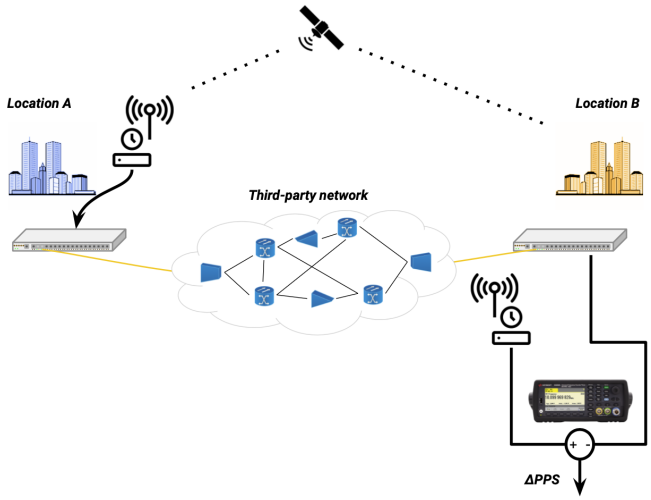

Figure 3.8: Schematic diagram of a experimental setup evaluating the performance of a long distance WR-based time service, assisted by a GNSS time reference.

than the WR link, and only the time average over a long-term is useful data.

This overview of the setup based on GNSS measurements is brief but provides a solid grasp of the overall methodology used in this study. Please consider that this section is just touching the basics. The finer details are saved for Chapter 6.

3.5 Conclusion

In this chapter, we have provided a comprehensive overview of the hardware platform, tools, laboratory equipment, and workflow that will be used to develop and validate the proposed solutions. Additionally, we have provided clear definitions of the concepts that will be used throughout the thesis to ensure that the author and reader share a common understanding. Overall, this methodology chapter sets the foundation for the rest of the thesis, providing a clear understanding of the research methodology, tools, and equipment used in our study. The three upcoming chapters build on this foundation to present a thorough analysis of the research problems and potential approaches to address them.

Chapter 4

Mechanisms for Improved Resiliency in the White Rabbit Switch: stability, holdover, and redundant capabilities

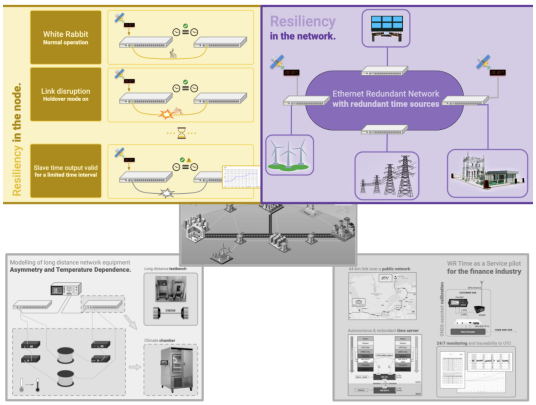

4.1 Introduction

The ability of White Rabbit (WR) to achieve sub-nanosecond accuracy has been proven under many different use cases. Some of which have been highlighted in the previous chapters, and many others that fall outside the scope of this study [67], [100]. Despite the proven effectiveness of WR, the technological landscape continues to evolve, demanding enhancements of the solution's robustness and reliability, particularly in challenging scenarios such as network reconfiguration or time source failure.

A critical feature that is required for the adoption of WR in new markets concerns the development of methods and algorithms that guarantee the stability and accuracy of the timing signals in the absence of an active time reference. The development of a new model of the White Rabbit Switch that improves the phase noise and stability of its time outputs is the basis for the development of a holdover mechanism atop. The reduced noise levels of the time signals, initially devised for more demanding scientific applications, serves as an ideal starting point for the reliability mechanisms, since the cleaner clock signal allows for reduced initial frequency error when no master time reference is visible. An overview of synchronization and syntonization in WR precedes the mentioned developements, which helps to accurately frame the problem that we aim to solve.

The chapter also outlines a series of mechanisms that have been developed to achieve seamless redundancy in a WR network with emphasis both timing and data. These mechanisms allows the WR devices to form a network that avoids single points of failure in selected topologies – an imperative in industrial networks, Smart Grid, and the Industrial Internet of Things. The implemented mechanisms enable devices to switch from one time reference to another without significantly impacting the quality of the time signals.

4.2 Background

4.2.1 Synchronization and Syntonisation in White Rabbit

One of the key features that enable WR to achieve its high level of accuracy is the close relationship between the physical layer and the time synchronization mechanism. The synchronization in WR is achieved by exchanging PTP messages between the master and slave devices, while the syntonization process is performed at the physical layer.

As detailed in Chapter 2, synchronization occurs when the internal time counters of both the slave and master devices display the same values at a given point in time. However, they will only continue displaying the same values if their internal clocks operate at the same frequency. To ensure this, a syntonization mechanism is needed between them. Often, PTP implementations only periodically correct the time counter values of the slave device but do not tune the frequency of the internal clock.

A functioning WR link relies on two simultaneous processes. First, the clock of

the slave node is syntonized to the clock recovered from the physical Ethernet signal sent by the master. Then, once syntonized, PTP frames are exchanged, enabling synchronization. Once synchronized, the slave phase-aligns with the master according to precise information about the link length contained in the PTP frame exchange. Once established, the link accuracy can be maintained by adjusting the clock's phase without altering the time counters. A prerequisite for phase alignment is that the devices must be accurately calibrated to account for any internal delays they may present.

The result of this process is that the slave clock remains phase-aligned at all times while the WR link is active. This clock also serves as the device's *reference clock*, which supports any applications running on it. These applications may be performing critical tasks that cannot tolerate disruptions in the clock they use, thus necessitating more reliable mechanisms. In this way, the applications can continue using the system's reference clock in case of system reconfiguration or lack of visibility of the master clock.

Since the most decisive element in maintaining WR synchronization is the phase tracking mechanism, the following section will describe it in detail. For more details in the general operation of the WR protocols, readers can refer to [98].

4.2.2 White Rabbit syntonization (DDMTD and SoftPLL)

This section provides an overview of the WR PLL design. The WR PLL can be considered as consisting of two main components: the Digital Dual Mixer Time Difference (DDMTD) and the SoftPLL. The DDMTD is implemented as physical programmable hardware modules that accurately measure the phase difference between two clock signals, while the SoftPLL serves as the core of the clock syntonization control system, processing the raw phase measurements and commanding the tuning of the clock frequencies.

The DDMTD is implemented in the FPGA and generates phase tags to measure the phase difference between the local clock and the master clock, which is recovered from the Ethernet bitstream. In Grandmaster operation, the DDMTD can also directly measure the phase difference of the external clock input. Two inputs of close frequencies are needed, and these are mixed and low-pass filtered in order to obtain a signal that has a lower frequency but maintains the same phase relationship as in the higher frequency ones. The use of slightly different frequencies to amplify the resolution of the measurement is a technique that is akin to a vernier caliper in the time domain. Further details on its implementation can be found in [68].

The phase tags are input into a control loop running on an embedded LM32 CPU, which implements a software-based Proportional-Integral (PI) controller. Figure 4.1

shows a simplified block diagram of the SoftPLL, which comprises two separate control loops: Helper and Main.

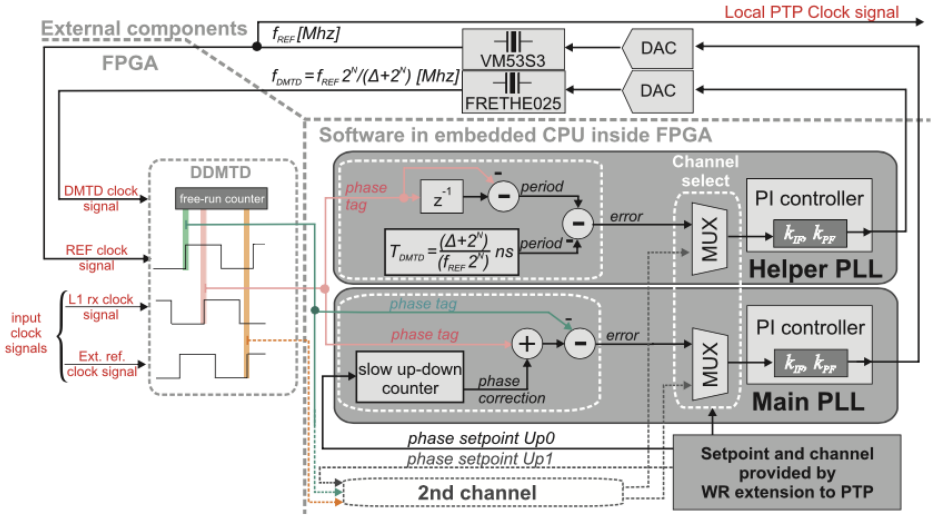

Figure 4.1: Overview of the WR phase-locked loop design. From [69]

The DDMTD requires a clock signal that is very close in frequency to the clocks whose phase relationship is being measured. To fulfill this requirement, the Helper PLL generates the DMTD clock. The Helper control loop functions by comparing the difference between consecutive phase tags to the expected ideal period of the DMTD signal, following the relation:

$$f_{DMTD} = f_{ref} \times \frac{2^N}{2^N + 1} \tag{4.1}$$

Here, f_{ref} represents the system's reference clock and N is a custom parameter that controls the frequency offset between f_{DMTD} and f_{ref}. Adjusting the N parameter involves balancing the precision and stability of the measurements [101], with higher Ns increasing the frequency with which measurements are done and the *magnifying* effect in the phase measurement, at the expense of lower bandwidth. With the current frequencies of operation and selected parameters of the DDMTD in the WRS devices ($N = 14$), the resolution of the phase tags is of 976 fs and these are generated every 3.814 kHz.

The primary objective of the main PLL is to generate a local copy of the clock signal recovered from the physical Ethernet signal of a WR master. This is accomplished by comparing the phase tags of the DDMTD reference clock channel with the phase tags of the recovered Ethernet clock. The PI control loop seeks to maintain a con-

stant phase difference between these two channels over time, taking into account a so-called *setpoint* that allows for programmable phase shifts between master and slave. The *setpoint* information is derived from the precise link delay measurement, which is obtained from the WR PTP packet exchange.

It is important to note that the system described is subject to various sources of noise at each stage, which may be amplified in subsequent stages if not carefully designed. The noise sources that contribute to the phase measurement mechanism's noise levels are listed as follows:

- The noise of the common clock of the DDMTD. This has been analyzed in depth by M. Rizzi in [101]. The analysis shows that the DDMTD in the typical configuration used in the WRS achieves 4 ps RMS single-shot precision.
- The low-voltage differential signaling (LVDS) input stage of the DDMTD clock into the FPGA, which introduces flicker noise.
- The DDMTD deglitching algorithm [98]. The D-type flip flops in which the DDMTD is implemented are sensitive to horizontal phase noise, or jitter. To minimize the impact of the glitches in the phase measurement, there is a deglitching algorithm that, although effective at filtering glitches and rejecting noise, has its limitations [68].
- The SoftPLL output precision is limited by the resolution of the digital-to-analog converter (DAC), which introduces quantization noise into the voltage signals that tune both the helper and the main clocks' frequency.
- The PI controller constants, if not correctly optimized, will degrade the phase noise profile of the device and introduce undesirable effects like poor stability or even divergence in the output phase and frequency that ultimately cause the clock to unlock from the reference.
- The contribution of the GTX transceiver both in the transmission side and in the clock recovery circuitry of the receiving end into the additive phase jitter.
- The use of the mixed-mode clock managerMMCM PLL (configurable clock generators/managers inside the Virtex-6 FPGA), which has significantly higher levels of noise than an external dedicated chip.
- The intrinsic noise of the main local oscillator, which, on top of it, is affected by the airflow of the fans inside the WRS enclosure.

The literature has extensively analyzed and characterized these sources of noise in WR devices [69], [70], [101]. In situations where the device loses communication with the master clock due to network malfunction or reconfiguration, it can only rely on past data to adjust the main clock and maintain holdover for as long as possible. To ensure accurate holdover operation, it is critical to minimize system noise. Otherwise, the cumulative noise of each stage will be incorporated into the holdover data and ultimately impact the holdover time signals performance.

For example, M. Lipinski's analysis of holdover requirements in his thesis [102] reveals that the standard WRS main oscillator can maintain sub-nanosecond accuracy for a maximum of 100 milliseconds after a network reconfiguration event is triggered. However, this time window is clearly insufficient for applications that require valid time signals during network reconfiguration. Additionally, the random processes suffered by the WRS local oscillator dominate over the effects that can be modeled and compensated to extend holdover time.

Based on these analyses, it is not advisable to use the standard WRS as a developing platform. Instead, the WRS-LJ[1] is the product designed by Seven Solutions that addresses the limitations mentioned above, and therefore serves as the hardware platform for our contributions.

4.3 Noise improvements in the WRS-LJ

The White Rabbit Switch Low-Jitter (WRS-LJ), as introduced in section 3.1.2, represents a new hardware iteration based on the v3.4 design of the WRS. It aims to confront the limitations of phase noise performance discussed in the literature [69], which restrict the applicability of WR in projects requiring more demanding short-term stability. The FPGA and software architecture modifications are kept to a minimum to facilitate maintainability and compatibility with future firmware releases of the WRS.

The two most significant factors in phase noise addressed by the WRS-LJ design are the elevated phase noise of the Grandmaster mode in the WRS (which is consequently propagated throughout the entire network) and the poor stability and intrinsic noise characteristics of the main oscillator control loop.

The high noise levels in Grandmaster mode can be attributed to the WRS requiring a 62.5 MHz signal to lock to an external reference, while the industry-standard solution is accepting 10 MHz signals. To provide the SoftPLL with the 62.5 MHz sig-

[1] `https://ohwr.org/project/wrs-lj-hw/wikis/home`

nal that it needs to lock onto, the 10 MHz signal is passed onto a FPGA input, and it serves as input to an internal PLL of the FPGA. However, the analysis by M. Rizzi in [69] shows that the internal PLL introduces considerable phase noise power located in the region between 10 kHz and 2 MHz. While this noise can be easily filtered out by a low-pass filter as typically integrated in analog PLLs, this is not the case in the digital implementation that uses the DDMTD and the SoftPLL. The unfiltered noise is ultimately aliased over the operating bandwidth of the SoftPLL. This issue is solved in the WRS-LJ by including an external AD9516-4 chip by Analog Devices[2] that can multiply the frequency of the 10 MHz output and successfully filter out the high-frequency components of the phase noise.

Regarding the selection of the main oscillator, the same analysis [69] demonstrated that the Mercury VM53S3 was responsible for the accumulation of phase noise between 1 Hz and 5 Hz, while it has also been observed that the air flow generated by the WRS fans substantially affected its performance. In contrast, the Connor-Winfield DOT050V-020M[3] in the WRS-LJ demonstrated superior performance. Experimental tests revealed that its casing was more robust against turbulent air currents, and encapsulating it in a polymer that provided thermal mass and protection against turbulent air had no noticeable effect. The new oscillator also had a narrower tuning range and better stability characteristics against aging, temperature, and voltage. Finally, the SoftPLL PI controller constants were optimized for the new oscillator's characteristics, achieving the lowest possible phase noise.

To showcase the improvements in terms of phase noise, the setup depicted in Figure 4.2 was configured. Two standard WRS were used (one grandmaster and one slave), while a deployment of up to nine WRS-LJ were analyzed (one Grandmaster and eight layers of cascaded slaves). The reference used for this setup is a Morion MV89 OCXO[4], and a Microsemi 3120A Phase Noise Test Probe was used to measure phase noise.

Integrating the phase noise over the entire spectrum supported by the measuring instrument (1 Hz to 100 kHz), as summarized in Table 4.1, reveals a tenfold improvement in RMS jitter. Moreover, examining the details of the results in Figure 4.3 shows that the improvement in the 1 Hz to 20 Hz Fourier frequencies ranges from 15 to 20 dB. These reduced noise levels have the unintended consequence of exposing previously hidden sources of noise: the periodic spikes visible between 20 Hz and 300 Hz are

[2] `https://www.analog.com/en/products/ad9516-4.html`
[3] `http://www.conwin.com/datasheets/tx/tx395.pdf`
[4] `https://www.morion-us.com/catalog_pdf/mv89.pdf`

introduced by the Ethernet management interface of the WRS and disappear when the device establishes a working connection.

Table 4.1: RMS Jitter results for WRS and WRS-LJ

(a) Standard WRS

Device	RMS Jitter (ps)
WRS GM	8.4
WRS SLAVE	10.0

(b) WRS Low Jitter

Device	RMS Jitter (ps)
WRS GM	0.9
WRS SLAVE 1	1.5
WRS SLAVE 2	1.8
…	…
WRS SLAVE 8	3.8

The conducted experiment effectively demonstrates the efficacy of the modifications made to the clocking electronics of the WRS. This substantial reduction in phase noise levels not only validates the proposed approach but also opens up new potential applications for WR devices. However, it serves as an intermediate step in our quest to enhance the reliability of the devices when the visibility of the master clock reference is lost.

The Connor-Winfield oscillator, which serves as the main clock of the WRS-LJ, offers considerably improved stability compared to its predecessor. However, as a voltage-controlled temperature-compensated crystal oscillator (VCTCXO), it still lags behind other oscillator families in performance and it is not stable enough to support holdover at the required frequency accuracy that a WR network requires.

4.4 WRITE expansion holdover board

From the analysis of the previous section, it was verified that the phase noise of the WRS-LJ is improved by the use of a better oscillator and an optimized control loop. Still, the stability of a VCTCXO is not enough to sustain holdover capabilities. To evaluate the performance of other commercially available oscillators with better performance, the hardware team at Seven Solutions developed the WRITE Holdover expansion board within the context of the Resilient Time Transfer work package of the WRITE project. This board is designed to accommodate an alternate oscillator as the main oscillator for the WRS-LJ, with functions identical to those of the surface-mounted main oscillator of the WRS-LJ. In other words, the oscillators from the expansion board must undergo phase measurement and tuning by the WR PLL, while

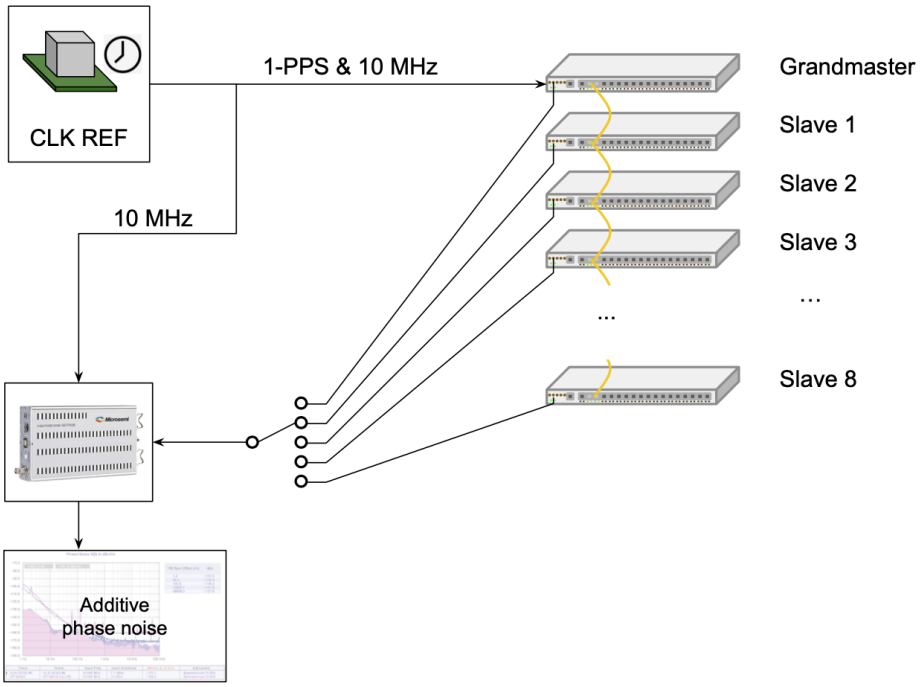

Figure 4.2: Experimental setup for determining the additive phase noise introduced by the WRS and the WRS-LJ

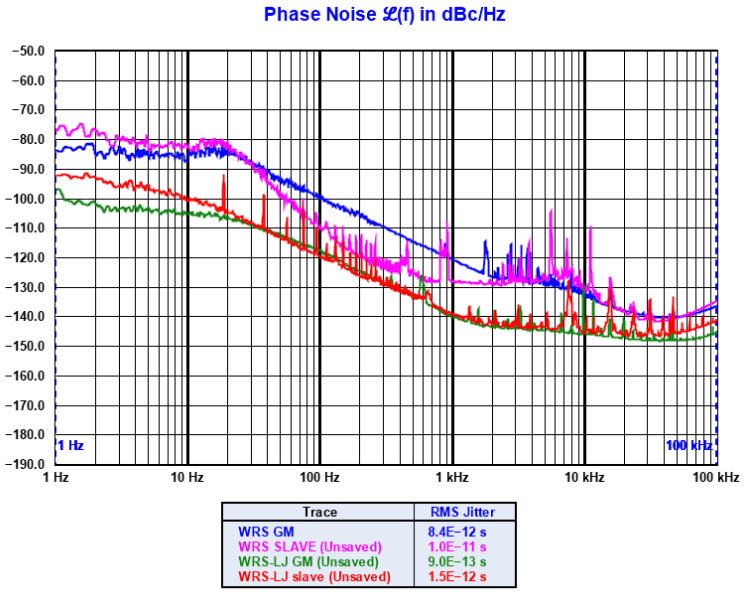

Figure 4.3: Comparison of phase noise levels for a Grandmaster and a Slave in both standard WRS and WRS-LJ.

the system's other modules will use the clock derived from this oscillator as the reference clock for all applications. The WRITE Holdover expansion board features physical footprints for up to eight different oscillators, a DAC that converts SoftPLL data into an appropriate voltage for clock tuning, various amplification stages that adapt the DAC output to the correct voltage range for each oscillator, and connectivity for Universal Asynchronous Receiver/Transmitter (UART)-capable oscillators. To accommodate the most power-demanding options, the expansion board is powered directly from the WRS-LJ power supply.

The clock signal originating from the expansion board oscillator is fed to the WRS-LJ input clock generator and distributor via a miniature coaxial Hirose u.FL cable. Remarkably, this implies that an advanced user can use the same cable to feed the WRS-LJ with any other clock source that meets the voltage level specifications and can be steered with the information provided by the expansion board's DAC output.

During the design phase of the WRITE holdover expansion board, we solicited recommendations from other partners of the WRITE project for a list of oscillators to consider supporting. We valued the inclusion of oscillators from various vendors and types, as well as the availability of resources among the partners. Since the goal is to implement an internal holdover solution by adapting the SoftPLL implementation, we excluded off-the-shelf holdover-capable oscillators that can be fed with a 1-PPS signal. Table 4.2 summarizes the features of the devices supported by the holdover expansion board. Stability and aging notation have been taken verbatim from the datasheets.

Of all the compatible devices, we had the Abracon AOCJY-10.000[5] (as depicted in Figure 4.4) and the Morion MV341[6] available for testing in our labs. Other partners of the WRITE project also tested the same holdover implementation that has been developed with their own oscillators. A summary of their results can be found in Deliverable D2 of the project [103].

4.5 The development of a holdover mechanism

The use of the improved clock circuitry of the WRS-LJ, along with the improved short-term stability of the experimental oscillators compatible with the WRITE expansion holdover board are two features that can be leveraged to improve the resiliency of a WR network by implementing an experimental holdover mechanism.

[5] `https://abracon.com/Precisiontiming/AOCJY.pdf`
[6] `http://www.morion-us.com/catalog_pdf/mv341.pdf`

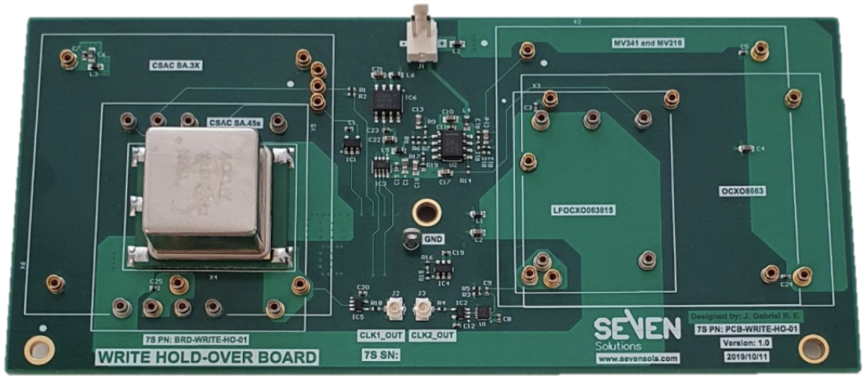

(a) Detail of the WRITE Holdover expansion board. The footprints for all the compatible options are visible.

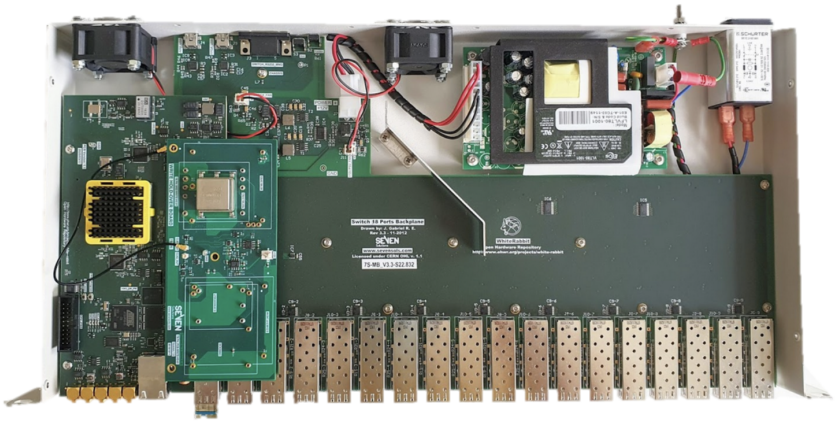

(b) View of the WRS-LJ with its top cover removed, showing the installed holdover expansion board and the Abracon OCXO.

Figure 4.4: The WRITE Holdover expansion board with a surface-soldered Abracon AOCJY-10.000 OCXO.

Table 4.2: Compatible oscillators with the WRITE Holdover expansion board and main stability characteristics

Model	Type	Temp. stability	Aging
Abracon AOCJY-10.000	OCXO	± 5 ppb (0 to 50 ºC)	± 100 ppb yearly
Microsemi CSAC SA.45s	Cesium chip-scale atomic clock	n/a	< 10 ppb yearly
Microsemi MAC SA.3x	Rubidium Miniature atomic clock	Up to ≤ 0.07 ppb	Up to ± 1 ppb yearly
Morion MV341	OCXO	Up to ± 1 ppb (0 to 55 ºC)	Up to ± 10 ppb yearly
Morion MV216	Double OCXO	Up to $<\pm 0.05$ ppb (0 to 55 ºC)	Up to $<\pm 5$ ppb yearly
IQD IQOV-164	OCXO	± 1 ppb (-30 to 75 ºC)	< 100 ppb yearly
Oscilloquartz 8661	OCXO	Up to < 4 ppb (0 to 60 ºC)	< 100 ppb yearly
Oscilloquartz OCXO 8663	OCXO	Up to < 1 ppb (0 to 60 ºC)	< 30 ppb yearly

This mechanism is intended to provide a usable time signal during the events in which a network event has interrupted the WR operation in a device, and awaits until the network is restore or during the network reconfiguration to use a different clock source.

In this section we will explore the factors that affect the performance of a holdover mechanism and the techniques used to increase holdover time, namely temperature, mechanical strain and other random processes. Following, we describe a commonly used model to account for the most relevant factors that can be adapted into the Soft-PLL implementation. Then, we will analyze how the mentioned factors and techniques can be controlled or implemented in our hardware platform. Afterwards we will provide an detailed overview of how the selected holdover strategies were implemented in the hardware platform, and the results obtained in our experimental environment.

4.5.1 Factors affecting holdover performance

Being precise in the description of factors that affect holdover performance is challenging because the precise effects of these factors depend greatly on the material of the resonator, manufacturing technology, and other aspects that fall outside of our area of expertise. While we can group these factors into broad categories, their specific impact cannot be determined with absolute certainty.

Short-term instabilities

Section 2.3 provided an overview of the different types of noise fluctuations that affect a time signal, but it did not explain the causes of said noise. The origins of these instabilities ultimately stem from intrinsic electronic noise, such as thermal noise, shot noise, and flicker noise present in the electronic components that are part of the oscillators and the systems in which they are integrated [39].

Temperature

The core element of a quartz oscillator is the piezoelectric resonator. Changes in temperature affect the resonant frequency of this element, consequently impacting the frequency stability of the oscillator. Various techniques exist to mitigate the impact of temperature changes, but a small component will always remain [104].

The effect of temperature can also be mitigated by placing temperature-sensitive components inside a thermally stable enclosure, also known as an *oven*. This gives rise to the term OCXO, which stands for *oven-controlled crystal oscillator*. OCXOs offer a >1000x improvement in stability compared to uncompensated oscillators [104]. For the purposes of this work, an OCXO is an excellent choice for achieving good frequency stability within a budget. However, it presents a challenge: the temperature inside the casing is no longer correlated with what we can measure inside the WRS enclosure, making it difficult to find correlations between the SoftPLL data and the device enclosure temperature. There are experimental workarounds for this issue, such as sensing the current drawn by the OCXO [105], but they have not been considered in our hardware platform.

Pressure and mechanical strain

Quartz resonators show dependence with atmospheric pressure due to changes in its elastic properties and accumulated mechanical stress. Therefore, the frequency of

an oscillator can be impacted by weather-related changes in atmospheric pressure, or by changes in altitude. Also, the stability of an oscillator can be largely impacted by accelerations that can be originated by shocks, acoustic noise and vibration such as the ones caused by a moving vehicle [104].

Aging

In addition to short-term noise (which, by definition, is modeled as a set of random processes) and medium-term dependencies (which can largely be attributed to environmental changes), aging is a mechanism that affects the oscillator on the scale of weeks, months, or years. According to [104], *many aging measurements have been reported, but few have included a detailed scientific or statistical study of the aging processes.* Aging is largely attributed to mass transfer due to adsorption and desorption of contaminants on the surface of the resonator, and the relief of mechanical stress within the mounting structure over time. Other effects with lesser impact are also reported. The aging rate of an oscillator is reported to be highest when it is new and evolves with an approximately logarithmic dependence [106].

4.5.2 Types of mechanisms to improve holdover

Of the factors considered in the previous section, only aging and temperature can be characterized with the information available for the WR PLL. Furthermore, the two oscillators available for testing on the holdover expansion board are of the OCXO type, meaning that the temperature inside the enclosure is kept constant with respect to the outside. In this situation, the SoftPLL cannot correlate the temperature inside the oven with the rest of the variables. Therefore, only the effect of aging can be measured and compensated for in this context.

Let us revisit the time error equation that was first described in section 2.3.3:

$$x(t) = x_0 + (y_0 - y_{0ref})t + \frac{(D - D_{ref})}{2}t^2 + \frac{[\phi(t) - \phi_{ref}(t)]}{2\pi f_{nom}} \tag{4.2}$$

If we assume that the reference clock is ideal (or, at least, that its degradation sources are negligible) all the "ref"-named variables in Equation 4.2 are equal to zero, and the equation remains:

$$x(t) = x_0 + y_0 t + \frac{D}{2}t^2 + \frac{\phi(t)}{2\pi f_{nom}} \tag{4.3}$$

where x_0 is the initial time error, y_0 is the fractional frequency offset with respect

to the reference (due to errors or finite frequency settability of the clock), D is the linear fractional frequency drift rate (used to model oscillator aging effects) and $\phi(t)$ agglutinates all the random phase components. Given that the value of D is susceptible to temporal variations due to factors such as aging and temperature, it is necessary to maintain a continuously updated estimation. This ensures that in the event of a holdover, the most recent estimate is available for accurate system performance.

When comparing the terms of the time error model with the situation of the WR PLL in the first instant that it loses visibility of its reference and holdover time starts, we find the following situation:

- The initial time error x_0 is ideally zero. At the moment when the slave loses visibility of the master clock, the WR link is considered to have been functional up until that point. Since calibration is not a perfect procedure and there is a finite level of precision, a few dozens of picoseconds of initial error can be expected under real-life circumstances. Therefore, there is no issue to be solved regarding the initial time error.

- The fractional frequency offset y_0 is proportional to the unbiased output variable of the SoftPLL that commands the frequency tuning of the main clock. In an ideal scenario without noise, this variable would take a constant value in the mid-range of the 16-bit register of the DAC. However, the real world is noisy, and the SoftPLL output variable also includes the additive effect of $\phi(t)$.

- To understand what the linear fractional frequency drift rate D of the time error model is concerning the WR PLL architecture, we need to consider a list that includes all the frequency words that the output has sent to the DAC over a certain amount of time. The parameter D represents the slope of a linear regression over that list of corrections.

Figure 4.5 offers illustrative examples of time error evolution in different scenarios. These scenarios relate to the fractional frequency drift D and the initial fractional frequency offset y_0, two of the fundamental parameters of the model in use. The term $y(t)$ represents the instant frequency offset when the random processes are not considered:

$$y(t) = \frac{\mathrm{d}x(t)}{\mathrm{d}t} = y_0 + Dt \tag{4.4}$$

In Figure 4.5a, the initial fractional frequency offset y_0 is presented. Here, the fractional frequency drift is non-existent ($D = 0$). However, the perfect determi-

nation of y_0 is impeded due to random processes and the limited resolution of the DAC. The scenario in Figure 4.5b offers a contrasting view where y_0 is perfectly determined, but a non-negligible fractional frequency drift rate D is introduced. This demonstrates the influence of the fractional frequency drift rate D on the time error evolution. Lastly, Figure 4.5c presents a composite case where both y_0 and D are constant and non-zero. This represents a more complex scenario where these two parameters simultaneously affect the evolution of time error.

Consequently, assuming an ideal scenario where temperature effects can be disregarded, there are still two variables that need to be determined in order to extend holdover time in the WRS-LJ: the one responsible for minimizing the fractional frequency offset y_0 and the variable that models the change of that fractional frequency offset as a linear regression, D. Let us now explore the commonly used methods to estimate these parameters in the industry before discussing our implementation.

Fixed holdover and DAC limitations

This simple approach consists of setting a fixed predetermined value in the actuator that controls the frequency of the clock. In the case of the SoftPLL, this is equivalent to setting a fixed DAC value. Of course, this approach does not minimize the fractional frequency offset (unless by chance), it does not consider the drift rate D, and it is not updated with information from the locked operation of the clock in any way. However, this approach is substantially better than leaving the control loop in charge, which would cause the frequency to drift further apart to the most extreme possible tuning positions.

Another slightly more sophisticated version of this approach relies on indefinitely using the last *good* known correction. Although this is better than using a constant value regardless of the previous state of the control system, this method will be greatly affected by the random component and the quantization error of the control system.

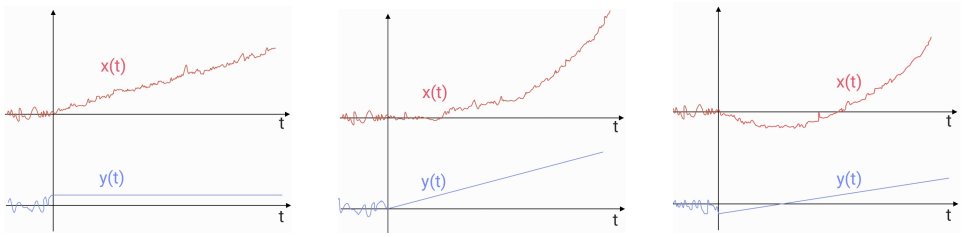

(a) $y = y_0, D = 0$ (b) $y_0 = 0, D > 0$ (c) $y_0 < 0, D > 0$

Figure 4.5: Illustration of time error evolution in different scenarios.

Additionally, determining the last good correction relies on failure detection, which is a challenging topic on its own.

Adaptive algorithms

The adaptive methods for frequency control involve adjusting the control signal of the OCXO based on historical data and predictive algorithms to maintain frequency accuracy and stability during holdover. For the operation of these methods, the time error of the last measurements must be available. Different approaches can be used for the prediction algorithm. Some solutions use extrapolation methods or prolong an interpolation function, while others are based on regression curves [107]–[109]. The predictive Kalman filtering algorithm is also used by many published solutions [110], but this solution is suboptimal in the presence of noise that is other than white and gaussian. In more recent years, some efforts are being invested in using Machine Learning techniques for extending holdover time in packet-based networks [111]. This last work claims to sustain accuracy over intervals of 1000 seconds "within reasonable limits" using SyncE-assisted PTP and OCXOs in two different temperature scenarios.

Indeed, the WR protocol offers a significant advantage in the context of holdover implementations due to its close integration of the physical layer syntonization and PTP packet exchange. This allows for a more precise model and abundant information about the phase measurement between the master and slave clocks, enabling the estimation of the desired parameters with a high degree of confidence. Consequently, WR-based solutions can potentially outperform other methods that rely on less integrated protocols or fewer data points.

Furthermore, the WR protocol's precision in time synchronization enables the consideration of the initial time error x_0 as negligible, which simplifies the holdover problem and allows for more efficient and effective solutions. Overall, the unique features and capabilities of the WR protocol make it well-suited for holdover implementations, providing a solid foundation for further advancements in this area.

4.5.3 WRS-LJ Experimental Holdover Implementation

In this section, we will discuss the modifications made to the SoftPLL in our proposed approach and the validation of these contributions through an experimental setup, building upon the concepts outlined earlier.

To briefly summarize the key findings from the previous section, the WR PLL control system of the WRS-LJ is at a stage where the initial time error, denoted as

x_0, at the onset of holdover mode is considered negligible. The subsequent challenges involve accurately determining the output frequency that minimizes the fractional frequency offset y_0 relative to the reference signal, and, upon successful estimation of this parameter, addressing the drift rate represented by D.

Software modifications

As outlined in section 4.2.2, the SoftPLL is the software running on the LM32 soft processor within the WRS-LJ FPGA. This software calculates the phase difference between its local clock and the master reference clock, utilizing the stream of phase tags generated by the DDMTDs at a rate of several thousand phase tags per second. Given the limited capabilities of a soft CPU and the scarce memory resources in the FPGA, it is impractical to store the vast amount of information consumed by the SoftPLL indefinitely. A potential future implementation of this mechanism will greatly benefit from access to a larger amount of memory and FPGA-assisted data processing to improve the accuracy of the model.

To keep the processing requirements minimal and efficient for the LM32, only the two primary variables within the main loop of SoftPLL operation are saved. These include the phase measurement error between the local clock and the master-recovered clock, as well as the output control values commanded by the SoftPLL for clock tuning.

Furthermore, to preserve storage space, the structure that maintains the history of past corrections and error measurements is implemented as a circular buffer with a capacity of 128 positions. This means that only 128 error measurements and 128 frequency corrections can be stored within the SoftPLL. The circular buffer is an optimal data structure for this application, as it retains the most recent data up to its capacity while overwriting older data. Its inherent *first in, first out* (FIFO) properties also ensure that data is consumed in the same order it was stored without requiring additional control mechanisms. Figure 4.6 provides a schematic representation of the circular buffers that store the SoftPLL holdover history.

Since phase tags and frequency corrections at the default rate of the WRS-LJ occur approximately twice every millisecond, the history buffer would be completely overwritten in less than 34 milliseconds with the current size of the circular buffers. To address this limitation, we have implemented a mechanism that only stores data in the buffer for one out of N executions (refer to Table 4.3). As the primary objective of the circular buffer is to store sufficient information for minimizing the fractional frequency offset y_0, it is essential to ensure that the window of stored data is extensive

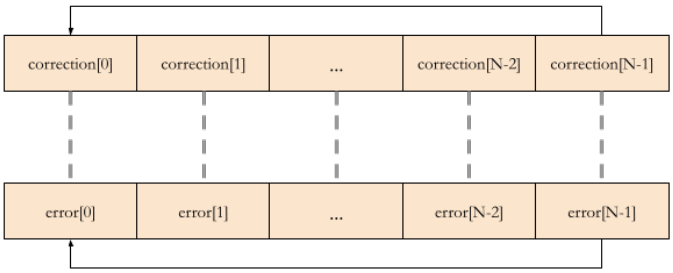

Figure 4.6: Schematic view of the circular buffers that are used to store the latest n measurements and outputs of SoftPLL history.

enough to eliminate any short-term transient or oscillatory behavior, thus ensuring that the SoftPLL operates under a "steady-state" condition.

Figure 4.7 shows a diagram of the SoftPLL functions involved in the main clock control loop, with a focus on the holdover operation. During normal WR operation, the SoftPLL operates normally, evaluating phase error and steering the DAC to compensate for it, as well as establishing a setpoint to phase align master and slave clocks. In addition to its normal operation, the holdover buffer is being populated with data from the measured error and the applied corrections. This operation is controlled by the *ho_update* function. Once holdover mode is triggered, the SoftPLL main control loop can no longer command the DAC, and the operation is only controlled by the *ho_update* function using just the data stored in the holdover buffer up to that moment. The software implementation includes placeholders for several holdover functions that may be needed by other partners of the WRITE project, but only *replaying* the circular buffer has been tested with the oscillators detailed in this section. Simpler holdover techniques such as fixed holdover were considered and evaluated, but quickly dismissed upon examining the first results.

Table 4.3: Examples of total buffer length with several undersampling parameters.

1:N Undersampling	Approx. Time Window of Circular Buffer
1	33 ms
16	536 ms
256	8.6 s
4096	137 s
65536	36.7 min

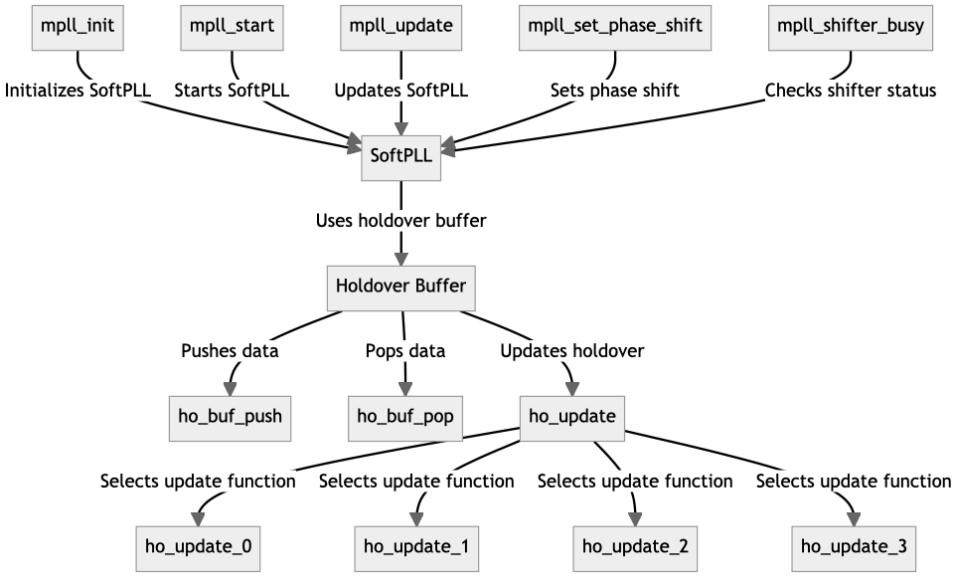

Figure 4.7: Diagram showing the most relevant functions of the main SoftPLL control loop, including the holdover implementation modifications.

In addition to the limitations imposed by storage capacity, the system also faces constraints in setting an accurate frequency value due to the limitations of the DAC detailed in section 4.5.2. The WRS-LJ is equipped with an AD5664 DAC[7] that, according to the following analysis, hinders the operation of the holdover implementation.

Given the critical role of the WRS/WRS-LJ DAC in managing the frequency of the clock, it is important to understand its limitations, which are amplified within the context of holdover operation. The current DAC in the SoftPLL architecture, despite its 16-bit resolution, guarantees accuracy only within ± 12 LSB (least significant bits). While the impact of this limitation might be less noticeable in a closed-loop system where errors are rapidly compensated, in a holdover operation, these errors can accumulate over time.

Let's consider the example of one of the compatible oscillators of the WR holdover expansion board, the Abracon AOCJY-10.000, which operates at a frequency of 10 MHz and has a pull range of 1.4 ppm peak-to-peak. Assuming ideal conditions, the oscillator can be adjusted to frequencies between 999993 and 10000007 Hz, offering a 14 Hz range. With a 16-bit DAC, this range is divided into 2^{16} steps, implying that

[7] `https://www.analog.com/media/en/technical-documentation/data-sheets/AD5624_5664.pdf`

each LSB ideally corresponds to a step of $214\,\mu$Hz in frequency.

However, the DAC's limitations introduce uncertainties in this scenario. Its limited precision means that no single data word will accurately represent the ideal frequency without introducing some amount of accumulating time error, even in a noiseless scenario. Moreover, the guaranteed accuracy of ± 12 LSB introduces an element of uncertainty, limiting the applicability of computations based on the history of previous data words.

When this ± 12 LSB uncertainty is factored in, the error rate per second becomes:

$$\text{error rate} = \frac{214\mu\text{Hz} \times 12}{10\,\text{MHz}} = 2.56 \times 10^{-10} \tag{4.5}$$

In a worst-case scenario, compliant with the DAC's specifications, this introduces an error of 256 ps/s, underscoring the impact of DAC limitations on holdover operations. This accumulating error greatly limits the usefulness of a fixed holdover implementation, which would reach $1\mu s$ in only 65 minutes considering just the DAC imperfections, with no other noise sources.

The proposed approach to overcome the limitations of DAC resolution involves repeatedly using the entire buffer of frequency corrections as the estimate for minimizing the fractional frequency offset. This method offers several advantages over more complex alternatives: it eliminates the need for calculations in either fixed or floating point and serves as a valid method for obtaining the average value of the entire buffer. Moreover, this approach ingeniously overcomes the limited resolution of the DAC. Although any calculations performed on the data within the buffer would be constrained by the 16-bit counter of the DAC, the repeated use of the stored 128 data positions as valid corrections ideally increases the granularity of the DAC input by 7 additional bits. Furthermore, since the stored corrections exhibit the same noise properties as those applied during locked operation, the non-linearities introduced by the DAC will be minimized by reusing data with consistent characteristics. Still, the virtually increased resolution can only address one of the two identified limitations introduced by the DAC, and the lack of accuracy in its four LSB imply that any data processing, even if as simple as calculating the average value of the whole buffer, can suffer from a bias that cannot be detected during holdover operation.

The software implementation also features a bus interface that allows users to enable or disable holdover mode, select the undersampling rate of the circular buffer, and enable or disable the learning process when the device is in a locked state. When holdover mode is activated, the phase tags generated by the DDMTD are no longer

processed, and the sole source of data becomes the stored history of measured errors and corrections.

In the following section, we present the results of a real world deployment of the proposed method for minimizing the fractional frequency offset y_0. Specifically, we tested two distinct OCXO models under various operating conditions. We then evaluate the performance of our method in relation to the intended use case and discuss its limitations as well as potential improvements.

4.5.4 Results

In this section, the experimental setup delineated in Section 3.4.1 is employed as the fundamental reference configuration. The system comprises a WRS-LJ acting as the Grandmaster, which obtains its external time output from a more stable external reference clock. The Grandmaster WRS-LJ supplies a time reference to a second WRS-LJ via a WR link. This second WRS-LJ, which is the device under test (DUT), is equipped with a holdover expansion board. To assess the time error, the time offset between the 1-PPS outputs of both the time reference and the DUT is measured using a Keysight 53230A Time and Interval Counter. An illustration of the experimental setup can be found in Figure 4.8.

Holdover is launched via a software command that triggers the holdover mode and simultaneously disables the WR network interface. Since one of the objectives of this study is the implementation of a holdover solution that can provide time signals with a degree of accuracy comparable to WR, once the time error of a run exceeds 100 ns the measurement is aborted and saved in that intermediate point.

In order to circumvent wraparound issues with the time counter instrument, an artificial time offset was introduced to the slave device. During the data processing phase, the measurements were then unbiased.

4.5.4.1 Abracon AOCJY-10.000

The Abracon AOCJY 10 MHz is a compact surface-mount device (SMD) OCXO with dimensions of 25 x 22 x 12 mm and an electrical frequency pull range of ± 0.7 ppm. Its steady-state power consumption is 1.4 Watts.

The test of this oscillator was performed in the laboratories of Seven Solutions during a period of COVID-19 prevention measures. The testing environment had windows and doors open when more than one person was present, and efforts were made to shield the devices from air flow. Therefore, the temperature of the equip-

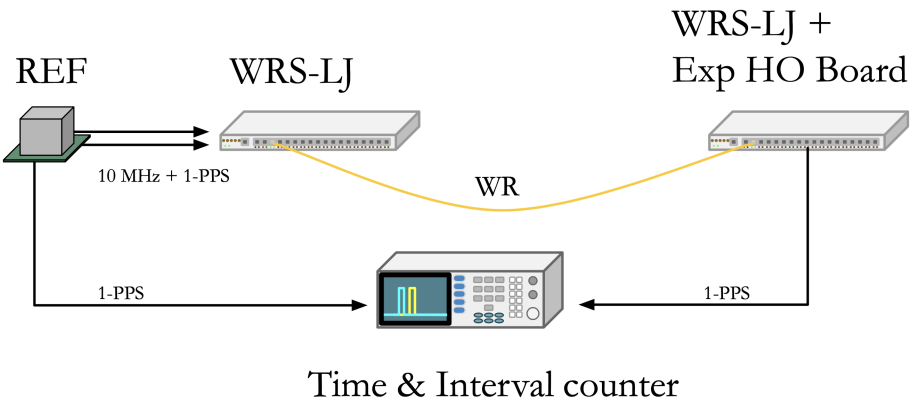

Figure 4.8: Experimental setup for the evaluation of holdover performance. The device under test is the slave WRS-LJ, that includes the expansion holdover board.

ment during the experiment is uncontrolled. Due to pandemic-related constraints, a Morion MV89 served as the time reference for this OCXO. To ensure stability, the reference was operated for over 24 hours in an enclosed space prior to use, while the remaining equipment was allowed to warm up for two hours. It was assumed that environmental conditions would remain relatively constant once the devices reached a thermally stationary state.

Various refresh rates were examined, incrementing the undersampling factor from 2^0 to 2^{14}, which corresponded to a 60-millisecond to 18-minute window. As a general guideline, time counter measurements were recorded until the time error exceeded 100 ns. Additionally, it is important to note that the rates were tested in ascending window length, ranging from 1:1 to 1:2^{14}.

First minute

Figure 4.9 illustrates the progression of all available measurements during the initial 60 seconds of holdover mode activation. The figure indicates that all analyzed rates can operate in holdover mode for at least 1 minute while maintaining the error below 1 WR system clock cycle (16 ns). Although there is no significant difference in the performance of most executions, the largest intervals exhibit inferior short-term performance. Two of the executions with the greatest error after one minute correspond to the two runs with the largest intervals between samples (rates 2^{14} and 2^{13}). Prior to these rates, the execution with rate 2^{12} does not exhibit a significantly larger error, but the slope of the frequency corrections is greater than those with shorter

intervals.

It is significant to highlight that in all trials, save one, the time error remained under 1 ns for a duration exceeding 55 seconds. Given that the typical time required for a WR synchronization from its initial state falls below 40 seconds, the experimental holdover solution employing this oscillator proves to be sufficiently robust. This solution is capable of maintaining a switchover operation whilst ensuring the synchronization impact on the time signals, and any potential slave devices connected to the device in holdover, stays within the sub-nanosecond range.

Disregarding the shortest and largest rates, which appear to exhibit worse performance, and focusing solely on the central rates, the results are more reproducible (Figure 4.10). While larger time windows seem to be noisier, it is also suspected that data from the shortest history windows were collected before the device's warmup was complete, as they all exhibit a similar short-term drift that vanishes in later runs.

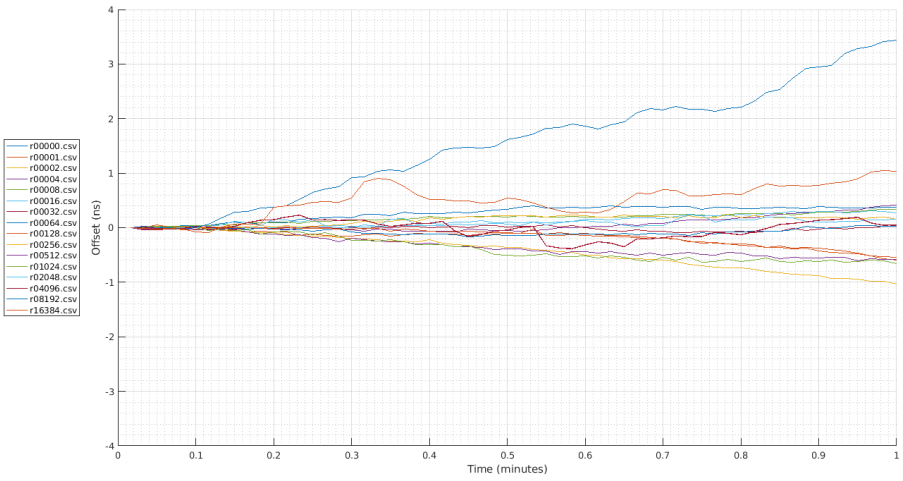

Figure 4.9: Time error in holdover mode using the Abracon OCXO during the first 60 seconds. Rates 1:N can be observed in the legend of the figure. rxxxxx represents rate 1:x, while r00000 represents an additional run of rate 1:1

First 10 minutes

Upon examining a more extensive time span of up to 10 minutes since the initiation of holdover operation (Figures 4.11 and 4.12), it becomes evident that the short-term noise is less prominent because other factors have become more significant. During this period, the shortest time windows again display a similar behavior. Considering that the shortest time windows were the first runs to be measured after

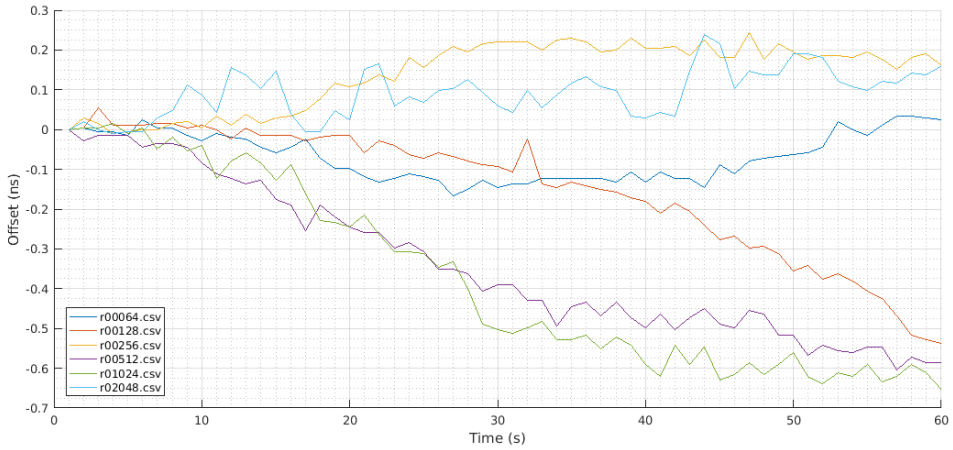

Figure 4.10: Time error in holdover mode using the Abracon OCXO during the first 60 seconds. Data shown represents the medium sized time windows (from 1:64 to 1:2048)

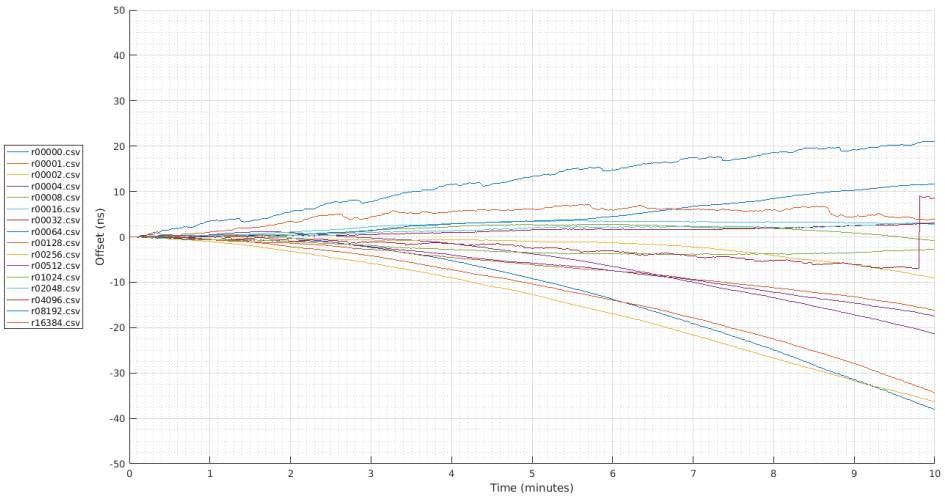

Figure 4.11: Time error in holdover mode using the Abracon OCXO during the first 10 minutes. The run with rate 1:4096 shows a sudden phase shift after 9.5 minutes that has been likely introduced during the postprocessing stage. For the sake of transparency, this run is included as observed.

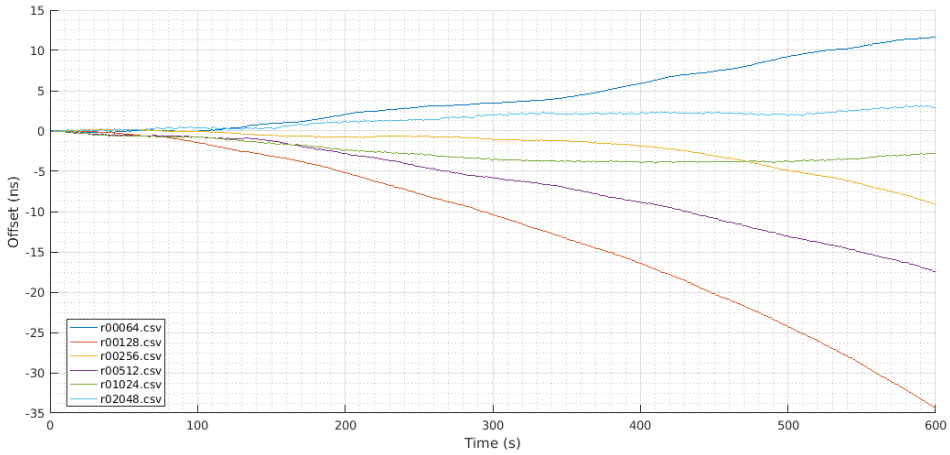

Figure 4.12: Time error in holdover mode using the Abracon OCXO during the first 10 minutes. The only data represented are the medium-sized history windows.

the setup and the warmup period, this suggests that the warmup period may not have been entirely completed when the measurements were taken (2-4 hours since power up). After the first 10 minutes, 56% of all holdover runs maintain a time error below 10 ns, which is reduced to 50% when considering only the medium-sized windows. The time error in different runs is dispersed with varying slopes, ranging from approximately 3 ns/minute to -3 ns/minute, with numerous intermediate slopes. According to our model and the DAC limitations, this can be explained as a correct estimation of initial frequency minimizing y_0, and an additional random component caused by the lack of accuracy of the DAC that results in uniformly distributed slopes in the TE in different runs. It is important to note that these measurements were conducted using a suboptimal reference clock in a room with uncontrolled temperature and air currents. These issues will be addressed in the subsequent measurements presented in Section 4.5.4.2.

The 10-minute window reveals a sudden phase shift in the experiment with rate 1:4096. We believe that the phase jump has been introduced due to a bug during the postprocessing stage where data is unbiased and PPS overlaps are fixed back together. However, we have decided to maintain the data from that run with the observed phase shift and transparently refer to the phase as an invitation to critically examine the limitations and uncertainties of the used methodology.

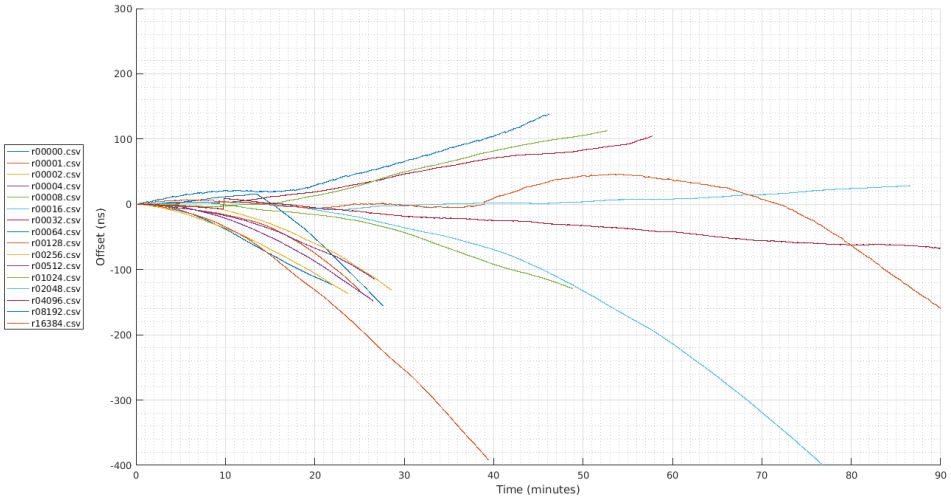

Figure 4.13: Time error in holdover mode using the Abracon OCXO up to 90 minutes.

First 90 minutes

The data from the subsequent 80 minutes of runs essentially corroborates the observations made during the analysis of the initial 10 minutes:

- The magnitude of the fractional frequency offset $y(t)$ appears to increase over time, with only three executions exhibiting a change in the sign of their $y(t)$ parameter throughout the run. Ideally, this change in sign would prevent the time error from perpetually increasing.
- Identifying a feasible drift rate D parameter seems improbable, as the drift appears to have opposite signs depending on the run. This contradicts our model unless the frequency changes are dominated by random processes that cannot be predicted or modeled. The behavior could also be partially attributed to an imperfect time reference that may introduce its own inaccuracies into these results.

Analysis of MTIE

The MTIE of all measurements taken with the Abracon OCXO (Figure 4.14) exhibits a similar pattern in how the MTIE increases with higher taus. Interpreting these MTIE results is challenging, as the behavior of most runs appears fairly consistent. However, it is noteworthy that, for taus greater than 20 seconds, the MTIE of all runs exhibits a strikingly similar slope. In contrast, greater variability is observed in

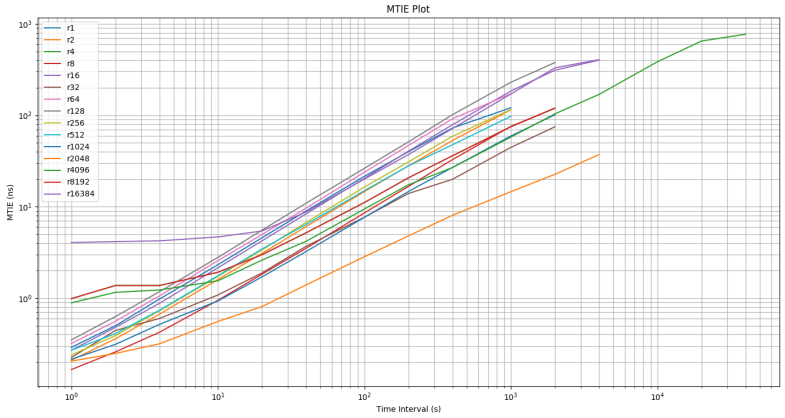

Figure 4.14: MTIE figures for all the Abracon AOCJY-10.000 holdover runs.

the smallest taus (1 to 10 seconds). This is likely an artifact caused by the time counter instrument when the 1-PPS signals overlap, and the instrument must compare the edge of one 1-PPS signal with the edge of the subsequent signal. The instrument then relies on its internal clock for an entire second, adversely impacting the results due to its inferior stability. Conversely, the data series without 1-PPS signal overlap does not display this behavior.

In addition to the MTIE, Table 4.4 demonstrates that the smallest time windows do not excel in either the short term or longer term. Windows derived from rates of 1:24 and larger take the longest time to reach a 100 ns TE. Coincidentally, an MTIE below 100 ns is the threshold below which a clock complies with the holdover specifications required for a Primary Reference Time Clock (PRTC), as defined in ITU-T G.8272 [112]. Moreover, it is evident that the largest analyzed windows exhibit the poorest short-term performance.

Discussion

The Abracon AOCJY is a OCXO that significantly enhances the holdover capabilities of WR devices and can be adapted easily into the existing hardware platforms. The mechanism for estimating the fractional frequency offset y_0 is sufficiently accurate to maintain a TE below 1 ns in most analyzed scenarios for over 60 seconds, which is adequate for performing network reconfigurations without disrupting the delivery of timing signals. Furthermore, the tests indicate that the drift rate of the fractional frequency is unlikely to be modeled as a constant, since different runs suggest opposite

Table 4.4: Summary of time elapsed to reach 1, 10, and 100 ns TE thresholds for all the analyzed runs. The color gradient of each cell is relative to the best and worst case in each column. The 100 ns threshold in the 1:2048 run was never exceeded during the length of the measurement.

Rate	seconds to 1 ns	seconds to 10 ns	seconds to 100 ns
1	85	444	1346
2	60	255	1346
4	83	420	1286
8	106	922	2517
16	118	1246	2727
32	252	859	3407
64	157	526	1418
128	86	294	1022
256	294	618	1512
512	140	251	1511
1024	119	1125	2829
2048	185	3815	Unknown
4096	201	1430	9207
8192	20	213	2336
16384	58	2409	5024

signs for the D parameter needed to compensate $y(t)$. However, readers should take into account that the experimental conditions for the Abracon tests were far from optimal, given the instability of the lab conditions and the use of a less-than-ideal reference, which may have influenced the results to some extent. These limitations are addressed in the following section, featuring another OCXO with similar technology but utilizing a more stable reference and a thermally stable environment.

With the current architecture of the SoftPLL and its limitations, this OCXO and its associated experimental holdover implementation can serve as a means to maintain accurate timing signals and ensure the operation of the WR device during short disruption events that trigger a network reconfiguration towards a different clock source. Its feasibility as a platform to extend holdover operations for longer times (i.e. 6, 12, 24 hours) is a topic yet to be proven.

While the experimental switchover solution demonstrates a meritorious capability in maintaining sub-microsecond accuracy for durations extending up to tens of minutes, it falls short in achieving this precision over an extended period of hours. Nevertheless, in the context of our specific application, this performance is a great improvement with respect to the initial state of the WRS. In comparison, superior solutions, such as those based on Rubidium (Rb) oscillators[8], may offer enhanced accuracy over longer durations. However, considering the balance between functionality, complexity, and the practicalities of our use case, the selected oscillator adequately fulfills our needs despite not being the optimal choice in an absolute sense.

4.5.4.2 Morion MV341

The MV341 is an OCXO manufactured by Morion, offering short-term stability of up to 1.5×10^{-13} at 1-second tau, a pulling range of ± 0.3 ppm, and consuming up to 3 W of power during steady-state operation.

Its evaluation was conducted using a setup identical in topology to the one depicted in Figure 4.8. However, for this OCXO, we could rely on the more stable performance of a Vremya-CH 1008[9] passive hydrogen maser, which provides a much more stable time reference compared to the one used for the Abracon OCXO. Additionally, instead of a lab with open windows and doors, the measurements taken in this section were carried out in a room with no air currents and a non-operational heating, ventilation and air conditioning (HVAC) system.

[8] `https://www.microchip.com/en-us/product/mac-sa53`

[9] `https://www.vremya-ch.com/index.php/en/products-en/passivehm-en/vch-1008-en/index.html`

In order to examine the differences caused by using a shorter or longer history data window, only two values were utilized for the history window on this occasion: 1:64, which overwrites the entire buffer in 2.1 seconds and shown in blue in the figures, and 1:1024, which takes 34 seconds to fill it and shown in red.

First minute

A total of 16 test runs were conducted, with minimum lengths of 12 minutes. In the first 60 seconds (Figure 4.15), there appears to be no discernible difference between the short and long windows used to estimate the initial fractional frequency offset. In comparison to the experiment with the previous OCXO, these results seem to be more reproducible: we observe that random noise fluctuations are weaker, and the data aligns better with the model. The figure demonstrates that the primary challenge, at least on this timescale, is determining the correct frequency value that minimizes y_0. The time error increases linearly with slopes arbitrarily determined to be between +1.4 ns/min and -2.0 ns/min in the worst cases. Assuming the error accumulation is ideally linear, this corresponds to a fractional frequency bound between 2.2×10^{-11} and -3.3×10^{-11} during the initial 60 seconds. Table 4.5 provides further details regarding the statistical values obtained for the fractional frequency offset.

Interestingly, despite the lower cost of the Abracon AOCJY, it demonstrated a slightly superior performance in terms of time error after a 60-second interval. Only the worst two runs of the Abracon run exceed the 2 nanosecond time error after 1 minute, while 7 out of the 16 Morion MV341 executions exceed this threshold.

First 10 minutes

The first ten minutes of execution (Figure 4.16) do not exhibit significant differences compared to the initial 60 seconds. The previously observed trend in the initial seconds of holdover run remains consistent throughout this period, despite minor frequency deviations due to random fluctuations. These fluctuations are evidenced by TE traces overtaking each other and can be confirmed by zooming in for a closer inspection (Figure 4.17). The fluctuations, particularly evident in the close-up figure, ultimately constrain holdover performance, as the oscillations are beyond the estimation capabilities of the current hardware. If the fluctuations were caused by oven artifacts, they could potentially be modeled. However, if these fluctuations are random, there is no viable solution.

It is important to note that there are no visible differences between longer and

Table 4.5: Summary of the statistic parameters of the Morion MV341 holdover executions. In this analysis, the observation time always starts at $t = 0$ for any tau.

	$\tau = 1$ minute	$\tau = 10$ minutes
Min y	-3.2×10^{-11}	-3.9×10^{-11}
Mean y	7.3×10^{-13}	-3.3×10^{-12}
Max y	2.4×10^{-11}	3.6×10^{-11}
Stdev y	1.6×10^{-11}	2.1×10^{-11}

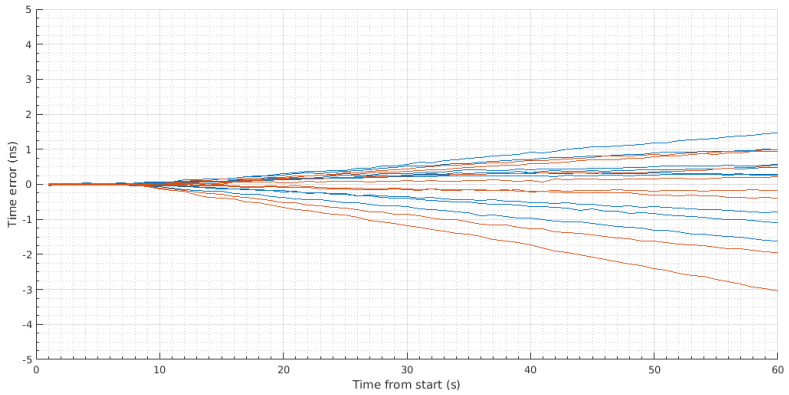

Figure 4.15: Evolution of the TE for the MV341 in holdover mode during the first 60 seconds since launch.

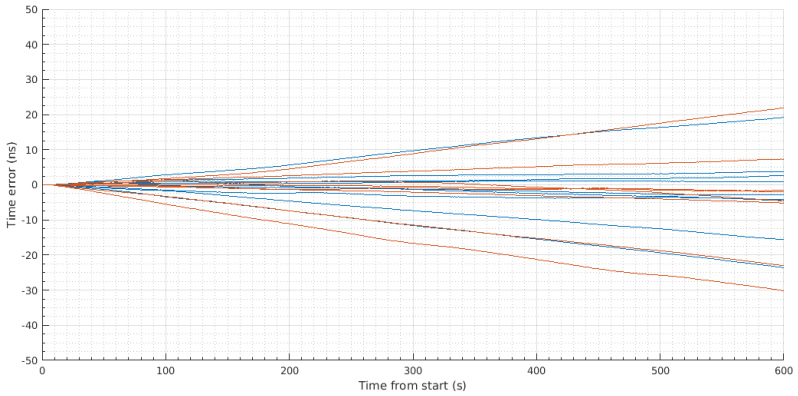

Figure 4.16: Evolution of the TE for the MV341 in holdover mode during the first 10 minutes since launch.

shorter windows using this OCXO.

Regarding the performance comparison with the Abracon AOCJY, when extending the evaluation window to 10 minutes, the performance dynamics between the Morion MV341 and the Abracon AOCJY 10 MHz oscillators shift. In this longer timeframe, the worst-case time error scenarios for both oscillators converge, exhibiting very similar values. This observation suggests that while initial performance differences may exist in shorter timeframes, these disparities tend to diminish over extended periods. This finding underscores the importance of considering the temporal dynamics of performance when selecting and implementing holdover mechanisms.

Long term

Regarding long-term experimentation, two holdover runs were executed for significantly extended durations (Figure 4.18) for the sake of completeness, although data from the previous sections suggest that the usefulness of the current holdover implementation is limited in these time frames. Regrettably, these runs were among the worst performers in terms of y_0 estimation, leading to a more substantial TE growth from the onset of execution. The sudden change of slope in one run's time error after 1150 seconds is not attributed to the holdover mechanism or reference but rather to the measuring instrument. The slope change results from an overlap between two 1-PPS signals, since the time counter needs to compare one second from one device with the next one from anocher, and one whole second elapses in between. Consequently, the poorer stability of the instrument's internal clock becomes apparent. Regardless of this issue, the two longest runs exceed the 100 ns threshold after no more than 40 minutes, and the 1 μs TE cap is exceeded after 4.7 and 5.4 hours.

As for the feasibility of estimating the drift of the fractional frequency offset, this also seems unfeasible with the Morion MV341. Although the longest run in 4.18 seems suitable for a polynomial fit, and the time error fitting with a grade 2 polynomial looks promising, the second-longest run displays an inflection point in the TE, ultimately dismissing the idea of estimating a constant (or at least, locally constant) fractional frequency drift rate.

Analysis of MTIE

As in the case of the previous OCXO, the Morion MV341 shows an almost identical rate of growth in the MTIE with higher taus among all the holdover runs. However, the MTIE values are significantly lower in the case of the Morion MV341. The best

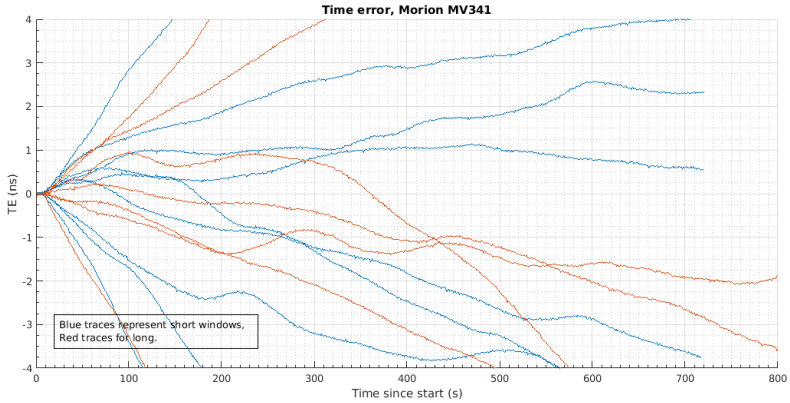

Figure 4.17: Closeup of the TE evolution during the first minutes of runs for the Morion MV341 OCXO. Although almost unnoticeable in the previous figure, there are random frequency fluctuations that change the instant rate of frequency offset.

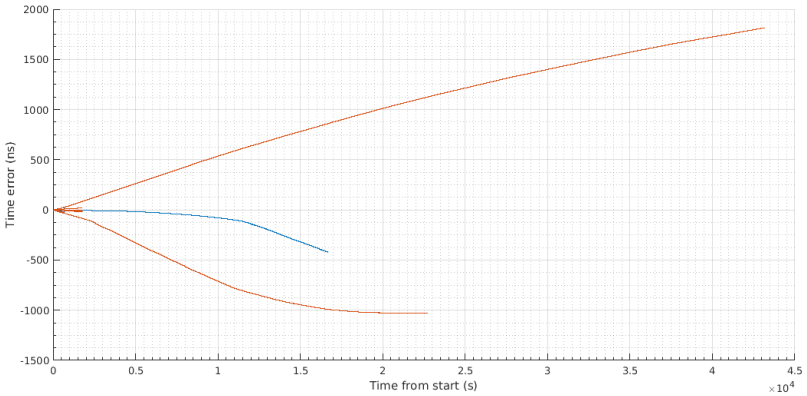

Figure 4.18: Evolution of three runs of the experiment that were run for longer periods (between 4.5 and 12 hours).

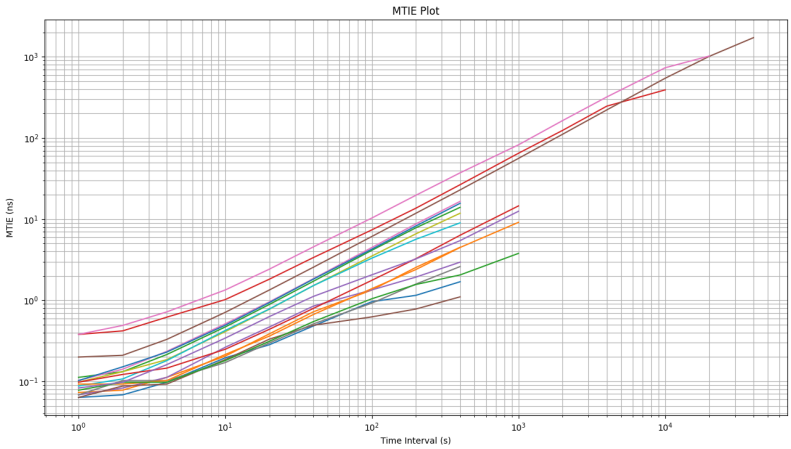

Figure 4.19: Calculated MTIE for all the holdover runs of the Morion MV341 OCXO.

run for the Abracon would be placed in the category of mediocre when compared with runs from the Morion OCXO. All the runs of this oscillator are better than 100 ns for a tau of 1000 seconds (which is the limit for holdover operation for a PRTC). In this case, the data has also suffered from the overlap in 1-PPS signal that deteriorates the data set when the measuring instrument has to rely on its internal oscillator for a whole second, instead of just a few nanoseconds.

Discussion

The results discussed in this section provide insights into the potential of the Morion MV341 as an OCXO to support holdover time in WRS and related devices. While the Morion MV341 shows on paper certain advantages in frequency stability compared to the Abracon AOCJY, the differences are not as pronounced as initially expected, particularly over extended periods. This suggests that both OCXOs can reliably deliver time signals during network reconfiguration or brief network disruptions. The parameter y_0 can be estimated with reasonable accuracy, though there is room for further refinement. This refinement may be related to both the software implementation and the resolution of the tuning electrical signal. The utility and feasibility of estimating the fractional frequency drift rate remain open questions, and further investigation is warranted to address these uncertainties.

4.6 Redundant timing system with low-latency data capabilities

As introduced in section 2.2, Smart grid integrates digital communication, information technology and automation into all the segments of an increasingly diverse electrical power grid, from generation to the consumer, with the intention of enhancing the efficiency, reliability and sustainability of the electrical power system.

The proliferation and seamless integration of the multifaceted technologies that constitute a smart grid necessitate the establishment of robust, reliable, and widespread communication network infrastructure. The smart grid ecosystem encompasses a multitude of interconnected devices, advanced metering infrastructure and diverse data analytic tools that impose the utilization of dependable and tried-and-tested network infrastructure like Ethernet. This study presents one of the practical use cases of such smart grid architecture, where we demonstrate that WR, with the necessary research and development, can operate despite the highly demanding requirements for redundancy and zero time recovery for both timing and data, thereby ensuring seamless system operations under all circumstances.

The data transmitted within a smart grid communications network encompasses not only control information for system management but also safety-critical signals that maintain system integrity. Lost or corrupted messages have potential repercussions extending beyond system inefficiency to personal injuries, substantial economic losses, and even widespread grid failure. This level of risk requires the implementation resilient mechanisms within the network infrastructure that can mitigate the chances of single points of failure and improve system fault tolerance.

In addition to the mentioned factors, the synchronization of devices that make up the smart grid is key for its effective operation. Maintaining a common time reference across all components facilitates a chear and coherent understanding of the sequence and timing of events within the network. Accurate synchronization ensures that these diverse elements can function as a cohesive whole, guaranteeing timely fault detection, rapid response to anomalies, and effective execution of preventative and corrective actions. In terms of time synchronization, the highest accuracy requirement comes from phasor measurement units (PMUs)[113]. For reasons also discussed in Chapter 2, the use of an alternative method to GNSS using terrestrial systems is recommended [114], and regarding fault tolerance and availability, IEC 61850 suggests the implementation of redundancy protocols to meet the demanding Smart Grid requirements, namely the parallel redundancy protocol (PRP) and the high-availability seamless redundancy

(HSR) protocol for seamless redundancy with zero time recovery. The criticality of smart grid data mandates this challenging recovery time in case of a failure, deeming the performance of other Ethernet redundant protocols insufficient [115].

This work aspires to achieve superior synchronization precision, essential for industrial applications, and to increase system fault tolerance by implementing the possibility to recover from a system failure. Furthermore, it becomes necessary for Control Systems and Smart grid to duplicate timing references and also data, as elaborated by H. Kirrmann et al. in their published contributions [116], [117]. During the time that this study was performend, just a few products offered this comprehensive set of features, with their time accuracy not exceeding beyond the lower limit of 30 nanoseconds based on standard PTP [118]–[120].

This section examines the details of the WR HSR protocol implementation in terms of data redundancy mechanisms, and of the time synchronization and switchover processes, followed by results for both timing and data redundancy configurations. The work described in this section presents a collaborative research effort conducted in conjunction with other esteemed researchers from the Department of Computer Engineering, Automatics and Robotics (ICAR, by its acronym in Spanish) from *Universidad de Granada*, where this author was previously affiliated. The findings of this research were published in the article titled *White Rabbit HSR: A Seamless Subnanosecond Redundant Timing System With Low-Latency Data Capabilities for the Smart Grid*, which appeared in the prestigious journal *IEEE Transactions on Industrial Informatics* [121].

4.6.1 Redundancy protocols and White Rabbit

Over the past few years, improving fault tolerance in industrial network has become increasingly important. Various methods and technologies have been adopted to not only recover from system outages and link or device failures, but also to improve the availability of network services. This fact remains relevant when applied to time-dependent services, as a consistent time notion across a domain is vital for precise capturing and triggering of events.

Two notable protocols facilitating these needs are PRP and HSR. PRP, used in tree topology networks, requires each node and link to be duplicated to ensure redundancy for time and data transfer. In contrast, HSR is applied in ring topologies and only needs an additional link to close the ring in terms of deployed infrastructure, without necessitating duplication of every node or link.

However, the implementation of HSR requires the development of complex hard-

ware and software that can handle the switching of duplicated frames. In terms of timing, to ensure interoperability, there is a dedicated profile for the use of PTP in HSR and PRP protocols (Annex C, [33]).

Given the cost-effectiveness of HSR in relation to PRP, the HSR implementation was chosen as the redundant protocol for timing and data redundancy. The guidelines and recommendations from the standards [33], [122], [123] have been followed during the development of the WR-HSR protocol presented in this study.

4.6.1.1 PTP clock types in WR-HSR

Before continuing with the details of the WR-HSR implementation, it is necessary to consider and give an overview of the different clock types in PTP. Each of these clock types can play a different role in the timing network:

- Ordinary clock (OC): this is a device with a single network connection that can be used as a master or as a slave for synchronization.
- Boundary clock (BC): a boundary clock has multiple network connections and can act as a master clock on some and a slave clock on others. BCs receive their timing information from the master clock, synchronize their local clock to it, and then distribute the timing information further down the network.
- Transparent clock (TC): a transparent clock has multiple network connections and measures the time taken for PTP events to pass through the device. It adjusts the correction field in the PTP event message to account for this delay (residence time), thereby improving accuracy in relation to a PTP-unaware device that might enqueue the PTP frames for undetermined, variable intervals. TCs do not synchronize themselves, and they do not influence the synchronization hierarchy because they do not modify the sync message content other than the correction field for the residence time.
- Hybrid clock (HY): a hybrid clock has multiple network connections and operates as a BC in some of them, and as a TC in others.

Likewise, PTP establishes two different methods for managing delay measurement in the network, namely end-to-end (E2E) and peer-to-peer (P2P). The choice between E2E and P2P often depends on the specific network architecture and requirements:

- End-to-end: in E2E, the master and slave clocks directly measure the round-trip delay. This method does not account for individual network path delays between intermediary devices.
- Peer-to-peer: in P2P, each pair of devices on the network measures the delay specific to their connecting link. This method provides a more precise measurement of path delay since it accounts for individual link delays, making it more suitable for complex networks with multiple hops between the master and slave clocks.

The WR implementation in the WRS was created as a PTP BC two-step clock device using the E2E mechanism to compute the delay of each link. Therefore, every node synchronizes to the clock information received from its immediate upstream WRS, and sends its own synchronized clock reference to the downstream devices. The implementation of PTP on top o the HSR protocol for the WRS, as suggested by [33], suggests the use of TC/HY clocks to ensure compatibility and scalability. Although the WR protocol provides specific mechanisms that may deem the P2P mechanism unnecessary, for the sake of compatibility, the WR implementation of the HSR protocol uses P2P TC and HY clocks [124].

4.6.1.2 Implementation of WR for redundant ring networks

The development of the WR-HSR protocol regarding timing is based on three main points: the syntonization of the entire ring using syntonization at the physical layer, the synchronization of HYs forming a ring topology using the exchange of WR messages, and the implementation of a switchover mechanism.

Since TCs do not synchronize and therefore there is no timing signal that can be compared to the master reference, the evaluation of the WR-HSR protocol was done using HYs instead of TCs. This is a fundamental difference with respect to standard PTP due to the utilization of WR, where all devices syntonize their clocks in order to have a common frequency reference across a network.

Section 4.2.1 introduced the details of the syntonization process in WR. The main difference between the standard tree topology and the ring implementation is that all ports of the ring, in terms of syntonization, are slaves except for the two grandmaster ports that are set as master. Furthermore, the SoftPLL must compare the local clock frequency to its right and left reference to guarantee the redundancy of the physical layer syntonization. For this reason, the syntonization process was modified to make the syntonization process possible in slave-slave links (although there is only one active

time reference, and the backup reference is continuously computed but not applied to the local clock). When the syntonziation process is accomplished in the entire ring, all the devices are *doubly syntonized* to the Grandmaster clock signal and they all share the same notion of time. Then, the synchronization process can compute the clock offset between the master and the slave.

Synchronization is carried out using the WR-PTP implementation. The delay between two adjacent nodes and the offset from the slave to the ring Grandmaster are computed from the PTP message exchange.

The initial phase of ring synchronization entails the assembly of a cascade of HY WRSs. Each of them will iteratively compute the offset to the master clock based on the message exchange until they reach sub-nanosecond accuracy. Subsequent to synchronization, their sole task is to monitor the modification in the clock phase to remain within the synchronization established limits. Once synchronized, the master's fiber link is connected to the last WRS's link of the ring that remains, effectively closing the link network. This enables all devices to receive PTP frames from their other port, which is considered the backup time reference. Here, the role of the backup link is to calculate the offset to the backup master by only considering the clock phase variations. During normal operation, only the offset as calculated from the primary time reference is applied to the local oscillator. The backup reference is computed but is only used to estimate an average offset that will be applied in the case of loss of the primary reference, ensuring continuous operation.

4.6.1.3 Switchover in the WR ring network

Introduced in section 2.6.3, switchover is, in the context of time synchronization networks, the process of transitioning between different sources of time or clock references. It involves switching from one primary time source to a secondary backup source in case of failure or degradation of the primary source.

IEC 61850 [123] mandates switchover to be carried out in less than 10 ms to minimize impact in the time signals during holdover operation. The switchover process would greatly benefit from the work of the previous sections of this chapter where the WRS hardware and software was improved with the focus on extending holdover time. However, due to external constraints, this section was developed prior to the holdover improvements. Despite the potential benefits, the stringent scheduling prevented the realization of these benefits. This highlights the challenges of managing academic research within industrial settings, where practical constraints such as version compatibility and predefined timelines dictate the research direction. Regardless,

the switchover process was implemented over the standard WRS that, as asserted in [102], counts with a holdover time of approximately 100 ms.

The mechanism for switchover in ring topologies was adapted from a version of the switchover mechanism suitable for parallel network topologies and developed at CERN [102]. This mechanism ensures that switchover happens in a matter of microseconds and the change from primary to backup source is performed before the holdover time is exceeded.

The ring topology imposes the propagation of this mechanism to all nodes in the ring as soon as possible, since only the nodes directly affected by the device or link failure will be aware of the event. All the remaining nodes would begin to drift until they receive a switchover alert and change their active reference. In collaboration with the main author of [121], three different mechanisms were considered for the dissemination of the switchover trigger message: sending control messages from the CPU, setting a timeout limit for the time a note is not receiving PTP frames, and sending Ethernet control symbols controlled by a FPGA module. The first and second methods were rejected after verifying that they would exceed the holdover time of the WRS oscillator. Hence, the switchover propagation mechanism is implemented as a FPGA module that inserts custom Ethernet control symbols to ensure its propagation in hundreds of nanoseconds. The author had the role of leading the data FPGA development side of the project and contributing to the experimental tests. Those interested in the details of the timing switchover implementation are encouraged to refer directly to our collaborative work in [121].

4.6.2 Data implementation of HSR

Effective data distribution is an essential feature of distributed control systems, as significat transmission delays or the failure to receive control messages could escalate into system failure. This section details the strategies to enhance data reliability for distributed applications. These approaches address and satisfy the demanding requirements for data dissemination in both industrial settigns and research facilities, made possible by the use of the HSR protocol for data distribution. All developments concerning data redundancy have been implemented exclusively within the WRS FPGA as a VHDL IP core.

4.6.2.1 HSR Overview

HSR operates based on the principle of duplicate send and duplicate accept, where data traffic is duplicated, and each duplicate is sent separately over two independent paths.

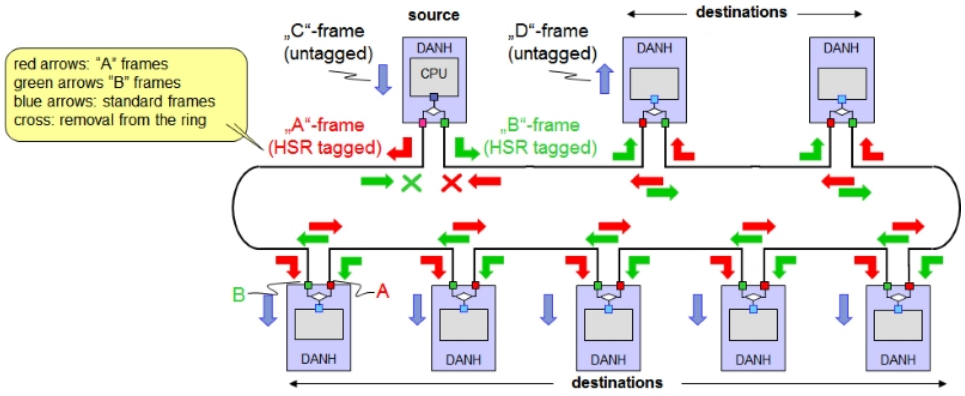

Figure 4.20: HSR network operation with multicast traffic. From Wikipedia, uploaded by H. Kirrmann, 2017, licensed GPL.

A key component of HSR is the Double Attached Node (DAN). Each DAN is connected to a ring network through two separate ports, and transmits every frame of data twice, once on each port. By doing so, the system ensures that even if one path encounters a fault, the data can still reach its destination via the alternate path. Conversely, a DAN must be designed to accept the first copy of a frame that it receives and reject any subsequent duplicates. A simplified version of this exchange is depicted in Figure 4.20.

RedBoxes are another essential element of the HSR protocol. They function as gateways between the HSR network and any conventional Ethernet networks. A RedBox also has two ports connected to the HSR ring but includes an additional port that connects to a regular Ethernet link. The HSR implementation of WR serves as a RedBox, where the WRS will convert the traffic from an additional connected device into HSR traffic when it enters the ring. More complex applications may require a QuadBox, an element that includes four ports for connection to two separate HSR rings.

The identification of duplicates is integral to the correct functioning of the HSR ring. This is achieved by embedding a sequence number in the HSR tag (see Figure 4.21) inserted in every frame of data sent by a DAN. This sequence number remains consistent across both the duplicated frames sent along different paths. As a DAN or RedBox receives incoming data, it checks the sequence number against a list of

previously received frames. If it identifies the sequence number as already having been received, the system knows that this is a duplicate frame and discards it. The DAN must also be in charge of removing the HSR tag and delivering the data frame to the upper layers.

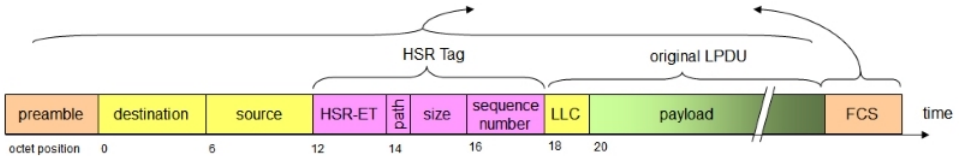

Figure 4.21: HSR frame format including the 6-byte HSR tag. From Wikipedia, uploaded by H. Kirrmann, 2017, licensed GPL.

The HSR tag is inserted before the frame original Ethertype and is composed of the following fields: a 16-bit Ethertype identifier (0x829F), a 4-bit path identifier, a 12-bit frame size field, and a 16-bit sequence number.

The HSR data components that hold all the HSR data activities have been implemented as a VHDL core placed transparently in the preexisting switch architecture. The core is called link redundancy entity (LRE) following the recommendations from [33]. Since it is only relevant for the operation of the ports involved in the ring, it only affects two of the eighteen ports in the WRS.

4.6.2.2 Link Redundancy Entity (LRE)

In this section, all the different units found inside the LRE (see Figure 4.22) are described in terms of function. The parts that are related to the operation of incoming traffic (*fast forwarding unit (FFU), dropper* and *untagger*) are first enumerated, while the ones that are related to outbound traffic (*tagger* and *arbiter*) are discussed in second place. The *fast switchover unit* (FSU) is discussed in a separate section since its operation, although implemented in the FPGA, is independent of the operation of the LRE.

Fast Forwarding Unit (FFU)

The FFU is the first module that a frame that comes from another node in the ring will find. Its purpose is the quick detection and decision of the destination of the frame. If the frame has to be forwarded to the next HSR node in the ring, the FFU will generate a copy of the frame that will be sent through the other LRE port, through the *arbiter*. This module introduces no additional delays in the incomming

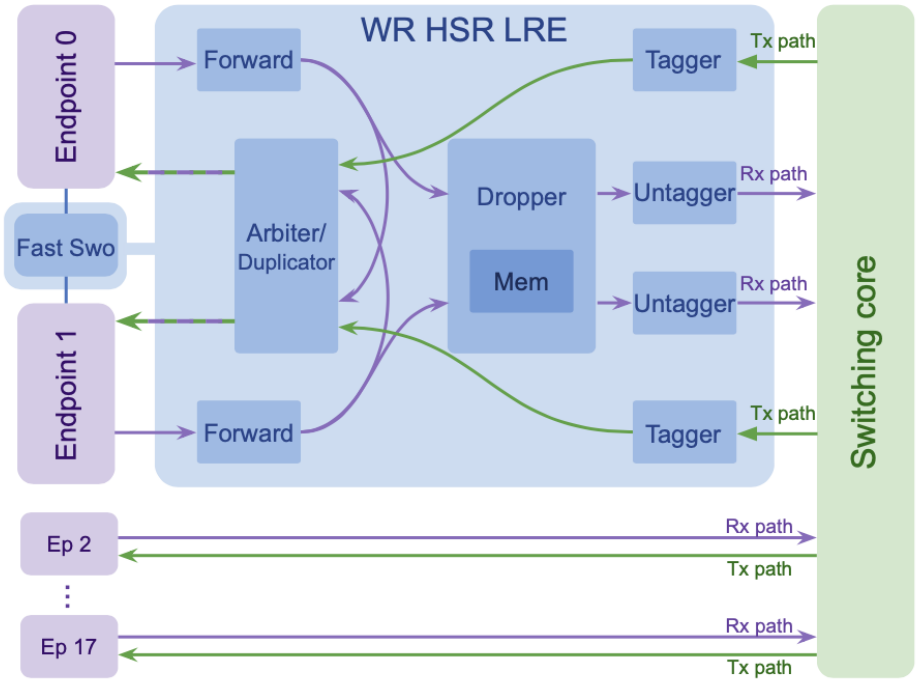

Figure 4.22: Diagram depicting the Link Redundancy Entity and its relation with its interfaces in the WRS FPGA project. From [121].

traffic path. The FFU acts in a cut-through manner, i.e., if the channel is available the data frame is bridged to the other LRE port as soon as enough information regarding the destination of the frame is extracted. The FFU considers destination and source of unicast, multicast, and broadcast Ethernet MAC addresses to make the bridging decision.

Dropper

The *dropper* is the module that contains the most complex logic within the LRE. Its purpose is letting the first copy of each HSR frame pair reach the upper layers while discarding the duplicate one. This entails that there is the need to implement a memory with shared access for both LRE ports in which source MAC addresses and sequence number information is written for every incoming frame. To accelerate the lookup time, it is implemented as a three-port content-addressable memory (CAM)[10].

The *dropper* memory can accomodate up to 32 devices in the ring. The maximum delay after requesting a lookup operation is of 136 ns. The memory stores the full 16-bit sequence number of the last received frame from a given device of the ring. The management of the frames is implemented in a window-like fashion: assuming that in a single ring there are no cases in which two frames coming from the same source can overtake each other, frames with a sequence number equal to the last accepted frame or slightly lower than the last accepted one will be discarded.

Untagger

The *untagger* removes HSR tags from accepted incoming frames to provide the upper layers with standard Ethernet frames transparently. The *dropper* commands the *untagger* whether a given frame has to be untagged. This operation is performed with zero additional delay.

Tagger

As suggested by its name, the operation of the *tagger* is reverse to the *untagger*. This module appends HSR tags to outbound frames. The tagger introduces the 6-byte tag described at the beginning of this section.

[10] `https://docs.xilinx.com/v/u/en-US/xapp1151_Param_CAM`

Arbiter

The *arbiter* is the unit that duplicates frames (outbound frames must come out of two physical network interfaces) and handles the demultiplexing from four incoming data paths to two outbound ones. Two of the incoming data paths come from each tagger, and there are two additional sources of traffic due to the FFU between the ports needed by the HSR protocol.

The arbiter introduces variable delays that depend on frame size and are impossible to bypass, since its normal operation involves setting traffic back to avoid collisions.

4.6.2.3 Fast Switchover Unit (FSU)

Although outside of the LRE, the Fast Switchover Unit (FSU) is in charge of a crucial task inside the WR-HSR operation. It is in charge of link failure detection, and consequently promptly commanding the device to switchover to the backup time reference, as well as send or forward the control message that is used to command and broadcast switchover to the rest of the ring.

When the FSU detects that one of the links attached to the ring is down, it sends a control symbol to the CPU inorder to activate the switchover mechanism. At the same time, that same control symbol is sent throught the functioning port to the rest of the nodes. This symbol is generated at the PCS layer to ensure its dissemination with the minimum possible delay.

When a node with two functioning ports detects a switchover control symbol, this is forwarded through the other port with no delay, as well as sent to the CPU, which will command the start of the switchover procedures.

This forwarding process is ultimately limited by the transmission speed in the fiber. In the case of a 6-WRS ring with 2-m fiber patches, the procedure takes no more than hundreds of nanoseconds.

4.6.3 Results

In this section, the results of the performed experiments regarding timing and data are presented. The experimental setup consists of six WRS arranged in a ring topology as depicted in Figure 4.23. The hardware version of the WRS is v3.4, while the firmware version used as the starting point of the development is v4.0.

The timing performance is evaluated by comparing the 1-PPS output between the Grandmaster of the ring and the slave under test. The custom firmware version

was calibrated before performing the tests using Axcen AXGE-1254-0531 and Axcen AXGE-3454-0531 SFPs with bidirectional transmission over a single fiber. These parameters remain valid for all the tests.

As in the case of other experiments in this chapter involving time offsets, these have been measured with a Keysight 53230a Universal Frequency Counter/Timer.

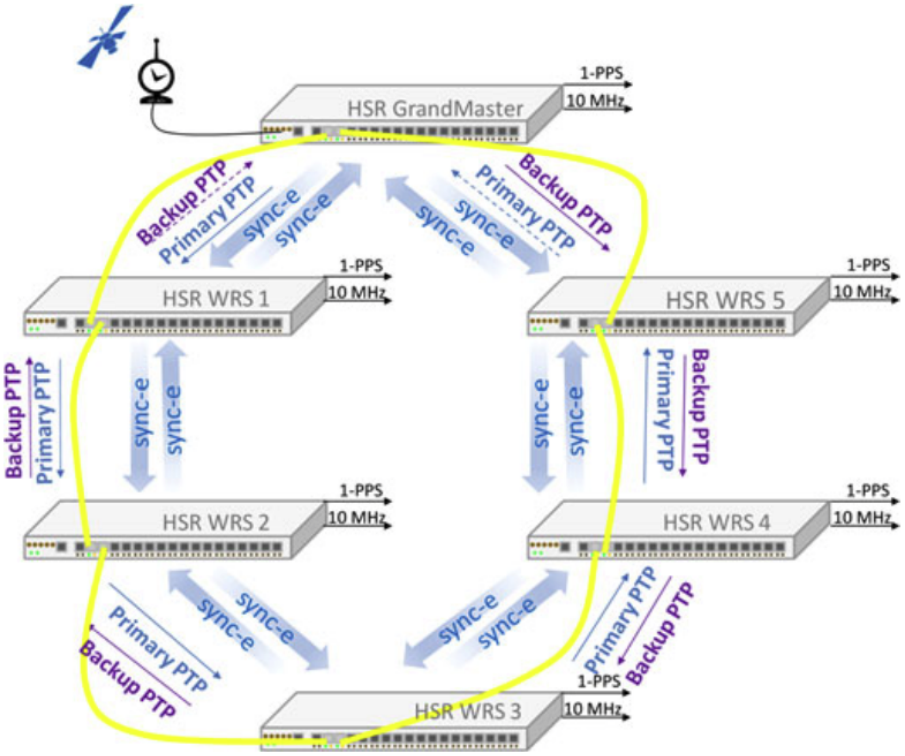

Figure 4.23: Experimental setup for the redundant topologies experiment. The ring is comprised of one grandmaster and six doubly-attached slaves. From [121].

4.6.3.1 Switchover timing results

Previous to the switchover experiments, all the nodes are synchronized one by one from the Grandmaster to the last WRS, and the correct synchronization is verified. After the last node is correctly synchronized to the Grandmaster from its "primary" time reference port, the ring is closed by connecting the last node, directly, back to the Grandmaster. The results show that the accuracy of the timing signals has been unaltered by the closure of the link. These can be found in [121] for more details.

Once the ring is stable and synchronized, a worst case scenario is simulated by

physically disrupting the link that connects the Grandmaster to the first slave. This represents the worst case because all the remaining WRSs must switchover from the primary to the backup reference. The results for the offset between master and slave after the switchover are very positive for all the switches in the ring. All the devices keep their synchronization accuracy below 1 ns, and the synchronization of their oscillators do not show any significant phase shift. The maximum phase shift, as shown in Table 4.6 and graphically represented in Figure 4.24, is of 170 ps. The obtained results demonstrate that the WR-HSR protocol implementation is capable of maintaining the paramount feature of WR, the sub-nanosecond accuracy, even after switching from a primary time difference to a backup one.

Table 4.6: Change of offset before and after switchover

Slave	Before Switchover	After Switchover	PPS Shift
Slave 1	150 ps	85 ps	65 ps
Slave 2	140 ps	310 ps	170 ps
Slave 3	122 ps	20 ps	102 ps
Slave 4	61 ps	-87 ps	149 ps
Slave 5	72 ps	178 ps	105 ps

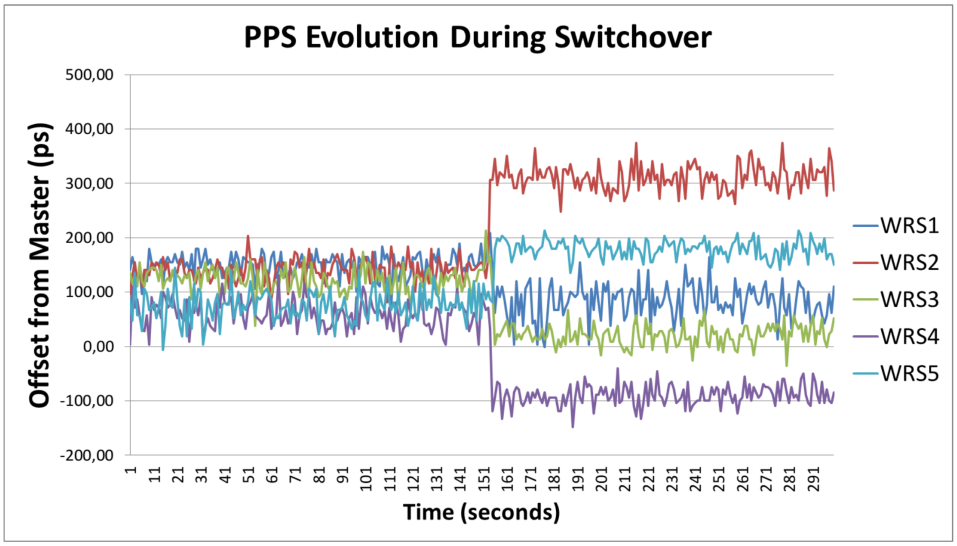

Figure 4.24: PPS offset between Grandmaster and each of the slaves during a worst case switchover scenario. From [121]

4.6.3.2 Data resiliency results

The setup for the data experiment starts from the ring setup used for the timing experiments, to which two PCs are connected, one each, to two adjacent WRSs. The PCs communicate with the WRSs with the help of copper SFP adapters that allow the exchange of Ethernet frames between the PC and the WRS. The latter act as RedBoxes for the PCs, and the topology is such that there is a three-hop link between the PCs via the shortest path of the ring, and a seven-hop link via the longest one. In order to evaluate the performance of the WRS HSR ring, three metrics were selected: latency, bandwidth, and frame loss.

Latency measurement is realized by comparing the transmission and reception timestamp of a frame that travels throughout the whole ring and comes back to the WRS that injected it.

Bandwidth has been determined by exchanging bursts of UDP frames between the two PCs using the network software tool *iperf*[11] and analysing the reports provided by the tool.

Frame loss is measured in two different ways: first, the bandwidth tests during normal operation of the closed ring must reveal whether frames have been lost during the length of the tests. In addition, the ring must demonstrate its zero-delay reconfiguration, i.e., no frame can be lost after failure of one link or device. This is proven by disconnecting one of the links during bursts of ping floods.

Latency results, summarized in Table 4.7, show that the simplest case of two WRSs introduce a latency of 1880 ns. Then, every additional hop adds 1430 ns to the total latency. These delays are composed in three parts: transmission delays due to the transmission medium (1 ns every 20 cm of fiber), forwarding delays between HSR ports (1008 ns), and fixed delays of the electronic circuitry (around 450 ns). The reduction of frame latency in the HSR implementation presents a reduction of 50% in contrast to the standard WRS firmware (from 2.2 [1] to 1.0 microseconds).

In relation to bandwidth, the results that can be found in Table 4.8 show that the ring of 6 WRSs under test can sustain traffic with a rate of up to 680 Mbps, a point after which frames start to be lost. Note that the physical duplication of all data frames leads to the insertion of twice as much frames in the ring as the amount generated by the PCs.

During the data tests, link failures were thoroughly simulated by breaking the ring in both paths while these were exchanging bursts of Internet Control Message Proto-

[11] `https://iperf.fr`

col (ICMP) ping messages. The results of these experiments show that no frames were lost due to any single-link failure, verifying the correct zero-delay reconfiguration requirement.

Table 4.7: Latency results for HSR WR Switches

	Two WRSs	Three WRSs	Three WRSs	Four WRSs	Five WRSs
Latency	1.88μs	3.31μs	4.74μs	6.16μs	7.59μs

Table 4.8: Bandwidth results for a HSR WRS ring with six devices

Frame Size	Bandwidth	Frames Sent	Frames Recv	%Lost
288 bytes	99.5 Mb/s	1 113 844	1 113 844	0.0%
288 bytes	199 Mb/s	2 229 670	2 229 670	0.0%
288 bytes	299 Mb/s	3 351 341	3 351 341	0.0%
288 bytes	400 Mb/s	8 126 971	8 126 971	0.0%
288 bytes	498 Mb/s	5 572 979	5 572 979	0.0%
288 bytes	603 Mb/s	6 798 765	6 798 765	0.0%
350 bytes	682 Mb/s	6 664 160	6 664 160	0.0%
372 bytes	700 Mb/s	5 306 220	5 170 491	2.56%

4.7 Conclusion

In this chapter, we have investigated the mechanisms responsible for the exceptional syntonization capabilities of the WR PTP implementation, the prevalent industry methods for ensuring reliability against network failures such as holdover, and the close relationship between these two aspects. We have identified the key components of the WR syntonization mechanism and their impact on holdover performance, and we have examined the main factors that enhance the stability of WR time signals. This analysis serves as a preparatory stage for the development of an experimental holdover solution and for understanding the challenges associated with its performance. From this analysis we have also obtained information about the current limitations of the system, and the components to improve in future iterations where the holdover mechanism is more deeply integrated within the system.

We have developed a novel holdover mechanism capable of extending the current WRS capabilities by a factor of over 200. While previous studies determined that the standard WRS could not maintain the sub-nanosecond accuracy for more than 100

ms, the worst obtained case in our data set maintains this accuracy for 20 seconds, and less than 10 ns error is routinely held for more than 10 minutes for selected configurations. Likewise, if we ignore the worst result after 60 seconds (which can be done considering that it is an extreme configuration case), all the rest of runs for the Abracon OCXO show that the sub-nanosecond holdover TE can be kept for at least 55 seconds, which is enough to perform a WR synchronization from scratch and successfully perform a switchover operation. This is achieved through the utilization of new hardware components and the adaptation of the software PLL that governs the node's reference frequency. Our implementation is centered around the selection and reuse of historical data once the connection to a master device is lost, ensuring the accurate delivery of timing signals during network reconfiguration and other disruptions even in the absence of a visible reference clock. The usefulness of the obtained implementation is demonstrated as a mechanism to ensure the delivery of accurate time signals during the network reconfiguration that follows a network disruption or a node failure.

Despite the considerable advancements made in holdover performance for WR devices, this holdover solution solution is not without limitations. The estimation of the initial frequency is limited by the current iteration of hardware of the devices, and restricts the usefulness of compensating other parameters such as temperature and drift rate. Future research could concentrate on addressing these limitations by enhancing the initial frequency estimation and increasing the estimators' complexity by incorporating new parameters, as well as improving the accuracy and precision of the DAC stage. Additionally, the development of mechanisms to detect failure events without corrupting data history and the integration of this holdover mechanism with existing switchover solutions warrant exploration.

Regarding the collaboration of the author in the implementation of redundancy mechanisms for data and timing, the integration of the HSR protocol within a WR network represents a significant advancement in the field of accurate synchronization for industrial applications and Smart grid. This development provides a robust solution for fault tolerant systems while ensuring sub-nanosecond synchronization.

The implementation of the HSR protocol for WR devices marks the first time that such a redundancy protocol has been applied to WR technology. This work has resulted in a system that maintains the typical WR accuracy even in the event of a switchover from a primary to a backup time reference. This level of accuracy significantly surpasses previous HSR PTP implementations, which typically achieve accuracy not below 30 ns.

Furthermore, the seamless switchover mechanism adapted for ring topologies ensures that the change from the primary reference to the backup one occurs within microseconds. This rapid switchover time is well below the holdover time of a standard WRS. The integration of the experimental holdover solution within this development is a work yet to be done, and the robustness of the HSR implementation would greatly benefit from the advantages of an extended holdover time as well as a more stable clock signal.

In conclusion, this chapter presents significant advancements in the field of reliability in WR networks, introducing new tools that can serve as the foundation for further innovative reliability mechanisms.

Chapter 5

Analysis and modeling of existing error and asymmetries in long distance network equipment

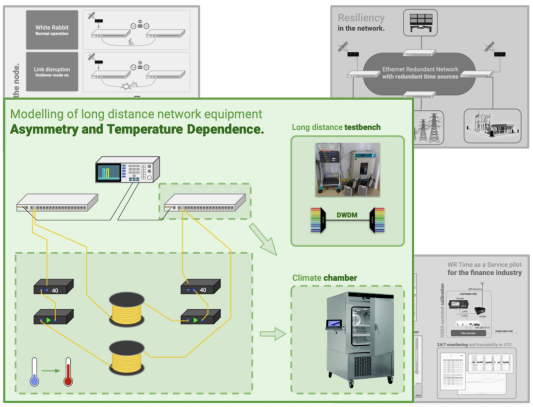

As stated in Chapter 2, National Metrology Institutes (NMIs) and research labs in related fields maintain frequency standards derived from atomic clocks and they have the need to compare the output of clocks in different geographical locations. To achieve that, before the current level of deployment of fiber optic networks, these institutions resorted to global navigation satellite systems (GNSS) techniques also overviewed in Chapter 2 like two-way satellite time and frequency transfer (TWSTFT) and phase comparison of GNSS signals. However, satellite comparison methods required at least 20 days to reach the most stringent requirements of uncertainty for microwave clocks [125].

The newer generation of optical clocks, which tick with higher frequencies than atomic frequency standards, slice time with finer granularity. The achievable uncertainty for comparison of optical clocks using satellite methods was deemed insufficient, and therefore the community resorted to optical fiber links [81], [126], [127]. The application of optical fibers for the purpose of frequency comparison and dissemination has been demonstrated in literature for microwave-frequency modulated lasers, and the direct transmission of optical carrier frequencies.

Citing [81], "a major challenge of the broad deployment of optical links is the difficulty (and the related high cost) to access dedicated fibre link between laboratories". The cost of deploying dedicated links for the dissemination of frequency standard and time services is unattainable, so the laboratories resorted to using fibre links embedded in existing public communication networks. However, the dissemination of metrology-grade frequency signals requires the propagation from one side of the link to another in exactly the same fiber so that phase noise can be measured and compensated. For the metrology purposes, this has been achieved through the means of bypassing the optical network components that do not comply with this requirement, and deploying alternate devices. Currently, the consortium in project CLONETS-DS [9] is in the effort of defining a network architecture for disseminating optical frequency standards on top of the European National Research and Educational Networks (NRENs).

Although White Rabbit (WR) does not require noise compensation mechanisms like the ones used for frequency dissemination, time transfer comes with its own specific challenges that does not affect the applications of frequency dissemination. Namely, accurate time transfer requires the precise measurement of transmission delay in the fiber in both directions. WR does not have an strict requirement of using the same fiber throughout all the optical path for its correct operation, but all the delay asymmetries need to be compensated via calibration.

In section 2.4.4, the review of existing long distance WR links mentioned that all the existing WR links over public networks rely on either a complete control on the network architecture, allowing the users to install, modify and calibrate their own equipment as needed, or on dividing a long link into several cascaded shorter links with intermediate WR regeneration stations.

Although operationally viable, these approaches are not scalable, and the generalization of deployments of long distance WR would not be feasible unless the WR service is treated as any other traffic by the optical network equipment of a public network to the maximum possible extent. As studied in 4.2.1, the performance of WR

is, in part, based on the physical Ethernet signal that encodes a clock. For the time transfer link to operate correctly, the clock that encodes the Ethernet signal must be unaltered. In this chapter, we aim to overcome this problem by analyzing the challenges that the deployment of WR links over public networks present, and aiming to address specific factors that impede the use of WR over public networks using the current link model and calibration mechanisms.

We begin this chapter by building on top of the analysis of factors limiting transmission in fiber optic medium from section 2.5, conducting an extensive review of the optical network components that are used to transport the optical signal throughout the network. In the context of fiber optic network, transport equipment refers to the physical hardware and devices used to transmit, receive, and manage data signals. This comprehensive review will provide the necessary background to understand the challenges of maintaining high-performance time synchronization over long distance fiber optic networks.

To put in practice the elements found in the analysis phase, we design and deploy a testbench, located in the facilities of Seven Solutions (nowadays, Safran Electronics and Defense Spain). This testbench will enable us to analyze the impact of various network configurations and components in terms of delay, asymmetry, and thermal dependence. The insights gained from the experimental setup will help us identify the factors affecting time accuracy performance and develop strategies to mitigate these issues. The results of this chapter will serve as one of the basis for the deployment of a pilot time service over a public network in the metropolitan area of Madrid that will be presented in the next chapter.

5.1 Optical network elements for long-distance links

The leap from 10 km WR links to running a WR link through an internet service provider (ISP) public network entails the use of multiple new elements that have to deal with the specifics of long-distance transmission described in section 2.5.1, as well as adaptations in other elements that were already in use. Depending on the need of amplification/regeneration (or absence thereof) and the use of unidirectional or bidirectional transmission, the different link cases can be classified into three main categories. Of them, two can be calibrated with proven and verified techniques (Figure 5.1) under the assumption that we have a certain degree of control over the network, while the remaining one poses the challenges that are being analyzed and modeled in this work.

- Bidirectional link without need for amplification. Since the fiber is shared in both transmission directions, the asymmetry in the link is only introduced by chromatic dispersion, which can be theoretically calculated or empirically determined in a laboratory. In the upcoming section, we discuss a scientific use case that makes use of this approach for WR links over distances of up to 173 km while maintaining accuracy under 2 ns.

- Unidirectional link without need for amplification. Under these conditions, the asymmetry of the link that causes the time transfer error is caused by the difference in the optical delay between the forward and reverse directions. Still, since there are no elements that force unidirectional transmission, the asymmetry can be determined by using the *fiber swapping technique*[128].

- Long-haul links that require amplification or regeneration. The difficulties introduced in this scenario are the topic of discussion of this chapter.

In the following section, we are going to review what elements come with the use of long-haul links, the purpose they serve, and the potential problems they may introduce for the operation of WR. Figure 5.2 contains pictures of most of the elements referred to in this study.

5.1.1 Equipment transceivers

SFPs (Small Form-factor Pluggable) are compact, standardized, hot-swappable devices that convert electrical signals into optical signals, and vice versa. Their optics are comprised of a transmitter that includes their laser source, and a photodiode that serves as a receiver.

WR typically is deployed using coarse wavelength division multiplexing (CWDM) SFPs rated for transmission of up to 10 km and wavelengths of 1490 nm and 1310 nm, these being two of the CWDM channels as specified in section 2.5 [77]. Although the performance of this kind of SFP is enough for keeping the accuracy below 1 ns in the distances for which they are rated, they are not without limitations. A publication from P. Boven [129] showed that the uncertainty of the wavelength of the SFP (that is, the deviation from their nominal value) between different units of the same model was too large for the WR use cases over long distances, stating that they would require calibration of each individual unit. Also, N. Kaur [128] found the spectrum of the CWDM SFPs to be very poor, spreading over tens of nm and non-compliant with their rated specifications.

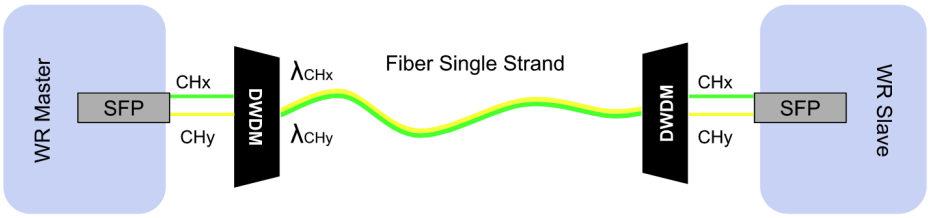

(a) Bidirectional link without need for amplification

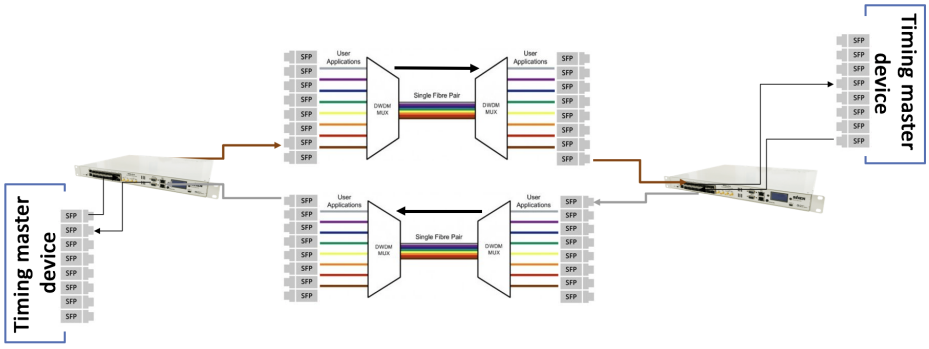

(b) Unidirectional link without need for amplification

Figure 5.1: WR links over long distance without need for amplification. Courtesy of Safran Electronics and Defense Spain.

As an alternative, the dense wavelength division multiplexing (DWDM) transceivers present a more pure spectrum (0.03 nm maximum for the units analyzed in [128]) and they implement a temperature control system in order to achieve the needed precision in their output wavelength. According to the results of [129], the dependence of SFP wavelength with temperature is two orders of magnitude smaller in DWDM when compared to their CWDM counterparts.

In addition to their better performance, the use of DWDM is a very frequent requirement when the deployments are made in a dark channel approach for the sake of compatibility.

Maximizing reach, accuracy and stability: WR and DWDM for SKA1 timing distribution

The Square Kilometre Array (SKA) [130] is an international project in charge of building the radiotelescope with the largest collecting area in the world. Its first phase consists of two radio telescopes, each of them made up of numerous receptors (artist depiction in Figure 5.3). Each of these receptors need to share a very precise time and

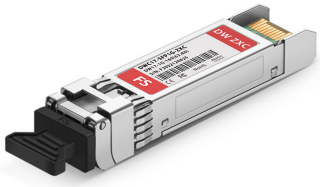

(a) A DWDM SFP transceiver with unidirectional (two fiber ports) interface.

(b) A MUX/DEMUX device providing access to eight DWDM channels as well as a monitoring CWDM channel.

(c) An OADM that allows insertion or extraction of one single channel from the common line fiber.

(d) EDFA in-line amplification module with its open case.

Figure 5.2: Manufacturer-provided pictures of some of the common elements in DWDM systems for long-haul operation. From Fiberstore.

frequency reference over distances of up to 173 km.

In a 2018 publication [129], the design of the time distribution for the SKA1 phase is described, resulting in, using the naming conventions previously introduced, a *bidirectional, dark fiber, long distance* time transfer system. The use of DWDM optics enable the designers to achieve the 2 ns accuracy objective *even in the longest links and assuming the worst temperature excursions* (173 km and 0 to 50 °C). The design proposes the use of DWDM optics in combination of 1-channel filters that allow to combine the stability and accuracy of the DWDM optics with the simplicity and ease of calibration of a bidirectional link. Figure 5.4 serves to illustrate the proposed topology of the links.

5.1.2 Optical transponder

In many cases, the DWDM "domain" does not start from the terminal equipment that generates and receives data. Instead, the terminal equipment has a low-cost SFP transceiver that generates a signal not suitable for long-distance transmission. In these cases, optical transponders are used to adapt the client-side interface to the DWDM domain.

A transponder is a device that has the ability to transform a signal into another wavelength or different standards [86]. For instance, they can convert a multi-mode signal into a wavelength suitable for single-mode transmission, convert bidirectional streams into unidirectional and vice versa, or simply change the wavelength of the stream. This is achieved by means of optical-electrical-optical (OEO) converters, that receive an optical signal, convert it into the electrical domain, and they use this electrical signal to modulate the laser of another transceiver [86].

Transponders can also integrate modules that provide the functionality that is characterized in this section by self-contained, discrete components such as amplification or dispersion compensation. The impact of transponders in a WR time link is not limited to the introduced, undetermined asymmetry, but also in terms of possible time indeterminism introduced by the OEO (variable delays depending on conditions of operation), and possible dependence with temperature.

5.1.3 Optical routing and switching

The DWDM signal that is generated from the WR devices must be aggregated with many other sources into one single fiber that contains many channels simultaneously, and along the way there might be situations that require disaggregating the channel, inserting it onto another fiber, or any other possibility that involves routing

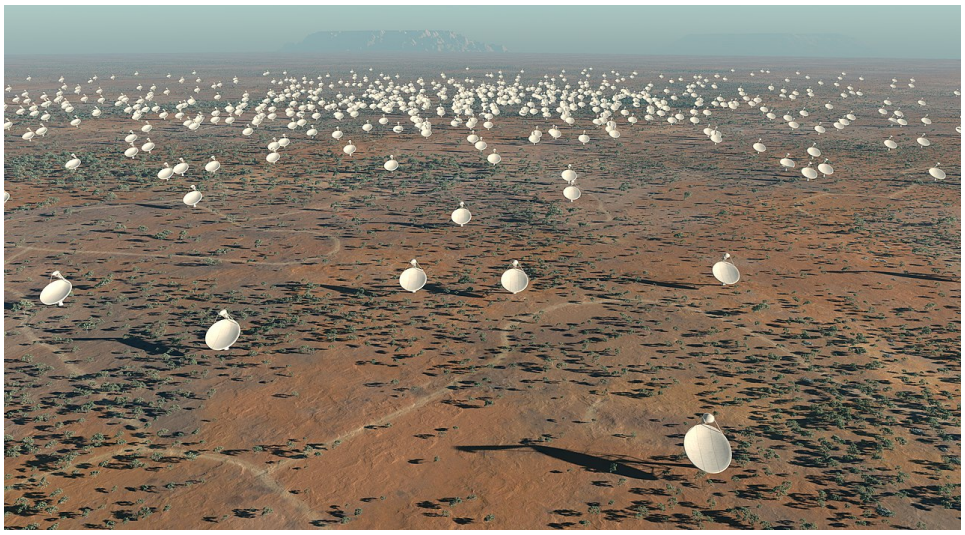

Figure 5.3: Artist's impression of the 5 km diameter central core of SKA antennas. From SKA Project Development Office and Swinburne Astronomy Productions, CC-BY 3.0.

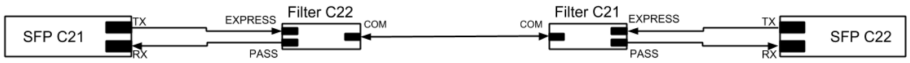

Figure 5.4: Schematic view of a DWDM bidirectional link. The output of the DWDM is designed for unidirectional fibers but the optical filters recombine and segregate the desired wavelengths

.

the optical signals. DWDM systems have the possibility of configuring the routing of the channels directly in the optical domain without conversion into an electrical signal, and depending on the specific manipulation that is to be achieved, there are multiple devices that cover all the necessary cases.

DWDM multiplexers/demultiplexers

A DWDM multiplexer/demultiplexer (MUX/DEMUX) is a passive device that can combine or separate mulitple optical signals from diverse DWDM channels into/from a single fiber. The principle of operation of these devices is based on difraction: thin-film filters, diffraction gratings or fiber Bragg gratings (FBG) are used to lead each wavelength component into the common fiber (or, in the case of demultiplexing, separating the multiple wavelength components into their own paths) [84]. Figure 5.5

depicts the principle behind the diffraction grating of a MUX/DEMUX.

In the case of a WR deployment, the two main problems that can be attributed to MUX/DEMUX devices are the uncontrolled asymmetry they introduce, and the potential thermal variations they may introduce that will be analized in this chapter.

Optical Add/Drop Multiplexers

Optical Add/Drop Multiplexers (OADMs) are devices used in optical networks to selectively add and drop specific wavelengths of light signales from a fiber optic link. They are useful when one channel needs to be diverted in a specific location, or when there is the need to give special treatment to one of the channels. This is the case in some of the experimental metrological deployments for long distance time and frequency dissemination, as in the case of the link between *Laboratoire de Physique des Lasers* (LPL) from *Université Paris 13* on the infrastructure of the French research network REFIMEVE shown in Figure 5.6 [81].

Reconfigurable Optical Switches

The capacity to redirect, insert and separate optical signals shown in the previous two components is not limited to passive, static devices. The development of wavelength selective switches (WSS) has made possible the manufacture of devices that can dynamically switch the desired number of signals that is routed from its common port to any of its N multi-wavelength ports, and vice versa. These 1xN WSS can be used as building blocks for more complex configurations, leading to devices like reconfigurable OADMs (ROADMs) or optical cross-connects (OXC) [79].

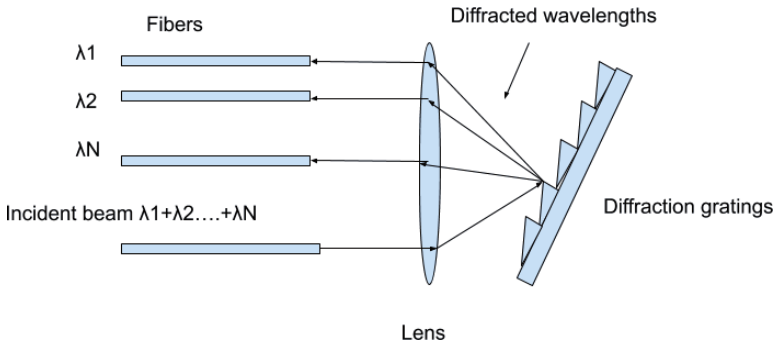

Figure 5.5: Principle of operation of diffraction gratings applied to signal routing. From Optically Multiplexed Systems: Wavelength Division Multiplexing. Creative Commons attribution

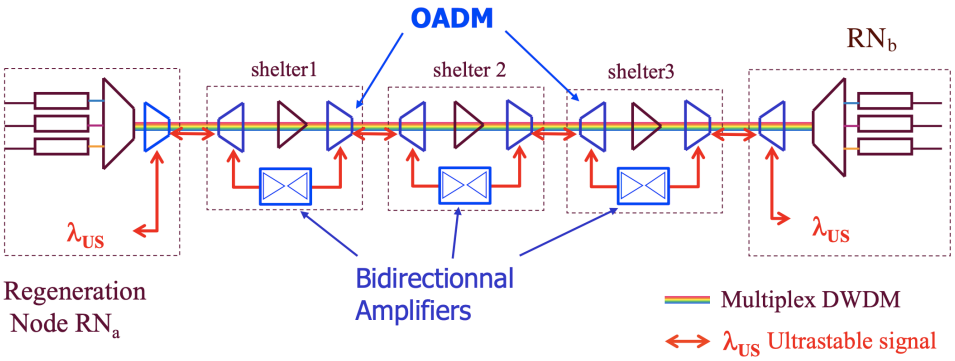

Figure 5.6: Bypassing unidirectional amplifiers using OADMs to redirect the frequency signal over custom bidirectional amplifiers in a metrological experiment. From [9].

The WSS achieves its function using a combination of optical components and technologies, including diffraction gratings, lenses, and microelectromecanical systems (MEMS). While the operation is similar to the one of static devices, the main difference introduced by the WSS is the inclusion of 2-D arrays of MEMS mirrors that can be individually operated to direct the individual wavelength channels to the desired output ports [131]. Other alternative methods of achieving the selective switching are Liquid Crystal, Liquid Crystal on Silicon, and phase arrays [132].

The use of reconfigurable devices in the paths of WR links does not only introduce unknown delays unless calibrated, but the nature of the switching mechanism entails that the apparent delay of the link can be dynamically changed, introducing a source of asymmetrical delay that can be subject to variations.

5.1.4 Amplification and regeneration

After a given distance, scattering and absorption mechanisms cause the optical power of the signal to be reduced. To maintain signal integrity, regeneration or amplification of the signal may be necessary.

Regeneration was done before the deployment of fiber amplifiers was popular, and it involved converting the signal from the optical domain to an electrical format, only to reconvert it to an optical format with a restored optical power. However, the development of optical amplifiers allowed to increase the level of the signal directly in the optical domain.

One of the most commonly used mechanisms to amplify the signals is by using erbium-doped fiber amplifiers (EDFA). These devices consist of a relatively short seg-

ment of optical fiber (approximately 10 to 100 m) that is doped with erbium, a rare earth element. This element exhibits optical properties that allow it to amplify optical signals in the region of the spectrum used by DWDM systems. A pump laser is used to provide the additional energy that is needed to amplify the signal. This laser excites the Erbium ions, and creates a population inversion. When the incoming optical signal reaches the EDFA, it interacts with the excited erbium ions, which drop back to a lower energy level and release their excess energy in the form of photons with the same wavelength, phase and direction as the optical signal [86].

It is known that some of the ions will decay back to its ground state without the presence of an stimulating photon. This phenomenon is known as spontaneous emission. This recombination is the origin of a broad spectral background of photons that get amplified along with the desired signal[86]. EDFAs are characterized by their noise figure (NF), which measures the degradation of the signal-to-noise ratio (SNR) between the input and the output of the amplifier. A typical NF value for a commercially available in-line EDFA is of 4.5 dB[1].

It is worth mentioning that EDFAs typically are designed so that they can only work in one direction because that is the predominant network architecture of its use case as in-line amplifier, although bidirectional EDFAs are commercially available[2] and are used for other applications such as the experimental frequency transfer link depicted in Figure 5.6.

As in the case of previous elements, the EDFAs are discrete network components that introduce an additional delay that will impact a WR link, may introduce an undetermined asymmetry, and the relation of the delay with temperature variations is a subject that will be studied in the upcoming sections. However, bidirectional EDFAs and other bidirectional amplification techniques are not without risks. The presence of bidirectional amplification by EDFAs was analyzed by [133] as a potential source for interactions between signals, and the coexistence between unidirectional data signals and other non-data services is a topic of active discussion within the community. Besides EDFA, the cited analysis includes other amplification methods based on non-linear effects such as Stimulated Raman Scattering and Stimulated Brillouin Scattering, whose description can be found in [84].

Regarding regeneration, although we are not considering it as a scalable technique for data networks, it is a valid approach to deploy WR links over networks that can be accessed, bypassing potential network equipment that may degrade the accuracy of the

[1] As per the specifications of the EDFA in the long distance testbench of this work.

[2] `https://www.keopsys.com/portfolio/bi-directional-fiber-amplifier/`

service. Several of the most relevant long distance deployments that were discussed in section 2.4.4 made use of WR devices as intermediate devices that consume the timing signal and regenerate it for the next hop. Figure 5.7 provides an schematic of the longest WR link in the world making use of the regeneration approach.

5.1.5 Dispersion compensation

Section 2.5.1.2 introduced the concept of dispersion, and Figure 2.17 showed the chromatic dispersion values for the most commonly deployed fibers. When accumulated dispersion is too high for ensuring the effective data transmission, network designers can resort to dispersion compensation modules (DCMs) to manage the levels of dispersion in the optical signal. DCMs are specially designed fiber components that exhibit dispersion characteristics opposite to that of the transmission fiber, thereby counteracting the accumulated dispersion effects [84].

DCMs employ fibers with a high negative dispersion coefficient, typically dispersion compensating fiber (DCF) or FBGs. DCF is a kind of fiber that is manufactured with different refractive index characteristics, resulting in a dispersion profile that is opposite to the transmission fiber, and also more pronounced (D around -100 ps/(km nm) opposite to around 17 ps/(km nm) of Corning SMF-28 Ultra at 1550 nm)[134]. This means that DCF can compensate approximately 6 unit lengths of transmission fiber dispersion per unit length of DCF. Besides, FBGs are periodic structures fabricated in the fiber core that reflect a specific wavelength range, creating a tailored dispersion profile.

DCF modules are not without drawbacks. Its attenuation parameter is higher, and the different fiber profile used in DCFs is more sensitive to other nonlinear effects and PMD. Additionally, their presence in a WR link can be very disruptive, since they amount to very large changes in the effective length of the optical link due to DCF modules containing several kilometers of fiber in a compact enclosure.

In the upcoming section, we show the experimental setup that was built in the premises of Seven Solutions to recreate a long-distance WR link, and the modeling work that was done to estimate the error that these techniques may cause into the timing signals.

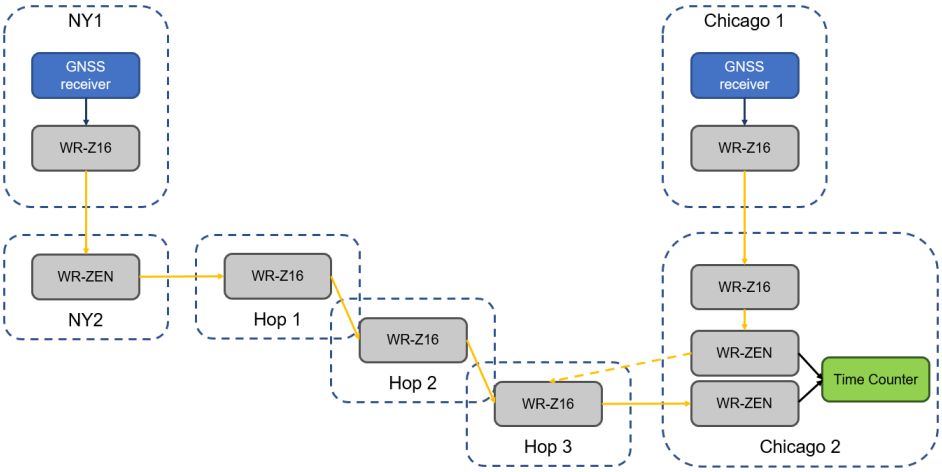

Figure 5.7: Schematic view of a WR link between two points in the metropolitan area of New York City and Chicago, using WR devices as intermediate regeneration stations. Courtesy of Safran Electronics and Defense Spain.

5.2 Temperature dependence and time error in unidirectional long distance WR links

With the purpose of verifying the applicability of the devices and components overviewed in the previous section, a long distance testbench was deployed under controlled conditions and using fiber spools to recreate a round trip link fiber spanning up to 100 km over unidirectional fibers. Figure 5.8 shows a depiction of the testbench. Two WRS were used as master and slave endpoints of the link, and the specifics of the long-distance testbench are detailed in the following list:

- Three pairs of DWDM SFPs rated for 80 km transmission over the channels 54, 57 and 59.
- Two pairs of MUX/DEMUX devices for DWDM channels from 53 to 60. Each pair is needed for aggregation and separation for each of the two fiber that comprises a unidirectional link.
- One add/drop multiplexer for a single channel.
- Two DCMs using a loop of DCF rated for compensation of 40 km of G.652 fiber.

Figure 5.8: Long distance testbench. Three WRS are shown on top of the rack, and MUX/DEMUX devices, OADM, EDFAs and DCMs follow from top to bottom. Spools of fiber are contained in the grey metal cases, and the device in the right side corner is a climate chamber that allows us to establish a precise temperature of operation.

- Two inline EDFAs rated for 17 dB gain over the DWDM wavelengths of operation.
- Several fiber spools (not all of them were shown in the picture) with lengths of 5, 20 and 50 km.
- A climate chamber capable of setting the desired temperature from 10 to 60 °C.
- Additional temperature probes for more precise monitoring of the temperature.
- Attenuators of varied values to achieve an adequate optical power level for the receiver side.

The first test that is carried out is an indirect analysis of the accuracy of the real wavelength of selected CWDM transceivers. We do not have access to the necessary equipment to directly observe the spectrum of the laser, but we can use the effect of dispersion in the WR link to empirically determine the time offset introduced by using different units of the same model of transceiver. In contrast, we show a mechanism to estimate the link asymmetry based on the precise wavelength of DWDM transceivers and the capabilities of WR to measure the link length with a precision of picoseconds.

Next, for the analysis of the delay and thermal variations of the network components, a reference setup has been deployed that contains most of the components simultaneously. Since the dimensions of the climate chamber are a limiting factor, we

only introduce two 50-km fiber spools inside the chamber, while the rest of the devices are kept under uncontrolled room temperature. After the analysis of the reference setup, each of the network components will be placed inside the climate chamber to analyze its contribution to the link total delay and its variability with temperature.

The analysis of the reference setup and all the network components involved serves not as a precise model of their behaviour that can be used to suppress their impact. Instead, the chosen apprach onsists on characterizing the thermal response of the WR link altogether, and subsequently propose a series of simpler setups that provide the information to isolate the contribution of each device isolated from the rest of the elements. Still, the results of this section can be consider an upper limit to the contribution to the error budget of a link.

5.2.1 Precision of CWDM/DWDM wavelength reproducibility

Several WR users [128], [129] have verified the lack of spectral purity and accuracy of the CWDM SFPs by observing its power spectral density (PSD) using an optical spectrum analyzer. Despite the unavailability of the requisite instrumentation in our laboratory, we can verify the poor accuracy of the transceivers and conducting a reproducibility test. The poor accuracy of the lasers in the CWDM SFPs can be indirectly quantified by measuring the difference in round-trip time and/or time error obtained by using different units from the same model for the transmission. The changes in round trip time, assuming constant temperature along the duration of the test, is attributed to variable errors between the nominal and the real transmission wavelength. Due to chromatic dispersion, signals transmitted with slightly different wavelengths experience a different time of flight.

In this experiment, a simple reproducibility test was set by evaluating the error introduced in the PPS output in a link consisting of two CWDM SFPs of wavelengths 1490 and 1550 nm over a 40 km link. The delay of the WR devices are calibrated, but the asymmetry parameter α is set to zero to suppress its influence in the determination of dispersion and reproducibility. Also, attenuators were inserted to accomodate the signal power to the range of the receiver. The results are presented in Table 5.1.

Although the number of samples is limited with just four different pairs of SFP, the worst case of time error introduced by the SFPs over the 40 km link is of 1651 ps of difference between pair 4 and pair 3, and the standard deviation of these four pairs is of 770 ps. With the observed results, the use of CWDM is not suitable for long-distance WR links due to the lack of reproducibility in the results, which could only be solved by calibrating the real wavelength of each unit. This process would be

Table 5.1: Reproducibility test of CWDM SFPs

	Run 1	Run 2	Run 3	Average
Pair 1	17467 ps	17663 ps	17647 ps	17652 ps
Pair 2	18252 ps	18294 ps	18305 ps	18284 ps
Pair 3	16663 ps	16687 ps	16685 ps	16678 ps
Pair 4	18312 ps	18335 ps	18340 ps	18329 ps

time consuming and would require the use of specialized lab instrumentation.

The obtained results are in line with the link error budget calculated for the long distance WR deployments in the SKA radiotelescope project [129], which estimated 1240 ps of error for 80 km links using these same transceivers. In the same conditions, the error budget is reduced in a 7.75-fold reduction to 160 ps of uncertainty over 80 km if DWDM transceivers are used.

5.2.2 Reference setup: 50 km unidirectional link

The reference setup for these tests contains all the elements that are part of the long distance testbench altogether (Figure 5.9). The master-to-slave path departs from the DWDM transceiver, goes through a 2-meter fiber patch to the first MUX. The signal leaves the MUX through the common transmission fiber that spans 50 km. Afterwards, it is amplified by an EDFA and is compensated for distortion by the DCM afterwards. The last step connects the output of the DCM to the slave WRS. The reverse path, from slave to master, is nominally identical although an unknown asymmetry in the link delay is expected. To accomodate the signal to the accepted range of the SFP receiver, attenuators were deployed after the amplification stage. The delay introduced by the attenuators was measured as a static, additional contribution to the round-trip time of 220-240 ps.

The SFPs (channel 57 for forward transmission and 59 for reverse, 1531.90 and 1530.33 nm, 1.57 nm apart) transmit their optical signal with a power of 0 dBm. The obtained data consist of the same set of measurements for not only the reference setup but for all the experiments: an array of logged temperatures over time, the round-trip time as provided by the WRS monitoring software, and the existing time offset between the two WR devices. The time offset is measured by comparing the 1-PPS output of the devices with a Keysight 53230A time and frequency counter. The climate chamber was configured for four 12-hour setpoints: 15, 30, 45, and finally 22 °C as room temperature. The results are graphically summarized in Figure 5.10.

Probably the most important limitation of this experiment is the difficulty to

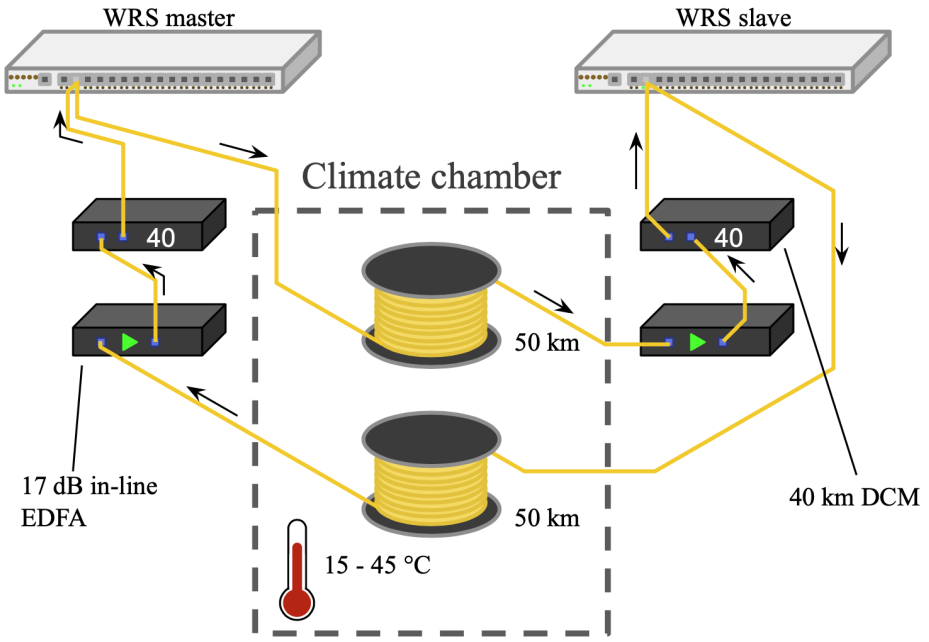

Figure 5.9: Schematic view of the reference setup of the long distance testbench.

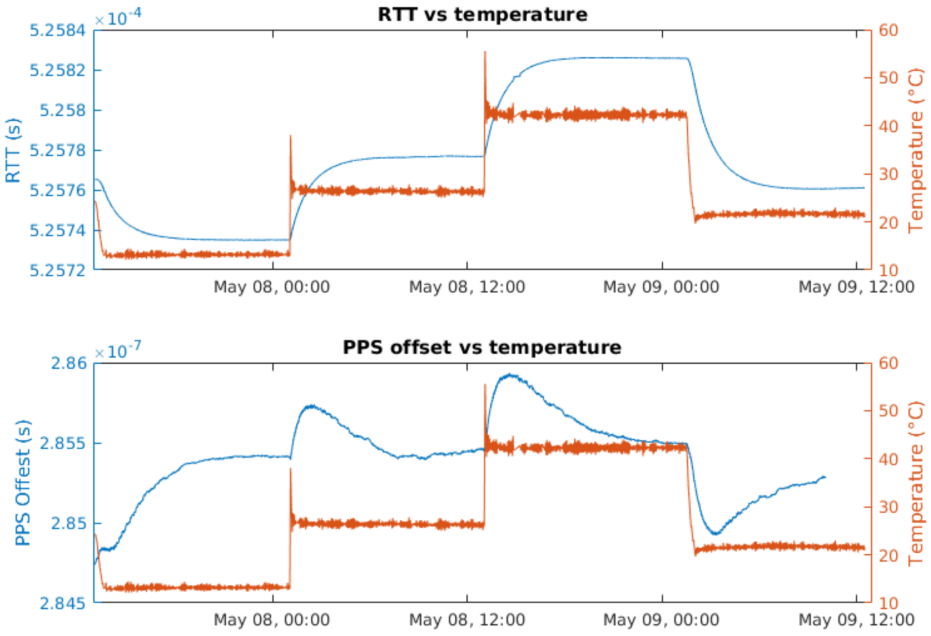

Figure 5.10: Variation of the measured RTT and the 1-PPS output of the reference setup. Data spans roughly 48 hours.

provide heat to the fiber spools in an even fashion. The outcome of placing two fiber spools (with no plastic jacket and uncased) into a climate chamber is substantially different than the temperature changes that an underground cable from a telecommunications operator may face up. Underground cables are buried beneath the soil, which provides a natural insulating layer that helps maintain a relatively stable temperature around the cable. Moreover, the soil surrounding underground cables has significant thermal mass, and this property helps dampem the temperature variations over time. In contrast, the fiber spools in this experiment are subject to extremely abrupt changes in temperature. Still, the evenness of the changes in temperature inside the climate chamber is a challenge: the topology of the spool leads to the core segments of the spool being protected by the outer layers, and heat gradually reaches the core of cylinder. This process has no equivalent in *real world* deployments.

Therefore, we can assume that the unevenness of the heat distribution is an extreme situation, a worst-case scenario that poses an upper bound to the uncertainty that can be expected in the much smoother and damped temperature changes that fibers in production networks suffer.

Although the metodology presents the limitations just mentioned, the results obtained are still significant and noteworthy. The evolution of the round-trim time agains temperature demonstrates that the round-trip time (RTT) is essentially a function of the temperature, and its evolution is directly related to the energy absorbed by the fiber spools. Over the two 50-km spools, the round-trip time has a range of 91 ns. In contrast, the time error introduced by these abrupt temperature changes is much more complex. Instead of being just a function of temperature, the time error depends on the distribution of heat throughout the spools, and only the final tail of every temperature step can be considered out of the transient regime. In any case, transients in the time error due to uneven temperature distribution are limited to 950 ps peak-to-peak after 50 km and two consecutive shifts of opposite sign, and they are kept below 500 ps for a single 20-degree shift. The limited impact is favoured by the use of close DWDM channels that limit the effect of dispersion. Compared to the tests of section 5.2.1, the impact in the time error of transmission using DWDM SFPs in the worst point of the temperature transient is smaller than the error introduced by different pairs of CWDM SFPs under stationary conditions.

In summary, from the 30 °C variation in the reference setup we estimate a linearized positive slope of 31 ps/(km K) or 3132.5 ps/K for the whole link, that leads to a 91.1 ns of span in the round-trip time. At the same time, the time error of the link has suffered a total span of 0.263 ns including the transients.

5.2.3 Impact of the WR devices

The topic of the impact of temperature in the WR devices was subject of research for the radioastronomy users who need the devices to operate under harsh temperature conditions. The analysis obtained a temperature coefficient of 8 ps/K for the WRS and 4 ps/K for the WR-ZEN [129]. A similar test was done in our long distance thermal testbench for the sake of reproducibility (Figure 5.11), with a temperature variation between 15 and 40 °C.

Our results are in the same order of magnitude than the data referenced in the literature, with a measured temperature coefficient of 14 ps per degree for the WRS. We cannot explain the difference between the results in the literature and our own, but it could be caused by slight differences in the methodology, since we ignore how the experiment of the referenced document was set up.

More precisely, the obtained data reflects a slope in the round-trip time of 4.49 ps/K over 25 K, and a 356 ps error over the same temperature gradient.

5.2.4 Impact of DWDM MUX/DEMUX modules

The setup for the MUX/DEMUX, as depicted in Figure 5.12, is limited to placing the case of the device inside the climate chamber, and connecting the minimum necessary fibers for the correct operation of the device under test. Several repetitions were made using different pairs of transceivers.

The results indicate no correlation between temperature and the round-trip time or the time error of the link. However, the delay introduced when using SFPs of channel 57 varies approximately 500 ps when compared to the delay of using SFPSs of channel 59. This can be easily attributed to internal asymmetries in the fiber segments

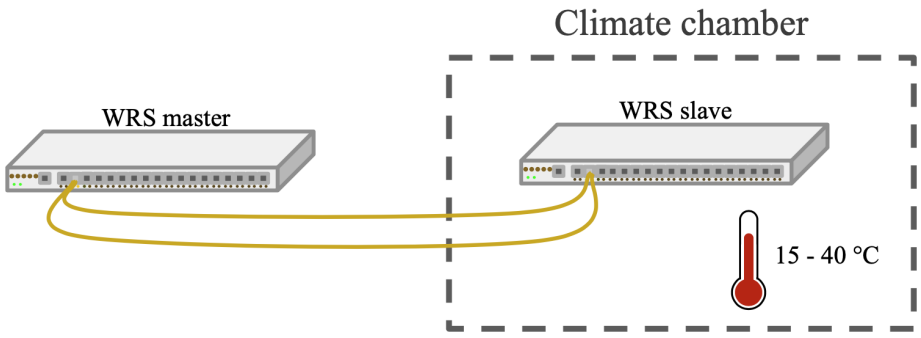

Figure 5.11: Schematic view of the thermal analysis of the WRS.

that connect the front panels of the MUX with the proper components. It is worth noting that 500 ps of time error is equivalent to a fiber delay of just 10 cm. Detailed data can be found in Table 5.2.

Table 5.2: Additive RTT and time error analysis with MUX/DEMUX devices

DWDM pair CH	Temp (ºC)	Time error (ps, uncalibrated)	Additive RTT (ps)
57	13	-1544	56812
59	13	-1033	57827
57	28	-1558	not available
57	44	-1543	58828
59	44	-1031	57733

5.2.5 Impact of EDFAs

The reference test of the erbium-doped amplifier, depicted in Figure 5.13, consists on a master-to-slave path that includes the 17-dB EDFA placed inside the climate chamber, while the slave-to-master path is a simple, direct fiber patch. This setup made necessary the use of attenuators that reduce the optical power to an adequate level for the working regime of the amplifier, as well as for the SFP receiver. The temperature was changed from 15 to 40 ºC.

The results, as graphically depicted in Figure 5.14, show that the thermal response of the EDFA is virtually unmeasurable within the introduced changes with our testbench. There is, however, a clear correlation between the round-trip time data and the WRS FPGA temperature log. This only reinforces the thesis of the signal properties being unaffected by the temperature change introduced by the climate chamber, since the 3 ºC difference in room temperature that suffer the WRSs are more notable than the 25 ºC temperature shift introduced by the climate chamber to the EDFA. The only measurable contribution of the EDFA in this testbench are the increase in round-trip time, which accounts to 87 additional nanoseconds with respect to the simplest setup (two WRS directly connected with DWDM SFPs), and an estimation of RTT dependence with temperature that results in 60 picoseconds of RTT change (Figure 5.15, experiment repeated months later as a verification of previous seemingly uncorrelated data.). Since the variation of time error with temperature would be a fraction of the overall changes in RTT, and these are barely within the resolution of the measuring instrument, we can conclude that the influence of temperature changes in the EDFA time error response are neglibible in the case that the EDFAs for each direction of a two-fiber link are located in the same environment, or kept under similar

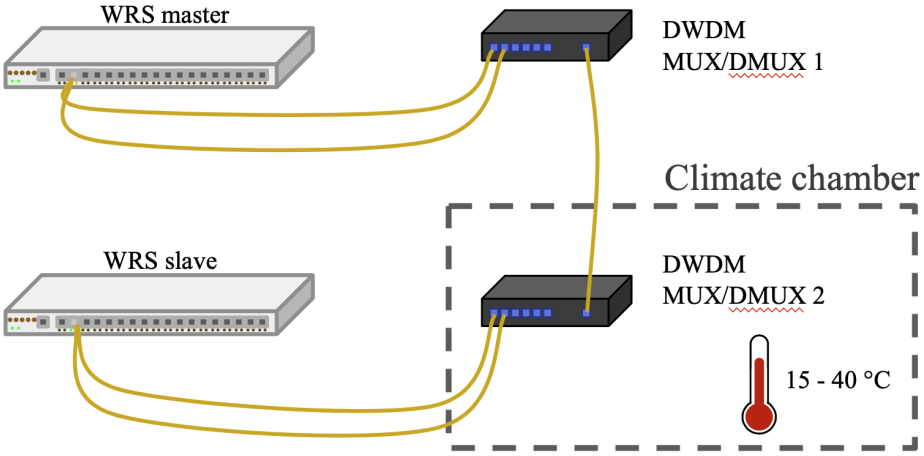

Figure 5.12: Schematic view of the thermal analysis of the DWDM MUX/DEMUX.

temperature conditions.

5.2.6 Impact of DCMs

Lastly, the analysis of the DCMs is performed in the same way as the previous analysis of the EDFA. The DCF is placed in the master-to-slave direction, while the slave-to-master path is left unaltered except for the use of attenuators. Figure 5.16 shows a schematic depiction of the setup.

The DCM module is the element that shows the greatest divergence in both RTT and time error along the temperature tests. Although expected due to the nature of the DCMs as a discrete spool of fiber with a distinct and opposite dispersion profile, this is particularly relevant since DCMs are often located asymetrically over the span of a long distance link. That is, a par of DCMs can be located in different places for each direction in the pair of fibers of a unidirectional link. These DCMs might suffer from different temperature conditions and even compensate for different fiber lengths. Therefore, they can introduce a large amount of offset error, which would require on-site data to be accounted for.

Results show (Figure 5.17) that a DCM that is rated for compensating the dispersion of 40 km of G.652 fiber introduces an additional 22.8 microseconds of link delay. Assuming roughly equal transmission speed, this is equivalent to an additional 4.5 km of effective link length. The dependence of the RTT with the temperature gives a coefficient of 199 ps/ºC. For the calculated additional link length, the thermal coefficient of the DCF is 46 ps/(km ºC), which is slightly higher than the coefficient obtained in

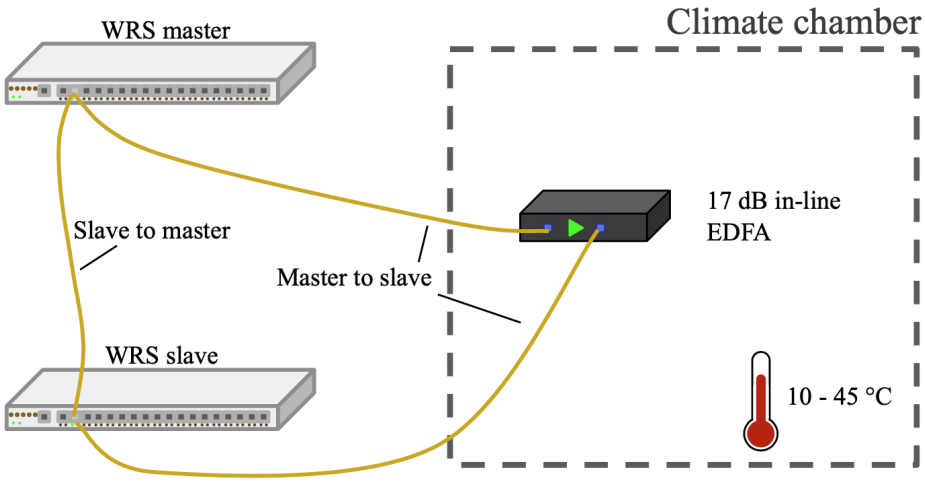

Figure 5.13: Schematic view of the thermal analysis of the EDFA.

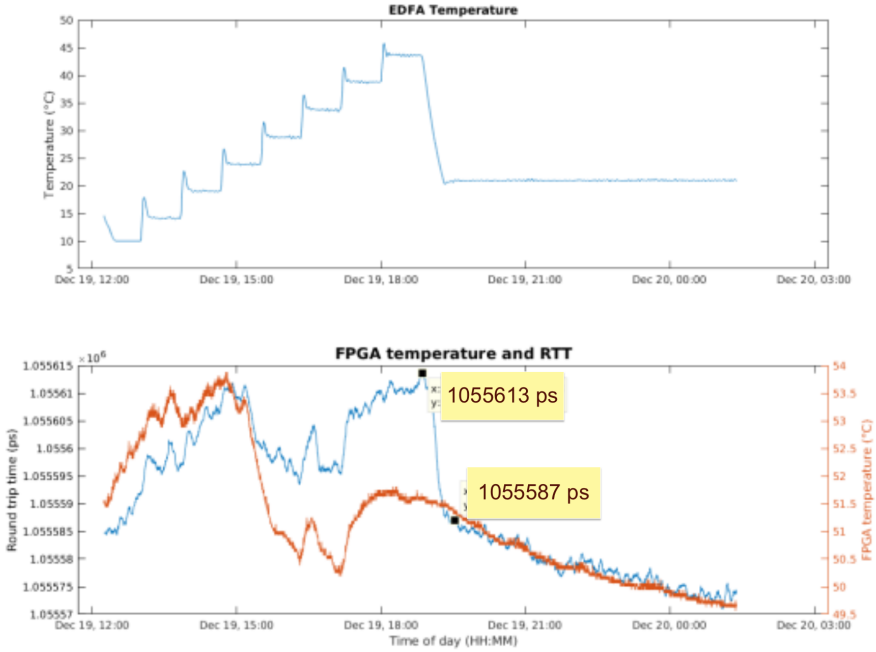

Figure 5.14: Results of the thermal analysis of the EDFA in terms of RTT. The changes of temperature in the FPGA are correlated to the observed changes in the RTT.

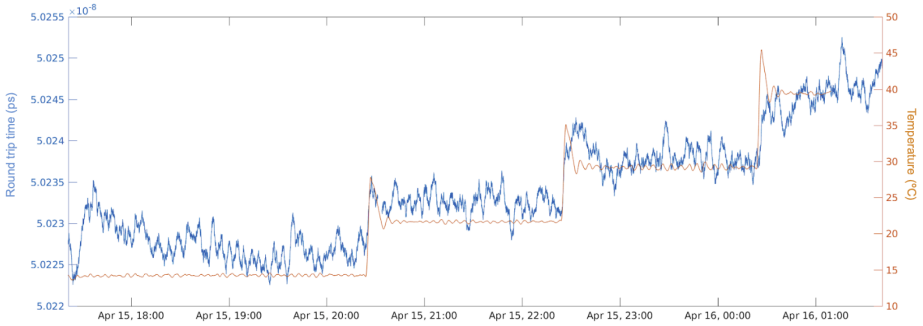

Figure 5.15: Detailed zoom-in of results of the thermal analysis of the EDFA in terms of RTT.

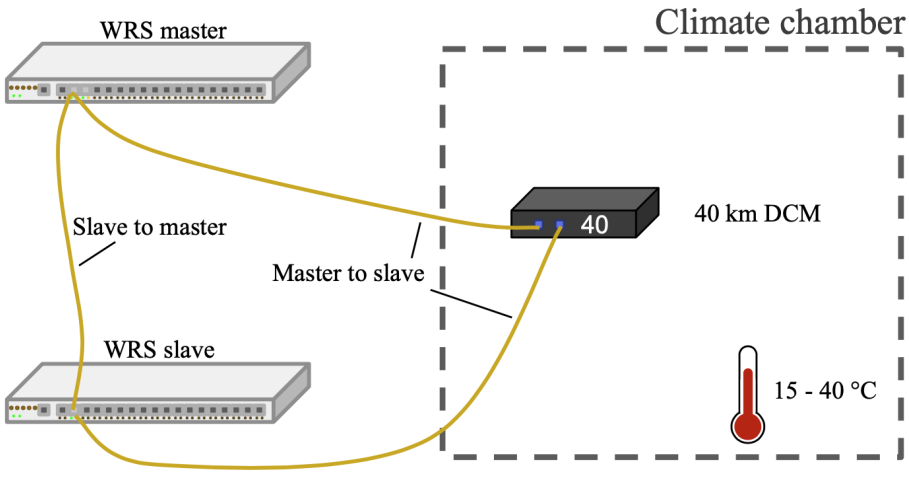

Figure 5.16: Schematic view of the thermal analysis of the DCM.

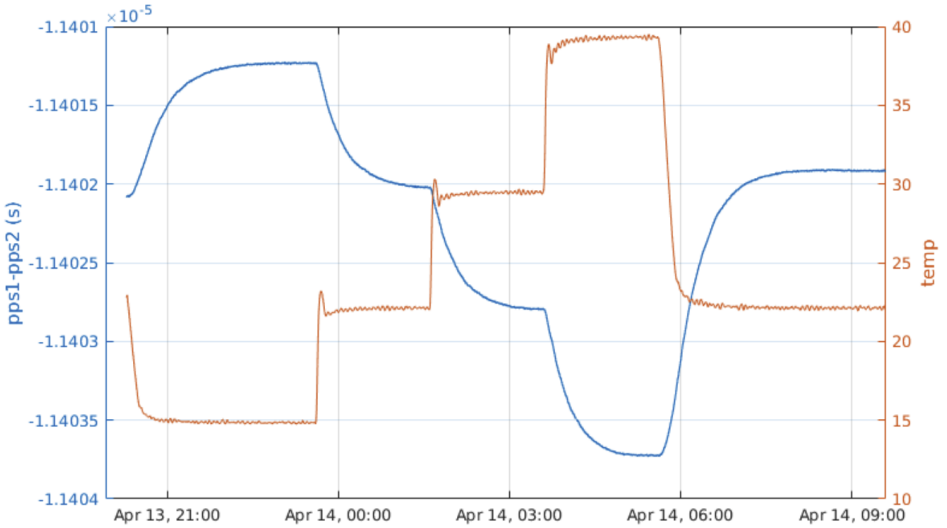

Figure 5.17: Results of the thermal analysis of the DCM in terms of time error.

the reference setup for the transmission fiber (31 ps/(km ºC)).

Nonetheless, the higher (and opposite in sign) dispersion values of the DCF causes that, although the dependence of RTT with temperature is approximately equivalent than in transmission fiber, this does not hold with the time error, and we find a higher dependence of the introduced asymmetry with temperature. Data shows that over the 25 ºC span in temperature, the time error introduced by the change in temperature amounts to -2.632 ns. This is equivalent to a coefficient of -100 ps/ºC, and it is the largest source of asymmetry found in the long distance testbench.

5.2.7 Summary

For the sake of readability, the results and coefficients obtained from all the experiments performed in the long distance testbench are summarized in Tables 5.3 and 5.4.

5.3 Conclusion

In this chapter, we conducted an extensive review of the phenomena that impede transmission in fiber optic networks, such as attenuation and dispersive mechanisms, along with a discussion of the optical network components that compensate or mitigate these effects. We proposed and deployed a testbench in the facilities of Seven

Table 5.3: Summary of RTT data and thermal coefficients.

Setup	ΔT (ºC)	RTT av. (ns)	Temp Coeff. RTT
Reference	30	525776	31 ps km^{-1} $ºC^{-1}$
WRS slave	25	950	4.5 ps $ºC^{-1}$
MUX/DEMUX	25	1059	negligible
40 km DCM	25	23751	46 ps km^{-1} $ºC^{-1}$
EDFA	25	1037	2.4 ps $ºC^{-1}$

Table 5.4: Summary of time error data and thermal coefficients.

Setup	ΔT (ºC)	Δ PPS (ns)	Temp Coeff. time error
Reference	30	0.263	0.175 ps km^{-1} $ºC^{-1}$
WRS slave	25	0.356	14 ps $ºC^{-1}$
MUX/DEMUX	25	negligible	negligible
40 km DCM	25	-2.632	-105 ps km^{-1} $ºC^{-1}$
EDFA	25	negligible	negligible

Solutions that enabled us us to analyze round trip time and time error, both in static and thermal scenarios.

Through the testbench setup, we performed a comprehensive analysis and characterized the behavior of the devices, including the calculation of a temperature coefficient for each of them. These coefficients can be used for future modeling and compensation in a highly managed WR network. Our experiments also demonstrated the significant improvement in long-distance performance when using temperature-stabilized DWDM equipment over low-cost CWDM transceivers.

One key validation from our analysis is that WR can effectively function over networks managed by telecom owners, relying on commercially available WR devices, without any additional modifications. In such scenarios, maintaining sub-nanosecond accuracy in long-distance WR deployments mandates a precise network model and strict control over its management to prevent alterations to the model. This is achievable by interested parties such as telecommunications companies or technological stakeholders with control over metro area networks, given they can supervise or manage the temperature of the end nodes.

In relation to temperature, the obtained analysis over very extreme conditions and abrupt temperature changes are promising. The conditions for the tests of this chapter are a worst case scenario. In real-world conditions, where network equipment is placed in actively controlled environments and the fiber cables are deployed underground and protected from rapidly changing weather conditions, it is feasible

to design a link with a target accuracy of less than 2 ns. Still, the applicability of the laboratory conditions are limited, and more valuable data will be obtained in the pilot deployment in the next chapter.

In contrast, the calibration of delay asymmetry in long-distance WR links over public networks poses a challenge that cannot be addressed simply by brute-forcing very complex network topologies that can dynamically change. These asymmetries, while having an impact on the accuracy, are stable, which is a crucial point of consideration. Their stability over time means they can be effectively compensated, allowing for a maintained level of accuracy. An alternative approach for this category of links could entail a black-box calibration that encompasses all the delays of each direction of the network within a single parameter. This approach will inherently require reliance on alternate time sources for calibration, such as those assisted by augmented GNSS timing techniques. In this case, the accuracy would inherently be constrained by the precision of the GNSS time equipment.

While we have successfully established that WR can operate effectively within telecom public DWDM networks, our analysis also revealed the presence of devices responsible for fault mitigation that can trigger network reconfiguration. Such network reconfigurations could significantly impact the timing system if not appropriately managed. It is imperative to have robust resiliency mechanisms in place to handle these scenarios, such as redundant time references and a robust holdover solution. These will ensure that the system continues to function effectively during network reconfiguration, thus preventing disruptions and maintaining accuracy at all times.

The insights gained in this chapter will be applied in the following one, where we will design, deploy, and evaluate an experimental time service over the Madrid metropolitan area. Overall, this research contributes to the understanding of the limitations of the current calibration approach of WR, and to characterizing the error budget that can be expected over public networks.

Chapter 6

Time as a Service for finance based on WR

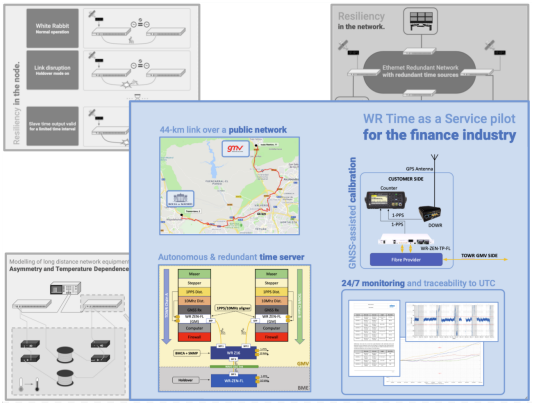

6.1 Introduction

In recent years, the financial industry has undergone significant changes as technology has made the action move from the raucous trading floors of Wall Street to the humming and blinking lights of suburban datacenters. There, amidst rows of relentless machines, the transactions happen at a pace that makes human intervention obsolete and quaint. One particular area of interest has been the use of high-precision time synchronization in the management and tracking of financial transactions. As the volume of these continues to grow, the need for precise and legally traceable timestamps has become crucial. Regulatory bodies are increasingly requiring financial institutions

to provide proof of accurate timestamps tied to a universally accepted reference.

White Rabbit (WR), a technology initially designed for synchronization and data transfer in scientific institutions, has found its way in the world of finance. Its unique ability to provide precise time synchronization over the optical media makes it a promising solution to the challenges currently faced by the industry. In this chapter, we explore the competitive and regulatory forces that push towards the adoption of precise time transfer systems in the financial industry, and we discuss the technical details that are specially relevant due to regulatory pressure. After that, we describe a pilot deployment that was made in the framework of the TOWR project, showcasing the operation of the resiliency mechanisms that have been previously described in this work, and demonstrating the feasibility of the use of WR over a public optical network to guarantee the access to traceable, precise and robust timing signals.

6.1.1 Overview of timing in the fintech industry

For most part of human history, trade primarily relied on human intervention, as agents engaged in the exchange of goods were required to be physically present in a centralized location, such as a bustling marketplace or a trading floor.

Nonetheless, as early as the advent of telecommunications networks, technological advancements have continually been employed to bridge the gap imposed by physical limitations. A notable example of this can be traced back to the 1930s, when a telegraph cable was strategically laid between Newfoundland, Canada, and the Portuguese Azores archipelago [135]. This connection enabled brokers to receive updates on the disparities of stock quotations on either side of the Atlantic, using the information on spread prices to their advantage.

In more recent times, with the development of modern information technology and the ubiquitous reach of computer networks, trading has moved from the trading floor to electronic markets. Still, the principle of operation of a stock market remains unchanged from its non-electronic predecessors. An electronic market transfers ownership of stocks from one investor to another, and the role of the market operator is to provide the platform to accomodate the transactions, ensuring that the integrity and fairness of the market is maintained. In the basic scenario of a limit order book, the market lists all the pending buy and sell orders, showing the number of shares and the price that buyers and sellers are willing to accept. *Bids* represent the buy orders in the order book, while *asks* represent the sell orders. When a buyer and a seller agree on a price, a trade occurs, and the respective orders are removed from the order book. The priority for orders to be executed is typically based on two parameters: price, and

then time.

The will to outperform other market actors in relation to those two metrics has resulted in ever-increasing technological developments. Since the advent of electronic trading and the transition from traditional trading floors, financial markets have undergone a significant transformation. The move to electronic markets has facilitated faster, more efficient trading, and this shift has sparked a technological arms race, with various actors seeking to develop the most effective algorithms and achieve the fastest possible response times. Algorithmic systems automatically place orders to buy and sell following predefined strategies without human intervention. By 2019,algorithmic trading accounted for 93% of total market volume in the United States [136]. More specifically, a subset of these algorithmic trading systems has drawn the attention of the public and regulators: high frequency trading (HFT).

High frequency trading is one of the key drivers behind this technological race. It can be loosely defined as a kind of algorithmic trading that involves a very high amount of operations in a very short-term horizon. HFT actors capitalize on minuscule price differences and fleeting market inefficiencies, making profits by rapidly entering and exiting positions. In this context, network latency is of critical importance as access to information a few nanoseconds before a competitor can be the key to a successful bid. Consequently, HFT firms have been investing heavily in technology infrastructure such as low-latency network connections and powerful trading servers located near exchanges (known as co-location). These investments aim to minimize the time it takes for trade orders to be transmitted, executed, and confirmed, allowing traders to react more quickly to changes in market conditions. The pursuit of faster execution has even led to the deployment of low-latency microwave transmission links and free space laser links, which can transmit data faster than traditional fiber-optic cables [137], but the performance is not limited to the latency parameter: response times of less than 100 nanoseconds were registered by 2018 in the German operator Deutsche Börse co-location [138].

However, this technological struggle has raised concerns among regulators and market participants. Back in 2010, the United States Securities and Exchange Commission (SEC) raised the question whether HFT has led to an unequal playing field where smaller firms and individuals do not have the resources to compete with better-funded institutions [139]. Additionally, the rapid pace of trading and the potential for algorithmic unintended consequences have raised concerns about market stability and the possibility of flash crashes, where markets experience sudden and severe price drops within a short period[140], [141]. Since then, the US market regulator has

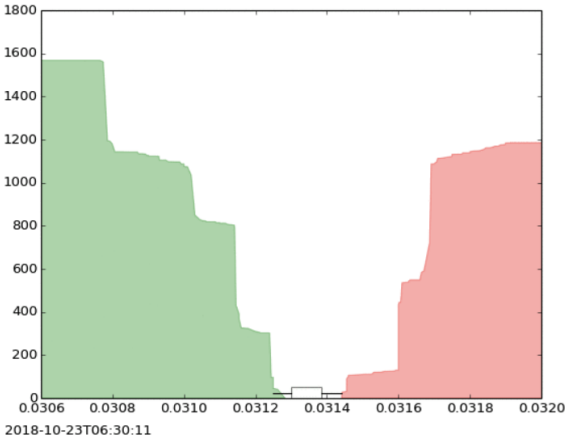

Figure 6.1: Order book depth chart from a currency exchange. The x-axis reflects the unit price, while the y-axis depicts the cumulative order depth. Bids are shown in green, while asks are shown in red. From "Order book depth chart", Wikimedia Commons, CC-SA-AT

proposed new rule changes but this is as of the time of writing an ongoing topic [142].

In the European side, the MiFID II (Markets in Financial Instruments Directive II) [143] and MiFIR (Markets in Financial Instruments Regulation) [144] pushed by the European Securities and Markets Authority (ESMA) came into force in July 2014 in the European Union. Among other goals, the directive aims at closer monitoring of algorithmic trading. The MiFID directive explicitly defines algorithmic trading and HFT firms subject to stringent risk controls including pre-trade risk checks, real-time monitoring, and post-trade reporting (Article 17 of MiFID II). The directive also obliges trading venues to provide equal access to the co-location services, as well as keeping transparent and non-discriminatory fees for these, among other provisions aimed at increasing transparency and ensuring market transparency. In particular, Regulatory Technical Standard (RTS) 25 [145] details the requirements relative to time synchronization with which HFT firms must comply. A selection of key provisions include:

P1 Clock synchronization requirements: RTS 25 mandates that trading venues and market participants including algorithmic trading firms synchronize their business clocks according to a common reference source. For trading venues in the EU, the reference time source is Coordinated Universal Time (UTC), or the timing signals of a satellite system, provided that the offset from UTC is accounted for.

P2 Granularity of timestamps: trading venues and market participants involved in HFT must recort their timestamps with an accuracy of at least one microsecond.

P3 Maximum divergence: RTS 25 sets limits on the maximum divergence between the trading venue or market participant clock and the reference time source. The limit is set to 100 microseconds for trading venues and HFT firms.

P4 Record-keeping and traceability: trading venues and market participants are required to maintein records of their clock synchronization processes, including the procedures and tools used to synchronize their clocks, and any deviations from the reference time source.

In the following section we provide an brief overview of the technical difficulties and implications of each of the requirements of the previous list.

6.1.2 Technical requirements of a financial timing service

Many of the topics of this section have been discussed before in this work. For the sake of simplicity, we will mention all of them and point to the dedicated section if necessary.

Synchronization with UTC

UTC is the global primary time standard currently in use, conformed by contributions from over 400 atomic clocks maintained by national metrology labs and other timing centers from all around the world such as the National Institute of Standards and Technology (NIST) in the United States, the Physikalische Technische Bundesanstalt (PTB) in Germany, and the Real Observatorio de la Armada (ROA) in Spain. Each of the physical realizations of UTC is refered to as UTC(k), and UTC is only calculated *a posteriori* from weighted averages of the times reported by the contributing timing centers.

UTC is, in fact, derived from International Atomic Time (TAI) after a correction of a given amount of leap seconds in order to keep the difference between UTC (based on fixed-length day) and the mean solar time at 0 Meridian below 0.9 s due to the irregularities in the Earth's rotation period.

The selection of the UTC(k) local representation to be considered as reference is a previous necessary step in the design of the timing system and is necessary to fully comply with the listed provision P1. In some cases, this selection may be determined by regulation mandating the use of a national representation of UTC. The design of

the timing system must include the capabilities to compare the time reference used at the customer facilities to a concrete UTC(k) implementation.

Divergence from UTC

As extensively discussed in Chapters 2 and 4, every physical clock is imperfect and does not "tick" with an ideal period. Even more, it will slightly change its rate over time due to several random processes. Consequently, two different clocks will inevitably drift from each other if left uncorrected. In order to align a time source to UTC, it is necessary to correct its timebase and also its frequency (synchronization and syntonization). To overcome this problem, a mechanism that repeatedly corrects the error must be established (section 2.4.3). As any observant reader could predict, the time service that is being described in this chapter is based on WR.

Traceability

The requirement of establishing a system of traceability to a realization of UTC implies that the measurement of a magnitude at the customer side can be related ot the time reference (UTC(k) or the timing signal of a satellite system) through a documented unbroken chain of calibrations, each contributing to the measurement uncertainty. The traceability of the measurement relies on three elements [146]:

- Calibration: the internal delays of the components of the timing system must be documented and compensated.
- Determination of uncertainty: every measurement must be accompanied by an indication of the quality of the result.
- Monitoring and storage: continuous polling of the system parameters and permanent storage of the timestamped relevant events serve as a way to detect malfunctioning and enable future auditing.

Another requirement not directly mentioned in the regulations but essential for any service is availability. A system aims to ensure a certain level of operational uptime, typically agreed in a service level agreement (SLA). Redundancy in parts of the system and contingency measures to minimize the impact of failures are examples of mechanisms to improve the availability of a service.

6.1.3 Time as a Service

From the previous discussion and the contents of the previous chapters, it is easy to conclude that the implementation of a traceable and robust time dissemination system is not trivial. For that reason, many Internet Service Providers and metrological centers are providing outsourced time solutions to customers in what is known as time as a service (TaaS). In this approach, the TaaS provider is in charge of transporting a traceable, highly-accurate time reference to the facilities of the customer, and guarantees its reliability and integrity. The rise and development of cloud computing and the increasingly prominent role of the data center has favoured the development of TaaS solutions that the cloud providers or data center management provide to its customers. This strategy allows data center compaines to provide different tiers of service according to application requirements, and the shared infrastructure between many customers is more cost-effective for businesses that can access highly-accurate time synchronization services without incurring the capital and operational expenses of building and maintaining their own infrastructure [147] [148].

Many actors involved in the dissemination of time, such as NMIs, have partnered with Internet Exchanges (IX, i.e. locations where ISPs and other network operators interconnect their networks to exchange traffic and provide better connectivity to their customers) to provide accurate, traceable TaaS solutions as part of the offerings of the IX. Many financial firms choose to co-locate their servers in data centers in IX points to connect directly to other financial firmas and trading platforms without having to go through third-party networks. This service allows the financial firms to outsource the time system to a reputable source that guarantees the integrity and reliability of time signal. Some examples of TaaS providing companies with customers in the finance sector include:

- NPLTime: this is a time service provided by the National Physical Laboratory (NPL) from the United Kingdom. It provides its customers with a time service distributed using IEEE 1588-2008, achieving an accuracy better than 100 ns traceable to UTC(NPL) along dedicated channels and PTP-compatible switches [149]. The time service includes monitoring and backup references in intermediate points of the link.

- Hoptroff Traceable Time as a Service (TTaaS): this network-delivered time synchronization solution achieves an accuracy of up to 100 ns using PTP. The service is designed to integrate with existing hardware and software infrastructure, not only oriented to timing but also to media broadcasting and gaming, among

others. It receives its satellite reference from GNSS systems, but also counts with terrestrial feeds from the NMIs from Sweden and the United States. The service includes telemetry and diagnostic solutions, and provides compliance for the regulatory requirements of reporting and monitoring. The timing infrastructure and data storage are secured under ISO 27001 information security standard [150].

- Equinix Precision Time, a time service provided by the global datacenter company Equinix that provides NTP and PTP time in all the areas served by Equinix private networks. The clock source is derived from GNSS sources and distributed using WR technology, but the *last mile* is disseminated only using NTP and PTP. For that reason, the expected accuracy is of 30-40 microseconds for NTP, and 1-10 microseconds for PTP. Redundancy capabilities are provided in all the supported regions. Although the service is compliant with MiFID, the typical accuracy levels are not guaranteed *as it depends on various factors such as geographic location, Equinix Fabric characteristics and customer-side equipment*, according to the service documentation [151].

- GTT time service: in this case, a reference clock is continually steered to UTC(NIST), the representation of UTC from NIST; via GPS time transfer [152]. It achieves time uncertainty of less than 15 ns at the point of the grandmaster clock. The time is distributed by the grandmaster close to co-location centers to customers using PTP or NTP. Holdover achieving 1 microsecond error is provided at the grandmaster side and traceability to UTC(NIST) is certified.

In this context, the next section describes the project that was carried out to deploy a project time service in the Madrid metropolitan area. After an overview of the project, we discuss the details of the link architecture, focusing on the time dissemination aspects from the site of the time source to the customer destination, and the associated resiliency and monitoring features.

6.2 TOWR project

The pilot described in this chapter relates to the TOWR project. The objective of TOWR is the development of a robust and accurate time distribution service via optical fiber using WR over the Madrid metropolitan area, keeping traceability to a UTC realization. This project was a joint effort between GMV, who was in charge of

the time source generation, and Seven Solutions as an expert in highly accurate time distribution via optical fiber.

The pilot was designed with the banking and finance applications in mind, in order to fulfill the business clocks synchronization requrements that have been discussed in section 6.1.1

In this context, a pilot project was carried out by Seven Solutions and GMV to provide time and frequency services in the Madrid metropolitan area combining GNSS time transfer, redundancy features, and state-of-the-art time distribution via optical fiber. A detailed description of the time dissemination side and the results of its operation can be found in the following section.

6.2.1 Link overview

The deployed time service is located in the northwest area of Madrid, connecting the site of the primary references to the datacenter of the customer. The fiber path has a length of 44 km between Tres Cantos and Las Rozas bordering the *Monte de el Pardo*. The primary references are located at GMV facilities, and the customer for this pilot project is Bolsas y Mercados Españoles (BME), the operator of all stock markets and financial systems in Spain. Figure 6.2 shows a map of the time source generation site, the destination, and the path of the 44-km link.

6.2.2 Time server hardware

One of the distinctive features of this time service is its aim ot mitigate the weaknesses of the satellite solutions. GNSS is not directly used to disseminate a time reference. Instead, the time is generated from autonomous atomic clocks, while GNSS is used only for fine corrections by means of time comparison to UTC.

The time server yhat provides the time reference for the system is a passive hydrogen maser (PHM). This atomic clock generates a very stable signal that is fed to a phase and frequency generator, while common-view GNSS time transfer (see section 2.4.3.1) is used to obtain the difference between the local time reference and the remote UTC timescale. The corrections are gradually applied to steer the output frequency of the maser toward UTC. Then, a distribution amplifier is used to provide signals for the WR grandmaster clock, which is the role performed by WR-ZEN.

To achieve full redundancy, all the hardware involved in the generation of the time signals at the time generation side is duplicated, from the PHM to the grandmaster WR devices. A brief schematic view of the time server hardware can be found in

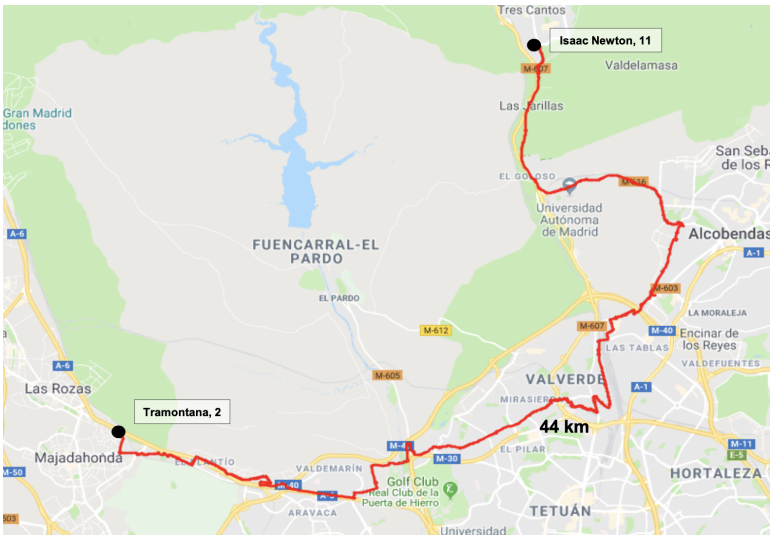

Figure 6.2: Map showing the time source generation site, the customer datacenter site, and the link path.

Figure 6.3, but readers are refered to the associated literature since the involvement of Seven Solutions in the implementation of the Time Server was not predominant [153]. The reference contains details on the time server hardware, calibration and operation.

6.2.3 Network time dissemination

Each of the hardware stacks in the time server ends up in a WR-ZEN node that acts as WR Grandmaster. Both are connected to a time distributing device (WR-Z16) via one link each. This device acts as a gateway between the time source site and

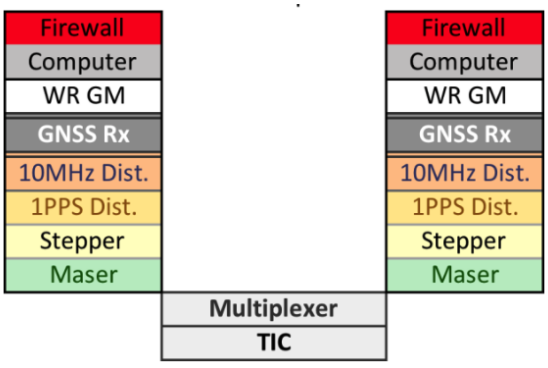

Figure 6.3: Schematic view of the time server hardware, demonstrating the redundant elements. Courtesy of Safran Electronics and Defense Spain.

the destination site, and it accomplishes two main roles: first, it performs the Best Master Clock Algorithm (BMCA, section 2.4.3.4) and selects the time server that is disseminated throughout the network based on quality parameters reported by the devices. In case of failure of the selected reference, the WR-Z16 performs a failover operation and disseminates time recovered from the backup reference. The other role of this device is to provide a fan-out for multiple customers or devices by means of WR, PTP and NTP protocols.

Ideally, the two time sources would be routed throughout the destination using different routes for the main and backup links. However, due to budget constraints, it was decided to keep one single link from the time source site to the destination, and the BMCA is done at the source site. This introduces a single point of failure in the system, but allows us to extend the time of operation of the pilot by reducing the costs of the link. Figure 6.4 depicts an overview of the link architecture.

Once the time signals are enconded as an Ethernet stream, they are sent to the customer site by means of a dedicated optical channel through a public network. The service provided by Colt, Private Wave, refers to the concept of dark channel introduced in section 2.5, although the access to the core network is not direct. The physical realization of the link is done in a DWDM unidirectional link, but the connectivity between the WR devices and the core network is done via a multi-mode 850 nm short-reach link to a transponder. This means that the information about the wavelength used and any other specifics of the optical transport network is unknown for the time service management. In any case, since the main section of the link is done using unidirectional links whose lengths, asymmetries and equipment are unknown, the calibration technique for this link must to be medium-agnostic. This will be further discussed later.

In addition, Chapter 5 performed an analysis on the sensitiveness of optical network components to temperature changes. In the TOWR link we expect changes to be both daily (day/night temperature difference) and seasonal. For this reason, the temperature in the endpoints of the link and local weather conditions in the surroundings of the link path were monitored. This effect is mitigated in the end nodes because data centers are temperature controlled, but weather can cause the round-trip time to vary as much as hundreds of nanoseconds, as proved in section 5.2.2. The results determined that over a 50 km unidirectional link every additional Kelvin causes an increase of 3.1 ns in the round-trip time, but less than 10 ps error in the 1-PPS signal of the WR devices. Translating these results into the 44 km link, and assuming a pessimistic 30 ºC difference for buried fiber between the hot and the cold season, the

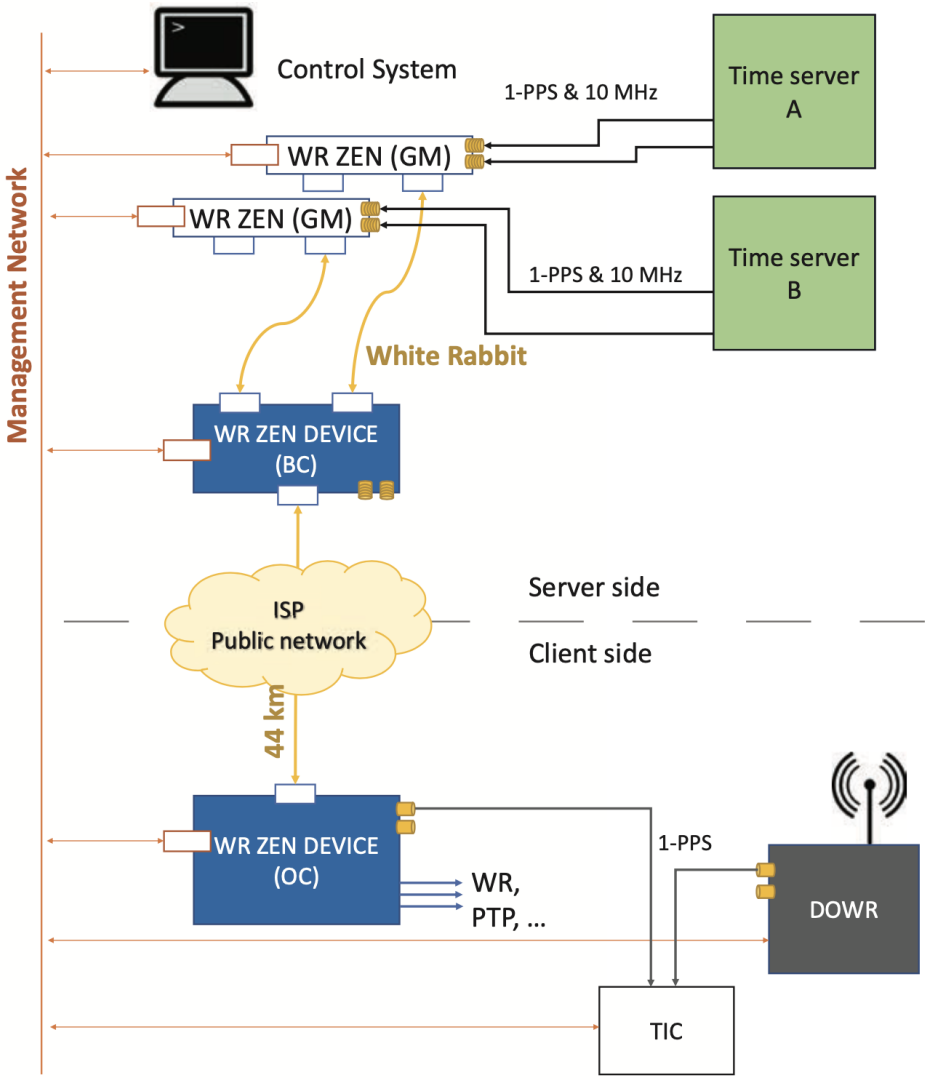

Figure 6.4: Overview of the TOWR link architecture. Both time servers provide a WR grandmaster with an independently generated time reference. A WR-Z16 acts as boundary clock, and selects the time reference that is transmitted through the network. The client side is continuously monitored and compared to a DOWR GNSS receiver. Courtesy of Safran Electronics and Defense Spain.

error in the time output is estimated at 231 ps. Thus, the impact of weather conditions in the timing performance remains below the uncertainty of the calibration described in the following section.

6.2.4 Customer side and Calibration

The link ends at the facilities where BME holds its datacenter. The endpoint of the link, after the transponder stage of the ISP, is a WR-ZEN that acts as slave of the WR link. Here, the time reference is reconstructed from the Ethernet stream, and the device generates different signals and protocols to accomodate the services required by the customer, such as WR links to other devices, 10 MHz and 1-PPS analog signals, PTP and NTP.

The internal delays of the WR devices and the delays introduced by the SFPs were calibrated prior to the deployment of the link, and so were the cables involved in monitoring and calibration. The only component of time error yet to be calibrated is the asymmetry between both unidirectional fibers. In order to align the time differences on both sides of the 44 km link, the calibration at the customer end was assisted by a WR-capable GNSS receiver that uses GPS time as a reference. This device, called DOWR[1], has its antenna, cables, and internal delays calibrated with respect to a setup certified by PTB. The calibration procedure takes several days due to processing and compensation of the ionospheric delay and has an uncertainty of ± 10ns. Since the time servers are steered daily to follow the disseminated reference UTC(PTB), and the difference between the GPS time and the German realization of UTC is published by time centers, the time output at the client side can be aligned to the one provided by the active time server, ultimately keeping the traceability back to UTC(PTB). The average measured asymmetry between the two fibers of the link was determined as 13.15 ns, which translates to an estimated asymmetry length of 5.37 m assuming an estimated group refractive index of 1.468 (from Equation 2.16 and its relationship to the RTT). The achieved accuracy is compliant with the provisions stated in P3 with a margin of more than three orders of magnitude with respect to the maximum acceptable divergence as defined in provision P3 from section 6.1.1.

It is important to note that the DOWR, the GNSS receiver used for the calibration of the link, does not support the common-view GNSS time transfer technique. This omission introduces a minor offset between the received time reference and UTC(PTB), the UTC realization from the German NMI that is selected as refer-

[1] `https://sevensols.com/dowr/`

ence for the TaaS time servers. While these deviations are minimal and can be rectified a posteriori through post-processing techniques, they constitute an unnecessary source of systematic uncertainty. It would be possible to eliminate these discrepancies from the onset by employing a GNSS receinver that is capable of supporting CV techniques. By utilizing simultaneous observations of the same GNSS satellite from two different locations, CV allows for the direct comparison of clocks at two locations, hence eliminating the error associated with the GNSS satellite clock and the delays in the GNSS signal propagation. This enhancement could facilitate an even more robust traceability chain to the primary source of time.

Figure 6.5 provides an example of normal operation of the link over 25 days. Neither the time server nor the WR time transfer deviates from its original calibration. The main contributor for the delay oscillations is the variable ionospheric delays introduced by the GNSS-based monitoring device, and average to zero over time.

6.2.5 Resiliency features

A key strength of the proposed solution resides in the autonomy and resiliency of the time servers, minimizing the system's dependence on GNSS. In the TOWR project, the role of GNSS is specifically limited to daily time transfer operations rather than its use as a continuous time source. This daily process involves checking and calculating the error with respect to the GNSS time reference, followed by a gradual steering process designed to minimize the difference. In addition, since the PHMs are highly stable sources of time, the corrections done in the steering process are predictable to some extent. Consequently, the reference not only becomes robust against jamming (the model of the time source [153] predict that the time servers can operate autonomously for up to a month before the time error reaches 10 ns), but also against spoofing. The location of the GNSS antenna is very accurately known, and a spoofing attack would need to have the necessary level of sophistication so that the spoofing signals serve as a valid solution at the location of the antenna, and that solution must be illegitimately correct over time. Also, a sudden change in the nature of the steering commands would raise warning alerts. This design enhances both the self-reliance and the resiliency of the time servers, effectively increasing the system's overall stability and security.

The second feature to increase the fault tolerance of the service against failures consists of fully duplicating the time servers. During normal operation, both stacks are independently functioning, and the frequency steering is calculated and applied individually. The boundary clock will swap the time reference that provides the time

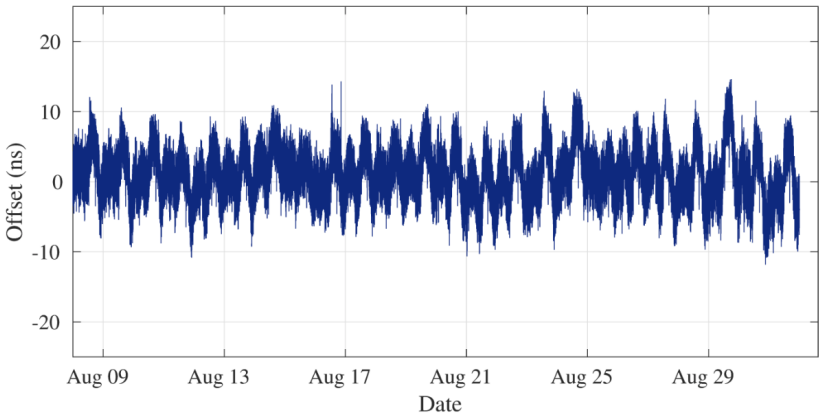

Figure 6.5: Offset between the customer node and the monitoring DOWR over 25 summer days. The 24-hour periodicity is due to the variable ionospheric delays introduced by the GNSS receiver. Courtesy of Safran Electronics and Defense Spain.

service when a contingency is detected using the BMCA. The ideal scenario for this situation is to have at least three autonomous time references, convert this situation to a Byzantine failure problem, and determine the consensus on the failing component via majority voting. However, the time server are costly pieces of equipment that cannot be implemented in multiple instances without considering budget constraints. In the current configuration of the pilot, the degradation of service is communicated by the grandmaster through the clock class parameter in the PTP *Announce* exchanged messages.

Figure 6.6 shows the behavior of the 1-PPS of the time signal at the customer side, in an event where the time server reference was changed four times over the course of 20 minutes. The graph shows that the 1-PPS is not received during reconfiguration (hence the large spikes, representing timeouts of the measuring instrument), and a slight time offset is appreciated in the mean time error levels between the active and the backup time server.

Finally, in the case of any kind of service disruption detected at the client side, the customer side WR-ZEN is equipped with a commercial solution to improve holdover time based on an oven-controlled crystal oscillator. This holdover oscillator is capable of providing a 1-PPS reference with up to 1.5-microsecond-bound time error for 24 hours usign a proprietary adaptive algorithm. Figure 6.7 shows data from a series of holdover executions. Since the link was disabled during the test, a DOWR in the client side is considered as reference. After 12 hours of holdover operation, the worst observed case stays below 100 ns, and the 1-microsecond threshold is broken after 26

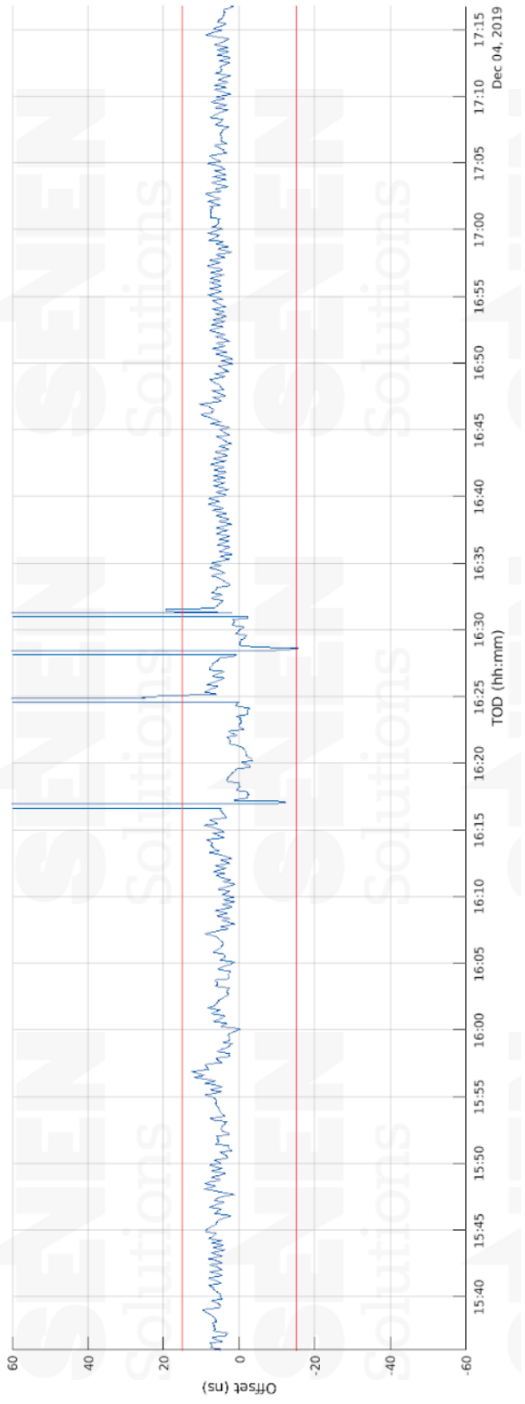

Figure 6.6: Time output at the customer site during a period of succesive time server failover events. Spikes in the 1-PPS signal represent lack of 1-PPS signal during reconfiguration. Holdover capabilities are disabled. Courtesy of Safran Electronics and Defense Spain.

hours in the worst observed case. In contrast, Figure 6.8 shows the typical operation of the link, where tests were limited to market holidays and weekends to comply with the established SLA and restored before market hours. The holdover mechanism ensures that the node equipment is ready to operate autonomously for days before reaching the 100-microsecond accuracy limit stated in provision P3.

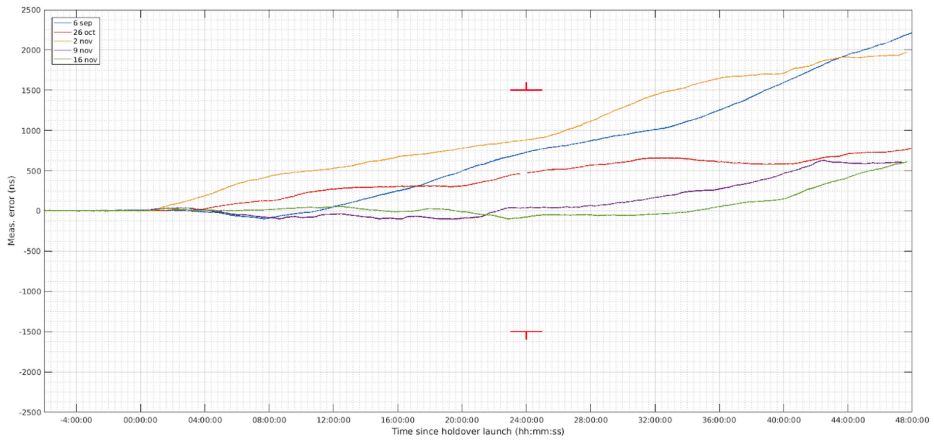

Figure 6.7: Comparison of five holdover executions of TOWR holdover in different weekends. The maximum expected error of 1.5 microseconds after 24 hours is highlighted in red. Courtesy of Safran Electronics and Defense Spain.

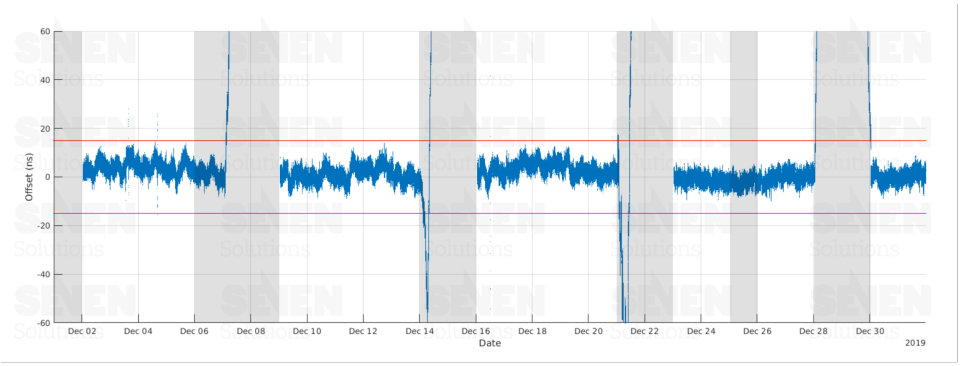

Figure 6.8: Daily time error reported for the month of December. Every Saturday the link is automatically disabled to evaluate the performance of the holdover solution. The limits of the time error established in the SLA (a maximum of 15 seconds error from initial calibration) are shown as horizontal red bars. Holidays and weekends when the market is closed are shown in grey. Courtesy of Safran Electronics and Defense Spain.

6.2.6 Service monitoring

As stated in P4 from section 6.1.1, the MiFID II directive mandates that financial firms must keep records of their clock synchronization and timestamping processes for at least five years. This obligation is critical for regulatory authorities to assess compliance and perform audits and forensics in the case of any relevant reportable event. The records should include details on the time sources used, the procedures employed for synchronization and timestamping, and any configuration changes, software updates, or any event that may have affected the synchronization process, but also documentation on the detected discrepancies and corrective actions taken, as well as records of regular monitoring and testing activities.

Regarding monitoring, entities subject to this regulation must conduct periodic assessments of the synchronization mechanism to verify its accuracy, reliability and traceability to UTC, implement automaetd tools to monitor the business clocks and identify deviations from the required level of accuracy, and develop contingency plans to address failures or disruptions promptly.

The pilot presented on this project implements a monitoring system that checks the health status of the link and the end devices at each side. All the parameters are polled every 2 seconds, and data is securely stored in external servers, available for record purposes and potential audits. Besides the GNSS steering that is done in the time servers and also has the function of monitoring with respect to an external reference, the time servers themselves are duplicated. In the server side, there is scientific-grade instrumentation dedicated to the continuous monitoring of the time signals compared to a GNSS-calibrated receiver.

Also, the SLA document establishes a series of metrics that can be objectively determined, such as the availability of the service, the tolerable offset to UTC, and the correct operation of the contingency mechanisms. The management system of the time service computes the data from which these metrics are derived, and the results are stored for future use. Additionally, periodic reports are shared with the customer regarding the performance and the detected events. Besides these reports, alerts are issued to the customer as soon as the agreed service level is compromised.

The monitoring system and its associated data is kept separate from the timing network, as depicted in orange in Figure 6.4. This segregation ensures that even in the event of unexpected disruption or system failure, the timing network remains unaffected and keeps providing service at the customer. Likewise, were the network disruption originated by a timing node or event, the monitoring system can still send

alerts and provide crucial information for troubleshooting and forensics purposes.

6.3 Conclusion

This chapter presents the motivation, architecture and results of the operation of a pilot time service deployed to the Spanish stock markets operator, following the TaaS in which the customer outsources the deployment and management of the timing network and equipment.

We began by examining the driving forces behind the finance industry's adoption of precise time systems for their operations. Our analysis revealed that the motivations come from both competitive and regulatory pressures. Importantly, the legal obligations of traceability and accuracy apply not only to the cutting-edge actors that benefit from it, but to all industry players in order to ensure transparency and fairness in the sector.

Addressing the technical challenges faced by operators due to the geographically extensive nature of finance firms, we identified that commonly deployed solutions struggle to achieve the required accuracy or can be susceptible to external interference. To overcome these limitations, our pilot scenario successfully implemented a solution that combined redundant atomic clocks aligned to UTC via GNSS time transfer with optical fiber time transfer using WR technology. This approach provides customer end nodes with highly accurate time capabilities deployed over a public ISP network.

The customer side was calibrated making use of GNSS techniques to within 10 ns of the observed reference, and WR time transfer transmitted the reference from the atomic time servers with an accuracy 1000 times more stringent than the limits imposed by regulation. Regarding holdover, the time can be kept up to 60 times more accurate than these same limits after 24 hours of autonomous operation of the slave side equipment. In summary, the deployed time service solution not only meets current requirements but is well-equipped to handle the demands of the foreseeable future.

Chapter 7

Conclusion

This concluding chapter synthesizes the primary findings, discussions, and contributions of the thesis. Initially, a summary of findings is provided, offering a succinct recapitulation of the outcomes from each chapter. Next, the main contributions section evaluates the extent to which the initial objectives have been fulfilled, providing a clear linkage between the goals set and the outcomes achieved.

Following this, a list of scientific publications related to this work in which the author has made a contribution is presented, signifying the research's reach and impact in the academic community. The chapter concludes by identifying potential future research directions, highlighting areas where the robustness and applicability of the WR technology can be further explored.

7.1 Summary of findings

The journey of this thesis, from the definition of the general concepts of time and frequency transfer to the broad and tangible applications thereof, as demonstrated by the redundancy protocol implementation for Smart grid or the pilot of a time service for the financial industry, has been both a rewarding challenge and an immense learning experience for me. I would like to start the conclusion chapter by acknowledging everybody who has supported me in any way during these last few years, and I would like to extend my deepest gratitude to the WR team at CERN and the whole community of the OHWR, whose pioneering work laid the foundation for my research.

This thesis presents a comprehensive study of the time and frequency transfer methods over long distance optical fiber links, with a focus on resilience and redundancy features as a vehicle to extend the adoption of WR in applications involving

critical infrastructure management. The work presented here builds upon a foundation of understanding timing concepts, methodology, and the technical details of the best-performing time transfer mechanisms over the optical medium. The WR protocol is emphasized as the preferred choice for the advancements made in this work, with a detailed examination of its implementation, performance, and challenges in long-distance deployments. The state-of-the-art study is complemented by an analysis of synchronization requirements in key sectors such as technology finance, telecommunications, and the power grid.

The research methodology, hardware platform, tools, and laboratory equipment were presented in Chapter 3. This sets the stage for the development and validation of the proposed solutions. The three subsequent chapters build on this foundation to present a thorough analysis of the research problems and potential approaches to address them. These three chapters deal with distinctively different areas of the challenges. Chapter 4 is related to improvements within the WR devices (hardware and code), while Chapter 5 is related to all the phenomena manifested outside the device, i.e., the optical signal and network considerations. Finally, Chapter 6 takes the lessons learned and the mechanisms put in practice under laboratory conditions, and takes them into operation using commercial equipment for a pilot TaaS deployment.

In Chapter 4, we researched the factors that limited the applicability of holdover solutions in the standard WRS, for which we conducted a deeper analysis of the phase syntonization and synchronization of the WR protocol. A new hardware platform, in which the author had a significant role in its development, was designed at Seven Solutions to address the hardware limitations of the standard design. Notably, this new platform, referred to as WRS-LJ, significantly improves the phase noise of the Grandmaster by 9 times (from 8.4 ps RMS jitter to 0.9 ps RMS), and by almost 7 times (from 10.0 ps RMS jitter to 1.5 ps RMS) in the first slave. This improvement, among other benefits, allows for an increased number of layered (or daisy-chained) devices to be deployed without increasing the noise with respect to the previous design.

The WRS-LJ served as the experimental platform for the development of a holdover solution. The challenge of minimizing the time error was approached using a quadratic model. While the precision in the estimation of the parameters was constrained by hardware limitations, the results obtained were still estimated to be at least 200 times better than the previous holdover capabilities of the WRS as determined in [102]. However, the holdover capabilities of the standard WRS default oscillator were so limited that they barely factor into the comparison. This level of performance is so far from the requirements of most applications that it is challenging

to consider it a viable holdover solution. In stark contrast, the current experimental setup, under the best configurations found (Abracon AOCJY 10 Mhz, undersampling and reproducing the same corrections done prior to the launch of holdover mode), demonstrates a time error of less than 1 ns after 1 minute in holdover mode using a relatively affordable OCXO. The importance of this improvement is key as it enables the device to perform a switchover, even from scratch, without a significant impact on the timing signals. This capability is crucial for maintaining the integrity and reliability of time and frequency transfer systems, particularly in scenarios where network disruptions or reconfigurations occur.

In terms of reliability, the implementation of a seamless redundancy protocol for ring topologies suitable for data and WR timing using standard WRSs demonstrated the feasibility of a time reference switchover with a sub-nanosecond impact on time accuracy. This included the implementation of a firmware capable of handling the redundancy features that ensure the reception of data and timing even in the event of a single network failure. The results showed that the switchover can be effectively performed in the worst-case scenario of a ring of six WRS devices without enhanced holdover capabilities. In the test scenario, the entire ring was able to switch to the backup reference in less than one microsecond since a link failure was detected. The timing impact of the switchover in the ring topology was limited to 170 ps with respect to its pre-switchover offset for the worst performing WRS in the ring. Remarkably, this experiment was performed using standard WRS without enhanced holdover capabilities.

Regarding data transmission, the redundant capabilities of the HSR protocol guarantee the delivery of data and control messages in the event of a link failure, with zero data loss. The switching mechanisms of the WRS have been optimized for operation in ring networks, halving the latency of switching in these topologies from 2.2 to 1.0 microseconds. The bandwidth results in the redundant scenario prove that the redundant data delivery can be sustained for bandwidths up to 682 Mb/s.

In Chapter 5, we performed an extensive review of the previous efforts to disseminate WR over long distances, and the latest developments to solve limitations on reach and asymmetry determination [75], [129]. Since the analyzed implementations are limited to private networks or laboratory conditions, we embarked on a full-scale survey of the phenomena that impede transmission in fiber netwokrs, with focus on attenuation and chromatic dispersion as factors that are present in substantial scale even in a well-engineered optical network. Results show that DWDM equipment is encouraged even for dark fiber timing networks due to the higher stability of the

DWDM transceiver lasers. The analysis of this factor was done exclusively by reviewing the literature due to the lack of specialized equipment, and a 7.7-fold reduction in the SFP uncertainty budget was found for 80 km links. To complement this study, the mechanisms and components that are deployed in optical networks to mitigate undesired effects and successfully transmit and route the optical signals were also analyzed. The result of this review elucidated the important of precise network modelling and strict network management control in long-distance deployments. For cases when this is not possible, an estimation of the impact of every possible element of a long distance testbench was modeled in terms of delay, asymmetry, and the changes induced in these two variables by thermal variability. From the analyzed elements of the testbench, thermal coefficients for the RTT delay range from negligible for discrete elements such as a DWDM multiplexer, to 46 ps $km^{-1}K^{-1}$ in the case of DCMs based on DCFs. The contribution of the end nodes is estimated as 14 ps K^{-1}. Propitiously, network equipment is usually located in climate-controlled rooms where the delay and asymmetry contributions can be considered static. Following this, our findings also underscore the possibility of using alternate calibration techniques to compensate for the unknown asymmetries in networks beyond the scope of our control with limited impact on timing performance.

Finally, in Chapter 6, the knowledge and insights collected from the previous chapters were applied to design, deploy and evaluate a pilot time service for the Spanish stock market operator. The pilot scenario successfully implemented a solution that disseminated time from redundant time servers to the remote datacenter of the customer with optical fiber using WR across a shared public network. This system provides the customer with resiliency features through extensive holdover capabilities at the time server and at the customer side, and also through redundancy in the time server hardware and dissemination. The service meets the current requirements mandated by regulation, and is well-equipped to handle future demands: during normal operation the traceability to UTC is guaranteed to a maximum time error of 15 ns. It is important to note that this 15 ns limitation is primarily due to the GNSS receiver used for calibration. By employing a multifrequency GNSS receiver, this limitation could be overcome, further improving the system's accuracy. The resiliency mechanisms ensure that the system can endure a loss of GNSS reference for weeks while keeping the time error below the requirements of current regulation. If the customer-side equipment loses the visibility of the time server, data demonstrates that the customer-end nodes are able to operate autonomously for up to 24 hours before the time error of the time signal reaches one microsecond.

It is important to highlight that all the development and mechanisms in this research are built on devices that make use of standard Ethernet network equipment. This means that the WR protocol can be used over existing telecom networks, maintaining the mentioned timing performance without needing a separate, dedicated infrastructure or the deployment of non-standard industry equipment. This major advantage of WR provides a cost-effective way to integrate it into current systems. In addition, the fact that the principles of operation of WR are part of the High Accuracy profile of the IEEE 1588-2019 standard is worth noting. This alignment with a well-known industry standard does not just attest to the robustness and relevance of WR but also forecasts a future where numerous vendors will offer interoperable products. A wide vendor support guarantees long-term accessibility, an aspect invaluable to system designers that need to ensure project longevity.

In conclusion, this thesis makes a significant contribution to the field of dissemination of highly-accurate time signals over optical links spanning long distances both in private and public networks. The work presented here has demonstrated the potential for advancements in resilience, redundancy, and performance required for the extensive use of WR in new use cases in the metropolitan and long-haul scale networks. Furthermore, the research has highlited areas for future research to address limitations and enhance the robustness of WR time transfer systems.

7.2 Main contributions

As we are in the middle of the conclusion chapter of this thesis, it is the perfect occasion to revisit our starting points—the original objectives that shaped this research. It is important to see how these align with the contributions that have been achieved. This is a chance to measure our progress and determine if the expectations we set for ourselves have been met.

1. We conducted a review of resiliency and redundancy mechanisms regarding time dissemination, with a particular focus on critical infrastructure. Our analysis centered on the implications and specific needs of this industry segment. We then expanded our analysis to include the specific conditions and challenges of long-distance WR links. This contribution is closely related to Objectives 1 and 3.

2. The noise response and clock circuitry of the WRS was improved in a new hardware design. Although not directly stated in the objective, the reduction of

noise and improvement in stability (summarized in section 4.3 is a requirement for achieving Objective 4. The adaptation of the HDL design for this new hardware to include new experimental oscillators is also an essential prerequisite for the implementation of any holdover solution (Objective 4). The results of this item have been published in the OHWR[1] as an open-source product, both hardware and firmware.

3. The main contributors to the degradation of the holdover time error were analyzed and modeled. An experimental implementation that estimates the initial fractional frequency offset was developed, improving its performance with any previous open source WR implementation and achieving at least 60 seconds of holdover time while maintaining the error below 1 ns. This is one of the goals of Objective 4. This contribution was done in the framework of the WRITE project, and, accordingly, the design and code is also available in the open source community.

4. In collaboration with other researchers, the implementation of a seamless redundant ring protocol for both time and data enabled the use of experimental seamless switchover was achieved. This contribution completes the achievement of Objectives 3 and 4.

5. In this same collaboration, we demonstrated the first use of WR in ring topologies, implementing novel clock roles for a WR network such as Transparent and Hybrid Clocks. This aligns with objectives 3 and 4.

6. The applicability of the redundant ring solution was demonstrated in a Smart grid testbench and published in a peer-reviewed scientific journal [1], satisfying Objective 5 for the Smart grid case.

7. A long distance testbench was deployed in laboratory controlled conditions. From it, a model of additive delay, time error, and temperature dependence of its components was derived. This translates directly to Objective 2.

8. A significant milestone was achieved with the deployment of a real-world use case in the finance sector, marking the first documented instance of a TaaS being deployed over a public network using WR [15]. This pioneering application served as a verification of the principles and solutions analyzed and implemented in previous chapters. The robustness of the double time server and the

[1] `https://ohwr.org/project/wrs-lj-hw/wikis/home`

resiliency features in the customer site node were thoroughly tested against alternative time sources and validated under real-world conditions. Importantly, these features were validated without any impact on the customer's perception of the time service. This achievement, which directly maps to Objective 5, highlights the potential of our work and serves as a precedent for future deployments of TaaS over public networks using WR.

When the objectives of this thesis were defined, we proposed a solution to the challenge of achieving high-accuracy long distance time dissemination over optical fiber links, aiming to exceed the performance of traditional satellite methods and network protocols like PTP or NTP. Our hypothesis was that WR links could be deployed with built-in fault-tolerance and resiliency, achieving accuracy and precision equivalent or superior to GNSS systems, all while utilizing existing telecommunication fiber networks. The multiple contributions enumerated in this section have affirmed the validity of this hypothesis.

7.3 Publications

This section presents a series of publications generated over the course of this doctoral research. The contributions listed below have undergone peer review and have been published in internationally recognized, high-impact academic journals and conferences.

7.3.1 Journal contributions

- J. L. Gutiérrez-Rivas, J. López-Jiménez, E. Ros and J. Díaz, "White Rabbit HSR: A Seamless Subnanosecond Redundant Timing System With Low-Latency Data Capabilities for the Smart Grid" in IEEE Transactions on Industrial Informatics, vol. 14, no. 8, pp. 3486-3494, Aug. 2018, doi: 10.1109/TII.2017.2779240.

- F. Ramos, J. L. Gutiérrez-Rivas, J. López-Jiménez, B. Caracuel and J. Díaz, "Accurate Timing Networks for Dependable Smart Grid Applications," in IEEE Transactions on Industrial Informatics, vol. 14, no. 5, pp. 2076-2084, May 2018, doi: 10.1109/TII.2017.2787145.

- F. Girela-López, J. López-Jiménez, M. Jiménez-López, R. Rodríguez, E. Ros and J. Díaz, "IEEE 1588 High Accuracy Default Profile: Applications and Chal-

lenges," in IEEE Access, vol. 8, pp. 45211-45220, 2020, doi: 10.1109/ACCESS.2020.2978337.

- J. Lopez-Jimenez, J. L. Gutierrez-Rivas, E. Marin-Lopez, M. Rodriguez-Alvarez and J. Diaz, "Time as a Service Based on White Rabbit for Finance Applications," in IEEE Communications Magazine, vol. 58, no. 4, pp. 60-66, April 2020, doi: 10.1109/MCOM.001.1900602.

- M. Jiménez-López, F. Girela-López, J. López-Jiménez, E. Marín-López, R. Rodríguez and J. Díaz, "10 Gigabit White Rabbit: Sub-Nanosecond Timing and Data Distribution," in IEEE Access, vol. 8, pp. 92999-93010, 2020, doi: 10.1109/ACCESS.2020.2995179.

- In preparation: a forthcoming paper, to be submitted to *IEEE Transactions on Instrumentation and Measurement*, will provide and extend a detailed account of our holdover implementation results. The specific title and list of authors are still under consideration at the time of writing this thesis.

7.3.2 Conferences

- J. López-Jiménez, J. Díaz and M. Rodríguez-Álvarez, "Impact of network component temperature variation on long haul White Rabbit links," 2018 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS), Geneva, Switzerland, 2018, pp. 1-6, doi: 10.1109/ISPCS.2018.8543085.

- E. F. Dierikx, Y. Xie, A. Savencu, J. Lopez and J. L. G. Rivas, "White Rabbit Multi-Point Time Distribution Network," 2021 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS), Gainesville, FL, USA, 2021, pp. 1-4, doi: 10.1109/EFTF/IFCS52194.2021.9604343.

7.4 Future work

The findings from this study open several ways for future research and development. Probably the most evident path is further improving the robustness of WR by combining the switchover and holdover implementations that have been proposed in this work. Unfortunately, due to external constraints, the development of the holdover implementation and the redundant network with seamless switchover capabilities

were done in different points in time, and using different versions of the hardware platform, therefore extending the estimated time for porting and adapting a solution implementing the combined features. However, as a developer within Safran Electronics and Defense, the author is in a unique position to push for the integration of these developed solutions into other commercial products and platforms. The ability to simultaneously track two WR references while having the improved holdover capabilities available provides a great starting point for developing a powerful and versatile failover/switchover mechanism with additional monitoring features.

Also related to holdover, there are improvements that can be done to improve the behavior of the estimators of the model. First, a more sophisticated holdover mechanism could be developed if the internal parameters of the OCXO control were visible to the WR system. If this were not possible, another approach would be to find different hardware components that would allow the development of our own, inhouse adaptive OCXO algorithm. The topic of early failure detection as a trigger condition for holdover mode is a field in which further efforts will result in tangible improvements, although this might need additional hardware or communication mechanisms to achieve a sizable population for majority voting.

In relation to the modeling and extension of the WR link model to more diverse and experimental scenarios, the use of WR over photonic bandgap fibers may be a topic worth studying, since these are being considered in certain experimental scenarios as a way to achieve lower link latencies. Photonic bandgap (hollow core) fibers show transmission speed close to the vacuum, but they are not cost effective at the present moment and their reach is limited due to high attenuation.

Another promising avenue for future research lies in the exploration of alternative time sources, particularly those based on Low Earth Orbit (LEO) satellites. Companies like Satelles[2] or Xona Space Systems[3] provide secure time and location services that could greatly enhance the resiliency of time systems for critical infrastructure. This could provide a robust and secure alternative or supplement to traditional GNSS-based time sources, offering potential benefits in terms of signal availability, global coverage, and resistance to jamming and spoofing. Investigating the integration of these emerging time sources with WR technology could open up new possibilities for enhancing the reliability and security of time services. This could be particularly valuable in critical infrastructure and financial applications where high levels of resiliency

[2] `https://safran-navigation-timing.com/testing-confirms-resilient-timing-success/`

[3] `https://www.xonaspace.com/pulsar`

and security are paramount.

Following the same line of more diverse scenarios for the dissemination of time based on WR, alternate media such as millimeter wave or free-space lasers are topics that are receiving attention for niche applications. The applicability of WR and the associated challenges are topics of interest and tangentially related to this work, but the fundamental differences with the guided optical medium favour a more bottom-up approach.

Apéndices: Introducción y Conclusiones en español

Apéndice A

Introducción

Este primer capítulo tiene el fin de presentar una visión general del contexto en el que se desarrolla el estudio presentado en esta memoria. Inicialmente en la sección A.1 se introduce brevemente los antecedentes y el contexto en el que se enmarca la tesis, se proporcionan nociones a alto nivel que facilite comprender tanto las dificultades existentes a día de hoy en el área de trabajo competente como las soluciones/mejoras propuestas e implementadas. La sección A.2 recoge formalmente los objetivos científicos que se ha propuesto conseguir, y a continuación en la sección A.3 se complementa esta introducción con un resumen de los proyectos nacionales e internacionales en los que se ha colaborado durante el desarrollo de este trabajo. Por último, se cierra esta introducción con una descripción esquemática del contenido que puede encontrarse en el resto de capítulos.

A.1 Motivación

La sincronización de tiempo precisa se ha convertido en una necesidad de relevancia con áreas de aplicación cada vez más amplias. Los niveles de precisión y exactitud que hace unos años eran requisitos propios de las infraestructuras científicas más exigentes del mundo se han ido abriendo paso progresivamente hacia nuevas aplicaciones, y a día de hoy la necesidad de conseguir sincronización trazable a unos pocos nanosegundos se puede encontrar en segmentos de la economía tan diversos como las redes de telecomunicaciones, los mercados financieros u otras infraestructuras críticas como las redes eléctricas inteligentes (Smart Grid) [1], [2].

Dentro de esta diversidad en campos de aplicación, se puede apreciar un punto en común en todos ellos: la extensión de estos sistemas a lo largo de áreas geográficamente

amplias o, en su defecto, la necesidad de interoperar con otros sistemas distantes. A menudo, los sistemas distribuidos en áreas extensas tienen la necesidad de que sus componentes operen de forma coordinada, para lo cual necesitan tener una noción de tiempo común.

Desde un punto de vista histórico, suele decirse que la idea de tiempo "moderna", esto es, de que en un área geográfica exista una única referencia maestra de tiempo a la que se ajustan todos los relojes (en contraposición al tiempo local solar que existía anteriormente) es una creación de la era industrial que surgió como respuesta a la necesidad de que los trenes dentro de una misma red de ferrocarriles compartan una hora común. Así, en la década de 1840 surgió en Inglaterra el concepto del *Railway Time* o *tiempo ferroviario* [3].Con el paso del tiempo el uso de la hora solar de cada localidad cayó en desuso debido a las ventajas que entraña el uso del tiempo estandarizado. En 1884 el tiempo medio de Greenwich o *Greenwich Mean Time* (GMT) fue elegido como estándar de tiempo universal en la Conferencia Internacional del Meridiano y se mantuvo como tal hasta 1972, cuando el estándar de Tiempo Universal Coordinado (conocido por su pseudoacrónimo UTC) tomó el relevo y el mundo dejó de lado el período de rotación de la Tierra como referencia para pasar a usar relojes atómicos.

Del mismo modo que a lo largo de estos dos últimos siglos la Física y la Metrología han progresado enormemente con el objetivo de conseguir referencias de tiempo cada vez más precisas y estables, el trabajo de multitud de científicos e ingenieros, incluyendo el presente estudio, ha consistido en desarrollar métodos y técnicas para diseminar el estándar de tiempo de uso común en cada época a las aplicaciones que lo requieren, maximizando las métricas técnicamente posibles de acuerdo con cada período histório y solución (precisión, extensión, fiabilidad, etc.).

Volviendo al presente, a lo largo de las últimás décadas han proliferado las infraestructuras de redes de comunicaciones, cuya capacidad en cuanto a ancho de banda, número de nodos a los que presta servicio y ubicuidad no ha dejado de crecer hasta el día de hoy. Estas redes son una tecnología instrumental para el desarrollo de cuantiosos sistemas distribuidos inteligentes. De la mano del crecimiento de la cantidad y complejidad de estos sistemas deriva la magnitud de los datos que atraviesan estas redes en las que, por diseño y debido a la necesidad de mantener compatibilidad hacia atrás, la importancia que otorgan a la sincronización o a la fiabilidad es bastante baja.

De hecho, en las redes Ethernet comúnmente en uso en la actualidad no existe garantía de que los datos enviados a través de la red llegarán a su destinatario en el mismo orden en el que se enviaron, ni se garantiza una latencia máxima. Aún más, ni siquiera se garantiza la entrega de datos al receptor [4]. Esto es lo que se conoce como

mecanismo de entrega *best effort* o de mejor esfuerzo.

Esto es una aparente contradicción ya que el funcionamiento normal de los sistemas distribuidos implica con frecuencia correlacionar datos de plataformas de sensores físicamente distantes, lanzar la puesta en marcha de un proceso en un intervalo de tiempo acotado, o simplemente monitorizar de forma precisa el estado global del sistema para poder realizar un análisis forense que recomponga ordenadamente todos los pasos recogidos. Para llevar a cabo tales tareas resulta imprescindible cumplir con el requisito de compartir una base de tiempos común. De lo contrario, los componentes del sistema no serían capaces de etiquetar eventos temporalmente, ni de determinar el tiempo que tarda un mensaje en llegar a su receptor. Sin estas premisas, es imposible garantizar la causalidad de los eventos que ocurren en el sistema.

Por lo tanto, los arquitectos de este tipo de sistemas deben incluir una solución específica para mantener la sincronización a lo largo de la red. Los dos métodos más comunes en la actualidad para mantener la sincronización de los equipos son las técnicas basadas en sistemas globales de navegación por satélite (GNSS, de *Global Navigation Satellite Systems*) o las técnicas basadas en el intercambio de paquetes en las previamente citadas redes *best effort*.

La primera de estas tecnologías es, en sí misma, un sistema distribuido: los sistemas globales de navegación por satélite como GPS (del inglés *Global Positioning System* o sistema de posicionamiento global) o GLONASS (de la transliteración del ruso de *Global'naya Navigatsionnaya Sputnikovaya Sistema* o sistema global de navegación por satélite). Estos sistemas son empleados para distribuir no solamente la localización del usuario final, sino también el tiempo. Esta fuente de tiempo ofrece importantes ventajas: su cobertura es global, el grado de precisión es relativamente alto para la inversión requerida y el mantenimiento de la fuente y su trazabilidad están garantizados por fuertes instituciones públicas. Sin embargo, las señales de GNSS, que son recibidas por el usuario final como una débil señal de radio que fue transmitida desde una constelación de satélites a varias decenas de miles de kilómetros de distancia, son vulnerables a toda una serie de posibles problemas o ataques deliberados. La señal es susceptible a turbulencias atmosféricas, debe ser obtenida a través de una antena ubicada en el exterior, y es vulnerable a interferencias o ataques de suplantación.[5]

De este punto débil del tiempo distribuido por GNSS, unido a la criticalidad de muchas de las infraestructuras que confían en él para su funcionamiento, surge la necesidad de los gestores de estos sistemas de provisionarse o complementar sus fuentes de tiempo con otros distribuidos a través de otros medios, y/o de proveerse de equipos que sean tolerantes a fallos y puedan gestionar varias fuentes simultaneamente.

De entre las fuentes y métodos de distribución alternativos a GNSS destacan aquellos que transmiten la información a través de las redes de comunicaciones. Dentro de esta familia existe una variedad de técnicas y protocolos que, dependiendo de su complejidad y de los requisitos para su implementación, son capaces de dar un nivel de exactitud y precisión variable. El rango de error que puede conseguirse con estos protocolos puede variar desde unos pocos milisegundos de error (como consiguen las implementaciones de software que existen en muchos dispositivos electrónicos de consumo) hasta las pocas decenas de picosegundos que pueden conseguirse con otros protocolos de intercambio de paquetes asistidas con hardware específico. White Rabbit, una extensión de PTPv2[6] que nació para cubrir la necesidad de una infraestructura científica de primer nivel (el Centro Europeo de Investigación Nuclear o CERN), es un ejemplo de estos últimos [7].

Sin embargo, existen dificultades para conseguir el mismo nivel de precisión cuando el tiempo se ha de distribuir a largas distancias por dos motivos fundamentales: por un lado, la naturaleza del medio de transmisión hace que no se puedan alcanzar grandes distancias usando el mismo modelo de enlace que en una red a escala local dado que lo imposibilitan fenómenos físicos como la atenuación y la dispersión de la señal. Por otro lado, al emplear otras técnicas y otros dispositivos para habilitar la transmisión de la señal hasta distancias más lejanas, se introducen nuevos elementos en el modelo del enlace que deben ser tenidos en cuenta. Existe una dificultad añadida a las anteriores si la transmisión se efectúa a través de una red de transporte gestionada por tercero (i.e. un proveedor de servicio o *ISP*) que puede alterar la red en cualquier momento. Si bien estas modificaciones suelen ser transparentes en cuanto al servicio de datos, que no se ve afectado, el impacto para un sistema de distribución de tiempo puede suponer la introducción de fallos en las señales de sincronización y/o la interrupción del servicio durante un intervalo de tiempo.

Con respecto a la robustez del sistema de sincronización, cabe considerar que un posible fallo en el sistema de distribución de tiempo de las infraestructuras y los casos de uso antes mencionados pueden dar lugar a mal funcionamiento o paradas completas de los sistemas y, en última instancia, a grandes pérdidas económicas y responsabilidades legales. Para mitigar este riesgo existen diversas aproximaciones que aumentan la resiliencia del sistema frente a fallos. Frente a la posibilidad de que una única fuente de tiempo se vea comprometida y se convierta en un punto único de fallo, se sugiere la implementación de técnicas de redundancia que permitan conmutar entre distintas fuentes de tiempo (*switchover* y *failover*). Por otro lado, en caso de fallo repentino del dispositivo o el sistema que se emplea como referencia de tiempo, se pueden imple-

mentar técnicas con hardware específico que garanticen que el error del oscilador local de un equipo no exceda de un determinado umbral (*holdover*). Esto es necesario para mantener la integridad la señal de tiempo durante el intervalo de conmutación a una fuente alternativa o, en caso de que no exista, para prolongar el estado de funcionamiento de forma autónoma, sin recibir correcciones de tiempo del exterior, durante el mayor tiempo posible hasta que el fallo se haya solventado.

Un sistema de tiempo que tenga en cuenta todos los conceptos y las dificultades que se han descrito como parte de esta sección de visión general se puede prestar como un servicio. Las compañías que lo requieren tienen necesidad de un sistema de sincronización complejo pero no posean la infraestructura o la experiencia necesarias para garantizar una señal que cumpla con todos los requisitos de precisión, exactitud, trazabilidad y alta disponibilidad. Este tipo de soluciones se conocen como *tiempo como un servicio* o *Time as a Service* ("TaaS", por sus siglas en inglés). Bajo este concepto, el tiempo es un producto suministrado por un tercero que se encarga de poner los medios para demostrar la integridad de la señal de tiempo en todo momento. Las soluciones TaaS son frecuentes en el campo de los servicios financieros. Dentro del marco de este trabajo, se tuvo la oportunidad de colaborar en el despliegue y mantenimiento de un proyecto piloto de un sistema de distribución de tiempo como un servicio "TaaS" basado en la generación de tiempo en un punto geográficamente distante al del cliente final, la distribución de una referencia de tiempo a través de la red pública de un ISP manteniendo la trazabilidad a UTC, y mecanismos de redundancia y resiliencia tanto en la fuente como en los nodos finales.

Para concluir esta introducción, cabe destacar que la naturaleza de este trabajo se centra en la extensión de la tecnología White Rabbit sobre redes de datos públicas, una tarea desafiante dada la variabilidad y la falta de control inherente a estas infraestructuras. A diferencia de otros trabajos presentes en la literatura, que se basan en tener un control completo de la red, este estudio se adentra en la complejidad de trabajar con redes públicas y cómo éstas pueden ser utilizadas para proporcionar sincronización de tiempo de alta precisión y robusta. A través de este enfoque, buscamos proporcionar una contribución única y valiosa a este campo de estudio que se define en los objetivos descritos en la siguiente sección.

A.2 Objetivos

Esta tesis aborda el desafío de la diseminación de tiempo a larga distancia con alta precisión a través de enlaces de fibra óptica, un servicio que tradicionalmente depen-

de de métodos satelitales o protocolos de red como PTP o Network Time Protocol (NTP), que no pueden alcanzar el mismo nivel de rendimiento. **Nuestra hipótesis se resume en que las implementaciones de enlaces WR pueden desplegarse utilizando infraestructura de redes de fibra de telecomunicaciones existentes, implementando mecanismos de tolerancia a fallos e incorporando características de resiliencia tanto en el nodo como en la red, logrando una precisión y exactitud igual o superior a la de los sistemas GNSS actuales**. Este enfoque ofrece una solución rentable que no depende de la implementación de infraestructura de red dedicada, es rápido de desplegar y alcanza el alto nivel de capilaridad que las redes de fibra de telecomunicaciones pueden proporcionar, así como cumple con los exigentes requisitos de múltiples industrias de infraestructura crítica como *Smart grid* o las financieras tecnológicas (*fintech*).

Nuestro objetivo es desarrollar métodos para desplegar enlaces WR de alta precisión a través de redes de acceso público, tratados como tráfico regular, mientras se mantiene la integridad de la señal óptica. Dada la naturaleza dinámica de las redes públicas, las características mejoradas de resiliencia y redundancia son necesarias para que los equipos WR resistan condiciones cambiantes y mantengan la integridad de la señal y la disponibilidad del servicio. Las líneas de trabajo seguidas para alcanzar estos objetivos están representadas gráficamente en la Figura 1.1.

Los objetivos científicos se especifican en la siguiente lista, que reproduce los mismos puntos que se definieron en el correspondiente Plan de Investigación cuando se inició este proyecto:

1. Estudio de los elementos ópticos utilizados en redes de larga distancia, así como de los problemas de distribución de la señal de sincronización en el medio óptico. Esto incluirá la investigación de los efectos de atenuación, dispersivos y cualesquiera otros que puedan afectar a la sincronización de tiempo en redes Ethernet ópticas.

2. Desarrollo de modelos de componentes y de red que permitan el despliegue del servicio de sincronización en redes de fibra óptica de larga distancia. Este modelo permitirá integrar información de componentes y equipos proporcionados u obtenidos en laboratorio para determinar su efecto en el sistema final.

3. Análisis de las topologías de red y evaluación de mecanismos de fiabilidad en caso de fallo parcial o total de la red. Este análisis se realizará para distintas aplicaciones, con especial énfasis en las telecomunicaciones y las infraestructuras críticas, y se diseñarán estrategias de mitigación de riesgos.

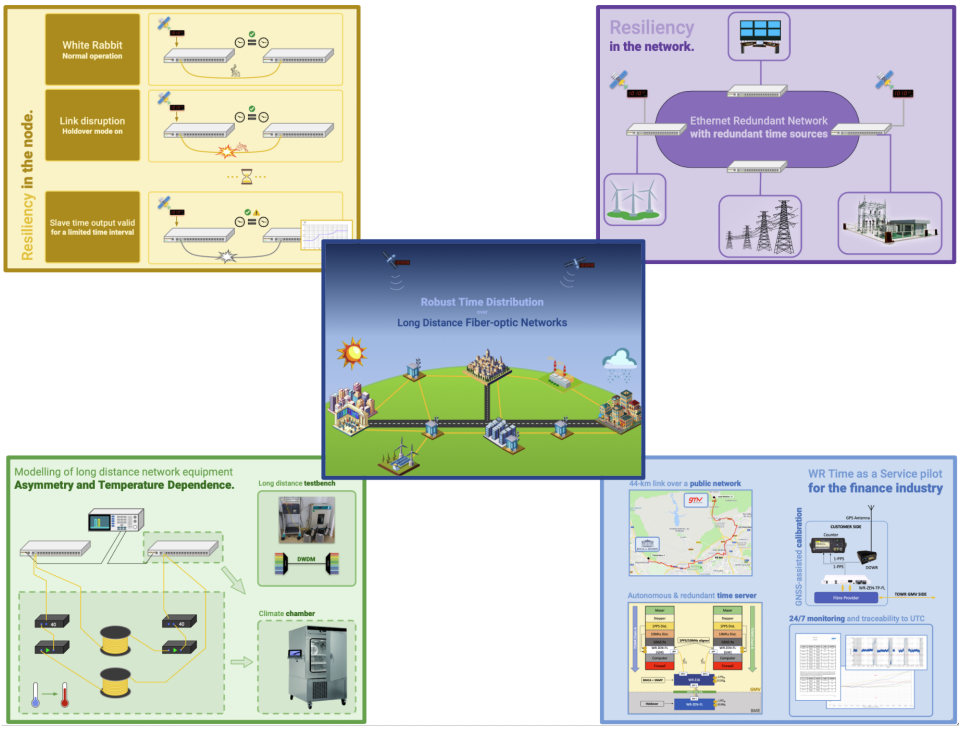

Figura A.1: Resumen gráfico del contenido de este trabajo.

4. Diseño e implementación de mecanismos de selección de topologías y funcionalidades de switchover/holdover para redes de larga distancia. Este objetivo incluye el desarrollo de software y descripción de hardware HDL para *matrices de puertas lógicas programables en campo* (FPGAs) de modo que se dote a los sistemas WR de estas funcionalidades aplicables a redes de larga distancia.

5. Aplicación y casos de uso. Validación experimental de las soluciones desarrolladas y verificación empírica con otros métodos de distribución de tiempo y frecuencia. Esto incluirá la definición de métricas de éxito, la realización de pruebas de campo y la recopilación y análisis de datos para demostrar la eficacia de las soluciones propuestas.

A.3 Marco de trabajo

Este trabajo de investigación se ha llevado a cabo en el marco de varios proyectos de investigación de ámbito nacional y europeo. En este apartado se proporciona una breve descripción de estos proyectos CLONETS, AMIGA, TOWR y WRITE. Además, esta tesis se ha financiado parcialmente por el Ministerio de Ciencia e In-

novación del Gobierno de España (MICIN) como parte del programa de Ayudas para contratos para la formación de investigadores en empresas (Doctorados Industriales) en su convocatoria 2017 dentro de la compañía Safran Electronics and Defense Spain (entonces conocida como Seven Solutions).

A.3.1 CLONETS y CLONETS-DS

El proyecto CLONETS (CLOck NETwork Services, con ID 731107 de la convocatoria H2020-INFRAINNOV-2016-2017 de la Comisión Europea) tiene por objeto preparar la transferencia de la nueva generación de tecnología a la industria y reforzar la coordinación entre las infraestructuras de investigación y las redes de telecomunicaciones de investigación y educación, con el fin de preparar el despliegue de una red sostenible y paneuropea que provea de servicios de tiempo de alto rendimiento a las infraestructuras científicas europeas [8].En este proyecto participa un consorcio formado entre institutos metrológicos europeos, instituciones académicas, y actores de la industria involucrados en la temática. Esta red se diseñará para que sea compatible con la distribución de tiempo y frecuencia a una multitud de agentes que necesitan servicios de menor rendimiento que la propia de las infraestructuras metrológicas, respondiendo a las necesidades creadas por desarrollos como la computación en la nube, el internet de las cosas y la industria 4.0.

El proyecto CLONETS tiene una serie de objetivos de alto nivel mediante los cuales se obtiene información de las infraestructuras de investigacion nacionales y sus necesidades de tecnologías de tiempo y frecuencia y se analizan posibles aplicaciones de una red de distribución de tiempo y frecuencia para otras áreas de la industria, enfocado en la distribución a través de redes de fibra óptica. También tiene como objetivo definir hojas de ruta y estrategias que conduzcan a la definición y construcción de la red paneuropea de distribución de tiempo y frecuencia.

Existe una continuación de este proyecto, llamado CLONETS-DS (con ID 951886 de la convocatoria H2020-INFRADEV-2018-2020 de la Comisión Europea [9]), enfocado en la elaboración de un estudio de diseño que recoja las necesidades de la comunidad científica y defina la arquitectura de la red.

A.3.2 TOWR

TOWR es un proyecto industrial cofinanciado por el programa NAVISP (Navigation Innovation and Support Programme) de la Agencia Espacial Europea con las compañías GMV Aerospace & Defence y Seven Solutions [10] e identificado con el

código de actividad NAVISP-El2-033. El objetivo del proyecto es el desarrollo de un servicio de distribución de tiempo mediante fibra óptica en el Área Metropolitana de Madrid. Esto se consigue mediante la generación de una base de tiempos propia a partir de osciladores muy estables, transferencia de tiempo GNSS para alinear los osciladores de referencia, y el uso de tecnología WR para distribuir la referencia de tiempo hasta el cliente final a través de una red de comunicaciones pública. El proyecto pone el foco en los mecanismos de resiliencia y en la trazabilidad de la señal para cumplir los requisitos de sectores como el financiero. El despliegue piloto del proyecto se llevó a término con el operador español de los mercados de valores, BME.

Revisando la sección A.2 puede verse que los objetivos de este proyecto se encuentran alineados en gran medida con los de esta tesis. Muchos de los resultados obtenidos se pudieron llevar a la práctica en el despliegue piloto y se describirán en detalle en el capítulo 6.

A.3.3 AMIGA

AMIGA6

AMIGA6 es un proyecto AMIGA (Análisis del Medio Interestelar de las Galaxias Aisladas) del Ministerio de Ciencia e Innovación con ID AYA2015-65973-C3-1-R liderado por el Instituto Andaluz de Astrofísica [11].

AMIGA6 se basa en los resultados de proyectos anteriores de AMIGA y forma parte del trabajo preparatorio que está realizando AMIGA para la explotación científica del radiotelescopio SKA [12]. Los ciclos de vida HI (nubes de hidrógeno atómico neutro), tanto en galaxias aisladas como en grupos densos, siguen siendo poco conocidos ya que sus columnas de baja densidad sólo son alcanzables por SKA. Los principales objetivos de AMIGA6 son: refinar los modelos de acreción de gas frío utilizando galaxias aisladas, y analizar el papel de la eliminación de mareas HI en la SF (formación de estrellas) suprimida en los grupos compactos de Hickson. Como preparación para el reto que supondrá la explotación de los datos de SKA, AMIGA6 complementa la ciencia fundamental con la investigación aplicada en 3 paquetes de trabajo, contribuyendo así a los consorcios de Big-Data SKA. AMIGA6 está liderado por la coordinadora de la participación española en el SKA, y reúne a todos los grupos españoles involucrados en el flujo de datos del SKA: Signal and Data Transport (SaDT), Central Signal Processor (CSP), y Science Data Processor (SDP). La relación con esta tesis se enmarca dentro del grupo de Transporte de Datos y Señales. En colaboración con la Universidad de Granada, se trabajó en una evaluación de las tecnologías de transferencia de tiempo y

las dificultades propias de la arquitectura de red del proyecto de radiotelescopio SKA.

AMIGA7

El subproyecto de AMIGA7 Procesado de Datos en Hardware (ID RTI2018-096228-B-C32 del Ministerio de Ciencia e Innovación) [13] investiga técnicas de optimización y síntesis de circuitos que implementan en hardware reconfigurable tanto algoritmos de procesado de señal del CSP como algoritmos científicos del SDP y los Centros Regionales del radio interferómetro de SKA (Square Kilometre Array). AMIGA7 abarca un período clave para SKA, tras el final de los diseños de los consorcios y el éxito de sus Critical Design Reviews, hasta el inicio de la construcción.

Los dispositivos reconfigurables están firmemente asentados en la arquitectura de los módulos del CSP de SKA, donde se utilizan profusamente. Los circuitos DSP suelen permitir el uso de formatos de coma fija que presentan mayores prestaciones y menor coste y consumo que sus equivalentes en coma flotante. El proyecto plantea la aplicación de técnicas de análisis de la sensibilidad de las señales según sus dependencias para optimizar métodos automáticos. La aplicación de estos métodos permitirá nuevos niveles de optimización de los módulos de CSP en su evolución de cara a la construcción de SKA.

Por otra parte, el posible uso de aceleradores de alta eficiencia energética basados en FPGAs para el cálculo científico en los Centros Regionales de SKA resulta de gran interés para reducir su alto coste energético. En este contexto, el proyecto continúa el desarrollo del entorno Witelo para la síntesis y generación automática de VHDL (Lenguaje de Descripción de Hardware VHSIC o *VHSIC Hardware Description Language*) para implementaciones de altas prestaciones en FPGA y soslayar así las limitaciones de las herramientas existentes. El objetivo es completar el entorno para permitir el procesado iterativo en 'stream' de mallas de datos no estructurados, optimizando los accesos a las memorias externas e implementando operaciones locales de reducción.

En un segundo frente, el proyecto evalúa la viabilidad y la eficiencia de diseños en FPGA de algoritmos científicos de SKA. En concreto, el objetivo prioritario son los algoritmos de inteligencia artificial usados en el análisis de las imágenes obtenidas con SKA dado su posible alto nivel de paralelismo.

A.3.4 WRITE

WRITE (White Rabbit for Industrial Timing Enhancement) es un proyecto de metrología europeo dentro de la iniciativa EMPIR (European Metrology Programme

for Innovation and Research) del que fueron partícipes los institutos metrológicos de seis estados europeos (Italia, Reino Unido, Francia, Suecia, Países Bajos y Finlandia) junto con compañías involucradas en el desarrollo de tecnologías de distribución de tiempo [14].

El objetivo general del proyecto es demostrar todos los pasos metrológicos necesarios para la adopción industrial de WR, incluyendo la mejora de los dispositivos y el estudio de una implementación efectiva en infraestructuras informáticas industriales ordinarias sin degradación del rendimiento de la técnica en comparación con los resultados obtenidos en laboratorios de investigación controlados o utilizando infraestructuras de fibra dedicadas.

Existen dos objetivos específicos del proyecto WRITE que se relacionan de una manera especialmente cercana con este trabajo. El primero de ellos consiste en el desarrollo de técnicas validadas de transferencia de tiempo redundante y resiliente a los usuarios finales de la industria mediante múltiples enlaces de tiempo desde la fuente hasta el usuario, y proporcionando fuentes de reloj alternativas; el segundo, el desarrollo de nuevos dispositivos WR con rendimiento mejorado y que interactúen mejor con los protocolos y normas industriales existentes.

A.3.5 Doctorado Industrial

Este proyecto de tesis se ha elaborado en el marco del Programa de Ayudas para la formación de doctores en empresas, conocido popularmente como Doctorado Industrial, y cofinanciado entre el Ministerio de Ciencia y Tecnología del Gobierno de España (en el inicio de la convocatoria en cuestión, el Ministerio de Economía, Industria y Competitividad) y la empresa Safran Electronics and Defense Spain (en el inicio de la convocatoria, Seven Solutions).

El plan de investigación presentado a la Universidad de Granada y que contiene los objetivos descritos en la sección A.2 fue empleado como base para el desarrollo de una memoria científico-técnica escrita en colaboración con la compañía Seven Solutions y evaluada por expertos independientes de la Agencia Estatal de Investigación.

Este programa, cuyo acceso se encuentra controlado por una rigurosa evaluación de las propuestas por parte de la Agencia Estatal de Investigación, permite al participante desarrollar un conjunto de capacidades y competencias singulares fruto de la combinación entre la investigación y el mundo académico con la experiencia y el enfoque práctico propios del mundo industrial. Esta memoria trata de poner en valor ambos aspectos para demostrar tanto las competencias académicas desarrolladas como las implicaciones prácticas de la investigación.

Cabe destacar que este doble enfoque, académico y práctico, abre la oportunidad de participar en iniciativas de investigación en el marco europeo, como los proyectos anteriormente mencionados, desde una perspectiva única en la que el investigador está involucrado doblemente, tanto en el desarrollo desde un punto de vista teórico, como en el más cercano a la industria y a la gestión de proyectos.

Asimismo, este tipo de proyectos exigen que los resultados de la investigación asociada no solo sean difundidos en la presente memoria de tesis doctoral o en publicaciones de revistas científicas, sino que estos resultados deben tener una componente industrial que permita su explotación mediante el desarrollo de productos y servicios. En este contexto, el investigador se encuentra no solo al frente del trabajo de investigación propio, sino también integrado en un equipo de desarrollo de productos actualmente disponibles en el mercado.

El programa del proyecto de Doctorado Industrial fue elaborado teniendo presente que la correspondencia con los plazos del programa de doctorado de la Universidad de Granada debía mantenerse en la medida de lo posible. Así, se elaboró una planificación temporal del proyecto de investigación a realizar, descrito como líneas de trabajo que culminan en hitos justificables con documentos entregables asociados. Este esquema se mantiene alineado con el Plan de Investigación presentado a la Escuela de Posgrado de la Universidad de Granada.

A.4 Estructura de los capítulos

Una vez presentada una visión general del campo de aplicación y la motivación para este trabajo, la organización de este documento se reanuda con la estructura que se detalla a continuación:

- En el Capítulo 2 se presenta al lector el estado de la técnica. Está conformado por una revisión de la literatura enfocada al trabajo preexistente y a un análisis de las tecnologías de las que se pretende hacer uso en los desarrollos descritos en los capítulos posteriores.
- El Capítulo 3 tiene el fin de presentar el *entorno de trabajo*: la metodología empleada, las plataformas que sirven como punto de partida y el universo de herramientas empleadas.
- El Capítulo 4 presenta las técnicas y métodos desarrollados para incrementar la fiabilidad de un nodo, enfocado en extender su capacidad de mantener la sincronización en los intervalos de tiempo en que se ha perdido la visibilidad de la

referencia y en el desarrollo de topologías redundantes que permiten conmutación sin tiempo de reconfiguración.

- El Capítulo 5 está enfocado en el análisis de las especificidades introducidas por el medio óptico propias de las transmisiones a larga distancia, y de los problemas que puede introducir en la señal de tiempo las técnicas empleadas por los proveedores de servicio de la red de datos.
- El Capítulo 6 describe el despiegue piloto de un servicio de tiempo *TaaS* en cuyo desarrollo y mantenimiento se participó, poniendo en valor en el mundo real los conceptos desarrollados en los capítulos anteriores. El trabajo presentado en este capítulo fue originalmente publicado en [15]. Aquí se presenta con mayor detalle y en el contexto más amplio de esta tesis doctoral.
- El Capítulo 7 recoge las principales contribuciones de este trabajo y señala algunas posibles líneas de trabajo futuras.

APÉNDICE A. INTRODUCCIÓN

Apéndice B

Conclusiones

Este capítulo final sintetiza las principales contribuciones, estudios y aportaciones de la tesis. En primer lugar, se ofrece un resumen de las aportaciones, que recapitula brevemente los resultados de cada capítulo. A continuación, la sección de contribuciones principales evalúa la medida en que se han cumplido los objetivos iniciales, estableciendo un vínculo directo entre las metas fijadas y los resultados obtenidos.

A continuación, se presenta una lista de publicaciones científicas relacionadas con este trabajo en las que el autor ha contribuido, como representación del alcance y la repercusión de las tareas de investigación en la comunidad académica. El capítulo concluye identificando posibles direcciones futuras de investigación, destacando áreas en las que se puede seguir explorando la solidez y aplicabilidad de la tecnología WR.

B.1 Resumen de los resultados

El recorrido de esta tesis, desde la definición de los conceptos generales de transferencia de tiempo y frecuencia hasta sus amplias y tangibles potenciales aplicaciones, como se demuestra en la implementación del protocolo de redundancia para la red eléctrica inteligente o la prueba piloto de un servicio de tiempo para la industria financiera, ha sido tanto un desafío gratificante como una inmensa experiencia de aprendizaje para mí. Me gustaría comenzar el capítulo de conclusiones agradeciendo a todos los que me han apoyado de alguna manera durante estos últimos años, y quisiera extender mi más profundo agradecimiento al equipo WR del CERN y a toda la comunidad de OHWR, cuyo innovador trabajo sentó las bases para mi investigación.

Esta tesis presenta un estudio integral de los métodos de transferencia de tiempo y frecuencia a través de enlaces de fibra óptica de larga distancia, con un enfoque

en características de resiliencia y redundancia como un medio para extender la adopción de WR en aplicaciones que involucran la gestión de infraestructuras críticas. El trabajo presentado aquí se sustenta sobre una base de definiciones de conceptos de sincronización, metodología y detalles técnicos de los mecanismos de transferencia de tiempo de mejor rendimiento a través del medio óptico. El protocolo WR se enfatiza como la elección preferida para los avances realizados en este trabajo, con un examen detallado de su implementación, rendimiento y desafíos en despliegues de larga distancia. El estudio del estado del arte se complementa con un análisis de los requisitos de sincronización en sectores clave como finanzas tecnológicas, telecomunicaciones y *Smart grid*.

La metodología de investigación, la plataforma de hardware, las herramientas utilizadas y la instrumentación de laboratorio se presentaron en el Capítulo 3. Este apartado sirve como cimientos para el desarrollo y validación de las soluciones propuestas. Los tres capítulos siguientes se apoyan en dichos cimientos para presentar un análisis exhaustivo de los problemas que se abordan en esta investigación y los enfoques potenciales para solucionarlos. Estos tres capítulos tratan áreas muy diferenciadas de los sistemas de transferencia de tiempo. El Capítulo 4 está relacionado con mejoras dentro de los dispositivos WR (hardware y código), mientras que el Capítulo 5 está relacionado con todos los fenómenos que se manifestados fuera del dispositivo, es decir, la señal óptica y las consideraciones de red. Finalmente, el Capítulo 6 toma las contribuciones de los capítulos anteriores y los mecanismos puestos en práctica en el laboratorio, y los lleva a un piloto en un entorno real utilizando equipos comerciales para una implementación piloto de TaaS.

En el Capítulo 4, investigamos los factores que limitaban la aplicabilidad de soluciones de *holdover* en el WRS estándar, para lo cual llevamos a cabo un análisis más profundo de la sintonización y sincronización de fase del protocolo WR. Se diseñó una nueva plataforma de hardware en Seven Solutions, en cuyo desarrollo el autor tuvo un papel significativo, para abordar las limitaciones de hardware del diseño estándar. Notablemente, esta nueva plataforma, denominada WRS-LJ, mejora significativamente el ruido de fase del Grandmaster en 9 veces (de 8.4 ps RMS jitter a 0.9 ps RMS), y casi 7 veces (de 10.0 ps RMS jitter a 1.5 ps RMS) en el primer esclavo. Esta mejora, entre otros beneficios, permite conectar un mayor número de dispositivos en cadena sin aumentar el ruido con respecto al diseño anterior.

El WRS-LJ sirvió como plataforma experimental para el desarrollo de una solución de *holdover*. El desafío de minimizar el error de tiempo se abordó utilizando un modelo cuadrático. Aunque la precisión en la estimación de los parámetros estaba li-

mitada por ciertos componentes de hardware, los resultados obtenidos se estimaron al menos 200 veces mejores que las capacidades de *holdover* anteriores del WRS, como se determinó en [102]. Sin embargo, las capacidades de reserva del oscilador WRS estándar predeterminado eran tan limitadas que apenas pueden ser comparadas. El rendimiento de *holdover* del oscilador del WRS estándar está tan lejos de los requisitos de la mayoría de las aplicaciones que es difícil considerarlo como una solución de *holdover* viable. En contraste, la configuración experimental de este trabajo, bajo las mejores configuraciones encontradas (Abracon AOCJY 10 MHz, submuestreo y reproduciendo las mismas correcciones realizadas antes del lanzamiento del modo *holdover*), muestra un error de tiempo de menos de 1 ns después de 1 minuto en modo de reserva usando un OCXO relativamente asequible. La importancia de esta mejora es clave, ya que permite al dispositivo realizar un cambio de referencia de tiempo, incluso desde cero, sin un impacto significativo en las señales de tiempo. Esta capacidad es crucial para mantener la integridad y fiabilidad de los sistemas de transferencia de tiempo y frecuencia, especialmente en escenarios donde ocurren interrupciones o reconfiguraciones dinámicas de la red.

En términos de fiabilidad, la implementación de un protocolo de redundancia con cero tiempo de recuperación para topologías en anillo adecuado para datos y sincronización WR usando WRSs estándar demostró la viabilidad del proceso de cambio de referencia de tiempo con un impacto de menos de un nanosegundo en la precisión de las señales de tiempo. Esto se consiguió la implementación de un firmware capaz de manejar las características de redundancia que aseguran la recepción de datos y temporización incluso en caso de un único fallo en la red. Los resultados mostraron que el cambio se puede realizar efectivamente en el escenario de peor caso de un anillo de seis dispositivos WRS sin capacidades de *holdover* mejoradas. En el escenario de prueba, todo el anillo pudo cambiar a la referencia de respaldo en menos de un microsegundo desde que se detectó un fallo en el enlace. El impacto en la sincronización de este cambio en la topología de anillo fue limitado a 170 ps con respecto a su estado previo para el WRS que tuvo el peor comportamiento de todo el anillo. Es destacable que este experimento se realizó utilizando WRS estándar sin capacidades de *holdover* mejoradas.

En cuanto a la transmisión de datos, las capacidades de redundancia del protocolo HSR garantizan la entrega de datos y mensajes de control en caso de un fallo en el enlace, garantizando cero pérdida de datos. Los mecanismos de cambio de WRS se han optimizado para operar en redes de anillo, reduciendo la latencia de cambio en estas topologías de 2.2 a 1.0 microsegundos. Los resultados de ancho de banda en el escenario

redundante demuestran que la entrega de datos redundantes se puede mantener para anchos de banda de hasta 682 Mb/s.

En el Capítulo 5, realizamos una extensa revisión de los esfuerzos previos para transmitir WR a largas distancias, y los últimos desarrollos para resolver limitaciones en alcance y determinación de asimetría [75], [129]. Dado que las implementaciones analizadas están limitadas a redes privadas o condiciones de laboratorio, nos embarcamos en un estudio a gran escala de los fenómenos que impiden la transmisión en redes de fibra, con enfoque en la atenuación y dispersión cromática como factores que están presentes incluso en una red óptica correctamente diseñada. Los resultados muestran que el uso de equipos DWDM es recomendado incluso para redes de temporización de fibra oscura debido a la mayor estabilidad de los láseres transceptores DWDM. El análisis de este factor se hizo exclusivamente revisando la literatura debido a la falta de equipos especializados, y se encontró una reducción de 7.7 veces en el presupuesto de incertidumbre SFP para enlaces de 80 km.

Para complementar este estudio, también se analizaron los mecanismos y componentes que se despliegan en redes ópticas para mitigar efectos no deseados y transmitir y enrutar con éxito las señales ópticas. El resultado de esta revisión arrojó luz sobre la importancia de un modelado preciso de la red y un estricto control de gestión de la red en implementaciones de larga distancia. Para casos en los que esto no es posible, se estimó el impacto de cada posible elemento de un banco de pruebas de larga distancia en términos de retardo, asimetría, y los cambios inducidos en estas dos variables por la variabilidad térmica. De los elementos analizados del banco de pruebas, los coeficientes térmicos para el RTT varían desde despreciables para elementos discretos como un multiplexor DWDM, hasta 46 ps $km^{-1}K^{-1}$ en el caso de DCM basados en DCF. La contribución de los nodos finales se estima en 14 ps K^{-1}. Afortunadamente, el equipo de red generalmente se encuentra en salas con control de temperatura donde las contribuciones de retardo y asimetría se pueden considerar estáticas. Gracias a esto, nuestros hallazgos también subrayan la posibilidad de utilizar técnicas de calibración alternativas para compensar las asimetrías desconocidas en redes más allá del alcance de nuestro control con un impacto limitado en el rendimiento de la temporización.

Finalmente, en el Capítulo 6, el conocimiento y las contribuciones recopiladas de los capítulos anteriores se aplicaron para diseñar, implementar y evaluar un servicio piloto de tiempo para el operador del mercado de valores español BME. El escenario piloto implementó con éxito una solución que diseminó el tiempo desde servidores de tiempo redundantes hasta el centro de datos remoto del cliente con fibra óptica utilizando tecnología WR a través de una red de un operador de telecomunicaciones.

Este sistema proporciona al cliente capacidades de resiliencia a través de implementaciones de *holdover* tanto en el servidor de tiempo como en el lado del cliente, así como mediante redundancia en el hardware del servidor de tiempo y en la diseminación. El servicio cumple con los requisitos actuales establecidos por la regulación y está bien equipado para absorber requisitos futuros: durante la operación normal, la trazabilidad a UTC está garantizada con un error de tiempo máximo de 15 ns. Es importante destacar que esta limitación de 15 ns se debe principalmente al receptor GNSS utilizado para la calibración. El empleo de un receptor GNSS multifrecuencia sería suficiente para superar esta limitación, mejorando aún más la exactitud del sistema. Los mecanismos de resiliencia aseguran que el sistema pueda soportar la pérdida de referencia GNSS durante semanas mientras mantiene el error de tiempo por debajo de los requisitos de la regulación actual. Si el equipo del lado del cliente pierde la visibilidad del servidor de tiempo, los datos demuestran que los nodos en el extremo del cliente pueden operar de manera autónoma hasta 24 horas antes de que el error de tiempo de la señal alcance un microsegundo.

B.2 Contribuciones principales

En este punto intermedio del capítulo de conclusiones de este trabajo, se vuelve necesario revisitar el punto de partida — los objetivos originales que dieron forma a esta investigación académica. Es esencial comprobar el alineamiento de estos con las contribuciones efectivamente llevadas a cabo. Este momento representa una oportunidad propicia para evaluar el avance alcanzado y verificar si se han satisfecho las expectativas previamente establecidas:

1. Realizamos una revisión de los mecanismos de resiliencia y redundancia en relación con la diseminación del tiempo, con un enfoque particular en la infraestructura crítica. Nuestro análisis se centró en las implicaciones y necesidades específicas de este segmento de la industria. A continuación ampliamos nuestro análisis para incluir las condiciones y desafíos específicos de los enlaces WR a larga distancia. Esta contribución está estrechamente relacionada con los Objetivos 1 y 3.

2. La respuesta al ruido y la circuitería del reloj del WRS se mejoraron en un nuevo diseño de hardware. Aunque no se indica directamente en los objetivos, la reducción del ruido y la mejora de la estabilidad (resumidos en la sección 4.3) son necesarios para alcanzar el Objetivo 4. La adaptación del diseño HDL para este

nuevo hardware para incluir nuevos osciladores experimentales también es un prerrequisito esencial para la implementación de cualquier solución de *holdover* (Objetivo 4). Los resultados de este punto han sido publicados en OHWR[1] como un producto de código abierto, tanto hardware como firmware.

3. Se analizaron y modelaron los principales contribuyentes a la degradación del error de tiempo en holdover. Se desarrolló una implementación experimental que estima el error inicial de frecuencia fraccionaria, mejorando su rendimiento con cualquier implementación WR de código abierto previa y logrando al menos 60 segundos de tiempo de holdover manteniendo el error por debajo de 1 ns. Este es uno de los objetivos del Objetivo 4. Esta contribución se realizó en el marco del proyecto WRITE y, en consecuencia, el diseño y el código también están disponibles en la comunidad de código abierto.

4. En colaboración con otros investigadores, se logró la implementación de un protocolo de anillo redundante con cero tiempo de recuperación para tiempo y datos, demostrando el cambio de referencia de tiempo sin interrupciones en las señales de tiempo del equipo esclavo. Esta contribución completa el logro de los Objetivos 3 y 4.

5. En esta misma colaboración, demostramos el primer uso de WR en topologías de anillo, implementando nuevos roles de reloj para una red WR, como relojes transparentes (TC) e híbridos (HY). Esto se alinea con los objetivos 3 y 4.

6. Se demostró la aplicabilidad de la topología redundante en anillo para Smart Grid en un banco de pruebas y se publicó en una revista científica revisada por pares [1], satisfaciendo el Objetivo 5 para el caso de Smart Grid.

7. Se desplegó un banco de pruebas de larga distancia en condiciones controladas de laboratorio. A partir de esto, se derivó un modelo de retraso aditivo, error de tiempo y dependencia de la temperatura de sus componentes. Esto se traduce directamente al Objetivo 2.

8. Se logró un hito significativo con el despliegue de un caso de uso en el mundo real en el sector financiero, marcando la primera instancia documentada de un TaaS siendo desplegado en una red pública utilizando tecnología WR [15]. Esta aplicación pionera sirvió como verificación de los principios y soluciones analizados e implementados en capítulos anteriores. La robustez del servidor

[1] `https://ohwr.org/project/wrs-lj-hw/wikis/home`

de tiempo duplicado y las características de resiliencia en el nodo del lado del cliente fueron probadas exhaustivamente contra fuentes de tiempo alternativas y validadas bajo condiciones del mundo real. Es importante destacar que estas características fueron validadas sin ningún impacto en la percepción del cliente del servicio de tiempo. Este logro, que se corresponde directamente con el Objetivo 5, resalta el potencial de nuestro trabajo y sirve de precedente para futuros despliegues de TaaS en redes públicas utilizando WR.

Cuando se definieron los objetivos de esta tesis, propusimos una solución al desafío de lograr la diseminación de tiempo de alta precisión a larga distancia sobre enlaces de fibra óptica, con el objetivo de superar el rendimiento de los métodos tradicionales de satélite y protocolos de red como PTP o NTP. Nuestra hipótesis era que los enlaces WR podrían ser desplegados con tolerancia a fallos y resiliencia integradas, logrando una precisión y exactitud equivalente o superior a los sistemas GNSS, todo esto mientras se utilizan las redes de fibra de telecomunicaciones existentes. Las múltiples contribuciones enumeradas en esta sección han afirmado la validez de esta hipótesis.

B.3 Publicaciones

Este apartado presenta una serie de publicaciones generadas durante el curso de la presente investigación doctoral. Las contribuciones enunciadas a continuación han sido objeto de revisión por pares y han sido publicadas en revistas académicas de alto impacto y reconocido prestigio internacional y en conferencias internacionales.

B.3.1 Publicaciones en revistas

- J. L. Gutiérrez-Rivas, J. López-Jiménez, E. Ros and J. Díaz, “White Rabbit HSR: A Seamless Subnanosecond Redundant Timing System With Low-Latency Data Capabilities for the Smart Grid” in IEEE Transactions on Industrial Informatics, vol. 14, no. 8, pp. 3486-3494, Aug. 2018, doi: 10.1109/TII.2017.2779240.

- F. Ramos, J. L. Gutiérrez-Rivas, J. López-Jiménez, B. Caracuel and J. Díaz, “Accurate Timing Networks for Dependable Smart Grid Applications,” in IEEE Transactions on Industrial Informatics, vol. 14, no. 5, pp. 2076-2084, May 2018, doi: 10.1109/TII.2017.2787145.

- F. Girela-López, J. López-Jiménez, M. Jiménez-López, R. Rodríguez, E. Ros and J. Díaz, “IEEE 1588 High Accuracy Default Profile: Applications and

Challenges," in IEEE Access, vol. 8, pp. 45211-45220, 2020, doi: 10.1109/ACCESS.2020.2978337.

- J. Lopez-Jimenez, J. L. Gutierrez-Rivas, E. Marin-Lopez, M. Rodriguez-Alvarez and J. Diaz, "Time as a Service Based on White Rabbit for Finance Applications," in IEEE Communications Magazine, vol. 58, no. 4, pp. 60-66, April 2020, doi: 10.1109/MCOM.001.1900602.

- M. Jiménez-López, F. Girela-López, J. López-Jiménez, E. Marín-López, R. Rodríguez and J. Díaz, "10 Gigabit White Rabbit: Sub-Nanosecond Timing and Data Distribution," in IEEE Access, vol. 8, pp. 92999-93010, 2020, doi: 10.1109/ACCESS.2020.2995179.

- En preparación: un próximo artículo, que será enviado a *IEEE Transactions on Instrumentation and Measurement*, proporcionará y ampliará un informe detallado de los resultados de nuestra implementación de holdover. El título específico y la lista de autores aún están bajo consideración en el momento de escribir esta tesis.

B.3.2 Publicaciones en conferencias

- J. López-Jiménez, J. Díaz and M. Rodríguez-Álvarez, "Impact of network component temperature variation on long haul White Rabbit links," 2018 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS), Geneva, Switzerland, 2018, pp. 1-6, doi: 10.1109/ISPCS.2018.8543085.

- E. F. Dierikx, Y. Xie, A. Savencu, J. Lopez and J. L. G. Rivas, "White Rabbit Multi-Point Time Distribution Network," 2021 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS), Gainesville, FL, USA, 2021, pp. 1-4, doi: 10.1109/EFTF/IFCS52194.2021.9604343.

B.4 Trabajo futuro

Las contribuciones de este estudio abren varias vías para futuras investigaciones y desarrollos. Probablemente el camino más evidente sea continuar mejorando la robustez de WR combinando las implementaciones de *switchover* y *holdover* que se han

propuesto en este trabajo. Lamentablemente, debido a restricciones externas, el desarrollo de la implementación de holdover y la red redundante con capacidades de *switchover* sin interrupciones se realizaron en diferentes momentos y utilizando diferentes versiones de la plataforma de hardware, lo que hace crecer el tiempo estimado para la adaptación y portabilidad de una solución que implemente estas características combinadas. Sin embargo, como desarrollador dentro de Safran Electronics and Defense Spain, el autor se encuentra en una posición única para impulsar la integración de estas soluciones desarrolladas en otros productos y plataformas comerciales. La capacidad de seguir en fase simultáneamente dos referencias WR mientras se dispone de capacidades de holdover mejoradas proporciona un gran punto de partida para desarrollar un mecanismo de *switchover* y *failover* potente y versátil con características de monitorización adicionales.

También relacionado con el holdover, hay mejoras que se pueden implementar para continuar mejorando el comportamiento de los estimadores del modelo. En primer lugar, se podría desarrollar un mecanismo de holdover más sofisticado si los parámetros internos del control OCXO fueran visibles para el sistema WR. Si esto no fuera posible, otro enfoque sería encontrar diferentes componentes de hardware que permitieran el desarrollo de un algoritmo OCXO adaptativo propio. El tema de la detección precoz de fallos como condición desencadenante del modo holdover es un campo en el que nuevos esfuerzos darán lugar a mejoras tangibles, aunque para ello podrían ser necesarios hardware o mecanismos de comunicación adicionales a fin de lograr una población considerable para la votación por mayoría.

En relación con el modelado y la ampliación del modelo de enlace WR a escenarios más diversos y experimentales, el uso de WR sobre fibras de núcleo hueco puede ser un tema digno de estudio, ya que estas se están considerando en ciertos escenarios experimentales como una forma de lograr latencias de enlace aún más bajas. Las fibras de núcleo hueco muestran una velocidad de transmisión cercana al vacío, pero no son rentables económicamente en la actualidad y su alcance es limitado debido a la alta atenuación.

Otra vía prometedora para futuras investigaciones radica en la exploración de fuentes de tiempo alternativas, particularmente aquellas basadas en satélites de órbita terrestre baja (LEO). Empresas como Satelles[2] o Xona Space Systems[3] ofrecen servicios de tiempo y ubicación seguros que podrían mejorar en gran medida la resistencia

[2] `https://safran-navigation-timing.com/testing-confirms-resilient-timing-success/`

[3] `https://www.xonaspace.com/pulsar`

de los sistemas de tiempo para infraestructuras críticas. Esto podría proporcionar una alternativa robusta y segura o un complemento a las fuentes de tiempo tradicionales basadas en GNSS, ofreciendo posibles beneficios en términos de disponibilidad de señal, cobertura global y resistencia a interferencias y falsificaciones. Investigar la integración de estas fuentes de tiempo emergentes con la tecnología WR podría abrir nuevas posibilidades para mejorar la confiabilidad y seguridad de los servicios de tiempo. Esto podría ser particularmente valioso en aplicaciones de infraestructuras críticas y financieras donde los altos niveles de resiliencia y seguridad son fundamentales.

Siguiendo la misma línea de escenarios más diversos para la difusión del tiempo basada en WR, medios alternativos como las ondas milimétricas o los láseres de espacio libre son temas que están recibiendo atención para aplicaciones de nicho. La aplicabilidad de la WR y los retos asociados son temas de interés y tangencialmente relacionados con este trabajo, pero las diferencias fundamentales con el medio óptico guiado favorecen un enfoque de desarrollo desde las bases, de abajo hacia arriba.

References

[1] F. Ramos, J. L. Gutiérrez-Rivas, J. López-Jiménez, B. Caracuel, and J. Díaz, «Accurate Timing Networks for Dependable Smart Grid Applications», *IEEE Transactions on Industrial Informatics*, vol. 14, no. 5, pp. 2076–2084, May 2018, Conference Name: IEEE Transactions on Industrial Informatics, ISSN: 1941-0050. DOI: 10.1109/TII.2017.2787145.

[2] S. Siu, W.-H. Tseng, H.-f. Hu, S.-Y. Lin, C.-S. Liao, and Y.-L. Lai, «In-Band Asymmetry Compensation for Accurate Time/Phase Transport over Optical Transport Network», en, *The Scientific World Journal*, vol. 2014, pp. 1–8, 2014, ISSN: 2356-6140, 1537-744X. DOI: 10.1155/2014/408613. [Online]. Available: http://www.hindawi.com/journals/tswj/2014/408613/ (visited on 05/10/2023).

[3] I. R. Bartky, «The Adoption of Standard Time», *Technology and Culture*, vol. 30, no. 1, pp. 25–56, 1989, Publisher: [The Johns Hopkins University Press, Society for the History of Technology], ISSN: 0040-165X. DOI: 10.2307/3105430. [Online]. Available: https://www.jstor.org/stable/3105430 (visited on 05/10/2023).

[4] S. Floyd and M. Allman, «Comments on the Usefulness of Simple Best-Effort Traffic», Internet Engineering Task Force, Request for Comments RFC 5290, Jul. 2008, Num Pages: 20. DOI: 10.17487/RFC5290. [Online]. Available: https://datatracker.ietf.org/doc/rfc5290 (visited on 05/10/2023).

[5] F. Dovis, Ed., *GNSS interference, threats, and countermeasures* (Artech House GNSS technology and applications series). Boston: Artech House, 2015, OCLC: ocn893017345, ISBN: 978-1-60807-810-3.

[6] «IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems», IEEE, Tech. Rep., ISBN: 9780738154008. DOI: 10.1109/IEEESTD.2008.4579760. [Online]. Available: https://ieeexplore.ieee.org/document/4579760/ (visited on 06/10/2023).

[7] J. Serrano, M. Cattin, E. Gousiou, *et al.*, «The White Rabbit project», 2013. [Online]. Available: https://cds.cern.ch/record/1743073 (visited on 05/10/2023).

[8] Z. Heckenast, J. Kronjaeger, and L. Lobo, *CLONETS positioning paper*, Sep. 2018. [Online]. Available: https://www.clonets.eu/clonets-misc.html?&no_cache=1&cid=6410&did=15020&sechash=d586f809 (visited on 05/10/2023).

[9] J. Vojtěch, L. Altmannová, G. Roberts, *et al.*, «CLONETS-DS – Clock Network Services-Design Study : Strategy and innovation for clock services over optical-fibre networks in Europe», in *2021 Conference on Lasers and Electro-Optics (CLEO)*, ISSN: 2160-8989, May 2021, pp. 1–2.

[10] European Space Agency, *TOWR - ESA's NAVISP Programmes*, Aug. 2018. [Online]. Available: https://navisp.esa.int/project/details/36/show (visited on 05/10/2023).

[11] Instituto de Astrofísica de Andalucía, *AMIGA6 – Analysis of the Interstellar Medium of Isolated Galaxies*, en-GB. [Online]. Available: https://amiga.iaa.csic.es/amiga6/ (visited on 05/10/2023).

[12] A. R. Taylor, «The Square Kilometre Array», en, *Proceedings of the International Astronomical Union*, vol. 8, no. S291, pp. 337–341, Aug. 2012, ISSN: 1743-9213, 1743-9221. DOI: 10.1017/S1743921312024039. [Online]. Available: https://www.cambridge.org/core/product/identifier/S1743921312024039/type/journal_article (visited on 05/10/2023).

[13] Instituto de Astrofísica de Andalucía, *AMIGA7 – Analysis of the interstellar Medium of Isolated GAlaxies*, en-GB. [Online]. Available: https://amiga.iaa.csic.es/amiga7/ (visited on 05/10/2023).

[14] E. L. English, B. Eglin, A. Parsons, *et al.*, «The WRITE (White Rabbit for Industrial Timing Enhancement) Project Update», in *Proceedings of the 52nd Annual Precise Time and Time Interval Systems and Applications Meeting*, 2021, pp. 146–166.

[15] J. Lopez-Jimenez, J. L. Gutierrez-Rivas, E. Marin-Lopez, M. Rodriguez-Alvarez, and J. Diaz, «Time as a Service Based on White Rabbit for Finance Applications», *IEEE Communications Magazine*, vol. 58, no. 4, pp. 60–66, Apr. 2020, Conference Name: IEEE Communications Magazine, ISSN: 1558-1896. DOI: 10.1109/MCOM.001.1900602.

[16] G. K. McMillan and P. H. Vegas, *Process / Industrial Instruments and Controls Handbook, Sixth Edition*, en. McGraw Hill Professional, Apr. 2019, Google-Books-ID: IsWIDwAAQBAJ, ISBN: 978-1-260-11798-1.

[17] D. Glass, *What Happens If GPS Fails?*, en, Section: Technology, Jun. 2016. [Online]. Available: https://www.theatlantic.com/technology/archive/2016/06/what-happens-if-gps-fails/486824/ (visited on 05/11/2023).

[18] M. V. Rafter, *U.S. Transportation Officials Seek Alternative Tech for GPS - IEEE Spectrum*, en, Apr. 2020. [Online]. Available: https://spectrum.ieee.org/us-transportation-officials-seek-alternative-tech-for-gps (visited on 05/12/2023).

[19] S. Zacks, *Introduction to Reliability Analysis* (Springer Texts in Statistics). New York, NY: Springer New York, 1992, ISBN: 978-1-4612-7697-5 978-1-4612-2854-7. DOI: 10.1007/978-1-4612-2854-7. [Online]. Available: http://link.springer.com/10.1007/978-1-4612-2854-7 (visited on 05/11/2023).

[20] *IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications*, ISBN: 9781504423267. DOI: 10.1109/IEEESTD.2017.7953616. [Online]. Available: https://ieeexplore.ieee.org/document/7953616/ (visited on 05/11/2023).

[21] J.-C. Lin, «Synchronization Requirements for 5G: An Overview of Standards and Specifications for Cellular Networks», *IEEE Vehicular Technology Magazine*, vol. 13, no. 3, pp. 91–99, Sep. 2018, Conference Name: IEEE Vehicular Technology Magazine, ISSN: 1556-6080. DOI: 10.1109/MVT.2018.2813339.

[22] 3rd Generation Partnership Project 2, *C.S0005-D: Physical Layer Standard for cdma2000 Spread Spectrum Systems*, Sep. 2005.

[23] S. Ruffini, S. Rodrigues, M. Lipinski, and J.-C. Lin, «Synchronization standards toward 5G», *IEEE Communications Standards Magazine*, vol. 1, no. 1, pp. 50–

51, Mar. 2017, Conference Name: IEEE Communications Standards Magazine, ISSN: 2471-2833. DOI: 10.1109/COMSTD.2017.7885238.

[24] H. Li, L. Han, R. Duan, and G. M. Garner, «Analysis of the Synchronization Requirements of 5g and Corresponding Solutions», *IEEE Communications Standards Magazine*, vol. 1, no. 1, pp. 52–58, Mar. 2017, ISSN: 2471-2825, 2471-2833. DOI: 10.1109/MCOMSTD.2017.1600768ST. [Online]. Available: https://ieeexplore.ieee.org/document/7885239/ (visited on 05/11/2023).

[25] D. Schneider, «The microsecond market», *IEEE Spectrum*, vol. 49, no. 6, pp. 66–81, Jun. 2012, ISSN: 0018-9235. DOI: 10.1109/MSPEC.2012.6203974. [Online]. Available: http://ieeexplore.ieee.org/document/6203974/ (visited on 05/10/2023).

[26] *Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments and amending Directive 2002/92/EC and Directive 2011/61/EU (recast) (Text with EEA relevance)Text with EEA relevance*, en, Legislative Body: OP_DATPRO, Mar. 2023. [Online]. Available: http://data.europa.eu/eli/dir/2014/65/2023-03-23/eng (visited on 05/11/2023).

[27] European Securities and Markets Authority, *Commission Delegated Regulation (EU) 2017/574 of 7 June 2016 supplementing Directive 2014/65/EU of the European Parliament and of the Council with regard to regulatory technical standards for the level of accuracy of business clocks*, 2017. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32017R0574.

[28] *IEC TR 63097 Smart Grid Standarditazion Roadmap*, en-US. [Online]. Available: https://syc-se.iec.ch/deliveries/iec-tr-63097-smart-grid-roadmap/ (visited on 05/11/2023).

[29] H. Farhangi, «The path of the smart grid», *IEEE Power and Energy Magazine*, vol. 8, no. 1, pp. 18–28, Jan. 2010, ISSN: 1540-7977. DOI: 10.1109/MPE.2009.934876. [Online]. Available: http://ieeexplore.ieee.org/document/5357331/ (visited on 05/11/2023).

[30] *IEEE/IEC International Standard - Measuring relays and protection equipment - Part 118-1: Synchrophasor for power systems - Measurements*, ISBN: 9781504453615. DOI: 10.1109/IEEESTD.2018.8577045. [Online]. Available: https://ieeexplore.ieee.org/document/8577045/ (visited on 05/11/2023).

[31] B. Moussa, M. Debbabi, and C. Assi, «Security Assessment of Time Synchronization Mechanisms for the Smart Grid», *IEEE Communications Surveys & Tutorials*, vol. 18, no. 3, pp. 1952–1973, 2016, ISSN: 1553-877X. DOI: 10.1109/COMST.2016.2525014. [Online]. Available: http://ieeexplore.ieee.org/document/7397831/ (visited on 05/11/2023).

[32] R. Mackiewicz, «Overview of IEC 61850 and benefits», in *2006 IEEE Power Engineering Society General Meeting*, Montreal, Que., Canada: IEEE, 2006, 8 pp. ISBN: 978-1-4244-0493-3. DOI: 10.1109/PES.2006.1709546. [Online]. Available: http://ieeexplore.ieee.org/document/1709546/ (visited on 05/11/2023).

[33] *IEC 62439-3:2021 Industrial communication networks - Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR)*, Dec. 2021.

[34] W. J. Riley, *Handbook of Frequency Stability Analysis*, en. U.S. Department of Commerce, National Institute of Standards and Technology, 2007, Google-Books-ID: k_n0nQAACAAJ.

[35] M. Lombardi, «Fundamentals of Time and Frequency», en, in *Mechatronics*, R. Bishop, Ed., CRC Press, Sep. 2005, pp. 10–1–10–18, ISBN: 978-0-8493-6358-0 978-1-4200-3724-1. DOI: 10.1201/9781420037241.ch10. [Online]. Available: http://www.crcnetbase.com/doi/10.1201/9781420037241.ch10 (visited on 05/11/2023).

[36] S. Bregni, *Synchronization of digital telecommunications networks*. Chichester ; New York: Wiley, 2002, OCLC: ocm46600898, ISBN: 978-0-471-61550-7.

[37] *Recommendation G.810: Definitions and terminology for synchronization networks*, Aug. 1996.

[38] «IEEE Standard 1139-2008 Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities», IEEE, Tech. Rep., ISBN: 9780738168555 9780738168562. DOI: 10.1109/IEEESTD.2008.4797525. [Online]. Available: http://ieeexplore.ieee.org/document/4797525/ (visited on 05/11/2023).

[39] D. Sullivan, D. Allan, D. Howe, and F. Walls, *NIST Technical Note 1337 Characterization of Clocks and Oscillators*. U.S. Department of Commerce, National Institute of Standards and Technology, 1990.

[40] D. Allan, «Time and Frequency (Time-Domain) Characterization, Estimation, and Prediction of Precision Clocks and Oscillators», *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 34, no. 6, pp. 647–654, Nov. 1987, ISSN: 0885-3010. DOI: 10.1109/T-UFFC.1987.26997. [Online]. Available: http://ieeexplore.ieee.org/document/1539968/ (visited on 05/12/2023).

[41] P. J. Teunissen and O. Montenbruck, Eds., *Springer Handbook of Global Navigation Satellite Systems*, en. Cham: Springer International Publishing, 2017, ISBN: 978-3-319-42926-7 978-3-319-42928-1. DOI: 10.1007/978-3-319-42928-1. [Online]. Available: http://link.springer.com/10.1007/978-3-319-42928-1 (visited on 05/18/2023).

[42] T. N. Bandi, «A Comprehensive Overview of Atomic Clocks and their Applications», *Biology, Engineering, Medicine and Science Reports*, vol. 9, no. 1, pp. 01–10, Jan. 2023, ISSN: 24546895. DOI: 10.5530/bems.9.1.1. [Online]. Available: https://www.bemsreports.org/index.php/bems/article/view/128 (visited on 05/18/2023).

[43] B. J. Bloom, T. L. Nicholson, J. R. Williams, *et al.*, «An optical lattice clock with accuracy and stability at the 10-18 level», en, *Nature*, vol. 506, no. 7486, pp. 71–75, Feb. 2014, ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature12941. [Online]. Available: https://www.nature.com/articles/nature12941 (visited on 05/20/2023).

[44] F. Riehle, «Towards a redefinition of the second based on optical atomic clocks», en, *Comptes Rendus Physique*, vol. 16, no. 5, pp. 506–515, Jun. 2015, ISSN: 16310705. DOI: 10.1016/j.crhy.2015.03.012. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1631070515000638 (visited on 05/20/2023).

[45] A. Gifford, «One-way GPS Time Transfer 2000», in *Proceedings of the 32th Annual Precise Time and Time Interval Systems and Applications Meeting*, Reston, Virginia, Nov. 2000, pp. 137–146.

[46] *Global Positioning System Standard Positioning Service Performance Standard, 5th Edition*, Apr. 2020.

[47] D. Hanson, «Fundamentals of two-way time transfers by satellite», in *Proceedings of the 43rd Annual Symposium on Frequency Control*, Denver, CO, USA: IEEE, 1989, pp. 174–178. DOI: 10.1109/FREQ.1989.68861. [Online]. Avail-

able: http://ieeexplore.ieee.org/document/68861/ (visited on 05/11/2023).

[48] S. Bisnath and Y. Gao, *Precise Point Positioning: A Powerful Technique with a Promising Future*, Apr. 2009. [Online]. Available: https://gge.ext.unb.ca/Resources/ion%20gnss+%202016/gpsworld.april09.pdf (visited on 05/12/2023).

[49] P. J. G. Teunissen and A. Khodabandeh, «Review and principles of PPP-RTK methods», en, *Journal of Geodesy*, vol. 89, no. 3, pp. 217–240, Mar. 2015, ISSN: 1432-1394. DOI: 10.1007/s00190-014-0771-3. [Online]. Available: https://doi.org/10.1007/s00190-014-0771-3 (visited on 05/12/2023).

[50] K. Wang and A. El-Mowafy, «LEO satellite clock analysis and prediction for positioning applications», en, *Geo-spatial Information Science*, vol. 25, no. 1, pp. 14–33, Jan. 2022, ISSN: 1009-5020, 1993-5153. DOI: 10.1080/10095020.2021.1917310. [Online]. Available: https://www.tandfonline.com/doi/full/10.1080/10095020.2021.1917310 (visited on 05/12/2023).

[51] P. Defraigne, P. Tavella, I. Sesia, *et al.*, «Demonstrator of Time Services based on European GNSS Signals: The H2020 DEMETRA Project», Monterey, California, Mar. 2017, pp. 127–137. DOI: 10.33012/2017.14982. [Online]. Available: https://www.ion.org/publications/abstract.cfm?articleID=14982 (visited on 05/12/2023).

[52] *Truck Driver with GPS Jammer Accidentally Jams Newark Airport*, en. [Online]. Available: https://www.scientificamerican.com/article/truck-driver-has-gps-jammer-accident-2013-08/ (visited on 05/11/2023).

[53] A. Darabseh, E. Bitsikas, and B. Tedongmo, «Detecting GPS Jamming Incidents in OpenSky Data», pp. 97–84. DOI: 10.29007/1mmw. [Online]. Available: https://easychair.org/publications/paper/Mp3x (visited on 05/11/2023).

[54] M. Lenhart, M. Spanghero, and P. Papadimitratos, «DEMO: Relay/Replay Attacks on GNSS signals», 2022, Publisher: arXiv Version Number: 1. DOI: 10.48550/ARXIV.2202.10897. [Online]. Available: https://arxiv.org/abs/2202.10897 (visited on 05/11/2023).

[55] D. Schmidt, K. Radke, S. Camtepe, E. Foo, and M. Ren, «A Survey and Analysis of the GNSS Spoofing Threat and Countermeasures», en, *ACM Computing Surveys*, vol. 48, no. 4, pp. 1–31, May 2016, ISSN: 0360-0300, 1557-7341. DOI: 10.1145/2897166. [Online]. Available: https://dl.acm.org/doi/10.1145/2897166 (visited on 05/11/2023).

[56] Timing Committee Telecommunications and TIming Group, *IRIG Serial Time Code Formats*, Aug. 2016.

[57] V. A. Dzuba and V. V. Flambaum, «Limits on gravitational Einstein Equivalence Principle violation from monitoring atomic clock frequencies during a year», 2016, Publisher: arXiv Version Number: 1. DOI: 10.48550/ARXIV.1608.06050. [Online]. Available: https://arxiv.org/abs/1608.06050 (visited on 05/11/2023).

[58] N. Huntemann, B. Lipphardt, C. Tamm, V. Gerginov, S. Weyers, and E. Peik, «Improved limit on a temporal variation of m_p/m_e from comparisons of Yb$^+$ and Cs atomic clocks», 2014, Publisher: arXiv Version Number: 1. DOI: 10.48550/ARXIV.1407.4408. [Online]. Available: https://arxiv.org/abs/1407.4408 (visited on 05/11/2023).

[59] S. Droste, F. Ozimek, T. Udem, *et al.*, «Optical-Frequency Transfer over a Single-Span 1840 km Fiber Link», en, *Physical Review Letters*, vol. 111, no. 11, p. 110 801, Sep. 2013, ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.111.110801. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.111.110801 (visited on 05/11/2023).

[60] C. Lisdat, G. Grosche, N. Quintin, *et al.*, «A clock network for geodesy and fundamental science», en, *Nature Communications*, vol. 7, no. 1, p. 12 443, Aug. 2016, ISSN: 2041-1723. DOI: 10.1038/ncomms12443. [Online]. Available: https://www.nature.com/articles/ncomms12443 (visited on 06/20/2023).

[61] P. Krehlik, L. Sliwczynski, L. Buczek, J. Kolodziej, and M. Lipinski, «ELSTAB—Fiber-Optic Time and Frequency Distribution Technology: A General Characterization and Fundamental Limits», *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 63, no. 7, pp. 993–1004, Jul. 2016, ISSN: 0885-3010. DOI: 10.1109/TUFFC.2015.2502547. [Online]. Available: http://ieeexplore.ieee.org/document/7332966/ (visited on 05/11/2023).

[62] IEEE, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, ISBN: 9781504463416. DOI: 10.1109/IEEESTD.2020.9120376. [Online]. Available: https://ieeexplore.ieee.org/document/9120376/ (visited on 05/11/2023).

[63] J. Martin, J. Burbank, W. Kasch, and P. D. L. Mills, «Network Time Protocol Version 4: Protocol and Algorithms Specification», Internet Engineering Task Force, Request for Comments RFC 5905, Jun. 2010, Num Pages: 110. DOI: 10.17487/RFC5905. [Online]. Available: https://datatracker.ietf.org/doc/rfc5905 (visited on 05/11/2023).

[64] O. Obleukhov, *Building a more accurate time service at Facebook scale*, en-US, Mar. 2020. [Online]. Available: https://engineering.fb.com/2020/03/18/production-engineering/ntp-service/ (visited on 05/11/2023).

[65] D. Mills, «Adaptive hybrid clock discipline algorithm for the network time protocol», *IEEE/ACM Transactions on Networking*, vol. 6, no. 5, pp. 505–514, Oct. 1998, ISSN: 10636692. DOI: 10.1109/90.731182. [Online]. Available: http://ieeexplore.ieee.org/document/731182/ (visited on 05/11/2023).

[66] F. Girela-Lopez, J. Lopez-Jimenez, M. Jimenez-Lopez, R. Rodriguez, E. Ros, and J. Diaz, «IEEE 1588 High Accuracy Default Profile: Applications and Challenges», *IEEE Access*, vol. 8, pp. 45 211–45 220, 2020, ISSN: 2169-3536. DOI: 10.1109/ACCESS.2020.2978337. [Online]. Available: https://ieeexplore.ieee.org/document/9023998/ (visited on 05/12/2023).

[67] M. Lipinski, E. Van Der Bij, J. Serrano, *et al.*, *White Rabbit Applications and Enhancements*, Geneva, Switzerland, Sep. 2018. [Online]. Available: https://ohwr.org/project/white-rabbit/uploads/5914348e620dd8a05efecac8a1485004/WhiteRabbitApplicationsEnhancements_Lipinski_ISPCS_2018.pdf (visited on 05/31/2023).

[68] P. Moreira, P. Alvarez, J. Serrano, and I. Darwazeh, «Sub-nanosecond digital phase shifter for clock synchronization applications», in *2012 IEEE International Frequency Control Symposium Proceedings*, Baltimore, MD, USA: IEEE, May 2012, pp. 1–6, ISBN: 978-1-4577-1820-5 978-1-4577-1821-2 978-1-4577-1819-9. DOI: 10.1109/FCS.2012.6243715. [Online]. Available: http://ieeexplore.ieee.org/document/6243715/ (visited on 05/11/2023).

[69] M. Rizzi, M. Lipinski, T. Wlostowski, *et al.*, «White rabbit clock characteristics», in *2016 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS)*, Stockholm, Sweden: IEEE, Sep. 2016, pp. 1–6, ISBN: 978-1-4673-9615-8. DOI: 10.1109/ISPCS.2016.7579514. [Online]. Available: http://ieeexplore.ieee.org/document/7579514/ (visited on 05/11/2023).

[70] M. Rizzi, M. Lipinski, P. Ferrari, S. Rinaldi, and A. Flammini, «White Rabbit Clock Synchronization: Ultimate Limits on Close-In Phase Noise and Short-Term Stability Due to FPGA Implementation», *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 65, no. 9, pp. 1726–1737, Sep. 2018, ISSN: 0885-3010, 1525-8955. DOI: 10.1109/TUFFC.2018.2851842. [Online]. Available: https://ieeexplore.ieee.org/document/8400550/ (visited on 05/12/2023).

[71] G. Daniluk, *White Rabbit calibration procedure version 1.1*, Nov. 2015. [Online]. Available: https://ohwr.org/project/white-rabbit/uploads/76cdbdbadccc9d6c54d5caf246550fbf/WR_Calibration-v1.1-20151109.pdf (visited on 06/10/2023).

[72] E. F. Dierikx, A. E. Wallin, T. Fordell, *et al.*, «White Rabbit Precision Time Protocol on Long-Distance Fiber Links», *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 63, no. 7, pp. 945–952, Jul. 2016, ISSN: 0885-3010. DOI: 10.1109/TUFFC.2016.2518122. [Online]. Available: http://ieeexplore.ieee.org/document/7383303/ (visited on 05/11/2023).

[73] H. Imlau, *Highly Accurate Time Dissemination and Network Synchronization at ISPCS 2019*, Portland, Oregon, Sep. 2019.

[74] Netnod, *Netnod's White Rabbit implementation achieves sub-nanosecond accuracy in live Swedish network | Netnod*. [Online]. Available: https://www.netnod.se/time-and-frequency/netnods-white-rabbit-implementation-achieves-sub-nanosecond-accuracy-in-live-swedish-network (visited on 05/11/2023).

[75] Namneet Kaur, «Long range time transfer over optical fiber links and cross-comparison with satellite based methods.», Ph.D. dissertation, SYRTE, Observatoire de Paris, Paris, Apr. 2018.

[76] B. L. Cozzens, *White Rabbit makes leap for time over fiber*, en-US, Section: Latest News, Sep. 2021. [Online]. Available: https://www.gpsworld.com/

white-rabbit-makes-leap-for-time-over-fiber/ (visited on 05/11/2023).

[77] ITU-T, *G.694.2 : Spectral grids for WDM applications: CWDM wavelength grid*, Dec. 2003. [Online]. Available: https://www.itu.int/rec/T-REC-G.694.2-200312-I/en (visited on 05/14/2023).

[78] ITU-T, *G.694.1 : Spectral grids for WDM applications: DWDM frequency grid*, Oct. 2020. [Online]. Available: https://www.itu.int/rec/T-REC-G.694.1/en (visited on 05/14/2023).

[79] A. Lord, Y. R. Zhou, R. Jensen, A. Morea, and M. Ruiz, «Evolution from Wavelength-Switched to Flex-Grid Optical Networks», in *Elastic Optical Networks*, V. López and L. Velasco, Eds., Series Title: Optical Networks, Cham: Springer International Publishing, 2016, pp. 7–30, ISBN: 978-3-319-30173-0 978-3-319-30174-7. DOI: 10.1007/978-3-319-30174-7_2. [Online]. Available: http://link.springer.com/10.1007/978-3-319-30174-7_2 (visited on 05/15/2023).

[80] ITU-T, *G.709 : Interfaces for the optical transport network*, 2020. [Online]. Available: https://www.itu.int/rec/T-REC-G.709-202006-I/en (visited on 05/15/2023).

[81] O. Lopez, F. Kéfélian, H. Jiang, *et al.*, «Frequency and time transfer for metrology and beyond using telecommunication network fibres», *Comptes rendus de l'académie des sciences, Physique*, vol. 16, pp. 531–539, 2015, Publisher: arXiv Version Number: 1. DOI: 10.48550/ARXIV.1507.04135. [Online]. Available: https://arxiv.org/abs/1507.04135 (visited on 05/14/2023).

[82] Japan National Institute of Information and Communication Technology, *Demonstration of World Record: 319 Tb/s Transmission over 3,001 km with 4-core optical fiber | 2021*, ja. [Online]. Available: https://www.nict.go.jp/en/press/2021/07/12-1.html (visited on 05/16/2023).

[83] Y. Hong, K. R. H. Bottrill, T. D. Bradley, *et al.*, «Low-Latency WDM Intensity-Modulation and Direct-Detection Transmission Over >100 km Distances in a Hollow Core Fiber», en, *Laser & Photonics Reviews*, vol. 15, no. 9, p. 2 100 102, Sep. 2021, ISSN: 1863-8880, 1863-8899. DOI: 10.1002/lpor.202100102. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1002/lpor.202100102 (visited on 05/16/2023).

[84] G. P. Agrawal, *Fiber-Optic Communication Systems*, en, 4th Edition. Wiley, Oct. 2010, Google-Books-ID: mTH0QTUJ8xIC, ISBN: 978-0-470-91851-7.

[85] N. Sklavos, M. Hübner, D. Goehringer, and P. Kitsos, Eds., *System-Level Design Methodologies for Telecommunication*, en. Cham: Springer International Publishing, 2014, ISBN: 978-3-319-00662-8 978-3-319-00663-5. DOI: 10.1007/978-3-319-00663-5. [Online]. Available: https://link.springer.com/10.1007/978-3-319-00663-5 (visited on 06/10/2023).

[86] G. Keiser, *Fiber Optic Communications*, en. Singapore: Springer, 2021, ISBN: 978-981-334-664-2 978-981-334-665-9. DOI: 10.1007/978-981-33-4665-9. [Online]. Available: https://link.springer.com/10.1007/978-981-33-4665-9 (visited on 05/15/2023).

[87] G. Ghosh, M. Endo, and T. Iwasaki, «Temperature-dependent Sellmeier coefficients and chromatic dispersions for some optical fiber glasses», *Journal of Lightwave Technology*, vol. 12, no. 8, pp. 1338–1342, Aug. 1994, ISSN: 07338724. DOI: 10.1109/50.317500. [Online]. Available: http://ieeexplore.ieee.org/document/317500/ (visited on 05/16/2023).

[88] G. Chauvel, «Dispersion in Optical Fibers», en, Anritsu, Tech. Rep., 2008. [Online]. Available: https://www.semanticscholar.org/paper/Dispersion-in-Optical-Fibers-Chauvel/0a6466906a91b868b3ed01b5b359b8a4e183034c (visited on 05/14/2023).

[89] Bo Xu, «Study of fiber nonlinear effects on fiber optic communication systems», es, Ph.D. dissertation, University of Virginia, 2003. [Online]. Available: https://www.proquest.com/dissertations-theses/study-fiber-nonlinear-effects-on-optic/docview/305300521/se-2?accountid=14542 (visited on 05/14/2023).

[90] Georges Sagnac, «L'éther lumineux démontré par l'effet du vent relatif d'éther dans un interféromètre en rotation uniforme», *Comptes rendus*, no. 157, pp. 708–710, 1913.

[91] J. Gersl, P. Delva, and P. Wolf, «Relativistic corrections for time and frequency transfer in optical fibres», *Metrologia*, vol. 52, no. 4, pp. 552–564, Aug. 2015, ISSN: 0026-1394, 1681-7575. DOI: 10.1088/0026-1394/52/4/552. [Online]. Available: https://iopscience.iop.org/article/10.1088/0026-1394/52/4/552 (visited on 05/16/2023).

[92] B. A. Forouzan and S. C. Fegan, *Data communications and networking* (McGraw-Hill Forouzan networking series), 4th ed. New York: McGraw-Hill Higher Education, 2007, OCLC: ocm62878618, ISBN: 978-0-07-325032-8 978-0-07-296775-3.

[93] N. J. Alpern and R. J. Shimonski, *Eleventh Hour Network+*, en. Elsevier, 2010, ISBN: 978-1-59749-428-1. DOI: 10.1016/C2009-0-20685-4. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/C20090206854 (visited on 05/12/2023).

[94] IEEE, «IEEE Standard for Local and metropolitan area networks: Media Access Control (MAC) Bridges», IEEE, Tech. Rep., 2004, ISBN: 9780738139821. DOI: 10.1109/IEEESTD.2004.94569. [Online]. Available: http://ieeexplore.ieee.org/document/1309630/ (visited on 05/11/2023).

[95] IEEE, «IEEE Standard for Local and Metropolitan Area Network–Bridges and Bridged Networks», IEEE, Tech. Rep., 2018, ISBN: 9781504449298. DOI: 10.1109/IEEESTD.2018.8403927. [Online]. Available: https://ieeexplore.ieee.org/document/8403927/ (visited on 05/11/2023).

[96] D. Allan, P. Ashwood-Smith, N. Bragg, *et al.*, «Shortest path bridging: Efficient control of larger ethernet networks», *IEEE Communications Magazine*, vol. 48, no. 10, pp. 128–135, Oct. 2010, ISSN: 0163-6804. DOI: 10.1109/MCOM.2010.5594687. [Online]. Available: http://ieeexplore.ieee.org/document/5594687/ (visited on 05/11/2023).

[97] IEEE, «IEEE Standard for Local and Metropolitan Area Networks–Link Aggregation», IEEE, Tech. Rep., 2020, ISBN: 9781504464284. DOI: 10.1109/IEEESTD.2020.9105034. [Online]. Available: https://ieeexplore.ieee.org/document/9105034/ (visited on 05/11/2023).

[98] T. Wlostowski, «Precise time and frequency transfer in a White Rabbit network», Master of Science Thesis, Warsaw University of Technology, Warsaw, 2011. [Online]. Available: https://ohwr.org/project/white-rabbit/wikis/Documents/Tom's-Master-thesis (visited on 05/17/2023).

[99] O. Oltu, P. Milea, and A. Simion, «Testing of digital circuitry using Xilinx chipscope logic analyzer», in *CAS 2005 Proceedings. 2005 International Semiconductor Conference, 2005.*, vol. 2, Sinaia, Romania: IEEE, 2005, pp. 471–474, ISBN: 978-0-7803-9214-4. DOI: 10.1109/SMICND.2005.1558829. [Online]. Avail-

able: http://ieeexplore.ieee.org/document/1558829/ (visited on 05/17/2023).

[100] E. Van Der Bij, *White Rabbit - Users of White Rabbit Technology*, en, Apr. 2021. [Online]. Available: https://ohwr.org/project/white-rabbit/wikis/WRUsers (visited on 06/02/2023).

[101] M. Rizzi, *Digital Dual Mixer Time Difference: Phase Noise and Stability*, Feb. 2017. [Online]. Available: https://ohwr.org/project/wr-low-jitter/wikis/Documents/DDMTD-report (visited on 06/02/2023).

[102] M. M. Lipinski, «Methods to Increase Reliability and Ensure Determinism in a White Rabbit Network. Metody zwiększenia niezawodności i zapewnienia determinizmu w sieci White Rabbit», Ph.D. dissertation, Warsaw U. of Tech., 2016. [Online]. Available: https://cds.cern.ch/record/2261452 (visited on 06/02/2023).

[103] E. Dierikx, E. L. English, J. L. Gutiérrez-Rivas, J. López-Jiménez, A. Parsons, and C. Rieck, *WRITE Deliverable D2: Good Practice Guide on Redundant and Resilient Time Transfer*, Nov. 2021. [Online]. Available: https://ohwr.org/project/write/wikis/uploads/5e9ffdc4e94e400a5b1dfdd72fdc6ec9/17IND14_WRITE_D2_20211130.pdf (visited on 06/02/2023).

[104] J. R. Vig, «Quartz crystal resonators and oscillators for frequency control and timing applications. A tutorial», *NASA STI/Recon Technical Report N*, vol. 95, p. 19 519, Aug. 1994, ADS Bibcode: 1994STIN...9519519V. [Online]. Available: https://ui.adsabs.harvard.edu/abs/1994STIN...9519519V (visited on 06/05/2023).

[105] «Study of a Novel OCXO Characterization Based on Oven Supply Currents for Enhanced Holdover Compensation», *proc. of the 22nd European Frequency and Time Forum (EFTF 2008)*, M. Frei, D. Manetti, E. Firouzi, *et al.*, Eds., 2008, Meeting Name: 22nd European Frequency and Time Forum (EFTF 2008).

[106] J. Vig and T. Meeker, «The aging of bulk acoustic wave resonators, filters and oscillators», in *Proceedings of the 45th Annual Symposium on Frequency Control 1991*, Los Angeles, CA, USA: IEEE, 1991, pp. 77–101, ISBN: 978-0-87942-658-3. DOI: 10.1109/FREQ.1991.145888. [Online]. Available: http://ieeexplore.ieee.org/document/145888/ (visited on 06/05/2023).

[107] F. Vernotte, J. Delporte, M. Brunet, and T. Tournier, «Uncertainties of drift coefficients and extrapolation errors: Application to clock error prediction», *Metrologia*, vol. 38, no. 4, pp. 325–342, Aug. 2001, ISSN: 0026-1394. DOI: 10.1088/0026-1394/38/4/6. [Online]. Available: https://iopscience.iop.org/article/10.1088/0026-1394/38/4/6 (visited on 06/05/2023).

[108] G. Busca and Q. Wang, «Time prediction accuracy for a space clock», *Metrologia*, vol. 40, no. 3, S265–S269, Jun. 2003, ISSN: 0026-1394, 1681-7575. DOI: 10.1088/0026-1394/40/3/306. [Online]. Available: https://iopscience.iop.org/article/10.1088/0026-1394/40/3/306 (visited on 06/05/2023).

[109] Y. Shmaliy and L. Arceo-Miquel, «Efficient predictive estimator for holdover in GPS-based clock synchronization», *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 55, no. 10, pp. 2131–2139, Oct. 2008, ISSN: 0885-3010. DOI: 10.1109/TUFFC.913. [Online]. Available: http://ieeexplore.ieee.org/document/4638900/ (visited on 06/05/2023).

[110] G. C. C.W.T. Nicholls, «Adaptive OCXO drift correction algorithm», in *Proceedings of the 2004 IEEE International Frequency Control Symposium and Exposition, 2004.*, Montreal, Canada: IEEE, 2004, pp. 509–517, ISBN: 978-0-7803-8414-9. DOI: 10.1109/FREQ.2004.1418510. [Online]. Available: http://ieeexplore.ieee.org/document/1418510/ (visited on 06/05/2023).

[111] R. Dutra, I. Freire, P. Bemerguy, A. Klautau, I. Almeida, and E. Medeiros, «An LSTM-based Approach for Holdover Clock Disciplining in IEEE 1588 PTP Applications», in *2021 IEEE Global Communications Conference (GLOBECOM)*, Madrid, Spain: IEEE, Dec. 2021, pp. 1–6, ISBN: 978-1-72818-104-2. DOI: 10.1109/GLOBECOM46510.2021.9685776. [Online]. Available: https://ieeexplore.ieee.org/document/9685776/ (visited on 06/05/2023).

[112] ITU-T, *G.8272 : Timing characteristics of primary reference time clocks*, Nov. 2018. [Online]. Available: https://www.itu.int/rec/T-REC-G.8272/en (visited on 06/05/2023).

[113] C. Hauser, D. Bakken, and A. Bose, «A failure to communicate: Next generation communication requirements, technologies, and architecture for the electric power grid», *IEEE Power and Energy Magazine*, vol. 3, no. 2, pp. 47–55, Mar. 2005, ISSN: 1540-7977. DOI: 10.1109/MPAE.2005.1405870. [Online]. Avail-

able: http://ieeexplore.ieee.org/document/1405870/ (visited on 06/10/2023).

[114] «IEEE Guide for Synchronization, Calibration, Testing, and Installation of Phasor Measurement Units (PMUs) for Power System Protection and Control», IEEE, Tech. Rep., ISBN: 9780738182957. DOI: 10.1109/IEEESTD.2013.6475134. [Online]. Available: http://ieeexplore.ieee.org/document/6475134/ (visited on 06/10/2023).

[115] «IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks», IEEE, Tech. Rep., ISBN: 9780738136639. DOI: 10.1109/IEEESTD.2003.94280. [Online]. Available: https://ieeexplore.ieee.org/document/1203093/ (visited on 06/10/2023).

[116] H. Kirrmann, K. Weber, O. Kleineberg, and H. Weibel, «HSR: Zero recovery time and low-cost redundancy for Industrial Ethernet (High availability seamless redundancy, IEC 62439-3)», in *2009 IEEE Conference on Emerging Technologies & Factory Automation*, Palma de Mallorca, Spain: IEEE, Sep. 2009, pp. 1–4, ISBN: 978-1-4244-2727-7. DOI: 10.1109/ETFA.2009.5347037. [Online]. Available: http://ieeexplore.ieee.org/document/5347037/ (visited on 06/10/2023).

[117] H. Kirrmann, K. Weber, O. Kleineberg, and H. Weibel, «Seamless and low-cost redundancy for substation automation systems (high availability seamless redundancy, HSR)», in *2011 IEEE Power and Energy Society General Meeting*, San Diego, CA: IEEE, Jul. 2011, pp. 1–7, ISBN: 978-1-4577-1000-1 978-1-4577-1002-5. DOI: 10.1109/PES.2011.6038906. [Online]. Available: https://ieeexplore.ieee.org/document/6038906/ (visited on 06/10/2023).

[118] Texas Instruments, «High-Availability Seamless Redundancy (HSR) Ethernet for Substation Automation», en, 2016. [Online]. Available: http://www.ti.com/lit/ug/tidub08/tidub08.pdf (visited on 06/10/2023).

[119] *HSR/PRP Switches and RedBoxes - Flexibilis*. [Online]. Available: https://www.flexibilis.com/products/hsrprp-switches/ (visited on 06/10/2023).

[120] SoC-e, *HPS - HSR-PRP Switch IP Core*, en-US. [Online]. Available: https://soc-e.com/products/hsr-prp-switch-ip-core-all-hardware-low-latency-switch-for-fpgas/ (visited on 06/10/2023).

[121] J. L. Gutiérrez-Rivas, J. López-Jiménez, E. Ros, and J. Díaz, «White Rabbit HSR: A Seamless Subnanosecond Redundant Timing System With Low-Latency Data Capabilities for the Smart Grid», *IEEE Transactions on Industrial Informatics*, vol. 14, no. 8, pp. 3486–3494, Aug. 2018, Conference Name: IEEE Transactions on Industrial Informatics, ISSN: 1941-0050. DOI: 10.1109/TII.2017.2779240.

[122] *IEC 61850-3:2013 Communication networks and systems for power utility automation - Part 3: General requirements*. [Online]. Available: https://webstore.iec.ch/publication/6010 (visited on 06/10/2023).

[123] *IEC 61850-5:2013 Communication networks and systems for power utility automation - Part 5: Communication requirements for functions and device models*. [Online]. Available: https://webstore.iec.ch/publication/6012 (visited on 06/10/2023).

[124] J. L. Gutierrez-Rivas, F. Torres-Gonzalez, E. Ros, and J. Diaz, «Enhancing White Rabbit Synchronization Stability and Scalability Using P2P Transparent and Hybrid Clocks», *IEEE Transactions on Industrial Informatics*, vol. 17, no. 11, pp. 7316–7324, Nov. 2021, ISSN: 1551-3203, 1941-0050. DOI: 10.1109/TII.2021.3054365. [Online]. Available: https://ieeexplore.ieee.org/document/9335233/ (visited on 06/10/2023).

[125] N. Christodoulia, S. N. Lea, G. Marra, *et al.*, «Mode-locked lasers for remote intercomparison of frequency standards over optical fibre networks.», en, NPL, UK, 2007, pp. 159–162, ISBN: 978-0-946754-49-6. [Online]. Available: https://eprintspublications.npl.co.uk/3973/ (visited on 05/14/2023).

[126] O. Terra, G. Grosche, and H. Schnatz, «Brillouin amplification in phase coherent transfer of optical frequencies over 480 km fiber», en, *Optics Express*, vol. 18, no. 15, p. 16 102, Jul. 2010, ISSN: 1094-4087. DOI: 10.1364/OE.18.016102. [Online]. Available: https://opg.optica.org/oe/abstract.cfm?uri=oe-18-15-16102 (visited on 05/14/2023).

[127] K. Predehl, G. Grosche, S. M. F. Raupach, *et al.*, «A 920-Kilometer Optical Fiber Link for Frequency Metrology at the 19th Decimal Place», en, *Science*, vol. 336, no. 6080, pp. 441–444, Apr. 2012, ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1218442. [Online]. Available: https://www.science.org/doi/10.1126/science.1218442 (visited on 05/14/2023).

[128] N. Kaur, «Transfert de temps à longue distance utilisant des liaisons à fibre optique et comparaison croisée avec des méthodes par satelliteires», These de doctorat, Paris Sciences et Lettres (ComUE), Apr. 2018. [Online]. Available: https://www.theses.fr/2018PSLEO002 (visited on 05/14/2023).

[129] P. Boven, «DWDM stabilized optics for white rabbit», in *2018 European Frequency and Time Forum (EFTF)*, Turin: IEEE, Apr. 2018, pp. 213–216, ISBN: 978-1-5386-5620-4. DOI: 10.1109/EFTF.2018.8409035. [Online]. Available: https://ieeexplore.ieee.org/document/8409035/ (visited on 05/14/2023).

[130] M. Banks, «Go ahead for the SKA radio telescope», *Physics World*, vol. 34, no. 8, pp. 8–8, Sep. 2021, ISSN: 0953-8585, 2058-7058. DOI: 10.1088/2058-7058/34/08/10. [Online]. Available: https://iopscience.iop.org/article/10.1088/2058-7058/34/08/10 (visited on 05/14/2023).

[131] C.-h. Tsai and J.-c. Tsai, «MEMS optical switches and interconnects», en, *Displays*, vol. 37, pp. 33–40, Apr. 2015, ISSN: 01419382. DOI: 10.1016/j.displa.2014.11.007. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0141938214000924 (visited on 05/15/2023).

[132] P. Wall, P. Colbourne, C. Reimer, and S. McLaughlin, «WSS Switching Engine Technologies», in *OFC/NFOEC 2008 - 2008 Conference on Optical Fiber Communication/National Fiber Optic Engineers Conference*, San Diego, CA, USA: IEEE, Feb. 2008, pp. 1–5. DOI: 10.1109/OFC.2008.4528672. [Online]. Available: http://ieeexplore.ieee.org/document/4528672/ (visited on 05/15/2023).

[133] J. Radil and J. Vojtěch, «Stimulated Brillouin Scattering and Raman Amplification in Standard Telco Fibers for Metrology Applications», in *2022 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS)*, ISSN: 2327-1949, Apr. 2022, pp. 1–3. DOI: 10.1109/EFTF/IFCS54560.2022.9850605.

[134] L. Grüner-Nielsen, S. N. Knudsen, B. Edvold, *et al.*, «Dispersion Compensating Fibers», en, *Optical Fiber Technology*, vol. 6, no. 2, pp. 164–180, Apr. 2000, ISSN: 10685200. DOI: 10.1006/ofte.1999.0324. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1068520099903243 (visited on 05/15/2023).

[135] J. W. Milnor and G. A. Randall, «The Newfoundland-Azores High-Speed Duplex Cable», *Transactions of the American Institute of Electrical Engineers*, vol. 50, no. 2, pp. 389–396, Jun. 1931, Conference Name: Transactions of the American Institute of Electrical Engineers, ISSN: 2330-9431. DOI: 10.1109/T-AIEE.1931.5055804.

[136] R. Kissell, *Algorithmic trading methods: applications using advanced statistics, optimization, and machine learning techniques*, 2nd ed. San Diego: Elsevier Inc, 2020, ISBN: 978-0-12-815630-8.

[137] Finextra, *Fixnetix microwaves London to Frankfurt route*, en, Jan. 2013. [Online]. Available: https://www.finextra.com/pressarticle/48135/fixnetix-microwaves-london-to-frankfurt-route (visited on 05/25/2023).

[138] A. Lohr, *Precise Timing in Financial Markets - Time Distribution in Deutsche Börse's T7 Trading Network*, en, CS 349F Technology for Financial Systems, Stanford, Oct. 2020. [Online]. Available: https://web.stanford.edu/class/cs349f/slides/CS349F_lec7_deutsche_precise_timing_in_finance.pdf (visited on 05/25/2023).

[139] United States Securities and Exchange Commission, *17 CFR Part 242 Concept Release on Equity Market Structure*, 2010. [Online]. Available: https://www.sec.gov/rules/concept/2010/34-61358.pdf (visited on 05/25/2023).

[140] A. A. Kirilenko, A. S. Kyle, M. Samadi, and T. Tuzun, «The Flash Crash: The Impact of High Frequency Trading on an Electronic Market», en, *SSRN Electronic Journal*, 2011, ISSN: 1556-5068. DOI: 10.2139/ssrn.1686004. [Online]. Available: http://www.ssrn.com/abstract=1686004 (visited on 05/25/2023).

[141] M. Baron, J. Brogaard, B. Hagströmer, and A. Kirilenko, «Risk and Return in High-Frequency Trading», en, *Journal of Financial and Quantitative Analysis*, vol. 54, no. 3, pp. 993–1024, Jun. 2019, ISSN: 0022-1090, 1756-6916. DOI: 10.1017/S0022109018001096. [Online]. Available: https://www.cambridge.org/core/product/identifier/S0022109018001096/type/journal_article (visited on 05/25/2023).

[142] A. Gordon, *Proposed SEC Rules Aim to Rein in High-Frequency Traders*, Apr. 2022. [Online]. Available: https://www.ai-cio.com/news/proposed-sec-

rules-aim-to-reign-in-high-frequency-traders/ (visited on 05/25/2023).

[143] *Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments and amending Directive 2002/92/EC and Directive 2011/61/EU (recast) Text with EEA relevance*, en, May 2014. [Online]. Available: http://data.europa.eu/eli/dir/2014/65/oj/eng (visited on 05/25/2023).

[144] *Regulation (EU) No 600/2014 of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments and amending Regulation (EU) No 648/2012 Text with EEA relevance*, en, Legislative Body: EP, CONSIL, May 2014. [Online]. Available: http://data.europa.eu/eli/reg/2014/600/oj/eng (visited on 05/25/2023).

[145] European Commission, *Annex to the Commission Delegated Regulation supplementing Directive 2014/65/EU of the European Parliament and of the Council with regard to regulatory technical standards for the level of accuracy of business clocks*, Jun. 2016. [Online]. Available: https://ec.europa.eu/finance/securities/docs/isd/mifid/rts/160607-rts-25-annex_en.pdf (visited on 05/25/2023).

[146] L. Lobo and P. Whibberley, *Time Traceability for the Finance Sector Fact Sheet*, 2016. [Online]. Available: https://www.esma.europa.eu/sites/default/files/consultations/2016/03/esma_cp_tr_ork_cs_bipm-npl_annex1.pdf (visited on 06/01/2023).

[147] R. Hunt, *Keeping Time With Amazon Time Sync Service | AWS News Blog*, en-US, Section: Amazon EC2, Nov. 2017. [Online]. Available: https://aws.amazon.com/blogs/aws/keeping-time-with-amazon-time-sync-service/ (visited on 06/01/2023).

[148] C. MacKenzie and R. Ramakrishnan, *What's on the Horizon for Time as a Service (TaaS) Solutions?*, en, Jan. 2023. [Online]. Available: https://blog.equinix.com/blog/2023/01/17/whats-on-the-horizon-for-time-as-a-service-taas-solutions/ (visited on 06/01/2023).

[149] D. Hicks, P. Whibberley, E. L. English, L. Lobo, T. Lee, and A. Austin, «NPLTime - UTC Traceable Time for the Financial Sector», Monterey, California, Feb. 2016, pp. 167–173. DOI: 10.33012/2016.13160. [Online].

Available: https://www.ion.org/publications/abstract.cfm?articleID=13160 (visited on 06/01/2023).

[150] *Hoptroff Traceable Time as a Service - Hoptroff Smart Timing*, en-GB. [Online]. Available: https://www.hoptroff.com/product-traceable-time-as-a-service-ttaas (visited on 06/01/2023).

[151] Equinix, Inc., *Equinix Precision Time Datasheet - Precise, Reliable, Secure Time as a Service*, 2022. [Online]. Available: https://www.equinix.com/content/dam/eqxcorp/en_us/documents/resources/whitepapers/wp_equinix_precision_time_en.pdf (visited on 06/01/2023).

[152] GTT Communications, Inc., *White Paper: Accurate, Traceable, and Verifiable Time Synchronization for World Financial Markets*, 2017. [Online]. Available: https://www.gtt.net/us-en/resources/whitepapers/accurate-traceable-and-verifiable-time-synchronization-for-world-financial-markets/ (visited on 06/01/2023).

[153] Inside GNSS, *GMV Uses GPS and Galileo to Provide Robust Timing for the Financial Sector*, en-US, Jun. 2019. [Online]. Available: https://insidegnss.com/gmv-uses-gps-and-galileo-to-provide-robust-timing-for-the-financial-sector/ (visited on 06/01/2023).