


Handbook on Selection and use of precise frequency and time systems

Edition of 2026



Handbook on Selection and use of precise frequency and time systems

Edition of 2026

Radiocommunication Bureau



© ITU 2026

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

Preface

This new version of the Handbook on Selection and Use of Precise Frequency and Time Systems has been developed by experts of ITU-R Working Party 7A, and invited experts, under the chairmanship of Dr Joseph Achkar (France) of ITU-R Radiocommunication Study Group 7 (Science Services). To carry out this work, a correspondence group was created, led by Felicitas Arias (France) and Andreas Bauch (Germany).

The Handbook consists of 10 Chapters which describe basic concepts, frequency and time sources and their use, measurement techniques, time and frequency dissemination and time scales, optical frequency standards, time as a service, and future prospects. The contents include detailed explanations and many references that can be consulted for additional details.

The technical content of this Handbook is intended for use by administrations of both developing and developed countries and by the Radiocommunication Bureau. The Handbook will also be useful to engineers in scientific and industrial organizations.

Mario Maniewicz

Director, Radiocommunication Bureau

Acknowledgements

The work of the Rapporteurs for various sections of the Handbook was outstanding, and we should like to thank the following administrations and organizations for their kind support and participation on the part of their experts:

France, Germany (Federal Republic of), United States of America, United Kingdom, Japan, Laboratoire Temps Espace, Observatoire de Paris (France), Physikalisch-Technische Bundesanstalt (Germany), University of Colorado (USA), Franche-Comté Electronique Mécanique Thermique et Optique – Sciences et Technologies (France), National Physical Laboratory (UK), NICT (Japan), Centre national d’études spatiales (France), Bureau international des poids et mesures (BIPM) and Conseil européen pour la recherche nucléaire (CERN), and to acknowledge the contributions and the helpful guidance and discussions with the following personalities:

Chapter 1	Judah Levine	University of Colorado, USA
	Andreas Bauch	Physikalisch-Technische Bundesanstalt, Germany
Chapter 2	Patrizia Tavella, Giulio Tagliaferro, Gianna Panfilo	Bureau international des poids et mesures, France
Chapter 3	Andreas Bauch	Physikalisch-Technische Bundesanstalt, Germany
	Judah Levine	University of Colorado, USA
Chapter 4		
§§ 4.1.1 to 4.1.3	Judah Levine	University of Colorado, USA
§ 4.1.4	Enrico Rubiola	Franche-Comté Electronique Mécanique Thermique et Optique – Sciences et Technologies, France
§ 4.2.1	Jean-Michel Friedt	Franche-Comté Electronique Mécanique Thermique et Optique – Sciences et Technologies, France
	Joseph Achkar	Laboratoire Temps Espace, Observatoire de Paris, France
§ 4.2.2	Paul-Eric Pottie	Laboratoire Temps Espace, Observatoire de Paris, France
Chapter 5	Felicitas Arias	Laboratoire Temps Espace, Observatoire de Paris, France
	Judah Levine	University of Colorado, USA
	Gianna Panfilo	Bureau international des poids et mesures, France

Chapter 6

§§ 6.1 to 6.3.1	Andreas Bauch	Physikalisch-Technische Bundesanstalt, Germany
	Peter Whibberley (retired)	National Physical Laboratory, UK
§§ 6.3.2 to 6.4.1	Maciej Lipinski	Conseil européen pour la recherche nucléaire, Suisse
§§ 6.4.2 to 6.4.5	Paul-Eric Pottie	Laboratoire Temps Espace, Observatoire de Paris, France
Chapter 7	Judah Levine	University of Colorado, USA
	Andreas Bauch	Physikalisch-Technische Bundesanstalt, Germany
	Jérôme Delporte	Centre national d'études spatiales, France
Chapter 8	Andreas Bauch	Physikalisch-Technische Bundesanstalt, Germany
Chapter 9	Tatsuya Ido	National Institute of Information and Communications Technology, Japan
	Jérôme Lodewick	Laboratoire Temps Espace, Observatoire de Paris, France
Chapter 10	Judah Levine	University of Colorado, USA

Our special thanks to the Chair of WP 7A Joseph Achkar for his leadership in this project.

Markus Dries
Chair, Study Group 7

CHAPTER 1

Introduction and basic concepts

TABLE OF CONTENTS

	<i>Page</i>
1.1 A short note on terms and vocabulary.....	2
1.2 Historical sketch.....	2
1.2.1 Frequency standards	2
1.2.2 Radio time dissemination	4
1.3 Basic principles of frequency standards.....	5
1.3.1 Quartz crystal frequency standards.....	5
1.3.2 Microelectromechanical Systems (MEMS).....	5
1.3.3 Atomic frequency standards	5
1.4 Basic metrological concepts.....	8
1.4.1 Frequency stability	8
1.4.2 Accuracy	8
References	9

1.1 A short note on terms and vocabulary

This Handbook primarily addresses users in the Radiocommunication and the Telecommunication Standardization Sector of the International Telecommunication Union. They represent a very important and large user community. Sources for standard frequency, one pulse-per-second (1pps) ticks, and time-of-day information are in abundant use, but their performance and even their operations are often described in ITU documents with terms that are not in all cases unambiguous and concordant with metrological practice. The Glossary, published as Recommendation ITU-R TF.686, contains some of the terms, but some nevertheless remain obscure. Among those are the terms “synchronization”, or “synchronized”.

Many applications require that frequency standards distributed in a system share a common rate (frequency). The signals may not carry any signature of “time” on them. Still the sources are often termed “synchronized”. Other applications require that 1pps signals of remote devices in a system tick simultaneously, but not necessarily in accordance with any internationally recognized time scale. This could be termed “internal synchronization”. On the other hand, Coordinated Universal Time UTC has been chosen by the global scientific community as a reference time scale, for example for comparisons of official time scales between national metrology laboratories (NMI) and to create national legal time scales. The use of UTC in the area of telecommunications is prescribed in the Radio Regulations (No. 1.14). “Global synchronization” could be the proper term to express that those clocks in a system tick “simultaneously” with reference to UTC and with a common time-of-day label attached. Clearly, uncertainty values need to be determined in all cases to make “simultaneous” a meaningful word.

Referring to the Glossary, a “clock” is “a device for time measurement and/or time display”. Through its output signals of some periodicity (1pps or standard frequency) it may provide the ruler for measurements of time interval to be made by a suitable instrument. “Clocks” in common (engineering) parlance, however, may not even have an option to enter and keep time-of-day information and later to display it. Nevertheless, “atomic clock” will be found in a wide sense in ITU documents and throughout this Handbook.

In the following, the distinction between “frequency standard” and “clock” shall be carried forward in the historical sketch which is divided in one part on frequency standards and another part on current dissemination services that fall under the realm of ITU. A first cursory description of available frequency standards follows, but scientific details and quantitative assessment follow in Chapter 3. The complementary Chapter 6 provides then details on the current dissemination services. Chapter 3 is preceded by a chapter dealing with (quantitative) characterization of frequency and time signals. It has been mostly inherited from the respective part of the existing Handbook as there have been no fundamental changes in measures and expressions of uncertainty and instability of signals.

1.2 Historical sketch

1.2.1 Frequency standards

The earliest frequency standards that took advantage of the development of electronic technology were based on an ensemble of precision inductors and variable capacitors that were connected in parallel to realize a resonant circuit. The National Bureau of Standards “wavemeter” [Cox, 1922] was constructed in this way. The quoted fractional accuracy was about 0.3%, and the resonant frequency could be adjusted from 3.5 kHz to 4.6 MHz.

Broadcast radio stations were assigned frequencies separated by 10 kHz. The transmission frequencies were of order 1 MHz and, by 1928, the stations were required to maintain the transmitted frequency within 500 Hz of the assigned value. The required fractional frequency accuracy was 0.05%, which could not be reliably realized with the L-C wavemeter design. The frequency references based on parallel L-C circuits were replaced by temperature-controlled quartz crystal oscillators starting in about 1927 [Cady, 1923]. The accuracy of these oscillators was several parts per million in uncertainty and improved to two parts in 10^{10} by 1931.

Quartz crystal oscillators rely upon the discovery of piezoelectricity by P. Curie in 1880 (Nobel prize 1903) and upon the invention of the first electronic amplifier (the triode) by Lee de Forest in 1907. Although the oscillation frequency is primarily determined by a mechanical resonance in the quartz crystal, the

characteristics of the electronic circuit also affect the frequency, and the earliest oscillators were based on a standard circuit. Neither the mechanical properties of the quartz nor the electrical properties of the circuit were easily reproducible, so that each device had to be individually calibrated.

Quartz crystal oscillators were widely used in communication systems and in frequency and time metrology laboratories starting in the 1920s. In the 1930s, in Germany, stationary quartz oscillators were perfected to the extent that they were instrumental to observe the seasonal variation of Earth rotation [Scheibe und Adelsberger, 1950]. This was possible by comparing free-running quartz oscillators to received time signals that were broadcast by terrestrial stations whose time marks were regulated based on astronomical observations.

In the years following, quartz oscillators were widely used during World War II and, since that time, a large amount of effort has been devoted to improving the resonator design, to optimizing the behaviour of the associated electronic circuit and to understanding the physical origin of the short-term frequency instabilities, the sensitivity to external factors and the ageing processes [Besson, 1977; Filler and Vig, 1993; Gerber and Ballato, 1985; Vig, 1991; Walls *et al.*, 1992; Walls and Gagnepain, 1992]. In the field of interest of this Handbook, high-quality quartz crystal oscillators are used as secondary frequency and time standards. Furthermore, they are present in most atomic frequency standards, where their frequency is controlled by the atomic resonance and where they are the physical source of the output signals.

The idea of using the resonance properties of an ensemble of isolated atoms (i.e. in a vapour) to realize a frequency standard and a clock is attributed to I. I. Rabi, proposed in 1939 (Nobel prize 1944). It builds on the fact that atoms are indiscernible particles which share identical properties wherever and whenever they are used as references. Here is the distinct difference to macroscopic references, such as the mechanical pendulum, oscillating circuits and even quartz resonators. Over the decades, research has strived to transfer the atomic properties in a more and more perfect manner without disturbance to a signal generator that provides the output of the frequency standard.

The first “atomic clock” (see notes above) was built at the National Bureau of Standards (NBS) during 1948–49, using a microwave absorption line of ammonia at a frequency of approximately 23.9 GHz. The natural linewidth of the transition was broadened by collisions of ammonia molecules in the gas cell (pressure broadening), by the thermal motion of the molecules (Doppler broadening) and by other effects, and the effective line quality factor Q (resonance frequency divided by linewidth) was about 50 000. Given these limitations, realizing the stability implied by the natural linewidth of a suitable transition would be possible if the interaction between the atoms and the probing field could be realized in a beam geometry rather than in a cell, and if the Doppler broadening could be reduced by cooling the atoms or by some other method, such as saturated absorption [Barger and Hall, 1969].

The most important discoveries which determine the structure of present operational caesium, rubidium and hydrogen atomic frequency standards are: a) the two separated field method of inducing a magnetic resonance and the hydrogen maser by N. F. Ramsey in 1950 and 1960, respectively (Nobel prize 1989); b) the optical pumping by A. Kastler in 1950 (Nobel prize 1966); and c) the maser effect in the ammonia maser by C. H. Townes, N. G. Basov and A. M. Prokhorov in 1955 (Nobel prize 1964). The physical principles are extensively laid down in [Vanier and Audoin, 1989], and the latest developments have been summarized in a complementing issue [Vanier and Tomescu, 2016].

One of the main objectives in the further development of atomic clocks has always been to extend the time during which the atoms interact with the probing field and to reduce thus the linewidth. Thereby, at the same time, a reduced frequency instability and a reduced frequency uncertainty could be achieved. When laser cooling was discovered and understood in the middle of the 1980s (Nobel Prizes 1997) [Chu, 1998; Cohen-Tannoudji, 1998; Philips, 1998], its importance for clock development was recognized immediately. “Laser cooling” is the common name for manipulating atomic eigenstates and the atoms’ velocity at the same time. It produces “cold” atoms with a thermal residual velocity of about 1 cm/s which corresponds to a temperature in the μK regime, starting out from an ensemble of atoms at room temperature.

The first caesium beam frequency standard was built at NBS in the early 1950s, and the first device of this kind began operation at the National Physics Laboratory (NPL, UK) in June 1955 [Essen and Parry, 1957]. The first commercial units were available around 1958. Soon afterwards, at the beginning of the 1960s, the production of rubidium frequency standards and of hydrogen masers also started.

The NPL caesium beam device was used between 1956 and 1958 to measure the unperturbed hyperfine transition frequency of the caesium 133 atom in the ground state in terms of the Ephemeris second defined as the time unit in these days. The result is included in the present definition of the unit of time as it was decided in 1967 by the Thirteenth Conférence générale des poids et mesures (CGPM): “The second is the duration of 9 192 631 770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the caesium 133 atom”. Following the decision of the 26th CGPM, the text of the definition is now: “The second, symbol *s*, is the SI unit of time. It is defined by taking the fixed numerical value of the caesium frequency $\Delta\nu_{Cs}$, the unperturbed ground-state hyperfine transition frequency of the caesium 133 atom, to be 9 192 631 770 when expressed in the unit Hz, which is equal to s^{-1} ” [BIPM, 2019]. This definition implies, as before, the exact relation $\Delta\nu_{Cs} = 9\,192\,631\,770$ Hz.

Laboratory caesium beam frequency standards have been constructed by several national metrology institutes and timing laboratories. Following the pioneering work of NPL, thermal beam devices provided the most accurate realization of the SI second until about 1996 [Vanier and Audoin, 2005]. At the time of writing in 2025, two devices of this generation are still in operation [Bauch, 2005]. The use of laser-cooling in this context was pioneered in France with the first so-called “caesium fountain” primary frequency standard, FO-1, being presented in 1995 [Clairon *et al.*, 1995]. The “fountain” configuration is advantageous for several reasons. The atoms are moving much more slowly than in a conventional thermal-beam design, and the interaction time between the atoms and the microwave field is correspondingly longer. The configuration also reduces the impact of several systematic frequency offsets. The result is - today, after 30 years of great effort - a combined Type A and Type B uncertainty of about 2×10^{-16} [current results available from the BIPM website, www.bipm.org BIPM Time Department Data Base and from the monthly Circular T]. The U.S. Naval Observatory (USNO) developed continuously operating rubidium atomic fountain clocks that have contributed to EAL (Echelle Atomique Libre) for a number of years [Peil *et al.*, 2016].

For many years, research has been pursued to develop atomic frequency standards based on an optical transition, either in ions stored in an electric trap, following W. Paul and H. G. Dehmelt (Nobel Prize 1989) [Wineland *et al.*, 1990] or in neutral atoms trapped in standing light-fields. Laser-cooling is indispensable in both approaches. Two very promising optical frequency standards are based on neutral Sr in an optical lattice [Nicholson *et al.*, 2015] and Yb^+ in a trap [Huntemann *et al.*, 2016]. An extensive summary of the status was provided in [Ludlow *et al.*, 2015] and is dealt with in Chapter 9. Considerable efforts are ongoing to prepare a decision of a future CGPM by which an optical transition is going to replace the hyperfine transition in caesium as the basis for the definition of the second [Dimarcq *et al.*, 2023].

The frequency stability and the reliability of manufactured atomic frequency standards using a resonance in the caesium, rubidium or hydrogen atom have also been considerably improved [Cutler, 2005]. Furthermore, in certain applications their volume has been reduced. The smallest device to date is the chip-scale atomic clock (CSAC) [Kitching *et al.*, 2000; Kitching, 2018]. Although this type of device is neither as stable nor as accurate as more conventional beam-tube designs, it is the configuration of choice whenever minimizing size, weight and power consumption are important design requirements.

1.2.2 Radio time dissemination

Already at the beginning of the last century, radio waves were recognized as being a suitable means of time transmission. The first tests regarding wireless transmission of time were carried out in 1903 by the United States Naval Observatory in Washington and in 1904 by the Geodetic Institute Potsdam. Similarly, scientists at the National Bureau of Standards (NBS) as it was named then reported on research into radio frequency propagation in 1905. The call sign WWV was assigned in 1919 [Sullivan, 2001]. The most well-known transmission tower for radio time was the Paris Eiffel-tower for some years where transmission started in 1910. From 1910 to 1916, the coastal radio station Norddeich already transmitted – as the first German time service – regular time signals which were based on time determinations of the Imperial Marine Observatory Wilhelmshaven. Then, starting in 1917, the large transmitting station Nauen (close to Berlin) emitted a time signal twice a day on long wave (3 900 m, ≈ 77 kHz). After these beginnings of wireless time transmission, a worldwide network of time signal stations has been set up in the course of time in different wavelength ranges – mainly on myriametric wave, long wave (corresponds to low frequency) and short wave. Monitoring and coordination of the transmission was in the hands of the Bureau international de l’heure (BIH), founded in 1911. In the years following 1960, transmission of coded time-of-day information on long-wave emissions

became available in the US, UK, Germany, Switzerland, Japan and recently in China. The characteristics of the transmissions are documented in Chapter 6 of this Handbook and are regularly updated in Recommendations ITU-R TF.583 and ITU-R TF.768.

In the past decades, the dissemination of time with terrestrial transmitters has been complemented and to some extent replaced by time transmission via satellites, subject of Handbook “Satellite time and frequency transfer and dissemination” [ITU, 2010] and Chapters 6 and 7. The evolution of Global Navigation Satellite Systems has also opened completely new possibilities for intercontinental precision time comparisons, position determination and global navigation.

1.3 Basic principles of frequency standards

In this section, a brief overview is given on the principle of the operation of the available frequency standards as well as the basic metrological concepts for their characterization, including stability and accuracy. A more detailed description of the devices including the characteristics of their performance is given in Chapter 3.

1.3.1 Quartz crystal frequency standards

A quartz crystal frequency standard is an electronic oscillator in which the main frequency determining element is a resonator made of a quartz crystal. The inductance and capacitance of the electronic circuit that excites the vibration of the quartz crystal also affects the oscillator frequency. The resonance frequency of the crystal itself is mainly determined by the macroscopic properties of the bulk material. It is therefore also dependent on all environmental influences that modify these dimensions and properties and can in practice not be deduced from fundamental properties of individual atoms. Therefore, the frequency of quartz crystal oscillators needs to be calibrated with respect to that of a more accurate frequency source. This calibration often includes the complete circuit in addition to the crystal itself for the reason mentioned above. High-performance units are widely used in frequency and time metrology laboratories as secondary frequency standards. Furthermore, the output signal of most atomic frequency standards is derived from a voltage-controlled crystal oscillator (VCXO), its frequency being controlled by the atomic resonance. A very important industrial application has evolved since the deactivation of Selective Availability in the US Global Positioning System GPS in 2000, the GPS- (or GNSS-) Disciplined Oscillator. For better hold-over characteristics, GNSS-DOs are also produced with a CSAC or a Rubidium gas-cell frequency standard as local oscillator (see Chapter 3).

1.3.2 Microelectromechanical Systems (MEMS)

The frequency reference of an electronic oscillator can also be based on a MEMS device, which is a mechanical system about 0.1 mm in size and generally made from silicon. The device exhibits a mechanical resonance, which can be sensed as a variable capacitance at the resonant frequency. Oscillators based on MEMS devices show great promise, especially when size and power consumption are important criteria. See, for example [Barrow *et al.*, 2013] and [Rocheleau *et al.*, 2013]

1.3.3 Atomic frequency standards

In contrast to quartz oscillators, the reference frequency in atomic frequency standards is mainly determined by the intrinsic properties of atoms of specifically selected elements, in other words, by fundamental constants which result from the basic interactions between elementary particles.

In our present state of knowledge in physics and astronomy, we postulate that the atomic properties are fixed and do not depend on space and time (within known relativistic effects). Therefore, it is possible to build and to distribute equipment which delivers the same frequency to a number of locations, provided that a given transition in a given element is observed, and that the relativistic effects (related to the local gravitational potential with respect to the geoid, for instance) are taken into account. An atomic reference frequency is then available locally in real time. It remains constant with time and equal to that of other atomic frequency standards situated at different places on the Earth and its close vicinity (within uncertainties discussed in Chapter 3). In the practical work, relying on commercial products, we are (at the time of writing in 2025) dealing almost entirely with three types of standards which make use of the properties of the elements hydrogen (the hydrogen maser), rubidium (the rubidium gas cell frequency standard) and caesium (the caesium beam

frequency standard). Standards based on an optical transition are used in a few laboratories. These will be discussed subsequently where future developments are dealt with in Chapter 3 and in Chapter 9.

1.3.3.1 Spectroscopic properties of interest

Let us consider an atom of a given element. It has quantum levels of well-defined energy. Let E_1 and E_2 be the energy of two of them, with $E_2 > E_1$. A transition between these two levels may occur under the effect of an electromagnetic radiation of frequency ν_0 . Energy conservation determines the value of ν_0 . This leads to Bohr's relation:

$$h\nu_0 = E_2 - E_1 \quad (1-1)$$

where h is Planck's constant. Therefore, a resonance may be observed. This resonance has a width $\Delta\nu$, which is given by Heisenberg's uncertainty relation:

$$\Delta\nu \Delta t \geq 1 \quad (1-2)$$

In general, the natural lifetime of the clock state is very long, so, Δt is determined by the duration of the interaction between the atomic species and the exciting fields. It may be limited by various physical phenomena or processes such as transit times and relaxation, and it ranges from approximately 1 ms in a rubidium cell and in a caesium thermal beam tube to 1 s in a hydrogen maser or an atomic fountain configuration. With the values of ν_0 involved (see Table 1) and the linewidth originating from these processes, the atomic line quality factor Q , $Q = \nu_0/\Delta\nu$, typically ranges between 10^7 and 10^{10} . The Rb hyperfine transition observed in a fountain configuration may nowadays serve as secondary representation of the second [Riehle *et al.*, 2018], and the accurate value stated in Table 1 was agreed upon in 2017. This accuracy is not accessible in a conventional rubidium gas cell frequency standard (see Chapter 3). In Table 1, the 1- σ uncertainty of the transition frequency is given in brackets referring to the last significant digits.

TABLE 1
Hyperfine transition frequency ν_0 of Hydrogen, Rubidium 87 and Caesium

Atom	Atomic mass	Hyperfine transition frequency (Hz)
H	1	1 420 405 751,770 (3)
Rb	87	6 834 682 610.904 312 6 (41)
Cs	133	9 192 631 770 ⁽¹⁾

⁽¹⁾ By definition of the second.

In currently common atomic frequency standards, the two levels of interest are determined by magnetic interaction inside the atom. It occurs between the magnetic moment of the unpaired electron of the alkali or alkali-like atoms in the ground state of its electron configuration and the magnetic moment of the nucleus. It is called the hyperfine interaction. This interaction is weak, and it leads to a small value of $E_2 - E_1$. Consequently, the resonance frequency ν_0 lies in the microwave frequency range, as shown in Table 1. As a practical consequence, the electronic system which controls the frequency of the associated quartz crystal oscillator can be made efficient, small and reliable.


1.3.3.2 Passive and active atomic frequency standards

The caesium beam and the rubidium cell frequency standards must be fed by a microwave probe signal in order to obtain the necessary information on the atomic resonance frequency. They behave as a resonator, and they are called passive frequency standards. As a resonator response we obtain a discriminating resonance signal I_s which is superimposed to the total signal I_t . I_s and I_t are measures of the number of particles contributing to the signal and the background, respectively. Small-size hydrogen masers with a compact cavity, but increased losses, so-called passive hydrogen masers (PHM), may be operated in a similar way. The electronic system that detects the resonance frequency and controls the frequency of a 5 or 10 MHz quartz crystal oscillator is shown schematically in Fig. 1. The probe microwave signal is synthesized from a VCXO and a frequency modulation is superimposed. This creates a modulation of the amplitude of the device

response. It is processed by synchronous detection to extract the error signal which is applied to control the VCXO.

FIGURE 1

Block diagram representation of the frequency control of a quartz crystal oscillator in a passive frequency standard




HDB-31-1

If a hydrogen maser has an appropriately sized cavity it will oscillate at the frequency of the clock transition if enough atoms in the upper clock state are injected into the cavity volume. In this case, it is called an active hydrogen maser (AHM). An electronic system implements the principles of heterodyne receivers and of phase-lock loops, as shown in Fig. 2. It phase-locks a quartz crystal oscillator at 5, 10 or 100 MHz to the maser oscillation.

FIGURE 2

Block diagram representation of the phase control of a quartz crystal oscillator by the H-maser oscillation



HDB-31-2

In both cases, passive and active, the output signals of the atomic frequency standards are derived from the controlled quartz crystal oscillator.

1.4 Basic metrological concepts

1.4.1 Frequency stability

In this section, the physical origin of the random frequency fluctuations or of the frequency drifts are briefly reviewed. Measures of stability are discussed in mathematical detail in Chapter 2. Theoretical values, based on the physics package of the oscillators, and typical numerical values are given in the respective sections of Chapter 3.

1.4.1.1 Definition

The term “frequency stability” is usually used instead of “frequency instability”. It is the spontaneous and/or environmentally caused frequency change within a given time interval. Generally, one distinguishes between stochastic frequency fluctuations and systematic effects such as frequency drifts – systematic variations, supposedly linear – with time. Systematic instabilities may be caused by environmental effects. Frequency stability is dependent on the measurement system bandwidth and/or on the sample time or integration time. Random or stochastic instabilities are characterized either in the time domain or in the frequency domain or both.

The separation of the contributions to the measured variance into stochastic and systematic effects is somewhat arbitrary. A stochastic contribution, often called a contribution of “type A” is generally estimated by a statistical technique, whereas a systematic effect, often called a contribution of “type B” is evaluated by an ancillary non-statistical method such as a calibration or a theoretical treatment.

For both types of contributions to the total uncertainty, the implicit assumption is that the data are stationary – that the estimation procedure does not depend on when it is performed. Unfortunately, this is often not the case, especially for estimates of the frequency drift, so that the elapsed time since the last estimate of the uncertainty budget is an important parameter. From a time-domain perspective, a stochastic contribution to the variance is one that changes rapidly with respect to the averaging time so that consecutive contributions are essentially uncorrelated, while the opposite is true for a systematic offset – it changes slowly with respect to the averaging time and so can be considered constant over relatively long intervals.

To further complicate the situation, the method that is used to evaluate a systematic effect may itself have some uncertainty, and this uncertainty will contribute to the overall uncertainty budget. This contribution is estimated by using the usual methods of propagation of errors.

1.4.2 Accuracy

1.4.2.1 Residual frequency offsets

Although one strives to design atomic frequency standards such that the ideal condition of atoms isolated and at rest is satisfied as much as possible, small frequency offsets cannot be avoided in current practice. The uncertainty on our knowledge of the magnitude of these frequency offsets determines the accuracy of atomic frequency standards. The physical origin of some of these frequency offsets is common to all types of frequency standards, e.g. the second order Doppler effect due to thermal motion of the atoms. Others are specific to a given type e.g. the buffer gas frequency offset and the light shift in a rubidium cell and will be dealt with in the relevant sections in Chapter 3.

1.4.2.2 Definition

In general, the accuracy is the degree of conformity of a measured or calculated value to its definition. In quantitative terms, it is expressed as the cumulative normalised uncertainty of the value realised relative to that given by the definition. Usually, this uncertainty has two different components. In the case of atomic frequency standards, one is related to random perturbations affecting the measurement of residual frequency offsets (e.g. that due to the measurement of the magnetic field applied). The other is a systematic component that is related to a lack of knowledge of some operating parameters (e.g. the inhomogeneity of the magnetic field). These two components must be combined with care. In general, the methods that are used to evaluate systematic contributions have uncertainties as well, which further complicates the overall uncertainty budget.

References

- BARGER R.L. and HALL J.H. [1969]. *Pressure Shift and Broadening of Methane Absorption Line at $3.39\ \mu$ Studied by Laser Saturated Molecular Absorption*, Phys. Rev. Lett., Vol. 22., 4-8.
- BARROW H.G., NAING T.L, SCHNEIDER R., ROCHELEAU T.O., YEH V., REN Z., NGUYEN C. T-C, [2013] *A real-time 32768-kHz clock oscillator using a 0.0154-mm^2 micromechanical resonator frequency-setting element*, Proc. IEEE 2012 Frequency Control Symposium, DOI 10.1109/FCS.2012.6243740.
- BAUCH A. [2005] The PTB primary clocks CS1 and CS2. *Metrologia*, Vol. 42, S43 – S54.
- BESSON R.J. [1977] *A New Electrodeless Resonator Design*. Proc. 31st Annual Symposium on Frequency Control, Atlantic City, N.J., USA, 147-152.
- BIPM [2019] *Le Système International d'Unités*, 9^e édition, Sèvres, France.
- CADY W.G. [1923] *United States Patent*, 1,472,583, issued 30 October 1923.
- CHU S. [1998] *The Manipulation of Neutral Particles*. Rev. Mod. Phys., Vol. 50, 685-706.
- COHEN-TANNOUDJI C. [1998] *Manipulating atoms with photons*. Rev. Mod. Phys., Vol. 50, 707-720.
- CLAIRON A., LAURENT P., SANTARELLI G., GHEZALI S., LEA S.N. and BAHOURA M. [1995] *A Caesium Fountain Frequency Standard: Preliminary Results*. IEEE Trans. Instr. and Meas., Vol. 44, 128-132
- COX R.T. [1922] *Standard Radio Wavemeter, Bureau of Standards Type R 70 B*, J. Opt. Soc. Am., Vol. 6, 162-168.
- CUTLER L.S. [2005] *Fifty Years of Commercial Caesium Clocks*. Metrologia, Vol. 42 S90 – S99.
- DIMARCQ N. *et al.*, (37 authors) [2023] *Roadmap towards the redefinition of the second*, Metrologia [in press, to be completed]
- ESSEN L. and PARRY J.V.L. [1957] *The Caesium Resonator as a Standard of Frequency and Time*. Philosophical Transactions of the Royal Society, Vol. 250, 45-69.
- FILLER R.L. and VIG J.R. [1957] *Long-Term Aging of Oscillators*. IEEE Trans. on Ultrasonics, Ferroelectrics, and Frequency Control, Vol. 42, 387-393.
- GERBER E.A. and BALLATO A. (Editors) [1985] *Precision Frequency Control*, 2 volumes, Academic Press, Orlando, Florida, USA.
- HUNTEMANN N., SANNER C., LIPPHARDT B., TAMM CHR., PEIK E. [2016] *Single Ion Atomic Clock with 3×10^{-18} Systematic Uncertainty*, Phys. Rev. Lett., Vol. 116, 063001.
- KITCHING J., KNAPPE S., VUKICEVIC N., HOLLBERG L., WYNANDS R., WEIDMANN, W. [2000]. *A Microwave Frequency Reference Based on VCSEL-Driven Dark Line Resonances in Cs Vapor*, IEEE Trans. Instr. Meas. Vol. 49, 1317-1317.
- KITCHING J. [2018], *Chip-Scale Atomic Devices*, Appl. Phys. Rev. Vol. 5, 031302.
- LUDLOW A.D., BOYD M.M., JUN YE, PEIK E. and SCHMIDT P.O. [2015], *Optical atomic clocks*, Rev. Mod. Phys. Vol. 87, 638 – 701.
- NICHOLSON T.L., CAMPBELL S.L., HUTSON R.B., MARTI G.E., BLOOM B.J., MCNALLY R.L., ZHANG W., BARRETT M.D., SAFRONOVA M.S., STROUSE G.F., TEW W.L., and YE J., [2015] *Systematic Evaluation of an Atomic Clock at 2×10^{-18} total uncertainty*, Nature Communications, Vol. 6., DOI 10.1038/ncomms7896
- PEIL S., HANSSEN J., SWANSON T., TAYLOR J. and EKSTROM C. [2016] *The USNO rubidium fountains*. J. Phys.: Conf. Ser. Vol. 723 012004.
- PHILIPS W.D. [1998] *Laser Cooling and Trapping of Neutral Atoms*. Rev. Mod. Phys., Vol. 50, 721-733.

- RIEHLE F., GILL P., ARIAS F. and ROBERTSSON L. [2018] *The CIPM List of Recommended Frequency Standard Values: Guidelines and Procedures*. *Metrologia*, Vol. 55(2) 188-200.
- ROCHELEAU T. O., NAING T. L., REN Z., and NGUYEN C. T-C. [2012] *Acoustic whispering gallery mode resonator with $Q > 109,000$ at 515 MHz*, Proc of the IEEE International Conference on MEMS Systems, DOI 10.1109/MEMSYS.2012.6170276.
- SCHEIBE A. und ADELSBERGER U. [1950] *Die Gangleistungen der PTR-Quarzuhren und die jährliche Schwankung der astronomischen Tageslänge*, *Zeitschrift für Physik*, Vol. 127, 416-428.
- SULLIVAN D. B. [2001] *Time and frequency measurement at NIST: The first 100 years*. Proc. 2001 IEEE Frequency Control Symposium and PDA Exhibition, 4-17.
- VANIER J. and AUDOIN C. [1989] *The Quantum Physics of Atomic Frequency Standards*, 2 volumes, Institute of Physics Publishing, Bristol.
- VANIER J. and TOMESCU C. [2016] *The Quantum Physics of Atomic Frequency Standards – Recent developments*. CRC Press, Taylor & Francis Group, Boca Raton, London, New York.
- VANIER J. and AUDOIN C. [2005] *The classical caesium beam frequency standard: fifty years later*. *Metrologia* Vol. 42, S31-S42.
- VIG J.R. [1991] *Quartz Crystal Resonators and Oscillators for Frequency Control and Timing Applications*. A Tutorial. Available from US Army Electronics Technology and Devices Laboratory. Fort Monmouth, N. J. USA.
- WALLS F.L. and GAGNEPAIN J.J. [1992] *Environmental Sensitivities of Quartz Crystal Oscillators*. IEEE Trans. on Ultrasonics, Ferroelectrics and Frequency Control, Vol. 39, 241-249.
- WALLS F.L., HANDEL P.H., BESSON R. and GAGNEPAIN J.J. [1992] *A New Model of 1/f Noise in BAW Quartz Resonators*. Proc. IEEE Frequency Control Symposium, Hershey, Pennsylvania, USA, 327-333.
- WINELAND D.J., ITANO W.M., BERGQUIST J.C., BOLLINGER J.J., HEINZEN D.J., MANNEY C.H., MOORE F.L., RAIZEN M.G. and WEIMER C.S., [1990] *Trapped-Ion Frequency Standards*. Proc. 22nd Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting, Vienna, VA, USA, 53-60.

CHAPTER 2

Characterisation: frequency domain, time domain

TABLE OF CONTENTS

	<i>Page</i>
2.1 Introduction	12
2.2 Model of the oscillator.....	12
2.2.1 The phasor model and the analytic signal	12
2.2.2 The low-noise oscillator	12
2.2.3 Spectrum of the low-noise oscillator.....	13
2.2.4 The high-noise oscillator	13
2.2.5 Spectrum of the high-noise oscillator.....	14
2.2.6 Effect of frequency multiplication	14
2.2.7 Demodulation of the noise processes	14
2.2.8 Standard definition of noise processes	15
2.2.9 Multiplicative and additive noise	16
2.2.10 Polynomial model	17
2.3 Characterisation: definitions and methods.....	18
2.3.1 Spectral domain.....	18
2.3.2 Time domain	19
2.3.3 Environmental	35
2.4 A bridge to the next Chapter.....	35
2.5 Appendix: Random processes.....	36
2.5.1 Introduction	36
2.5.2 Definition of a random process	36
2.5.3 Stationary random processes	36
2.5.4 Non-stationary random processes	36
2.5.5 Auto-correlation function.....	36
2.5.6 Power spectral density.....	36
2.5.7 Linear filtering of random processes.....	37
References	37
ITU Texts	39

2.1 Introduction

This Chapter is an introduction to the methods used to characterise the random phase and frequency instabilities of oscillators in both the frequency domain and the time domain.

The characterisation of phase and frequency instabilities in precision oscillators is already covered in Recommendation ITU-R TF.538 and extensively in the literature. For example, references [Allan, 1987; Allan *et al.*, 1988; IEEE, 1988; Lesage *et al.*, 1979; Rutman, 1978 and Stein, 1985] are all review papers covering topics similar to those discussed in the present Chapter.

In addition, the IEEE standard IEEE 1139-2022 also recommends the statistical characterization of oscillators and frequency standards [IEEE 2022].

This Chapter is aligned with Recommendation ITU-R TF.538-4 (07/2017) – Measures for random instabilities in frequency and time (phase).

The originality of this particular presentation is the fact that the different topics are presented in a unified approach using electrical engineering concepts such as the analytic signal, random processes, linear operators, transfer functions or the sampling theorem instead of following the conventional approach inherited from the historical development of time and frequency characterisation concepts.

A particular effort has been made in order to give to the reader the necessary insight into the theoretical aspects of the phase and frequency noise characterisation methods and at the same time to give the practical tips and algorithms necessary for their application. The next chapter covers similar topics to the above, but with more experimental guidance.

2.2 Model of the oscillator

2.2.1 The phasor model and the analytic signal

The signal of the sinusoidal oscillator is a narrow-band, band-pass signal, centred about the carrier frequency ν_0 . It can be modelled by a phasor, i.e. by a vector rotating in the complex plane at the carrier frequency:

$$Y(t) = A_Y(t) \exp(j2\pi\nu_0 t) \quad (2-1)$$

A is the nominal amplitude, $\gamma(t)$ is the complex envelope and $\exp(j2\pi\nu_0 t)$ the phasor. The real signal $s(t)$ is the real part of the phasor signal $\psi(t)$:

$$Y(t) = s(t) + j\bar{s}(t) \quad (2-2)$$

where $s(t)$ is the Hilbert transform of $\bar{s}(t)$, i.e. the projection of the phasor on the real axis.

It can be shown that the phasor representation of a bandpass signal is a particular case of an analytic signal [Bedrosian, 1962; Bernier *et al.*, 1985].

The complex envelope $\gamma(t)$ of the analytic signal is a low-pass complex random process that represents the information modulated on the carrier. In the case of a telecommunication band-pass signal the low-pass message signal is translated on purpose about a carrier frequency by modulation. In the case of an oscillator signal, on the other hand, the complex envelope is a random process characteristic of a certain type of oscillator: quartz oscillator, atomic frequency standard, etc. A review of random modulation can be found in [Papoulis, 1983].

The complex envelope can be separated either into an in-phase $p(t)$ and a quadrature $q(t)$ low-pass processes or into a unity carrier, an amplitude low-pass noise process $\varepsilon(t)$ and a phase low-pass noise process $\varphi(t)$:

$$\gamma(t) = p(t) + jq(t) = [1 + \varepsilon(t)] \exp(j\varphi(t)) \quad (2-3)$$

2.2.2 The low-noise oscillator

In the case of a precision or low-noise oscillator, the mean square value of the phase noise process is much smaller than unity, and the mean square of the amplitude noise process is much smaller than unity,

$$E\{\varphi^2(t)\} \ll 1$$

$$E\{\varepsilon^2(t)\} \ll 1$$

and where $E\{\}$ is the statistical expectation operator and the phase noise exponential can be linearised:

$$\exp(j\varphi(t)) \approx 1 + j\varphi(t)$$

yielding to lowest order:

$$\Psi(t) = A [1 + \varepsilon(t) + j\varphi(t)] \exp(j2\pi\nu_0 t) \quad (2-4)$$

Under these conditions the in-phase component, $p(t)$, can be identified as a unity carrier plus the amplitude noise while the quadrature component, $q(t)$, is identified as the phase noise.

$$\begin{aligned} p(t) &= 1 + \varepsilon(t) \\ q(t) &= \varphi(t) \end{aligned} \quad (2-5)$$

The phase noise and the amplitude noise are orthogonal when the complex envelope is represented in the complex plane. This property enables the independent demodulation of the phase and amplitude noise processes by synchronous demodulation as explained below, but the property holds true only in the low-noise case.

2.2.3 Spectrum of the low-noise oscillator

Note that the definitions relevant to the autocorrelation functions and to the power spectral densities are given in the appendix on random processes and follow the notations used in [Wozencraft *et al.*, 1965]. In particular, $S_{\varphi\varphi}(f)$ refers to the two-sided power spectral density of the random process $\varphi(t)$, $S_{\varphi\varphi}^+(f)$ refers to the one-sided power spectral density of the random process $\varphi(t)$, while $R_{\varphi}(\tau)$ refers to the autocorrelation function of the random process $\varphi(t)$. Note, moreover, that outside the present chapter of the handbook, the autocorrelation function of $\varphi(t)$ is noted $R_{\varphi}(\tau)$ and the corresponding spectral density $S_{\varphi}(f)$. The ITU and IEEE standard is to use the one-sided spectral densities. The autocorrelation function of the complex envelope associated with the low-noise oscillator signal as defined by equation (2-5) in [Bernier *et al.*, 1985]:

$$R_{\gamma\gamma}(\tau) = 1 + R_{\varepsilon\varepsilon}(\tau) + R_{\varphi\varphi}(\tau) \quad (2-6)$$

The power spectral density of the oscillator signal $s(t)$ is the Fourier transform of the low-pass autocorrelation function in equation (2-6) translated about the carrier frequency, because of equations (2-1) and (2-2), yielding:

$$S_{ss}^+(f) = \frac{A^2}{2} (\delta(f - \nu_0) + S_{\varepsilon\varepsilon}(f - \nu_0) + S_{\varphi\varphi}(f - \nu_0)) \quad (2-7)$$

where $\delta(f - \nu_0)$ is a Dirac function representing the carrier and where $S_{\varepsilon\varepsilon}(f - \nu_0)$ and $S_{\varphi\varphi}(f - \nu_0)$ are the two-sided power spectral densities of the amplitude and phase noise processes.

The low-noise oscillator therefore appears on the spectrum analyser as a carrier at the nominal frequency ν_0 superposed to the amplitude and phase two-sided power spectral densities translated about the carrier frequency. In telecommunication terms this spectrum is that of a carrier AM modulated and simultaneously PM modulated with a low modulation index by the amplitude and phase random processes characteristics of the oscillator.

The equations in this section are true only if the signal spectrum is symmetrical about ν_0 . If not, then equation (2-6) has a cross term, and equation (2-7) can be written as equation (2-8) {[Bernier 1985, Equation 14]}. For $|\nu| < \nu_0$:

$$\frac{1}{2} [S_{ss}^+(\nu_0 + f) + S_{ss}^+(\nu_0 - f)] = \frac{A^2}{4} [S_{\gamma\gamma}(f) + S_{\gamma\gamma}(-f)] = \frac{A^2}{2} [\delta(f) + S_{\varepsilon\varepsilon}(f) + S_{\varphi\varphi}(f)] \quad (2-8)$$

which says that the average sideband level = AM level + PM level. Note that $S_{\varphi\varphi}(f) = \mathbf{L}(f)$, see equation (2-25).

2.2.4 The high-noise oscillator

In the case of a high phase noise oscillator, the mean quadratic value of the phase noise process is large:

$$E\{\varphi^2(t)\} \geq 1$$

and the exponential term of the complex envelope cannot be linearised. In this case the phase noise process contributes to both the in-phase and quadrature components of the complex envelope yielding:

$$\begin{aligned} p(t) &= (1+\varepsilon(t)) \cos \varphi(t) \\ q(t) &= \sin \varphi(t) \end{aligned} \quad (2-9)$$

In this case, the amplitude and phase noise processes are not orthogonal anymore in the complex plane. The lack of orthogonality complicates the problem of independent demodulation of the processes.

2.2.5 Spectrum of the high-noise oscillator

Assuming that the amplitude noise is negligible and that the phase noise is high, it can be shown [Bernier *et al.*, 1985] that the autocorrelation function of the complex envelope can be separated into three terms corresponding respectively to the carrier and to the in-phase and quadrature components [Bernier *et al.*, 1985]:

$$R_{\gamma\gamma}(\tau) = \exp(-R_{\varphi\varphi}(0)) + r_p(\tau) + r_q(\tau) \quad (2-10)$$

where:

$$\begin{aligned} r_p(\tau) &= [\exp(-R_{\varphi\varphi}(0))] [\cosh(R_{\varphi\varphi}(\tau)) - 1] \\ r_q(\tau) &= [\exp(R_{\varphi\varphi}(0))] \sinh(R_{\varphi\varphi}(\tau)) \end{aligned} \quad (2-11)$$

The one-sided power spectral density of the band-pass signal $s(t)$ is therefore [Bernier *et al.*, 1985]:

$$S_{ss}^+(f) = \frac{A^2}{2} \left(\exp(-R_{\varphi\varphi}(0)) \delta(f - \nu_0) + G_p(f - \nu_0) + G_q(f - \nu_0) \right) \quad (2-12)$$

where G_p and G_q are the Fourier transforms of r_p and r_q .

As the mean quadratic value $R_{\varphi\varphi}(0) = E\{\varphi^2(t)\}$ of the phase noise increases beyond unity, the carrier vanishes exponentially, i.e. the signal becomes incoherent, and the in-phase and quadrature spectral components both tend towards a Gaussian shape regardless of the original shape of the phase noise power spectral density. This is known in telecommunication theory as Woodward's theorem [Bernier *et al.*, 1985; Blachman *et al.*, 1969; Papoulis, 1983]. In the case of a low-pass white phase noise process, the carrier disappears completely and the spectrum becomes Gaussian when the mean squared value of the phase noise exceeds about 8 rad^2 [Bernier *et al.*, 1985]. The above is a simple stationary case. However, in any practical oscillator the phase noise is indeed non-stationary. The non-stationary case is treated in [Walls *et al.*, 1975].

2.2.6 Effect of frequency multiplication

When a carrier frequency is multiplied by n in a frequency multiplier, the amplitude noise $\varepsilon(t)$ can be kept very low by the use of limiters but the phase noise $\varphi(t)$ is multiplied by n . The mean squared value of the phase noise process is therefore multiplied by n^2 in the ideal case. In practice, the phase noise level may be higher than this fundamental lower limit and the amplitude noise level depends on the actual characteristics of the multiplier device. If the multiplication factor, n , is high enough, the mean squared value of the phase noise process will eventually reach and exceed unity. If this is the case the original low-noise oscillator signal becomes a high-noise oscillator signal, i.e. becomes more or less incoherent as described above. See Chapter 4 for more details.

2.2.7 Demodulation of the noise processes

The application of both frequency and time domain characterization methods described in the next sections implies that the low-pass phase noise process $\varphi(t)$ of the oscillator under test can be extracted by demodulation. Extraction of the phase noise is simpler in the case of a low phase noise oscillator.

For the characterization of oscillators in the time domain, the demodulation method is heterodyne down-conversion to a relatively low IF frequency followed by zero-crossing detection. The zero-crossings are then "time tagged" or the IF period is averaged and sampled using a digital period/frequency counter or a time-tag counter.

In the case of spectral domain characterization methods, the demodulation technique is synchronous detection. Synchronous detection of a band-pass signal with the local oscillator in-phase with the carrier yields the in-phase component of the complex envelope i.e. $p(t)$. Conversely the synchronous detection of a band-pass signal with the local oscillator in quadrature with the carrier yields the quadrature component of the complex envelope i.e. $q(t)$.

As shown above the amplitude noise and the phase noise processes are orthogonal in the complex plane, and thus can be easily separated by synchronous detection in the case of a phase noise process with a small mean squared value.

Moreover, at low Fourier frequencies the phase noise of any real oscillator is usually modelled by a non-stationary process of the “power-law” type, see sections about the polynomial model and about non-stationary processes below, whose mean value can be defined only locally. Therefore, the only way to maintain the local oscillator in quadrature with the carrier for the purpose of the synchronous detection of the phase noise is by means of phase locked loop techniques. The phase of the local oscillator must be corrected continuously in order to track the slow random phase drift of the oscillator under test with respect to the reference oscillator.

2.2.8 Standard definition of noise processes

2.2.8.1 Amplitude and phase noise processes

The basic noise processes that characterise the oscillator are the amplitude and phase noise processes $\varepsilon(t)$ and $\varphi(t)$ defined above in the phasor model of the oscillator. As mentioned above the amplitude and the phase noise processes in a low noise oscillator are orthogonal in the complex plane and they can be detected independently with a very good rejection of each other. Therefore, for the purpose of frequency stability analysis the amplitude noise can be neglected and it is the phase noise process that is relevant. Caution should be exercised here. AM to PM conversion can take place in some equipment. Even linear band-pass filtering produces AM to PM conversion if the transfer function is not perfectly symmetric with respect to the carrier frequency [Bernier *et al.*, 1985]. For a more detailed discussion refer to Chapter 4.

2.2.8.2 Time error process

The time error process $x(t)$ can be defined as the normalised phase difference accumulated between an oscillator used as a clock and a reference oscillator considered as the reference clock. It is related to the phase noise process by:

$$x(t) = \frac{\varphi(t)}{2\pi\nu_0} \quad (2-13)$$

where the nominal frequency ν_0 is expressed in [Hz] and the time error in [s].

2.2.8.3 Instantaneous frequency process

The normalised instantaneous frequency deviation $y(t)$ is defined as the derivative of the time error process $x(t)$:

$$y(t) = \frac{dx(t)}{dt} = \frac{1}{2\pi\nu_0} \frac{d\varphi(t)}{dt} \quad (2-14)$$

$y(t)$ has no units. It is a frequency “deviation” in the sense that being associated with the complex envelope it describes not the instantaneous frequency of the oscillator signal $s(t)$ but the frequency deviation with respect to the nominal carrier frequency ν_0 .

The instantaneous frequency $\nu_0(t)$ in [Hz] is related to $y(t)$ by the following equation:

$$y(t) = \frac{\nu(t) - \nu_0}{\nu_0} \quad (2-15)$$

For the purpose of frequency stability analysis, it is the normalised frequency deviation $y(t)$ that is the most relevant quantity. The normalised frequency deviation is conserved after frequency multiplication or division, and it also enables the direct comparison between the frequency instability levels of two oscillators of different nominal carrier frequencies.

2.2.9 Multiplicative and additive noise

2.2.9.1 Multiplicative noise

When the oscillator signal is amplified, transmitted, multiplied or is subjected to any other kind of processing, it can be degraded by either or both multiplicative and additive noise.

Multiplicative noise is actually a random modulation process that can be represented in a form similar to the complex envelope of the oscillator signal itself.

Assume that the signal of the oscillator is amplified and that the amplifier introduces random AM and PM modulations of the signal. The random modulation mechanism may be, for example, flicker noise in the polarisation currents of the transistors which modulates the gain and the phase-shift at the carrier frequency. The modulation function associated with the multiplicative noise process has the form:

$$M(t) = 1 + \varepsilon_m(t) + j\varphi_m(t) \quad (2-16)$$

The complex envelope of the oscillator signal at the output of the amplifier is the complex envelope at the input multiplied by the modulation function. Assuming that the phase noise processes of both the original signal and of the amplifier are both low-noise processes,

$$E\{(\varphi_r(t))^2\} \ll 1 \text{ and } E\{(\varphi_m(t))^2\} \ll 1$$

the complex envelope of the low-noise oscillator at the input is:

$$\gamma_i(t) = 1 + \varepsilon_i(t) + j\varphi_i(t) \quad (2-17)$$

and, neglecting quadratic terms, the complex envelope of the oscillator signal at the output is:

$$\gamma_o(t) = M(t)\gamma_i(t) = 1 + \varepsilon_i(t) + \varepsilon_m(t) + j(\varphi_i(t) + \varphi_m(t)) \quad (2-18)$$

The above result shows that in the low phase noise case, where the amplitude and phase noise processes are orthogonal, the multiplicative amplitude and phase noise processes simply add to the original amplitude and phase noise processes.

2.2.9.2 Additive noise

The case of additive noise can be treated as follows: assume a band-limited white noise process $n(t)$ of single-sided power spectral density N_0 . According to Rice's representation theorem [Papoulis, 1983; Wozencraft *et al.*, 1965] any band-pass random process can be decomposed into an in-phase and quadrature components with respect to an arbitrary carrier frequency ν_0 :

$$n(t) = \sqrt{2}(n_p(t) + jn_q(t)) \exp(j2\pi\nu_0 t) \quad (2-19)$$

Assuming that the band-pass process is a band limited white noise process of bandwidth $2B$ and single-sided spectral density N_0 and assuming, moreover, that its power spectral density is symmetrical with respect to the arbitrary carrier frequency ν_0 , it can be shown that $n_p(t)$ and $n_q(t)$ are statistically independent low-pass white noise processes of single sided power spectral density N_0 and of bandwidth B [Wozencraft *et al.*, 1965].

Rice's decomposition is in fact an analytic signal representation. Normalising with respect to the carrier peak amplitude A of the oscillator signal one obtains:

$$n(t) = A \left(\frac{\sqrt{2}}{A} n_p(t) + j \frac{\sqrt{2}}{A} n_q(t) \right) \exp(j2\pi\nu_0 t) \quad (2-20)$$

Comparing term to term with the analytic signal (equation (2-4)) of the low-noise oscillator, the contributions to amplitude and phase noise produced by additive noise are found to be two identical low-pass white noise processes of bandwidth B .

$$S_{\varepsilon\varepsilon}^+(f)|_{\text{additive}} = S_{\varphi\varphi}^+(f)|_{\text{additive}} = \frac{N_0}{\frac{1}{2}A^2} = \frac{1}{S/N (1 \text{ Hz})} \quad (2-21)$$

The one-sided power spectral density of both processes is equal to the inverse of the signal to noise ratio in a one Hertz bandwidth. Here the noise power in a 1 Hz bandwidth is defined as the power spectral density N_0 of the additive noise $n(t)$ and the signal power is defined as the power of the sinusoidal carrier of peak amplitude A .

2.2.10 Polynomial model

The conventional way of modelling the noise of oscillators is by the use of the polynomial model of $S_{yy}^+(f)$.

$$S_{yy}^+(f) = \sum_{\alpha=-2}^{\alpha=2} h_{\alpha} f^{\alpha} \quad (2-22)$$

The phase/time process $x(t)$ being the integral of the frequency process $y(t)$ there is a direct correspondence between the power spectral densities of $x(t)$ and of $y(t)$:

$$S_{xx}^+(f) = \frac{S_{yy}^+}{(2\pi f)^2} = \frac{1}{(2\pi)^2} \sum_{\beta=-4}^0 h_{\alpha} f^{\beta} \quad \text{where} \quad \beta = \alpha - 2 \quad (2-23)$$

Table 2 summarises the different types of noise processes encountered in the polynomial model of oscillators. The model is convenient for classifying the different types of noise processes and for the translation of one type of frequency stability measurement into another.

Note that, in general, for a power law process $z(t)$ of power spectral density $S_{zz}^+(f) = kf^{\theta}$, the process is non-stationary when $\theta \leq -1$ in the sense that the mean squared value of the process is given by:

$$E\{z^2(t)\} = R_{xx}(0) = \int_0^{\infty} S_{xx}^+(f) df \quad (2-24)$$

which is undefined when the integral diverges at the origin. Note that here is used the generic process $z(t)$ and the exponent θ in order to state general properties that apply equally well to the time/phase process $x(t)$ and to the frequency process $y(t)$ with their respective exponents β and α .

The integral in equation (2-24) does not converge at high Fourier frequencies for $\theta \geq -1$ because pure power-law processes are not band limited. This problem is only an artefact of the model, though, since the actual phase and frequency noise processes of oscillators are always band limited.

TABLE 2

Classification of power-law processes in the polynomial model

	Exponent α of $S_{yy}^+(f)$	Exponent β of $S_{xx}^+(f)$
White noise PM	2	0
Flicker noise PM	1	-1
White noise FM	0	-2
Flicker noise FM	-1	-3
Random walk FM	-2	-4

While f^{-1} is often called “flicker noise”, a single integration of a white noise process yields an f^{-2} power law process, which is called a “random walk” or a one-dimensional Brownian motion process.

In [Greenhall, 1983] it is shown that the power-law random processes found in oscillators can be characterised by power-spectral densities even when they are non-stationary ($\theta \leq -1$).

2.3 Characterisation: definitions and methods

2.3.1 Spectral domain

2.3.1.1 Basic definitions

In the spectral domain four measures are typically used to describe the oscillator: the power spectral density of $y(t)$, the power spectral density of $x(t)$, the power spectral density of $\varphi(t)$ and $L(f)$ a spectral measurement related to the concept of spectral purity.

The basic spectra are summarised in Table 3. The question whether the quantity of interest is $y(t)$, $x(t)$ or $\varphi(t)$ depends upon the particular application. Most often, however, spectral measurements are used to characterise the spectral purity of the oscillator; therefore, it is the spectral density of $\varphi(t)$ that is usually given.

TABLE 3

Basic measurements in the spectral domain

Measurement	Symbol
PSD of $y(t)$	$S_{yy}^+(f)(\text{Hz}^{-1})$
PSD of $x(t)$	$S_{xx}^+(f)(\text{s}^2 \text{ Hz}^{-1})$
PSD of $\varphi(t)$	$S_{ff}^+(f)(\text{rad}^2 \text{ Hz}^{-1})$
Script $L(f)$	$\mathbf{L}(f)(\text{Hz}^{-1})$

The spectral purity measure $\mathbf{L}(f)$ is defined as:

$$\mathbf{L}(f) = \frac{1}{2} S_{\varphi\varphi}^+(f) \quad (2-25)$$

The relationship between the two measurements is justified below.

2.3.1.2 Spectral purity concepts

The spectral purity of an oscillator refers to the noise found on each side of the carrier when the signal of an oscillator is measured directly on a spectrum analyser.

Recalling that the power spectral density of the signal $s(t)$ of a low phase noise oscillator is given by equation (2-7), assume that the power spectral density of the oscillator signal is measured, normalised by the carrier average power $\frac{1}{2} A^2$ and translated from ν_0 to the origin into a low-pass spectrum.

$$\mathbf{L}(f) = \frac{S_{ss}^+(f+\nu_0)}{\frac{1}{2} A^2} = \delta(f) + S_{\varepsilon\varepsilon}(f) + S_{\varphi\varphi}(f) \quad (2-26)$$

This is to satisfy the definition of the spectral purity measure $\mathbf{L}(f)$ if the amplitude noise $\varepsilon(t)$ can be made negligible with respect to the phase noise $\varphi(t)$. Then the spectral purity measure becomes identical to a measure of phase noise for all Fourier frequencies except the origin. This is often not true for frequency synthesisers and other devices where the oscillator signal has passed through a number of stages of amplification.

Note, however, that commercial spectrum analysers do not have the necessary resolution nor the dynamic range for the characterisation of the phase noise by a direct measurement of the band-pass signal $s(t)$. Moreover, the amplitude noise is often not negligible. In other words, the actual measurement of $\mathbf{L}(f)$ is not feasible in practice. Therefore, although commercial oscillators are almost always specified by $\mathbf{L}(f)$ the actual measurement involves the demodulation of $\varphi(t)$ by phase locked loop techniques and the evaluation of its power spectral density using a digital low-pass spectrum analyser. The measurement of $S_{ff}^+(f)$ is then divided by a factor of two and labelled $\mathbf{L}(f)$.

2.3.2 Time domain

2.3.2.1 Introduction

Time-domain, in contrast to frequency-domain, deals with the effects of averaging over different values of time, τ . For useful time-domain measures, Fourier transform relationships exist between the Fourier frequency, f , and the time domain parameter, τ . Observing the dependence of a variance as τ is varied is often very insightful as to the oscillator characteristics.

Time-domain characterisation conventionally uses statistics computed from a discrete time series of average frequency samples obtained from a digital counter which is used to measure the beat note between the oscillator under test and a reference oscillator. The same discrete time series can, of course, be used in an FFT, for example, to compute the frequency domain-spectrum.

The approach of this presentation, on the other hand, consists of defining time domain variances in terms of the mean squared value of the continuous stationary random process obtained by the application of a linear operator, specific to the time domain variance of interest, to the instantaneous frequency process $y(t)$. Since $y(t)$ averaged over τ is $(x(t) - x(t - \tau))/\tau$, these variances can also be expressed in terms of $x(t)$.

This approach enables the establishment of the different variances in a simpler and more intuitive way than heretofore. The fact that, in practice, the variances are estimated over a finite number of samples taken from the underlying continuous process only affects the uncertainty of the estimation. The problem of the uncertainty on the practical estimation of the variances can be treated separately from the definition of the variance themselves and is already extensively investigated in the literature.

2.3.2.2 Basic concepts

2.3.2.2.1 Model of the frequency counter

Assume that an Oscillator Under Test (OUT) is to be characterised in the time domain. The conventional method is to heterodyne the original bandpass signal down to an arbitrary low beat frequency ν_0 by mixing the OUT signal with the signal of a reference oscillator much more stable than the OUT. The low frequency beat signal reproduces faithfully the original frequency fluctuations of the OUT, which is typically measured with a digital frequency counter.

The usual digital counter measures the frequency beat signal by counting the zero crossings of the signal. Zero crossing typically means the time-point where the signal voltage transverses positively through zero volts. For every sampling period, T , the counter delivers a sample ν_k which is the instantaneous frequency $\nu_b(t)$ [Hz] averaged over the gate opening time, τ .

In most frequency counters the sampling period, T , is larger than the averaging time, τ . There is therefore a dead-time, $T - \tau$ during which the counter computes the average frequency, resets the internal counters/registers, and waits for the next zero crossing of the signal. The next count can start only at the next zero crossing of the signal therefore in a frequency counter the dead-time is at least one period of the signal.

The dead time can be important as will be explained below because it biases the statistics when the samples are used for time domain measurements.

If the nominal average frequency ν_b is subtracted from the samples ν_k and if they are normalised with respect to the OUT mean frequency ν_0 , one obtains the normalised samples y_k .

$$y_k = \frac{\nu_k - \nu_b}{\nu_0} \quad (2-27)$$

The discrete random process, or time series, $y_0, y_1, y_2, \dots, y_k$ coming from the frequency counter is therefore equal to the continuous random process $y(t, \tau)$, i.e. the moving average of $y(t)$ over τ as defined in next section, sampled with a period, T .

The time domain characterisation methods of the frequency stability are usually defined in terms of statistics performed over the discrete process $y_0, y_1, y_2, \dots, y_k$ coming from the frequency counter. It will be shown, however, that they are more easily understood when they are also defined in terms of the underlying continuous random processes.

2.3.2.2.2 Moving average operator

As mentioned above time domain variances are traditionally defined in terms of statistics performed on the discrete time series $y_0, y_1, y_2, \dots, y_k$ coming from the frequency counter. The samples obtained from the counter are a periodic sampling of the instantaneous frequency averaged over the gate opening interval, τ .

The frequency samples from the counter may be considered as discrete samples of the moving average of $y(t)$:

$$y_0, y_1, y_2, \dots, y_k = y(t_0, \tau), y(t_0 + T, \tau), y(t_0 + 2T, \tau), \dots, y(t_0 + kT, \tau)$$

where T is the sampling interval and where:

$$\text{ma}(\tau)\{y(t)\} = y(t, \tau) = \frac{1}{\tau} \int_{t-\tau}^t y(\alpha) d\alpha \quad (2-28)$$

defines the causal moving average operator. Note that is meant by “causal” the fact that the operator applied at time t is defined only in terms of the present and past values of the process. The impulse response of the moving average operator $\text{ma}(\tau)\{\}$ is a positive pulse function of duration and of amplitude $1/\tau$. The squared modulus of its transfer function is:

$$|H_{ma}(f)|^2 = \frac{\sin^2(\pi f \tau)}{(\pi f \tau)^2} \quad (2-29)$$

The moving average operator describes the averaging process taking place inside the frequency counter and will be used below for the analysis of the time domain variances.

2.3.2.2.3 Increment operator

The increment operator $\Delta(\tau)\{\}$ is defined as follows:

$$\Delta(\tau)\{x(t)\} = x(t) - x(t - \tau) \quad (2-30)$$

It can be applied recursively. When it is applied once the result is called the “first difference”. When it is applied twice:

$$\Delta^{(2)}(\tau)\{x(t)\} = \Delta(\tau)\{\Delta(\tau)\{x(t)\}\} = x(t) - 2x(t - \tau) + x(t - 2\tau) \quad (2-31)$$

the result is called the “second difference”, and so on.

The squared modulus of the transfer function associated with the increment operator $\Delta(\tau)\{\}$ is

$$|H_{\Delta}(f)|^2 = 4 \sin^2(\pi f \tau) \quad (2-32)$$

Note that the transfer function of the increment operator is equivalent to a first-derivative operation at low Fourier frequencies close to the origin:

$$|H(f)|^2 = 4 \sin^2(\pi f \tau) = 4 (\pi f \tau)^2 \quad \text{for } f \ll \frac{1}{\pi \tau} \quad (2-33)$$

As a consequence, the first difference $\Delta(\tau)\{z(t)\}$ of a power-law process $z(t)$, of a power spectral density $S_{zz}^+(f) = kf^\theta$, yields a process of power spectral density $4k \sin^2(\pi f \tau) f^\theta$ that behaves like a power-law process of exponent $\theta+2$ in the region of low Fourier frequencies close to the origin.

Therefore, for instance, the increment operator will transform a non-stationary flicker power-law process with exponent $\theta = -1$ into a stationary process with $\theta = +1$ behaviour in the region of low Fourier frequencies.

The resulting process is no longer a “power-law” process because the increment operator is not equivalent to the first-derivative operator for all Fourier frequencies but nevertheless it is stationary.

The capability of the increment operator, applied once or several times, to transform a non-stationary power-law process into a stationary process by its high-pass filtering action is fundamental in the definition of time domain frequency stability measurements.

2.3.2.3 Basic time domain measurements

2.3.2.3.1 The true variance

At first sight the logical way to characterise frequency stability of the OUT would be to evaluate the true variance of the frequency samples:

$$I^2(\tau) = E\{y_k^2\} \quad (2-34)$$

The time series y_k is a sampling of the continuous moving average process $y(t, \tau)$; therefore, the mean squared value of the time series is equal to the mean squared value of the $y(t, \tau)$. Using equations (2-28), (2-29) and (2-76) one obtains immediately:

$$I^2(\tau) = \int_0^\infty S_{yy}^+(f) \frac{\sin^2(\pi f \tau)}{(\pi f \tau)^2} df \quad (2-35)$$

which relates the true variance to the power spectral density of $y(t)$.

Unfortunately, the integral in equation (2-35) is infinite for non-stationary power-law noise processes, i.e. for $\alpha \leq -1$.

In other words, the time series y_k coming from the counter is non-stationary since the underlying continuous process $y(t, \tau)$ is non-stationary. The mean squared value of $y(t, \tau)$, which defines the true variance, diverges, i.e. becomes infinite when it is averaged over an infinite time because of the non-stationarity of the models representing the oscillators being characterised.

It is this observation, made in the early days of atomic frequency standards, that led to the definition of other, more appropriate, measures of frequency stability in the time domain.

2.3.2.3.2 The Allan variance

The traditional or classical Allan variance is defined as the average two-sample variance of the frequency samples y_k measured without a dead-time, i.e. with $T = \tau$.

$$\sigma_y^2(\tau) = \frac{1}{2} E\{(y_k - y_{k-1})^2\} \quad (2-36)$$

Observing that the action of taking the difference of two adjacent samples measured without a dead-time is equivalent to a first-difference operation applied over τ it follows that the Allan variance can be defined as the mean squared value of the continuous process $\Delta(\tau)\{y(t, \tau)\}$.

$$\sigma_y^2(\tau) = \frac{1}{2} E\{(\Delta(\tau)\{y(t, \tau)\})^2\} \quad (2-37)$$

In terms of linear operators, the original process $y(t)$ is filtered by a moving average operator, modelling the action of the counter, followed by a first difference operator. Using equations (2-24), (2-29), (2-32) and (2-76) one obtains immediately:

$$\sigma_y^2(\tau) = 2 \int_0^\infty S_{yy}^+(f) \frac{\sin^4(\pi f \tau)}{(\pi f \tau)^2} df \quad (2-38)$$

The above integral is defined and finite for all the power-law processes of Table 3. The process $\Delta(\tau)\{y(t, \tau)\}$, whose mean squared value defines the classical Allan variance, is indeed stationary although the original unfiltered process $y(t)$ was non-stationary. This comes from the high-pass filtering action of the first-difference operator.

If the counter producing the samples y_k has a dead time, then the moving average operation is taken over the gate opening interval τ and the first difference operation is taken over the sampling interval T . Therefore, the same definition in equation (2-36) applied to the frequency samples y_k now yields a different result since by applying again (2-24), (2-29), (2-32) and (2-76) as before, is obtained:

$$\sigma_y^2(\tau, T) = 2 \int_0^\infty S_{yy}^+(f) \frac{\sin^2(\pi f T) \sin^2(\pi f \tau)}{(\pi f \tau)^2} df \quad (2-39)$$

instead of equation (2-38).

The presence of a dead-time in the sampling of the average frequency introduces a bias in the statistics and does no longer yield the true Allan variance when the definition in equation (2-36) is applied to the samples. The bias can be small in some cases. In the case of white noise FM, it is zero. In general, it needs to be considered if it is intrinsic to a measurement system and cannot be avoided. See § 2.3.2.4.2.

2.3.2.3.3 The modified Allan variance

The modified Allan variance was originally defined in terms of the samples x_k of the phase/time process $x(t)$ [Allan *et al.*, 1981]:

$$\text{mod } \sigma_y^2(n\tau_0) = \frac{1}{2} \frac{1}{(n\tau_0)} E \left\{ \left(\frac{1}{n} \sum (x_{i+2n} - 2x_{i+n} + x_i) \right)^2 \right\} \quad (2-40)$$

In order to make easier a comparison with the classical Allan variance, the definition in equation (2-40) can be rewritten in terms of frequency samples:

$$\text{mod } \sigma_y^2(n\tau_0) = \frac{1}{2} E \left\{ \left[\frac{1}{n} \sum_{i=1}^n \left(\frac{1}{n} \sum_{k=1}^n y_{i+k+n} - \frac{1}{n} \sum_{k=1}^n y_{i+k} \right) \right]^2 \right\} \quad (2-41)$$

where the y_k are frequency samples measured without a dead time and averaged over the sampling interval τ_0 .

This new formulation given in equation (2-41) can be interpreted as follows. The inner summations group the elementary samples n by n into super samples averaged over $n\tau_0 = \tau$. The outer summation, on the other hand, is a discrete moving average (dma) operation, performed over the averaging interval $n\tau_0 = \tau$, and applied to the first difference of the super samples.

Therefore, the modified Allan variance can be defined very simply in terms of the underlying continuous processes as follows:

$$\text{mod } \sigma_y^2(\tau) = \frac{1}{2} E \{ [dma(n, \tau_0) \{ \Delta(\tau) \{ y(t, \tau) \} \}]^2 \} \quad (2-42)$$

where the discrete moving average operator applied over the averaging interval $n\tau_0 = \tau$ is defined as:

$$dma(n, \tau_0) \{ z(t) \} = \frac{1}{n} \sum_{i=0}^{n-1} z(t - i\tau_0) \quad (2-43)$$

The squared modulus of the transfer function associated with the discrete moving average operator is:

$$|H_{dma}(f)|^2 = \frac{\sin^2(\pi f \tau)}{n^2 \sin^2\left(\frac{\pi f \tau}{n}\right)} \quad (2-44)$$

where the averaging interval τ is n times the elementary sampling interval τ_0 .

This last result shows that the modified Allan variance is identical to the classical Allan variance except for a supplementary discrete moving average operation.

Using the transfer functions in equations (2-24), (2-29), (2-32) and (2-76) is obtained immediately:

$$\text{mod } \sigma_y^2(n, \tau) = 2 \int S_{yy}^+(f) \frac{\sin^2(\pi f \tau)}{(n\pi f \tau)^2 \sin^2\left(\frac{\pi f \tau}{n}\right)} df \quad (2-45)$$

This is the expression generally found in the literature. The advantage of the modified over the classical Allan variance is the fact that it produces distinct slopes as a function of τ for each of the power law noise processes defined in the polynomial model. The disadvantage, on the other hand, is the fact that the modified Allan variance depends not only on the averaging time τ but also on the number n of elementary samples used to build the interval τ and, indirectly, also on the sampling interval τ_0 .

This is simply due to the fact that, like in any digital processing operation, aliasing occurs when the sampling rate does not satisfy the condition of the sampling theorem, in other words when the sampling rate is smaller than the Nyquist frequency which is defined as twice the bandwidth of the signal.

Assuming, on the other hand, that the sampling rate $1/\tau_0$ is higher than twice the bandwidth B of the process $y(t)$, then the sampling theorem is satisfied and the discrete moving average operator yields approximately the same result as the continuous moving average operator.

It can be verified that when $n \rightarrow \infty$ for constant τ , which is equivalent to allowing the sampling rate $1/\tau_0$ to tend to ∞ , the discrete transfer function in equation (2-44) tends to the continuous transfer function in equation (2-35).

Therefore, by applying the condition of the sampling theorem, i.e. for $1/\tau_0 > 2B$, where B is the bandwidth of $y(t)$, an analogue of the modified Allan variance can be established using the continuous moving average operator in equation (2-28) instead of the discrete moving average operator in equation (2-43), yielding:

$$\text{cmod } \sigma_y^2(\tau) = \frac{1}{2} E \left\{ \left[\Delta(\tau) \left\{ ma^{(2)}(\tau) \{y(t)\} \right\} \right]^2 \right\} \quad (2-46)$$

instead of equation (2-42). Similarly, the modified Allan variance can be established in the spectral domain using the transfer function of the continuous moving average operator in equation (2-29) instead of the transfer function of the discrete moving average operator in equation (2-44), yielding:

$$\text{cmod } \sigma_y^2(\tau) = 2 \int_0^\infty S_{yy}^+(f) \frac{\sin^6(\pi f \tau)}{(\pi f \tau)^4} df \quad (2-47)$$

instead of equation (2-45).

The “continuous” modified Allan variance, noted cmod, is the limit of the modified Allan variance measured without truncation, (see section on truncation below), and in compliance with the sampling theorem.

The results about the cmod Allan variance summarised in Tables 6 to 8 were computed analytically from equation (2-47) and establish that the cmod variance has distinct slopes for different power-law processes and, at the same time, is independent of n , of τ_0 and of the system bandwidth B as shown and verified in [Bernier, 1987].

The cmod variance therefore corrects the two drawbacks of the classical Allan variance: the fact that the slope of the classical variance is the same for both white PM and flicker PM and the dependence on the system bandwidth.

Moreover, the results summarised in Tables 6 to 8 also show that even if the condition of the sampling theorem is not satisfied, the conventional modified Allan variances have the same properties as the cmod Allan variance for all noise processes except white PM on the condition that $n \gg 1$. In the case of white PM the dependence of the modified Allan variance on the system bandwidth is the same as that of the classical Allan variance.

After this important conclusion, let us define yet another formulation of the modified Allan variance, which does not bring new insight, but that is very convenient for the efficient practical estimation of the modified Allan variance.

Define $w(t)$ as the integral of $x(t)$:

$$w(t) = \int_{t_0}^t x(\alpha) d\alpha \quad (2-48)$$

Now, by the use of equation (2-48), the moving average of $x(t)$ can be expressed as:

$$x(t, \tau) = \frac{1}{\tau} \Delta(\tau) \{w(t)\} \quad (2-49)$$

By applying equation (2-48) a second time and by taking the first difference one obtains the continuous modified (cmod) Allan variance, i.e. the modified Allan variance in which the supplementary moving average operator is a continuous operator, yielding:

$$\text{cmod } \sigma_y^2(\tau) = \frac{1}{2} E \left\{ \left(\frac{1}{\tau^2} \Delta^{(3)}(\tau) \{w(t)\} \right)^2 \right\} \quad (2-50)$$

This is the continuous case. In the discrete case, on the other hand, when the samples come out from the counter with a sampling interval τ_0 , if one defines w_k as the discrete sum of all the past x samples,

$$w_k = \sum_{i=0}^k x_i \quad (2-51)$$

then the discrete moving average required for the computation of the discrete modified Allan variance is given by:

$$\frac{1}{n} \sum_{i=0}^{n-1} x_{k-i} = \frac{1}{n} (w_k - w_{k-n}) \quad (2-52)$$

As a consequence, the discrete modified Allan variance can be defined in terms of the third difference of $w(t)$:

$$\text{mod } \sigma_y^2(n\tau_0) = \frac{1}{2} \frac{1}{(n^4 \tau_0^2)} E\{(w_k - 3w_{k-n} + 3w_{k-2n} - w_{k-3n})^2\} \quad (2-53)$$

This last result can be used for the efficient estimation of the modified Allan variance as shown in the section about algorithms below.

2.3.2.3.4 The time interval error

The time interval error (TIE) is a notion used in the predication of time scales and also in telecommunications for the synchronisation of digital networks [Kartaschoff, 1987]. It constitutes a good example of a direct application of the Allan variance.

Assume that at time t two oscillators are both synchronised (time difference set to zero) and syntonised (frequency offset set to zero). The time interval error $TIE(t, \tau)$ is the time error $x(t)$ accumulated between the two oscillators, considered as clocks, at a time $t + \tau$ in the future.

There are different possible ways of defining the TIE depending upon which estimator of the initial frequency offset is chosen. And, as a consequence, the statistical properties of the process $TIE(t, \tau)$ also depend upon which estimator is chosen.

Suppose the frequency offset is estimated by averaging the instantaneous frequency $y(t)$ during the interval τ preceding the instant t . The frequency estimator at time t is therefore $y(t, -\tau)$ and the TIE becomes:

$$TIE(t, \tau) = x(t + \tau) - x(t) - \tau y(t, \tau) \quad (2-54)$$

The first term of equation (2-54) is the free running non-stationary time error random process $x(t)$ at the time $t + \tau$ in the future. The second term is the initial time difference $x(t)$. Subtracting it is equivalent to synchronising the oscillators at time t . The third term is the time difference accumulated after an interval τ because of the initial frequency offset $y(t, \tau)$. Subtracting it is equivalent to syntonising the oscillators at time t .

It can be shown that the TIE defined by equation (2-54) is identical to the second difference of $x(t)$:

$$TIE(t, \tau) = x(t + \tau) - x(t) - \Delta(\tau)\{x(t)\} = \Delta^{(2)}(\tau)\{x(t + \tau)\} \quad (2-55)$$

which is stationary for all the power-law processes found in the polynomial model. Therefore, although the exact behaviour of the TIE cannot be predicted, since it is a random process, its mean squared value can be computed.

The mean squared value of the TIE is equal to the mean square value of the second difference of $x(t)$ which can be expressed as a simple function of the classical Allan variance:

$$E\{TIE^2(t, \tau)\} = 2\tau^2 \sigma_y^2(\tau) \quad (2-56)$$

The mean squared value of the second difference of $x(t)$ is in fact the second structure function of $x(t)$. Structure functions are discussed in the section below about other time domain variances.

Equation (2-26) does not imply that optimum prediction procedures were used – only a particular procedure. If a clock has white FM noise and optimum prediction is used, then the TIE is given by $\tau \sigma_y(\tau)$. In this case the optimum frequency estimate is the mean frequency from the infinite past.

A generalisation of the TIE is reported in [Bernier, 1988] in which a predictor of order n of $x(t)$ is defined for which the mean squared value of the error on the prediction is given by the structure function of order $n + 1$. For example, if the TIE corrects only the effect of the initial frequency offset as in equation (2-55) the predictor is of order 1 and the error on the prediction is the structure function of order 2 given by equation (2-56). If the predictor corrects not only for the initial frequency offset but also for the initial frequency drift, then the predictor is of order 2 and the error on the prediction is the structure function of order 3 and so on. Prediction modelisation of time-scales is an important topic and is treated extensively in the literature [Tavella *et al.*, 1991; Allan, 1987].

2.3.2.3.5 The time variance

The time variance $\sigma_x^2(\tau)$ [Allan *et al.*, 1991] is defined as:

$$\sigma_x^2(\tau) = \frac{1}{3}\tau^2 \text{ mod } \sigma_y^2(\tau) = \frac{1}{6}E\left\{\Delta^{(2)}(\tau)\{\bar{x}(t)\}\right\}^2 \quad (2-57)$$

where $\bar{x}(t)$ is the average time error over an interval τ ending at t .

The time variance $\sigma_y^2(\tau)$ is now used, among others, in the telecommunications industry for specifying network performance. It is a kind of TIE mean squared value estimator similar to the one defined in previous section. More detailed properties of the time variance can be found in [NIST, 1990].

2.3.2.3.6 The dynamic Allan variance

The dynamic Allan variance (DAVAR) [Galleani *et al.*, 2009] is a measure of the variation of the random frequency instability in the time domain, defined as:

$$\sigma_y^2(t, \tau) = \frac{1}{2(T_w - 2\tau)} \int_{t - \frac{T_w}{2} + \tau}^{t + \frac{T_w}{2} - \tau} \mathcal{E}\left[(\bar{y}(t' + \tau) - \bar{y}(t'))^2\right] dt' \quad (2-58)$$

where $\mathcal{E}$ denotes the expected value obtained as an ensemble average over an infinite number of experimental realizations, and T_w is the length of the analysis window. The square-root of dynamic Allan variance is dynamic Allan deviation (DADEV).

An estimate of the DAVAR can be obtained by:

$$\sigma_y^2(n, k) = \frac{1}{2k^2\tau_0^2(N_w - 2k)} \sum_{m=n - \frac{N_w}{2}}^{m=n + \frac{N_w}{2} - 2k - 1} (x_{m+2k} - 2x_{m+k} + x_m)^2 \quad (2-59)$$


where $t = n\tau_0$, $\tau = k\tau_0$, $T_w = N_w\tau_0$, and N_w is assumed to be even.

Experimental evidence shows that the instability of clocks and oscillators, whose standard measure is the Allan variance, can change with time due to several factors, such as temperature, humidity, vibrations, gravitational effects, and radiations. The dynamic Allan variance (DAVAR) measures such variations with time of the clock instability.

Figure 3 shows the average frequency deviation $\bar{y}(t)$ of a white frequency noise whose variance increases suddenly. The DAVAR, whose square root $\sigma_y(t, \tau)$ is shown in the mesh plot, reveals the change of variance in the noise, whereas the Allan variance, whose square root $\sigma_y(\tau)$ is shown in the side plot, averages out the change of variance.

The DAVAR was also introduced in the IEEE Standard for Definitions of Physical Quantities for Fundamental Frequency and Time Metrology – Random Instabilities [IEEE 2022].

FIGURE 3
DAVAR of a change of variance (mesh plot)



2.3.2.3.7 A long-term estimator TheoBR

A long-term estimator for the Allan variance (AVAR) is the Theory variance #1 [Howe *et al.*, 2003], known as Theo1, and with bias removed, known as TheoBR variance [Howe *et al.*, 2006], which is Theo1 with a slight bias removed relative to sample AVAR. Theo1 averages every permissible squared second-difference of time errors x_i in a given data run N_x and reports frequency stability 50% longer than the longest possible τ of AVAR. Thus, Theo1 and TheoBR represent a new species of descriptive statistics with substantially higher equivalent degrees of freedom than the overlapping version of sample AVAR.

Theo1 definition

$$Theo1(m, \tau_0, N_x) := \frac{1}{0.75(N_x - m)(m\tau_0)^2} \times \sum_{i=1}^{N_x - m} \sum_{\delta=0}^{\frac{m}{2}-1} \frac{1}{\left(\frac{m}{2} - \delta\right)} \left[\left(x_i - x_{i-\delta+\frac{m}{2}} \right) + \left(x_{i+m} - x_{i+\delta+\frac{m}{2}} \right) \right]^2 \quad (2-60)$$

for m even, $10 \leq m \leq N_x$, $\tau = 0.75m\tau_0$

Theo1 deviation, known as Theo1-dev, is unbiased relative to the Allan deviation for White-FM noise. Theo1-dev is biased slightly for other noise types. This bias can be automatically removed using an algorithm that computes the average bias between Theo1 and Allan variance in an overlapping τ region of the two and uses it to correct Theo1. Thus, it follows that:

$$TheoBR = \frac{E[*AVAR]}{E[*Theo1]} \times Theo1 \quad (2-61)$$

where *AVAR and *Theo1 are computed for a range of $m\tau_0$ where sample AVAR and Theo1 overlap and are chosen so as to represent the bias at large τ for AVAR rather than at small τ . The bias correction is a constant, not dependent on τ .

TheoBR definition

$$TheoBR(m, \tau_0, N_x) := \left[\frac{1}{n+1} \sum_{i=1}^n \frac{AVAR(m=9+3i, \tau_0, N_x)}{Theo1(m=12+4i, \tau_0, N_x)} \right] \times Theo1(m, \tau_0, N_x) \quad (2-62)$$

where $n = \left\lfloor \frac{0.5N_x}{3} - 3 \right\rfloor$

The TheoBR was also introduced in the IEEE Standard for Definitions of Physical Quantities for Fundamental Frequency and Time Metrology – Random Instabilities [IEEE 2022].

TheoH definition

Hybrid Theory variance, known as TheoH variance [Howe, 2006], is a hybrid statistic (hence, the “H”) that combines values of the overlapping version of AVAR in the short-term and TheoBR variance to estimate the long-term frequency stability out to $\frac{3}{4}$ of the data run length N_x , which as mentioned is 50% longer than the data length usable by AVAR.

$$TheoH(m, \tau_0, N_x) := \begin{cases} AVAR(m, \tau_0, N_x), & \text{for } 1 \leq m \leq \frac{k}{\tau_0} \\ TheoBR(m, \tau_0, N_x), & \text{for } \frac{k}{0.75\tau_0} \leq m \leq N_x - 1, m \text{ even} \end{cases} \quad (2-63)$$

where k is $\tau \leq 0.2T$, and $AVAR(m, \tau_0, N_x)$ has sufficient confidence.

The usual upper τ -limit of AVAR is 20% of the data run time T . Since TheoBR is more computationally intense than AVAR, only large values of m are included in the above definition, not all values of m . Usually it is not necessary to calculate all values of m ; octave or decade spans of m are sufficient to characterize random noise. To delineate AVAR from TheoBR values, an intentional small gap for $\frac{0.2T}{\tau_0} < m < \frac{0.8T}{3\tau_0}$ is inserted but can be filled if one wants to compute more or all m values.

The square-root of TheoH variance is TheoH deviation, also known as TheoH-dev.

2.3.2.3.8 Other time domain measurements

Other time domain measurements are described in the literature: structure functions [Lindsey *et al.*, 1976; Greenhall, 1983], the Hadamard variance [Rutman, 1978], the High-Pass variance [Rutman, 1978] and others.

As shown above for the Allan variance, the principle of time domain variances is always the same: the variance is defined as the mean squared value of a stationary process produced by the action of a linear operator applied to the original non-stationary low-pass frequency process $y(t)$ or phase process $x(t)$.

In particular, the structure function of order n of the process $z(t)$ is defined as the mean squared value of the increment of order n of the process [Lindsey *et al.*, 1976]:

$$D_z^{(n)}(t, \tau) = E \left\{ \left(\Delta^{(n)}(\tau) \{z(t)\} \right)^2 \right\} \quad (2-64)$$

For example, assume a random walk FM noise process. The power spectral density of the frequency $y(t)$ is a power-law of the Fourier frequency f with exponent -2 and the power spectral density of the time function $x(t)$ is a power law with exponent -4 . As shown by equation (2-33), each application of the increment operator increases the exponent by $+2$ at Fourier frequencies close to zero. Therefore, the first increment of $y(t)$ is stationary and the second increment of $x(t)$ is also stationary. The structure functions $D_y^{(1)}(t, \tau)$ and $D_x^{(2)}(t, \tau)$ are therefore independent of t and constitute time domain variances i.e. mean squared values depending on the time difference parameter τ like the classical Allan variance. In fact, the Allan variance is directly related to the second structure function of $x(t)$:

$$\sigma_y^2(\tau) = \frac{1}{2\tau^2} D_x^{(2)}(\tau) \quad (2-65)$$

which explains the result in equation (2-56).

2.3.2.3.9 The multi-variance analysis

The multi-variance analysis should be considered as an estimator of the PSD of $y(t)$ rather than a classical variance measurement. Since the PSD of $y(t)$ may be modelled as a sum of power laws (from f^{-2} to f^{+2} frequency noise), the goal of this method is to measure each noise level. The principle consists in using several variances over the same signal. These different variances must be chosen in order to get an accurate measurement for low frequency noises (Allan variance, Picinbono variance) as well as for high frequency noises (modified Allan variance, Time variance). Moreover, the drifts may also be measured by this method.

The relevant parameters (noise levels and drift coefficients) are determined by a weighted least square method: knowing the response of each variance for each type of noise (see Tables 9 and 10), and weighting each

variance measurement by the inverse of its confidence interval, the noise levels and drift coefficients are calculated as the most probable parameters in the sense of the least squares.

This method is more sensitive than a single variance measurement since the accuracy of the measurement is at least as good as the most accurate variance of the set of variances. The choice of the variance thus determines the sensitivity of the multi-variance analysis: the best set of variance would be a set containing variances specially designed for each relevant parameter. But this choice may be adapted to each case [Vernotte *et al.*, 1995].

The main advantage of the multi-variance analysis is the determination of the uncertainty domain over each parameter measurement. The confidence intervals are estimated from the dispersion of the variance results by using a singular value decomposition. Moreover, since the assessment of the precision leads to associate an eigen vector to each noise level or drift coefficient, it is possible to calculate the angle between each eigen vector. A small angle between two eigen vectors implies that the parameters corresponding to these vectors yield similar effects over the variance responses. In this connection, the separability may be defined as the aptitude of a variance or a set of variances to distinguish between two noises. Thus, the determination of the angles between the eigen vectors corresponding to the parameters yields a quantitative estimation of the separability of the multi-variance analysis and provides a criterion to estimate the measurement reliability. Only the types of noises and drifts that do exist in the signal are then measured. For a full description and measurement results of the multi-variance analysis see [Vernotte *et al.*, 1993] and [Walter, 1992].

2.3.2.4 Pitfalls

2.3.2.4.1 Effect of zero-crossing detection

Time domain measurements are estimated using either phase or frequency samples coming from a digital counter. The counter detects the zero-crossings of the signal which is usually a beat note between the OUT and a reference oscillator.

For the measurement of precision oscillators, it is usually necessary to use a dedicated zero-crossing detector between the signal and the input of the counter.

The reason is that if the slew rate of the signal at the input of the counter is not high enough, the noise in the trigger circuitry of the counter will contribute to the measured Allan variance. Note that, paradoxically, the higher the time resolution of a counter is, the higher the bandwidth of the trigger circuitry and the higher the RMS amplitude of the additive noise is in the trigger stage of the counter.

The design of the dedicated zero-crossing detector must be adapted to the characteristics of the signal to be measured (e.g. amplitude bandwidth, beat note frequency). If the transitions at the output of the zero-crossing detector are sharp enough, then the noise of the trigger circuitry in the counter becomes negligible.

However, it is then the noise contribution of the zero-crossing detector itself that determines the noise floor of the measuring system. The noise due to zero-crossing detection is an additive wideband noise. Therefore, it appears as band-limited white phase noise in the Allan variance measurement, and it is characterised by a τ^{-1} behaviour of the Allan deviation $\sigma_y(t)$.

2.3.2.4.2 Effect of dead-time

As mentioned above, the Allan variance can be defined by equation (2-36) in terms of frequency samples only if there is no dead-time in the sequence of measurements.

The bias in the Allan variance measurements introduced by the dead-time was studied extensively and Tables can be found in the literature [Barnes *et al.*, 1989; Lesage *et al.*, 1979; NIST, 1990]. As mentioned above in the section on the Allan variance, however, the availability of modern time-tagging counters that can sample $x(t)$ and the recent possibility of transforming a conventional time interval counter into a time-tagging counter by post-processing software and an auxiliary “picket fence” reference signal [Greenhall, 1989] make it easy to avoid a dead-time in the measurements.

Besides, grouping neighbouring frequency samples averaged over an interval τ_0 is equivalent to acquire one sample averaged over an interval $n\tau_0$ only if there is no dead-time in the sequence of measurements. The bias

introduced by the grouping of samples in the presence of dead-time is analysed in [Lesage, 1983]. However, the same remark as before applies: recent techniques make it easy to avoid a dead-time.

2.3.2.4.3 Effect of system bandwidth

The Allan variance transfer function in integral equation (2-38) goes to zero as f^{-2} for $f \rightarrow \infty$. Therefore, for a pure power-law frequency process $y(t)$ with exponent $\alpha \geq 1$ the integral does not converge at high Fourier frequencies. Note that this happens only if the process $y(t)$ is a pure power-law process, i.e. if its bandwidth is infinite.

In a real oscillator the bandwidth of the process $y(t)$ is finite because the signal is band-limited and the actual upper bound of integral equation (2-38) is the bandwidth B of $y(t)$. Therefore, the integral is always bounded on the high frequency side.

For $\alpha < 1$ the f^{-2} high-frequency cut-off behaviour of the transfer function is sufficient by itself to limit the bandwidth and to make the integral converge. The Allan variance is therefore independent of the bandwidth of the frequency process $y(t)$.

For $\alpha \geq 1$, on the other hand, the f^{-2} high-frequency cut-off behaviour of the kernel is not sufficient by itself to limit the bandwidth and to make the integral converge. Therefore, the Allan variance is a function of the integral upper boundary, i.e. of the bandwidth B of $y(t)$.

The real-life power-law processes found in oscillators are always low-pass filtered and are well modelled by power-law behaviour only at low Fourier frequencies. For $\alpha \geq 1$, therefore, the Allan variance is a function not only of the bandwidth B but also of the specific shape of the low-pass filter used to limit the bandwidth. This has been taken into account in tables published in the literature [NIST, 1990].

Note, moreover, that the bandwidth B , as it appears in the Allan variance measurement, is usually limited by the low-pass bandwidth of the measuring system used after demodulation rather than by the bandwidth of the original band-pass oscillator signal.

The fact that the Allan variance of a white phase noise process is not only a function of the power spectral density of the white phase noise but also a function of the measuring system bandwidth constitutes a weakness of the classical Allan variance.

On the other hand, the transfer function of the continuous modified Allan variance, as defined by equation (2-47), has a f^{-4} high-frequency cut-off behaviour because there are two cascaded moving average operators in its definition. In this case, then, the transfer function is sufficient to limit the bandwidth for all power-law processes in the polynomial model. The modified Allan variance is therefore independent of the bandwidth for all power law processes in the polynomial model if the condition of the sampling theorem is satisfied.

If the condition of the sampling theorem is not satisfied however, aliasing effects make the modified Allan variance dependent on the bandwidth in the case of white phase noise. For the other power-law processes, including flicker of phase, the high-frequency cut-off behaviour of the transfer function defined by equation (2-45) is sufficient for convergence.

2.3.2.4.4 Truncation effects

The main lobe of both the classical and of the modified Allan variance transfer functions is centred about the Fourier frequency:

$$\frac{1}{2\tau}$$

Therefore, if one attempts to compute the classical or modified Allan variance for averaging times too small i.e. for:

$$\tau < \frac{1}{2B}$$

it happens that the band-pass transfer function of the variance is centred on a Fourier frequency higher than the low-pass bandwidth of the signal $y(t)$.

Therefore, the resulting Allan variance goes rapidly to zero as τ is further decreased. In this region of truncation, the variance depends strongly on the shape of the low-pass filter used to limit the bandwidth of the signal and the value of the variance will be biased low.

Note also that, in the case of the modified Allan variance, if the signal is sampled in compliance with the sampling theorem is obtained:

$$\tau_0 < \frac{1}{2B}$$

Therefore, truncation occurs for small values of n and disappears only for $n \gg 1$ since $\tau = n\tau_0$ [Bernier, 1987].

2.3.2.5 Algorithms

2.3.2.5.1 Frequency averaging by phase sampling

Both the classical and the modified Allan variances can be defined either in terms of frequency samples or in terms of phase samples.

Because $y(t)$ is simply the derivative of $x(t)$ it is very easy to transform any definition that uses one quantity into an equivalent definition in terms of the other quantity.

The two following identities are very useful for the calculation of the time domain variances:

$$ma(\tau)\{y(t)\} = \frac{1}{\tau} \Delta(\tau)\{x(t)\} \quad (2-66)$$

$$\Delta(\tau)\{y(t)\} = \frac{1}{\tau} \Delta^{(2)}(\tau)\{x(t)\} \quad (2-67)$$

Assuming that the phase/time function $x(t)$ is sampled with a sample period τ , the time series

$$x_1, x_2, x_3, \dots, x_n$$

enables one to compute the series of frequency samples averaged over τ using the first difference of the x samples:

$$y_1, y_2, y_3, \dots, y_n$$

where:

$$y_n = \frac{1}{\tau} (x_n - x_{n-1}) \quad (2-68)$$

The second difference of $x(t)$ yields the first difference of the frequency necessary for the computation of the Allan variance.

$$\Delta y_1, \Delta y_2, \Delta y_3, \dots, \Delta y_n$$

where:

$$\Delta y_n = y_n - y_{n-1} = \frac{1}{\tau} (x_n - 2x_{n-1} + x_{n-2}) \quad (2-69)$$

Regarding the hardware, averaging directly the frequency in a digital counter inevitably produces a dead-time since, after a sample is delivered, the counter must wait at least one period of the signal waiting for the next zero-crossing before the next count is started.

On the other hand, if it is the time error function $x(t)$ that is sampled the above results show that the frequency samples averaged over τ without a dead-time can be computed easily, and the value of τ can be selected in the software only constrained by $\tau = n\tau_0$, where n is an integer.

In practice the time error function $x(t)$ can be sampled by time-tagging the zero-crossings of the signal using a “time-tagging” counter. If a time-tagging counter is not available, though, [Greenhall, 1989] has published an algorithm, called the “picket fence” algorithm, for computing directly samples of $x(t)$ using a conventional time interval counter.

2.3.2.5.2 Computation of classical Allan variance

The classical Allan variance can be estimated from N successive samples of $x(t)$, with sampling interval τ_0 , using the following equation:

$$\sigma_y^2(\tau) = \frac{1}{2\tau^2(N-2n+1)} \sum_{i=0}^{N-2n} (x_{k-i} - 2x_{k-i-n} + x_{k-i-2n})^2 \quad \text{where} \quad \tau = n\tau_0 \quad (2-70)$$

This is an “overlapping” estimate over N samples of equation (2-36) computed using the transformation equation (2-68). The formula is causal and therefore usable in real-time calculations. Assume that x_k is the last sample acquired, defining the present, then equation (2-70) involves only past samples in the calculation.

The uncertainty on the true mean square value due to the calculation of the estimation from a finite number of samples was determined in [Lesage *et al.*, 1973] and [Lesage *et al.*, 1976]. Lesage and Audoin results are summarised in [Allan *et al.*, 1988] and [IEEE, 1988].

Note, however, that the reason why the variance is computed from discrete samples is an artefact coming from the use of a digital counter to collect the data. In general, the Allan variance can be defined as the mean square value of the continuous process $\Delta(\tau)\{y(t, \tau)\}$ and it could be also estimated by taking the mean square value of the continuous process observed during a finite observation interval. The uncertainty on the estimation of the Allan variance from the mean squared value of $\Delta(\tau)\{y(t, \tau)\}$ taken over a finite time interval is treated in [Greenhall, 1983]. The effect of sampling the continuous process $\Delta(\tau)\{y(t, \tau)\}$ for the estimation of the Allan variance is treated in [Walter, 1994]. It yields a better confidence for the estimate.

The estimator in equation (2-70) also has the property that for each value of τ only five registers are required in the computer in order to update the Allan variance in a continuous way (see Table 4).

TABLE 4

Registers allocation for the classical Allan variance computation

Register 1	x_i
Register 2	$x_i + n$
Register 3	$x_i + 2n$
Register 4	Sum of variances
Register 5	Number of variances

Assume that the sampling period of $x(t)$ is τ_0 . For each value of $\tau = n\tau_0$ for which one wants to compute the Allan variance, there are five registers to reserve.

From the stream of $x(t)$ samples coming every τ_0 from the counter one picks up one every n samples. The spacing between samples is therefore $\tau = n\tau_0$. From the three values spaced $n\tau_0$ apart, the squared second-difference is computed. Registers 4 and 5 are updated. The contents of register 1 are replaced with the next x_i spaced τ_0 later, etc. for 2 and 3, and registers 4 and 5 are updated again.

Note, finally, that in practice it is always the standard deviation, i.e. the square root of the variance, that is specified and not the variance itself. The practical consequence is that on a log-log plot the slopes associated with the standard deviation are half the slopes associated with the variance.

2.3.2.5.3 Computation of modified Allan variance

At first sight the computational load associated with the modified Allan variance is much larger, especially for large values of n , because of the summation over n in the definition in equation (2-40). In fact, however, the computational load is essentially the same as for the classical variance except for one more register to reserve.

From the modified Allan variance formulation in equation (2-53) an N samples overlapping estimate of the modified Allan variance can be defined:

$$\text{mod } \sigma_y^2(n\tau_0) = \frac{1}{2} \frac{1}{(n^4\tau_0^2)} \frac{1}{N-3n+1} \sum_{i=1}^{N-3n+1} (w_{k-i} - 3w_{k-i-n} + 3w_{k-i-2n} - w_{k-i-3n})^2 \quad (2-71)$$

where the averaging time τ is n times the elementary sampling interval τ_0 .

This last result demonstrates that for each value of τ only six registers are required in the computer in order to update the modified Allan variance in a continuous way (see Table 5).

One common register is required to store w_k i.e. the sum of all the previous x_k coming out from the counter as defined by equation (2-51).

TABLE 5

Registers allocation for the modified Allan variance computation

Register 1	w_i
Register 2	$w_i + n$
Register 3	$w_i + 2n$
Register 4	$w_i + 3n$
Register 5	Sum of variances
Register 6	Number of variances

From the stream of w_k samples updated every τ_0 in the common register, one picks up one every n samples. As before the spacing between samples is $\tau = n\tau_0$. The same procedure is followed as in computing equation (2-70) only the third-difference squared is used instead [Greenhall, 1992].

In conclusion, it can be seen that by defining the time series w_i , which is simply the sum of all the previous x_k , the computation effort to get the modified Allan variance is essentially the same as for the classical Allan variance. Software is available from NIST for computing these variances.

2.3.2.5.4 Summary

For computational purposes, Tables 6 and 7 give symmetrical formulae for the estimation of the classical and modified Allan variances on the basis of either frequency or phase elementary samples. The spectral formulae are also given on the basis of either the phase noise or frequency noise power spectral densities.

TABLE 6

Symmetrical formulae for classical Allan variance

$\sigma_y^2(t) = 2 \int_0^\infty S_{yy}^+(f) \frac{\sin^4(\pi f \tau)}{(\pi f \tau)^2} df$
$\sigma_y^2(t) = \frac{2}{v_0^2} \int_0^\infty S_{\phi\phi}^+(f) \frac{f^2 \sin^4(\pi f \tau)}{(\pi f \tau)^2} df$
$\sigma_y^2(n\tau_0) = \frac{1}{2} \frac{1}{(n^4 \tau_0^2)} \frac{1}{(N - 2n + 1)} \sum_{i=0}^{N-2n} (x_{k-i} - 2x_{k-i-n} + x_{k-i-2n})^2$
$\sigma_y^2(n\tau_0) = \frac{1}{2} \frac{1}{(N - 2n + 1)} \sum_{i=0}^{N-2n} \left[\frac{1}{n} \sum_{j=0}^{n-1} y_{k-i-j} - \frac{1}{n} \sum_{j=0}^{n-1} y_{k-i-j-n} \right]^2$

TABLE 7

Symmetrical formulae for modified Allan variance

$\text{mod } \sigma_y^2(n, \tau) = 2 \int_0^\infty S_{yy}^+(f) \frac{\sin^6(\pi f \tau)}{(n \pi f \tau)^2 \sin^2\left(\frac{\pi f \tau}{n}\right)} df$
$\text{mod } \sigma_y^2(n, \tau) = \frac{2}{v_0^2} \int_0^\infty S_{\phi\phi}^+(f) \frac{f^2 \sin^6(\pi f \tau)}{(n \pi f \tau)^2 \sin^2\left(\frac{\pi f \tau}{n}\right)} df$
$\text{mod } \sigma_y^2(n \tau_0) = \frac{1}{2} \frac{1}{(n^4 \tau_0^2)} \frac{1}{(N - 3n + 1)} \sum_{i=0}^{N-3n} (w_{k-i} - 3w_{k-i-n} + 3w_{k-i-2n} - w_{k-i-3n})^2$
$\text{mod } \sigma_y^2(n \tau_0) = \frac{1}{2} \frac{1}{(N - 3n + 1)} \sum_{i=0}^{N-3n} \left[\frac{1}{n} \sum_{l=0}^{n-1} \left(\frac{1}{n} \sum_{j=0}^{n-1} y_{k-i-j-l} - \frac{1}{n} \sum_{j=0}^{n-1} y_{k-i-j-l-n} \right) \right]^2$

The relationship $S_{yy}^+ = \frac{f^2}{v_0^2} S_{\phi\phi}^+$ used in the transformation is explained simply by the fact that $y(t)$ is the derivative of $x(t)$.

The formulae found in Tables 6 and 7 are derived from equations (2-38), (2-45), (2-70) and (2-71). They are subject to the same restrictions and definitions. Although the variances can be estimated from elementary frequency samples averaged without a dead time over τ_0 , it is important to note that the estimation of the variances using phase samples is less computationally expensive as explained in the text.

2.3.2.6 Applications

Originally the classical Allan variance was introduced as a means to characterise precision frequency sources at a time when spectral analysis at very low Fourier frequencies was not possible. Nowadays it is technically feasible to sample $x(t)$ at a sampling rate higher than the Nyquist rate and to compute numerically from the samples the power spectral density of $x(t)$ or of $y(t)$ at all Fourier frequencies. Given the power spectral density, it is possible to identify the noise processes and to compute every kind of time domain variances.

Nevertheless time-domain variances are still useful tools to characterise oscillators and to identify the noise processes. One reason is the required acquisition and processing time for a given degree of confidence. On the one hand time-domain variances can be updated continuously as new data is acquired. On the other hand, spectral analysis requires the sampling of many slices of data, with each slice having a duration equal to the inverse of the required resolution, before an accurate power spectral density estimation is obtained.

Another reason is that most practical measuring systems sample $y(t)$ or $x(t)$ without respecting the sampling theorem. In this case spectral analysis of the data would be ruined by aliasing while the classical Allan variance is still meaningful.

For the purpose of the identification of the noise processes defined in the polynomial model, the modified Allan variance is superior to the classical Allan variance because it yields different slopes for all the useful noise process models. This is to be compared to the classical Allan variance that yields the same τ^{-2} dependence for both white PM and flicker PM. Moreover, on the condition that the sampling theorem is satisfied, the modified Allan variance is independent of the measurement bandwidth B for all power-law processes while the classical Allan variance is a function of the bandwidth B for both white PM and flicker PM. In the particular case of flicker PM, the modified Allan variance is independent of the system bandwidth even if the sampling theorem is not satisfied.

2.3.2.7 Conversion between time and frequency domains

Table 8 shows the correspondence between the log-log slopes associated with the spectral domain and time domain measurements of power-law processes.

TABLE 8

Correspondence of log-log slopes between spectral and time domains for power-law noise processes

	Slope of log-log plot			
	Spectral domain		Time domain	
	$S_{yy}^+(f)$	$S_{\varphi\varphi}^+(f)$	$\sigma_y^2(\tau)$	$\text{mod } \sigma_y^2(\tau)$
Slope	a	$b = a - 2$	m	m'
Random walk FM	-2	-4	1	1
Flicker FM	-1	-3	0	0
White FM	0	-2	-1	-1
Flicker PM	1	-1	-2	-2
White PM	2	0	-2	-3

As mentioned above the modified Allan variance yields different slopes for the different noise processes. This property facilitates the identification of the processes. This is to be compared with the classical Allan variance that yields the same τ^{-2} slope for both white PM and flicker PM.

On the other hand, the classical Allan variance is a function of the bandwidth B for white PM and flicker PM. The reasons for this dependence are explained above. As a consequence, the Allan variance for a white PM process allows the determination of the mean squared value $N_0 B$ of the process and does not allow the determination of the spectral density N_0 unless the bandwidth B of the system is properly calibrated. On the other hand, Allan variance is independent of the measurement bandwidth B for all power-law processes on the condition that the sampling theorem is satisfied. The modified Allan variance measured under this condition is noted “cmod” in Table 9 and the coefficients A_c to E_c are to be used for the conversion (see Table 10).

TABLE 9

Conversion factors for translation of power-law processes

	$S_{yy}^+(f)$	$\sigma_y^2(\tau)$	$\text{mod } \sigma_y^2(\tau)$	cmod $\sigma_y^2(\tau)$
Random walk FM	$h_{-2} f^{-2}$	$A h_{-2} t$	$A_m h_{-2} t$	$A_c h_{-2} t$
Flicker FM	$h_{-1} f^{-1}$	$B h_{-1}$	$B_m h_{-1}$	$B_c h_{-1}$
White FM	h_0	$C h_0 t^{-1}$	$C_m h_0 t^{-1}$	$C_c h_0 t^{-1}$
Flicker PM	$h_1 f$	$D h_1 t^{-2}$	$D_m h_1 t^{-2}$	$D_c h_1 t^{-2}$
White PM	$h_2 f^2$	$E h_2 t^{-2}$	$E_m h_2 t^{-3}$	$E_c h_2 t^{-3}$

TABLE 10
Coefficients to be used in Table 9

$A = \frac{2}{3}\pi^2$	$A_m = A_c$	$A_c = \frac{11}{20}\pi^2$
$B = 2 \ln(2)$	$B_m = B_c$	$B_c = \frac{1}{8}(27 \ln(3) - 32 \ln(2))$
$C = \frac{1}{2}$	$C_m = C_c$	$C_c = \frac{1}{4}$
$D = \frac{1.038 + 3 \ln(2\pi B\tau)}{4\pi^2}$	$D_m = \begin{cases} D & \text{small } n \\ D_c & \text{large } n \end{cases}$	$D_c = \frac{3}{8\pi^2}(8 \ln(2) - 3 \ln(3))$
$E = \frac{3B}{4\pi^2}$	$E_m = E\tau_0$	$E_c = \frac{1}{8\pi^2}$

If the modified Allan variance is measured without satisfying the sampling theorem, then aliasing occurs. The modified variance measured under this condition is noted “mod” in Table 9 and the coefficients A_m to E_m are to be used for the conversion (see Table 10). Contrary to all the other coefficients in Table 10, these are not exact analytical values. The exact modified Allan variance depends on the sampling period τ_0 and on the number of samples n through equation (2-45). The coefficients shown are asymptotic values based on the assumption that for random walk FM, flicker FM and white FM the aliasing effect is negligible and that the discrete modified variance is equal to the continuous variance given by equation (2-47). For flicker PM it was verified in [Bernier, 1988] by numerical integration that the discrete modified variance behaves like the classical variance for small values of n and like the continuous modified variance for large values of n . In the case of white FM, it was verified in [Bernier, 1987] by numerical integration that the discrete modified variance is equal to the classical Allan variance for $n = 1$ and follows the -3 slope of the continuous modified Allan variance even for small values of n . Conversion factors for non-integer exponent power-law processes are computed theoretically in [Walter, 1994] and fully confirm the coefficients presented here for integer exponents.

The process of translating from the frequency domain to the time domain is an exact one. However, the inverse process, going from the time domain to the frequency domain, is only approximate unless there is only one power-law process [Vernotte, 1993].

2.3.3 Environmental

The above measures can be used to characterise the environment and hence the effect of particular environmental parameters of an oscillator. In recent years the applications of precision oscillators have expanded very rapidly and the question of the characterisation of environmental influences is becoming increasingly important. The physics of environmental influences on frequency standards is treated in [Audoin *et al.*, 1990; Mattison *et al.*, 1976; Papoulis, 1983 and Walls, 1990]. The characterisation and the specification of environmental sensitivities is treated in [Beard *et al.*, 1989; Breakiron, 1989; Brendel *et al.*, 1989; Dragonette *et al.*, 1991; Gagnepain, 1989; Garvey 1989; Sydnor, 1989; Walls, 1990 and IEEE, 1994].

2.4 A bridge to the next Chapter

Most of the formalism and notations used in this chapter are borrowed to the field of electrical engineering and telecommunication theory. The present chapter, therefore, does not follow the conventions normally used by the time and frequency specialists who have developed concepts and notations of their own. This section is a bridge from the notations used in the present chapter and the standard IEEE and ITU-R notations borrowed from the specialised time and frequency formalism that are used in the other chapters.

In the standard time and frequency notation, for example, the autocorrelation function of $\varphi(t)$ is noted $R_\varphi(\tau)$ and there is no distinction between the one sided and the two-sided power spectral densities. Both are noted $S_\varphi(f)$.

Instead of equations (2-3) and (2-5), the standard model of the oscillator noise uses directly the projection of the phasor model on the real axis, i.e. the real signal instead of the analytic one:

$$u(t) = U_0(1 + \varepsilon(t)) \sin(2\pi\nu_0 t + \varphi(t)) \quad (2-72)$$

where $u(t)$ is the oscillator output voltage and U_0 the nominal amplitude.

2.5 Appendix: Random processes

2.5.1 Introduction

In the domain of frequency stability characterisation, the quantities of interest are essentially band-limited, low-pass and band-pass random processes such as $x(t)$, $y(t)$ and $n(t)$ defined above.

2.5.2 Definition of a random process

A random process $x(t)$ is the statistical ensemble $\{\Omega\}$ of all the possible sample functions $x(\omega, t)$ that share the statistical properties that define the random process $x(t)$. The random process is to the sample function what the random variable is to a random number. A good introduction to random processes can be found in [Wozencraft *et al.*, 1965].

The processes $\varepsilon(t)$, $\varphi(t)$, $x(t)$, $y(t)$, $n(t)$ etc. defined in previous sections are all random processes.

2.5.3 Stationary random processes

A random process is stationary if its statistical properties are invariant to a translation of time and have a finite mean. This also means that its statistical properties are independent of the origin of time. A wide sense stationary process has a finite mean for statistical properties, but these are not invariant to a translation in time.

2.5.4 Non-stationary random processes

A random process is non-stationary if its statistical properties are a function of time. The phase noise $\varphi(t)$ and the frequency noise $y(t)$ of oscillators generally contains non-stationary terms. For instance, the phase noise process $\varphi(t)$ usually has flicker or random walk components that makes its mean value and mean squared value undefinable. Only local averages, that are a random function of time, can be defined.

It is precisely the non-stationarity of the oscillator noise processes that ruled out the use of classical statistical methods for their characterisation in the time domain and that led to the development of specific methods such as the Allan variance. In all the above, it should be kept in mind that stationarity or non-stationarity is a property of models and not of oscillators.

2.5.5 Auto-correlation function

The auto-correlation function of a real stationary random process $x(t)$ is defined as:

$$R_{xx}(\tau) = E\{x(t)x(t + \tau)\} \quad (2-73)$$

where $E\{\}$ is the statistical expectation operator. In the case of a non-stationary or wide sense stationary real process the autocorrelation function is also a function of the time t and not only of the time difference τ .

2.5.6 Power spectral density

The double-sided power spectral density of the process $x(t)$ is defined as the Fourier transform of its auto-correlation function.

$$S_{xx}(f) = \int_{-\infty}^{+\infty} R_{xx}(\tau) \exp(-j2\pi f\tau) d\tau = 2 \int_0^{+\infty} R_{xx}(\tau) \cos(2\pi f\tau) d\tau \quad (2-74)$$

The single sided power spectral density $S_{xx}^+(f)$ is the two-sided power spectral density folded over the origin so that only the positive Fourier frequencies are used.

$$S_{xx}^+(f) = \begin{cases} 2S_{xx}(f) & \text{for } f \geq 0 \\ 0 & \text{for } f < 0 \end{cases} \quad (2-75)$$

2.5.7 Linear filtering of random processes

By definition, when a random process $z(t)$ is filtered by a linear operator of impulse response $h(t)$, the output $w(t)$ is the convolution of the impulse response $h(t)$ with the input $z(t)$. It can be shown that the power spectral density of the output $w(t)$ is given by [Wozencraft *et al.*, 1965].

$$S_{ww}^+(f) = S_{xx}^+(f)|H(f)|^2 \quad (2-76)$$

where $|H(f)|^2$ is the squared module of the transfer function $H(f)$ which is defined as the Fourier transform of the impulse response $h(t)$ of the linear operator.

References

- ALLAN D.W. [1966] Statistics of Atomic Frequency Standards. Proceedings IEEE, Vol. 54, 221-230.
- ALLAN D.W. [1987] Time and Frequency (Time-Domain) Characterization. Estimation, and Prediction of Precision Clocks and Oscillators. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, Vol. UFFC-34, no. 6, 647-654.
- ALLAN D.W., BARNES J.A. [1981] A Modified 'Allan Variance' with Increased Oscillator Characterization Ability. Proceedings 35th Annual Frequency Control Symposium. Ft. Monmouth NJ, pp. 470-475.
- ALLAN D.W., HELWIG H., KARTASCHOFF P., VANIER J., VIG J., WINKLER G.M.R., YANNONI N.F., [1988] Standard Terminology for Fundamental Frequency and Time Metrology. Proceedings 42nd Annual Symposium on Frequency Control, pp. 419-425.
- ALLAN D.W., WEISS M.A., JESPERSEN J.L., [1991] A Frequency-Domain View of Time-Domain Characterization of Clocks and Time and Frequency Distribution Systems. Proceedings 45th Annual Symposium on Frequency Control, pp. 667-678.
- AUDOIN C., DIMARCQ N., GIORDANO V., VIENNET J. [1990] Physical Origin of the Frequency Shifts in Cesium Beam Frequency Standards: related Environmental Sensitivities. Proceedings 22nd PTTI Meeting, Vienna (Virginia), pp. 419-440.
- BARNES J.A., ALLAN D.W. [1989] Variances Based on Data with Dead-Time Between the Measurements: Theory and Tables. Technical Note 1318, National Bureau of Standards.
- BEARD R., DANZY F., POWERS E., WHITE J. [1989] Test and Evaluation Methods at the NRL Clock Test Facility. Proceedings 43rd Annual Symposium on Frequency Control, pp. 275-288.
- BEDROSIAN E. [1962] The Analytic Signal Representation of Modulated Waveforms. Proceedings IRE, Vol. 50, 2071-2076.
- BERNIER L.G. [1987] Theoretical Analysis of the Modified Allan Variance. Proceedings 41st Annual Frequency Control Symposium, Philadelphia, pp. 116-121.
- BERNIER L.G. [1988] Linear Prediction of the Non-Stationary Clock Error Function. Proceedings 2nd European Time and Frequency Forum, Neuchâtel CH, pp. 125-137.
- BERNIER L.G., GARDIOL F.E. [1985] The Analytic Signal Representation of Oscillators with Application to Frequency Stability Analysis. Proceedings 39th Annual Frequency Control Symposium, Philadelphia, pp. 127-131.
- BLACHMAN N.M., MCALPINE G.A. [1969] The Spectrum of a High-Index FM Waveform: Woodward's Theorem Revisited. IEEE Transactions on Communication Technology, Vol. COM-17, no. 2, 201-208.
- BREAKIRON L. [1989] Measurement of the Frequency Response of Cesium Clocks to Temperature Fluctuations. Proceedings 43rd Annual Symposium on Frequency Control, pp. 296-297.
- BRENDEL R., EL HASSANI C., BRUNET M., ROBERT E. [1989] Influence of Magnetic Field on Quartz Crystal Oscillators. Proceedings 43rd Annual Symposium on Frequency Control, pp. 268-274.

- DRAGONETTE R.A., SUTER J.J. [1991] Barometric Pressure-Induced Frequency Offsets in Hydrogen Masers. Proceedings 45th Annual Symposium on Frequency Control, pp. 586-590.
- GAGNEPAIN J. [1989] Characterization Methods for the Sensitivity of Quartz Oscillators to the Environment. Proceedings 43rd Annual Symposium on Frequency Control, pp. 242-247.
- GALLEANI L., TAVELLA P. [2009] The Dynamic Allan Variance. IEEE Trans. Ultra. Ferro. Freq. Contr., Vol. 56, no. 3, pp. 450-464.
- GARVEY R. [1989] Testing and Specification of Environmental Sensitivities in Cesium and Precision Quartz Signal Sources. Proceedings 43rd Annual Symposium on Frequency Control, pp. 263-267.
- GREENHALL C.A. [1983] A Structure Function Representation Theorem with Applications to Frequency Stability Estimation. IEEE Transactions on Instrumentation and Measurement, Vol. IM-32, no. 2, 364-370.
- GREENHALL C.A. [1989] A Method for Using a Time Interval Counter to Measure Frequency Stability. IEEE Transactions. on Ultrasonics. Ferroelectrics and Frequency Control, Vol. UFFC-36, no. 5, 478-480.
- GREENHALL C.A. [1992] A Shortcut for Computing the Modified Allan Variance. IEEE Frequency Control Symposium, pp. 262-264
- HOWE D.A., PEPPLER T.K. [2003] Very Long-term Frequency Stability: Estimation using a Special-purpose Statistic. Proc. 2003 Joint Mtg. IEEE Intl. Freq. Cont. Symp. and EFTF Conf., 233-238.
- HOWE D.A., MCGEE-TAYLOR J., TASSET T. [2006] TheoH Bias-removal Method. Proc. 2006 IEEE Intl. Freq. Cont. Symp., 788-792.
- HOWE D.A. [2006] TheoH: a hybrid, high-confidence, statistic that improves on the Allan deviation. Metrologia, 43, S322-S331.
- IEEE [1988] Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology. IEEE Std 1139-1988.
- IEEE [1994] Guide for Measurement of Environmental Sensitivities of Standard Frequency Generators. IEEE Std 1193-1994.
- IEEE [2022] IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology – Random Instabilities. IEEE Std 1139-2022.
- KARTASCHOFF P. [1987] Reference Clock Parameters for Digital Communications Systems Applications. Proceedings 19th PTTI Meeting, pp. 515-549.
- LESAGE P. [1983] Characterization of Frequency Stability: Bias due to the Juxtaposition of Time-Interval Measurements. IEEE Transactions on Instrumentation and Measurement, Vol. IM-32., no. 1, 204-207.
- LESAGE P., AUDOIN C. [1973] Characterization of Frequency Stability: Uncertainty due to the Finite Number of Measurements. IEEE Transactions on Instrumentation and Measurement, Vol. IM-22, 157-161.
- LESAGE P., AUDOIN C. [1976] Corrections to: Characterization of Frequency Stability: Uncertainty due to the Finite Number of Measurements. IEEE Transactions on Instrumentation and Measurement, Vol. IM-25, 270.
- LESAGE P., AUDOIN C. [1979] Effect of Dead Time on the Estimation of the Two-Sample Variance. IEEE Transactions on Instrumentation and Measurement, Vol. IM-28, no. 1, 6-10.
- LESAGE P., AUDOIN C. [1979] Characterization and Measurement of Time and Frequency Stability. Radio Science, Vol. 14, no. 4, 521-539.
- LINDSEY W.C., CHIE C.M. [1976] Theory of Oscillator Instability Based Upon Structure Functions. IEEE Proceedings, Vol. 64, 1662-1666.

- MATTISON E.M. [1990] Physics of Systematic Frequency Variations in Hydrogen Masers. Proceedings 22nd PTI Meeting, Vienna (Virginia), pp. 453-464.
- NIST [1990] Characterization of Clocks and Oscillators. NIST Technical Note 1337.
- PAGNELLI C.J., Cashin W.F. [1991] Measurement of Precision Oscillator Phase Noise Using the Two-Oscillator Coherent Down-Conversion Technique. Proceedings 23rd PTI Meeting, Pasadena CA, NASA Conference Publication no. 3159, pp. 189-207.
- PAPOULIS A. [1983] Random Modulation: A Review. IEEE Transactions Acoustics, Speech, and Signal Processing, Vol. ASSP-31, no. 1, pp. 96-105.
- RILEY W.J. [1990] The Physics of the Environmental Sensitivity of Rubidium Gas Cell Atomic Frequency Standards. Proceedings 22nd PTI Meeting, Vienna (Virginia), pp. 441-452.
- RUTMAN J. [1978] Characterization of Phase and Frequency Instabilities in Precision Frequency Sources: Fifteen Years of Progress. Proceedings IEEE, Vol. 66, no. 9, 1048-1075.
- STEIN S.R. [1985] Frequency and Time – Their Measurement and Characterization. Precision Frequency Control, Vol. 2, edited by E.A. Gerber and A. Ballato, Academic Press, New-York, 1985, pp. 191-416.
- SYDNOR R. [1989] Environmental Testing at the Jet Propulsion Laboratory's Frequency Standards Laboratory. Proceedings 43rd Annual Symposium on Frequency Control, pp. 242-247.
- TAVELLA P., PREMOLI [1991] Characterization of Frequency Standard Instability by Estimation of their Covariance Matrix. Proceedings PTI Meeting.
- VERNOTTE F., LANTZ E., GROSLAMBERT J., GAGNEPAIN J.J. [1993] Oscillator noise analysis: multi-variance measurement. IEEE Transactions on Instrumentation and Measurement, Vol. IM-42, no. 2, 342-350.
- VERNOTTE F., MCHUGH M., ZALAMANSKY G. [1995] Cut-off frequencies and noise power law model of spectral density: adaptation of the multi-variance method using the structure function approach. Proceedings EFTF 95, Besançon, France, pp. 373-376.
- WALLS F.L. [1990] Environmental Sensitivities of Quartz Crystal Oscillators. Proceedings 22nd PTI Meeting, Vienna (Virginia), pp. 465-486.
- WALLS F.L. and DEMARCHI A. [1975] RF Spectrum of a Signal After Frequency Multiplication; Measurement and Comparison with a Simple Calculation. IEEE Transactions on Instrumentation and Measurement, IM-24, pp. 210-217.
- WALTER T. [1992] A multi-variance analysis in the time domain. Proceedings 24th PTI.
- WALTER T. [1994] Characterizing Frequency Stability: A Continuous Power-Law Model with Discrete Sampling. IEEE Transactions on Instrumentation and Measurement, Vol. IM-43, no. 1, 69-79.
- WOZENCRAFT J.M., JACOBS I.M. [1965] Principles of Communication Engineering, Wiley & Sons, New-York.

ITU Texts

Handbook Satellite Time and Frequency Transfer and Dissemination

Recommendation ITU-R TF.583 – *Time codes*

Recommendation ITU-R TF.686 – *Glossary and definitions of time and frequency terms*

Recommendation ITU-R TF.768 – *Standard frequencies and time signals*

Recommendation ITU-R TF.538 – *Measures for random instabilities in frequency and time (phase).*

CHAPTER 3

Local frequency and time sources

TABLE OF CONTENTS

	<i>Page</i>
3.1 Introduction	42
3.2 Quartz crystal frequency standards.....	42
3.2.1 The resonator.....	42
3.2.2 The oscillator.....	43
3.2.3 Frequency instability of quartz crystal oscillators.....	44
3.3 General introduction to AFS.....	45
3.4 Performance characteristics of traditional AFS	46
3.4.1 Instability of atomic frequency standards	46
3.4.2 Accuracy of AFS.....	47
3.4.3 Reproducibility, resettability	48
3.5 The rubidium gas cell frequency standard.....	49
3.6 The hydrogen maser	51
3.7 The caesium beam AFS	54
3.8 Miniaturized atomic frequency standards.....	57
3.8.1 Getting rid of the discharge lamp	57
3.8.2 Getting rid of the microwave resonator.....	57
3.8.3 Getting rid of the issue of double-resonance excitation	58
3.9 Advanced frequency standards based on laser-cooling and optical transitions	58
3.9.1 Atomic fountains.....	58
3.9.2 Other advanced AFS under development.....	58
3.9.3 Optical frequency standards	59
3.9.4 Applications of advanced frequency standards	59
References	59
ITU Texts	62

Overview of Chapter 3

Whereas Chapter 1 covers the basic physics and operational concepts of precision frequency standards, Chapter 3 covers time and frequency sources in more detail. The user, considering environmental conditions and local constraints and opportunities, will be shown in this Chapter different options for obtaining sources of time and frequency.

Typically, the sources described here are frequency sources. None of them will, intrinsically, deliver time-of-day information, and synchronization to a reference time scale, such as UTC needs to be provided externally. In this regard, it is useful to appreciate that the frequency sources discussed here are tied to fundamental physics, from which the unit of time and frequency (the SI second) is defined. In contrast, time, which is the integral of frequency, has as its constant of integration an artefact number agreed upon by international standards bodies or as determined by the user. From such is obtained UTC, UTC(k), or some other self-consistent time scale as detailed in Chapter 5. The dissemination of time and frequency is topic of Chapter 6.

Chapter 3 may seem, in part, repetitious with material in Chapter 1, but the intent here is to help the user decide which of the sources given may best fill the users' needs and requirements.

Local frequency and time interval sources

3.1 Introduction

Precision time and frequency standards based on mechanical resonances in quartz material and atomic resonance phenomena are widely used in scientific and technological applications including data networks, telecommunication systems, navigation, metrology, the world's timing systems, and scientific research such as astronomy and spectroscopy. Chapter 1 has already explained the basic principle of these devices and in Chapter 2 the terminology to describe their performance was laid down. Here, each device is described in a separate section in some detail (see [Vanier and Audoin, 1989], and [Vanier and Tomescu, 2016] for a full account including recent developments). Understanding the physics will better help the user see how a particular standard may perform in a given environment. Tables are added which contain relevant performance data. The most severe limitations to performance due to environmental effects are mentioned. The largest part of this Chapter deals with atomic frequency standards (AFS) with atomic resonance frequencies in the microwave range. So-called optical frequency standards are briefly discussed in § 3.9.3 and in detail in Chapter 9.

3.2 Quartz crystal frequency standards

Quartz crystal oscillators are the workhorses in time and frequency metrology. The piezo-electric effect of quartz (SiO_2) is employed to excite mechanical vibrations by applying electric oscillating fields to a piece of suitably cut single crystal of the material which constitutes the crystal resonator. The material is very convenient. It can be grown in large quantities, at low cost, with sufficient chemical purity and high lattice perfection. It exhibits low electrical losses.

Conceptually, the resonance frequency of the resonator is used to discipline the electric oscillator. Thus, the components are described one-by-one. A detailed technical and historical survey is available from [Lombardi, 2011].


3.2.1 The resonator

Mechanical deformations in the quartz material, such as flexure, extension, and shear, can be used. They may propagate in the bulk of the material and resonance occurs when the acoustic waves of a given type fulfil boundary conditions. The quartz resonator design is aimed at minimizing the perturbation of the intrinsic properties by environmental factors. The “cut” – the orientation of the piece of quartz with respect to the crystallographic axes – is chosen such that the influence of temperature and stress upon the oscillating frequency is minimized. The mechanical properties of the quartz plate are used in an electrical circuit thanks to the piezoelectric effect, which couples electric fields and mechanical displacements in non-centrosymmetric crystals (such as quartz). Figure 4 shows the electrical equivalent circuit of a piezoelectric resonator. In frequency and time metrology applications, the resonance frequency ν_0 is usually 5 or 10 MHz. In other

applications, it can be as high as 1 GHz. The quality factor Q varies inversely with the resonance frequency. Typically, one has $Q = \frac{K}{\nu_0}$ with $K \approx 10^{13}$ Hz.

FIGURE 4

Electrical equivalent circuit of a piezoelectric resonator



HDB-31-4

In an advanced design, called BVA for “Boîtier Vieillissement Amélioré” [Besson, 1977], the electrodes are not deposited on the faces of the vibrating plate, as is most often done. The electrodes are placed on auxiliary plates, at a distance of a few micrometres from the active one. Then, the mass loading of the main plate by the electrodes is suppressed, as well as any stress release and material migration at the interface between quartz and metal.


3.2.2 The oscillator

The resonator is inserted in a feedback loop, as shown schematically in Fig. 5. In steady state, the closed loop phase shift is 0 or π and the oscillation amplitude is a constant. For a given phase perturbation, the related fractional frequency change is proportional to $1/Q_L$, where Q_L is the loaded quality factor of the resonator in the network. Typically, $Q_L \sim 10^6$ at 5 MHz. Two additional capacitors enable the adjustment of the frequency of the oscillator, a fixed capacitor and a varactor (summarized as C_L in Fig. 4). The latter is used for fine tuning with an external DC voltage source. This feature is needed in all applications where quartz oscillators are steered by an atomic reference, i.e. in all devices discussed further on.

In all high-quality quartz crystal frequency standards, the resonator and other temperature sensitive elements are placed in a controlled oven whose temperature is set as closely as possible at a point where the resonator frequency does not depend on temperature.

FIGURE 5

Schematic representation of a quartz crystal oscillator, A: amplifier, Z: impedance



HDB-31-5

3.2.3 Frequency instability of quartz crystal oscillators

A detailed discussion of the measures of instability is contained in Chapter 2. A concise treatment of the properties of quartz oscillators and of phase noise can be found in Chapter 4 and in [Rubiola, 2009].

3.2.3.1 Random frequency fluctuations

Close to the carrier, i.e. at Fourier frequencies smaller than approximately the resonator half-bandwidth, $\nu_0/2Q$, quartz crystal oscillators show flicker noise of frequency, which originates in the quartz resonator and in the circuitry of the feedback loop. Then the phase noise power spectral density (PSD) varies as f^{-3} . This noise determines the frequency stability for sampling times between about 1 and 10^3 s. In that range, the time domain frequency stability measure, $\sigma_y(\tau)$ (see Chapter 2), is a constant. Typical values for a 5 MHz state of the art quartz crystal oscillator are $\sigma_y(\tau)$ between 8×10^{-14} and 3×10^{-13} in this region of τ .

Lower quality quartz crystal oscillators are also used in many digital devices and are also used in all computers to maintain the time of day and for other purposes such as scheduling a process to run at some specific time. These oscillators generally exhibit flicker phase noise for sampling times less than a few hundred seconds. The time deviation, $\sigma_x(\tau)$, is typically about $0.5 \mu\text{s}$ at these averaging times. Flicker frequency noise is the dominant noise type for averaging times larger than about 10^3 s.

At Fourier frequencies larger than $\nu_0/2Q$, the oscillator is perturbed by flicker phase noise until it becomes smaller than the white phase noise. The flicker phase noise (PSD varying as f^{-1}) is a result of the processing of the flicker noise of the frequency in the oscillator feedback loop.

At Fourier frequencies larger than approximately 100 Hz, for 5 MHz quartz crystal oscillators, the phase noise is white noise. It results from the noise added in the amplifying output circuit. The phase noise PSD within the bandwidth of the output stage is given by:

$$S_\phi(f) \approx \frac{FkT}{p} \quad (3-1)$$

where k is the Boltzmann constant, T is the thermodynamic temperature, P the power generated in the oscillation loop and F is the noise factor of the output amplifier.

3.2.3.2 Systematic effects

For sampling times larger than $\approx 10^3$ s in the case of a 5 MHz quartz crystal oscillator, the frequency stability is perturbed by ageing effects. They yield a slow fractional frequency change of the order of 10^{-11} per day. This rate is not constant and can vary by almost an order of magnitude over a period of a few days. These


ageing effects are caused by the change of the properties of the electronic components, the stress relief and the diffusion of impurities in the quartz crystal, etc. Furthermore, the frequency of quartz crystal oscillators is affected by a change of the external conditions, such as supply voltage, temperature, humidity, pressure and vibrations. The measured frequency change is the sum of these effects, and the sign of the change occasionally reverses.

3.3 General introduction to AFS

The small value of $E_2 - E_1$ that is true for all AFS with clock frequencies in the microwave region implies that $(E_2 - E_1) \ll kT$, where kT is the thermal energy. Therefore, at thermal equilibrium, the two energy levels are almost equally populated. It follows that this thermal equilibrium must be broken in order to be able to observe a change in the atomic properties or an exchange of a detectable amount of energy when the atomic transition occurs. Two different methods are implemented to overcome thermal equilibrium. The first one is related to the fact that the atoms show opposite values of their effective magnetic moment whether they occupy one or the other of the two energy levels that are considered. Consequently, the deflection of atoms in a beam passing an inhomogeneous and intense magnetic field depends on their internal state and they can be separated [Gerlach and Stern, 1924; Gerlach, 1925]. This method of state selection is used in caesium beam AFS and in hydrogen masers. It is not very efficient because it discards atoms that are not initially in the proper clock state, where “proper” depends on the chosen geometry. For example, there are 16 equally populated magnetic sub-levels in the ground state of caesium 133, and magnetic selection immediately discards one half. The remaining ones are in different magnetic sub-states, so that finally only about 1/8 of the atoms in the beam after state selection are available for the microwave interrogation of the so-called clock transition.

The other method is based on optical pumping [Kastler, 1950]. In addition to the clock transition that is based on the hyperfine interaction levels 1 and 2 in the ground state, the atoms have excited levels of much larger energy, such as level 3 shown in Fig. 6. An optical radiation, having an appropriate wavelength, can transfer atoms from level 1 to level 3, for instance. Spontaneous decay back to both levels 1 and 2 occurs in an extremely short time, typically < 50 ns. An atom that decays back to level 1 is pumped to level 3 again. But an atom that decays to level 2 is effectively trapped there because the natural lifetime of level 2 is very long. The net result of optical pumping is that the population of one of the ground-state levels (level 2 in the example given) increases to the detriment of that of the other level (level 1). This is the principle of state preparation in the commercially available rubidium AFS and in the optically pumped caesium beam AFS. When all the atoms have been pumped into level 2, no further absorption takes place, and the atoms become transparent to the pump radiation. If an external frequency induces transitions from level 2 back to level 1, then the atoms start absorbing the pumping radiation again, and the transmitted intensity of the pump radiation decreases.

FIGURE 6
Principle of optical pumping



The clock transition frequency is determined through Bohr's relation:

$$h\nu_0 = E_2 - E_1 \quad (3-2)$$

where h is Planck's constant. The width of the observed atomic resonant response, in an active as well as in a passive atomic frequency standard which can be traced back to Heisenberg's uncertainty relation:

$$\Delta\nu \Delta t \geq 1 \quad (3-3)$$

In all microwave atomic frequency standards in use the natural lifetime of the clock state is very long, so, Δt is determined by the duration of the interaction between the atomic species and the probing field. Details will be given for the individual standards below. Before a general chapter describes the measures of the performance of the standards.

3.4 Performance characteristics of traditional AFS

3.4.1 Instability of atomic frequency standards

It is quite common to distinguish the short-term and medium-term frequency stability, which are determined by random processes, and the long-term stability which depends on systematic effects.

3.4.1.1 Short-term frequency stability

In both categories of atomic frequency standards, active and passive, the short-term frequency stability is that of the quartz crystal oscillator. This occurs for sampling times τ smaller than the control loop time constant, which is typically of the order of 1 – 10 s.

3.4.1.2 Medium-term frequency stability

For τ larger than the control loop time constant ($\approx 1 - 10$ s), up to a limit which depends on the Atomic Frequency Standard (AFS) considered (≈ 1 day for caesium AFS, $\approx 10^4$ s for hydrogen AFS, $\approx 10^3$ s for rubidium AFS), the frequency stability is determined by the quality factor of the atomic resonance and by the signal to noise ratio of its observation.

3.4.1.2.1 Medium-term frequency stability in passive frequency standards

The frequency stability is directly related to the precision of the measurement of the extremum of the atomic resonance pattern. It can be shown that the time-domain frequency stability measure $\sigma_y(\tau)$ is given by [Vanier and Audoin, 1989]:

$$\sigma_y(\tau) = \frac{1}{Q \frac{S}{N(\tau)}} \quad (3-4)$$

where $S/N(\tau)$ is the amplitude signal to noise ratio for the sampling time τ and Q the atomic line quality factor.

In the caesium beam and the rubidium cell frequency standards, the noise is the shot noise of the flux of particles detected, atoms or photons, respectively. It is obtained:

$$\frac{S}{N(\tau)} \approx \frac{I_s}{I_t^{1/2}} \tau^{1/2} \quad (3-5)$$

where I_s is the particle flux measuring the line amplitude and I_t is the total flux of particles detected. It is thus obtained:

$$\sigma_y(\tau) = \frac{I_t^{1/2}}{Q I_s} \tau^{-1/2} \quad (3-6)$$

The frequency stability varies as $\tau^{-1/2}$, which is related to the fact that the perturbing frequency noise is white noise (random and uncorrelated). In the rubidium cell frequency standard, usually $I_s \ll I_t$. In the commercial caesium beam frequency standard, I_s is almost equal to I_t and equation (3-6) gives:

$$\sigma_y(\tau) \approx \frac{1}{Q I_s^{1/2}} \tau^{-1/2} \quad (3-7)$$

In the passive hydrogen maser, it follows that:

$$\frac{S}{N(\tau)} \approx \frac{1}{K} \left(\frac{P}{kT} \right)^{1/2} \tau^{1/2} \quad (3-8)$$

where kT is the thermal noise energy in the cavity mode and P the power delivered to the cavity by the atoms. K is a constant, larger than unity, which depends on the operating parameters. It is then obtained, in that case:

$$\sigma_y(\tau) \approx \frac{K}{Q} \left(\frac{kT}{P} \right)^{1/2} \tau^{-1/2} \quad (3-9)$$

3.4.1.2.2 Medium-term frequency stability in active frequency standards

In all types of oscillators, one distinguishes schematically between the frequency noise generated in the oscillation sustaining circuit and the frequency noise added to the oscillation in subsequent amplifying circuits. The noise generated in the oscillation feedback loop is white frequency noise (in quartz crystal oscillator, this white frequency noise is usually masked by flicker noise FM). The related time-domain frequency stability measure is given by [Vanier and Audoin, 1989]:

$$\sigma_y(\tau) \approx \frac{1}{Q} \left(\frac{kT}{2P} \right)^{1/2} \tau^{-1/2} \quad (3-10)$$

In active hydrogen masers, this white frequency noise predominates the frequency stability for τ ranging between about 20 and 10^4 s. The noise added in the amplifying output circuit is white phase noise, as shown by equation (3-1). It is obtained:

$$\sigma_y(\tau) \approx \frac{1}{v_0} \left(\frac{3FkTf_h}{2\pi P} \right)^{1/2} \tau^{-1} \quad (3-11)$$

where f_h is the noise bandwidth of the frequency stability measurement equipment.

3.4.1.3 Long-term frequency stability

For τ larger than ≈ 1 day for caesium AFS, $\approx 10^4$ s for hydrogen AFS and $\approx 10^3$ s for rubidium AFS, the frequency stability measure $\sigma_y(\tau)$ decreases with reduced slope and reaches a constant “flicker floor” value. It may even increase with sampling time τ at longer times.

There is no theory that relates the long-term frequency changes to a fundamental physical effect such as shot noise or thermal noise. Instead, they are the result of slow changes of a number of frequency offsets, given in the detailed descriptions of the AFS, which perturb the idealized resonance frequency defined by equation (3-2). Most of these frequency offsets depend on the way the atomic transition is observed and thus on the design or type of frequency standard. Their change is determined by ageing effects and by variations of the environmental conditions. The rate of ageing and the sensitivity to external perturbations are directly related to the design choices made.

3.4.2 Accuracy of AFS

3.4.2.1 Residual frequency offsets

Although one strives to design atomic frequency standards such that the ideal condition of atoms isolated and at rest is satisfied as much as possible, small frequency offsets cannot be avoided in current practice. The uncertainty on our knowledge of the magnitude of these frequency offsets determines the accuracy of atomic frequency standards. The physical origin of some of these frequency offsets is common to all types of frequency standards, e.g. the second order Doppler effect due to thermal motion of the atoms. Others are specific to a given type, e.g. the buffer gas frequency offset and the light shift in a rubidium cell and will be dealt with in the respective following sections.

3.4.2.2 Definition

In general, the accuracy expresses the agreement between the result of a measurement and the true value of the measurand (see Glossary, preceding Chapter 1). For a quantitative statement, it is expressed as the cumulative normalized uncertainty of the value realized relative to that given by the definition. Usually, this

uncertainty has two different components. In the case of atomic frequency standards, one is related to random perturbations affecting the measurement of residual frequency offsets (e.g. that due to the measurement of the magnetic field applied). The other is a systematic component that is related to a lack of knowledge of some operating parameters (e.g. the inhomogeneity of the magnetic field). These two components must be combined with care [BIPM, 2008]. In general, the methods that are used to evaluate systematic contributions have uncertainties as well, which further complicates the overall uncertainty budget.

3.4.2.3 Primary and secondary frequency standards

Laboratory caesium beam frequency standards are purposely built to realize as accurately as possible the definition of the second. They are usually termed primary clocks as they provide the primary realization of the unit of time with highest accuracy¹.

Their accuracy is the normalized uncertainty of the measured or estimated frequency difference between the realized value of the hyperfine transition frequency and the unperturbed transition frequency given in Table 1. The best values achieved to date (2025) with devices employing a thermal atomic beam are about 1×10^{-14} , employing laser-cooled atoms almost 1×10^{-16} (see § 3.9) have been reported. The quoted values are based on an enumeration of the magnitudes of known systematic effects and an estimate of the uncertainty of the contribution of each of these offsets to the frequency of the device. Although the value of the unperturbed hyperfine transition frequency in hydrogen, ν_H , is known with a rather small relative uncertainty, of the order of 2×10^{-12} , the frequency realized by a practical hydrogen maser has systematic offsets caused by the design of the resonant cavity and other factors. These offsets often change slowly with time, so that the AHM output frequency can have both significant offset and drift.

The accuracy of rubidium gas-cell AFS is comparatively poor because of various effects, as described below. The output frequency must be calibrated against that of a more accurate atomic frequency standard. The rubidium gas-cell frequency standard is thus used as a secondary frequency standard.

Quartz crystal oscillators may have a relative offset of the order 10^{-6} from nominal without calibration. Because of this, quartz crystal oscillators are almost always calibrated (e.g. when used as internal reference in synthesizers and time interval counters) against a standard frequency and are used as secondary frequency standards.

3.4.3 Reproducibility, resettability

The concept of reproducibility has two aspects (see Glossary):

- a) with respect to a set of independent devices of the same design, it is the ability of these devices to produce the same value;
- b) with respect to a single device, put into operation repeatedly without adjustments, it is the ability to produce the same value.

The resettability (see Glossary) is the ability of a device to produce the same value when specified parameters are independently adjusted under stated condition of use.

It is worth mentioning that the accuracy figure of a given atomic frequency standard (e.g. 3×10^{-12}) is an upper bound of its reproducibility and of its resettability (e.g. 5×10^{-13}) as well as of its long-term frequency stability (e.g. 2×10^{-12} for the life of the device).

¹ Note that in ITU-T a different use of the term “primary clock” is common.

3.5 The rubidium gas cell frequency standard


Frequency standards involving optically pumped atomic vapour in closed cells have been successfully used in numerous applications where the requirements regarding frequency instability in the short and long term cannot be fulfilled with quartz oscillators, but where constraints on space, power consumption, or cost prevent the use of other types of atomic frequency standards. The rubidium gas cell frequency standard is an example. Development of this kind of frequency standard started in the late 1950s in the USA, and the first commercial model became available in the 1960s. The history was recently compiled by [Riley, 2019]. The reference transition used is the ^{87}Rb hyperfine transition at $\nu_0 = 6.84$ GHz. The simplified level scheme is shown as Fig. 7 a).

In the rubidium gas cell atomic frequency standard (in short: Rb AFS) the state selection and detection of the clock transition is achieved using an optical-pumping method. Optical excitation on the D2 line to the $^2\text{P}_{3/2}$ level (E_3 in Figs 6 and 7 a)) requires a radiation source at 780 nm wavelength. Light from a discharge lamp (filled with the isotope ^{87}Rb) passes a filter cell (called the hyperfine filter) which contains ^{85}Rb isotope vapour before it excites ^{87}Rb atoms in a cell filled with buffer gas inside a microwave resonator. Atoms excited to the E_3 states in the cell then relax to both ground state levels either by spontaneous emission or by collisions with the buffer gas. The buffer gas, a mixture of inert gases featuring elastic collisions with the Rb atoms, increases the interaction time, T , of the atoms with the microwave field sustained in the resonator by reducing the (inelastic) collision rate of the atoms with the walls of the cell.

The action of the isotopic filter cell alters the spectrum of the light allowing only ^{87}Rb atoms in the lower hyperfine level (E_1 in Figs 6 and 7 b)) to be optically pumped. This level is thus depopulated, and the cell becomes transparent. If microwave radiation at frequency ν_p close to ν_0 is applied to the atoms, the lower hyperfine level is again populated, and optical absorption starts again. In resonance, $\nu_p = \nu_0$, the signal at the photo detector shows a minimum, as illustrated in Fig. 7 c). The spectral width of the resonance feature is typically 500 Hz due to limits on T caused by both buffer gas and residual wall collisions. The signal at ν_p is synthesized from a VCXO and modulated. Synchronous detection yields a control signal, U_R , to steer the VCXO. The centre frequency of the resonance line deviates considerably from the value of ν_0 of unperturbed ^{87}Rb atoms. Frequency shifts are caused by the presence of magnetic fields, collisions with the buffer gas, and the simultaneous interaction of atoms with optical and microwave radiation. The total shift may reach about 10^{-9} in relative units and the gas composition and light intensity is initially adjusted so as to yield optimum performance with a minimum of environmental sensitivity. The latter two effects, however, are *per se* not stable in time, as the lamp ages (change in its spectrum and its intensity) and the gas composition in the filter cell and the resonance cell itself change somewhat with time due to diffusion and outgassing. If helium is present in the environment, it can diffuse into the filter cell and significantly change the performance of the device. Both effects also change with the ambient temperature and temperature gradients.

FIGURE 7

a) Energy levels of ^{87}Rb ; b) Rubidium gas cell frequency standard, schematic representation;
c) Detector signal I_D versus frequency ν_p of the microwave radiation



HDB-31-7

As the individual contributions to offsets in the output frequency may become large and are difficult to trace, the Rb AFS cannot be considered as a primary frequency standard. Instead, it must be calibrated against a known frequency reference. In addition, these standards tend to be more environmentally sensitive than caesium and hydrogen standards. If calibrated by the manufacturer or by a standards laboratory, Rb AFS can transport frequency at about the 1×10^{-11} level of uncertainty. Their intrinsic short-term stability, however, can often be better than that of caesium standards because of the higher signal to noise ratio available in many of them. Vibrations and AC magnetic fields should be avoided during operations to maintain this good short-term frequency stability.

Rb AFS are available at different levels of performance, which basically depend on the size of the units and the sophistication of the control of operational parameters. In Table 11 small module-size units (volume less than 10^{-3} m^3) – here the hyperfine filter is often integrated into the resonance cell, containing natural Rb– and instrument size units (typical. volume of 10^{-2} m^3) from various manufacturers are compared. Space qualified versions are available that serve as frequency references in many space applications, in particular on board of GNSS satellites. A recent review on the topic was provided by [Schmittberger, Marlow and Scherer, 2021].

A nowadays wide-spread use of Rb AFS is in combination with a GNSS receiver. Output frequency and 1PPS output (if available) are controlled with reference to the received GNSS signals with a long time constant ($> 10\,000$ s). Applications of GNSS disciplined oscillators have been discussed in [Lombardi, 2016; EURAMET, 2016; Piester *et al.*, 2019] and the subject is taken up in Chapter 8.

TABLE 11

Performance data of rubidium AFS

Type	A	B	C	D	E	F	G
Power source	AC+DC	4.5 – 32 V DC	15 V + 5 V	24 V	12 or 24 V	24 V	AC + DC
Frequency offset at delivery	5×10^{-11}	5×10^{-11}		5×10^{-11}	5×10^{-11}	5×10^{-11}	5×10^{-11}
Temperature sensitivity over temperature range		1×10^{-10} , 80 K	3×10^{-10} , 60 K	1×10^{-10} , 60 K	2×10^{-10} , 60 K	1×10^{-10} , 80 K	3×10^{-10} , 50 K
Magnetic field sensitivity		7E-11/G, max 2 G		1E-10 / G	1E-10 / G	2E-10 / G	4E-11/ G, max 2 G
Aging/month	5×10^{-11}	5×10^{-11}	2×10^{-10}	3×10^{-11}	3×10^{-11}	5×10^{-11}	5×10^{-11}
Aging/year	5×10^{-10}	5×10^{-10}	2×10^{-9}			$< 10^{-10}$	5×10^{-10}
ADEV 1s	2E-11	5×10^{-11}	1.4×10^{-11}	5×10^{-11}	3×10^{-11}	2×10^{-11}	3×10^{-11}
ADEV 10s	1e-11	1.6×10^{-11}		1×10^{-11}	8×10^{-12}	1×10^{-11}	1×10^{-11}
ADEV 100s	2e-12	5×10^{-12}		3×10^{-12}	3×10^{-12}	2×10^{-12}	3×10^{-12}
Interface	No	RS232	No	RS232	RS232	RS232	RS232
1 PPS output	Yes	Yes	No	No	Yes	Yes	Yes
Power (Watt) normal ops./ warm-up	10	10 / 14	8 / 32	10 / 25	10 / 30	15 / 50	25 / 45
Warm up (minutes)		8	4	15	12	7	8

3.6 The hydrogen maser

The ground state hyperfine splitting of the hydrogen atom corresponds to a transition line at 1 420 MHz. In the active maser stimulated emission inside a high-Q cavity is used to detect the atomic transition. An outline of the hydrogen maser is given in Fig. 8. Atomic hydrogen is generated in a radio-frequency driven discharge fed by (molecular) hydrogen gas. The beam of atomic hydrogen emerges from an orifice and is magnetically state selected in a hexapole or quadrupole magnet. Hydrogen atoms in the upper energy state E_2 are directed into a storage bulb made of fused silica whose inner surface is coated with Teflon (Dupont trade name). The bulb is surrounded by a high-Q microwave cavity tuned to the atomic resonance frequency. While the atoms bounce around inside the bulb, they decay from level E_2 to level E_1 . Doing so, they release their energy through stimulated emission caused by the microwave field which the atoms themselves produce. The emitted radiation is added in phase to the pre-existing radiation due to the thermal background radiation in the cavity, a process that results in amplification, explaining the acronym MASER (Microwave Amplification by Stimulated Emission of Radiation). The Teflon coating ensures mostly elastic wall collisions. The average interaction time of the atoms with the microwave field is thus about one second, despite of almost 10^5 collisions with the wall of the bulb. A magnetic shield protects the interaction region from the ambient magnetic field, and a

homogeneous field of about $0.1 - 1 \mu\text{T}$ is created in this region. Pumps are used to evacuate the device and maintain a pressure smaller than 10^{-5} Pa , despite of the constant hydrogen flux.

In the **active** hydrogen maser (AHM), the radiation emitted by the atoms themselves is used for control of the quartz frequency, but the amount of energy available from the atoms is extremely small. Hence the associated electronics for sensing this L-band signal of this quantum transition are quite elaborate. Some of the radiation, of the order $10^{-13} - 10^{-14} \text{ W}$, is coupled out of the cavity and fed to a low-noise microwave receiver. There the signal is mixed with a signal at a multiple frequency of the VCXO (e.g. 1 400 MHz), and the VCXO is phase-locked to the atomic radiation.

The origin of the frequency instability of the AHM at short averaging times is the instability of the free-running VCXO, cavity thermal noise, and externally added noise. Despite of the complexities, with well-designed electronics, the excellent signal to noise available provides the best short-term stability of any frequency standard in the rf-region commercially available (see Table 12).

FIGURE 8
Hydrogen maser
 a) energy levels of the hydrogen atom; b) schematic representation




TABLE 12

Performance data of active and passive hydrogen masers

	Active	Passive
Potential deviation of the output frequency from the nominal value	3×10^{-13} (at delivery with manufacturer calibration)	10^{-12} (at delivery with manufacturer calibration)
$\sigma_y(\tau = 1 \text{ s})$	1×10^{-13}	7×10^{-13}
Flicker floor	$3 \times 10^{-15} \rightarrow 3 \times 10^{-16}$	1×10^{-15}
$\sigma_y(\tau = 1 \text{ month})$ with cavity auto-tuning without cavity auto-tuning	4×10^{-16} 3×10^{-13}	2×10^{-15} Not applicable
Temperature (per K)	3×10^{-15}	1×10^{-14}
Magnetic field (per 10^{-4} T), ($10^{-4} \text{ T} = 1 \text{ G}$)	5×10^{-15}	2×10^{-14}

The long-term frequency stability is controlled largely by the frequency stability of the cavity: The cavity-pulling of the frequency is a rather severe effect in an active device, in part because of the high Q -value of the cavity. It is extremely challenging to construct a cavity of enough mechanical stability of the individual components and their joints. The AHM has therefore been traditionally considered as useful mainly for applications where its superior short- and medium-term performance is a must, for example in VLBI ground stations. Over the years, however, several methods of “auto-tuning” the cavity in the maser have been investigated and are now incorporated in many commercial devices. Some of the masers have shown a frequency ageing of as low as 10^{-16} per day compared with primary caesium clocks. The residual drift is probably caused by the ageing of the Teflon coating of the bulb, whereby the “wall-shift” factor of several 10^{-12} (caused by the collisions of the atoms with their enclosure) is changed. Because of cavity-pulling and wall-shift, the best accuracy in the determination of the unperturbed transition frequency of the hydrogen atom is about 2×10^{-12} . Though typically very small, the long-term frequency drift is still a problem in most masers.

Because of the strong impact of the mechanically defined resonance frequency of the cavity on the maser output frequency, the devices are typically sensitive to mechanical shocks and ambient temperature changes. It is only possible to achieve the long-term stability reported above if the interaction volume is surrounded by a very effective thermal shield. An active maser typically occupies a volume of 0.5 m^3 , weighs 80 kg and may cost a factor of three to five times that of a high-performance caesium clock.


In the *passive* hydrogen maser (PHM), the cavity volume and thus the complete model can be made smaller than in an active maser. Since it is operated with a properly loaded cavity of lower Q , no self-sustained oscillation is possible. Instead, a microwave probing field is applied to the cavity, and the atomic resonance is detected using a microwave receiver. The signal processing is like that in a caesium clock. The short-term stability is about an order of magnitude worse than that of an active maser at $\tau = 1 \text{ s}$, and the stability only improves as $\sigma_y(\tau) \sim \tau^{-1/2}$. At $\tau = 1000 \text{ s}$, an active maser is typically more than 10 times more stable than a passive maser. However, because of the cavity servo, the passive maser is, in general, less sensitive to the environment. The passive hydrogen maser is the primary on-board clock technology for the European satellite navigation system Galileo. Its development started in the late nineties in Italian and Swiss firms, inheriting knowledge of the Russian company KVARZ (Nizhny Novgorod). The first device of that kind was flown on-board GIOVE-B launched in 2008. It was the first time ever that such technology was put in orbit.

In [Vanier and Tomescu, 2016] several approaches of improving the performance of hydrogen masers are presented. Among those is the development of the cryogenic hydrogen maser because these devices promise the best short-term stability of any oscillator available. It is the cavity thermal noise that determines substantially the observed frequency instability. Due to the complexity of the construction and operation, however, no commercial application has ever been achieved.

3.7 The caesium beam AFS

A Cs AFS is shown schematically in Fig. 9. A caesium atomic beam effuses from an oven which contains a few grams of ^{133}Cs . Nature is favourable as it provides only this one isotope as stable against radioactive decay. The beam passes through a state-selecting magnet, named polarizer, which deflects only atoms at the energy level E_2 in the desired direction. In a microwave cavity with two arms the atoms are irradiated twice with the pulse of a microwave field, following Norman Ramsey's idea of the separated-field excitation of the atoms [Ramsey, 1950; Ramsey, 1990].

FIGURE 9
Caesium atomic beam frequency standard



HDB-31-9

When the resonance condition (see equation (3-2)) is fulfilled, the atoms are transferred to the state E_1 . The analyser magnet deflects only these atoms to a hot wire detector which ionizes the caesium atoms on its surface. In commercial caesium clocks the ions are accelerated into a mass spectrometer and directed to the first dynode of a secondary electron multiplier (SEM). The output signal I_D is proportional to the number of ions being deflected by the analyser magnet. A resonance feature appears in the output signal I_D when the frequency v_p of the probing microwave field is swept across the value v_0 . Its line width W is a few hundred Hertz, given by $W = \langle v \rangle / ((2+\epsilon) \times L)$, where $\langle v \rangle$ is the mean velocity of the atoms in the beam and ϵ is a small numerical

factor depending on the cavity geometry and the width of the velocity distribution. L is the cavity length as defined in Fig. 9. The VCXO is steered in the same way as in a Rb standard. The value of ν_0 of unperturbed caesium atoms is 9 192 631 770 Hz per definition of the second.

Laboratory devices, referred to as primary caesium clocks or frequency standards, have been developed in some institutes and are specifically designed to facilitate the evaluation of all systematic frequency shifts [Vanier and Tomescu, 2016]. These devices have been operated with the aim of realizing the definition of the second with the highest accuracy possible. In addition to size and immobility, they differ from commercial devices by a few details: The ion current of a few picoamperes is measured directly. Due to the length of the cavity (about 1 m or even longer), the line width W is below 100 Hz. They used to determine the scale unit of International Atomic Time (TAI) for decades (see Chapter 5). With the advent of laser-cooled “fountain clocks”, however, their importance has decreased, and only two are still operational world-wide [Bauch, 2005].

The following applies in principle to all three types of microwave atomic frequency standards. In regard of the question, however, to what extent the SI second can be realized with a caesium clock it seems appropriate to elaborate here the various frequency offsets to be considered. They have been considered in detail [De Marchi, 1987; Vanier and Audoin, 1989]. In brief, the most significant of them are related to:

- i) the existence of additional energy levels very close to those defining the “clock transition”, levels E_1 and E_2 in Fig. 6. The hydrogen and the alkali atoms require an applied static magnetic field to separate several otherwise degenerated sub-levels (C-field in Fig. 9). This field shifts quadratically the resonance frequency ν_0 . Its value is a compromise between two opposite requirements. It must be as small as possible to limit the sensitivity of the transition frequency to field fluctuations, but it must be large enough to avoid as much as possible the overlapping of neighbouring transitions that occur between the different magnetic sublevels. In Cs AFS, for instance, the value of this field is chosen around 7 μ T, 1 T = 10^{-4} Gauss). The related ν_0 frequency shift is close to 2 Hz, i.e. $\approx 2 \times 10^{-10}$ in relative units. A relative variation of the magnitude of this field equal to 5×10^{-4} thus gives a relative change of the transition frequency of $\approx 1 \times 10^{-13}$. This requires the implementation of:
 - a) efficient magnetic shields in order to attenuate sufficiently the ambient magnetic field and its fluctuations, and
 - b) a stable current source to feed the coils producing the magnetic field.

The same conclusion holds for the other atomic frequency standards. The value of the static field can be measured by probing more sensitive transitions whose frequency depends linearly on its magnitude.
- ii) atomic motion. It gives rise to a possible residual first order Doppler effect and to an unavoidable, relativistic, second order Doppler effect. The related normalized frequency shifts are about 1×10^{-13} in the Cs beam frequency standard, for instance.
- iii) distortion of the atom wave function. Collisions between alkali or alkali-like atoms as well as collisions with a buffer gas or the walls of the container disturb the energy levels and cause a frequency shift. Similarly, collisions with the wall of a containing vessel perturb the atom hyperfine interaction and also give a frequency shift. These effects are negligible in the Cs AFS. The frequency offset is of the order of 10^{-11} in the hydrogen maser and of 10^{-9} in the Rb AFS.
- iv) state selection. A frequency shift occurs when the atoms are subjected to an extraneous radiation. This leads, for example, to the light shift, which is specific to the Rb AFS and could become a problem in devices employing optical pumping (see below).
- v) coupling to the interrogation microwave field. Because of this coupling, a sufficiently accurate tuning of the microwave cavity is required. Otherwise, a cavity frequency pulling arises. It is usually smaller than 1×10^{-13} . Proper microwave amplitude can greatly reduce the cavity pulling effects.
- vi) interrogation process. In all passive AFS the measured resonance frequency is sensitive to the quality of the frequency modulated signal that is necessary to probe the transition. Furthermore, then, some of the frequency offsets depend on the amplitude of the applied microwave field.

The frequency instability of the Cs AFS is at short averaging times determined by the shot-noise originating in the number fluctuations in the atomic beam and in the long term by control of the operational parameters.

The physics of Cs AFS is intrinsically less sensitive to environmental perturbations than that of Rb AFS. The output of Cs AFS can nevertheless be affected by adverse magnetic fields, by temperature and temperature gradients. If for example, a commercial Cs AFS is turned upside down, its frequency will change – probably due to temperature gradients more than due to the magnetic field of the earth or that of gravity forces.

Since 1992, a commercial Cs AFS is on the market which proved to be much less sensitive to environmental changes than heretofore observed. This improvement of nearly an order of magnitude in the long-term stability has come because of studying the physics of the transition and using its unique non-linearity coupled with digital electronics servo techniques [Cutler, 2005]. As an example, the C-field is measured periodically, and a feed-back loop stabilizes it at its design value. Very minor improvements to the specifications of this device have been introduced up to the time of writing in 2025.

Investigations into whether state selecting magnets can be dispensed with and optical pumping used instead for state selection and detection of the atoms have been in progress for some time [Audoin, 1992; Vanier and Tomescu, 2016]. The feasibility of this method has been demonstrated and used in two primary clocks [Makdissi and de Clercq, 2001; Shirley *et al.*, 2001] which have been decommissioned in the meantime, and in a few bread board versions of a compact Cs AFS. The use of state selection and detection by optical-pumping techniques would also contribute to the performance of the configuration because such techniques make use of a larger fraction of the flux of atoms than can be realized with the traditional magnetic deflection method of state selection. Recently, a commercial device, specifically designed to become constituent of a so-called Enhanced Primary Reference Time Clock as defined in Recommendation ITU-T G.8272.1 [ITU-T, 2024], has become available. Its technical specifications are very similar to those in Table 13 column “Digitally controlled”, but a longer operation time has been announced.

Commercial clocks weigh about 25 kg and fit into standard 19-inch-wide cabinets (modular versions also exist). They can be simultaneously supplied with AC and DC at a power consumption of about 50 W. Primary laboratory clocks or frequency standards are not transportable or available for purchase. Table 13 shows performance data of three models of commercially available caesium clocks and the data of a laboratory primary caesium clocks as operated in a national standards laboratory [Bauch *et al.*, 2005].

TABLE 13

Performance data of Cs AFS

Unit	Standard	High performance	Digitally controlled	Primary laboratory
Potential deviation of the output frequency from the nominal value	7×10^{-12}	7×10^{-12}	5×10^{-13}	1×10^{-14}
$(\sigma_y(\tau = 1 \text{ s}))$	1.2×10^{-11}	1.2×10^{-11}	$< 5 \times 10^{-12}$	4×10^{-12}
$(\sigma_y(\tau = 100 \text{ s}))$	2.7×10^{-12}	2.7×10^{-12}	8.5×10^{-13}	$< 10^{-13}$
Flicker floor ($\sigma_y, \min$)	n. s.	3×10^{-14}	5×10^{-15}	$< 1 \times 10^{-15}$
$\sigma_y(\tau = 1 \text{ month})$	n. s.	3×10^{-14}	5×10^{-15}	3×10^{-15}
Temperature, valid for full range 0 – 50 °C	n. s.	1×10^{-13}	8×10^{-14}	+
Magnetic field (max. 2×10^{-4} T), all orientations	n. s.	1×10^{-13}	8×10^{-14}	+

+ operation in shielded environment only; n. s. not specified, sources: product specifications, download 2020-03-02.

3.8 Miniaturized atomic frequency standards

This section provides a snapshot on current trends in the development of commercial AFS for a widest range of applications. High accuracy applications in metrology and basic science rely on the developments discussed in § 3.9.

The underlying physical principles are not particularly new, but at the time of writing of the first issue of this Handbook, developments had not left the status of purely academic research. Over the years, a wealth of studies and approaches followed. Nevertheless, the commercial market at the time of writing in 2025 is still dominated by the clock designs presented before, the Rb AFS and the Cs beam AFS. But this need not remain so forever. Two differing and sometimes mutually exclusive trends are noted:

- Reducing the volume and power consumption compared to the Rb AFS at the price of similar or even reduced performance. This trend culminated in the so-called Chip scale Atomic Clock CSAC which has proved a commercial success with thousands of units produced each year.
- Avoiding the limitations on frequency stability at short and long averaging times common to the legacy Rb AFS. Those are on one side related to the double resonance excitation with simultaneous presence of light from a discharge lamp and the microwave field. The other issue, relaxation and frequency shift due to the need of a buffer gas, is not addressed. Confinement or cooling of atoms and thus avoiding buffer gas adds substantial complexity, and examples of this are discussed in § 3.9.2.

All new developments rely on the availability of reliable narrow-linewidth semiconductor laser diodes. As they are available at wavelength 780 nm for Rb and 850 nm or 894 nm for Cs, both elements are used in each design. Frequency doubling of laser radiation from sources that are in widest use in telecommunications is relatively easy and facilitates the newer standards relying on rubidium. Another advantage is the lower value of the Rb nuclear spin and thus the lower number of sublevels in the ground-state than in Cs. A very detailed and systematic description of the various approaches, including a rich bibliography, can be found in [Vanier and Tomescu, 2016; Schmittberger, Marlow and Scherer, 2021]. A very much simplified treatment is offered in the following.

3.8.1 Getting rid of the discharge lamp

The approach closest to the legacy Rb AFS is replacement of the discharge lamp by a laser source. The continuous double-resonance excitation is kept in which the optical field polarizes the atoms by optical pumping, and the microwave field sustained in the resonator drives the hyperfine transition. The operation remains simple, with a minimum addition of new optical components that are necessary e. g. in the pulsed mode of excitation (see § 3.8.3). The use of a laser leads to a better signal-to-noise ratio in the detection of the hyperfine resonance due to more efficient pumping and reduced background light. The laser frequency, however, needs to be stabilized by a second feedback-loop based on observing the absorption of the laser radiation in the cell. In [Bandi *et al.*, 2014] short-term frequency instability as low as $1.4 \times 10^{-13} \times (\tau/s)^{-1/2}$ for $\tau < 300$ s was demonstrated, and ways to achieve good long-term stability were proposed.

3.8.2 Getting rid of the microwave resonator

Given the fact of the hyperfine transition frequencies being ≈ 6.8 GHz for Rb and ≈ 9.2 GHz for Cs, the microwave resonator to confine the RF-field cannot be made very small. In particular, in view of further miniaturization, the use of coherent population trapping (CPT) has been studied since many years. CPT is accomplished by irradiation of the atoms with two laser radiation fields separated in frequency by the hyperfine transition frequency, i.e. a second laser field coupling levels E_2 and E_3 in Fig. 6 is added. At exact resonance coherence is created and the atoms get into a non-absorbing “dark” state as was discussed as early as 1976 [Alzetta *et al.*, 1976]. In practical realizations, the current driving the laser is modulated at $\nu_0/2$, the value depending on the atom of choice. In 2005, a commercial product was available [Vanier *et al.*, 2005] and a first review article on the subject was published [Vanier, 2005]. Broadest application of CPT today is in the CSACs, in which CPT in micro-machined cells is used. After the first demonstration at the National Institute of Standards and Technology (NIST) in 2004 [Knappe *et al.*, 2004], commercial production started in 2011. A full-length review, including a description of the many (partially unexpected) applications is available in [Kitching, 2018]. The commercial product comes with the following specifications (extract only from

Microchip CSAC Datasheet_00003478.pdf, accessed 2025-10-22): Relative frequency offset from nominal at shipment: $\pm 5 \times 10^{-11}$, frequency instability at $\tau = 1$ s: 3×10^{-10} , long-term aging: $\approx 10^{-9}$ / month. Another manufacturer combines the product with a single-frequency GPS receiver, so that the 1 PPS output follows UTC(USNO) within ± 15 ns modulo an unspecified systematic offset in time. The frequency drift is suppressed through the lock to GPS, and in hold-over mode the specifications of the free-running CSAC apply (Jackson Lab GPSDO specs sheet, accessed 2025-10-22).

3.8.3 Getting rid of the issue of double-resonance excitation

Pulsed Optical Pumping (POP) of atoms in an absorption cell has been demonstrated as one way to avoid laser-light-induced frequency shifts that is unavoidable in CW double-resonance operation. Light-shifts originate from real photons that induce virtual transitions among the hyperfine energy levels of the atoms. Using POP, a laser pulse induces a population imbalance among the ground-state hyperfine levels by optical pumping. Then the laser is switched off, and the population is altered by irradiation with microwave radiation. The state population is probed afterwards with a second laser pulse and fluorescence detection. The microwave excitation can mimic Ramsey's separated-field excitation of the atoms, described before, but now in the time domain. Then the recorded linewidth of the clock transition has a width of about the invers of the separation of two microwave impulses in time – at least if relaxation can be neglected. Due to the pulsed operation, the phase evolution of the local oscillator from which the microwave probing signal is synthesized remains unobserved. In consequence certain requirements on local oscillator phase-noise need to be obeyed, otherwise the frequency instability of the POP AFS will be compromised by the so-called Dick-effect [Vanier et Tomescu, 2016; Wynands and Weyers, 2005]. Pioneering work on POP AFS was done at the Italian Istituto Nazionale di Ricerca Metrologica [Micalizio *et al.*, 2012], and recently an uptake by the industry was reported, with the aim to develop a space qualified POP frequency standard [Micalizio *et al.*, 2019].

3.9 Advanced frequency standards based on laser-cooling and optical transitions

3.9.1 Atomic fountains

When the development of thermal beam primary clocks had revealed more or less fundamental limits preventing further improvements, the first fountain clocks employing laser-cooled atoms appeared in the early 1990's [Gibble and Chu, 1993; Clairon *et al.*, 1995; Wynands and Weyers, 2005]. The U.S. Naval Observatory (USNO) developed continuously running rubidium atomic fountain clocks that have contributed to EAL (Echelle Atomique Libre) for a number of years [Peil *et al.*, 2016].

Fountain clocks are stationary, complex devices that at least now and then require human intervention for tuning and adjusting the required lasers.

3.9.2 Other advanced AFS under development

The application of laser-cooling has reached a certain level of maturity. Several research groups and commercial enterprises made efforts to develop turn-key frequency standards based on laser-cooled atoms but avoiding the bulky installation of a fountain. The idea is to produce a sample of laser-cooled atoms at rest and irradiating it in free fall for several ten milliseconds with microwave radiation, before recapturing the atoms and read out the atomic state population. Thus, the idea is similar as that used in POP described before. At the time of writing, there is little experience available in the open literature about long-term frequency stability, reliability of operation and ease of handling of the two commercial products that were presented at conferences. The first in time was MuClock, based on isotropic laser cooling of the atoms inside the microwave cavity, named HORACE (*HOR*loge à *Re*froidissement d'*A*tomies en *CE*llule = clock based on atoms cooled from vapour cell) [Esnault *et al.*, 2011]. MuClock is a current product of Muquans, France. Similar ideas but a different technique of atom preparation and detection have been realized in portable cold Rb87 clocks offered by SpectraDynamics, Inc. USA [Ascarrunz *et al.*, 2018].

A completely different idea is pursued since a couple of years at NASA's Jet Propulsion Laboratory. In the following section, the use of radiofrequency ion traps will be mentioned in the context of optical frequency standards. In order to be localized ("trapped"), ions do not require laser-cooling, unlike neutral atoms. Thus, the use of a cloud of stored mercury (Hg) ions in a trap allows microwave spectroscopy of the 40.5 GHz ground state

hyperfine transition of ^{199}Hg ions. The ions are confined in a linear radiofrequency trap and buffer-gas cooled. The state preparation and detection of the clock transition is achieved by optical pumping with a Hg discharge lamp. The medium-term goal is the demonstration of long-term operation and qualification for use as frequency standard in a spacecraft [Hoang *et al.*, 2019; Prestage and Weaver, 2007]. Work on transportable ion-trap frequency standards based on the 12.6 GHz hyperfine transition in $^{171}\text{Yb}^+$ has been reported by at least two groups [Mulholland *et al.*, 2018; Tallant *et al.*, 2019]. Here the complexity is going to be even larger as several tuneable narrow linewidth lasers are needed for the operations.

It should be noted that similar developments have also been reported out of People's Republic of China. Thus, the examples given do not represent the full set of developments world-wide, and likely the section could be updated within a five-years' time frame. For the time being it is difficult to decide which approach will prove as most successful and find its way in high accuracy applications. In general, such devices could become interesting for the generation of reference time scales in telecommunications applications, in the ground segment of GNSS and in the ground segment of space science systems in general.

3.9.3 Optical frequency standards

A new type of AFS has been intensively investigated for more than 30 years, so-called optical clocks. The status of development can be found in [Ludlow *et al.*, 2015]. They also rely on an atomic transition as a reference, however, not in the microwave range, but in the visible spectral range. The clock oscillator used in these devices is thus a frequency-stabilized laser. By increasing the clock frequency by five orders of magnitude (from approx. 10^{10} Hz in a Cs AFS to approx. 10^{15} Hz) considerable stability and accuracy gains have been demonstrated. The short-term stability is improved due to the considerably higher number of oscillations per unit of time. Enhanced accuracy is obtained because some of the disturbing external influences on the atom cause a characteristic shift of fixed magnitude which has, however, less influence at the higher frequency of the optical clock than for a clock operating in the microwave range. The field of research is very active; and a variety of atomic references have been proposed and studied. The demonstrated advances triggered considerations towards a redefinition of the second in terms of an optical transition frequency [Riehle *et al.*, 2018]. It would be necessary to identify one "best" candidate and –following good metrological practice– realize a couple of standards and verify that they agree at the expected level. See Chapter 9 for a detailed presentation of the field.

3.9.4 Applications of advanced frequency standards

At the time of writing in 2025, the use of the advanced AFS presented in §§ 3.9.1 and 3.9.3 is de facto limited to National Metrology Institutes (NMI) and advanced research centres. As mentioned before, a forerunner in timing applications is USNO which included a set of continuously running rubidium atomic fountain clocks in their clock ensemble [Peil *et al.*, 2016; BIPM, 2020]. In other NMIs, such as LNE-SYRTE (France) and PTB (Germany), caesium fountain frequency standards serve as reference for the realization of their local reference time scales, UTC(OP) for France and UTC(PTB), respectively. In the very long term, these time scales are steered to remain close to the international reference UTC (see Chapter 5). The Japanese National Institute of Communications Technologies (NICT) uses an optical frequency standard as intermediate reference for the generation of its time scale UTC (NICT) [Hachisu *et al.*, 2018], and others have shown interest to follow this example [Yao *et al.*, 2018]. Applications in fundamental science have been explained in [Ludlow *et al.*, 2015] and references therein. An up-to-date view on the field is provided in Chapter 9.

References

- ALZETTA G., GOZZINI A., MOI M., and ORRIOLS G. [1976] *An Experimental Method for the Observation of R. F. Transitions and Laser Beat Resonances in Oriented Na Vapour*. *Nuovo Cimento* Vol. B 36, 5-20.
- ASCARRUNZ F.G., DUDIN Y.O., DELGADO ARAMBURO M.C., SAVORY J. and JEFFERTS S.R. [2018] *Long-term frequency instability of a portable cold 87Rb atomic clock*. Proc. 2018 Precise Time and Time Interval Meeting, ION PTTI 2018, Reston, Virginia, USA, 107-110.

- BANDI T., AFFOLDERBACH C., STEFANUCCI C., MERLI F., SKRIVERVIK A.K., and MILETI G. [2014] *Compact High-Performance Continuous-Wave Double-Resonance Rubidium Standard With $1.4 \times 10^{-13} \times \tau^{1/2}$ stability*. *IEEE Trans. on Ultrasonics, Ferroelectrics, and Frequency Control*, Vol. 61, 1769-1778.
- BAUCH A. [2005] *The PTB primary clocks CS1 and CS2*. *Metrologia*, Vol. 42, S43-S54.
- BESSON R.J. [1977] *A New Electrodeless Resonator Design*. Proc. 31st Ann. Symp. Freq. Control, pp. 147-152.
- BIPM [2008] *Evaluation of measurement data – Guide to the expression of uncertainty in measurement*, Joint Committee for Guides in Metrology (JCGM/WG 1). JCGM 100:2008, corrected version 2010, for further update see: BICH *et al.*, [2012] *Revision of the Guide to the Expression of Uncertainty in Measurement Metrologia*, Vol. 49, 702.
- CLAIRON A., LAURENT P., SANTARELLI G., GHEZALI S., LEA S.N. and BAHOURA M. [1995] *A Caesium Fountain Frequency Standard: Preliminary Results*. *IEEE Trans. Instr. and Meas.*, Vol. 44, 128-132.
- DE MARCHI A. [1987] *New Insights into Causes and Cures of Frequency Instabilities (Drift and Long Term Noise) in Caesium Beam Frequency Standards*. Proc. 41st Annual Frequency Control Symposium, Philadelphia, USA, 53-58.
- DEHMELT H. [1981] *Coherent Spectroscopy on Single Atomic System at Rest in Free Space*. *Journal de Physique Colloques*, Vol. 42(C8) 299-306.
- ESNAULT F.X., ROSSETTO N., HOLLEVILLE D., DELPORTE J., DIMARCQ N. [2011] *HORACE: A compact cold atom clock for Galileo*, *Advances in Space Research*, Vol. 47 (5) 854-858.
- EURAMET [2016] *Guidelines on the use of GPS disciplined oscillators for frequency and time traceability*. EURAMET Technical Guide No. 3, Version 1.0.
- GERLACH W. und STERN O. [1924] *Über die Richtungsquantelung im Magnetfeld*. *Annalen der Physik*, Vol. 74, 673-699.
- GERLACH W. [1925] *Über die Richtungsquantelung im Magnetfeld II. Experimentelle Untersuchungen über das Verhalten normaler Atome unter magnetischer Kraftwirkung*. *Annalen der Physik*, Vol. 76, 163-197.
- GIBBLE K. and CHU S. [1993] *Laser-Cooled Cs Frequency Standard and a Measurement of the Frequency Shift due to Ultracold Collisions*. *Phys. Rev. Lett.*, Vol. 70, 1771-1774.
- HACHISU H., NAKAGAWA F., HANADO Y. and IDO T. [2018] *Months-long real-time generation of a time scale based on an optical clock*. (online) [WWW.Nature.com/Scientificreports](https://www.nature.com/scientificreports/) | (2018) 8:4243 | DOI:10.1038/s41598-018-22423-5.
- HOANG T.M., CHUNG S.K., LE T., PRESTAGE J.D., YI L., TJOELKER R.J., and YU N. [2019] *Performance of Micro Mercury Trapped Ion Clock*. Proc. 2019 Joint Conf. of the IEEE International Frequency Control Symposium and European Frequency and Time Forum (EFTF/IFC), Orlando, FL, USA.
- KASTLER A. [1950] *Quelques suggestions concernant la production optique et la détection optique d'une inégalité de population des niveaux de quantification spatiale des atomes. Applications à l'expérience de Stern et Gerlach et à la résonance magnétique*. *Journal de Physique et le Radium*, Vol. 11, 255-265.
- KATORI H., TAKAMOTO M., PAL'CHIKOV V.G. and OVSIANNIKOV V.D. [2003] *Ultrastable Optical Clock with Neutral Atoms in an Engineered Light Shift Trap*. *Phys. Rev. Lett.*, Vol. 91, 173005.
- KITCHING J. [2018], *Chip-Scale Atomic Devices*, *Appl. Phys. Rev.* Vol. 5, 031302.
- KNAPPE S., VISHAL S., SCHWINDT P. D. D., HOLLBERG L., KITCHING J., LIEW L., MORELAND J. [2004] *A microfabricated atomic clock*. *Applied Physics Letters*. Vol. 85(9), 1460-1462.

- LOMBARDI M. [2011] *The Evolution of Time Measurement, Part 2: Quartz Clocks*. *IEEE Instrumentation & Measurement Magazine*, October 2011, 41-48.
- LOMBARDI M. [2016] *Evaluating the frequency and time uncertainty of GPS disciplined oscillators and clocks*. *NCSLI Measure*. Vol. 11(3-4), 30-44.
- LUDLOW A.D., BOYD M.M., JUN YE, PEIK E. and SCHMIDT P.O. [2015] *Optical Atomic Clocks*. *Reviews of Modern Physics*, Vol. 87(2), 637-701.
- MAKDISSI A. and DE CLERCQ, E. [2001] *Evaluation of the Accuracy of the Optically Pumped Caesium Beam Primary Frequency Standard of BNM-LPTF*. *Metrologia*, Vol. 38(5) 409-425.
- MICALIZIO S., GODONE A., LEVI F. and CALOSSO C. [2010] *Medium-Long Term Frequency Stability of Pulsed Vapor Cell Clocks*. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, Vol. 57 (7), 1524-1534.
- MICALIZIO *et al.* [2019] *Preliminary characterization of a Rb Pulsed Optically Pumped clock for space applications*. Proc. 2019 Joint Conf. of the IEEE International Frequency Control Symposium and European Frequency and Time Forum (EFTF/IFC), Orlando, FL, USA.
- MULHOLLAND S., DONNELLAN S., BARWOOD G.P., GENTLE D., HUANG G., KLEIN H.A., PATEL P., WALSH G., BAIRD P.E.G. and GILL P. [2018] *A Portable Microwave Clock Using Laser-Cooled Trapped $^{171}\text{Yb}^+$ Ions*. Proc. 2018 of the IEEE International Frequency Control Symposium Squaw Creek Resort, Lake Tahoe, CA, USA.
- PEIL S., HANSSEN J., SWANSON T., TAYLOR J. and EKSTROM C. [2016] *The USNO rubidium fountains*. *J. Phys.: Conf. Ser.* Vol. 723, 012004.
- PIESTER D., BAUCH A., POLEWKA T., STALIUNIENE E. and TEICHEL K. [2019] *Disciplined oscillators for traceable frequency and time in metrology and financial sectors*. *NAVIGATION*, Vol. 66, 661-672.
- PRESTAGE J. D. and WEAVER G.L. [2007] *Atomic Clocks and Oscillators for Deep-Space Navigation and Radio Science*. *Proceedings of the IEEE*, Vol. 95(11) 2235-2247.
- RAMSEY N.F. [1950] *Molecular Beam Resonance Method with Separated Oscillating Fields*. *Physical Review*, Vol. 78, 695-699.
- RAMSEY N.F. [1990] *Experiments with separated oscillatory fields and hydrogen masers*. *Reviews of Modern Physics*, Vol. 62, 541-552.
- RILEY W.J. [2019] *A History of the Rubidium Frequency Standard*, IEEE UFFC-S History, <http://ieee-uffc.org/about-us/history/a-history-of-the-rubidium-frequency-standard.pdf>.
- RIEHLE F., GILL P., ARIAS F. and ROBERTSSON L. [2018] *The CIPM List of Recommended Frequency Standard Values: Guidelines and Procedures*. *Metrologia*, Vol. 55(2) 188-200.
- RUBIOLA E. [2009] *Phase Noise and Frequency Stability in Oscillators*. *The Cambridge RF and Microwave Engineering Series*, Cambridge, New York, Cambridge University Press.
- SCHMITTBERGER MARLOW B.L. and SCHERER D.R. [2021], *A review of commercial and emerging atomic frequency standards*, *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, Vol. 68(6).
- SHIRLEY J.H., LEE W.D. and DRULLINGER R.E. [2001] *Accuracy Evaluation of the Primary Frequency Standard NIST-7*. *Metrologia*, Vol. 38(5), 427-458.
- TALLANT J., NOBLE J., GUAN D., DAO N. and OVERSTREET K.R. [2019] *Progress towards a commercial $^{171}\text{Yb}^+$ Microwave Atomic Frequency Standard*. Proc. 2019 Joint Conf. of the IEEE International Frequency Control Symposium and European Frequency and Time Forum (EFTF/IFC), Orlando, FL, USA.
- VANIER J. and AUDOIN C. [1989] *The Quantum Physics of Atomic Frequency Standards*. Bristol and Philadelphia, Adam Hilger.

- VANIER J. [2005] *Atomic clocks based on coherent population trapping: a review*. *Applied Physics B*, Vol. 81, 421-443.
- VANIER J., LEVINE M.W., KENDIG S., JANSSEN S., EVERSON C., and DELANEY M.J. [2005] *Practical Realization of a Passive Coherent Population Trapping Frequency Standard*. *IEEE Trans. on Instrumentation and Measurement*, Vol. 54(6), 2531-2539.
- VANIER J. and TOMESCU C. [2016] *The Quantum Physics of Atomic Frequency Standards – Recent developments*. CRC Press, Taylor & Francis Group, Boca Raton, London, New York.
- WYNANDS R. and WEYERS S. [2005] *Atomic Fountain Clocks*. *Metrologia*, Vol. 42(3), S64-S79.
- YAO J., PARKER T.E., ASHBY N. and LEVINE J., [2018], *Incorporating an Optical Clock into a Time Scale*. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, Vol. 65 (1), 127-134.

ITU texts

Recommendation ITU-T G.8272.1 (01/2024) – *Timing characteristics of enhanced primary reference time clocks*

CHAPTER 4

Measurement techniques (Metrology)

TABLE OF CONTENTS

	<i>Page</i>
Introduction	64
4.1 Frequency measurement techniques	64
4.1.1 Direct measurement of time (phase)	65
4.1.2 Direct measurement of frequency	67
4.1.3 Heterodyne measurements	67
4.1.4 Measurements of phase noise.....	70
4.2 Measurement techniques for transmission and reception links	81
Introduction	81
4.2.1 Software-Defined Radio measurements in time metrology	81
ITU References.....	88
References for Section 4.2.1	89
4.2.2 Optical fibre links: measurement methods.....	90
References for Section 4.2.2.....	106

Introduction

There are many measurement techniques for characterizing the performance of frequency sources. These techniques differ widely in the type of information they provide including the uncertainty and resolution of the data. The techniques can be divided into two broad categories:

- 1) methods for characterizing the frequency stability of an oscillator that compare the output of the device under test to a reference source using direct measurements, heterodyne measurements or phase measurements in the time domain;
- 2) methods that are based on evaluating the performance of the device by using methods based on a description of the stability in the Fourier domain by using the Power Spectral Density and similar metrics.

The two methods provide different perspectives on the same behaviour, and the choice between them often depends on the application that will use the output of the device, and whether the application is more sensitive to fluctuations in the amplitude of the output signal or to its frequency and phase, which are often determined more by the time of a zero-crossing and are much less sensitive to the amplitude of the output.

4.1 Frequency measurement techniques

In general, the output of a frequency source is given by:

$$V(t) = (V_o + \varepsilon) \sin(2\pi\nu_o t + \varphi) \quad (4-1)$$

The amplitude is the sum of two contributions: an average value, ν_o , and a noise contribution, ε , which is a random variable that is specified statistically. The explicit time dependence of the output is specified by the average frequency, ν_o , and a phase offset, φ , which might be a constant, but is usually a stochastic variable that is specified statistically. The configurations where ε or φ are purely deterministic parameters that specify amplitude or phase modulations respectively are not considered here.

In this first section is considered the case where the statistical behaviour of the output of a frequency source is dominated by the stochastic contribution of the phase term, φ , and the contribution of the amplitude noise, ε , can be ignored or treated as a small perturbation on the output because it is much smaller than ν_o . This simplification is particularly useful when the frequency source is used as a source of time or time interval because these applications generally use a “tick” that is defined as the instant of time when the output voltage goes through zero with a positive slope. The times defined by these “ticks” are sensitive to the phase noise, but are much less sensitive to ε . The tick is often realized as a pulse, and the leading edge of the pulse specifies the “on time” marker. The width of the pulse and the trailing edge are not used.

In the following discussion, *time*, is measured in seconds and fractions, and indicates the *difference* between the corresponding ticks of two clocks and not the more commonly used meaning of the time of either one of them. If the outputs of the two devices are sine-wave voltages rather than ticks, the time difference can also be specified as the fraction of the period of the device considered to provide the reference signal for the measurement. If the nominal frequency of the reference signal is ν_r , the period is $T_r = 1/\nu_r$. If the time difference between the two signals is Δt , the phase difference, measured in radians is given by $\Delta\varphi = 2\pi(\Delta t/T_r)$. The phase difference is expressed as the value mod 2π even when the time difference exceeds the period of the reference signal.

The frequency of the device under test is the evolution of the time difference of this device over a specified time interval with respect to the device that provides the reference signals for the measurements. This definition is the ratio of two times and is therefore dimensionless.

The relationship between the frequency defined in the previous paragraph and the more common oscillation frequency of the physical device can be illustrated by considering the evolution of the time difference between two noiseless oscillators, a reference device whose parameters are identified with the subscript r , and a device under test identified with subscript d .

$$V_r(t) = V_{or} \sin(2\pi\nu_r t) \quad (4-2)$$

$$V_d(t) = V_{od} \sin(2\pi\nu_d t) \quad (4-3)$$

The time intervals between consecutive ticks of the reference device and device under test are $\Delta t_r = \frac{1}{\nu_r}$ and $\Delta t_d = \frac{1}{\nu_d}$ respectively, and the time of each clock is incremented by these values on each tick. On each tick, the time difference between the device under test and the reference device changes by:

$$\frac{1}{\nu_r} - \frac{1}{\nu_d} = \frac{\nu_d - \nu_r}{\nu_r \nu_d} \quad (4-4)$$

The elapsed time for this change is specified by the reference device as $1/\nu_r$, the period of the reference device, so that the evolution of the time difference divided by the elapsed time, which is one cycle of the reference, is:

$$\frac{\nu_d - \nu_r}{\nu_d} \quad (4-5)$$

which is the ratio of two frequencies and is dimensionless. The two frequencies are generally almost the same, so that equation (4-5) can be interpreted as $\Delta\nu/\nu$, the fractional frequency offset of the device under test with respect to the reference frequency. The evolution of the time difference between the two devices is a function of this fractional frequency difference, and not on the magnitudes of the physical frequencies of the devices. This fraction is unchanged if both frequencies are scaled by any constant value.

The frequency difference in equation (4-5) is a signed quantity. A positive value implies that the tick of the device under test is occurring earlier than the tick of the reference device. In other words, if the tick of the device under test is already earlier than the reference device, then the difference in time will continue to increase. If the tick of the device under test is later than the tick of the reference device, the difference will decrease and the two ticks will move closer to coincidence.

4.1.1 Direct measurement of time (phase)


The direct measurement of the time difference between two signals uses a time-interval counter, which reports the time interval, measured by counting the ticks of a reference frequency, between the tick outputs of the two devices as shown in Fig. 10. If the outputs of the two devices are sine-waves, then the ticks are generated by each device when its output voltage goes through zero with a positive slope. If the outputs of the two devices are pulses, then the tick is the leading edge of the pulse; the width of the pulse and the trailing edge are not used.

With the configuration shown in Fig. 10, the counter reports $t_1 - t_2$ which is the time interval from a tick from device 1 to the *next* tick in device 2 (green arrow) and is always positive definite in most hardware configurations. The value of $t_2 - t_1$ might be more appropriate in situations when the ticks are assigned specific time values so that the stop pulse occurs before the start pulse with the same time value. It can be measured by interchanging the start and stop connections or by subtracting the value from the period of the signals, which are usually nearly equal.

The start pulse from device 1 and the stop pulse from device 2 do not occur at the same time. If the frequencies of the two oscillators are different, then the time difference between two signals varies systematically, and a correction factor must be applied to the reading of the counter to report the time difference between the two signals referenced to the same instant. If device 1 is considered to be the reference device whose ticks define the time, then the reading of the counter must be adjusted by the factor $(\nu_2 - \nu_1)/\nu_1$. See equation (4-5).

FIGURE 10

The time interval $t_1 - t_2$ measured by a simple time interval counter and shown by the green arrow



HDB-31-10

The counter may report an ambiguous value if the two ticks are almost simultaneous, and the ticks from device 2 may have to be delayed if this is a problem. This will happen if the time difference is either almost zero or very close to a full period. This situation is particularly serious when the time difference between the two devices is comparable to the rise times of the leading edges of the ticks or if the leading edges of the pulses have complex structure due to reflections in the cables connecting the counter to the frequency sources.

The value reported by the counter is the number of ticks of the reference signal that are counted between the start and stop signals, and the accuracy of the reference signal affects this value. A simple counter will have an ambiguity of one count of the reference signal, but more sophisticated counters provide some interpolation and can increase the resolution of the measurement by one or two orders of magnitude. The resolution also can be increased by measuring the time difference over longer time intervals. Averaging the time difference in this way is not appropriate if there is a deterministic frequency offset between the two devices, and it may be difficult to interpret and not statistically robust if the statistical properties of the time difference are not white phase noise.

If the frequencies of the two devices are sufficiently different, then the measurement of the time difference may eventually exceed a period of the reference signal, and the reported values will be periodic. The ambiguity associated with this periodicity can be removed by dividing both input frequencies by some integer, which increases the time before roll-over is reached. A common solution is to divide the input frequencies so that the signals are converted to 1 Hz. The interval between roll-over events can be increased by adding a separate counter that keeps track of the number of roll-over events and adds this count to the reported time difference. The interval between roll-over events will be determined by the dynamic range of this auxiliary counter.

4.1.1.1 Estimating the frequency from the time data

The configuration in the previous section can also provide an estimate of the frequency of device 2 with respect to device 1. When the time-difference measurements described in the previous section are repeated at known equally spaced intervals, the frequency difference between the two devices can be estimated by using the evolution of these time differences divided by the interval between them. The measurements at the endpoints of the averaging interval may be affected by noise in the measurement of the time of the end-point ticks. If this noise contribution has a variance σ_x^2 , this variance will make a contribution to the variance of the frequency of $\sigma_y^2 = 2\sigma_x^2/\tau^2$, where τ is the time interval from the start pulse to the stop pulse. The variance of this frequency estimate can be reduced by increasing the measurement interval, but the contribution of the variance of the frequency will also contribute, and the optimum measurement interval must consider both contributions.

The measurement technique described in the previous paragraph gives equal weight to all the contributions to the evolution of the time difference during the measurement interval, which may include multiple periods of the reference. This type of measurement is often called a “ Π counter” as a graphical analogy between the Greek letter and the weighting function. The Π counter reports the average of the time evolution over the measurement interval, and this is appropriate if the time differences are characterized mostly by white phase noise over that averaging time. It will be less useful if the time differences include a significant frequency-noise contribution, and especially if there are significant fluctuations with periods that are multiples or sub-multiples of the measurement interval. The simple average of the time evolution is not optimum in these situations, and more complex averages, such as the “ Λ counter” which uses a triangular weighting function, or a “ Ω counter” which uses a parabolic weighting function may be more appropriate. These different counters can be implemented in the front-end measurement hardware or in post-processing of the data. The choice of post-processing adds complexity and usually requires a short measurement interval to facilitate the subsequent averaging, but it is more versatile, since the same measurements can be post-processed with all three methods. Ideally, the measurement interval should be short enough so that there are at least two measurements per cycle of the highest Fourier frequency that is of interest. It is often easier to satisfy this requirement with the heterodyne methods described in § 4.1.4.

4.1.2 Direct measurement of frequency

A time-interval counter can also be configured to measure the frequency of an input signal. In this configuration, the start and stop pulses are generated by an internal oscillator, and the counter measures the number of ticks of the input signal between a start pulse and the next stop pulse. The display shows frequency directly in Hz when the interval between the start and stop pulses is 1 s, but most devices can be configured to use longer measurement intervals and scale the display accordingly so as to show the frequency in Hz.

The resolution of this method depends on the number of counts of the input signal during the measurement time and implicitly assumes that the measurement interval is much longer than the period of the input signal. The measurement technique becomes less useful for measuring very low frequencies because the number of counts is quite small in this case. This weakness can be addressed by using longer measurement intervals or by reversing the functions of the two signals so that the ticks of the reference oscillator are counted and the input signal provides the start and stop pulses. This method measures the period of the input signal instead of its frequency. In many devices, the display is scaled by actual period of the reference oscillator to show the period of the input signal in seconds.

This technique measures the average of the frequency of the input signal over the measurement interval, and this averaging will produce ambiguous results if there are significant fluctuations in the input signal with periods comparable to the measurement interval. This ambiguity will be particularly serious if there are significant fluctuations in the input frequency with periods that have an exact or nearly exact relationship to the measurement interval.

4.1.3 Heterodyne measurements

In the heterodyne technique, the signal under test with nominal frequency ν is mixed with a reference signal with nominal frequency ν_o , and the beat frequency, $\nu_b = \nu - \nu_o$ is measured. The reference frequency is chosen so that the beat frequency is much lower than the frequency of the device under test, and the frequency resolution is improved by the factor ν_o/ν_b . The fluctuations in the frequency (or the phase) of the signal under test are preserved in the heterodyne process and therefore become a much larger fraction of the lower heterodyne frequency. This improves the measurement of time differences.

The beat signal can be obtained by connecting the two signals to a non-linear element such as a double-balanced mixer or by digitizing the input signal and computing the beat frequency digitally by using a digitally constructed reference signal. In either case, the output of the mixing process is the sum and the difference of the signal and reference frequencies, and the difference frequency is extracted by a lowpass filter. The lowpass filter can be constructed by using either analogue or digital methods; the choice between these two implementations depends on how the mixing process was realized. In either case, the design of the lowpass filter is not critical because the sum and difference frequencies are very far apart.

The frequency of the beat signal depends on the absolute value of the difference between the frequency of the input and the frequency of the reference, and additional information is needed to decide whether the input frequency is above or below the frequency of the reference. This ambiguity can be resolved by making a small change to the reference frequency and noting how this change affects the beat frequency. A step in the phase of the reference frequency can provide the same information. The ambiguity can also be resolved by comparing the phase of the beat frequency with the phase of a second beat frequency obtained by mixing the input with the reference signal phase shifted by one-quarter of a period.

The beat frequency is a sine wave in principle if the input signals were sine waves, but the process of constructing the beat frequency usually introduces significant harmonic distortion, which complicates the analysis of the time variation of the beat signal. However, the phase relationships among the harmonics are usually quite stable so that it is straightforward to generate “ticks” from the zero-crossings of the beat signal, which are analogous to the ticks described in previous sections and can be analysed by using the methods described in the previous sections. If the reference signal is not a sine wave, then it will contain harmonics of the fundamental frequency, and it is possible to use the beat frequency between the input signal and one of the harmonics to generate the ticks. The beat frequency with the desired harmonic is extracted by the lowpass filter. This technique is often used when it is difficult to generate the appropriate reference frequency.

An important application of the heterodyne method is to measure a frequency that is too high to be measured by conventional hardware. The heterodyne method translates the signal frequency to a lower value where it can be easily measured, and where the fluctuations in the input frequency can be more easily observed.

This method is routinely used to measure optical frequencies, which are too high to be measured by conventional hardware. An “optical comb” is routinely used for this purpose. The comb consists of a series of optical frequencies that differ in frequency by a much lower frequency, which can be measured by conventional hardware. The beat frequency between the frequency of the optical signal under test and the nearest “frequency tooth” of the comb is measured by conventional hardware, and this beat frequency preserves the statistical characteristics of the optical input.

The beat frequency must be high enough so that the anticipated noise in the signals do not drive the beat frequency through zero. A common configuration is to choose a reference frequency so that the beat frequency is twice as high as the highest Fourier component in the input signal that is of interest.

If the input signal consists of a range of frequencies, there can be two beat signals –one generated by the component of the input below the reference frequency and the second by the component above the reference frequency. The reference frequency is chosen so that only one of these possible frequencies is present in the input. This condition will be satisfied if the beat frequency is greater than one-half of the range of frequencies in the input. If the range of frequencies in the input signal is large enough, this condition will require a beat frequency that may be too high to measure. A second heterodyne process is used in this case to reduce the beat frequency to a lower value that can be measured.

The phase of the beat frequency depends on the phase of the reference signal and the phase shift introduced by the lowpass filter, and these contributions are usually unknown. The heterodyne method is sensitive to changes in the relative phase of the input and reference signals but not to the absolute value of either one of them. It is useful for measuring fluctuations in the time of the input signal, but not the time itself.

4.1.3.1 Homodyne measurements

Homodyne measurements measure the characteristics of the beat frequency between two inputs that have the same frequency. The beat signal is the sum of a constant value based on the static phase difference between the two signals and a fluctuating contribution driven by the phase noises of the two devices.

One method of realizing the homodyne configuration is to lock the frequency of the oscillator under test to the reference frequency with a loop filter that has a known transfer function. (The reference frequency can be a lowpassed copy of the signal from the device under test.) The phase noise of the two signals can be computed by combining the output of the homodyne measurements with the transfer function of the loop filter. The loop filter typically has a lowpass transfer function, which reduces the sensitivity of the beat frequency to fluctuations that have a Fourier frequency less than the unity-gain frequency of the loop. These fluctuations can be observed by measuring the feedback signal in the loop rather than the beat frequency output.


4.1.3.2 The dual-mixer configuration

The dual-mixer configuration is used to measure the time difference between two oscillators that have nearly the same nominal frequency, ν . Although a heterodyne technique could be used in principle, the beat frequency would be too small to measure in any reasonable time. For example, the fractional frequency difference of atomic clocks is often less than 10^{-11} , so that the period of the beat frequency would be many years.

The frequency of a third, reference oscillator, ν_r , is adjusted so that the beat frequency between the reference oscillator and the two input signals is a small value –typically of order 10 Hz. This offset frequency is often realized by a frequency synthesizer locked to one of the input oscillators with a frequency offset, but it can also be generated by a purely digital technique that uses a third independent oscillator. The frequency constructed in this way would be compared to one of the input signals, and the digital synthesizer would be programmed to maintain the nominal value for the beat frequency.

The beat frequencies between the synthesized reference frequency and each of the inputs is computed by a dual-mixer as in the heterodyne method. The outputs of the two mixers are passed through lowpass filters to remove the sum frequency component of the mixing process, and the zero-crossings of the two difference frequencies are used to generate ticks. Finally, the time difference between these two ticks is measured by a conventional time interval counter. The phase of the reference frequency is common to both channels and its fluctuations are attenuated in the difference. See Fig. 11.

FIGURE 11
Dual mixer system



HDB-31-11

The time of a measurement is determined by counting the cycles of the oscillator that is used as the reference for the measurement system. The oscillator that drives the offset frequency generator (oscillator 1 in the Figure), is usually used for this purpose. The offset frequency may be used to measure multiple oscillators simultaneously. Each measurement channel would use the same offset frequency, but would have a separate mixer, lowpass filter, and time interval counter.

The offset frequency is offset below the frequency of its reference, so that the phase of the reference signal is retarded in time with respect to both input signals. The phase of the offset frequency comes into quadrature with the reference frequency at some instant, and the mixer/lowpass filter produces a start tick at that instant. The phase of the offset frequency comes into quadrature with the frequency of oscillator 2 at some later time, and that generates the stop pulse for the time interval counter. The time interval counter measures the time difference modulo 1 cycle of the offset frequency and the integer number of cycles that have elapsed since the last measurement are recorded by the separate integer cycle counters.

To illustrate the method, consider two oscillators with frequencies ν_1 and ν_2 and phases ϕ_1 and ϕ_2 . The output signals of the two oscillators are described by equations similar to equation (4-1):

$$V_1(t) = (V_{01}) \sin(2\pi\nu_1 t + \phi_1) \quad (4-6)$$

$$V_2(t) = (V_{02}) \sin(2\pi\nu_2 t + \phi_2) \quad (4-7)$$

The reference frequency is described in the same way. It is usually generated from one of the input frequencies by using a frequency synthesizer that introduces a small frequency offset, but this is not fundamental.

$$V_r(t) = (V_{0r}) \sin(2\pi\nu_r t + \phi_r) \quad (4-8)$$

The two mixers multiply the signals in equations (4-6) and (4-7) with the signal in equation (4-8) and extract the signals at the two difference frequencies, V_{1d} and V_{2d} :

$$V_{1d} = \frac{V_{01}V_{0r}}{2} \cos(2\pi(\nu_1 - \nu_r)t + (\phi_1 - \phi_r)) \quad (4-9)$$

$$V_{2d} = \frac{V_{02}V_{0r}}{2} \cos(2\pi(\nu_2 - \nu_r)t + (\phi_2 - \phi_r)) \quad (4-10)$$

The two difference frequencies are $V_{b1} = \nu_1 - \nu_r$ and $V_{b2} = \nu_2 - \nu_r$. The phase differences are $\phi_{b1} = \phi_1 - \phi_r$ and $\phi_{b2} = \phi_2 - \phi_r$. The two signals generate ticks when the corresponding argument of the cosine function is $(2n + 1.5)\pi$ for any integers, n and m :

$$t_{n1} = (n + 0.75) \frac{1}{\nu_{b1}} - \frac{\phi_{b1}}{2\pi} \frac{1}{\nu_{b1}} \quad (4-11)$$

$$t_{m2} = (m + 0.75) \frac{1}{\nu_{b2}} - \frac{\phi_{b2}}{2\pi} \frac{1}{\nu_{b2}} \quad (4-12)$$

The time difference is

$$t_{n1} - t_{m2} = \left(n \frac{1}{\nu_{b1}} - m \frac{1}{\nu_{b2}} \right) + 0.75 \left(\frac{1}{\nu_{b1}} - \frac{1}{\nu_{b2}} \right) - \frac{1}{2\pi} \left(\frac{\phi_{b1}}{\nu_{b1}} - \frac{\phi_{b2}}{\nu_{b2}} \right) \quad (4-13)$$

The first term on the right side of equation (4-13) is the explicit time dependence of the difference. The second term is the definition of a tick as the time when the output crosses 0 with a positive slope, and the third term is the time difference between the two signals when $t = 0$. If oscillator 1 provides the reference time for the measurement, in one period of the signal from oscillator 1, the time difference evolves by a factor proportional to

$$\frac{\nu_{b2} - \nu_{b1}}{\nu_{b2}} = \frac{\nu_2 - \nu_1}{\nu_{b2}} \quad (4-14)$$

since the difference in the beat frequencies is the same as the difference in the signal frequencies. Comparing equation (4-14) with equation (4-5) shows that the initial time difference between the two oscillators has been magnified by $\nu_2/\nu_{b2} \cong \nu/\nu_b$ since the frequencies of the two signals are almost the same. For example, if the signal frequencies are 5 MHz and the beat frequency is 10 Hz, the magnification factor is 5×10^5 . The two integer cycle counters record the values of n and m , which will be equal initially, but will slowly diverge. The start and stop pulses that control the time-interval counter occur at different times, and the time difference must be adjusted as in the previous discussion, and this correction must include the difference in the two integer cycle counters. See equation (4-5).

The value of the beat frequency is a compromise. Reducing the beat frequency increases the magnification factor and therefore the resolution of the measurement but also increases the $1/f$ noise of the mixer. A beat frequency of approximately 10 Hz is often chosen as a good balance between these considerations.

4.1.4 Measurements of phase noise


In the following sections, is considered the more general case where the fluctuations in the amplitude of the output signal cannot be ignored. The output signal will be affected by both amplitude noise (AM) and phase noise (PM). The output signal in this case is a more general form of 0.1.

$$v(t) = V_0[1 + \alpha(t)] \cos[2\pi\nu_0 t + \phi(t)] \quad (4-15)$$

where V_0 is the peak amplitude, ν_0 is the carrier frequency, $\alpha(t)$ is the random fractional amplitude, and $\varphi(t)$ is the random phase. The physical dimension of $\alpha(t)$ is one, and $\varphi(t)$ is an angle, measured in radian². Figure 12 shows the clock signal as observed with an ideal oscilloscope. In the presence pure AM noise (lower-left side), the peak amplitude fluctuates by a random amount equal to $V_0\alpha(t)$, but the phase is not perturbed. Oppositely, a signal affected by pure PM noise (lower-right side) has fluctuating phase, but the peak amplitude is exactly equal to V_0 , with no fluctuations. Real signals are affected by both AM and PM noise, generally not in the same amount. Modulation theory teaches us that the bandwidth of $\varphi(t)$ and $\alpha(t)$ is strictly limited to ν_0 . However, in most cases the bandwidth of the system is no more than 1-10 % of ν_0 . The practical implication is that cycle-to-cycle noise cannot be observed, and it takes several cycles of the carrier for the fluctuations to show up. This is why is needed a dual time base oscilloscope in Fig. 12, to enlarge a small region of the carrier delayed by τ^* after the trigger.

FIGURE 12

A comparison of the effects of amplitude and phase noise



Note to Fig. 12: The blue ovals indicate where the largest fluctuations occur, while the pink arrows point to the stable parameters.

Phase noise is best described in terms of one-sided power spectral density (PSD) of the random phase $\varphi(t)$, denoted with $S_\varphi(f)$, as a function of the modulation frequency (or Fourier frequency) f . The quantity $S_\varphi(f)$ has the physical dimension time but, owing to the presence of the square angle, the appropriate unit is rad^2s , or equivalently rad^2/Hz . It is often expressed in decibels, i.e. $\text{dB rad}^2/\text{Hz}$, using $10 \log_{10}[S_\varphi(f)]$. The PSD of $\alpha(t)$, denoted with $S_\alpha(f)$, is used to describe the fractional-amplitude noise. It has the physical dimension of time, and the measurement unit is s, or equivalently $1/\text{Hz}$.

The PSD is estimated using the Welch algorithm [Welch P.D., 1967]. The one-sided PSD ($f > 0$) is:

² Strictly speaking, the SI unit of angle has the dimension of one, but it has the name of radian, symbol rad.

$$\widehat{S}_\varphi(f) = \frac{2}{T} \langle |\Phi_T(f)|^2 \rangle_m$$

where the hat accent means “estimated” (most readers may ignore it), $\Phi_T(f)$ is the discrete Fourier transform of $\varphi(t)$ truncated over the time T , and the angular parentheses $\langle \rangle_m$ denote the average on m realizations. The factor 2 is necessary for energy conservation, after deleting the negative frequencies.

For historical reasons, manufacturers prefer expressing phase noise in terms of $\mathcal{L}(f)$, defined as:

$$\mathcal{L}(f) = \frac{1}{2} S_\varphi(f) \quad (\text{definition, IEEE Std 1139}) \quad (4-16)$$

generally given in dBc/Hz, i.e. $10 \log_{10} \mathcal{L}(f)$. Of course, $\mathcal{L}(f)$ and $S_\varphi(f)$ express the same physical quantity but for the measurement unit. It follows from equation (4-16) that $\mathcal{L}(f)$ makes use of a unit of angle equal to $\sqrt{2} \text{ rad} = 81^\circ$ instead of 1 rad, but no angle associated to $\mathcal{L}(f)$ is mentioned in the literature ever. Because $\mathcal{L}(f)$ is not an SI quantity, at some point it must be abandoned in favor of $S_\varphi(f)$.

In our experience there is a lot of confusion with $\mathcal{L}(f)$, in part related to the obsolete and wrong definition which most people still have in mind:

$$\mathcal{L}(f) = \frac{\text{Power in 1 Hz bandwidth at a frequency } f \text{ off the carrier frequency } f_0}{\text{Carrier power}} \quad (\text{wrong definition}) \quad (4-17)$$

The definition in equation (4-17) suffers from several flaws:

- It is a conceptually incorrect representation of phase noise because the sideband power originates from any combination of AM and PM noise, not from PM noise alone.
- It is experimentally incorrect because all commercial phase noise analysers use true phase detectors. None of them measures the “SSB power”.
- It fails to describe phase noise at low f , where slow phenomena make the phase swing exceed 1 rad² in 1 Hz bandwidth.

A model which is found useful to describe the phase noise of oscillators and components is the *polynomial law*, or *power law* [Barnes JA, (1971), Halford D., (1968)], shown in Fig. 13.

$$S_\varphi(f) = \sum_{n=-4}^0 b_n f^n \quad (4-18)$$

where the coefficients b_n are the parameters which describe the corresponding noise process. The powers of f are often called “slope” as f^n is a straight line in log-log scale. All such processes are found in oscillators, and steeper terms, like $1/f^5$ and $1/f^6$ are possible. In contrast, two-port devices, like amplifiers, are limited to white and flicker phase noise processes. The difference is that steeper phenomena, like white frequency noise (b_{-2}/f^2 , b_{-3}/f^3 , etc.) represent a diverging “time error”, which is seen in oscillators, not in the input-to-output delay of devices. [Rubiola E., 2010] and Boudot R. *et al.*, 2012] provide details on the noise in oscillators and amplifiers.

FIGURE 13

Polynomial law of phase noise





Figure 14 shows an example of phase noise PSD (10 GHz DRO). From the right-hand to the left, white phase noise, white frequency noise, and a large flicker-frequency noise region are found. Flicker of phase is not visible (this is typical of microwave oscillators), and the transition between flicker and random walk frequency noise is visible at a few Hz frequency.

FIGURE 14

Example of phase noise spectrum, with polynomial law and artefacts marked on the plot



HDB-31-14

Note to Fig. 14: The oscillator is a Synergy Microwave GSDRO100, 10 GHz output frequency with a dielectric resonator inside, measured with a Rohde Schwarz FSWP-26 phase noise analyser. The small spurs at 30 Hz, 600 Hz, and 450 kHz are likely due to residual interferences in the power lines. The increased slope below 20 Hz reveals the transition from white FM noise to FM random walk at lower frequencies, which is a typical behaviour of oscillators. The jump at 1 MHz is clearly an artefact of the analyser. Its origin is ascribed to a change in frequency resolution and number of average spectra. The origin of the bump at 1-100 kHz is not understood. Measurement done by [Y. Gruson and E. Rubiola, 2025].

4.1.4.1 Basic measurement methods and techniques


A) Absolute (oscillator) and differential (added noise) measurements

Figure 15 (A) shows the absolute measurement of the phase of a device under test (DUT) with respect to a reference device. The phase noise detector is still unspecified, the signal $v(t)$ is either analogue or digital, and k is the calibration factor that relates output voltage to the phase difference. It is assumed that the noise of the reference is negligible compared to the DUT noise. Alternatively, two equal oscillators can be compared, and it is assumed that each contributes half the phase noise measured. Figure 15 (B) shows the measurement of a two-port device as the difference between output and input phase. Manufacturers often refers to this measurement as “added noise”. The oscillator is common mode, and its noise cancels. However, it is necessary that the DUT delay is small relative to the period of the fluctuations of the reference oscillator. When this is not the case, the oscillator noise will affect the output because the delay de-correlates the two signal branches from the oscillator.


FIGURE 15

Basic methods for the measurement of phase noise

(A) Absolute measurement (oscillators)



(B) Differential measurement (two-port devices)



HDB-31-15


There is an alternate use of the differential scheme of Fig. 15 (B), where a reference delay τ is inserted in lieu of the DUT, with the purpose of measuring the auto-correlation of the oscillator's phase noise. The delay unit can be a delay line (e.g. an optical fibre, a cable, a SAW device) or a resonator. The equations are given in [Volyanskiy *et al.*, Section 2C, 2008]. Interchanging the role of resonator and delay, frequency stability of a resonator (the delay unit) can be measured using a more stable oscillator as the reference. When a sufficiently stable oscillator is not available, two equal resonators can be inserted in the “DUT” and “REF” place of Figure 15 (B) to reject the oscillator noise [Rubiola *et al.*, 2000]. When the two-resonator configuration is used, the output is a measurement of the total noise, and the contribution of each cannot be separated.

B) The saturated-mixer phase detector

The principle is shown in Fig. 16. A double balanced mixer (DBM) saturated at both inputs by signals in quadrature converts the phase $\phi(t)$ into a voltage $v(t) = k_\phi \phi(t)$. The mixer is followed by a low-pass filter, a low-noise amplifier (LNA), an analogue-to-digital converter (ADC), and digital hardware and software, which calculates the average fast Fourier transform (FFT). The filter removes the $2\nu_0$ signal at the mixer output and creates a path for the current at $2\nu_0$ in the mixer internal diodes. The conversion coefficient k_ϕ is typically 100-500 mV/rad, depending on the mixer technology and on the signal power. The frequency range is limited by the mixer. Mixers are available in the whole RF/microwave spectrum, from < 1 MHz to 50 GHz and beyond. The reader interested in understanding mixers can refer to the tutorial [Rubiola, E., 2006].

FIGURE 16

Principle of the double-balanced-mixer technique



HDB-31-16

The mixer requires that the two inputs are at the same frequency and in quadrature. This comes at the cost of having a reference oscillator at the same frequency of the DUT or introducing a frequency synthesizer. In absolute measurements, the quadrature condition is achieved by phase-locking the reference to the DUT. The appropriate control bandwidth is long enough to enable the measurement of the error signal v , and short enough to keep the quadrature condition within 0.1-0.2 rad. Conversely, the quadrature condition in differential measurements needs a variable phase shifter or a line stretcher.


Experience suggests that the background white noise is limited by the LNA, and that the mixer's noise is lower. At room temperature, the noise of best amplifiers is of 1.2-1.5 nV/ $\sqrt{\text{Hz}}$, including the thermal 0.9 nV/ $\sqrt{\text{Hz}}$ noise of the 50 Ω input load. For example, with 200 mV/rad conversion coefficient and 1.25 nV/ $\sqrt{\text{Hz}}$ amplifier noise, the background is 6.25 nV/ $\sqrt{\text{Hz}}$, thus $S_\phi = 3.9 \times 10^{-17} \text{ rad}^2/\text{Hz}$ (this is the b_0 coefficient of the polynomial law in equation (4-18)), i.e. -164 dBrad²/Hz. The flicker phase noise background (the b_{-1} coefficient) is generally limited by the mixer, depending on technology and frequency range. For reference, the flicker can be -120 to -130 dBrad² for a microwave mixer in favourable conditions, and -130 to -140 dBrad² for a RF mixer, degrading if the power is lower than optimum.

C) Direct digitization of the RF signal

The principle is shown in Fig. 17. The input and reference signals are digitized and down converted to DC by multiplication with reference signals from a numerically controlled oscillator (NCO). Phase and amplitude are extracted from the I-Q data stream with the CORDIC algorithm (See Reference [Meher *et al.*, 2009] for a review). Acquisition, NCO, I-Q detection, and possibly the CORDIC algorithm, require the processing speed of a FPGA.

A question arises: why not simplify the scheme by feeding one ADC with the reference signal, and removing the other ADC? The reasons are: (i) low-noise, high-speed ADCs need a clock frequency in a specific range, which may not match the available reference frequency, (ii) it is necessary that the instrument adapt itself to the input frequency with appropriate choice of the sampling frequency, in order to avoid the dark zone around the Nyquist, and (iii) the use of odd sampling frequencies, instead of round values like 100 or 125 MHz, helps managing the spurs.

FIGURE 17

Principle of the direct digitization technique

HDB-31-17

The digital phase detector has two main advantages over the saturated DBM.


- The I-Q detection and the CORDIC algorithm enable to extract $\varphi(t)$ in a wide range, even millions of cycles, only limited by the computing hardware. So, there is no need of phase lock in absolute measurements nor of a variable phase shifter in differential measurements.
- The NCO.
- The presence of two separate NCOs enables the measurement of oscillators at different frequencies, for example 100 MHz input vs. 10 MHz reference.

Frequency range and noise are limited by the ADC technology and sampling frequency. Albeit ADCs are available to > 10 GHz, practical implementations prefer ≈ 100 MHz clock frequency, exploiting multiple Nyquist zones to achieve a typical 200 MHz maximum input frequency. For reference, the noise of an ADC at low input frequency (10 MHz) can be -100 to -105 dBrad² flicker, plus -150 to -160 dBrad²/Hz white.

4.1.4.2 More about the architecture of commercial instruments

Virtually all the commercial instruments make use of a dual channel scheme, where two equal branches are used to measure simultaneously the same DUT. Because DUT and reference oscillators are separated devices, and the phase detectors as well, it is sound to assume that their noise is statistically independent. By contrast, the DUT is common to the two branches, thus it is fully correlated, and captured by the statistical process. Thus, a correlation and averaging algorithm can reduce the background noise by a factor $\approx 1/\sqrt{m}$, where m is the number of averaged spectra.


FIGURE 18

Saturated-mixer dual-channel phase noise analyser

HDB-31-18

Figure 18 shows a saturated-mixer instrument based on this idea. Because a general-purpose instrument is expected to measure oscillators at arbitrary frequencies, two frequency synthesizers are inserted between the reference oscillators (generally, OCXOs) and the mixer. The oscillators are phase-locked to the input for the reasons, and with the criteria explained in § 4.1.4.1-b. Figure 18 can be easily modified to implement differential measurements (Fig. 15). In this case the PLLs are open, one synthesizer (or an external oscillator) is used as the common-mode oscillator, and two variable phase shifters (line stretchers) are needed to ensure the quadrature condition. Instruments of this type are manufactured by Arcale (acquired NoiseXT), Berkeley Nucleonics Corp., Keysight (it has its own instruments, and recently acquired Anapico), Maury Microwaves (acquired Holzworth), and Quantic Wenzel.

FIGURE 19
Direct-digitization dual-channel phase noise analyser



HDB-31-19

The direct-digitization version of the dual-channel instrument, shown in Fig. 19, is a straightforward application of Fig. 17 in the dual-channel configuration. Instruments of this type are produced by Arcale (NoiseXT) and Microchip. They are limited to 200 MHz maximum input frequency.

Rohde Schwarz produces phase noise analysers working up to 50 GHz input frequency. This is achieved with microwave down conversion before digitization and numerical processing. The internal architecture [Feldhaus G *et al.*, 2016] differs from that shown before.

After the heuristic reasoning, the mathematics of the single-channel noise rejection is explained. A more detailed treatise is available in [Rubiola E. *et al.*, 2010].

Referring to Fig. 18 or Fig. 19, the signals at the input of the FFT analyser, and their Fourier transforms (uppercase), are:

$$x(t) = [\varphi(t) - \psi(t)] \leftrightarrow X(f) = [\Phi(f) - \Psi(f)] \quad (4-19)$$

$$y(t) = [\varphi(t) - \theta(t)] \leftrightarrow Y(f) = [\Phi(f) - \Theta(f)] \quad (4-20)$$

The trivial gain of the system is implied. The quantity $\varphi(t)$ is the random phase of the oscillator under test, with no additional terms, while the random phases $\psi(t)$ and $\theta(t)$ account for the instrument noise (e.g. references, amplifiers, mixers, ADCs). It is sound to assume that $\varphi(t)$, $\psi(t)$ and $\theta(t)$ are statistically independent because they come from separate hardware. Their Fourier transforms are statistically independent as well. Only $\varphi(t)$ appears in both $x(t)$ and $y(t)$. The cross PSD is given by:

$$S_{yx}(f) = \frac{2}{T} Y(f)X^*(f) \quad (4-21)$$

where T is the measurement time, the superscript “*” stands for complex conjugate, and the factor “2” fixes the scale factor from the two-sided Fourier transform into to the one-sided PSD. Letting the frequency f implied for simpler notation, the above formula is expanded as:

$$S_{yx} = \frac{2}{T} (\Phi\Phi^* - \Phi\Psi^* - \Theta\Phi^* + \Theta\Psi^*) \quad (4-22)$$

The dual-channel FFT analyser measures the cross spectrum $\langle \widehat{S}_{yx} \rangle_m$ averaged on m data records of $x_T(t)$ and $y_T(t)$ acquired simultaneously over a time interval T :

$$\langle S_{yx} \rangle_m = \frac{2}{T} [\langle \Phi\Phi^* \rangle_m - \langle \Phi\Psi^* \rangle_m - \langle \Theta\Phi^* \rangle_m + \langle \Theta\Psi^* \rangle_m] \quad (4-23)$$

Thanks to the statistical independence mentioned, the cross terms $\langle \Phi\Psi^* \rangle_m$, $\langle \Theta\Phi^* \rangle_m$ and $\langle \Theta\Psi^* \rangle_m$ converge to zero for large m , and there remains only $\langle \Phi\Phi^* \rangle_m$. In conclusion, the estimated phase noise is:

$$\widehat{S}_\varphi = \frac{2}{T} \langle S_{yx} \rangle_m \quad (4-24)$$

The measurement process takes a time mT , not counting housekeeping (dead time for data transfer, calculations, etc.). Of course, for finite m , the terms $\langle \Phi\Psi^* \rangle_m$, $\langle \Theta\Phi^* \rangle_m$ and $\langle \Theta\Psi^* \rangle_m$ are not completely averaged out and set the measurement limit.

The estimation is a powerful concept because the simple average is not the one and only option, and other estimators can be considered. It is worth mentioning that $S_\varphi(f)$ is a real and positive quantity because $\Phi\Phi^*$ is obviously real and positive. By contrast, the Fourier transform is a complex quantity, thus $\langle \Phi\Psi^* \rangle_m$, $\langle \Theta\Phi^* \rangle_m$ and $\langle \Theta\Psi^* \rangle_m$ are complex, and consequently $\langle S_{yx} \rangle_m$ is complex. The background noise is equally split between real part and imaginary part. Therefore, equation (4-24) can be replaced with the real-part estimator:

$$\widehat{S}_\varphi(f) = \frac{2}{T} \langle \Re\{S_{yx}(f)\} \rangle_m \quad (4-25)$$

This results in higher sensitivity because the unnecessary part of the background noise is removed, at no cost in terms of hardware and computation complexity. A problem with equation (4-25) is that negative values are incompatible with the logarithmic scale (dBrad²/Hz), thus they must be removed. This results in a measurement bias (over-estimation of S_φ). It can be proved that equation (4-25) is the optimum estimator for white noise, that is, the estimator that converges to $S_\varphi(f)$ with the lowest m , or equivalently in the shortest measurement time. The background noise is:

$$S_\varphi(f) = \frac{2}{T} \frac{\sqrt{S_x(f) S_y(f)}}{\sqrt{2m}} \quad \text{averaging limit} \quad (4-26)$$

Because the background is equally split between $\Re$ and $\Im$ and the useful signal goes in $\Re$, the imaginary part $\langle \Im\{S_{yx}(f)\} \rangle_m$ is a useful estimator of the background noise.

Sadly, most instruments use the absolute-value estimator:

$$\widehat{S}_\varphi = \frac{2}{T} | \langle S_{yx} \rangle_m | \quad (4-27)$$

The drawbacks of equation (4-27) are that (i) it takes four-fold longer measurement time for the same rejection of the background noise of equation (4-25) because of the noise in the imaginary part, (ii) it is biased, because the average converges from positive values, and (iii) it hides negative, nonsensical outcomes. The latter point is critical in the case of low-noise oscillators, where the cross spectrum may contain large regions of negative values. This is discussed in [Nelson *et al.*, 2024; Gruson *et al.*, 2020].

4.2 Measurement techniques for transmission and reception links

Introduction

This subchapter focuses on measurement methods applied to the transmission medium itself –specifically, radio and optical links– where propagation effects, modulation schemes, and signal processing techniques determine overall accuracy and robustness.

Two principal approaches are examined. The first concerns software-defined radio (SDR) techniques applied to two-way satellite time and frequency transfer (TWSTFT) and related time metrology applications, where flexible, programmable architectures enable improved signal analysis and calibration. The second explores optical-fibre-based measurement techniques, addressing time and frequency transfer through amplitude, coherent, and heterodyne detection methods. Together, these methodologies illustrate how advanced digital signal processing and optical technologies are extending the limits of precision in transmission–reception metrology.


4.2.1 Software-Defined Radio measurements in time metrology

This section describes the principle and applications of software-defined radio (SDR) digital signal processing (DSP) for time metrology.

4.2.1.1 Software-defined radio principle

Software-Defined Radio (SDR) signal processing aims at implementing as much signal processing of radio frequency signals in the digital, discrete-time and discretized signal, domain [Report ITU-R SM.2152 (2009)]. The ideal SDR receiver implementation would sample, using radio frequency grade Analogue to Digital Converters (ADCs), the radio frequency band at a rate of at least twice the highest useful spectral component. This approach is often impractical because radio frequency grade ADCs lack resolution and hence dynamic range when detecting a weak signal next to a strong interference source, but most significantly samples a wide area of the spectrum that might not contain useful information for a given application. Hence, a common trade-off of SDR is to first convert the radio signal to baseband by using an analogue Local Oscillator (LO) and a mixer and then sampling the signal close to baseband. A typical approach is to convert from radio frequency band to baseband by using In-Phase and in-Quadrature (IQ) detectors in which a second copy of the LO, shifted by 90° is used to create the I and Q components matching the complex number representation. However, the analogue implementation of the IQ mixer is plagued by imperfections such as IQ imbalance in which the gains of I and Q channels are not equal and the phase between I and Q is not exactly 90° . Hence, a benefit of the SDR approach is to consider an intermediate step where the radio frequency band is brought close to baseband but at an Intermediate Frequency (IF) sufficiently high compared to the maximum bandwidth of interest using a single (real) mixer, sampling the signal sufficiently fast (at least twice the IF plus the useful bandwidth) and then performing the complex transposition in the digital domain. Indeed, by multiplying the real samples with $\cos\left(\frac{f_{\text{IF}}}{f_s} n\right)$ and $\sin\left(\frac{f_{\text{IF}}}{f_s} n\right)$ with f_{IF} the intermediate frequency, f_s the sampling frequency and $n \in \mathbb{N}$ the sample index, the resulting I and Q components exhibit ideally equal gain and quadrature component without requiring additional compensation. This technique, used for example in the frequency transposition of the Sentinel-1 C-band spaceborne Radio frequency Detection And Ranging (RADAR) [Piantanida *et al.*, 2019], will benefit to metrological application by avoiding the unwanted amplitude to phase conversion of mixers with the associated imperfections when analysing the spectral characteristics of a Device Under Test (DUT) (Fig. 20).

FIGURE 20
Software-defined radio receiver architectures



HDB-31-20

Note to Fig. 20:

Left: The ideal SDR receiver implementation sampling the raw radio frequency signal, at the expense of limited dynamic range and the risk of hiding a weak signal (green arrow) in the strong background noise (purple arrow) within the sampled frequency band (blue).

Middle: the practical implementation of SDR using an analogue transposition from radio frequency band to baseband followed by an IQ detector mixing with the LO and its quadrature copy to feed a dual channel ADC clocked at the sampling frequency f_s , plagued by IQ imbalance issues.

Right: the preferred SDR approach where a single mixer transposes the radio frequency signal by f_{LO} , generating an even spectrum; and the sampled signal is digitally transposed (red) to baseband by shifting by the additional IF f_{IF} with the signal image removed by low-pass filtering as shown.

All charts represent power P as a function of frequency f .

In an emitter context, the same strategy aims at reducing the number of analogue components in the transmitter chain, by feeding the Digital to Analog Converter (DAC) output to the antenna, possibly with a frequency mixer to transpose the baseband signal to radio frequency band. However, the unwanted spectral components will be radiated in the analogue domain. Hence, care must be taken to properly oversample and filter or remove the unwanted spectral components in the analogue domain.

Thus, the benefits of SDR, especially for Time and Frequency applications, include:

- *Stability:* A digital algorithm does not drift over time and is perfectly reproducible from one digital setup to another, without requiring the fine tuning of analogue parameters that would be needed when duplicating analogue measurement setups.
- *Flexibility:* The measurement parameters and feedback control are dynamically updated based on measurement results (e.g. Kalman filtering). Additionally, data logging and remote control are naturally included in networked computers. These are common requirements for long-term metrological monitoring.
- *Reconfigurability:* A single hardware board can be used for multiple applications, such as time and frequency transfer or phase-noise characterization.

However, the SDR also comes with some drawbacks:

- The *stability* of the sampling system is only as good as the LO clocking the ADCs and the jitter of the track and hold. Once the signal has been digitized, all subsequent processing is asynchronous digital processing but the stability of the analogue to digital conversion timing will define the metrological quality of the sampled signals.
- *Limited dynamic range* due to the discrete levels of the ADC. An ADC with N -bit resolution provides at most $6 \times N$ dB dynamic range between the weakest and strongest signal that can be digitized. In practical applications, the Effective Number of Bits (ENOB), which is usually less than the number of bits in the device will define this dynamic range.
- *Limited bandwidth:* Although digital systems can be clocked at hundreds of MHz or even GHz, each signal processing operation requires multiple cycles so that decimation is needed for complex

operations, and only a few MHz loop bandwidths will be available for closed loop control. Hence, the benefit of an SDR is more significant in the long term, for applications that require only low feedback loop bandwidth, or where analogue drifts must be removed to support stable digital processing. The SDR is less useful in a high-feedback loop-bandwidth where analogue processing remains dominant. Analogue processing can reach microwave frequencies when using acoustic or dielectric signal processing components.

SDR processing combines a hardware platform with rapid data transfer to the host computer. Digital signal processing (DSP) frameworks, based on High Performance Computing (HPC) benefit from the parallelization of executed instructions (Single Instruction/ Multiple Data – SIMD). Digital signal processing accepts the raw IQ stream and represents the characteristics of the radio frequency signal including Forward Error Correction (FEC) and hardware imperfection correction algorithms, such as a Costas loop for frequency correction, and a symbol synchronization for clock correction). It is a challenge to maintain this complex set of software, which includes drivers and signal processing code. The maintenance task is complicated because the code includes many dependencies and relies on multiple libraries. The time spent on software maintenance is significantly greater than the time spent developing applications and novel algorithms for Time and Frequency (TF) metrology.


4.2.1.2 Software-defined radio techniques applied to the TF field

Software-Defined Radio provides the long-term stability of digital signal processing and the accuracy of the IQ digital demodulation at the expense of bandwidth and dynamic range. Furthermore, the digital implementation is well suited for distributed processing in which multiple emitters or receivers are synchronized with delay compensation of the communication link between them. White Rabbit is one high stability implementation of TF distribution over a digital network and is being incorporated in some SDR frameworks³. Most time transfer schemes, starting with Global Positioning System (GPS) are based on spectrum spreading using pseudo-random sequences. The Linear Feedback Shift Register (LFSR) implementation of the pseudo-random sequence generators is perfectly suited to digital software implementation (Fig. 21 (A)). Just as the Pulse Repetition Interval (PRI), T , defines the range ambiguity in RADAR systems, the code length divided by the decimation factor defines the unambiguous timing capability (Fig. 21 (B)). For example, a 10 000-chip long code running at 2.5 MS/s leads to a PRI of $T = 4$ ms, which is much less than the ground to geo-stationary satellite two-way trip of about 260 ms, making the ranging measurement ambiguous.


³ <http://kb.ettus.com> Using Ethernet-Based Synchronization on the USRP%E2%84%A2N3xx Devices.

FIGURE 21

Software-defined radio implementation of a spread-spectrum time transfer system



(A)



(B)

HDB-31-21

Note to Fig. 21: (A) SDR emitter principle implemented as a GNU Radio Companion flowchart, with the generation of a maximum-length Galois Linear Feedback Shift Register (GLFSR), here configured with a 12-bit shift register for generating a 4 095-chip long code. Time of flight measurement with subsampling period resolution is demonstrated by spreading energy with adjacent samples in the low-pass filter before decimating: the single sample time delay introduced before decimation leads to sub-sampling period delay after the correlation computer as complex-conjugate product of the Fourier transforms of the reference and delayed-measurement signals. Despite the cross-correlation not being implemented as a GNU Radio block, the inverse Fourier transform of the complex conjugate products of the Fourier transforms of the reference and delayed signals is used. The Fourier transforms require converting the continuous streams to vectors for performing the integration. All these mathematical properties are readily implemented with the associated GNU Radio blocks as shown in (A). In (B), the asymmetric shape of the correlation is fitted with a polynomial function for fine identification of the cross-correlation peak position. In that case, the computation is performed offline as post-

processing of the recorded IQ streams. In this demonstration, $4\,095/3 = 1\,365$ samples after decimation by the low-pass filter broadcast at 2 MS/s led to a PRI of $T = 682.5\ \mu\text{s}$ as seen on the X-axis of the right chart as the delay between successive correlation peaks.

It may not be possible to use digital techniques for collecting and analysing timing information to generate a physical signal representing the original TF information that can be shared between SDR platforms. Just as the ADC timing determines all subsequent sample timestamps during asynchronous post-processing, generating a timing signal requires the same clock that was used to sample the input signals.

One major benefit of using an SDR to record signals is that the recordings can be used for post-processing and assessing the performance of various algorithms on the same datasets. However, the size of the datasets generated by SDR can become significant and can challenge the communication bandwidth to mass storage media. As an example, consider sampling the global navigation satellite system (GNSS) L1 band at 2 MS/s. If the resulting complex quantities are stored as 8-bit integers on two channels for direction of arrival analysis or anti-spoofing/anti-jamming processing, the file size is $2 \times 2 \times 2 = 8\ \text{MB/s}$ or 480 MB/min or 28.8 GB/h. The post-processing with the assumption that each sample timestamp is given by the sample index multiplied with the sampling period, requires that not a single sample is lost during the acquisition. This is challenging for a general-purpose multitasking operating system, and requires that the SDR acquisition software either runs with highest priority to avoid being pre-empted by lower priority processes, or that it runs within the real-time extension framework of the operating system. In either configuration, it will always be limited by hardware transfer bandwidth. The 2 MS/s is a conservative estimate for the legacy GPS L1 C/A code and will quickly increase with more complex modulations schemes and broader spectra as used in other GNSS frequency bands.

4.2.1.2 Use of SDR in TF applications

The identification of harmful interference to time transfer, the capacity to generate physical signals after digitization of the collected data (1 Pulse Per Second (PPS) output), the flexible and robust digital alternative to analogue modems, which can improve two-way TF transfers, and a technology well suited to synchronizing computer networks, are a few examples covered in this section.

A) SDR in Global navigation satellite system


Global navigation satellite systems are one of the most widely used means for TF dissemination. The digital communication scheme that uses Code Division Multiple Access (CDMA) for time transfer from Medium Earth Orbiting (MEO) satellites (20 000 km) makes the system prone to spoofing and jamming attacks [Uhrich *et al.*, 2019]. Such denial of service or malicious information injection is made even easier with the advent of SDR and the reconfigurable hardware compatible with sub-2 GHz GNSS L-bands. Timing information can be degraded with no detectable impact on the position solution by adjusting the clock tuning parameters in the navigation messages generated by the spoofing source. A SDR with access to raw IQ information collected from two antennas receiving the same signal, is ideally suited to determine the angle of arrival of the jamming or spoofing source and combining the data from the two antennas to cancel interfering signals emitted from that direction. The technique is an implementation of a Controlled Radiation Pattern Antenna (CRPA).

As an illustration of a SDR framework, GNU Radio is an opensource set of libraries that can be scheduled from a Python script and can process synthetic or real signal collected from hardware interfaces. The flowchart for processing GNSS signals and for identifying satellites by using Doppler-code maps to identify satellites in the acquisition phase is illustrated in Fig. 22. It is designed with GNU Radio Companion, the graphical interface generating the Python script. In this example, the classical loop over the frequency offset that compensates for the offset of the LO and for Doppler shifts are manually handled with a slider, and potential satellite code sequences are run in parallel. The code cross-correlation is implemented in the frequency domain for efficiency since the Fast Fourier Transform, *FFT*, reduces the computation load of cross-correlating the code sequence $p_i(t)$ representing satellite i with the noisy received signal $s(t)$ as $xcorr(s, p_i)(\tau) = iFFT(FFT(s) \cdot FFT^*(p_i))$ for time offset τ . The magnitude of the step of the frequency offset is set to a value much smaller than the inverse of the chip period of 1 ms and is typically 10 to 100 Hz. The frequency range must include the Doppler shift of $\pm 5\ \text{kHz}$ induced by the satellite motion as well as the LO offset, which depends on the SDR receiver quality.

This flowchart might be used as an introduction to the gnss-sdr framework [Fernández-Prades *et al.*, 2011] implementing the full GNSS analysis chain, from acquisition to tracking and Position-Velocity-Time (PVT) solution extraction. Since gnss-sdr products include the time offsets between the LO and GNSS oscillators, the resulting information can be used to steer the clock driving the ADC timestamping each new recorded sample and hence correct for LO fluctuations.

FIGURE 22

GNU Radio Companion flowchart for processing GPS L1 records by correlating with various Gold codes and identifying which satellite is visible with which Doppler shift



HDB-31-22

Note to Fig. 22: In this example the correlation between two Gold codes are computed, resulting in negligible correlation when two different codes are used and a sharp correlation peak with a width equal to the sampling period when running the same code on both File Sources. If one of the File Sources is replaced with a real signal record, e.g. the file `3genuine_69dB_11h15_1` on <https://iqengine.org> → GNU Radio SigMF Repo → Space folder, an additional frequency shift for manually compensating for Doppler shift and LO offset must be introduced with the Xlating FIR Filter block of GNU Radio. Notice how the iqengine.org hold both the data file and a meta file complying with the SigMF standard for describing the data organization and format.

Since GNSS signal decoding to extract PVT information involves multiple processing layers to convert the radio frequency signal to a useful information, starting from an existing opensource software makes it easy to add missing functionalities. For example, CRPA: gnss-sdr by CTTC (Centre Tecnològic de Telecomunicacions de Catalunya) in Spain implements all GNSS bands and constellation signal analysis, and also provides a rigorous framework based on GNU Radio to provide an extensible framework for adding custom processing features [Feng *et al.*, 2021].

SDR is suited for collected radio frequency signals and generating an IQ data stream. Once the IQ streams have been generated at the output of the ADC, the generation of a hardware output signal from these data might be desirable. As a demonstration of this principle, the expected output of a GNSS receiver is an accurate 10 MHz continuous wave and a 1 PPS timing signal that is steered by the GNSS disciplined oscillator. The gnss-sdr outputs the time difference between the local clock determined by the ADC sampling time and the GNSS constellation clock, and this information is used to steer the same clock used by the ADC to run a

counter to generate the 1 PPS output. Since the ADC is most often triggered by a Field Programmable Gate Array (FPGA) running reconfigurable hardware, the use of opensource gateware allows for adding the 1 PPS generation functionality to produce a physical signal as output of the SDR processing [Rabus *et al.*, 2021].

B) SDR in two-way satellite time and frequency transfer

Two-Way Satellite Time and Frequency Transfer (TWSTFT) is a long-range technique using a geostationary satellite for comparing remote clocks by exchanging orthogonal messages in a two-way protocol and recovering the time of flight [Recommendation ITU-R TF.1153-4 (08/2015)]. Although the satellite moves around its parking position at a rate of up to a few m/s, so that the time of flight varies by a few ns/s the contribution of the satellite motion is cancelled in the TWSTFT double-difference comparison [ITU-R Handbook on Satellite Time and Frequency Transfer and Dissemination (2010)] if the two signals broadcast from the two remote sites reach the satellite at close enough time. Sub-nanosecond comparisons can be realized by synchronizing the two remote stations to sub-100 ms uncertainty, e.g. using computer Network Time Protocol (NTP) synchronization.


TWSTFT is a primary technique for the generation of Coordinated Universal Time (UTC). About 20 timing laboratories around the world continuously operate TWSTFT links for remote time and frequency comparisons. The use of SDR receivers in TWSTFT ground stations (SDR TWSTFT) has been shown to have significantly reduced both diurnal and noise measurements [Huang *et al.*, 2016, Jiang *et al.*, 2018].

The successful experience with SDR technology for receivers in time metrology has opened the way to build a complete autonomous system based on an SDR emission-reception platform. Doing so would be consistent with a recommendation of the CCTF in 2021 [CCTF Recommendation on TWSTFT (2021)]. A fully SDR implementation of the emitter and receiver signal processing is described in [https://github.com/oscimp/amaranth_twstft] where spectrum spreading is achieved with pseudo-random number generator sequences, selected as orthogonal to identify each emitter at each receiver site. Oversampling the cross-correlation peak allows for improving the timing resolution from the sampling rate of 5 MS/s defined by the satellite channel bandwidth: the improvement factor is determined by the Signal to Noise Ratio (SNR) and allows reaching sub-nanosecond timing resolution and uncertainty. A post-processing approach is selected to avoid the limitations of inadequate computational power for real time calculation. Each record is post-processed to sequentially identify as many emitters as needed to recover the satellite velocity, the time of flight and possibly drift rate. The resulting information is shared over a computer network for TWSTFT calculation. The loopback time difference is measured between the sites and is used to remove the contribution of the communication channel to the delay.

As an example of such a measurement result, TWSTFT measurements were made over a period of three and a half weeks by using SDR emitter and receiver platforms developed in France by LNE-OP (formerly LNE-SYRTE) and LNE-LTFB. There were two 5-min measurement sessions during each odd UTC hour. The initial results were slightly better than those obtained with commercial analogue modems, as shown in Fig. 23. This validates the developments obtained in the laboratory, as well as their robustness [Achkar *et al.*, 2024; Friedt *et al.*, 2024]. The next objective is to work on improving the performance, including implementing Delay Locked Loop (DLL) and Phase Locked Loop (PLL) tracking.

FIGURE 23

Evolution of time differences (left) and corresponding frequency instabilities (right) of the two-way link OP-LTFB equipped with SATRE modems and SDR emitter-receiver platforms over a 3.5-week measurement



Note to Fig. 23: The “all SDR” measurement includes all SDR measurements at beginning and end of odd UTC hours, “ITU SDR” only uses 2-min records once every odd UTC hour similar to SATRE schedules. Note that the links (left) are not calibrated and that offsets (+7 134 ns for SATRE link and −47 ns for SDR link) have been applied in order to facilitate the comparison of the different plots.

C) Distributed SDR synchronization

SDR emitters and receivers are controlled by digital hardware that is well suited for synchronization over computer networks. One protocol relevant to time and frequency applications is PTPv2/White Rabbit where two-way messages are shared over fibre networks for compensating for the communication delays, similar in principle to TWSTFT. Thus, spatially separated SDR sources and sinks become coherent, allowing for direction of arrival in distributed RADAR systems or pattern detection using cross-correlation in Very Large Baseline Interferometry (VLBI) applications [Boven *et al.*, 2019; Hedekvist *et al.*, 2022].

DSP can benefit from upper Nyquist zone analysis by using intentional aliasing to remove analogue components that would otherwise be needed to convert the relevant frequencies to baseband. This is desirable since analogue components can be a source of drift and additive noise including phase to amplitude noise conversion [Matusko *et al.*, 2023].

Conclusion

In this section was discussed the software-defined radio digital signal processing for time metrology, emphasizing the benefits of stability, flexibility and reconfigurability. The principle of SDR, its advantages and drawbacks for the time and frequency area are described. Concrete applications of SDR to time metrology are detailed. Finally, a few examples on the use of SDR in time and frequency applications are provided. Relevant references in the covered field supported this contribution.

ITU References

- Recommendation ITU-R TF.1153-4 (08/2015) – The operational use of Two-Way Satellite Time And Frequency Transfer employing pseudorandom noise codes
- Report ITU-R SM.2152 (2009) – Definitions of Software-Defined Radio (SDR) and Cognitive Radio Spectrum (CRS)

Handbook on Satellite Time and Frequency Transfer and dissemination (2010)

References for § 4.2.1

- Achkar J, Meyer E, Chupin B, Meyer F, Chiu O, Lours M, Friedt JM. Two-Way Satellite Time and Frequency Transfer using an opensource, openhardware Software-Defined Radio platform. *Accepted for presentation at the 4th URSI AT-RASC*, Gran Canaria, Spain, 19-24 May 2024
- Boven P, Tour C, Koelmeij JCJ, Smets R, Szomoru A. Dwingeloo Telescope has Fringes (Again). *14th European VLBI Network Symposium & Users Meeting (EVN)*, 2019, DOI 10.22323/1.344.0156.
- CCTF 2021 Recommendation TWSTFT: On Developing and Utilizing the Next-Generation Modems and Technologies for Two-Way Satellite Time and Frequency Transfer (TWSTFT) Used in UTC Generation. https://github.com/oscimp/amaranth_twstft
- Feng W, Friedt JM, Goavec-Merou G, Meyer F. Software-Defined Radio Implemented GPS Spoofing and Its Computationally Efficient Detection and Suppression. *IEEE Aerospace and Electronic Systems Magazine* 36 (3), 2021, DOI 10.1109/MAES.2020.3040491
- Fernández-Prades C, Arribas J, Closas P, Avilés C, Esteve L. GNSS-SDR: an open source tool for researchers and developers. *Proc. of the 24th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS)*, Portland, Oregon, Sept. 19-23, 2011, pp. 780-794, <https://github.com/gnss-sdr/gnss-sdr>
- Friedt JM, Chupin B, Lours M, Meyer E, Chiu O, Meyer F, Daniau W, Achkar J. Results of a Software-Defined Radio (SDR) implementation of Two-Way Satellite Time and Frequency Transfer (TWSTFT) emitter and receiver system. *Submitted to the 37th European Frequency and Time Forum*, Neuchâtel, Switzerland, 25-27 June 2024.
- Hartl P. Present state of long-distance time transfer via satellites with application of the MITREX-modem. *MILCOM 1986 – IEEE Military Communications Conference: Communications-Computers: Teamed for the 90's*, 1986, DOI 10.1109/MILCOM.1986.4805775. <https://timetech.de/new-trends-in-two-way-time-and-frequency-transfer-via-satellite>
- Hedekvist PO, Ebenhag SC, Rieck C, Wallberg D, Bergroth M. Fiber Optic Time Transfer from UTC(k) to a VLBI Antenna in a Coherent Communication Network. *Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS)*, 2022, DOI 10.1109/EFTF/IFCS54560.2022.9850949
- Huang YJ, Fujieda M, Takiguchi H, Tseng WH, Tsao HW. Stability improvement of an operational Two-Way Satellite Time And Frequency Transfer system. *Metrologia* 53 881, 2016, DOI 10.1088/0026-1394/53/2/881
- Jiang Z, Zhang V, Huang YJ, Achkar J, Piester D, Lin SY, Wu W, Naumov A, Yang Sh, Nawrocki J, Sesia I, Schlunegger C, Yang Z, Fujieda M, Czubla A, Esteban H, Rieck C, Whibberley P. Use of SDR receivers in TWSTFT for UTC computation. *Metrologia* 55 685, 2018, DOI 10.1088/1681-7575/aacbe6
- Matusko M, Ryger I, Goavec-Merou G, Millo J, Lacroûte C, Carry E, Friedt JM, Delehay M. Fully-digital platform for local ultra-stable optical frequency distribution. *Rev. Sci. Instrum.* 94, 034716, 2023, <https://doi.org/10.1063/5.0138599>
- Piantanida R *et al.*, Sentinel-1 Level 1 Detailed Algorithm Definition. European Space Agency, 2019.
- Rabus D, Goavec-Merou G, Cabodevila G, Meyer F, Friedt JM. Generating a timing information (1 PPS) from a software-defined radio decoding of GPS signals. *Proc. Joint EFTF/IFCS*, 2021.
- Uhrich P, Abgrall M, Riedel F, Chupin B, Achkar J, Rovera GD. A Powerful Signal nearby L1 Frequency Band Jamming GNSS Stations in Observatoire de Paris. *ITU Journal: ICT Discoveries* 2 (1), 2019, <https://www.itu.int/en/journal/2019/001/Pages/01.aspx>

4.2.2 Optical fibre links: measurement methods

The choice of detection methods in optical communication systems is linked to the way information is emitted. Different modulation formats impose different requirements on the receiver, and understanding this relationship is fundamental to evaluating the strengths and weaknesses of each detection scheme. In this subchapter is provided a progressive overview of the main detection techniques in fibre-optic communication applied to time and frequency transfer, beginning with a short reminder of the key background concepts that underpin signal propagation in optical fibres. This serves as a foundation for discussing how modulation and detection are intertwined in time and frequency transfer.

First, the amplitude detection is examined, the simplest and most widely deployed and commercialized approach, which relies on intensity modulation of the optical carrier. Then are examined coherent detection techniques, which have become the standard for high-capacity data transmission over long-haul optical links but are still under active investigation for time and frequency transfer. Finally, heterodyne methods are discussed, which are particularly relevant in the context of ultra-precise frequency transfer and time-frequency metrology.

4.2.2.1 Background

A) Wave propagation, phase and group velocities

Consider a monochromatic spectral component of an optical wave propagating along the $+z$ direction. A local (quasi-)monochromatic waveform can be written as $s(t, z) = f(\omega t - \beta(\omega) z)$, where ω is the angular frequency, $\beta(\omega)$ is the propagation constant and $f(\cdot)$ is the waveform shape in the appropriate argument.

The *phase velocity* v_ϕ and the *group velocity* v_g are defined as:

$$v_\phi(\omega) = \frac{\omega}{\beta(\omega)} = \frac{c}{n(\omega)} \quad (4-28)$$

$$v_g(\omega) = \frac{d\omega}{d\beta} = \frac{c}{n(\omega) + \omega \frac{dn}{d\omega}} \quad (4-29)$$

To describe dispersion effects for a band of frequencies around a carrier ω_0 it is useful to Taylor-expand $\beta(\omega)$ about ω_0 :

$$\beta(\omega) = \beta_0 + \beta_1(\omega - \omega_0) + 1/2 \beta_2(\omega - \omega_0)^2 + \dots \quad (4-30)$$

with the commonly used coefficients

$$\beta_0 \equiv \beta(\omega_0) \quad (4-31)$$

$$\beta_1 \equiv \left. \frac{d\beta}{d\omega} \right|_{\omega_0} = \frac{1}{v_g(\omega_0)} \quad (\text{group delay per unit length}) \quad (4-32)$$

$$\beta_2 \equiv \left. \frac{d^2\beta}{d\omega^2} \right|_{\omega_0} \quad (\text{group-velocity dispersion, GVD}) \quad (4-33)$$

The leading effect of GVD on a narrow band pulse of spectral width $\Delta\omega$ propagating a distance L is a temporal broadening approximately given by:

$$\Delta T_{\text{broad}} \approx |\beta_2| L \Delta\omega \quad (4-34)$$

B) Relation between β_2 and the dispersion parameter D

Telecom literature frequently uses the chromatic-dispersion parameter D in units of ps/(nm·km) [Agrawal, 2021; Senior, 2008; Kaminow and Li, 2002; Rijckenberg *et al.*, 2007]. The two quantities are related by:

$$D(\lambda) = -\frac{2\pi c}{\lambda^2} \beta_2(\omega) \quad (4-35)$$

with $\omega = 2\pi c/\lambda$.

Table 14 provides typical values for CD parameters:

TABLE 14

Typical chromatic-dispersion values for SMF-28e fibre [ITU-T]; [Senior, 2008]; [Corning Incorporated].; [Optokon, a.s.]. The exact dispersion at 1 310 nm depends on the fibre's zero-dispersion wavelength λ_0 . The zero-dispersion region is typically ≈ 3.5 near 1 285-1 330 nm)

Wavelength (nm)	Chromatic dispersion (ps/(nm·km))
1 310	~ 0
1 550	$\sim 17-18$

Historical note

Long-haul optical links were once routinely equipped with dispersion compensation fibre (DCF) spools for chromatic dispersion compensation. With the development of coherent detection and digital signal processing, such passive compensation has been removed from backbone systems.

B) Polarization mode dispersion (PMD)

Optical fibres exhibit slight birefringence, meaning that the refractive index differs – by approximately 10^{-7} – between two orthogonal polarization axes of the electric field. As a result, in addition to chromatic dispersion, polarization mode dispersion (PMD) must also be considered when evaluating link performance. Polarization mode dispersion is characterized by the PMD coefficient D_{PMD} in ps/ $\sqrt{\text{km}}$. The corresponding differential group delay (DGD) is:

$$\tau_{\text{DGD}} \approx D_{\text{PMD}} \sqrt{L} \quad (4-36)$$

Early monomode fibres (manufactured in 1980s-1990s) often had PMD coefficients as high as 0.5 ps/ $\sqrt{\text{km}}$, and, even after improvements, 0.2 ps/ $\sqrt{\text{km}}$ was common. At the time, process control on core geometry, residual stress, and cabling was less sophisticated. Since the late 1990s, better MCVD/PCVD fabrication, improved tension control during draw, and optimized coating/cabling techniques have brought typical values for modern low-PMD fibres down to 0.04-0.06 ps/ $\sqrt{\text{km}}$. Standards such as ITU-T G.652.D and G.655 now specify $\text{PMD} \leq 0.2$ ps/ $\sqrt{\text{km}}$ as a worst-case, but most manufacturers routinely deliver much better (0.05 ps/ $\sqrt{\text{km}}$).

Table 15 summarizes representative values of the PMD coefficient for modern standard optical fibres.

TABLE 15

Typical PMD values for SMF-28e fibre [ITU-T]; [Senior, 2008]; [Corning Incorporated]

Wavelength (nm)	PMD (ps/ $\sqrt{\text{km}}$)
1 310	0.04-0.06
1 550	0.05-0.07

B) How dispersion couples with laser frequency noise and modulation bandwidth

Chromatic dispersion alone simply makes different frequency components travel at different group velocities. Two practical mechanisms make this important for system performance:

- *Laser frequency fluctuations*: if the optical carrier frequency ω_0 fluctuates (instantaneous deviation $\delta\omega(t)$), the propagation constant seen by the field changes by

$$\delta\beta \approx \beta_1 \delta\omega + 1/2 \beta_2 \delta\omega^2 + \dots$$

The accumulated phase at the receiver over distance L therefore contains terms proportional to $\beta_1 L \delta\omega$ and $\beta_2 L \delta\omega^2$. Practically, this means laser frequency noise is converted into phase and timing fluctuations at the receiver; the conversion strength scales with the dispersion coefficients (mainly

β_1, β_2) and with link length L . In other words, dispersion converts *frequency noise* into *phase / timing / amplitude* fluctuations.

- **Modulation (signal) bandwidth:** a modulated signal occupies a finite spectral width $\Delta\omega_{\text{sig}}$. Because group delay varies with frequency (via β_1 and higher-order terms), different spectral components of the modulated waveform acquire different delays. The result is intersymbol distortion (pulse broadening and overlap) that increases with β_2 , $\Delta\omega_{\text{sig}}$ and distance L . Thus, for a given fibre and length, higher data rates (larger signal bandwidth) cannot be supported unless dispersion is compensated or equalized.


In optical fibre links, the choice of detection technique is intrinsically linked to both the transmission distance and the target data rate, since these parameters determine the relative impact of dispersion, noise, and sensitivity limits. Chromatic dispersion imposes a fundamental trade-off between the achievable data rate (or modulation bandwidth) and the transmission distance, since both larger spectral widths $\Delta\omega$ and longer fibre lengths L increase the distortion. In addition, laser frequency noise becomes a critical factor in time-frequency metrology and in applications requiring ultra-low phase or timing noise, since dispersion converts carrier-frequency fluctuations directly into arrival-time and phase errors at the receiver. These conversion mechanisms emphasize the need for dispersion management techniques, or alternatively coherent detection combined with DSP.

4.2.2.2 Time-delay demodulation in fibre links

An essential aspect of optical fibre links, often overlooked in simplified treatments, is that the fibre itself acts as a delay line demodulator. Any optical field propagating along the link accumulates a phase proportional to its travel time, and spurious reflections or Rayleigh back scattering cause additional delayed replicas of the transmitted signal to be superimposed at the receiver. As a consequence, the detected signal contains not only the intended contribution from the nominal propagation path, but also interference terms from delayed copies of the field.

FIGURE 24

Schematic representation of fibre-link demodulation



HDB-31-24

Note to Fig. 24: Light propagating through the fibre undergoes multiple paths due to back scattering and reflections. At the photodetector, the superposition of delayed replicas of the optical field results in interference, which manifests as frequency mixing in the case of frequency transfer and as temporal distortions in time transfer.

For a monochromatic source of angular frequency ω_0 , the electric field at the output of the link of length L is, in its simplest form,

$$E(t) = E_0 e^{i(\omega_0 t - \beta L)} \quad (4-37)$$

If, however, part of the optical field is back-scattered or reflected from an intermediate point along the fibre, it will reach the detector with an additional time delay τ . The resulting detected intensity at the photodiode is therefore:

$$I(t) \propto |E_0 e^{i\omega_0 t} + E_1 e^{i\omega_0(t-\tau)}|^2 \quad (4-38)$$

which expands to:

$$I(t) \propto |E_0|^2 + |E_1|^2 + 2 \Re\{E_0 E_1^* e^{i\omega_0 \tau}\} \quad (4-39)$$

The last term illustrates the fundamental demodulation mechanism: a delayed replica of the signal mixes with the direct field, creating beat notes whose spectral content depends on the distribution of delays τ across the link.

If multiple back-scattered or reflected components are present along the link, each located at a different point i with an associated delay τ_i and complex amplitude E_i (including both attenuation and any local optical gain), the total detected intensity becomes:

$$I(t) \propto |E_0 e^{i\omega_0 t} + \sum_{i=1}^N E_i e^{i\omega_0(t-\tau_i)}|^2 \quad (4-40)$$

Expanding this expression gives:

$$I(t) \propto |E_0|^2 + \sum_{i=1}^N |E_i|^2 + 2 \Re\{\sum_{i=1}^N E_0 E_i^* e^{i\omega_0 \tau_i}\} + 2 \Re\{\sum_{i < j} E_i E_j^* e^{i\omega_0(\tau_i - \tau_j)}\} \quad (4-41)$$

The cross terms $E_0 E_i^* e^{i\omega_0 \tau_i}$ represent the beating between the main signal and each delayed replica, while the $E_i E_j^*$ terms describe beating among delayed components themselves. Because $|E_i|$ incorporates both propagation losses and any optical gain (e.g. from bidirectional amplifiers), the relative strength of these contributions varies along the link. The resulting interference produces a complex spectral signature in frequency-transfer measurements and a correspondingly intricate temporal response in time-transfer experiments.


In frequency transfer, this manifests as a complex radio-frequency (RF) spectrum at the photodiode output, with side-bands corresponding to the beating between the main carrier and delayed replicas. Narrow-band RF filtering is typically used to isolate the desired heterodyne beat associated with the full round-trip signal.

In time transfer, the same principle results in multiple transitions in the detected waveform. Spurious reflections or scattering introduce secondary steps, temporally shifted with respect to the main rising edge. In this case, the challenge is to identify the genuine arrival time of the pulse, which requires time-domain filtering around the expected transition region.

These delay-induced demodulation effects highlight why careful link engineering is necessary: stray reflections, connectors, or insufficient isolation at components can severely distort both the frequency-domain spectrum and the time-domain response. The phenomenon is intrinsic to fibre-based TF transfer, and understanding it provides the key to interpreting both complex beat spectra and distorted timing traces.

FIGURE 25

Illustration of the two filtering strategies for time or frequency transfer detection methods: Frequency transfer (narrow-band carrier analysis, continuous signal, no time selection) and time transfer (broadband edge detection, pulse signal, no frequency selection)




A) Distinction between time transfer and frequency transfer

While the same opto-electronic conversion chain can be employed for both time transfer and frequency transfer, the requirements at the signal-processing stage differ fundamentally (see Fig. 26).

FIGURE 26

Schematic comparison between frequency transfer (narrow-band carrier analysis) and time transfer (broadband edge detection)



HDB-31-26

For frequency transfer, the quantity of interest is a well-defined carrier tone (typically in the MHz or GHz range). Accordingly, the photo-detected signal is passed through a narrow band filter centred on the carrier frequency, which suppresses noise outside the useful band and maximizes signal-to-noise ratio (see Fig. 25). The signal is continuous and there is no time selection, and no dead-time. In this case, the quality of transfer is primarily determined by the stability of the carrier phase and the rejection of out-of-band perturbations.

In contrast, time transfer relies on detecting the precise timing of sharp transitions in the received signal. Here, the fastest possible rise time is required, implying that the detection chain must preserve a wide electrical bandwidth. Instead of frequency-domain filtering, the emphasis is on time-domain processing: the signal is analysed over a carefully selected time span around the expected transition (see Fig. 25). The signal is pulsed, and there is no frequency selection. Such methods allow extraction of the timing information with high precision but also place stringent requirements on the front-end electronics and overall system bandwidth.

Time transfer methods, especially those relying on digital edge detection, require careful processing to achieve high precision. In particular, deglitching is often necessary to remove outliers or spurious events that could degrade synchronization performance. A prime example of such advanced processing is found in the White Rabbit (WR) technology [Lipinski *et al.*, 2010].

It is important to highlight here the fundamental relationship between time and frequency transfer. Frequency information can be recovered from a time transfer by differentiating the timing edges to obtain the successive time intervals, corresponding to the period of the time base. Conversely, integrating the transferred frequency yields a time error rather than an absolute time value, since the integration constant remains unknown. Any interruption in the frequency transfer introduces a new, undefined offset in the time position. In this sense, time transfer provides a more complete and self-contained reference, whereas frequency transfer alone cannot unambiguously establish absolute time without additional information.

B) Analogue versus digital detection methods

In fibre-based time and frequency transfer, the detection of the transmitted signal can be implemented using either analogue or digital techniques. Both approaches are compatible with time transfer and frequency

transfer, but they differ fundamentally in the way the signal is processed and in the aspects of the signal they exploit.

C) Analogue methods

Analogue methods operate on the continuous waveform of the optical or electrical signal. For frequency transfer, analogue detection typically involves narrow-band filtering of a continuous-wave (CW) signal, where phase or frequency fluctuations are directly measured. For time transfer, analogue detection can be performed on pulsed signals by extracting the temporal envelope, often requiring high-bandwidth detection but without discrete event processing.

D) Digital methods

Digital methods, by contrast, rely on the identification of discrete signal events –primarily the edges or transitions of the waveform. In time transfer applications, digital techniques effectively perform time-of-arrival detection of pulses, while in frequency transfer they may track phase increments via successive edge events. Because edges carry broadband information, digital detection can be seen as inherently broadband, aligning these methods with the family of time-detection techniques, even when the ultimate goal is frequency transfer. This perspective emphasizes that digital techniques are fundamentally tied to temporal resolution rather than narrow-band spectral filtering.

E) Front-end electronics

Once the optical field is converted into an electrical signal by the photodiode, several approaches are possible for further processing. The most straightforward case is *direct detection*, already described in § 4.1.2, where the photocurrent is directly related to the incident optical power (see equation (4-18)). Alternatively, *heterodyne methods* (see § 4.1.3) can be employed, where the received optical field is mixed with a local oscillator laser. In this case, the intermediate frequency (IF) signal produced at the output of the photodiode retains the amplitude and phase information of the received field,

$$i(t) \propto E_{\text{rx}}(t)E_{\text{LO}}(t) \quad (4-42)$$

enabling highly sensitive measurements of frequency and phase deviations. A further possibility, relevant in the context of precise timing, is to use dual-mixer time measurement systems – either in analogue or digital form (see § 4.1.3.2) – which provide sub-nanosecond resolution in time stamping and delay evaluation.

In all cases, the output of the photodiode must be conditioned by suitable front-end electronics. A trans-impedance amplifier (TIA) is typically used to convert the weak photocurrent into a usable voltage signal while providing the required bandwidth. Depending on the application, the front end may also include low-noise amplification stages, equalization, or specialized circuitry optimized for high dynamic range. The design of this electronic front end has a direct impact on the achievable sensitivity, linearity, and ultimately the stability of the time or frequency transfer.

4.2.2.3 Amplitude detection

A) Principle of operation

This is the simplest optical detection technique: the optical carrier is directly detected in intensity. The drive current of the laser diode is modulated, typically using an On-Off Keying (OOK) format.

Photons incident on a photodiode generates a photocurrent proportional to the incident optical power. Let the optical power be:

$$P(t) = P_0[1 + m \cdot s(t)] \quad (4-43)$$

where P_0 is the average optical power, m is the modulation index, and $s(t)$ is the normalized data signal (e.g. ± 1 or $0/1$).

The photocurrent is then:

$$i_{\text{ph}}(t) = R \cdot P(t) \quad (4-44)$$

where R is the photodiode responsivity in A/W. Since the optical power is proportional to the square of the optical field $E(t)$, the detected signal follows:

$$P_{\text{det}}(t) \propto |E(t)|^2 \quad (4-45)$$

B) Common modulation formats (digital methods)

The simplest form is NRZ (Non-Return-to-Zero), where the signal remains constant during the bit period and changes only between bits. Other formats include RZ (Return-to-Zero), where the signal returns to zero between bits. This implies shorter duty cycle, a better timing recovery but also a broader spectral occupancy. Manchester coding and multi-level amplitude modulation offer various trade-offs in spectral efficiency and detection complexity. A wide variety of modulation formats has been developed. For a comprehensive review, interested readers may refer to specialized literature [Kaminow and Li, 2002; Senior, 2008].

In the case of analogue transmission, the most common formats are sine-wave and square-wave modulation. The essential distinction with respect to digital techniques lies in the continuity of the modulation waveform: in analogue schemes the signal is continuous in time, whereas in digital communication the modulation is structured by coding schemes (e.g. 8b/10b) and embedded within a protocol framework.

C) Amplitude-modulation TF transfer in practice

Such amplitude modulation methods are convenient for time transfer and RF frequency transfer. A variety of engineering solutions implement time and frequency transfer over optical fibre using intensity (amplitude) modulation on the physical layer. Among commercially available solutions, notable examples include Ethernet-based *White Rabbit* (WR) multi-point networks, point-to-point time/frequency transfer over fibre links, and bidirectional point-to-point Wavelength Division Multiplexing (WDM) approaches (see § 6.3.2.5).

Table 16 summarizes representative offerings and reported performance for RF transfer at short averaging time (Allan deviation at $\tau = 1$ s), together with the stated or typical fibre reach. These AM/IM schemes encompass both digital implementations (e.g. WR, see § 6.3.2.5) and analogue intensity-modulated links (e.g. ELSTAB, see § 6.4.1). The various implementation details and architectures are discussed more extensively in Chapter 6. Here, the intent is not to provide an exhaustive survey but rather to illustrate typical performance ranges available from current technologies. The reported ranges correspond to links without optical amplification. When amplification is employed, and depending on the link architecture, the achievable distance can exceed 500 km.

TABLE 16
Representative AM/IM-based time/frequency transfer solutions over fibre

Product	1-s stability $\times 10^{-13}$	Range [km] (no amplifiers)
(1) White Rabbit [Włostowski 2011]; [Tour <i>et al.</i> , 2017]; [Dierikx <i>et al.</i> , 2016] [Kaur, 2017]	$< \sim 10$	~ 10 -120
(2) TimeTech OFC 10568/10569 [TimeTech GmbH 2015]	12	5
(3) TimeLink Micro TMO3000 [TimeLink Microsystems 2023]	100	50
(4) Linear Photonics / MACOM		
TimeLink DL [MACOM / Linear Photonics 2023]	< 0.8	25
(5) ELSTAB [Krehlik <i>et al.</i> , 2016] [PIKTime Systems 2024]	< 3	60-100

Note to Table 16: “1 s stability” is the Allan deviation (ADEV) at $\tau = 1$ s of the relative frequency stability of the transferred signal. Ranges are as stated or typically achievable per the referenced documentation (single span, per hop, without optical amplification). Principles: (1) Ethernet medium with physical layer1 syntonization + PTP, intensity on fibre, 10 MHz / 1 PPS. (2) Analogue IM link, 10 MHz / 1 PPS. (3) Digital/analogue time and RF (SFP-based). (4) Analogue Time-over-fibre (RF or PPS). (5) 10 MHz and 1-PPS, bi-directional fibre noise cancellation.

D) Main limitations

Range limitations

A first limitation to be considered is the achievable range of intensity-modulated direct-detection systems. The detected photocurrent is proportional to the square of the received optical field (see equation (4-46)). As a consequence, the link attenuation is effectively squared, which in logarithmic units corresponds to a doubling of the loss in dB when compared to heterodyne methods (see § 4.2.2.5). In practice, the minimum detectable optical power depends not only on the transimpedance amplifier (TIA) performance but also on the photodiode technology. For standard PIN photodiodes with typical TIAs, detection thresholds are in the range of -40 to -30 dBm. Avalanche photodiodes, which provide internal gain, are commonly employed in long-reach SFP modules. Devices marketed with the “EX” (extended-range) designation can extend the optical reach from approximately 40 km to 80 km, and, in optimized configurations, up to about 160 km. This scaling imposes stringent constraints on the maximum transmission distance between hops, both for analogue and for digital modulation formats, as the fundamental power penalty cannot be mitigated by coding or digital signal processing. Consequently, the achievable range and link budget are determined by the available optical power, the receiver sensitivity, and the overall optical link budget, which must balance transmission losses against the attainable optical gain.

AM–PM conversion

In time-frequency metrology, amplitude fluctuations in the detected signal are partially converted into phase fluctuations, causing instabilities and biases. A simple expression for the induced phase variation is:

$$\varphi(t) = \alpha \cdot \frac{\Delta P(t)}{P_0} \quad (4-46)$$

where α is the AM-PM conversion coefficient in radians per unit relative amplitude.

A practical difficulty is that AM-PM is Device- and operating-point dependent. AM-PM conversion depends strongly on photodiode type, bias, temperature, optical spot position, wavelength, and photocurrent. Many devices show nulls (zero crossings) and sharp peaks as photocurrent changes, so any values are only representative [Williams *et al.*, 2008]; [Li *et al.*, 2012]. In Table 17 is include the approximate photocurrent (see equation (4-45)) assuming a typical responsivity at 1 550 nm $R = 0.8$ A/W.

TABLE 17

Representative ranges of AM-PM conversion coefficient α (rad per fractional optical power change) for PIN (InGaAs) and GaAs photodiodes. Photocurrent estimated with $R = 0.8$ A/W

Input power	Optical power	Photocurrent	PIN (InGaAs)	GaAs / fast PDs
(dBm)	(mW)	nA	α (rad per $\Delta P/P$)	α (rad per $\Delta P/P$)
-40	10^{-4}	80	0.01-0.1	0.02-0.2
-30	10^{-3}	800	0.01-0.2	0.02-0.3
-20	10^{-2}	8.0×10^3	0.02-0.4	0.03-0.6
-10	10^{-1}	8.0×10^4	0.05-1.0	0.1-1.5
0	1.0	8.0×10^5	0.1-2.0	0.2-3.0

Limited modulation depth and dynamic range

The modulation index m is constrained by both the linearity of the laser and the photo-detector. Excessive modulation ($m > 1$) can drive the laser or detection front-end into a nonlinear regime, generating harmonic distortion and intermodulation products that degrade the fidelity of the transmitted signal. Such distortions can broaden the spectrum, introduce spurious frequency components, and complicate subsequent filtering or demodulation. Conversely, choosing too small a modulation index ($m \ll 1$) reduces the amplitude of the modulated signal relative to the noise floor, thereby lowering the signal-to-noise ratio (SNR). In practice, for analogue optical links, modulation indices in the approximate range of $0.5 \leq m \leq 1$ are often used to balance

linearity and SNR while minimizing distortion. The optimal value depends on the specific laser, TIA, and photodiode characteristics, as well as the required dynamic range and link budget.

Chromatic dispersion

Chromatic dispersion causes a frequency-dependent group delay, distorting the signal.

To evaluate the effect, following [Weiershausen and Schneiders, 2011], the approximate relationship between maximum achievable bit rate $B_{\max}$ and fibre length L is:

$$B_{\max}^2 \cdot L < \frac{1}{\pi|D| \cdot \Delta\lambda} \quad (4-47)$$

where D is the dispersion coefficient in ps/(nm·km) and $\Delta\lambda$ is the source spectral width.

In practice, time and frequency transfer links with digital encoding use 1.25 GBs/s or less data rate. Time and frequency transfer links with Analogue modulation use 5-100 MHz.

Direct detection is low-cost, robust, and ideally suited for short and medium-haul links. It remains a dominant technique for local time and frequency dissemination, with fractional-frequency instabilities as low as a few parts in 10^{-13} at 1-second averaging.

4.2.2.4 Coherent telecom techniques

A) Principle of coherent detection

In contrast to amplitude-modulation detection, where information is retrieved from the intensity of the received optical field, coherent techniques rely on mixing the incoming signal with a local oscillator (LO). This allows access not only to the amplitude of the optical field but also to its phase and frequency. In this sense, the detection is “coherent” since the phase relationship between the signal and the LO is preserved. A convenient framework to describe coherent transmission is the in-phase (I) and quadrature (Q) representation of the field:

$$E(t) = I(t) \cos(\omega_0 t) + Q(t) \sin(\omega_0 t)$$

where ω_0 is the carrier frequency, and $I(t)$ and $Q(t)$ are the baseband modulating signals. By detecting both quadratures, coherent systems can exploit amplitude, phase, frequency, and even polarization degrees of freedom.


B) Modulation formats

Similar to non-coherent systems, coherent communication may employ amplitude, frequency, or phase modulation. However, with access to both quadratures, coherent systems can implement richer modulation formats. Figure 27 illustrates three of the most commonly used modulation formats.

FIGURE 27

Constellation diagrams for common coherent modulation formats

(a) BPSK with two symbols along the I -axis, (b) QPSK with four quadrature points, (c) 16-QAM with a 4×4 grid



Binary phase-shift keying

The simplest phase modulation format is Binary Phase Shift Keying (BPSK), where the carrier phase is shifted between two states (0 and π) to encode a binary symbol:

$$E(t) = A \cos(\omega_0 t + \pi b(t))$$

with $b(t) \in \{0,1\}$ representing the transmitted bit.

Quadrature phase-shift keying (QPSK)

By modulating both quadratures independently, one can encode two bits per symbol using four phase states separated by $\pi/2$.

Quadrature amplitude modulation (QAM)

In QAM, both the amplitude and phase of the carrier are modulated. A symbol corresponds to a point in the I - Q plane, and higher-order constellations such as 16-QAM, 64-QAM, or 256-QAM allow encoding multiple bits per symbol. These formats achieve higher spectral efficiency at the cost of increased sensitivity to noise and phase errors.

Overview of common modulation formats

Table 18 summarizes the most common coherent modulation formats used in fibre-optic communications and their typical achievable data rates in modern systems.

TABLE 18

Common coherent modulation formats and indicative achievable data rates (per channel)

Modulation format	Bits per symbol	Typical data rates
BPSK	1	Up to ~ 100 Gb/s
QPSK	2	100-200 Gb/s
8-QAM	3	200-400 Gb/s
16-QAM	4	400 Gb/s - 1 Tb/s
64-QAM	6	Beyond 1 Tb/s (lab demonstrations)

These values are indicative and depend on channel spacing, symbol rate, forward error correction (FEC) overhead, and the transmission distance. For detailed discussion, see [Winzer *et al.*, 2018]; [Agrell *et al.*, 2016].

C) Coherent detection architecture

Many coherent detection schemes have been proposed in the literature, reflecting the diversity of approaches to exploiting amplitude, phase, frequency, and polarization of the optical carrier. A comprehensive review of these methods is given in [Ip *et al.*, 2008] which summarizes the principal approaches and their relative merits. Among these, the configuration that has become standard in modern optical communications is the coherent transceiver for polarization-division multiplexed (PDM) systems, employing advanced vector modulation formats and DSP techniques. A representative architecture is illustrated in Fig. 28 [Faruk and Savory, 2017].

At the transmitter, two independent data streams corresponding to the x- and y-polarizations are first processed by the DSP. This processing may include encoding, modulation, precompensation for linear and nonlinear transmission impairments, and pulse-shaping filtering. The digitally processed signals are then converted to analogue form using four digital-to-analogue converters (DACs), representing the in-phase (I) and quadrature (Q) components for each polarization. These analogue signals drive two IQ modulators (IQMs), which are fed by a single laser source split through a polarization beam splitter (PBS). The outputs of the IQMs are subsequently combined with a polarization beam combiner (PBC) and launched into the optical fibre.

At the receiver, a phase- and polarization-diversity front-end maps the incoming optical signal into four electrical signals corresponding to the I and Q components of each polarization (x and y). This is achieved using two PBSs and a pair of 90 degrees optical hybrids, followed by balanced photodetection, yielding:


$$\begin{bmatrix} I_x \\ Q_x \\ I_y \\ Q_y \end{bmatrix} = \begin{bmatrix} \text{Re}(E_x E_{LO}^*) \\ \text{Im}(E_x E_{LO}^*) \\ \text{Re}(E_y E_{LO}^*) \\ \text{Im}(E_y E_{LO}^*) \end{bmatrix} \quad (4-48)$$

where E_x and E_y denote the dual-polarization electric fields of the received signal, E_{LO} is the local oscillator field, and $\text{Re}(\cdot)$ and $\text{Im}(\cdot)$ represent the real and imaginary components, respectively.

The detected signals are amplified using transimpedance amplifiers (TIAs) before digitization by analogue-to-digital converters (ADCs) for further DSP-based recovery.

FIGURE 28

Configuration of DSP-based coherent transceiver from [Faruk and Savory, 2017]



HDB-31-28

Legends to Fig. 28:

DAC: digital-to-analogue converter	IQM: IQ modulator
PBS: polarization beam splitter	PBC: polarization beam combiner
LO: local oscillator	TIA: trans-impedance amplifier
ADC: analogue-to-digital converter	DSP: digital signal processing

D) Coherent-modulation time and frequency transfer in practice

While intensity modulation has been widely utilized for time and frequency transfer over optical fibres, coherent modulation techniques may offer enhanced performance, particularly in terms of spectral efficiency and noise resilience. These methods leverage phase and quadrature components of the optical field, enabling more precise control and measurement of time and frequency signals. To date, coherent modulation formats have been less explored for time and frequency transfer and remain largely in the research domain.

In practical implementations, coherent modulation schemes often employ Mach-Zehnder modulators (MZMs) to encode time and frequency information onto optical carriers. For instance, dual-drive MZMs can facilitate quadrature amplitude modulation (QAM), allowing for the transmission of multiple bits per symbol, thereby increasing the data rate and improving the signal-to-noise ratio (SNR) over long distances. This approach is particularly beneficial in urban environments where fibre links are subject to various noise sources [Li *et al.*, 2019].

An example of such an application is the work developed in [Lu *et al.*, 2021], which demonstrates fibre-optic time-frequency transfer in gigabit Ethernet networks over urban fibre links. Their system utilizes coherent detection techniques to achieve high-precision synchronization across multiple nodes.

Another notable example is the sub-ps stability time transfer method based on optical modems, using analogue signal imprinted by an electro-optic modulator [Frank *et al.*, 2018]. This method employs simple BPSK modulation format and coherent homodyne detection to achieve sub-picosecond timing stability.

Beyond these examples, the integration of optical frequency combs has further advanced coherent modulation techniques. Optical frequency combs generate a series of equally spaced frequency lines, which can be used as a reference for time and frequency transfer. When combined with high-order modulation formats like 256-QAM, these systems can achieve ultra-high data rates and precise synchronization over optical fibres (Ali *et al.*, 2025).

E) Limitations of coherent techniques

Coherent systems face several challenges:

Laser linewidth

A critical limitation arises from the finite linewidth of practical lasers. The linewidth determines the phase noise of the optical carrier, which directly affects phase-modulated formats such as QPSK and higher-order QAM. The phase error variance is proportional to the product of the laser linewidth $\Delta\nu$ and the symbol duration T_s :

$$\sigma_\phi^2 \approx 2\pi\Delta\nu T_s$$

As the modulation format order increases, the constellation points become more densely packed in the I - Q plane, making the system more sensitive to phase noise. For quadrature phase-shift keying (QPSK), linewidths on the order of a few MHz are typically tolerable. However, advanced formats such as 16-QAM require linewidths below ~ 100 kHz, and for 64-QAM or beyond, even narrower linewidths (down to tens of kHz) are necessary. Narrow-linewidth external-cavity lasers or integrated cavity designs are therefore essential for high-order coherent systems.

Polarization mode dispersion (PMD)

Single-mode fibres are not perfectly isotropic: slight asymmetries and environmental fluctuations cause the two polarization modes to propagate at different group velocities. This effect, known as polarization mode dispersion, results in pulse broadening and depolarization of the signal. PMD is typically quantified by the differential group delay (DGD), denoted τ_{DGD} , which grows with the square root of the fibre length:

$$\tau_{DGD} \approx D_{PMD}\sqrt{L}$$

where D_{PMD} is the PMD coefficient (typically 0.05 – 0.2 ps/ $\sqrt{\text{km}}$ for modern fibres) and L is the transmission length in km. While PMD can be partially compensated using digital signal processing, it remains a limiting factor for very high symbol rates and long-haul links.

Digital signal processing (DSP) complexity

Coherent receivers rely on powerful DSP algorithms to perform carrier recovery, polarization demultiplexing, and equalization of chromatic dispersion and PMD. While DSP enables remarkable flexibility and resilience, its implementation requires very high-speed electronics.

While these algorithms enhance system performance, they introduce complexities in time and frequency transfer applications, as they introduce timing jitter, which arises from factors such as sampling clock jitter, algorithmic delay, and Phase-Locked Loop (PLL) dynamics (Faruk and Savory, 2017). Recent reports show less than 1 ps skew error for a 1 200-km 32-GBd DP-16QAM optical transmission experiment [Diniz *et al.*, 2016]. In addition to jitter, DSP algorithms introduce processing delays, that may fluctuate and therefore can affect the accuracy of the transferred signals. For instance, in multi-channel or networked DSP platforms, computational resources are shared. High processing loads or dynamic reallocation of resources (e.g. due to varying traffic) can modulate latency, adding a data-flux-dependent jitter component.

To the best of our knowledge the impact on time jitter and time offset in time-transfer applications has not yet been thoroughly investigated or consistently reported in the literature.

Fibre non linearities

At high optical powers, the fibre refractive index exhibits nonlinear behaviour through the Kerr effect. This leads to self-phase modulation (SPM), cross-phase modulation (XPM), and four-wave mixing (FWM). In coherent systems, these effects distort both amplitude and phase, creating inter-channel interference in wavelength-division multiplexed (WDM) systems. The severity of nonlinear impairments increases with optical power and transmission distance. The following formulas provide the basis for quantifying the impact of fibre nonlinearities on time and frequency transfer.

Self-phase modulation (SPM)

Self-phase modulation induces a nonlinear phase shift in the transmitted signal, causing a frequency chirp. It can lead to frequency inaccuracy and timing jitter. The instantaneous frequency shift $\Delta\omega(t)$ for a Gaussian pulse is given by:

$$\Delta\omega(t) = \frac{4\pi L n_2 I_0}{\lambda \tau^2} t \exp\left(-\frac{t^2}{\tau^2}\right) \quad (4-49)$$

where L is the fibre length, n_2 is the nonlinear refractive index, I_0 is the peak intensity, λ is the wavelength, and τ is the pulse duration.

Cross-phase modulation (XPM)

Cross-phase modulation occurs when the intensity of one channel modulates the phase of another, leading to inter-channel interference. The induced phase shift $\Delta\phi_{\text{XPM}}$ between two channels is:

$$\Delta\phi_{\text{XPM}} = \gamma L P_{\text{signal}} P_{\text{pump}} \quad (4-50)$$

where γ is the nonlinear coefficient, L is the fibre length, and P_{signal} and P_{pump} are the powers of the signal and pump channels, respectively.

Four-wave mixing (FWM)

Four-wave mixing is a third-order nonlinear process in which three optical waves at frequencies f_1 , f_2 , and f_3 interact to generate a fourth wave at frequency f_4 :

$$f_4 = f_1 + f_2 - f_3 \quad (4-51)$$

Four-wave mixing generates new spectral components through nonlinear interactions between different frequency components. The efficiency of FWM depends on the fibre's nonlinear properties, power levels, and phase matching. For continuous-wave signals, the generated idler power P_4 at fibre length L can be expressed as:

$$\eta_{\text{FWM}} = \frac{P_4(L)}{P_1 P_2 P_3} = \frac{9\gamma^2 L_{\text{eff}}^2}{4} \text{sinc}^2\left(\frac{\Delta\beta L}{2}\right) \quad (4-52)$$

where P_1, P_2, P_3 are the powers of the interacting waves, and $P_4(L)$ is the generated idler power, $\gamma = \frac{2\pi n_2}{\lambda A_{\text{eff}}}$

is the fibre nonlinear coefficient, $L_{\text{eff}} = \frac{1-e^{-\alpha L}}{\alpha}$ is the effective interaction length, accounting for fibre loss α , $\Delta\beta = \beta_1 + \beta_2 - \beta_3 - \beta_4$ is the phase mismatch, with β_i the propagation constant at frequency f_i , and $\text{sinc}(x) = \frac{\sin x}{x}$ represents the reduction in FWM efficiency due to phase mismatch. Perfect phase matching ($\Delta\beta = 0$) maximizes FWM efficiency, while dispersion-induced phase mismatch reduces it.

The phase noise introduced by FWM is certainly more difficult to evaluate. It can be quantified by the power spectral density (PSD) of the generated idler waves [Anthur *et al.*, 2013]. The impact on timing stability depends on the efficiency of FWM and the spectral separation between channels.

In summary, although coherent detection enables advanced modulation formats and unprecedented data rates in high-speed telecommunications, its use for time and frequency transfer remains limited –likely due to the increased complexity of detection and the time jitter and asynchronous operation introduced by DSP. Its

performance is fundamentally constrained by the linewidth of the lasers, fibre impairments such as PMD and nonlinearities, and the computational overhead of DSP.

4.2.2.5 Heterodyne techniques

As already introduced in § 4.2.2.3, once the optical field is converted into an electrical signal by the photodiode, different strategies can be employed for signal processing. In particular, heterodyne detection methods (§ 4.1.3) offer significant advantages compared to direct detection. By mixing the received optical field with a local oscillator (LO) laser, an intermediate frequency (IF) signal is generated at the photodiode output,

$$i(t) \propto E_{\text{rx}}(t)E_{\text{LO}}(t) \quad (4-53)$$


which carries both the amplitude and the phase fluctuations of the transmitted signal. This approach enables highly sensitive measurement of frequency deviations and phase noise.

The key advantage of heterodyne detection lies in its efficiency at very low received power levels. Because the LO provides a strong reference field, the detection process does not rely on optical pre-amplification, thereby avoiding noise figure penalties. This makes heterodyne techniques a method of choice in long-haul and bidirectional optical links, where the end-to-end transmission losses can easily exceed -30 dB.

Beyond their role in signal recovery, heterodyne methods naturally extend into optical frequency transfer, where phase-coherent information carried by the optical field is preserved and analysed at radio or intermediate frequencies. In this context, they provide the fundamental building block for a large class of ultra-precise frequency dissemination schemes.

FIGURE 29

General principle of heterodyne detection. (a) The received optical field is mixed with a local oscillator to generate an intermediate frequency (IF) signal that retains amplitude and phase information. (b) In optical frequency transfer, super-narrow filtering is achieved using tracking oscillator




Among ultra-precise dissemination techniques, optical frequency transfer is a prominent example that relies fundamentally on heterodyne detection. In the following, the focus is on this case and discuss its specific limitations in comparison with amplitude-based methods.

A) Context: principle of principle optical frequency transfer

Optical frequency transfer through fibre links relies on precise measurement of the phase difference between the emitted optical field and the same field after propagation along the fibre. The concept is often described in terms of a strongly imbalanced Michelson interferometer, where the fibre link acts as one arm, see Fig. 30. The signal returning after a round trip accumulates the phase fluctuations introduced by the fibre, which can then be detected and actively compensated [Ma *et al.*, 1994]; [Foreman *et al.*, 2007]. More information is provided in Chapter 6.

FIGURE 30

Simplified sketch of an optical frequency transfer system. AOM: acousto-optic modulator.
FM: Faraday mirror. Polarization control devices are omitted



HDB-31-30

A key point is that, unlike in telecommunication systems, the laser source is typically left unmodulated: no amplitude, frequency, phase, or polarization modulation is applied. Instead, the phase of the pure optical carrier is preserved, and its stability is exploited for frequency metrology. The measurement principle is thus based on detecting the Doppler shifts imprinted by fibre-induced fluctuations on the returning light.

B) Experimental implementation

In practice, optical frequency transfer systems employ frequency shifters, most commonly acousto-optic modulators (AOMs), at the fibre ends [Jiang *et al.*, 2008]; [Terra *et al.*, 2010]. This ensures that the beat note corresponding to the light that has travelled the entire round trip appears at a well-defined frequency offset, separating it from spurious contributions such as stray reflections or Rayleigh scattering. These parasitic signals beat at different frequencies and can be suppressed by narrow-band radio-frequency filtering.

Large-scale networks, such as the French REFIMEVE infrastructure, further incorporate tracking oscillators to follow the optical beat note [Cantin *et al.*, 2021]. Typically, loop bandwidths on the order of 10-100 kHz are used, which allows selective detection and robust rejection of unwanted noise contributions [Lopez *et al.*, 2010].

Dual-polarization coherent receivers, with similar architecture of coherent telecommunication, were also employed successfully and demonstrate robust operation (Clivati *et al.*, 2020).

C) Optical frequency transfer in practice

Significant progress in optical frequency transfer has been achieved in research laboratories worldwide, enabling robust and ultra-stable frequency dissemination over long-haul fibre links.


In typical field deployments, heterodyne detection techniques allow for very low received power levels, with optical budgets in the range of -55 dBm to -60 dBm. For a launch power of 0 dBm and assuming a fibre attenuation of approximately 25 dB per 100 km, this corresponds to an achievable link span of about 250 km without intermediate amplification.

Under these conditions, state-of-the-art systems demonstrate fractional frequency instabilities on the order of 5×10^{-16} at one-second integration time within a 1 Hz measurement bandwidth over a 250 km fibre link.

D) Limitations

FIGURE 31

Dependence of the laser coherence length with laser linewidth, highlighting the need for kilohertz-class or narrower lasers for 100 km-scale coherent links



HDB-31-31

The performance of optical frequency transfer links is closely tied to the properties of the laser source. The line-width of the laser determines the coherence length, which in turn limits the maximum achievable link length. A rough scaling is that transmitting over 100 km requires a laser line-width of 1 kHz or less. The relationship between laser line-width $\Delta\nu$, link length L , and achievable phase stability can be approximated as:

$$L \lesssim \frac{c}{2n\pi \Delta\nu \tau} \quad (4-54)$$

where c is the speed of light, n the refractive index of the fibre, and τ the round-trip time. Figure 31 illustrates this dependence. For link lengths between 100 km and 10 000 km, ultra-stable lasers are required. While such sources are now commercially available, their cost remains prohibitive for widespread deployment.

Due to the use of narrow-linewidth, unmodulated lasers, the impact of chromatic dispersion is negligible compared to intensity-modulated systems. The remaining technical limitations arise from:

- *Polarization Mode Dispersion (PMD)*: Although greatly reduced in metrological transfer compared to telecom signals, PMD sets a residual stability floor, typically at the level of 10^{-20} - 10^{-21} in fractional frequency stability for long-haul links [Xu *et al.*, 2018].
- *Interferometer stability*: The fibre-based interferometers used to detect the round-trip phase fluctuations can themselves introduce instability. With careful design, their contribution can be suppressed to the level of 10^{-18} down to 10^{-21} , as demonstrated in state-of-the-art systems [Xu *et al.*, 2018].

As a result, modern optical frequency transfer links are capable of transmitting ultra-stable signals with fractional instabilities well below 10^{-19} at one-second integration, making them invaluable for applications in time-frequency metrology, fundamental physics, and geoscience.

ITU texts

ITU-T. n.d. *Recommendation ITU-T G.652 – Characteristics of a Single-Mode Optical Fibre and Cable*.

References for § 4.2.2

- Agrawal, Govind P. 2021. *Fiber-Optic Communication Systems*. 5th ed. Wiley.
- Agrell, Erik, Magnus Karlsson, Alan R. Chraplyvy, Peter J. Winzer, David J. Richardson, *et al.*, 2016. “Roadmap of Optical Communications.” *Journal of Optics* 18 (6): 063002. <https://doi.org/10.1088/2040-8978/18/6/063002>
- Ali, B. *et al.*, 2025. “Optical Frequency Comb-Based 256-QAM WDM Coherent Optical Fiber Communication.” *Photonics* 4 (2): 16. <https://doi.org/10.3390/photonics4020016>
- Anthur, Aravind P., Regan T. Watts, Tam N. Huynh, Deepa Venkitesh, and Liam P. Barry. 2013. “A General Phase Noise Relationship for Four-Wave Mixing.” *arXiv:1308.0914*. <https://arxiv.org/abs/1308.0914>
- Barnes JA, Chi AR, Cutler LS *et al.*, “Characterization of Frequency Stability,” *IEEE Transact Instrum Meas* 20(2) pp.105-120, May 1971
- Boudot R, Rubiola E, [Phase noise in RF and microwave amplifiers](#), *IEEE Transact. Ultrason., Ferroelec. Freq. Control* 59(12) pp.2613-2624, December 2012
- Cantin, E., M. Tønnes, R. L. Targat, A. Amy-Klein, O. Lopez, and P.-E. Pottie. 2021. “An Accurate and Robust Metrological Network for Coherent Optical Frequency Dissemination.” *New Journal of Physics* 23 (5): 053027. <https://doi.org/10.1088/1367-2630/abc79e>
- Clivati, Cecilia, Paolo Savio, Silvio Abrate, *et al.*, 2020. “Robust Optical Frequency Dissemination with a Dual-Polarization Coherent Receiver.” *Optics Express* 28 (6): 8494–511. <https://doi.org/10.1364/OE.378602>
- Corning Incorporated. n.d. *Corning SMF-28e / SMF-28e+ Optical Fiber – Product Information / Datasheet*
- Dierikx, Erik F. *et al.*, 2016. “White Rabbit Precision Time Protocol on Long-Distance Fiber Links.” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 63 (7): 945–52. <https://doi.org/10.1109/TUFFC.2016.2518122>
- Diniz, Júlio César Medeiros, Edson Porto da Silva, Molly Piels, and Darko Zibar. 2016. “Joint IQ Skew and Chromatic Dispersion Estimation for Coherent Optical Communication Receivers.” *Signal Processing in Photonic Communications 2016* (Vancouver, Canada), July. <https://orbit.dtu.dk/en/publications/joint-iq-skew-and-chromatic-dispersion-estimation-for-coherent-optical-communication-receivers>
- Donley E (chair) *et al.*, *IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities*, IEEE, New York (NY, USA), December 2022. DOI: 10.1109/IEEESTD.2022.9973001
- Exail. 2025. *Regeneration Laser Station*. <https://www.exail.com/product/regeneration-laser-station>.
- Faruk, Md Saifuddin, and Stephen J. Savory. 2017. “Digital Signal Processing for Coherent Transceivers Employing Multilevel Formats.” *Journal of Lightwave Technology* 35 (5): 1125-41. <https://doi.org/10.1109/JLT.2016.2649399>
- Feldhaus G, Roth A, A 1 MHz to 50 GHz direct down-conversion phase noise Analyzer with cross-correlation, Proc. European Freq. Time Forum, Article no. 1012, York, UK, 4-7 April 2016

- Foreman, S. M., A. D. Ludlow, M. H. G. de Miranda, J. E. Stalnaker, S. A. Diddams, and J. Ye. 2007. "Coherent Optical Phase Transfer over a 32-Km Fiber with 1 s Instability at 10^{-17} ." *Physical Review Letters* 99 (15): 153601. <https://doi.org/10.1103/PhysRevLett.99.153601>
- Frank, F., F. Stefani, P. Tuckey, and P. E. Pottie. 2018. "A Sub-Ps Stability Time Transfer Method Based on Optical Modems." *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 65 (7): 1121–28. <https://doi.org/10.1109/TUFFC.2018.2833389>
- Gruson Y, Rus A, Rohde UL, Roth A, Rubiola E, Artifacts and Errors in Cross-Spectrum Phase Noise Measurements, *Metrologia* 57(5) 055010, October 2020. DOI: 10.1088/1681-7575/ab8d7b
- Halford D, A general mechanical model for $|f^\alpha|$ spectral density random noise with special reference to flicker noise $|1/f|$, *Proc. IEEE* 56(3) pp. 251-258, March 1968. DOI: 10.1109/PROC.1968.6269
- Kaur N, Frank F, Pottie P-E, Tukey P [2017] *Time and frequency transfer over a 500 km cascaded White Rabbit network* 2017 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS), Besancon, France, 2017, pp. 86-90, doi: 10.1109/FCS.2017.8088808
- Krehlik P, Śliwczynski Ł, Buczek Ł, Kolodziej J and Lipinski M [2016] *ELSTAB-fiber optic time and frequency distribution technology – a general characterization and fundamental limits* in *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 63, no. 7, pp. 993-1004, 2016.
- Ip, E., A. P. T. Lau, D. J. F. Barros, and J. M. Kahn. 2008. "Coherent Detection in Optical Fiber Systems." *Optics Express* 16 (2): 753–91. <https://doi.org/10.1364/OE.16.000753>
- Jiang, H., F. Kéfélian, P. Lemonde, *et al.*, 2008. "Long-Distance Frequency Transfer over an Urban Fiber Link Using Optical Phase Stabilization." *Journal of the Optical Society of America B* 25 (12): 2029–35. <https://doi.org/10.1364/JOSAB.25.002029>
- Jiaxing Time-transfer Optoelectronics Co., Ltd. 2025. *Jiaxing Time-Transfer Optoelectronics Co., Ltd. - Home*. <https://www.time-transfer.com/>
- Kaminow, Ivan P., and Tingye Li, eds. 2002. *Optical Fiber Telecommunications IV*. Academic Press.
- Li, F. *et al.*, 2019. "Optical i/q Modulation Utilizing Dual-Drive MZM for Fiber-Wireless Integration." *Optics Letters* 44 (17): 4235–38. <https://doi.org/10.1364/OL.44.004235>
- Li, X., D. Huang, S. A. Diddams, *et al.*, 2012. "Investigation of Amplitude-to-Phase Conversion in High-Speed Photodiodes." *2012 IEEE International Frequency Control Symposium (IFCS)*, 1-4. <https://doi.org/10.1109/FCS.2012.6243721>
- Lipinski, M., J. Serrano, *et al.*, 2010. "White Rabbit: Sub-Nanosecond Timing Distribution over Ethernet." *IEEE Transactions on Nuclear Science* 57: 1300-1308.
- Lopez, O., A. Haboucha, B. Chanteau, C. Chardonnet, A. Amy-Klein, and G. Santarelli. 2010. "Cascaded Multiplexed Optical Link on a Telecommunication Network for Frequency Dissemination." *Optics Express* 18 (16): 16849-57. <https://doi.org/10.1364/OE.18.016849>
- Lu, Z. *et al.*, 2021. "Fiber-Optic Time-Frequency Transfer in Gigabit Ethernet Networks over Urban Fiber Links." *Optics Express* 29 (8): 11693-703. <https://doi.org/10.1364/OE.422727>
- Ma, L.-S., P. Jungner, J. Ye, and J. L. Hall. 1994. "Delivering the Same Optical Frequency at Two Places: Accurate Cancellation of Phase Noise Introduced by an Optical Fiber or Other Time-Varying Path." *Optics Letters* 19 (21): 1777-79. <https://doi.org/10.1364/OL.19.001777>
- MACOM / Linear Photonics. 2023. *Linear Photonics Components and Systems Overview*. Product brochure. <https://www.rellpower.com/wp/wp-content/uploads/2023/08/Linearizer-Photonics-Product.pdf>

- Meher PK, Valls J, Tso-Bing J *et al.*, 50 Years of CORDIC, *IEEE Transact Circuits Syst I* 56(9) pp. 1893-1907, September 2009. DOI: 10.1109/tcsi.2009.2025803.
- Nelson CW, Hati A, Howe DA, A collapse of the cross-spectral function in phase noise metrology, *Rev Sci Instrum* 85(3) 024705, March 2024. DOI: 10.1063/1.4865715.
- Optokon, a.s. n.d. *Single Mode Fiber g.652.d – Product / Technical Data*.
- Owen D, Good Practice to Phase Noise Measurement, NPL, Teddington, UK, May 2004. ISSN: 1368-6550.
- PIKTime Systems. 2024. *OSTT-4 Optical Systems for Time and Frequency Transfer*. Product datasheet. <https://piktime.com/offer/ostt-4/>
- Rijckenberg, G.-J., A. M. J. Koonen, and A. Ng'Oma. 2007. "Chromatic Dispersion Tolerance of the Optical Frequency Multiplication Technique for Radio over Single Mode Fibre." *Proceedings of the 2007 Conference on Optical Fiber Communication and the National Fiber Optic Engineers Conference (OFC/NFOEC)*.
- Rubiola E, *Phase noise and frequency stability in oscillators*, Cambridge University Press 2010. ISBN: [978-0-521-15328-7 Paperback](#)
- Rubiola E, Gros Lambert J, Brunet M, and Giordano V, Flicker noise measurement of HF quartz resonators, *IEEE Transact Ultrason, Ferroelec. Freq. Control* 47(2) pp.361-368, March 2000. DOI: 10.1109/58.827421.
- Rubiola E, *Tutorial on the double-balanced mixer*, arXiv:physics/0608211 [physics.ins-det], August 2006. DOI: 10.48550/arXiv.physics/0608211.
- Rubiola E, Vernotte F, *The cross-spectrum experimental method*, [arXiv:1003.0113 \[physics.ins-det\]](#), February 2010. Still not published in a peer-reviewed journal.
- Rubiola E, Vernotte F, The companion of Enrico's Chart for phase noise and two-sample variances, *IEEE Transact. Microw. Theory Techn.* 71 (7) pp. 2996-3025, July 2023. The Enrico's Chart for Phase Noise and Two-Sample Variances, URL: <https://doi.org/10.5281/zenodo.4399218>
- Senior, John M. 2008. *Optical Fiber Communications: Principles and Practice*. 3rd ed. Prentice Hall.
- Terra, O., G. Grosche, and H. Schnatz. 2010. "Brillouin Amplification in Phase Coherent Transfer of Optical Frequencies over 480 Km Fiber." *Optics Express* 18 (15): 16102-11. <https://doi.org/10.1364/OE.18.016102>
- TimeLink Microsystems. 2023. *TMO1510 – Transport of RF Signals over Optical Fiber*. Product web page. http://www.timelinkmicro.info/wp-content/uploads/2025/02/fp2158a4_TMD1510_TMDO1510_en.pdf
- TimeTech GmbH. 2015. *10568/10569 Time and Frequency Optical Links (Datasheet)*. ION PTTI archive. https://www.ion.org/ptti/upload/files/1039_Datasheet_10568_10569.pdf
- Tour, C. van *et al.*, 2017. *Sub-Nanosecond Time Accuracy and Frequency Distribution Using White Rabbit*. ngVLA Memo #22. https://library.nrao.edu/public/memos/ngvla/NGVLA_22.pdf
- Volyanskiy K, Cussey J, Tavernier H, Salzenstein P, Sauvage G, Larger L, Rubiola E, Applications of the optical fibre to the generation and to the measurement of low-phase-noise microwave signals, *Journal Optical Society of America B: Optical Physics* 25(12) pp.2140-2150, December 2008. DOI: 10.1364/JOSAB.25.002140.
- Weiershausen, Werner, and Malte Schneiders. 2011. "Transport Solutions for Optically Amplified Networks." In *Optically Amplified WDM Networks*, edited by Andrew D. *et al.*, Ellis. Elsevier. <https://doi.org/10.1016/B978-0-12-374965-9.10011-1>
- Welch PD, The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms, *IEEE Transact Audio Electroacoust* 15(2) pp. 70-73, June 1967.

- Williams, P. A., J. J. McFerran, and J. Taylor. 2008. “Amplitude-to-Phase Conversion in Photodiodes for Femtosecond Lasers.” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 55 (8): 1639–46. <https://doi.org/10.1109/TUFFC.2008.841>
- Winzer, Peter J., David T. Neilson, and Alan R. Chraplyvy. 2018. “Fiber-Optic Transmission and Networking: The Previous 20 and the Next 20 Years.” *Optics Express* 26 (18): 24190–239. <https://doi.org/10.1364/OE.26.024190>
- Włostowski, Tomasz. 2011. “Precise Time and Frequency Transfer in a White Rabbit Network.” PhD thesis, Warsaw University of Technology. https://white-rabbit.web.cern.ch/documents/Precise_time_and_frequency_transfer_in_a_White_Rabbit_network.pdf
- Xu, D., W.-K. Lee, F. Stefani, O. Lopez, A. Amy-Klein, and P.-E. Pottie. 2018. “Studying the Fundamental Limit of Optical Fiber Links to the 10^{-21} Level.” *Optics Express* 26 (8): 9515–27. <https://doi.org/10.1364/OE.26.009515>

Where to learn more on the topic of § 4.1.4:

- Samuel R. Stein, Chapter 12 in *Frequency and Time – their measurement and characterization*, Precision Frequency Control, edited by E. A. Gerber and A. Ballato, New York, Academic Press, 1985.
- David W. Allan, *Time and Frequency (Time-domain) characterization, estimation and prediction of precision clocks and oscillators*, IEEE Trans. On UFFC, vol. 34, No. 6, 1987, pp. 647-654.
- S. Stein, D. Glaze, J. Levine, J. Gray, D. Hilliard, D. Howe, and L. Erb, Performance of a high accuracy phase measurement system, Proceedings of the 36 Annual Symposium on Frequency Control, 1982, pp. 314-320.

These papers and many others are reprinted in *Characterization of Clocks and Oscillators*, edited by D. B. Sullivan, D. W. Allan, D. A. Howe, and F. L. Walls, NIST Technical Note 1337, March 1990. This publication is available online from tf.nist.gov.

The relevant notations and several basic concepts about phase noise and frequency stability presented here are found in [Rubiola E *et al.*, 2023; Donley E (chair) *et al.*, 2022]. Such readings are recommended to everyone, together with the reference chart [The Enrico’s Chart for Phase Noise and Two-Sample Variances, URL: <https://doi.org/10.5281/zenodo.4399218>]. The booklet [Owen D., 2004] is recommended too, albeit elderly.

CHAPTER 5

Time Scales

TABLE OF CONTENTS

	<i>Page</i>
List of Acronyms.....	112
5.1 Introduction	113
5.1.1 Universal time	113
5.1.2 Ephemeris Time	114
5.1.3 International Atomic Time	114
5.1.4 Coordinated Universal Time	115
5.1.5 UTCr: the rapid realization of UTC	116
5.2 Time scales in general relativity	117
5.2.1 Coordinate systems in general relativity	117
5.2.2 The 1991 IAU Resolution A4	118
5.2.3 The 2000 IAU Resolution B1.9.....	119
5.2.4 International Atomic Time	119
5.2.5 Other time scales	119
5.3 Generation of time scales.....	120
5.3.1 Expected qualities	120
5.3.2 Timing data	122
5.3.3 Time scale algorithms	125
5.3.4 Accuracy of the scale interval of a time scale	132
5.3.5 Examples	132
5.4 Dissemination of time scales	134
5.4.1 Synchronizing the time scale of a laboratory	134
5.4.2 Paper time scales	136
References	132
ITU Texts	133

List of acronyms

BeiDou	Chinese GNSS
BIH	Bureau international de l'heure
BIPM	Bureau international de poids et mesures
CCDS	Consultative committee for the definition of the second, now CCTF
CCTF	Consultative committee for time and frequency
CGPM	General Conference of Weights and Measures
CIPM	Comité international des poids et mesures
EAL	Echelle Atomique Libre
ET	Ephemeris time
Galileo	European GNSS
GLONASS	GLOBAL NAVIGATION Satellite System (RUS)
GNSS	Global navigation satellite system
GPS	Global positioning system (US)
IAG	International Association of Geodesy
IAU	International Astronomical Union
IERS	International Earth Rotation and Reference Systems Service
NIST	National institute for standards and technology (US)
NPL	National Physical Laboratory (UK)
NTSC	National timing centre of China
OP	Observatoire de Paris
PFS	Primary frequency standard
PL	Acronym for the designation of the Polish national time scale
PTB	Physikalisch-Technische Bundesanstalt (DE)
RMO	Regional Metrology Organization
SU	Acronym used for historical reasons in the designation of the Russian national time scale
TAI	International Atomic Time
TCB	Coordinated Barycentric Time
TCG	Coordinated Geocentric Time
TT	Terrestrial time
TWSTFT	Two-way satellite time and frequency transfer
USNO	United States Naval Observatory
UTC	Coordinated Universal Time
UTC _r	“Rapid” UTC
UT1	Universal Time corrected by polar motion, rotational time

5.1 Introduction

The true nature of time has no rational explanation; we simply feel that time never stops or reverses. But, apart from philosophical considerations, experience indicates that any event can be localized by specifying three space coordinates and one temporal coordinate, generally denoted (x, y, z, t) . A time scale is thus defined as the time axis of a system of coordinates. In the Glossary, a time scale is defined as *a family of time codes for a particular coordinate time that provide an unambiguous time ordering of events*. The phrase “time scale” is also used as a shorthand for the algorithm that provides the names of events.

At first sight, establishing a time scale seems a simple task, as any evolving system allows the transformation of the measure of time variation into the measure of a dimensional quantity. But scientists ask for metrological properties, and reliable, stable, and accurate reference time scales are required. In addition, there is no absolute time as conceived by Newton in classical mechanics. In reality, all temporal phenomena are affected by gravitational fields and velocities with respect to the observer. Time scales must thus be defined in the framework of general relativity, as explained in the Resolution 2 “On the definition of time scales” adopted by the 26th General Conference of Weights and Measures in 2018 [Comptes rendus CGPM, 2018].

Conventionally, one distinguishes between two different types of time scale: integrated time scales and dynamic time scales:

- For integrated time scales, the primary data is a unit of duration, that is, of time interval, defined from a physical phenomenon. The time scale is constructed by fixing a conventional origin and by accumulating units of duration without dead-time and without interruption. This concept for the construction of a time axis was applied to the duration of the day, leading to the definition of Mean Solar Time. The present worldwide reference time scale is Coordinated Universal Time UTC, derived from International Atomic Time, TAI. TAI is an integrated time scale; it is obtained by the accumulation of atomic seconds defined as a number of periods of the radiation corresponding to a given transition of the caesium atom.
- For dynamic time scales, the primary data results from the observation of a dynamic physical system, described by a mathematical model in which time is a parameter that unambiguously identifies the configuration of the system. The time measurement thus becomes a position measurement, and the unit of time is defined as a particular duration. Universal Time corrected for the effects of polar motion, represented by UT1, and Ephemeris Time ET, are dynamic time scales, based respectively on the rotation of the Earth on its axis and on its revolution around the Sun.

In the past, a number of time scales, either dynamic or integrated, have been defined. The associated unit of duration was then used to define the second of the International System of Units (SI). The change from one definition to another has been motivated by the desire to improve accuracy and by other considerations such as the usefulness of the definition in practical metrology. In these changes, the condition of continuity between successive definitions has been maintained. A brief summary is given in the following sections.

5.1.1 Universal Time

Universal Time is affected by irregularities due to the motion of the polar axis of the Earth, and by the irregular velocity of rotation of the Earth. UT1 is the form of Universal Time corrected for the polar motion, thus affected by the irregularities of the Earth’s rotation. UT1 is proportional to the angle of rotation of the Earth on its axis. The coefficient of proportionality is chosen so that 24 hours of UT1 are close to the mean duration of the day, and the phase is chosen so that 0 h UT1 corresponds, on average, to midnight in Greenwich.

The associated unit of time is the second of mean solar time. Its definition is not very precise: it is *the fraction 1/86 400 of the mean solar day*. This was the SI second till 1960. Astronomers estimated that it could be estimated with a fractional uncertainty of order 10^{-9} , this level of accuracy being achieved after a decade of astronomical observations, analysis and averaging.

The UT1 was the reference time scale, and thus the basis of legal time, until 1972. It still provides a record of the Earth’s rotation for geophysics, astronomy, etc.

5.1.2 Ephemeris Time

Ephemeris Time (ET) is a dynamic time, derived from the theory of the Earth's orbital revolution around the Sun: it is realized by an expression for the mean longitude of the Sun. This expression and the initial phase of ET were chosen so that ET and UT1 were in approximate coincidence in 1900.

The associated unit of time is the ephemeris second, defined as *the fraction 1/31556925,9747 of the tropical year for 1900 January 0 at 12 hours ephemeris time*. This was the SI second from 1960 to 1967.

Through observations of planets and of the moon, it has been possible to obtain the time differences between ET and UT1 since 1630, with a precision of several tens of seconds for the 17th century, and several seconds for the 19th century; ET thus constitutes a reference for studying the Earth's rotation in the past.

5.1.3 International Atomic Time

International Atomic Time (TAI) is an integrated time available since 1955. The unit of time is the atomic second, which became the SI second in 1967 and is still in use. Its definition, adopted by the 13th Conférence générale des poids et mesures in 1967, is as follows:

The second is the duration of 9 192 631 770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the caesium 133 atom.

The atomic second can be realized in a laboratory. The best commercial caesium clocks, now widely used in timing centres, produce it with a stated accuracy of order 1×10^{-13} .

Following the decision of the 26th CGPM, the text of the definition is now: “The second, symbol s, is the SI unit of time. It is defined by taking the fixed numerical value of the caesium frequency $\Delta\nu_{\text{Cs}}$, the unperturbed ground-state hyperfine transition frequency of the caesium 133 atom, to be 9 192 631 770 when expressed in the unit Hz, which is equal to s⁻¹” [Comptes rendus CGPM, 2018]. This definition implies, as before, the exact relation $\Delta\nu_{\text{Cs}} = 9\,192\,631\,770$ Hz.

Some laboratories maintain primary frequency standards based on caesium; these give an ultimate accuracy of realization of order 1×10^{-16} (see Chapter 3). The progress in the development of optical frequency standards with fractional-frequency uncertainties of order 10^{-18} resulted in the recommendations by the Comité consultative du temps et de la fréquence (CCTF) and the Comité international des poids et mesures (CIPM) that some transitions be accepted as secondary representations of the second [Riehle *et al.*, 2018, Margolis *et al.*, 2024].

The function of a clock is to provide a continuous, ordered set of markers with an associated numbering. This constitutes a time scale. The number associated with a given marker is designated as the “reading of the clock”. Since physical devices can fail, laboratories are inevitably led to keep several clocks. Clock data are collected and treated together through a combination of their readings in order to generate an ensemble time. On a worldwide scale, such an ensemble time is International Atomic Time, TAI.

The definition of TAI was first provided by the Comité consultatif pour la définition de la seconde (CCDS) (name of the committee at that time, now CCTF) in 1970. It reads as follows:

International Atomic Time (TAI) is the time reference coordinate established by the Bureau International de l'Heure on the basis of the readings of atomic clocks operating in various establishments in accordance with the definition of the second, the unit of time of the International System of Units.

The origin of TAI was defined to coincide with UT1 on 1 January 1958.

An important consequence of improved accuracies in the realization of the atomic second is that relativistic effects are significant. In this context, the definition of the second must be understood as the definition of proper time, *i.e.* strictly speaking, the user must be in the neighbourhood of the clock and at rest with respect to it. When comparing two realizations of the SI second on the surface of the Earth, differences of a few parts in 10^{13} may appear due to the different gravitational fields to which the clocks are subjected. The definition of TAI was thus extended as follows, in a declaration of the CCDS [CCDS Report, 1980] during its 9th Session held in 1980:

TAI is a coordinate time scale defined in a geocentric reference frame (origin of the frame at the centre of the Earth) with the SI second as realised on the rotating geoid as the scale unit.

In 1988, the responsibility for maintaining TAI was transferred from the Bureau international de l'heure (BIH) to the Bureau international des poids et mesures, BIPM.

TAI is in use since the 1970s; however, its formal definition and that of UTC, including the responsibilities of relevant organizations were adopted by the 26th CGPM only in 2018, in its Resolution 2 [Comptes rendus CGPM, 2018]. This definition of TAI is in agreement with the Resolutions A4 (1991) and B1.9 (2000) of the International Astronomical Union (IAU).

The Resolution states that

“TAI is a continuous time scale produced by the BIPM based on the best realizations of the SI second, and is a realization of TT as defined by IAU Resolution B1.9 (2000),

- *in the transformation from the proper time of a clock to TAI, the relativistic rate shift is computed with respect to the conventionally adopted equipotential $W_0 = 62636856.0 \text{ m}^2\text{s}^{-2}$ of the Earth's gravity potential, which conforms to the constant LG defining the rate of TT,*
- *as stated in the IAU Resolution A4 (1991), $TT - TAI = 32.184 \text{ s}$ exactly at 1 January 1977, 0h TAI at the geocentre, in order to ensure continuity of TT with Ephemeris Time,*
- *UTC produced by the BIPM, based on TAI, is the only recommended time scale for international reference and the basis of civil time in most countries,*
- *UTC differs from TAI only by an integral number of seconds as published by the BIPM,*
- *users can derive the rotation angle of the Earth by applying to UTC the observed or predicted values of $UT1 - UTC$, as provided by the IERS,*
- *UTC provides a means to measure time intervals and to disseminate the standard of frequency during intervals in which leap seconds do not occur,*
- *traceability to UTC is obtained through local real-time realizations “UTC(k)” maintained by laboratories contributing data to the calculation of UTC, identified by “k”;*

and confirms that

- “1 *International Atomic Time (TAI) is a continuous time scale produced by the BIPM based on the best realizations of the SI second. TAI is a realization of Terrestrial Time (TT) with the same rate as that of TT, as defined by the IAU Resolution B1.9 (2000),*
- 2 *Coordinated Universal Time (UTC) is a time scale produced by the BIPM with the same rate as TAI, but differing from TAI only by an integral number of seconds...”*

5.1.4 Coordinated Universal Time

Coordinated Universal Time, UTC, was introduced in 1972. It represents a compromise between TAI and UT1, and is defined by the following system of equations valid for any date t :

$$UTC(t) - TAI(t) = n \text{ seconds } (n \text{ integer}) \quad (5-1)$$

and

$$|UTC(t) - UT1(t)| < N \text{ s}$$

The quantities $UTC(t)$ and $TAI(t)$ differ, for any date t , by an integral number of seconds, equal to 37 from 1st January 2017. The International Earth Rotation and Reference Systems Service, IERS, which is responsible for the calculation and publication of UT1 under the form of the values of $UT1 - UTC$, decides on the date of adjustment of seconds based on the predicted divergence between the time scales UT1 and TAI. The adjustments of an interval of N seconds are introduced at the end of a month, preferably at the end of June or December.

By definition, UTC has the same metrological properties as TAI, which is an atomic time. In addition, it follows the rotation of the Earth to within a number of seconds. A maximum value of $N = 0.9 \text{ s}$ was established in 1972, sufficient, at the epoch of its definition, for the purposes of astronomical navigation, where UT1 is

required in real time. At that epoch, the only access to UT1 was via the time signal emissions broadcasting UTC. The situation has changed in the last 20 years, as UTC is disseminated today using new technologies on Earth and from space (see § 5.4, and ITU-R Handbook Satellite Time and Frequency Transfer and Dissemination (2010)). In this new context, the international organizations agreed on enlarging the maximum value of $|UTC - UT1|$ to maintain UTC continuous during centuries [CGPM 2026, Recommendation ITU-R TF.460]

UTC is the general basis for the distribution of time around the world. Local times are derived from UTC by a shift of local offsets of multiples of 15 minutes (which can change from standard time to daylight-saving time), decided by the administration of individual countries, or regional groups. All time signals, at whatever level are synchronized to these local times, and thus to UTC.

The reference time scales TAI and UTC are calculated and distributed by the Time Department of the BIPM. These are deferred-time, paper time scales, without physical representation, and depend on weeks of data collection and analysis to achieve ultimate metrological quality. National time laboratories thus keep other time scales, which are carefully compared *a posteriori* with TAI and UTC at each new issue. These are local time scales, TA(k), and the local representations of UTC, UTC(k), where k designates the acronym of the laboratory.

In 2025, 19 TA(k) were maintained. They were generated from ensembles of clocks carefully kept on one site, such as is the North American A.1(MEAN) from the USNO and the AT1 from the NIST, the German TA(PTB), or at several different sites within the country, such as is done for the French TA(F), or maintained by institutes in different countries such as the Polish TA(PL) which also includes clocks from Lithuania. Some of these atomic time scales are steered to primary frequency standards maintained at the laboratory, as TA(F) and TA(PTB). Others are free-running. They are known through values of time differences with respect to a physical clock which is also kept in the laboratory, and they have no physical representations. Their scientific purpose is to provide a stable local reference.

In 2025, 87 UTC(k) were in operation throughout the world. They are generally realized by the output of a clock, with or without frequency correction, and thus correspond to a physical signal accessible in real time. The UTC(k) are not free-running but are closely steered to follow UTC, with the goal to be within ± 100 ns [CCDS Report, 1993]. The UTC(k) time scales of a number of laboratories are routinely kept within 10s of ns of UTC. The UTC(k) provide real-time synchronization, in particular they are used as references for broadcast time signals.

For most demanding applications, such as millisecond pulsar timing, the BIPM issues atomic time scales retrospectively. These are designated TT(BIPMxx) where xx are the two last digits of the year of computation [Guinot, 1988, Petit, 2003]. The successive versions of TT(BIPMxx) are both updates and revisions: they may differ for common dates. They are computed using data from all available primary and secondary standards having contributed to the computation of TAI.

5.1.5 UTCr: the rapid realization of UTC


Since 1988, UTC has been calculated with one-month data batches at five-day intervals. Extrapolation of values over 10 to 45 days based on prediction models is necessary in many applications. Therefore, UTC is not adapted for real and quasi-real time applications, and it was recognized that a more rapid realization would be of benefit to a variety of applications. For these reasons, the BIPM time department started in 2012 to implement the computation of UTCr, a rapid realization of UTC published every week and based on daily data submitted daily [Petit *et al.*, 2014]. In 2025, 67 laboratories representing 63% of the clock weight in UTC contribute data for the UTCr calculation. The results are published every Wednesday before 18 h UTC on the BIPM web page. The stability of UTCr was expected to be about comparable to that of UTC, albeit slightly worse because the number of participating clocks would necessarily be smaller and because, in general, a deferred solution (here UTC) is expected to be better than a rapid solution (UTCr). In order to achieve a similar performance, it was decided to use the same algorithm (frequency prediction, weighting scheme) as in UTC and to apply it in a similar manner with a calculation interval covering approximately one month. Since UTCr is designed to be a prediction of UTC, an important goal is to minimize the time difference [UTCr – UTC]. For this reason, a steering algorithm has to be implemented. In Fig. 32 the differences between UTC and UTCr are reported up to September 2025. These differences are obtained as a weighted average of the individual

differences between UTC and UTCr for each laboratory k , computed at the date t_j ; the weights W_{kj} are the global weight of the laboratory k in the UTCr calculation at the date t_j :

$$D(t_j) = \sum_{k=1}^N W_{kj} ([UTC_r - UTC(k)](t_j) - [UTC - UTC(k)](t_j)) \quad (5-2)$$

FIGURE 32

Values of UTCr-UTC until September 2025



HDB-31-32

5.2 Time scales in general relativity

5.2.1 Coordinate systems in general relativity

In general relativity, time scales are considered as one of the coordinates of four-dimensional space-time reference systems.

Due to the curvature of space-time, the scale units of these coordinates do not, in general, have a globally constant relation to locally measurable (proper) quantities [Misner *et al.*, 1973, Brumberg 1991, Petit and Wolf, 2005]. In the framework of Newtonian mechanics, it is always possible to define coordinates in such a way that their scale units are everywhere equal to measured distances and durations. This is impossible in general relativity, where the relation between measured quantities and coordinate scale units depends on the position in space-time of the measuring observer. For time scales, this implies that the relation between a coordinate time interval and the locally realized second, using an atomic clock for example, depends on the position of the clock.

In principle, one is free to use any set of coordinates for the description of space-time. However, it turns out that by defining several overlapping systems of coordinates, each valid in a restricted region, the treatment of practical problems and the relationship between coordinates and measurable quantities can be greatly simplified [IAU, 1992]. Such definitions provide several time coordinates, each valid in a particular region of space-time, with the relation between them given by relativistic coordinate transformations.

A system of coordinates in general relativity is defined by its metric tensor $g_{\alpha\beta}(x^\lambda)$ (Greek indices run from 0 to 3) which is position- and time- dependent and must be known for the whole region of space-time within which the coordinate system is used.

The need to define several relativistic systems of space-time coordinates, in particular barycentric and geocentric ones, was recognized by the International Astronomical Union IAU in its 1991 Resolution A4 [IAU, 1991, 1992]. This Resolution includes definitions of barycentric and geocentric coordinate time scales and so provides the theoretical basis for the definition of TAI. The Resolution B1.9 adopted by the IAU in 2000 [IAU, 2001] gives a definition of Terrestrial Time TT different from that of 1991 to avoid any source of uncertainty coming from the definition and realization of the geoid.

5.2.2 The 1991 IAU Resolution A4

The International Astronomical Union approved Resolution A4 at its General Assembly held in Buenos Aires in August 1991. The complete text of this Resolution is in the 1992 IAU Information Bulletin 67. It contains several recommendations of importance for the definition and realization of coordinate time scales. They are explained in the following sections.

5.2.2.1 Recommendation I

Recommendation I explicitly refers to the general theory of relativity to provide the theoretical background for the definition of space-time reference frames. It specifies the form of the metric to be used for coordinate systems centred at the barycentre of an ensemble of masses:

$$ds^2 = -c^2 d\tau^2 = g_{\alpha\beta}(x^\lambda) dx^\alpha dx^\beta \quad (5-3)$$

$$ds^2 = -\left(1 - 2\frac{U}{c^2}\right) (dx^0)^2 + \left(1 + 2\frac{U}{c^2}\right) [(dx^1)^2 + (dx^2)^2 + (dx^3)^2] \quad (5-4)$$

where ds is an infinitesimal space-time line element, τ is proper time as realized by an ideal clock, c is the velocity of light in vacuum, U is the sum of the gravitational potentials of the ensemble of masses and of a tidal potential generated by masses external to the ensemble, the latter potential vanishing at the barycentre. The four space-time coordinates are defined to be $(x^0 = ct, x^1, x^2, x^3)$. The Einstein summation convention is used, implying summation over repeated indices. It should be noted that equation (5-4) gives only the first terms of a series, which is sufficient for the present level of observational accuracy. Higher order terms may be added as necessary. For time and frequency applications, this will be the case when clock stabilities reach some parts in 10^{19} .

5.2.2.2 Recommendation II

Recommendation II states that the space coordinate grid with its origin at the centre of mass of the Earth should show no global rotation with respect to a set of distant extragalactic objects, that the time coordinates for all coordinate systems should be derived from a time scale realized by atomic clocks operating on the Earth, and that the basic physical units of space-time are the second of the International System of units (SI) for proper time and the SI metre for proper length. This Recommendation should also apply to clocks on board terrestrial satellites.

5.2.2.3 Recommendation III

Recommendation III defines the scale units and origins of all time coordinates and designates the solar system barycentric time coordinate and the geocentric time coordinate as Barycentric Coordinate Time, TCB, and Geocentric Coordinate Time, TCG, respectively. It should be noted that these time coordinates exhibit secular differences between themselves and with respect to TAI.

5.2.2.4 Recommendation IV

Recommendation IV defines Terrestrial Time (TT) as a geocentric coordinate time scale differing from TCG by a constant rate, the scale unit of TT being chosen so that it agrees with the SI second on the rotating geoid. This constant rate is presently estimated to $6,9692904 \times 10^{-10}$ with an uncertainty of 1×10^{-17} (1σ).

As the theoretical time scales TCB, TCG and TT are completely defined by the IAU Resolution A4, the path to realized time scales is immediate.

5.2.3 The 2000 IAU Resolution B1.9

The XXIVth General Assembly of the IAU, in its meeting in Manchester in August 2000 adopted the Resolution B1.9, that recommends redefining Terrestrial Time (TT) as a time scale differing from TCG by a constant rate $dTT/dTCG = 1 - L_G$, where L_G is a defining constant. Its numerical value was chosen to conform to the value for the gravity potential on the geoid recommended by the International Association of Geodesy (IAG) in 1999. Recommendation IV of 1991 Resolution IAU A4 defined L_G equal to U_G/c^2 , where U_G is the geopotential at the geoid.

5.2.4 International Atomic Time

According to its definition, International Atomic Time (TAI) is a realization of TT, apart from an offset of 32.184 s introduced for historical reasons to maintain continuity with the previous astronomical definition. It is obtained by combining the data from an ensemble of about four hundred atomic clocks spread world-wide. To achieve this, it is necessary to compare those clocks using a convention of coordinate synchronization [Allan and Ashby, 1986]. This is defined as follows [Klioner, 1992]:

Two events fixed in some reference system by the values of their coordinates (t_1, x_1, y_1, z_1) and (t_2, x_2, y_2, z_2) are considered to be simultaneous with respect to this reference system if the values of the corresponding time coordinates are equal: $t_1 = t_2$. Two clocks are considered to be synchronized with respect to some reference frame if they simultaneously (in the above sense) produce the same time markers.

In the vicinity of the Earth a geocentric, non-rotating reference frame, as defined in Resolution A4, Recommendation II of the IAU, is used for the synchronization of clocks, and in particular for the calculation and dissemination of TAI.

TAI is based on the best realizations of the SI second [CGPM, 2018]. To do this, data from the most accurate primary standards and secondary realizations of the second are individually corrected for the gravitational frequency shift arising from the elevation of the laboratory above the geoid and are then combined to form the scale unit of TAI.

5.2.5 Other time scales

Atomic coordinate time scales like TA(k), UTC, UTCr, UTC(k) and TT(BIPMxx) are time coordinates closely related to TAI and are provided for different purposes. The scales TCG and TCB are related to TT and hence TAI by relativistic transformations (see notes to Recommendations III and IV of the IAU Resolution A4, IAU 1992, and Resolution B1.9 [IAU, 2001]). Up to now international reference time scales and their national representations are introduced. In general, time scales are established in different applications and in institutions of quite different kinds. Even if the use of UTC is recommended in ITU-R (Recommendations ITU-R TF.460 and ITU-R TF.486), for practical reasons operators of infrastructure systems generate local signals for distribution based on their equipment (this aspect is addressed further in Chapter 8).

Atomic time scales have been developed for the internal support of satellite navigation systems. The four operational GNSS maintain internal independent time scales: GPS Time, for the American Global Positioning System GPS, GLONASS Time for the Russian GLObal Navigation Satellite System GLONASS, Galileo System Time GST for the European navigation system and BeiDou System Time BDT for the Chinese navigation system. They are generated with a high update rate (of the order of several minutes) from an ensemble of clocks kept in the master control station of the system and, in case of GPS, in onboard satellites, and steered on a local representation of UTC: UTC(USNO) for GPS Time, UTC(SU) for GLONASS Time, UTC(NTSC) for BeiDou System Time, and to a prediction of UTC based on contributions of some European

laboratories maintaining local representations of UTC for Galileo System Time. For operational needs, the time scales are kept close to UTC (modulo 1 s), often to better than 10 ns. These system times include in most cases an integral offset with respect to UTC. This offset avoids discontinuities in the system time that would otherwise occur at every leap second insertion. Values of these offsets as well as of the differences between UTC and the predictions of UTC(k) broadcast by GNSS are available at the BIPM website (www.bipm.org), BIPM Time Department Database, only for GPS and Glonass system times) until end of April 2024, and for the four constellations (GPS, Glonass, Galileo and BeiDou) since May 2024.

5.3 Generation of time scales

The practical problem is to generate a time scale from an ensemble of atomic clocks maintained in one or several laboratories. The efficient combination of the readings of the participating clocks requires:

- the definition of the expected qualities of the time scale,
- the characterization of the available timing data,
- the design of an algorithm for data treatment.

5.3.1 Expected qualities

The departure from an accumulation of ideal SI seconds on the rotating geoid can be estimated through computation of its “normalised frequency departure” at date t , commonly referred to as “frequency”, defined as:

$$y(t) = \frac{v(t) - v_0}{v_0} \quad (5-5)$$

where $v_0 = 1$ Hz, and $v(t)$ is the reciprocal of the scale unit, expressed in SI second, of the time scale for date t .

Actual physical clocks have offsets with respect to the “ideal” time scale (in fact to a reference time scale). These offsets can be characterized as the combination of a static (or very slowly varying) deterministic set of parameters and stochastic contributions that can only be characterized statistically. The deterministic parameters of each clock are used to correct the data at every measurement epoch. The stochastic contributions of each clock are assumed to be uncorrelated and randomly distributed and are minimized by combining their data to obtain a reliable, stable and accurate time scale. A separate, but important, point is the delay of access to the scale. For some purposes access must be immediate, for others a significant delay may be tolerable.

5.3.1.1 Reliability

Individual physical clocks may fail with immediate interruption of the time scales they deliver. Reliability thus calls for redundancy and eventually for national or international collaboration between laboratories maintaining atomic clocks.

The simplest solution to this problem is to replace the clock which has failed by another one. This is what is normally done in time laboratories which generate a UTC(k). Usually the UTC(k) is directly linked to the output of a physical clock, generally the best one of the ensemble on site, this being designated the “master clock” by some laboratories. Its output is controlled by small, predetermined frequency corrections and time steps, often through a microphase stepper, so as to steer UTC(k) to UTC. A change of the physical clock does not influence the output time scale if the microphase stepper is suitably programmed to handle the change.

More often, reliability is ensured by using an ensemble of clocks and computing an ensemble time. This time rarely has a physical realization. In computing such an ensemble time, it is necessary to minimize the perturbations that result as clocks enter and leave the ensemble. Clearly the more clocks at the disposal of an ensemble, the less is the effect as one clock enters or leaves. For this reason, there has been a general increase in the number of clocks in a given ensemble. The commercialization of the 5071A clock successively by the firms Hewlett-Packard, Symmetricom and Microsemi, has led to a steady increase in the number of industrial caesium clocks, together with the incorporation of hydrogen-maser clocks to the individual time scales. This results in about 410 atomic clocks contributing to TAI in 2025. One of the major consequences is a gain in reliability for TAI.

5.3.1.2 Stability

The stability of a time scale may be defined as its ability to maintain a constant scale interval, even if it differs from the ideal one. A measure of stability thus consists in the estimation of the dispersion of the frequency values $y(t)$ with time. Some statistical tools have been developed to estimate stability; they are developed in Chapter 2. They are efficient for the characterization of the usual types of random noise which affect clock signals. The most common stability estimation tool is the two-sample, or Allan variance $\sigma_y^2(\tau)$ which depends upon the observational, or sampling, time τ .

The stability of an ensemble time scale depends upon the stabilities of the contributing clocks and the design of the algorithm used to generate it. The algorithm must, in particular, correctly handle any change in clock behaviour. The general considerations are detailed in § 5.3.3, but the central idea is to generate a time scale more stable than any of the contributing clocks. This can be realized, but generally only for a given range of averaging times τ .

In principle, the concept of stability applies only to free-running time scales. A UTC(k) is, by definition, steered so it is affected by intentional frequency steps and its short- and middle-term stability generally degraded. In addition, a crucial problem comes from the fact that time scale frequency values are always estimated or measured with respect to the frequency of another time scale or physical clock. The stability analysis of such comparative measures leads to the evaluation of the coupled stability of the two time scales. The statistics of the measurement process that compares the two time scales may also be important, and, in some cases, this effect may make the largest contribution to the statistics of the resulting estimate. Two cases may arise:

- The frequency of the time scale under test is evaluated by comparison with a time scale of significantly better quality, such as those realized by primary frequency standards. The observed instability then may be ascribed totally to the time scale under test.
- The two time scales which are compared are of similar quality. A technique for noise decoupling is then necessary. If the time scales involved in the comparison may be assumed to be completely independent, the N -cornered-hat technique produces an estimate of the intrinsic stability of each element. If the independence is not verified, variances and covariances should be handled together for a complete analysis [Tavella and Premoli, 1994]. Complete independence of the two time scales is often difficult to realize in practice.

5.3.1.3 Accuracy

The accuracy of a time scale may be defined as its ability to maintain a mean scale interval as close as possible to its definition. For time scales which realize TT, the mean scale interval should be as close as possible to the SI second on the rotating geoid.

For primary and secondary frequency standards, accuracy is estimated by an uncertainty budget obtained through the evaluation of the known physical effects that modify the output frequency with respect to the definition. When it is not possible to constitute such an uncertainty budget, accuracy is evaluated by comparison of the duration of the scale interval with the best realization of the SI second provided by primary frequency standards. It is of course necessary to compensate for the effect of the gravitational red shift on the primary standard frequency results in order to convert its realized SI second to the geoid (null height). The accuracy of a time scale is generally given by a frequency difference between the time scale and the primary frequency standard, evaluated for averaging times corresponding to the best stability of the time scale, including the uncertainty of the primary frequency standard.

Improvement in the accuracy of a time scale is generally carried out outside the main algorithm, which deals only with stability optimization. This may be done by steering the frequency of the time scale to the frequency of a primary standard or to a more stable reference time scale. The method that is used to realize this steering must be chosen to minimize the degradation of the short-term stability of the time scale.

5.3.1.4 Delay of access

The delay of access to a time scale is linked to the quality of raw timing data and to the scientific purposes the time scale is supposed to fulfil.

Raw timing data is acquired according to a basic measurement cycle, whose duration ranges from several minutes to several hours, and is affected by measurement noise. Depending on the level of this noise, it may be useful to smooth raw measurements by accumulating data over several successive basic samples of measurement (see § 5.3.2). This delays the access to the resulting time scale. In addition, it is useful to observe the behaviour of contributing clocks for a long period, both before and after the moment to which the data applies, in order to make the best use of their data. This also delays the access.

What constitutes an acceptable delay of access to a time scale depends on its use. For a reference time scale, such as UTC, the requirement is extreme reliability and long-term stability. To match this purpose, the reference time scale relies on a large number of clocks of different types, located in different parts of the world. Data must therefore be collected and handled correctly, which takes time. The delay is thus considerable but is acceptable because of the ultimate quality obtained. For scientific studies inside a laboratory, however, it may be necessary to produce the time scale in near real time, immediately after the clock measurements, even if this may compromise the long-term qualities of the scale.

5.3.2 Timing data

5.3.2.1 General form of timing data

Timing data takes the form of time differences between clocks. An atomic clock delivers a series of electrical pulses separated from one another by a very stable time interval. An interval between pulses of 1 s is common but is not a requirement and other intervals are used. Each pulse is an event with an associated number, a kind of label which is tied to it. This associated number is the reading of the clock for that particular event: for example, it may read as 2022 March 13 11 h 27 min 13.0 s. It can also be designated as the date of the event; its origin is arbitrary and is chosen to be convenient, but it is incremented by a constant value at each new pulse. Clock readings vary continuously and rapidly, so they can only be “caught in flight”. However, counters are available: they can be started with a given pulse coming from one clock and stopped with the pulse with the same label coming from another clock. A counter thus measures time differences which are proper quantities. These are thus measurable and expressed in SI units.

Denote $h_i(t)$ and $h_j(t)$ the time coordinates of the pulse labelled t , delivered by clock H_i , and of the pulse, with the same label, delivered by clock H_j , in a given reference frame. The interval of coordinate times:

$$x_{ij}(t) = h_j(t) - h_i(t) \quad (5-6)$$

is needed for the generation and dissemination of coordinate time scales.

At the present level of accuracy in clock comparisons, the coordinate quantity $x_{ij}(t)$ can be approximated using the interval of proper time issued from a counter, taking into account signal propagation delays for clocks separated by large distances [Petit and Wolf, 1994]; it thus can be expressed in SI units. In addition, in current practice, one does not specify a reference frame, and one designates the time coordinate $h_i(t)$ as the “reading of clock H_i at date t ”, which, strictly, is not correct. For sake of conformity with the existing literature, the same designation is used in this text. However, the actual meaning of equation (5-6), which involves only coordinate quantities, should not be forgotten.


The quantities $x_{ij}(t)$ are the basic measurements used for time scale generation. They are obtained by time transfer methods, implemented between clocks located on the same site or in remote locations. Generally, a non-redundant network of time links is used, a given clock being compared only once with all the others at each date.

5.3.2.2 Comparison of clocks located on the same site

For computation of some of the TA(k) kept by national timing centres, all the contributing clocks are located on the same site. This is the case for the NIST, the SU, and the USNO. On each site, one clock is designated as the reference clock for the ensemble, and the times of the other clocks are measured with respect to it. It

FIGURE 34

Comparison between clocks located on two different sites



HDB-31-34

There exist several methods for making distant time comparisons. The introduction of time transfer methods using satellite systems led to a major improvement in the precision, accuracy, and coverage of world-wide time metrology. They have accuracies of order several ns (1σ) or even better for calibrated equipment. Clock comparisons are possible with sub-nanosecond accuracy with optic fibre links. As of 2025 three methods are fully operational:

- 1 Comparisons using satellites of GNSS, routinely by the US Global Positioning System, GPS, but also using signals from the Russian Glonass, the European Galileo and the Chinese BeiDou. [For information on these comparisons refer to Chapter 6, § 6.2.6 GNSS Signals in High-Accuracy Applications]. The ITU-R Handbook on Satellite Time and Frequency Dissemination [ITU Handbook 2010] provides a comprehensive description of the GPS and Glonass satellite systems (Chapters 4 and 5 respectively), as well as the principles of time transfer via GPS (Chapters 12 and 13).

To facilitate the exchange of data among timing laboratories and between these laboratories and the BIPM, the Common GNSS Generic Time Transfer Standard (CGGTTS) was defined in 1994 [Allan and Thomas, 1994] for the use of GPS measurements. The standard recognized the capabilities of the receiving equipment of that epoch. This standard has been updated to follow the evolution of GNSS, the outcome of new constellations and the development of new types of receivers [Defraigne & Petit, 2015]. The CGGTTS computation procedure is based on the analysis of measured pseudoranges for 13-minute tracks, the interval that was originally necessary for a single-channel time receiver to get a full navigation message. Nowadays, these files can be produced directly by the receiver. But they can also be constructed from the RINEX (Receiver Independent Exchange) format files (for a detailed description refer to the IGS website igs.org, IGS formats and standards). Solutions with sub-

nanosecond precision can be achieved with Precise Point Positioning (PPP) solutions from RINEX data [Petit and Jiang, 2008].

- 2 Two-way comparisons using telecommunication satellites (Two-way Satellite Time and Frequency Transfer, TWSTFT) are routinely used for clock comparison in the computation of UTC at about ten institutes. Although the TWSTFT solution is easy and is based on geometric considerations, the cost of the equipment and its rather complex operation are factors that limit the expansion of the TWSTFT network. The ITU-R Handbook on Satellite Time and Frequency Dissemination [ITU Handbook 2010], (Chapter 14), and [Recommendation ITU-R TF.1153] provide a comprehensive description of the TWSTFT technique and its use for time transfer. Significant improvement of the 1-day stability of the time comparison can be achieved with the operation of the newly developed Software Defined Radio (SDR) receivers (see § 4.2.1 of this Handbook and Report ITU-R SM.2152), replacing the reception part of the time-transfer hardware with high-speed digitization and software data processing.
- 3 Time comparisons over optical fibre provides a higher performance alternative to satellite-based time transfer, with demonstrated instability below 100 ps and sub-ns accuracy. Comparisons via optical fibre are possible across continental distances. Intercontinental comparisons including optical fibre stretches could be possible implementing combination with space techniques. Calibration with pico-second accuracy is possible (see Chapter 6 for a detailed description).

The characterization of the delays in the time transfer equipment, usually called “calibration” is essential to the accuracy of time transfer and time dissemination. The current scheme for obtaining such “receiver calibrations” includes so-called Group 1 laboratories (G1) in each Regional Metrology Organization (RMO), which are periodically visited by BIPM equipment. These G1 laboratories operate several reference stations and are responsible for the calibration of the other laboratories of their RMO. UTC links between laboratories calibrated with this scheme receive a calibration uncertainty of 1.5 ns for G-G1 links and 2.5 ns for G1-G2 links. In 2025, 80% of the stations participating in UTC had been calibrated under this scheme. Some of the remaining stations had received a calibration of their equipment by the manufacturer before installation, and in such cases an uncertainty not lower than 5 ns is assigned in the computation of UTC. Finally, there are still some time laboratories that have never been calibrated. For these stations an “uncertainty” of 20 ns is currently assigned in the BIPM Circular T, but this should in the future be replaced by “not calibrated”. Guidelines for GNSS equipment calibration issued by the BIPM are available at the BIPM website (bipm.org, BIPM Time Department Database, Characterization of time transfer equipment). Results of calibrations of equipment for reception of GPS and Galileo signals can be accessed also at the BIPM Time Department Data Base.

The calibration of TWSTFT is different, it consists in the calibration of the link connecting the two stations involved in a comparison. A TWSTFT link calibration campaign is carried out using a TWSTFT mobile station or eventually a GNSS travelling system when a mobile station is not available. The TWSTFT calibration guidelines prepared by the CCTF Working Group on TWSTFT as well as the status of TW calibration exercises are available at the BIPM Time Department Data Base.

5.3.3 Time scale algorithms

5.3.3.1 Definition of an average time scale

To match the definitions of time scales given in the introduction and in § 5.2, the reading of an atomic time scale TA may be theoretically written, at date t , as the weighted average of the readings of the contributing clocks:

$$TA(t) = \sum_{i=1}^N \omega_i(t) h_i(t) \quad (5-9)$$

(Since the reading of a clock includes both deterministic and stochastic contributions, the weighted average defined in the preceding equation must be computed with care. The weighted average of the deterministic contributions of an ensemble of clocks is unlikely to be zero. See discussion below.) The basic mathematical definition given in equation (5-9) plays a fundamental role in the development of time scale algorithms, which is detailed in the following. Some ensemble time scales, such as GPS time, do not make use of a similar explicit average definition but use a Kalman filter approach, which is described in § 5.4.2.1.

Relative weights $\omega_i(t)$, $i = 1, \dots, N$, are introduced in order to discriminate between clocks according to their errors in predicting the ensemble time. They satisfy the following relation:

$$\sum_{i=1}^N \omega_i(t) = 1 \quad (5-10)$$

The time of a clock in general differs from the ensemble-average time scale because of both systematic deviations as well as random deviations. The weighting function in equation (5-9) is typically chosen to optimise stability and is not dependent on the systematic deviations (e.g. frequency offset, frequency drift) but only on the random deviations where is considered a deviation systematic if its variation is very slow relative to the measurement interval. Hence, if the weight of a clock is changed or if a clock is added or subtracted, the resulting time scale computation will be adversely affected, since the weighting function in equation (5-9) apply to systematic deviations as well [Guinot, 1987]. It follows that equation (5-9) is not satisfactory for actual clock ensembles, for which changes of weights are unavoidable. Equation (5-9) must thus be modified as follows:

$$TA(t) = \sum_{i=1}^N \omega_i(t) [h_i(t) + h'_i(t)] \quad (5-11)$$

where $h'_i(t)$ is a time correction added at date t to the reading of clock H_i , and designed to ensure time and frequency continuity of TA at a previous date t_0 when weights have been changed. The correction can be written by using the evolution of the time of the clock with respect to the ensemble (the “frequency” of the clock), and the evolution of this frequency:

$$h'_i(t) = x_i(t_0) + y_{ip}(t) \times (t - t_0) + \frac{1}{2} \times d_{ip}(t) \times (t - t_0)^2 \quad (5-12)$$

where $x_i(t_0)$ is known, since it results from the computation of TA at date t_0 , where $y_{ip}(t)$ is the predicted frequency of clock H_i , relative to TA, over the interval $[t_0, t]$ and $d_{ip}(t)$ is the predicted frequency drift of the clock H_i over the interval $[t_0, t]$. The frequency $y_i(t)$ of clock h_i , relative to TA, over the interval $[t_0, t]$ can be estimated from:

$$y_i(t) = \frac{[TA(t) - h_i(t)] - [TA(t_0) - h_i(t_0)]}{t - t_0} \quad (5-13)$$

The frequency estimated in equation (5-12) often has a contribution due to white frequency noise, so that the frequency used in the prediction term in equation (5-12) is often the average of the estimate obtained in equation (5-12). The averaging time is generally different for different clocks. It is determined by the domain over which the fluctuations in the value obtain from equation (5-12) are dominated by white frequency noise. The AT1 algorithm that is used at NIST includes an exponential filter that is used for this purpose. [Levine and Parker, 2002] The frequency drift of the clock h_i over the interval $[t_0, t]$ can be estimated by using the least square approach to the frequency values of the clock H_i with respect to a reference that in the case of UTC is TT(BIPM) [Azoubib *et al.*, 1977, Petit, 2003]. This method assumes that the frequency variation of the clock is dominated by white frequency noise and linear frequency drift. The frequency drift of many masers in some cases does not satisfy this requirement.

It is also possible to characterize the evolution of the time of a clock with respect to the ensemble by using auto-regressive methods (ARMA and ARIMA). The current reading of a clock is computed based on a linear combination of previous data. The optimum linear combination may have no clear connection to the more physical frequency and frequency drift parameters of the previous equations. These techniques were studied and discarded some time ago but have had a more recent renaissance because they may be more suitable for time scales based on the artificial intelligence.

Until TA is computed at date t , $y_i(t)$ is unknown. It is thus necessary to predict it according to the past behaviour of clock H_i . This predicted frequency is denoted $y_{ip}(t)$ and the frequency drift is denoted by and appear in equation (5-12).

Equations (5-6), (5-10) and (5-11) lead to the following system of equations, assuming no measurement noise:

$$\begin{cases} \sum_{i=1}^N \omega_i(t) \times x_i(t) = \sum_{i=1}^N \omega_i(t) \times x_i(t_0) + \sum_{i=1}^N \omega_i(t) \times y_{ip}(t) \times (t - t_0) + \\ \quad + \sum_{i=1}^N \omega_i(t) \times \frac{1}{2} \times d_{ip}(t) \times (t - t_0)^2 \\ x_i(t) - x_j(t) = x_{ij}(t) \end{cases} \quad (5-14)$$

System in equation (5-14) is deterministic with N equations and N unknowns. The solution is unique and the results are the time differences $x_i(t)$, $i = 1, \dots, N$, which give access to TA for date t . The difference between clock H_j and TA is explicitly given by:

$$x_j(t) = \sum_{i=1}^N w_i(t) \times [h'_i(t) - x_{ij}(t)] \quad (5-15)$$

System (5.14) can be found in time laboratories using more clocks to obtain the reference as for example TAI [Panfilo *et al.*, 2019] and UTCr [Petit *et al.*, 2014] at BIPM. The assumption that the measurement noise is negligible relative to the clock noise may have to be re-considered for hydrogen masers and other stable clocks when data from these clocks is used in constructing the UTCr time scale.

Since the definition of the time scale, and thus the resulting system of equations, nearly always takes the same form, the specificity of a given algorithm lies in the choices made for:


- the length of the time interval $[t_0, t]$,
- the weights attributed to clocks,
- the way clock frequencies are predicted, and the way frequency drift is dealt with,
- the method used to detect and remove outliers and other errors without degrading the scale computation.

These choices are closely related to the purposes for which the time scale is designed.

The decreasing dead time of fountains measurements makes possible other time scale algorithms. The fountain is used to steer a clock (usually a H-Maser) that provides UTC(k). A steering algorithm based on the information given by the fountain inside the laboratory and by TAI (monthly) is applied to a H-maser to provide UTC(k). Examples are given by UTC(PTB) [Bauch *et al.*, 2012], UTC(OP) [Rovera *et al.*, 2016], UTC(NPL), UTC(USNO), UTC(SU) etc. These time scales show the best performance in being close to UTC. As it is reported in Fig. 35 the difference between UTC - UTC(k) for the laboratories listed remains inside the interval of $[-4 \text{ ns}, 4 \text{ ns}]$ for more than three years.

In these cases the steering procedures used are the key issue of the algorithm implemented; the steering applied to the clock used as reference to maintain its frequency as close as possible to the PFS frequency and the steering applied to maintain the UTC(k) realization close to UTC [Bauch *et al.*, 2012; Rovera *et al.*, 2016].

FIGURE 35
Values of UTC – UTC(k) for six laboratories contributing data to the BIPM



HDB-31-35

Because of the progressive development of Secondary Frequency Standards and in view of the possible redefinition of the second [Dimarcq *et al.*, 2024] the future development of time scale algorithm will be probably based on the use of optical clocks. Some publications are already available, and the results show very good performance [Yao *et al.*, 2018, Yao *et al.*, 2019, Milner *et al.*, 2019 and Hachisu *et al.*, 2018]. Optical clocks do not run continuously at present, so that a time scale that incorporates data from an optical clock usually also has conventional microwave-based devices that act as a steered flywheel, similar as in a fountain-based time scale. This combination is likely to be common for some time.

5.3.3.2 Length of the basic interval of computation

In some laboratories, TA(lab) is a result of a measurement of the time differences of an ensemble of clocks. In this configuration then two basic durations have been previously defined:

- τ_0 , duration of the basic measurement cycle.
- T_0 , minimum duration over which raw data should be averaged in order to smooth out the measurement noise sufficiently to reach the intrinsic qualities of the clocks being compared. Orders of magnitude for T_0 are several minutes to several hours inside a laboratory, 12 hours to 1 day between two laboratories linked via short-distance GPS common views, and several days between two laboratories linked via long-distance GPS common views.

Also defined are two dates:

- t_0 , date for which TA is known.

- t , a following date ($t > t_0$) for which smoothed timing measurements are available and for which TA is to be computed by solving equation (5-14).

The update interval $T = t - t_0$ is, in general, of the same order of magnitude and slightly longer than T_0 . Its length is thus directly linked to the quality of timing data. It is for example:

- $T = \text{few hours}$ for UTC(k), which only uses timing data taken on site,
- $T = 1 \text{ day}$ for UTCr, which uses timing data from all over the world, the maximum baseline between laboratories being of order 6 000 km,
- $T = 5 \text{ days}$ for TAI, which uses timing data from all over the world, the maximum baseline between laboratories being of order 6 000 km.

Another requirement is efficient characterisation of the behaviour of the participating clocks in order to weight them correctly and efficiently predict their frequencies relative to TA (see the following sections). It is thus often necessary to observe clocks over a duration longer than T . One then has two possibilities as follows.

Consider an integer n greater than 1.

A memory of the last n intervals of duration T is retained. The time scale is delivered in near real-time, with a delay no greater than T , but it is based only on the past behaviour of contributing clocks. There are no re-processing and no post-processing. The weights and the frequency prediction are valid for an interval of duration T . The resulting algorithm is thus dynamic and adaptive at intervals of T .

The advantage of this approach is that the time scale is accessible in real time. The disadvantage is that it may not be possible to take into account the abnormal behaviour of one clock before its data are included in the scale. A stable clock which suddenly presents a frequency jump can thus affect the time scale before the anomaly is detected.

The $(n + 1)$ dates included in the interval are treated as a whole. This delivers a deferred-time time scale, computed in post-processing. The weight and the frequency prediction of a given clock are valid for an interval of duration nT . They are changed for the following interval of duration nT , but are equal for all dates included in a given interval of duration nT . The clock behaviour observed during the whole interval of computation is taken into account. The resulting algorithm is dynamic and adaptive a posteriori at intervals of duration nT .

The advantage is the possibility of taking into account any abnormal behaviour of clocks occurring during this period. The disadvantage is the access to the time scale in deferred time for the $(n + 1)$ dates included in the interval of computation.

TAI is computed by the same process with $T = 5 \text{ days}$ and $nT = 30 \text{ days}$ ($n = 6$), for UTCr $T = 1 \text{ day}$ and $nT = 7 \text{ days}$.

In the case of national time scales UTC(k) making use of fountains in their algorithms the steering procedure of the frequency with respect to a fountain called f_{fou} is updated each day and the other f_{TAI} is updated each month after the publication of Circular T .

Most of the algorithms used in national laboratories adopt this choice, updating TA(k) in real time or in near real-time without post-processing. In addition, some algorithms, as those for UTC and UTCr at the BIPM, choose an iterative procedure to evaluate weights and frequency predictions: this takes the form of successive recomputations of TA for the same interval, with detection of outliers at each step, until results converge [Panfilo and Arias, 2019; Petit *et al.*, 2014].

5.3.3.3 Weighting procedure

5.3.3.3.1 General ideas

Since time scale algorithms are designed to optimise frequency stability, each clock should be weighted according to its own frequency stability. The weight attributed to a given clock is thus basically chosen to be inversely proportional to its frequency variance σ_i^2 .

$$\omega_i = \frac{1/\sigma_i^2}{\sum_{k=1}^N 1/\sigma_k^2}, \quad i = 1, \dots, N \quad (5-16)$$

The reason for this is that, if the contributing clocks are independent and if weights are not artificially limited, the frequency variance of the resulting time scale may be written as:

$$\frac{1}{\sigma_{TA}^2} = \sum_{i=1}^N \frac{1}{\sigma_i^2} \quad (5-17)$$

which means that the time scale is, in principle, more stable than any contributing element. The choice of the variance type (classical, filtered, or Allan) depends on the purposes for which the time scale is generated, and can thus differ according to the algorithm which is considered. However, there are two limiting factors as follows.

The frequencies of clock H_i , used for the computation of its frequency variance, are estimated over an interval of duration τ . According to equation (5-17), the stability of the resulting time scale is optimised for averaging times close to τ . It is thus of pre-eminent importance to determine for which τ values the contributing clocks present their best stabilities, and to define what objective of stability the time scale should fulfil. In other words, the optimisation of both short-term and long-term stability could call for contributions from different types of clocks, treated according to different procedures in the algorithm.

The frequencies of clock H_i , used for the computation of its frequency variance, are estimated by comparison with a reference. Very often this reference is the time scale itself, because its stability is assumed to be better than that of the contributing clocks.

In addition to the fundamental aspects that have just been discussed, the weighting procedure should obey some other rules. The most important is the implementation of an upper limit of weight, necessary practically to make the time scale rely on the best clocks and yet avoid giving a predominant role to any one of them. Another is an objective criterion to safeguard the time scale against the possible abnormal behaviour of some clocks. It is important to stress that the existence of an upper limit of weight improves the reliability but invalidates equation (5-17). It leads to a time scale that is not as stable as a scale constructed with weights determined purely from statistical considerations.

To illustrate, the examples of the algorithms used for UTC and UTCr are used.

5.3.3.3.2 Weighting procedure in UTC

As already noted, UTC operates in post-processing, treating as a whole measurements sequence taken over a basic period $nT = 30$ days. Measurements are available every $T = 5$ days, on MJDs ending with 4 and 9. The time scale is updated for each of the six dates t included in the one-month period under consideration: $t = t_0 + mT$, with $m = 1, 2, 3, 4, 5, 6$. The date t_0 is the last date of the previous calculation interval, for which the time scale is maintained and not updated. The separation between updates is thus 5 days, but the gap between calculations is 30 days.

In UTC, the weight $w_i(t)$ of clock H_i is constant over the calculation interval I (30/35 days) of computation: it is thus valid for the seven dates $t = t_0 + mT$, with $m = 0, 1, 2, 3, 4, 5, 6$, continuity at t_0 being ensured by the clock frequency prediction. It can be written $w_i(I)$ and obeys equation (5-16), where σ_i^2 is calculated as:

$$\sigma_i^2 = \frac{\sum_{j=1}^M \left(\frac{M+1-j}{M} \right) e_{i,j}^2}{\sum_{j=1}^M \left(\frac{M+1-j}{M} \right)} \quad (5-18)$$

where $|y_{i,j} - \hat{y}_{i,j}|^2$ is the difference between the real frequency and the predicted one, i identifies the clock, j the calculation interval and M the number of available measurements which can vary from 5 to 12 given that 5 is the minimum number of months requested to observe the behaviour of a clock prior to its introduction into the UTC calculation and one year (the standard period of observation) of the frequency data of the clock H_i results from an exponential filter used to de-weight the past behaviour of the clocks. These are the frequencies computed over interval I and the five previous two-month intervals. As the frequency over the interval I is not yet known, an iterative process is used [Panfilo *et al.*, 2014] which begins with the weights obtained at the previous month computation, ending at date t_0 ; this gives an indication of the behaviour of each clock during the interval I and thus makes it possible to refine weights in the following iterations.

The UTC weight determination uses clock measurements covering a full year, so annual frequency variations and long-term drifts can lead to de-weighting. In addition, the choice of 30 days, corresponds to a good

averaging time for the detection of frequency anomalies. Sampling over 30 days thus allows optimisation of TAI stability in the long-term.

In UTC there is an upper limit of weights which corresponds to an A/N with $A = 6$ and N the number of clocks participating in UTC calculation.

To conclude, the weights obey the same rules, in particular: optimisation of the stability and detection of abnormal behaviour. The specific choices that have been made match the available timing data and fulfil the fundamental requirement of access to a time scale in real time or deferred time.

5.3.3.3 Weighting procedure in UTCr

The weighting algorithm in UTCr is harmonized to the UTC weighting algorithm. The same algorithm used in UTC has been implemented in UTCr. The maximum weight of a clock is set at $3.8/N$; where N is the number of clocks with a non-null a priori weight; a test for ‘abnormal behaviour’ is implemented, by checking the difference between the predicted and the real frequency of the clocks.

Other laboratories use weight limits. For example, NIST limits the weights of any clock to 30% independent of N .

5.3.3.4 Frequency prediction

5.3.3.4.1 General ideas

The way the frequency of clock H_i is predicted depends on its statistical characteristics and on the duration for which the prediction should be valid. There are several pure cases:

- The predominant noise is modelled as white frequency noise: this is the case for commercial caesium clocks for averaging times τ ranging from 1 day to 10 days. The most probable frequency, estimated over an interval of duration τ , for the following τ interval is then given by the mean of the frequency values observed over a number of previous intervals of duration τ .
- The predominant noise is modelled as random walk frequency modulation: this is the case for most commercial caesium clocks for averaging times τ ranging from 20 to 70 days. The most probable frequency value for the following τ interval is then the last frequency value estimated over the previous interval of duration τ .
- The predominant frequency deviation is modelled as a linear drift: this is the case for some hydrogen masers for averaging times τ longer than several days. The most probable frequency for the following τ interval is then the last frequency calculated over the previous interval of duration τ corrected by a term deduced from the estimated frequency drift.

To optimise an ensemble, it is thus necessary to have a good knowledge of the behaviour of the contributing clocks, and to be astute in selecting suitable modes of frequency prediction for the different clock types.

To illustrate, the algorithm used in UTC and UTCr is considered.

5.3.3.4.2 Frequency prediction in UTC and UTCr

As already noted, UTC is calculated in post-processing, treating as a whole measurements sequence taken over a basic period $nT = 30$ days. As with its weight, the predicted frequency of clock H_i is constant over one month interval I of computation; it is thus valid for the seven dates $t = t_0 + mT$, with $m = 0, 1, 2, 3, 4, 5, 6$, and can be written $y_{ip}(I)$.

In UTC the predicted frequency used for the present calculation interval is equal to a quadratic prediction of the frequency obtained over the previous interval. The frequency drift affecting in particular the H-Masers is taken into account in the prediction model.

In the UTCr algorithm, where the quadratic model is used, the frequency drift of a clock is taken to be the drift value obtained for that clock in the most recently available monthly UTC computation.

5.3.4 Accuracy of the scale interval of a time scale

Improvement of the accuracy of a time scale is generally carried out outside the main algorithm, which deals only with optimising stability.

For TAI, it is carried out by frequency steering the free-running time scale derived from the stability algorithm of UTC. The frequency corrections are smaller than the frequency fluctuations of the time scale to minimize a degradation of its stability. They are decided after comparison of the frequency of the computed time scale with a combination of the frequencies of primary frequency standards, continuously operating or occasionally evaluated, all around the world [Azoubib *et al.*, 1977; Petit, 1999]. In this exercise, the effect of the gravitational red shift on the primary standard frequencies is taken into account.


The steering of UTCr to UTC is done by replacing the past values of [UTCr - Clock] by the values [UTC - Clock] when they become available after each monthly UTC computation. This ensures that the past values of the clock data, used to compute the predictions $h_i(t)$, never diverge between UTCr and UTC in the long term. However, the steering of UTCr near the start of every month can introduce a time-step in UTCr - UTC(k) relative to UTC - UTC(k) for the same time, which affects the results until the next applied steering that means after the next UTC calculation month so it has been decided to apply it gradually over four weeks, instead of all at once.

5.3.5 Examples

5.3.5.1 Stability of UTC/TAI

To estimate the frequency stability of UTC/TAI the best references available are considered, which currently are the Terrestrial Time as realized by the BIPM (called TT) and the Rubidium clocks provided monthly as clocks by the USNO for UTC. In Fig. 36, the stability of EAL/UTC/TAI is reported to highlight the long term (at about 40 days) performance at the level of 2×10^{-16} . In the plot the short term of the red line is dominated by the noise of the link used to compare clocks and the very long term of both blue and red lines by the steering to SI second.

FIGURE 36
Frequency stability of EAL




5.3.5.2 Stability of some independent time scales

In Fig. 37 the stability of TAI-TA(USNO) and TAI-TA(PTB) are reported as example the free time scales.

The minimum values of TA(PTB) and TA(USNO) are respectively 3.6×10^{-16} and 3.5×10^{-16} .

FIGURE 37




5.3.5.3 Steering of some local representations of UTC

The examples of the stability analysis of some UTC-UTC(k) are reported in Fig. 38. The scope is to show the effect of the steering procedure to maintain UTC(k) close to UTC. The “bump” affecting the frequency stability of the data is evident in each line and is the effect of the steering procedure applied by the individual laboratories.

FIGURE 38

Examples of stability analysis of various UTC – UTC(k) ending in September 2025



HDB-31-38

5.4 Dissemination of time scales

A pre-requisite for the dissemination of time data is the accurate calibration of the equipment that is used as part of the dissemination process. Calibrating receiving equipment can be very difficult, because the receiver is often embedded in a larger device that uses the data. For example, a synchro-phasor, which is a device that is used to characterize the phase of the voltage on a power line, often contains an embedded GNSS receiver. The output of the receiver is often not available, so that validating the performance of the GNSS receiver is difficult. As a second example, the application that applies a time stamp to a financial transaction is often combined with the system that calibrates and validates the accuracy of the clock that is used by the time stamp process. Again, it is difficult to characterize the accuracy of the system clock and the statistical performance of the synchronization algorithm. These considerations are very important in any practical design, but are outside of the scope of the current discussion, which assumes that the clocks can be observed separately from the application that uses them, and that any hardware that is used in the process has been calibrated with an uncertainty that is consistent with the requirements of the application. Refer to Chapter 6 and to [ITU Handbook 2010] for complete information on time dissemination.

5.4.1 Synchronizing the time scale of a laboratory

The time and the frequency of a time scale are free parameters because time scale algorithms operate only on the measurements of time differences between clocks. Therefore, the overall time and frequency must be determined by comparing the time scale to another one whose characteristics are known. The comparison can be realized by transmitting the ensemble time of the scale to be calibrated to the reference time scale, and the frequency of the scale to be calibrated is estimated by the evolution of these time differences. The

characteristics of the channel that is used to transmit the time differences is important, and the comparison of the time difference between the two scales is generally limited by the accuracy of the determination of the channel delay, while the stability of the delay is important for determining the frequency of the scale to be calibrated. The uncertainty in the measurement of time differences is characterized by a Type A uncertainty, which can be of order 100 ps (1σ) for a channel that is realized with two-way satellite time transfer, and a Type B uncertainty, which is of order a few ns, and depends on the details of the process that is used to calibrate the channel delay.

There are other methods for comparing the ensemble times of two time scales in addition to transmitting the time data between the two scales. For example, the two ensemble times can be compared by using any clock that is simultaneously measured with respect to the two ensembles, even if the clock that is used in the comparison is not a member of either ensemble. The « common view » technique is an example of this method. The difference between two time scales, UTC(k) and UTC(k') is estimated by measuring the difference between each local time scale and the signal transmitted by a GNSS satellite.

$$UTC(k) - UTC(k') = \{UTC(k) - \{t(SV_j) + \delta_j\}\} - \{UTC(k') - \{t(SV_j) + \delta'_j\}\} \quad (5-19)$$

where:

$t(SV_j)$: time signal transmitted by satellite j

δ : path delay between the source and the receiver.

The common view method most often uses the signals transmitted by GNSS satellites because the satellites transmit information that can be used to estimate the path delay, but the method can use any signal, even if it is not intended for this application. Signals from LORAN stations and even television signals can also be used in common view without any active cooperation between the source and the receivers.

The common view method can be extended to the configuration where the signals from a single transmitter cannot be observed simultaneously at both receiving stations. This “all in view” method depends on the fact that the times transmitted by the GNSS satellites of a constellation are combined to compute a “GNSS system time”, $t(\text{GNSS})$ by the operators of the system. An estimate of the offset between the time transmitted by the satellite and the system time of the constellation is broadcast by the satellite in a “navigation message” and a more precise estimate of this difference is also computed after the fact. Then, the time difference between UT(k) and UTC(k') can be estimated by receiving the signals from two satellites, j and j' and combining these measurements with the offset of the satellite clock from the GNSS system time.

The all in view method effectively replaces the common view of a single physical transmitter with the common view of the ensemble time scale of multiple transmitters. In addition to the effects discussed in the previous section, the method depends on the accuracy of the relationship between the time transmitted by a satellite and the system time of the ensemble, and the method is usually used with post-processed ephemerides for this reason. The broadcast ephemeris is generally not accurate enough for this method.

The common view and all in view methods are also used by timing laboratories to distribute UTC(k). The laboratory publishes the difference UTC(k) - $t(\text{GNSS})$ or UTC(k) - $t(SV_j)$ (SV_j stands for satellite vehicle), and any user can combine these data with similar measurements made by the user's receiver to estimate the difference between the user's clock and UTC(k). The BIPM also publishes UTC - $t(\text{GNSS})$, which can be used for the same purpose (BIPM website, bipm.org, BIPM Time Department Database).

Two scales can also be compared by means of a portable clock that is carried between the two scales. This method may involve effects of General Relativity and is usually limited to comparisons of time scales that are not too far apart and are at similar gravitational potentials. The method is somewhat easier to implement if only the frequencies of the two scales are being compared. The clock need not be kept operating during the travel between the sites involved, and relativistic effects are not important.

The generation of a time scale and the method that is used to compare its ensemble time with respect to another scale are usually independent processes. As an example of this point, International Atomic Time (TAI) could be defined as the weighted mean of the time scales of timing laboratories, denoted by UTC(k), instead of the weighted mean of single « master clocks » at these facilities, and this method was actually used many years ago. The channels linking the timing laboratories to the BIPM would be the same in either configuration.

However, some laboratories realize UTC(k) by means of an ensemble of clocks that are located at different places. The channels that link these clocks together are then an integral part of the ensemble algorithm, and the statistics of the channel must be considered in the design of the algorithm.

5.4.2 Paper time scales

The previous discussion implicitly assumed that a time scale had a physical realization. That is, there was a physical signal that realized the ensemble time and frequency. This physical signal might be the output of one of the clocks in the laboratory (the “master clock”) or it might be realized by a phase stepper that generated a physical output based on the results of the ensemble calculation.

There are a number of important time scales that have no physical realization and therefore cannot be disseminated by any of the methods that have been discussed in the previous sections. In general, these time scales are disseminated implicitly as the difference between the time scale and another time scale or clock that generates a physical signal. The “system time” of a GNSS ensemble has already been mentioned as an example of this case. The system time is disseminated implicitly by transmitting the time difference between its time and the time of the clock in a GNSS satellite. A receiver receives the time based on the signal derived from the satellite clock and can re-construct the system time by combining this measurement with the information that is transmitted in the data channel. Other parameters in the data channel of the GPS constellation can be used to determine a prediction of UTC(USNO) from the broadcast signal, and all GNSS constellations transmit data to determine the prediction of the relevant UTC(k). The BIPM publishes values of the time differences between UTC and the system times of the various GNSS constellations (BIPM website, www.bipm.org, BIPM Time Department Database)

The time scales generated by the BIPM, Coordinated Universal Time, UTC, and International Atomic Time, TAI, are also paper time scales and have no physical realization. They are disseminated implicitly by the BIPM through Circular T, which gives the difference UTC - UTC(k) for all of the laboratories that contribute data to the calculation of these scales. As in the previous example, the UTC time scale is estimated by combining the data published in Circular T with the physical time signals transmitted by a laboratory, UTC(k). (As discussed in the previous paragraph, the signals from a GNSS satellite can be used in the same way.) The BIPM publishes Circular T after some delay, so that UTC is not available in real time. The TAI time scale differs from UTC by an exact number of integral seconds, and is estimated in the same way as UTC.

The UT1 time scale, which is related to the position of the Earth in space, is estimated by the IERS based on astronomical observations. It is also a paper time scale, and it is disseminated by the IERS in Bulletin A, which gives UT1 - UTC. Other laboratories publish similar information in various formats.

5.4.2.1 The Kalman filter

The Kalman algorithm characterizes each clock in an ensemble by a state vector. The vector for each clock has a component giving the time difference, x , measured in seconds, a dimensionless frequency, y , which describes the evolution of the time difference, and a drift parameter, d , with units of 1/s, which describes the evolution of the dimensionless frequency. (A simpler implementation does not include a drift component in the state vector.) The parameters are relative to the ensemble average time, which generally has no physical representation. In the absence of any measurements, the evolution of the state of clock j during a time interval τ is described by:

$$x_j(t_k) = x_j(t_{k-1}) + y_j(t_k)\tau + \frac{1}{2}d_j(t_k)\tau^2 + \xi_j \quad (5-20)$$

$$y_j(t_k) = y_j(t_{k-1}) + d_j(t_{k-1})\tau + \eta_j \quad (5-21)$$

$$d_j(t_k) = d_j(t_{k-1}) + \varphi_j \quad (5-22)$$

$$\tau = t_k - t_{k-1} \quad (5-23)$$

The noise contributions, ξ , η , and φ are uncorrelated variables with a mean of zero and a known standard deviation. The noise parameters are used to construct the covariance matrix. The diagonal terms of this matrix are derived from the standard deviations of the contributions, and the off-diagonal terms describe the impact of the noise contribution to any state to the variance of the other states.

The state of the system is measured at time t_k . Any component of the clock state could be measured in principle, but the most common measurement is the difference in the time states between each clock and one of them, which is designated as the “master” device. These measurements may have some noise contributions, in addition to the noise terms in the previous equations. The results of the measurements are compared with the predictions based on the evolution of the state components, and the result is the “innovation” – the difference between the predictions of the evolution equations and the results of the measurements. The Kalman gain modifies the state of each clock, taking into account the various noise contributions, and the result is an updated state vector, which then evolves until the next measurement.

The Kalman algorithm has a number of advantages compared to other algorithms. It can handle more complex noise contributions, while most other algorithms are more limited. It can accept measurements of any component of the state vector, and the measurement can be with respect to an external device that is not part of the ensemble, such as a primary reference oscillator, which provides only frequency information. On the other hand, its performance depends on an accurate description of the noise processes, and it cannot easily handle noise terms that are not Gaussian or that are not stationary. It is more difficult to add or drop member clocks, and it has no inherent method for handling outliers. The method is often implemented by using matrix algebra, and numerical round-off and truncation may be serious concerns for large ensembles. It can be difficult to compute a robust estimate of the drift component, since its contribution to the evolution of the time difference between clocks may be masked by the frequency noise contribution. This problem can be particularly difficult in ensembles that include both cesium standards and hydrogen masers; this difficulty is not unique to the Kalman method.

References

- ALLAN D.W. and ASHBY N. [1986] *Relativity in Celestial Mechanics and Astrometry*, Ed. Kovalevsky J. and Brumberg V.A., Reidel Dordrecht, 299-313.
- ALLAN D.W and THOMAS, C. [1994] *Technical directives for standardization of GPS time receiver software*, Metrologia Vol. 31, 69-79.
- AZOUZIB J., GRANVEAUD M. and GUINOT B., [1977] *Estimation of the scale unit duration of time scales*, Metrologia, Vol. 13, 87-93.
- BAUCH A., WEYERS S., PIESTER D., STALIUNIENE E. and YANG W., [2012] *Generation of UTC(PTB) as a fountain-clock based time scale*, Metrologia, Vol. 49, 180.
- Bureau international des poids et mesures [2025] *Le Système international d'unités/The International System of Units* [Brochure], 9^e Edition.
- BRUMBERG V.A. [1991] *Essential Relativistic Celestial Mechanics*, Adam Hilger, Bristol.
- CCDS Report [1980] BIPM Com. Cons. Déf. Seconde, Vol. 9, p S 15.
- CCDS Report [1993] BIPM Com. Cons. Déf. Seconde, Vol. 12, p S 69.
- CGPM [2018] *Proceedings of the 26th meeting of the General Conference on Weights and Measures* (November 2018), p. 475.
- DEFRAIGNE P., PETIT G. [2015] *CGGTTS-Version 2E: an extended standard for GNSS Time Transfer*, Metrologia, Vol. 52 G1.
- DIMARCQ N., GERSTVOLF M., MILETI G., BIZE S., OATES C.W., PEIK E., CALONICO D., IDO T., TAVELLA P., MEYNADIER F., PETIT G., PANFILO G., BERTHOLOMEW J., DEFRAIGNE P., DONLEY E.A., HEVEKVIST P.O., SESIA I., WOUTERS M., DUBÉ P., FANG F., LEVI F., LODEWYCK J., MARGOLIS H.S., NEWELL D., SLYUSAREV S., WEYERS S., UZAN J.P., YASUDA M., YU D.H., RIECK C., SCHNATZ H., HANADO Y., FUJIEDA M., POTTIE P.E., HANSEN J., MALIMON A. and ASHBY N. [2024] *Roadmap towards the redefinition of the second*, Metrologia, Vol 61, 012001.
- GUINOT B. [1987] *Some properties of algorithms for atomic time scales*, Metrologia, Vol. 24, 195-198.

- GUINOT B. [1988] *Atomic time scales for pulsar studies and other demanding applications*, Astronomy and Astrophysics, Vol. 192, 370-373.
- HACHISU H., NAKAGAWA F., HANADO Y. AND IDO T. [2018] *Months-long real-time generation of a time scale based on an optical clock*, Scientific Reports 8, 4243
- International Astronomical Union [1991] *IAU transactions Vol. XXIB*, Proc. 21st Gen. Assembly Buenos Aires, Kluwer Acad. Publ., Dordrecht, Boston, London.
- International Astronomical Union [1992] *Information Bulletin* 67, p. 7.
- International Astronomical Union [2001] *IAU transactions*, Vol. 24, Astronomical Society of the Pacific, Issue 2, 33-59.
- KLIONER S. A. [1992] *The problem of clock synchronization: a relativistic approach*, Celest. Mech. Dyn. Astron. Vol. 53, 81-109.
- LEVINE, J. AND PARKER T.E. [2002], *The Algorithm used to realize UTC(NIST)*, Proc. 2002 IEEE Intl. Freq. Cont. Symp., 537-542.
- LEVINE, J. [2012] *The statistical modeling of atomic clocks and the design of time scales*, Review of Scientific Instruments, Vol. 83, 021101-1.
- MARGOLIS H.S., PANFILO G., PETIT G., OATES C., BIZE S. [2024] *The CIPM list "Recommended values of standard frequencies", 2021 update*, Metrologia, Vol. 61, 035005
- MILNER W. R., ROBINSON J.M., KENNEDY C.J., BOTHWELL T., KEDAR D., MATEI D.G., LEGERO T., STERR U., RIEHLE F. *et al.*, [2019] *Demonstration of a timescale based on a stable optical carrier*, Phys. Rev. Lett. 123, 173201
- MISNER C.W., THORNE K.S. and WHEELER J.A. [1973] *Gravitation*, W.H. Freeman and Company, San Francisco.
- PANFILO G., HARMEGNIES A., TISSERAND L. [2014] *A new weighting procedure for UTC*, Metrologia, Vol. 5, 285-292.
- PANFILO G., ARIAS F. [2019] *The Coordinated Universal Time (UTC)*, Metrologia, Vol. 56, 042001.
- PETIT G. [1999] *Use of primary frequency standards for estimating the duration of the scale unit of TAI*, Proc. 31st PTTI Meeting, 297-304.
- PETIT G., JIANG Z. [2008], *GPS Precise point positioning for TAI computation*, IJNO Article ID 562878, doi:10.1155/2008/562878.
- PETIT G. and WOLF P. [1994] *Relativistic theory for picosecond time transfer in the vicinity of the Earth*, Astronomy and Astrophysics, Vol. 286, 971-977.
- PETIT G. [2003] *A new realization of Terrestrial Time*, Proc. 35th PTTI Meeting, 307-317.
- PETIT G. AND WOLF P. [2005] *Relativistic theory for time comparisons: a review*, Metrologia Vol. 42, S138-S144.
- PETIT G. AND JIANG Z. [2007] *GPS All in View time transfer for TAI computation*, Metrologia Vol. 45 35.
- PETIT G., ARIAS E. F., HARMEGNIES A., PANFILO G. AND TISSERAND L. [2014] *UTC_r: a rapid realization of UTC*, Metrologia, Vol. 51, 33-39.
- RIEHLE F., GILL P., ARIAS E.F., ROBERTSSON L. [2018] *The CIPM list of recommended standard values: guidelines and procedures*, Metrologia, Vol. 55, 188-200.
- ROVERA G D, BIZE S, CHUPIN B, GUÉNA J, LAURENT PH, ROSENBUSCH P, UHRICH P AND ABGRALL M. [2016] *UTC(OP) based on LNE-SYRTE atomic fountain primary frequency standards*; Metrologia, Vol. 53, 81-88.
- TAVELLA P. and PREMOLI A. [1994] *Estimating the instabilities of N clocks by measuring differences of their readings*, Metrologia, Vol. 30, 479-486.

YAO J., PARKER T.E., ASHBY N., AND LEVINE J. [2018] *Incorporating an optical clock into a time scale*, IEEE Trans. Ultrason. Ferroelectr. Freq. Control 65, 127.

YAO J., SHERMAN J.A., FORTIER T., LEOPARDI H., PARKER T., MCGREW W., ZHANG X., NICOLODI D., FASANO R. *et al.*, [2019] *Optical-clock-based time scale*, Phys. Rev. Applied 12, 044069

ITU texts

Recommendation ITU-R TF.486 – *Use of UTC as reference in standard frequency and time signals emissions*

Recommendation ITU-R TF.460 – *Standard frequency and time signal emissions*

Recommendation ITU-R TF.1153 – *The operational use of two-way satellite time and frequency transfer employing PN codes Recommendation*

Report ITU-R SM.2152 – *Definitions of Software-Defined Radio (SDR) and Cognitive Spectrum (CRS)*

Handbook “Satellite time and frequency transfer and dissemination” [ITU, 2010]

CHAPTER 6

Time and frequency dissemination

Table of Contents

	<i>Page</i>
6.1 Introduction.....	142
6.2 Radiocommunication Services.....	142
6.2.1 VLF Services	143
6.2.2 LF Services	143
6.2.3 HF Broadcast	143
6.2.4 Television Broadcast	143
6.2.5 GNSS (Broadcast)	143
6.2.6 GNSS Signals in High-Accuracy Applications	144
6.2.7 Services Based on Other Geostationary Satellites	145
6.2.8 Services Based on Meteorological Satellites	145
6.2.9 Commercial Satellite Services	145
6.2.10 TWSTFT.....	146
6.3 Standardized Wired Services	146
6.3.1 Services Based on Telephone	146
6.3.2 Solutions and Techniques Used in Computer and Telecommunication Networks.....	147
6.4 Specialized High Accuracy Fibre-Based Services	159
6.4.1 Modulation-based: ELSTAB	159
6.4.2 Optical Carrier-based.....	160
6.4.3 Building Scientific and Societal Infrastructure.....	176
6.4.4 Implementation of Optical Fibre Networks	177
6.4.5 Outlook: Fibre Sensing	178
References	178
ITU Texts	184

6.1 Introduction

The worldwide need for synchronization of local timing signals or syntonization of locally produced frequency signals with accepted national or international standards is ongoing, even growing in terms of required reliability and accuracy. Since the ITU and other scientific and technical organizations have recommended UTC as the appropriate international reference for time and frequency for most applications, it is relevant to briefly recall in this Handbook the various sources that exist for precise UTC information and the means for gaining convenient access to them.

Although the BIPM has the responsibility for establishing, maintaining, coordinating, and generally overseeing the UTC time scale, users throughout the world generally access local approximations to UTC through various national time and frequency dissemination services that are coordinated to be within close – or at least well-documented - agreement with the international UTC time scale. These services maintain and disseminate local realizations of UTC, UTC(*k*) in institute “*k*”. The level of accuracy needed at the user’s site dictates the selection of a source of time and frequency reference signals and their dissemination methods from among several alternatives. Many different systems, techniques and services existing worldwide have been used successfully to satisfy a large variety of time and frequency requirements. They include both those dedicated for time and frequency needs and others that are designed primarily for other functions, such as navigation or communications. Radiocommunication and wired time and frequency dissemination services have been compiled in this Chapter.

The term synchronization is often used with heterogeneous meaning, even within the same user group (see § 1.1). A one pulse-per-second (1 pps) generator driven by a standard frequency source produces after switch-on a signal with an undefined offset between the output 1 pps rising edge and any reference signal, whether UTC or its local realization, UTC(*k*) in institute “*k*”. Using a special function (often denoted SYNC) the output 1 pps can be *synchronized* to an external input 1 pps so that the time interval between input and output is within one full period of the driving standard frequency. By additional means (e.g. a microphase stepper), the phase of the driving and reference frequency sources can be aligned to improve synchronization of the output 1 pps rising edge with the input reference. Another usage of *synchronization* implies that the time offset between a local and a reference signal is *known* with some uncertainty, but no measures are taken to restrict it. In the ITU-T community and its associated literature, the term “network synchronization” has been in use for decades to mean adjusting the frequency of oscillators to a reference value, which is *de facto* syntonization. Only with the 5G mobile telecommunication on the horizon has synchronization in the context of ITU-T community included not only frequency but also time and phase adjustments to be aligned with the reference.

A historical note on time and frequency dissemination is included in § 1.2.

6.2 Radiocommunication services

As described in Recommendations ITU-R TF.374-6 and ITU-R TF.460-6, the ITU has allocated the following specific frequencies for time-and-frequency dissemination: 19.95-20.05 kHz; 2 495-2 505 MHz; 4 995-5 005 MHz; 9 995-10 005 MHz; 14 990-15 010 MHz; 19 990-20 010 MHz; and 24 990-25 010 MHz to the standard-frequency and time-signal service. The ITU Radio Regulations also include language which permits the use of specified portions of the 14-90 kHz region of the spectrum for time and frequency broadcasts on a secondary basis. In addition, several other frequencies were allocated to the standard-frequency and time-signal satellite service, but these have never been used for this purpose. Existing services are documented in Recommendation ITU-R TF.768, with its annually updated Annex (Recommendation ITU-R TF.768-7 Annex⁴). This Annex contains the characteristics of standard-frequency and time-signal emissions in allocated bands, and characteristics of stations emitting with regular schedules with stabilized frequencies outside of allocated bands. Recommendation ITU-R TF.583 complements the information given by describing the coding of time information in the broadcast signals. Over the last decades, the number of dedicated radiocommunication services has declined and reception of signals from GNSS has taken over in many applications.

⁴ It is accessible from the ITU-R Study Group 7 website (select Working Party 7A\Related Activities or directly <https://www.itu.int/oth/R0A08000007/en>, accessed 2026-01-09).

The description of various radio-based dissemination methods for time and frequency applications is given in this Chapter, with the restriction to frequency bands and services of actual use, now or in the past. While the ITU Handbook “Satellite time and frequency transfer and dissemination” [ITU-R Handbook 2010] describes the use of satellite signals for time comparisons, which is relevant in the context of the generation of UTC but in general not for dissemination to the broader public, such techniques are mentioned below for completeness.

6.2.1 VLF Services

VLF (very-low-frequency) broadcasts in the 10-30 kHz range are useful primarily for frequency applications. Stable propagation characteristics and long-distance coverage make such signals useful for frequency comparisons at the 1×10^{-11} level or better. Operating services primarily exist in Russia, Japan, and the US (Recommendation ITU-R TF.768). In recent years the use of VLF broadcasts for time and frequency comparisons has declined due to the emergence of other systems and techniques.

6.2.2 LF services

Various dedicated LF (low-frequency) services for time dissemination in the frequency range 40-162 kHz are operated in cooperation with institutes entrusted with the dissemination of legal time and serve millions of users. The most prominent services are WWVB in the US, MSF in the UK, DCF77 in Germany, JJY in Japan, BPC and BPL in China, ALS162 in France and RTZ and RBU in Russia. Recommendation ITU-R TF.583 contains an up-to-date compilation of the codes used for time dissemination.

Due to their long wavelengths, LF radio signals can pass through walls and buildings and be received indoors. They are diffracted over obstacles and travel beyond the horizon, following the surface of the Earth in the ground-wave propagation mode (Recommendation ITU-R P.368-9). This mode has the advantage of relatively stable propagation delays, when compared to sky wave propagation, in which signals are reflected from the ionosphere (Recommendation ITU-R P.684-7). Stable propagation delays are essential for use of the received signal as an accurate and stable source of reference frequency. Existing services provide holdover frequency that can be accurate enough as a backup for GNSS signals in critical applications if GNSS signals are unavailable [Shermann *et al.*, 2021].

Signals from terrestrial navigation systems such as LORAN-C at 100 kHz still play an important role in some regions. In Western Europe, Japan, and the US most of the stations have been decommissioned. However, in other regions, LORAN-C stations are still in operation. In some cases, they have been upgraded to provide improved performance, including the addition of a data channel, and the enhanced service is then usually termed eLoran.

6.2.3 HF broadcast

HF (high-frequency) broadcast services make use of the internationally allocated frequencies for this purpose at 2.5, 5, 10, 15, 20 and 25 MHz. Other HF services use frequencies in other bands to reduce mutual interference. Such services provide modest accuracy performance but offer advantages in terms of wide geographical coverage, convenience of use and inexpensive user equipment.

6.2.4 Television broadcast

The use of television signals for time dissemination used to be very effective and was common in some countries. It proved especially advantageous within a limited local region because of its simplicity, relatively low cost, and high accuracy. For decades it was used within Europe to bridge between the institutes operated on both sides of the Iron Curtain [Tolman *et al.*, 1967]. With the development of digital broadcast, however, the properties of received signals no longer allow this use to be continued because of the variable and unpredictable propagation and signal-processing times.

6.2.5 GNSS (Broadcast)

Recommendation ITU-R TF.374-6 explicitly refers to the frequency bands commonly used by the GNSS. In addition to positioning, GNSS signals can be used for time determination and synchronization applications. Signals from the satellites of the Global Positioning System (GPS) – the first GNSS – have been used since the early 1980s for time comparisons. The primary purpose of all GNSS, namely GPS, GLONASS, Galileo

and BeiDou, is positioning and navigation. GNSS systems rely on precise timing, as the satellite ranges used to calculate position are derived from propagation time measurements of the signals transmitted from each satellite. A basic description of the design and operation of a GNSS is given in § 7.5 of this Handbook. The signals broadcast by GNSS satellites are derived from onboard atomic clocks (caesium beam clocks, rubidium gas cell clocks, passive hydrogen masers) and contain timing and orbit information in their broadcast navigation message. GNSS timing receivers have been developed, which use the received signal to discipline inbuilt oscillators to GNSS system time or, preferentially, to the prediction of UTC contained in the navigation message. Such GNSS DO (for Disciplined Oscillators) usually deliver a one pulse-per-second (1 pps) electrical signal or even a set of output signals (standard frequency signals, signals for telecommunication applications and further wire-based dissemination, see § 6.4) that represent, with some uncertainty, either the GNSS system time or UTC as predicted in the GNSS navigation message. The use of such devices in calibration laboratories, telecommunication and power facilities is very common. Many receivers are embedded in servers providing time (including time-of-day) information via network protocols such as NTP or PTP (see below). There is an ongoing debate on how traceability to national or international standards can be established with reasonable effort using such instrumentation [Lombardi, 2016, Matsakis *et al.*, 2018, Defraigne *et al.*, 2022] which is taken up in Chapter 8.

With the widespread use of GNSS DO and their relevance in many application sectors in mind, the following describes the basis of their operation in some more detail.

In the GNSS ground segment, the common GNSS system time (GNSS_T) is generated from an ensemble of atomic clocks. The offset between GNSS_T and some selected available realization UTC(*k*) of UTC is also determined in the GNSS ground segment and serves to build the quantity [Broadcast_UTC_{GNSS} – GNSS_T] for broadcast by the satellites. In a GNSS DO the local user time scale TS_{user} is typically realized in such a way that the offset between TS_{user} and Broadcast_UTC_{GNSS} is close to zero (for timing applications) or at least kept constant (for frequency applications). Different GNSS rely on different UTC(*k*) and different algorithms to determine the offset between the specific UTC(*k*) and the respective GNSS_T and to predict its evolution into the future.

To relate TS_{user} to UTC, the difference UTC – Broadcast_UTC_{GNSS} need to be added. This latter can be obtained from the following relation

$$\begin{aligned} \text{UTC} - \text{Broadcast_UTC}_{\text{GNSS}} = \\ [\text{UTC} - \text{UTC}(k)] + [\text{UTC}(k) - \text{GNSS_T}] - [\text{Broadcast_UTC}_{\text{GNSS}} - \text{GNSS_T}] \end{aligned} \quad (6-1)$$

The quantity UTC – UTC(*k*) is provided by BIPM in Circular T as explained in Chapter 5. The second term can be provided by a UTC(*k*) timing laboratory which operates a calibrated GNSS receiver. BIPM publishes daily offsets in Section 4 of its Circular T.

Reception of GNSS signals suffers from delays introduced by the ionosphere and troposphere, and by multipath and receiver noise. A wide range of performance in terms of frequency instability and offset of the 1 pps output signal from UTC can be obtained, depending on the sophistication of the equipment and the efforts to determine the signal delays in the GNSS reception chain. Typical values achieved after several thousand seconds of averaging are relative frequency instability of 10^{-12} and time offset from UTC below 100 ns.

6.2.6 GNSS signals in high-accuracy applications

High-accuracy applications often involve the comparison of two time scales at distant sites rather than the dissemination in a somewhat hierarchical manner of a reference time to a user. A good example is the collaboration of about 80 timing institutes across the world in realizing UTC, cooperating with and guided by the BIPM. The reception of GPS signals has been used since the 1980s following the Common-View scheme (GPS CV) [Allan and Weiss, 1980], or later its variant called GNSS All-In-View (AV) [Petit and Jiang, 2008a]. In early 2022, all the time links between laboratories “*k*” which realize a replica of UTC, named UTC(*k*), are based on GNSS, either GPS only (in the case of 87% of the links), or combined with Two-Way Satellite Time and Frequency Transfer (TWSTFT) (to be described later). Signals from the other GNSS have been used and will be used more extensively in the future, based on the same methods. In CV, two stations simultaneously

observe the same GNSS satellite. At each site the difference between local clock (T) or local time reference and GNSS_T is recorded at the same instant using the same satellite using dedicated GNSS time receivers.

$$\text{User A observes: } A = T_a - \text{GNSS_T} \quad (6-2)$$

$$\text{User B observes: } B = T_b - \text{GNSS_T} \quad (6-3)$$

By computing the difference between the two sets of numbers:

$$A - B = (T_a - \text{GNSS_T}) - (T_b - \text{GNSS_T}) = T_a - T_b \quad (6-4)$$

the differences between the two local clocks can easily be calculated. In the AV method, averages over all available observations at a given epoch are made first before the difference is formed. For long-distance links this brings the advantage of a larger number of useful observations. Common errors in the received signal, however, due to incorrect space vehicle clock and orbit modelling no longer cancel. Thus, AV typically includes post-processing using products provided by the International GNSS Service (IGS) [Johnston *et al.*, 2017]. CV and AV traditionally involve code-based measurements, and for inter-continental links the BIPM reports values for the standard uncertainty representing the link instability between 1.5 ns and 5 ns. This value is often smaller than the uncertainty for accurate time transfer which is governed by the uncertainty of the determination of the signal delays in the receiving equipment [BIPM, 2021].

The highest performance GNSS-based approach involves Precise Point Positioning (PPP). Contrary to GNSS CV and AV based on code measurements only, PPP is a one-way method based on dual-frequency code and carrier phase observations taken at one location [Petit and Jiang, 2008b]. Combined with IGS products (precise orbit and clock predictions) the position of the receiving station and the time offset between local clock and the time reference of the satellite clock products is obtained. In early 2025 the required products are available routinely only for GPS. Initially, a frequency transfer performance of order 10^{-15} at one-day averaging was achieved, limited by the noise of the code measurements used to determine carrier phase ambiguities in the long term. The next achievement preserves the integer nature of the carrier phase ambiguities and is termed Integer ambiguity PPP (IPPP) [Petit *et al.*, 2015, Petit, 2021]. Application of IPPP for some UTC(*k*) links has been studied but has not yet become a BIPM routine in 2025. After a week of averaging, frequency comparison is feasible with 10^{-16} uncertainty, decreasing further for averaging times up to one month. The accuracy of time comparisons remains linked to the calibration uncertainty and does not immediately improve.

6.2.7 Services based on other geostationary satellites

Other satellite systems were used for transmitting time signals in a one-way mode in the past. A Standard-Frequency and Time-Signal (SFTS) Dissemination Service using the Indian National Satellite programme (INSAT) system was provided by National Physical Laboratory (Delhi) for about 20 years up to the late 1990s. The service consisted of a train of 5 kHz burst signals, which were frequency modulated on the carrier. With the growing popularity of GPS receivers, the cost for maintaining the service became prohibitive.

Occasionally communication satellites have been used for time and frequency transfer. Design and function of such satellite services have been reported in Chapter 6 of the Handbook “Satellite time and frequency transfer and dissemination” [ITU, 2010].

6.2.8 Services based on meteorological satellites

Between 1974 and 2004, the U.S. Geostationary Operational Environmental Satellite System (GOES) transmitted a time code referenced to the UTC(NIST) time scale. The time code was disseminated continuously from two geostationary satellites located normally at 75° and 135° West longitude. Satellite position data were also transmitted to users so that suitable automatic receivers can compute the signal path delay and correct their 1 Hz outputs accordingly. Specified time code accuracy as delivered to the user was 100 μs [Lombardi and Hanson, 2005].

6.2.9 Commercial satellite services

The STL (Satellite Time and Location) service has been developed by Satelles, acquired by Iridium. The development was motivated by the observation that since GPS had been introduced, many of its applications, such as positioning and time synchronization, had gained greater importance, in particular though its abundant

use in sectors supporting the critical infrastructure of a nation [Sherman *et al.*, 2021]. Specifically, it was noted that some of those applications were hindered by the insufficient performance of GPS within indoor environments and scenarios like urban canyons. Because of the low signal power of GNSS, the signals are prone to jamming, and the proliferation of software defined signal generation devices increased the risk of spoofing as well. As a complement or alternative to GPS, STL's key feature is the use of the Iridium constellation of satellites, which was originally launched to provide voice and data communications to satellite phones and pagers with worldwide coverage. The constellation consists of 11 satellites on each of the 6 orbital planes, with an orbit height of about 780 km and inclination of 86.4° (almost polar orbit). The STL service uses burst signals specifically designed for position, navigation, and timing applications Novick A, Johnson P [2024]. Due to the low orbit height, the signal strength can be higher by up to a factor of 1000 compared to GPS. Cryptographic features are included in the signal to offer increased security. STL has been identified as a possible building block for improved resilience of time and frequency transfer services [Sherman *et al.*, 2021]. Limitations to its usefulness arise from the small footprint of the satellite antennas on the ground and potential issues due to low satellite coverage at low latitudes, substantially less than GNSS [Reid *et al.*, 2016]. No wide-spread experience can be reported as of early 2025. Other commercial satellite services may in the future play a similar role.

6.2.10 Two-way satellite time and frequency transfer

For completeness, it is briefly mentioned Two-Way Satellite Time and Frequency Transfer (TWSTFT) whose operational use has been described in Recommendation ITU-R TF.1153-4. A detailed description can also be found in the Handbook "Satellite time and frequency transfer and dissemination" [ITU, 2010]. It has been developed into a technique for time and frequency transfer between laboratories that contribute with their atomic clocks and, in some cases, with their primary frequency standards, to the realization of UTC, as an independent alternate to GNSS-based comparisons. The TWSTFT technique is thus used in two areas: for true time transfer, by comparing local time scales such as local realizations of UTC, and for accurate frequency comparisons between atomic clocks. TWSTFT is typically performed in networks with multiple participating stations using code division multiple access (CDMA) and the need to standardize:

- measuring procedures,
- data processing, and
- formats for the exchange of data and relevant information between participating stations and interested bodies, such as the BIPM

was noted at an early stage. Since the two-way technique is essentially a point-to-point communication system, it should not be regarded as a general dissemination technique. In high-risk critical infrastructure management, it can constitute a GNSS-independent method that provides better immunity to jamming and spoofing [Sherman *et al.*, 2021]. For general applications its cost and operational complexity are prohibitive.

6.3 Standardized wired services

This section describes wired services that provide standardized and scalable time and/or frequency dissemination methods. These services can complement, replace and/or backup the radiocommunication services described in § 6.2. While many of these services are less accurate than those described in § 6.4, the services described in this section allow to reach hundreds, thousands or even millions of users.

6.3.1 Services based on telephone

Several timing centres all over the world have established services designed to disseminate coded time information over telephone lines in an automated mode. Typically, computers and other automated systems are programmed to dial such services as needed, receive an ASCII time code and reset the local clock. Some systems are designed for automatic correction for the path delay involved through the telephone system. In such case, accuracies in the range of 1-10 ms are possible, but often the services are just used for ensuring time-of-day (TOD) to about 1-second accuracy. The available time dissemination services and time signals emissions, including services based on telephone, are published by BIPM on its Time Department Data Base, Sub-section Dissemination services [BIPM, 2026]. In 1990, several time-keeping institutes in Europe agreed upon a common format for time dissemination via the public telephone network in the frame of a EUROMET

(what is now called EURAMET, the European Regional Metrology Organization) project. This “European telephone time code” is today still being distributed by eight institutes. The European telephone time code *inter alia* contains DUT1 (the offset between UT1 and UTC with a resolution of 0.1 s) as defined in Recommendation ITU-R TF.460-6. The anticipated increase of DUT1 beyond 0.9 s will make a change of the coding necessary.

6.3.2 Solutions and techniques used in computer and telecommunication networks

6.3.2.1 Introduction

Computer and telecommunication networks have evolved into mainly packet networks. Packet networks are primarily asynchronous in nature and thus pose a major challenge for time and frequency transfer. On the other hand, they provide highly scalable solutions maximizing the reach of time and frequency dissemination. The time and/or frequency transfer through such networks is either supported as an additional feature or transparent to the network. The terms time synchronization and frequency synchronization (syntonization) are commonly used in the context of packet networks when referring to the time and frequency transfer through such networks⁵.


Synchronization

In packet networks, synchronization of network elements is achieved through two-way time transfer based on packet timestamping. The basic principle is illustrated in Fig. 39 a). A network device A transmits a packet and records its transmission time (t_1) with the local time of device A. The same packet is received by a network device B which records the reception time (t_2) using its local time. This process is repeated in the other direction to obtain t_3 and t_4 .

⁵ In ITU-R Recommendations and in this Handbook, the terms ‘synchronization’ and ‘syntonization’ are used as defined in Recommendation ITU-R TF.686 in the context of time and frequency transfer. In the context of telecommunication networks, the term ‘synchronization’ can have a broader meaning (e.g. data synchronization) and both frequency and time are relevant, thus in ITU-T Recommendations the term “time synchronization” is used to highlight the timing context, as defined in Recommendation ITU-T G.8260. The term ‘frequency synchronization’ is used to refer to syntonization, see Recommendation ITU-T G.8261.

FIGURE 39

Principles of two-way time transfer and frequency transfer in packet networks



HDB-31-39

Assuming that all timestamps are available and that the medium is symmetric, the one-way delay introduced by the network between the two devices can be calculated as follows:

$$\text{one_way_delay} = [(t_4 - t_1) - (t_3 - t_2)]/2 \quad (6-5)$$

The offset between the local time of device A and the local time of device B can be calculated as follows:

$$\text{offset} = \text{time_of_B} - \text{time_of_A} = t_2 - (t_1 + \text{one_way_delay}) \quad (6-6)$$

The offset information allows correcting the internal time of one of the devices to be aligned with the internal time of the other, with some uncertainty which results from

- the precision of the timestamps
- the asymmetries of the data path
- the frequency offset between the two network elements.

The method above is used by Network Time Protocol (NTP), Precision Time Protocol (PTP), White Rabbit (WR) and High Accuracy (HA) PTP, all described in the following sections.

Frequency synchronization (syntonization)

In packet networks, syntonization of network devices can be achieved in two ways. One method consists in timestamping transmitted and received packets with the local times of the transmitter and the receiver respectively (see Fig. 39 b)). Based on these timestamps, the rate of reception of packets is compared with the rate of their transmission. The calculated frequency offset between the transmitter and the receiver is used to correct the frequency of the receiver. Such a frequency correction is typically used along with the time offset correction when synchronizing the two network devices, as described above.

Another method for syntonization of network devices in packet networks uses physical layer syntonization. Figure 40 illustrates this concept in which network device A encodes a data stream transmitted on its interfaces with a frequency derived from a frequency source (e.g. frequency standard). Network device B recovers the


frequency from the received data stream on an interface directly connected to network device A, using a Clock and Data Recovery (CDR) circuit. A Phase-locked loop (PLL) is used by network device B to obtain a frequency syntonized to the frequency source of network device A. This frequency is then used to transmit data on interfaces of network device B and is typically available directly through various outputs.

The first method of syntonization is used by NTP and PTP, while the latter is used by Synchronous Ethernet (SyncE), WR and HA PTP.

The time and frequency disseminated over packet networks, described in the following sections is typically obtained from a GNSS-disciplined frequency standard or directly from an institute “k” realizing a UTC(k).

FIGURE 40

Physical layer syntonization concept (modified figure from [Jobert *et al.*, 2013])




HDB-31-40

6.3.2.2 Network Time Protocol (NTP)

Network Time Protocol (NTP) is the most common method for synchronizing personal computers, servers, networked devices and appliances over the Internet public network and it is one of the oldest Internet protocols in current use. In existence since 1985 (NTPv0), its version 1 (NTPv1) was brought to attention of a wider engineering community with [MILLS D, 1991], its latest version 4 (NTPv4) is specified in [IETF 2010] and accompanying documents. The reference implementation and detailed documentation is available from [NTP ORG, 2025].

FIGURE 41

Typical architecture of interconnecting NTP clients and servers



HDB-31-41

The protocol allows a device implementing NTP client functionality (client) to synchronize its time to an approximation of UTC, typically to within a few milliseconds, by exchanging NTP packets with one or more devices implementing NTP server functionality (server). A server typically includes NTP client functionality. The synchronization process through packets exchange is initiated by the client communicating directly with the server using User Datagram Protocol (UDP) over Internet Protocol (IP) version 4 or version 6 (IPv4 or IPv6). The evaluation of time offset between a client and a server is performed using the two-way time transfer technique described before in Fig. 39 a). No network support for the protocol is assumed between the server and the client, thus packet delay variations and path asymmetries due to different routing in forward and backward direction are expected. The timestamps are typically generated in software, which reduces their precision and accuracy. The timescale used by NTP is based on UTC and the NTP timestamps express the UTC time as the number of seconds and fractions of a second that have elapsed since 1 January 1900. As such, care must be taken to maintain the monotonicity of the perceived time during the leap second events, as described in detail in the documentation published by [NTP ORG, 2025].

As depicted in Fig. 41, NTP uses a hierarchical, semi-layered architecture. It assumes a number of widely accessible primary time servers (stratum 1), each connected to a time source (stratum 0) which is synchronized to UTC, typically via GNSS or a direct connection to a UTC(k) time scales maintained by a national laboratory. A stratum N server ($1 < N < 16$) can establish a client-server communication with many stratum N-1 servers and peer-to-peer communication with equal stratum servers. The algorithms specified by NTP take advantage of multiple connections and aim at:

- Maintaining multiple connections and correcting the asymmetrical delays of each connection;
- Applying filtering to reduce errors;
- Detecting abnormal behaviours, such as byzantine failures;
- Combining time sources (weighted average value).

These algorithms mitigate errors that may result from network disruptions, server failures, and possible hostile actions. While some security mechanisms are specified in NTPv4 [IETF 2010], their limitations and flaws

have been identified. Therefore, Network Time Security was introduced in [IETF, 2020]. It specifies a mechanism for using Transport Layer Security (TLS) and Authenticated Encryption with Associated Data (AEAD) to provide cryptographic security for the client-server communication of the NTP. It allows confirming the identity of the servers that are exchanging time information and protect the transmission of that time information across the network.

6.3.2.3 Precision Time Protocol (PTP)

The IEEE1588 standard, titled “IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems”, defines the Precision Time Protocol (PTP). It allows precise in-band synchronization in distributed systems interconnected with packet networks. The first edition [IEEE, 2002] of this standard was developed for test and measurement (instrumentation and data acquisition), industrial automation and power generation applications, interconnected using IPv4 networks [Eidson, 2006]. Updates included in the second edition [IEEE, 2008] were motivated by adaptation of the standard by other industries, notably telecommunication, and over other transport protocols such as IPv6 and IEEE 802.3 [Jobert *et al.*, 2013]. The third edition of the standard [IEEE, 2019] consolidated its growing use across many industries and in scientific installation [P1588, 2020].

The IEEE1588 standard evolved into a flexible framework which can be fine-tuned to the specific needs of different industries. It defines:

- The PTP that follows the principles of two-way time transfer described in § 6.3.2.1 and exchanges properties of participating PTP devices, through PTP messages.
- PTP device types and their behaviour, e.g. Boundary Clock, Ordinary Clock, and Transparent Clock.
- A distributed distance-vector algorithm (BMCA, see section 9.3 in [IEEE, 2019]) that allows the selection of a PTP device that is the source of time/frequency (the Grandmaster) for the other PTP devices, and determination of the synchronization hierarchy.
- The means to extend the basic behaviour of the protocol.
- Mappings of the protocol to different transport standards (e.g. IPv4, IPv6 and IEEE802.3/Ethernet).
- The concept of PTP Profiles, which allow configuring the protocol to meet the needs of a particular industry, as well as defining new extensions and performance requirements. Industry-specific PTP Profiles are defined in separate documents published by standard-defining organizations (SDOs), such as ITU-T for telecommunication (e.g. PTP Profiles in Recommendations ITU-T G.8265.1, G.8275.1 and G.8275.2), SMPTE for professional broadcast (PTP Profile in ST2059-2 [SMPTE, 2021]), AES and IEEE for Audio/Video (e.g. PTP Profiles in AES67 [AES, 2023] and IEEE 802.1AS [IEEE, 2026]), IEC and IEEE for Power Industry (e.g. PTP Profiles in IEC 62439-3 [IEC, 2021] and IEEE C37.238 [IEEE, 2017]), Avnu and AUTOSAR for Automotive (e.g. Auto-Ethernet-AVB-Func-Interop-Spec or AUTOSAR FO R20-11 [AUTOSAR, 2020]) and LXI and AIA for Test and Measurement (e.g. PTP Profile in LXI 1588 Profile [LXI, 2008]).
- Default PTP Profiles which are meant to cover a wide variety of typical industry requirements.

In a typical PTP network, time is distributed from a single Grandmaster through a packet network to one or more Ordinary Clocks. The Grandmaster is typically connected to an external source of time and frequency providing an approximation of UTC, such as a GNSS-disciplined frequency standard. In a such case, the PTP disseminates the “timescale PTP” (see section 7.2 in [IEEE, 2019]) which is defined as “*the elapsed time since the PTP epoch measured using the second defined by International Atomic Time (TAI)*” where PTP epoch is 1 January 1970 00:00:00 TAI. The “timescale PTP” is monotonic, and the protocol distributes information about leap seconds so that the PTP devices can provide users information about UTC.

An Ordinary Clock serves either as the Grandmaster or synchronizes its local time to that of the Grandmaster. Between the Grandmaster and the Ordinary Clocks, the network equipment can be either transparent to PTP, forwarding or routing but not processing PTP messages, or it may assist PTP, providing on-path support. There are two types of such assisting PTP devices: a Boundary Clock and a Transparent Clock.


A Boundary Clock (introduced in the first edition [IEEE, 2002]) has multiple PTP Ports. It participates in the BMCA and synchronizes its local time to that of the Grandmaster. It can also serve as the Grandmaster.

A Transparent Clock (introduced in the second edition [IEEE, 2008]) has multiple PTP Ports but it neither participates in the BMCA, nor synchronizes its local time to that of the Grandmaster. Instead, it evaluates the time it takes this Transparent Clock to forward or route certain types of PTP message and introduces this information in the PTP messages so that a receiving Boundary or Ordinary Clock can better estimate the one-way path delay introduced by the network, excluding the Transparent Clock. A Transparent Clock cannot serve as the Grandmaster.

As depicted in Fig. 42, a PTP network consists of a PTP device serving as a Grandmaster and one or more PTP devices, which can be Ordinary Clocks (OCs), Boundary Clocks (BCs), Transparent Clocks (TCs), or any combination thereof. These can be connected directly or through PTP-unaware networking equipment. Each Boundary and Ordinary Clock typically implements a servo loop which synchronizes their local time to that of the adjacent Boundary or Ordinary Clock, higher in the hierarchy established by the BMCA, following the principles of a two-way time transfer, see § 6.3.2.1. Effectively, in a PTP network, all the Boundary and Ordinary Clocks that belong to the same PTP domain (a logic grouping within a PTP network) synchronize their local time to a single Grandmaster.

FIGURE 42

A PTP network with different PTP device types and PTP unaware segments



HDB-31-42

In some networking topologies (e.g. ring) where a large number of Boundary Clocks would be chained, the fact that each Boundary Clock implements a servo mechanism with certain characteristics (e.g. loop bandwidth) might result in degraded or unstable synchronization performance. Such topologies are better served by Transparent Clocks.

Notably, the IEEE1588 standard does not define stringent performance requirements for PTP devices and therefore does not guarantee any particular synchronization performance. While it purposely accommodates simple and lightweight implementations (e.g. for embedded devices, sensors or even wall-clocks), the standard includes mechanisms, recommendations, guidelines and examples that allow implementations to achieve accuracy ranging from sub-microsecond to sub-nanosecond, in the case of the High Accuracy Default PTP Profile, see § 6.3.2.5. The main factors that improve synchronization performance to better than achievable by NTP include:

- full on-path support (Boundary and Transparent Clocks),
- well-defined timestamp generation and wide-spread support of hardware timestamping,
- hardware latencies and medium asymmetry correction.

While the IEEE1588 standard does not define nor guarantee synchronization performance, such requirements are defined by some of the PTP Profiles. Notable examples include ITU-T profiles described in the next section.

6.3.2.4 Distribution of time and frequency in telecommunication networks

The International Telecommunication Union Telecommunication Standardization Sector (ITU-T) is responsible for technical specifications (ITU-T Recommendations) giving shape to global communication infrastructure, including Recommendations concerning network synchronization and time distribution performance in telecommunication networks. Since the 90s, these Recommendations have been under the responsibility ITU-T Study Group 15 [ITU-T SG15]. A short description of high-level concepts is provided in this Handbook for completeness. For more details refer to Technical Flyers in [ITU-T SG15].

Historical background

Syntonization, the adjustment of a frequency to match another frequency, has been important for telecommunication networks since the introduction of the Plesiochronous Digital Hierarchy (PDH) in transport networks and digital switches in Public Switch Telephone Networks (PSTN) in the 1970s. The digital switches had to be locked to a common reference frequency to prevent buffer overflows (so-called slips) due to frequency offsets, driving performance requirements. These requirements were inherited by newer technologies which allowed transport with higher bandwidth over optical fibre networks: Synchronous Digital Hierarchy (SDH) and Synchronous Optical Networks (SONET) introduced in the 1980s, as well as Optical Transport Network (OTN) in the 2000s, and Synchronous Ethernet in the 2010s. Since the introduction of the Global System for Mobile Communications (GSM) in the 1990s, the evolution of mobile networks and the related transport network is the key driver of requirements in telecommunication networks, initially for frequency and later for time/phase dissemination. In mobile networks, the needs are mainly related to the correct generation of radio signals. For an overview of the history of synchronization in telecom network, see chapter 1 of [Jobert *et al.*, 2013].

SyncE: Frequency transfer


Synchronous Ethernet (SyncE) introduces physical layer syntonization to transport frequency over the Ethernet medium in packet networks, the concept described in Fig. 40 and § 6.3.2.1. It also defines the exchange of Synchronization Status Messages (SSM) with information about the frequency reference. In terms of frequency transfer specifications, it is compatible with SDH and SONET, as well as IEEE 802.3 and its applicable interfaces. Relevant ITU-T Recommendations can be found in Fig. 44.

PTP telecom profiles: Frequency and time/phase transfer

ITU-T has specified three PTP Profiles to cover different deployment scenarios. The PTP Telecom Profile for Frequency allows frequency transfer using PTP messages over network segments without SyncE support, using the method explained in § 6.3.2.1, Fig. 39 b). The two PTP Profiles for Time/Phase allow frequency and time transfer combining SyncE-based frequency transfer and PTP-based time transfer, see Fig. 43. The first profile assumes full timing support (PTP and SyncE) from the network while the other allows partial timing support. Relevant ITU-T Recommendations can be found in Fig. 44.

FIGURE 43

PTP synchronization and SyncE syntonization in the PTP Profile for Time/Phase with full timing support



HDB-31-43


Primary references: Frequency and time sources

Mobile communication relies on frequency, and increasingly often time/phase, provided by telecommunication networks. Adjacent mobile equipment might be connected to different sources of frequency and time/phase while requiring certain level of frequency and time/phase alignment among them. As such, ITU-T not only provides recommendations for the frequency and time/phase transfer over telecommunication networks but also specifies performance of frequency and time source with respect to a common reference, i.e. UTC. These specifications define characteristics of frequency sources called Primary Reference Clocks (PRC) and Enhanced Primary Reference Clocks (ePRC), as well as frequency and time/phase sources called Primary Reference Time Clocks (PRTC), Enhanced Primary Reference Time Clocks (ePRTC) and Coherent Network Primary Reference Time Clocks (cnPRTC). Relevant ITU-T Recommendations can be found in Fig. 44.

The above-mentioned references for telecommunication networks, such as PRC, ePRC, PRTC, ePRTC, cnPRTC, effectively rely on the dissemination, reception and exchange of standard frequency and time signals as defined by ITU-R.

FIGURE 44

Summary of ITU-T Recommendations relevant for frequency and time/phase transfer over packet networks⁶



⁶ Published by ITU-T Q13/15, URL from November 2025: https://www.itu.int/en/ITU-T/studygroups/2022-2024/15/Documents/flyers/Flyer_ITU-T_Overview-Question13.pdf


6.3.2.5 White Rabbit and High Accuracy PTP

White Rabbit (WR, [WR, 2021; Wlostowski, 2011; Lipinski *et al.*, 2011]) is an open-source implementation of the IEEE1588 standard which provides sub-nanosecond synchronization over Ethernet fibre-optic links. WR combines PTP (see § 6.3.2.4) and its two-way time transfer with

- physical layer syntonization (similar to SyncE described in § 6.3.2.2) to ensure accurate frequency transfer without increased PTP message rate and allow phase comparison between received and transmitted frequencies (see Fig. 43). Unlike the PTP Profile for Time/Phase with Full Timing Support defined in [Recommendation ITU-T G.8275.1], in WR, the PTP synchronization and physical layer syntonization are coherent and follow the same hierarchy (see and compare Figs 43 and 45);
- phase detection, in particular Digital Dual Mixer Time Difference (see § 4.3.2 and [Moreira, 2010]), to improve timestamping precision of PTP messages (t_1, t_2, t_3, t_4 in § 6.4.2.1);
- link delay model which allows correction of medium and hardware asymmetries through their online estimation and/or careful calibration [Daniluk, 2015].

FIGURE 45

PTP synchronization, physical layer syntonization, phase detection in WR/HA



HDB-31-45

These techniques require support from the protocol and impose functional requirements on the PTP device hardware implementation. They are specified in [Cota, 2011] which defines the “White Rabbit” PTP Profile (WR Profile), describes its PTP enhancements and the required hardware support and implementation.

The above constituted an input to the P1588 Working Group responsible for IEEE1588 standard development. The study of the enhancements introduced by WR resulted in their generalization and inclusion into the third edition of the IEEE1588 standard [IEEE, 2019], a process detailed in [Lipinski, 2022]. This edition defines in Annex I.5 a new “High Accuracy Request-Response Default PTP Profile” (HA Profile). This profile mandates support of a number of WR-inspired optional features (see Annex L, sections 16.7, 16.8, 17.6 and 9.2.2.2 in [IEEE, 2019]), as well as support of calibration procedures (see Annex N in [IEEE, 2019]). As described in § 6.3.2.4, the IEEE1588 standard does not include performance requirements. However, [IEEE, 2019] provides in Annex M an example of the HA Profile implementation that achieves synchronization performance equivalent to that of the WR Profile. In fact, the two profiles, WR and HA, require identical hardware support

and differ only in terms of the protocol. The implementation of both profiles is made available by CERN (European Organization for Nuclear Research) and collaborators under open-source licenses (hardware, firmware, software, gateware – see [WR, 2025]) and the term “White Rabbit” is now used to refer to this implementation of both profiles.

White Rabbit requires full on-path support with Boundary or Ordinary Clocks inter-connected by direct fibre optic links. It assumes single-mode single fibre and uses a different wavelength in each direction within the fibre. Typically, the IEEE 802.3ah 1000BASE-BX10 standard is followed with optical modules (Small Form-Factor Pluggable, SFP) transmitting at 1 490 and 1 310 nm, allowing to reach up to 20 km. White Rabbit is also suitable for Coarse Wavelength Division Multiplexing (CWDM) and Dense Wavelength Division Multiplexing (DWDM) as demonstrated in [McKenzie, 2024; Kaur, 2022; Boven, 2018], with fibre links up to 170 km. The longest WR link to date is the 950 km link between Espoo and Kajaani in Finland [Dierikx, 2016]. It is realized in a dual fibre using the same wavelength for communication in both directions and, after calibration, was verified to agree with GPS Precise Point Positioning (PPP) measurements within ± 2 ns, as described in [Dierikx, 2016].

In terms of time transfer, the performance of White Rabbit is limited by:


- the repeatability of time/phase offset between WR devices over resets or link re-establishing, which has been measured in [Hazell, 2024] to be peak-to-peak below 22 ps,
- uncertainty of calibration which is impacted by repeatability of time/phase offset and part-to-part variation,
- temperature dependence of parameters which WR assumes to be constant, such as chromatic dispersion, transmission wavelengths or hardware delays. Their magnitude depends on the link length, type of SFP used and WR device design [Boven, 2018].

On up to 10 km links using 1 490 nm and 1 310 nm or up to 100 km links using adjacent DWDM channels, accuracy better than 100 ps per link can be expected.

Figure 46 depicts the relative frequency instability, expressed by the Modified Allan Deviation (MDEV), and phase noise measured in [Rizzi, 2018] at the 10 MHz output of a WR Grandmaster (GM) connected to a frequency standard, and 10 MHz output of a WR Boundary Clock (BC) connected to this Grandmaster with 10 km link using 1490 nm and 1310 nm. Two classes of devices are distinguished: Class I – basic and Class II – low jitter, see Fig. 46.

FIGURE 46


Relative frequency instability (Modified Allan Deviation, MDEV) and phase noise measured in [Rizzi, 2018] for the Grandmaster (GM) and Boundary Clock (BC) connected to it over 10 km fibre. Two classes of devices are distinguished: Class I – basic and Class II – low jitter



This performance can be improved and longer DWDM links used, as illustrated in [Kaur, 2017] where a cascade of five WR devices (modified PLL bandwidth and packet rate) interconnected using 125 km DWDM links shows results depicted in Fig. 47.

FIGURE 47

Relative frequency instability (Allan Deviation, ADEV) at the output of a cascaded 500 km WR link, compared with the output of a GPS disciplined oscillator, reported in [Kaur, 2017]



HDB-31-47

6.3.2.6 Other solutions

Other solutions for time and frequency transfer in computer and telecommunication networks exist, which are typically proprietary and non-standard. Examples are provided below.

Dynamic Synchronous Transfer Mode (DTM)

The Dynamic Synchronous Transfer Mode (DTM), described in [Danielson, 2014] is a transmission system employing time division multiplexing and a circuit-switching network technology. It is designed to provide a guaranteed quality of service for streaming video services but can be used for packet-based services as well. The DTM builds on the Time Division Multiplex principle, using 8 kHz frame rate to be compatible with telecommunication infrastructure and 64-bit data-slots. Frames are continuously exchanged between nodes and the starting channels of the frames in both directions hold among other things node time, local time, leap second information, and correction factors. Synchronization between network elements is achieved using the two-way time transfer technique described in Fig. 39 a) in combination with asymmetry estimation which takes advantage of the meshed topology of the network. DTM has been implemented in 15 DVB-T networks in Sweden, the largest covering 800 nodes. With a bi-directional exchange of information between nodes, two-way time transfer can be achieved. With this technique in a controlled environment, in which all network nodes and links are under control of one operator, time transfer at a level of a few ten-nanosecond level is achievable.

Precision TimeNet

Precision TimeNet [Tecknexus, 2024] is an overlay synchronization technique which can run over public managed IP networks. As its basis, it follows the principles of the two-way time transfer technique described in Fig. 39 a). It uses high and adjustable packet rates (500–40 000+ packet per second) and sizes to handle packet delay variations, as well as intelligent filtering algorithms with adjustable filter bands. It takes advantage of meshed network topology to mitigate asymmetries and includes mechanisms to detect and mitigate synchronization deterioration due to reconfiguration of underlying networks. It is used for Digital Terrestrial Television networks as well as telecommunication networks as basis for 5G infrastructure, e.g. in Sweden and Turkey.

6.4 Specialized high accuracy fibre-based services


6.4.1 Modulation-based: ELSTAB

ELSTAB is a commercial high-accuracy time and frequency transfer and distribution technology using a single-mode fibre. This technology was developed by AGH University of Science and Technology in Kraków, Poland [Krehlik *et al.*, 2016]. It allows stabilization of the propagation delay of the fibre link using two variable and matched electronic delay lines, as depicted in Fig. 48. An ELSTAB-based time and frequency transfer link is composed of local and remote modules interconnected by a fibre link. Such a link may include several specialized bi-directional optical amplifiers necessary to compensate losses due to the fibre attenuation, and thus extend the link range from up to 100 km (without amplifiers) to 700–800 km (with amplifiers) [Krehlik *et al.*, 2016].

The link delay stabilization employs a bi-directional transmission of an intensity-modulated light within a single fibre. In the local module, the input frequency signal (10 MHz in standard configuration) and time signal (1 PPS or 100 PPS) are combined in the coder (see Fig. 48). The integrated signal passes the variable delay block, then is converted to the optical domain, and feeds the fibre through the optical multiplexer. In the remote module, the signal is retransmitted back to the local module where, after passing the second variable delay block and the decoder, it is then phase-compared with the forward-direction signal. Its output tunes the delay lines to make the round-trip propagation delay constant. This simultaneously makes the delay of the time/frequency signals delivered to the output terminals of the remote module constant, which is possible because the delay fluctuations are nearly the same in both directions of the same optical fibre. The variable delay blocks consist of continuously tuned (analogue) delay lines, covering 100 ns of compensation, and switched (digital) delay modules providing additional 2000 ns of compensation. The combined compensation range of 2 100 ns is fully seamless in operation. Such tuning range is adequate for compensating delay fluctuations caused by seasonal temperature variations over distances exceeding 1 000 km. The optical bandwidth used by the system corresponds to two ITU channels separated by 100 GHz.

FIGURE 48

Simplified block diagram of the ELSTAB system




The propagation delay introduced by the link is not only constant and stable but can also be accurately determined performing the measurement of the round-trip delay and the asymmetry caused by the chromatic dispersion of the fibre. Both measurements can be performed using the PPS REF(erence) and PPS RET(urned) terminals available at the local module [Sliwczynski *et al.*, 2019] (respectively PPS IN and PPS RET in Fig. 48). If required, the delay between the input and output PPS signals (respectively PPS IN and PPS OUT in Fig. 48) can be aligned with 1 ps resolution by setting the appropriate value in the advancer located in the remote module.

The ELSTAB technology can provide time distribution accuracy of a few to several dozens of picoseconds depending on the particular link length and the number of concatenated links involved (see, e.g. [Sliwczynski *et al.*, 2019] for calibration campaign results). The frequency distribution performance (ADEV) is provided at a level of about 1.5×10^{-13} at 1 s averaging, improving to a few times 10^{-17} at one day. As an example, for the ELSTAB system consisting of 11 amplifiers and having total length of 615 km, described in [Krehlik *et al.*, 2016], the reported performance is depicted in Fig. 49.

FIGURE 49

Relative frequency instability measured at the remote output of ELSTAB links: 615 km with 11 amplifiers (left) and 60 km with no amplifiers (right)



A dispersion insensitive version of the ELSTAB system is also available in order to simplify the link calibration. In this solution, the asymmetry caused by fibre's chromatic dispersion is mitigated by using additional optical path, increasing the occupied optical bandwidth to three ITU channels [Sliwczynski *et al.*, 2024]. The performance of this dispersion insensitive version is only slightly worse compared to the standard one: the ADEV is degraded to $5.5 \cdot 10^{-17}$ from 3.5×10^{-17} after 10^4 s averaging for 120 km long link with one amplifier.

The ELSTAB technology is used worldwide by many National Metrological Institutes and other time laboratories, among other places, in USA, Germany, United Kingdom, Brazil, South Korea and Poland. It is also considered for applications in telecommunication systems [Sliwczynski *et al.*, 2020].

6.4.2 Optical carrier-based

6.4.2.1 Ultra-stable optical frequency transfer

Ultra-stable optical frequency transfer refers to the family of techniques used to convey the frequency of an ultra-stable laser over long-haul, single-mode optical fibres. In these schemes, it is the frequency of the optical carrier itself – not an embedded modulation – that is preserved and disseminated. Such transfers rely primarily on heterodyne detection methods (see § 4.7).

This subsection first outlines the physical origins of noise and fluctuations introduced by fibre propagation and clarifies the nature of the active compensation applied to mitigate them. It then reviews the principal link architectures and optical amplification strategies in current use (without attempting an exhaustive survey). A comprehensive review covering optical frequency transfer over both fibre and free-space links was recently published [Caldwell *et al.*, 2025]. This work provides an excellent, carefully documented, and nearly exhaustive treatment of the subject, making it a valuable reference for interested readers. Finally, the typical performance levels and achievable ranges demonstrated in state-of-the-art systems are summarized.

6.4.2.2 Delay and delay variations

The one-way propagation delay of light in a fibre of physical length L and refractive index n is:

$$\tau = \frac{nL}{c} \quad (6-7)$$

where c is the speed of light in vacuum. In practice, both the refractive index and the physical length of the fibre are not constant but vary over time, leading to delay fluctuations. The key contributions can be summarized as follows:

The *optical length* nL varies with time. The dominant contribution arises from fluctuations of the refractive index n , while variations of the physical length L provide a second-order term. These fluctuations are primarily driven by acoustic perturbations on short time scales and by thermal effects on intermediate and long time scales. The principal coefficients characterising these dependencies are [Cohen and Fleming, 1979; Stowe *et al.*, 1982; Imaoka and Kihara, 1992; Froggatt and Moore, 1998]:

$$\begin{aligned} \frac{dn}{dT} &\simeq 10^{-5} \text{ K}^{-1} && \text{(thermo-optic coefficient),} \\ \frac{dn}{d\epsilon} &\simeq 0.2 && \text{(strain-optic coefficient),} \\ \frac{1}{L} \frac{dL}{dT} &\simeq 5 \times 10^{-7} \text{ K}^{-1} && \text{(thermal expansion),} \\ \Delta n = n_{\text{slow}} - n_{\text{fast}} &\simeq 10^{-7} && \text{(fibre birefringence),} \end{aligned} \quad (6-8)$$

and where longitudinal strain is defined as $\epsilon = dL/L$. These parameters quantify how temperature variations, mechanical strain, and environmental perturbations modulate the effective propagation delay. An additional term arises from the fibre's birefringence, since the counter-propagating optical waves traverse the link with mutually orthogonal polarizations.

Phase fluctuations $\varphi(t)$ in the optical carrier directly induce frequency fluctuations,

$$f(t) = \frac{1}{2\pi} \frac{d\varphi(t)}{dt} \quad (6-9)$$

which can equivalently be expressed as a Doppler shift. If the optical path length varies with an effective velocity v , the received frequency becomes:

$$v_{\text{RX}} = v_{\text{TX}} \left(1 - \frac{v}{c}\right) \quad (6-10)$$

since variations in propagation delay τ_D correspond to

$$\frac{d\tau_D}{dt} = \frac{v}{c} \quad (6-11)$$

The mechanical coupling of the fibre to the ground renders it sensitive, via the Doppler effect, to ground vibrations. At short measurement intervals, ranging from approximately 1 ms to several tens of seconds, the PSD is dominated by Doppler-induced fluctuations and multiple scattering processes. For longer integration times, typically beyond 100–1 000 s, environmental effects such as temperature and humidity variations become the primary contributors to phase noise or phase wandering.


6.4.2.3 Free-running phase noise

It is useful to describe the phase noise of a free-running fibre link, i.e. the uncompensated link. The single-sideband phase noise, PSD, can be expressed as:

$$S_{\phi, \text{free}}(f) = \int_0^L \delta \varphi \left(z, t - \tau + \frac{z}{v} \right) dz \quad (6-12)$$

where L is the one-way physical length of the fibre, $v = c/n$ is the group velocity, and $\delta\varphi(z, t)$ is the local phase noise density (in rad^2/Hz per unit length). This formulation emphasises that the fibre does not contribute uniformly along its span: $\delta\varphi$ is position dependent because the environmental excitation spectrum varies with location. For instance, fibre segments in urban corridors or near railways and highways experience stronger acoustical and mechanical perturbations than rural sections, whereas submarine fibres exhibit markedly reduced noise [Fujieda *et al.*, 2010; Clivati *et al.*, 2018; Marra *et al.*, 2022; Tønnes, 2022; Husmann *et al.*, 2024; Zhou *et al.*, 2025]. Examples are given in Fig. 50.

FIGURE 50



Note to Fig. 50: Literature reports value ranging 10^{-4} to 10^3 $\text{rad}^2/\text{Hz}/\text{km}$ between submarine link and link in urban area, as in Tokyo, London or Shangai (from [Clivati *et al.*, 2018]).

One observes from equation (6-12) that the integrated phase fluctuations grow quadratically with the total link length, i.e. $S_{\varphi, \text{free}} \propto L^2$. This scaling should be regarded only as a rule of thumb, since the linear noise density is not constant along the fibre span.

In ultra-stable optical frequency transfer, several important points follow. The phase noise spectrum is typically dominated by a combination of white phase noise and white frequency noise, particularly for Fourier frequencies below about 10 Hz [Tønnes, 2022]. The modulation index, defined as the peak-to-peak amplitude of the phase variation, can become much greater than unity ($\gg 1$) on long-haul fibre links, indicating large excursions in optical phase.

Active compensation systems act to cancel the fibre-induced phase fluctuations, such that:

$$\varphi_c(t) = -\Phi(t) \quad (6-13)$$

where $\Phi(t)$ represents the one-way fibre's accumulated phase noise and $\varphi_c(t)$ is the applied correction. Within a closed-loop configuration, the counter-propagating optical waves are phase-aligned at the detection point, but only up to an unknown static phase offset. This means the compensation stabilises the variation of the link delay but does not reveal its absolute value. The accumulated time error, $x(t)$, is obtained by integrating the relative frequency fluctuations over time:

$$x(t) = \int_0^t \frac{\delta\nu(t')}{\nu_0} dt' \quad (6-14)$$

where $\delta\nu(t') = \nu_{\text{RX}}(t') - \nu_{\text{TX}}$ and ν_0 is the nominal optical frequency.

NOTE – For terrestrial fibre links spanning hundreds of kilometres, typical uncompensated delay variations can range from a few picoseconds over short time scales (seconds to minutes) up to several nanoseconds over diurnal cycles. These correspond to phase excursions of many tens to hundreds of radians at optical frequencies (~ 200 THz).


6.4.2.4 Stabilization techniques

Underlying hypothesis

The following discussion assumes that the transmission path is reciprocal, meaning that the linear phase noise accumulated in the forward direction is equal to that accumulated in the reverse direction. Specifically, it is assumed that the servo loop effectively divides the round-trip phase noise by two, thereby compensating only for reciprocal noise. Consequently, any non-reciprocal noise or asymmetrical phase shifts cannot be corrected by this stabilization mechanism. All results and analyses presented here are therefore valid only under the assumption of path reciprocity. Deviations from this assumption lead to unmitigated statistical and systematic errors.

Active noise cancellation

The basic principle of fibre-link stabilization is illustrated schematically in Fig. 51.



The optical path is arranged as a strongly imbalanced Michelson interferometer: one arm contains the long-haul fibre under test, while the short reference arm is kept local. Faraday mirrors, rather than simple mirrors, are used to automatically align the polarization states upon reflection, maximising interference contrast and reducing polarization-induced fading [Pistoni and Martinelli, 1991]. Acousto-optic modulators (AOMs) are inserted to frequency-shift the optical signals, enabling heterodyne detection of the returning light. The beat note is compared to a stable reference in a phase-locked loop (PLL), which drives the AOM frequency to cancel phase fluctuations induced by the fibre. This configuration provides a direct measurement of the round-trip phase noise, which is used for active noise cancellation on the forward path. The details of heterodyne detection, tracking oscillators, and narrow-bandwidth detection strategies have already been described in § 4.1 and are directly applicable here. Using Faraday mirrors at the remote end ensures automatic realignment of the polarization.

Under these conditions, the detected signal contains both the fibre-induced round-trip phase noise and the laser's intrinsic noise auto-correlation, expressed as:

$$\Phi_{\text{det}}(t) = \Phi_{\text{RT}}(t) + \Phi_{\text{Laser}}(t) - \Phi_{\text{Laser}}(t - 2\tau) \quad (6-15)$$

where $\Phi_{\text{RT}}(t)$ represents the fibre's round-trip noise and $\Phi_{\text{Laser}}(t)$ is the laser phase at the emission time. The subtraction term $\Phi_{\text{Laser}}(t - 2\tau)$ reflects the fact that the same laser field is sampled again after a full round trip, producing an auto-correlation of the laser phase noise separated by twice the one-way delay τ . The correction is implemented through frequency shifting of the optical signal. By monitoring the correction signal applied to the AOMs, one effectively records the free-running fibre noise with the opposite sign. The compensation process operates continuously. The feedback is delayed by at least twice the one-way propagation time. This represents an unusual situation in which the optical delay significantly exceeds the electronic delay, which is typically the case for links longer than a few kilometres.

According to servo loop theory, the feedback loop gain is set to be smaller than $1/(4\tau)$. Under this limit, determined by the propagation speed of light in the medium, and assuming no other technical limitations dominate, the phase noise power spectral density of the compensated link is given by (Newbury *et al.*, 2007):

$$S_{\Phi,c}(f) = \frac{(2\pi f\tau)^2}{3} S_{\Phi,\text{free}}(f) \quad (6-16)$$

where $S_{\Phi,\text{free}}(f)$ is the PSD of the free-running (uncompensated) link, τ is the one-way propagation delay, and f is the Fourier frequency.

FIGURE 52

End-to-end phase noise measurement on an 86-km link. The servo peak at $1/4\tau$ is clearly seen (from [Jiang *et al.*, 2008])





Figure 52 shows the experimentally recorded phase noise power spectral density of a folded 86-km fibre link in an urban environment. The folded configuration enables simultaneous observation of both the free-running fibre noise and the compensated noise at the local end. The red trace corresponds to the uncompensated fibre noise, while the black trace represents the noise after compensation. The phase noise exhibits a rapid increase, scaling approximately as f^{-2} between 10 Hz and 100 Hz, and as f^{-3} from 100 Hz up to a few kHz. Following the servo bump at $1/4\tau$, the compensated link displays an almost flat PSD, characteristic of white phase noise in a coherent optical fibre link.

End-to-end measurements, illustrated in Fig. 53

After the photodiode, the signal passes through the front-end electronics and a narrow-band filter, often followed by a tracking oscillator (see § 4.1). For the evaluation of mid- and long-term performance, the filtered beat notes are recorded using a dead-time-free frequency counter, typically operated with a one-second gate time.

Two main counter operating modes are commonly employed. In Π -mode, the counter measures the full bandwidth determined by its sampling rate, which for most modern counters ranges from approximately 1 kHz to several MHz. In Λ -mode, the counter applies a weighted averaging function with a triangular shape (resembling the Greek letter Λ); this mode effectively filters out noise above $1/(2t_g)$, where t_g is the gate time.


FIGURE 53
End-to-end measurement



Note to Fig. 53: Laser light is transmitted from the local end to the remote end via the primary (top) fibre. A second fibre link (bottom) then returns the light to the local end. By generating a beat note between the original local laser and the twice-transferred light, the stability and accuracy of the link can be evaluated.

In the literature, frequency stability is reported either as the Allan deviation (ADEV) or, in some cases, as the modified Allan deviation (MDEV) (see Chapter 2). The impact of the counter's operating mode and the choice of data processing on the reported stability is illustrated in Fig. 54.

FIGURE 54
End-to-end stability of uncompensated and compensated optical fibre link, for counters operated in Π - and Λ -mode. Modified Allan deviation (MDEV) and Allan Deviations (ADEV) are computed from the same data set




One can clearly observe the effect of the reporting mode: in Π -mode, the measurement bandwidth is 500 Hz, whereas in Λ -mode it is 0.5 Hz. The narrower bandwidth in Λ -mode reduces the noise power, increases the signal-to-noise ratio, and improves the short-term stability. For pure white phase noise, a reduction by a factor of 10 is expected, which is nearly achieved in Fig. 54. As also expected for white phase noise, the Allan deviation (ADEV) decreases proportionally to τ^{-1} , while the modified Allan deviation (MDEV) follows a $\tau^{-3/2}$ dependence.

Two-way techniques

An alternative to active compensation schemes is the use of passive cancellation techniques. Inspired by satellite methods introduced by Hansson in the late 1980s, one can arrange for reciprocal transmission of light along the same fibre [Hansson, 1989; Lopez *et al.*, 2010; Calosso *et al.*, 2015]. The experimental setup is depicted in Fig. 55. At both the local and remote ends are ultra-stable lasers and Michelson interferometers. The light from each laser propagates in opposite directions along the fibre. At each end, the beat note between the local laser and the transmitted light is detected, filtered in a narrow band, tracked, and counted. The data are then exchanged during post-processing.

FIGURE 55

Schematic of a passive optical frequency transfer link. AOM: acousto-optic modulator. FM: Faraday mirror. PC: Polarization controller



HDB-31-55

At the photodiodes, the phases of the two beat-notes Φ_1 and Φ_2 are given by:

$$\Phi_1 = \Phi_{L1} - (\Phi_{L2} + \Phi_{21}) \quad \text{and} \quad \Phi_2 = \Phi_{L2} - (\Phi_{L1} + \Phi_{12}). \quad (6-17)$$

where Φ_{L1} and Φ_{L2} denotes the optical phase of the laser 1 and 2, Φ_{12} and Φ_{21} represents the accumulated phase noise on the fibre from west to east and east to west respectively. One can compute the sum and difference of these observables. The half-difference reads:

$$\frac{\Phi_1 - \Phi_2}{2} = (\Phi_{L1} - \Phi_{L2}) + \frac{\Phi_{12} - \Phi_{21}}{2} \quad (6-18)$$

which corresponds to the phase comparison between the two distant ultra-stable lasers.

The half-sum is given by:

$$\frac{\Phi_1 + \Phi_2}{2} = \frac{\Phi_{12} + \Phi_{21}}{2} \quad (6-19)$$

which represents the one-way fibre noise under the assumption of reciprocity. Thus, from these two measurements, the system allows simultaneous evaluation of both the laser comparison and the free-running fibre noise.


In such a two-way setup, the light travels only half the distance compared to an active compensation system. In other words, there is no true round-trip; the two one-way signals are later combined in post-processing to form an effective two-way measurement. As a result, the phase noise power spectral density rejection factor is four times greater and is given by:

$$S_{\Phi,c} = \frac{(2\pi f\tau)^2}{12} S_{\Phi,\text{free}} \quad (6-20)$$

Without entering into details, it is also possible to implement a unidirectional system. A simplified sketch is displayed in Fig. 56. In this configuration, no Michelson interferometer is required; instead, the setup relies on two fibre couplers and polarization controllers to optimize the beat notes, along with the same set of acousto-optic modulators (AOMs) to perform heterodyne detection and resolve multiple-path ambiguities.

FIGURE 56

Sketch of a passive optical frequency transfer link, so called uni-directional two-way




AOM: acousto-optic modulator. PC: Polarization controller.

Note to Fig. 56: The data link to perform the comparison in post-processing is not shown for sake of clarity.

Each photodiode detects two beat notes: one arising from the interference between the local laser and the light transmitted along a given one-way path, and another corresponding to the beat note of the local laser with the light propagating along the round-trip path. Passive cancellation of propagation noise can be achieved by exchanging data, as in the previously described scheme. In addition, because the light from each laser returns to its point of origin after completing the loop, this configuration allows the propagation noise to be cancelled without the need for data exchange. Since the light propagates in two distinct fibres, the noise correlation between the two signals is reduced, leading to significantly lower noise rejection. Detailed descriptions of such implementations can be found in [Bercy *et al.*, 2014].

FIGURE 57

Stability performance comparison between active and passive stabilization scheme for a 92 km long link



HDB-31-57

ANC: Active Noise Cancellation. TWB: Two-Way Bi-directional. TWU: Two-Way Uni-directional [Bercy *et al.*, 2014].

Figure 57 shows a comparison of two passive and active noise cancellation techniques implemented on the same fibre setup. The unidirectional two-way configuration illustrates the impact of reduced correlation between the fibre noise in the two fibres, even when they are housed within the same cable. The plot demonstrates excellent agreement between the experimental results and theoretical predictions.


Active and passive noise compensation techniques can be effectively combined to create so-called hybrid link architectures. Such configurations may be employed to compare two links at the crossover point between networks, to simplify the complexity of end-to-end measurements, or to investigate the fundamental performance limits of fibre links [Xu *et al.*, 2018; Xu *et al.*, 2021].

6.4.2.5 Cascaded links

A fundamental limitation in long-distance optical frequency transfer arises from the finite speed of light in the fibre, which sets the position of the $1/\tau$ servo peak. Since the phase noise scales at least as f^2 , and in some cases as f^3 , this limitation not only constrains the bandwidth of the servo loop but also determines the total amount of noise that must be compensated. Consequently, the longer the fibre link, the more noise needs to be corrected, and the compensation becomes increasingly inefficient.

FIGURE 58


Principle of cascaded link. A long-haul link is subdivided into a sum of smaller spans



A practical approach to mitigate this problem is the implementation of a daisy-chain, or so-called cascaded link, in which a long fibre span is subdivided into shorter segments, typically ranging from 200 to 600 km each. This strategy reduces the accumulated phase noise in each segment and shortens the effective propagation time, improving the efficiency of noise suppression, see Fig. 58. Cascaded links have been successfully demonstrated in both the RF and optical domains, and they are compatible with C- and L-band operation [Fujieda *et al.*, 2010; Lopez *et al.*, 2010; Husmann *et al.*, 2021]. Moreover, the technique can be implemented alongside standard amplification schemes such as EDFAs, Raman amplifiers, and fibre Brillouin amplifiers, and even using on-chip interferometers for integrated setups [Koke *et al.*, 2019; Akatsuka *et al.*, 2020]. The concept can be extended to multiple input/output, by designing a multi-branch Michelson interferometer [Cantin *et al.*, 2021].

FIGURE 59


Schematic of a cascaded optical frequency transfer link



AOM: acousto-optic modulator; FM: Faraday mirror; PC: polarization controller.

Figure 59 presents a simplified schematic of a cascaded link. A cascaded link typically includes at least two Michelson interferometers within a so-called repeater laser station. One interferometer is used to phase-lock the remote laser to the incoming light from span N , while the other phase-locks the light for span $N+1$. The laser locked to the incoming signal emits light backwards, allowing the stabilization of the link to be performed at the local end. The polarization of the light exiting the link must be adjusted to maximize the beat note. Typically, a small portion of the remote laser's output is preserved for user access. To ensure efficient operation, such systems generally incorporate remote control and monitoring, automatic operation, and automatic polarization control. Thanks to the heterodyne detection technique, these systems can operate with very low input power, down to approximately -60 dBm.

FIGURE 60

PSD of phase noise for a 4-spans cascaded link (from [Chiodo *et al.*, 2015])


HDB-31-60

Due to the cascaded nature of the operation, the residual phase noise from span N is transferred as input noise to span $N+1$. Figure 60 presents the experimental observation of this effect. This phenomenon imposes a practical limit on the number of repetitions that can be implemented. Simulations indicate that a maximum reach of approximately 2 200 km is feasible (Tønnes, 2022). In practice, such distances have already been achieved – for example, on the link connecting Paris (France) to Torino (Italy) [Clivati *et al.*, 2022].

These technologies have been transferred to industry and are now commercially available as off-the-shelf components. Typical performance for an industrial system is shown in Fig. 61.

FIGURE 61

Relative frequency deviation (Allan deviation and Modified Allan Deviation) at the output of the first industrial grade repeater laser station for a 2-spans cascaded link of 680 km [Guilloou-Camargo *et al.*, 2018]



HDB-31-61

6.4.2.6 Amplification techniques

Optical amplification plays obviously a critical role in long-haul coherent optical frequency transfer and cascaded link systems, as it compensates -partly- for fibre attenuation. Among the most widely used technologies are the erbium-doped fibre amplifier (EDFA), the semiconductor optical amplifier (SOA) or booster optical amplifier (BOA), Raman amplifiers (RA), and Fibre Brillouin Amplifiers (FBA). Each of these technologies offers distinct performance characteristics in terms of achievable gain, noise figure, and operational bandwidth. A comparison of their typical performance is provided in Table 19, which summarizes indicative values reported in the literature and observed in practice.

TABLE 19

Comparison of common optical amplification technologies

Technology	Gain (dB)	Noise figure (dB)	Bandwidth (GHz)
EDFA (C-band)	10-20	4-6	3 750-5 000
EDFA (L-band)	10-20	4-6	10 000
SOA/BOA	10-20	6-8	6 250-12 500
RA	15-25	4-7	12 500
FBA	40-60	< 2	0.01

EDFA: Erbium-Doped Fibre Amplifier

SOA/BOA: Semiconductor/Booster Optical Amplifier / Booster Optical Amplifier.

RA: Raman Amplifier

FBA: Fibre Brillouin Amplifier

Note to Table 19: Values are indicative and may vary depending on implementation.


Interested readers may consult several standard references for detailed descriptions of *classical* (telecom) optical amplifier architectures, pumping schemes, and the associated energy-level diagrams. Erbium-doped fibre amplifiers are extensively treated in Desurvire's classic monograph [Desurvire, 1994], as well as in Keiser's textbook on fibre-optic communications [Keiser, 2015]. Distributed Raman amplification is well covered in Ter-Mikirtychev's comprehensive text [Ter-Mikirtychev, 2014]. Details on FBA designs are reported in [Grosche *et al.*, 2009].

These amplification schemes differ not only in performance but also in their practical implementation. Phase preservation (low added phase noise) and bidirectionality are crucial. Solutions reported in the literature include fibre Brillouin amplification (using the fibre itself as the gain medium, efficient and bi-directional for small pump powers) [Grosche *et al.*, 2009], in-line Raman amplification (distributed gain) (Clivati *et al.*, 2015), and specialized bidirectional EDFA schemes or repeater laser stations that recover phase without breaking the link [Lopez *et al.*, 2010]. CLONETS reports discuss these tradeoffs [CLONETS Consortium, 2018].

EDFAs remain the standard choice for long-haul telecommunications despite their reduced gain thanks to their compatibility with C- and L-band transmission. SOAs and BOAs, while offering lower gain and higher noise figures, are valued for their compactness and ease of integration into photonic circuits [Pinkert *et al.*, 2015]. Raman amplifiers provide distributed gain directly within the transmission fibre, offering improved gain and noise performance [Clivati *et al.*, 2015]. In all these cases, due to the bidirectional operation and the necessity to avoid self-oscillation, the achievable gain with bi-directional version of these amplifiers is usually between 10 to 20 dB, much below the standard of regular telecom techniques (20-25 dB).

Brillouin amplifiers, with their exceptionally low noise figures, narrow gain bandwidths, and high gain, are best suited for ultra-stable optical frequency transfer where spectral purity is paramount [Grosche *et al.*, 2009]. Figure 62 reports the achievable gain for EDFA and FBA. Together, these technologies provide a range of options to optimize performance for specific link configurations and operational requirements.

FIGURE 62



6.4.2.7 Compatibility with parallel data traffic

Erbium-doped fibre amplifiers (EDFAs), semiconductor optical amplifiers (SOAs), and bulk optical amplifiers (BOAs) provide localized, lumped gain. They can therefore be deployed in-field on dedicated wavelength channels using standard wavelength-division multiplexing (WDM) filtering without significant interaction with adjacent services.

In contrast, Raman amplification (RA) and fibre Brillouin amplification (FBA) employ the transmission fibre itself as the gain medium. When such distributed amplification is introduced on an operational link, even if the metrological channel is spectrally isolated by WDM components, some coupling remains between the optical gain experienced by data traffic and that applied to the metrological signal. Furthermore, FBA offers gain only within a very narrow spectral band. Because of this limited bandwidth, fibre Brillouin amplifiers are generally considered incompatible with simultaneous high-capacity data traffic on parallel channels.

6.4.2.8 Performances

Stability


The stability of coherent optical frequency transfer links typically ranges from 10^{-14} to 10^{-17} at a one-second integration time, measured within a 1 Hz bandwidth, for link length ranging from 50 km to 1 000 km. These links generally average down as fast as $\tau^{-3/2}$ for at least 100 s to a few thousand seconds, after which the improvement slows. The precise behavior depends on factors such as the link configuration (fully bidirectional or quasi-bidirectional), the launched optical power, the link length, the characteristics of the free running noise, and the overall quality of the deployed instrumentation. In practice, the most meaningful metric for instrumentation quality is the ratio between the free-running noise and the compensated noise. For this reason, it is highly desirable for authors to publish both curves when reporting their implementations. Over longer time scales, diurnal perturbations, seasonal variations, and annual terms become apparent [Tønnes *et al.*, 2022]. Such analysis requires also a precise understanding of the effect of missing data in the time series recordings. It is likely that further analysis and extended data acquisition will reveal additional features of geophysical or environmental interest.

Accuracy

The accuracy of these techniques, both active and passive, has been explored by several groups [Cantin *et al.*, 2021; Koke *et al.*, 2021;]. An accuracy budget was reported for experiments on fibre spools, where the effect of polarization mode dispersion (PMD) could be quantified and identified as a limitation at the level of 1.2×10^{-20} [Xu *et al.*, 2021]. Figure 63 illustrates the free-running link, the two-way compensated signal, the post-processed two-way compensation, and the noise floor of the system, highlighting the effects of temperature, humidity, and polarization. The plot highlights the limiting effect of the PMD and the contribution of the interferometer sensitivity to temperature.

For in-field implementations, particularly in the context of optical clock comparisons over a fibre network, additional terms must be considered, such as time mismatches between recordings and other sources of uncertainty. Interested readers may consult [Tampellini *et al.*, 2017; Tønnes, 2022]. Comparisons between optical frequency transfer in fibre and in free space further reinforce confidence in the accuracy budget, demonstrating uncertainties well below 10^{-19} [Beloy *et al.*, 2021; Shen *et al.*, 2022].

FIGURE 63

Evidence of interferometer noise and PMD noise limitation [Xu *et al.*, 2021]

HDB-31-63

Fundamental limits


The fundamental limits of optical frequency transfer in fibre are determined by the finite speed of light in the medium and the requirement of path reciprocity. The finite speed of light constrains the noise rejection bandwidth and leads to accumulation of noise over longer integration times [Newbury *et al.*, 2007]. Deviations from perfect reciprocity arise from small frequency shifts between the two counter-propagating beams; the resulting effect is extremely small, below 10^{-23} . Polarization effects are the dominant contribution at the level of 10^{-19} – 10^{-20} [Xu *et al.*, 2021]. Ultimately Sagnac effect, as a non-reciprocal terms, set the second fundamental limit of optical fibre link. Technical limitations are primarily associated with the sensitivity of the reference arm of the Michelson interferometer to temperature fluctuations. In a well-designed interferometer with proper temperature stabilization, this effect can be reduced to below 10^{-21} .

Ranges

The operational range of optical fibre links has steadily expanded over the past decades. Starting from just a few kilometres in the late 1990s, these links experienced rapid growth throughout the 2000s, and this expansion accelerated between 2010 and 2020, in parallel with the widespread deployment of optical fibre networks across Europe. By the late 2020s, significant infrastructure was established in China, further extending the global reach of ultra-stable optical frequency dissemination. Figure 64 illustrates the reported ranges of optical fibre links over this period.

FIGURE 64

Reported coherent fibre links range over time in the literature [Beauvoir *et al.*, 1998; Terra *et al.*, 2010; Predehl *et al.*, 2012; Chiodo *et al.*, 2015; Schioppo *et al.*, 2022; Lindvall *et al.*, 2025]



HDB-31-64

While individual terrestrial fibre links are typically limited to approximately 2 000–3 000 km (excluding subsea link, for which the phase noise PSD is considerably lower), there is no fundamental barrier to concatenating multiple links. This allows the realization of continental-scale networks.

The next frontier lies in establishing long-haul submarine links capable of connecting the five continents and remote locations, such as Tahiti in the central Pacific. Such links are critical for the synchronization and syntonization of geodetic observatories, which underpin the realization of precise global reference frames and advance fundamental astronomy.

6.4.2.9 Fibre network for time and frequency dissemination

From fibre links to fibre network

Extending ultra stable optical frequency dissemination from individual fibre links to full-scale fibre networks introduces a set of unique technical, organizational, and economic challenges.

Availability and cost of fibre

Long-term access to optical fibre is essential but costly. The main challenge lies in establishing a sustainable economic model: the effective cost per user is determined by the total expense of maintaining the fibre network divided by the number of beneficiaries. Minimizing rental costs, for instance through data sharing, and engaging the widest possible scientific and industrial community, is therefore critical to ensure network viability.

6.4.2.10 Dedicated frequency channels

In France, this challenge is addressed through the deployment of dedicated frequency channels, allowing parallel transmission of ultra-stable signals and conventional data traffic on the same fibre using dense wavelength-division multiplexing (DWDM). Such networks are typically hosted on National Research and Education Network (NREN) backbones. Implementing these systems requires modifications to the telecom

infrastructure, including the installation of low-noise, bi-directional optical amplifiers within NREN shelters – a complex and labour-intensive task.

However existing telecommunications networks cannot always be adapted for critical applications beyond conventional data transmission, and not only due to the unidirectionality of long-haul fibre links. But emerging multicore optical fibres (MCFs) offer a transformative opportunity: they enable the coexistence of advanced applications with standard data traffic, paving the way for truly multi-functional networks. This approach is particularly promising for long-haul submarine cables. Recent work has proposed and demonstrated the faithful transfer of ultra-stable optical signals through 25 km deployed MCF in a manner fully compatible with the unidirectional architecture of long-haul systems, achieving a fractional frequency instability of 3×10^{-19} at 10 000 [Hoghooghi *et al.*, 2025].

Technology maturation and system vision

Achieving a fully operational network demands not only achieving maturity of technology and knowledge transfer to reach Technology Readiness Level (TRL) 9, but also a comprehensive system vision. This encompasses the capacity for production, installation, operation, and both preventive and corrective maintenance of the network.

Network supervision

Mandatory network supervision is essential, including integration into a Network Operating Centre (NOC) to monitor performance, detect anomalies, and ensure reliability over long-term operation.

Data management and community engagement

Finally, an effective fibre network requires robust capabilities for data acquisition, storage, processing, reduction, and archiving, in compliance with FAIR principles (Findable, Accessible, Interoperable, Reusable) [WILKINSON *et al.*, 2016]. Up-to-date documentation, live monitoring, and organized archives are critical. Beyond infrastructure, fostering community adoption through active management and support is necessary to ensure that the network delivers scientific and industrial value and is widely utilized.

6.4.3 Building scientific and societal infrastructure

The design and deployment of fibre-based optical networks must encompass a sufficiently broad range of scientific domains to ensure that the infrastructure delivers value to a critical mass of users. One of the primary beneficiaries of such networks is the domain of reference frames and precision measurements. Fibre networks support chronometric geodesy, providing tools to monitor sea level rise induced by global warming, and enable precision measurements for both applied and fundamental research, ranging from tests of variations in fundamental constants to applications in cosmology and astronomy. Precision timing, increasingly relevant to accelerator physics and high-energy experiments, also benefits from optical fibre networks, paving the way for future tests of general relativity through multi-messenger astronomy when major facilities are interconnected.

Fibre infrastructures also provide a natural platform for advancing quantum telecommunication technologies, including continuous-variable quantum key distribution (CV-QKD) and twin-field QKD (TF-QKD), supporting secure communications at unprecedented scales. In addition, the networks foster progress in photonic development, encompassing laser technology and high-speed optical communications, driving both scientific innovation and industrial applications in next-generation optical technologies. Another key application lies in environmental sensing, where fibre networks enable novel measurements for Earth, ocean, and atmospheric sciences, including high-precision temperature monitoring and geophysical observations, which will be further discussed in the concluding section.

Beyond their scientific utility, these infrastructures carry significant societal benefits. Optical fibre dissemination provides a robust backup to GNSS-based systems, enhancing the resilience of critical infrastructure for positioning, navigation, and timing (PNT). By combining broad scientific applicability with tangible societal impact, fibre networks constitute a transformative infrastructure for the 21st century.

6.4.4 Implementation of optical fibre networks

The most prominent case is in Europe. Since the first optical clock comparison in 2015 between SYRTE and PTB, several long-haul fibre links have been cross-connected [Lisdat *et al.*, 2016]. By 2025, France, Italy, the UK, Germany, Poland, and Switzerland (notably at CERN) are interconnected [Lindvall *et al.*, 2025]. The network continues to expand: Belgium, the Netherlands, Australia, the Czech Republic, and Spain are expected to join soon.

In France, the REFIMEVE network covers the entirety of metropolitan France and incorporates over 5 000 km of fibre for RENATER, dedicated to optical frequency transfer. More than 30 research laboratories are currently connected. REFIMEVE was recognized as a National Research Infrastructure in 2021.

The Italian Quantum Backbone (IQB) is a nationwide fibre-optic infrastructure that links the Italian National Metrology Institute (INRiM) with +15 research centres and laboratories throughout Italy, spanning more than 1 800 km. Initially conceived for the distribution of high-accuracy time and frequency references, the IQB has since evolved into a flexible platform that now supports cutting-edge experiments in quantum communication, distributed sensing, and a broad range of advanced photonic applications.

Poland is also advancing long-haul optical frequency links. The primary objective of the National System for Generation and Distribution of the Reference Optical Carrier in Poland is to establish an ultra-stable laser source that can be directly referenced to an optical atomic clock, and to distribute this signal through a dedicated network of fibre-optic links operated by PSNC, over +3 000 km. NLPQT consortium encompass seven research laboratories and qualified industrial partners.

In UK the QSNET consortium is developing a network of atomic and molecular clocks with unprecedented sensitivity for probing potential variations in the fine-structure constant (α) and the proton-to-electron mass ratio (μ). By achieving this level of precision, QSNET aims to impose significantly tighter constraints on a broad spectrum of fundamental and phenomenological theories beyond the Standard Model, as well as on proposed models of dark matter.

The Czech Republic has also established a non-commercial fibre-based infrastructure. Designed as a shared national platform, this network leverages optical fibres to deliver high-precision timing, coherent frequency transfer, and an advanced vibrational sensing capability. By allocating classical and quantum signals to separate wavelength bands, such as the C-band and O-band, the Czech infrastructure minimizes interference and supports multiple concurrent applications over the same fibre lines.

The REFIMEVE network in France and IQB in Italy have established frequency distribution links connecting major cities and a cross-border link, but across Europe, however, significant gaps remain in the continent-wide infrastructure. These gaps arise partly from limited investment in certain countries, but more critically from the absence of a centralized European body responsible for interconnecting the existing national time and frequency networks. GÉANT has initiated dialogue with numerous European National Metrology Institutes (NMIs) to gather their perspectives on the proposed Core Time/Frequency Network (C-TFN). Now GÉANT Association aims to establish a fibre infrastructure of long-haul, cross-border fibre routes for the distribution of time and frequency throughout Europe, fostering transnational connectivity.

If one includes all fibre-based technologies (optical frequency transfer, ELSTAB, WR), Europe is projected to encompass approximately 30 000 km of dedicated or shared with data traffic fibre network in the next five years.

In China, fibre network expansion is also very rapid. In China, the High Precision Ground-Based Time Service System (HP-GBTSS) project has been launched as a major national scientific and technological infrastructure. One of the important tasks of the HP-GBTSS project is to build a 10 000-kilometre-level fibre optic network for coherently transmitting optical frequency signals [Zhou *et al.*, 2025; Zhang *et al.*, 2025].

Together, European and Chinese developments indicate a bifurcated yet convergent path: a maturation of technical infrastructures, policy recognition, and scientific-industrial coordination that promise widespread deployment of ultra-stable optical networks globally.

6.4.5 Outlook: fibre sensing

Coherent networks as distributed environmental sensors

Coherent fibre-optic systems are inherently sensitive to phase fluctuations that remain invisible in simple intensity-based detection. Remarkably, the very corrections applied by digital signal processing (DSP) for carrier recovery can reveal tiny phase anomalies imprinted onto the optical field such as those caused by seismic waves or environmental perturbations. Exploiting the state of polarization (SOP) as an additional sensing axis further enhances this capability [Mazur *et al.*, 2022; Carver and Zhou, 2024]. As a result, recent years have seen a surge of interest from both academic research groups and major manufacturers in turning coherent telecommunication networks into large-scale distributed sensors for environmental and infrastructure monitoring.

For example, the EU project SENSEI (Smart European Networks for Sensing the Environment and Internet quality (SENSEI Consortium 2024)) pursues this vision by developing photonic interferometric techniques – using low-phase-noise lasers and on-chip integration – to repurpose live telecom fibres for high-precision deformation monitoring within the network control layer, enabling applications such as seismic hazard detection and self-healing infrastructures. At the same time, scientific demonstrations have captured seismic events using standard telecom links and coherent detection, illustrating how “dark” or continually active fibres can be harnessed for seismic monitoring at scale [Yan *et al.*, 2019; Guerrier *et al.*, 2025].

This transformational approach positions coherent fibre infrastructure not only as the backbone of high-speed digital communication but also as a continent-spanning sensor array, with profound implications for Earth science, civil protection and network resilience [Hara *et al.*, 1999].

Emerging applications: frequency transfer as a large-scale environmental sensor

As the optical frequency transfer technique is inherently phase-sensitive, the correction applied to suppress fibre-propagation noise does not only reveal the optical phase itself but also encodes environmental perturbations acting on the fibre. Acoustic vibrations, seismic waves, and even slow variations from temperature and humidity leave characteristic imprints on the measured phase. This transforms a technique originally devised for metrological frequency dissemination into a powerful sensing tool [Marra *et al.*, 2022].

In recent years, this capability has been demonstrated both in laboratory and real-field conditions. Seismic signals have been retrieved from terrestrial fibre networks and submarine telecommunication cables, showing remarkable agreement with conventional seismometers and extending coverage to otherwise uninstrumented areas [Marra *et al.*, 2018]. Beyond natural events, first studies now also target anthropogenic noise sources, exploiting the same interferometric signals for environmental monitoring at the city or regional scale [Noe *et al.*, 2023; Donadello *et al.*, 2024].

Major research groups and network operators are actively pursuing this new paradigm, often referred to as “*Network as a Sensor*”, in which already-deployed optical fibre infrastructure provides a continent-scale seismic and environmental observatory [Mazur *et al.*, 2022; Guerrier *et al.*, 2025]. Pilot projects range from controlled laboratory demonstrations to transoceanic cable experiments, with promising prospects for global monitoring of earthquakes, ocean dynamics, and infrastructure integrity. Early work in related areas, such as distributed sensing for submarine power cable maintenance [Hara *et al.*, 1999], anticipated this vision. Today, advances in coherent detection and digital signal processing enable scalable implementations across operational telecommunication networks.

This convergence of metrological frequency transfer with environmental sensing represents an exciting frontier: the same coherent infrastructures that distribute ultrastable optical frequencies for clocks and scientific facilities may, at the same time, provide a new window into the Earth system, with applications spanning geophysics, oceanography, and infrastructure monitoring.

References

AES [2023] AES67-2023: AES standard for audio applications of networks – High-performance streaming audio-over-IP interoperability.

- AKATSUKA T. *et al.*, [2020] *Optical Frequency Distribution Using Laser Repeater Stations with Planar Lightwave Circuits*, Optics Express 28 (7): 9186–97. <https://doi.org/10.1364/OE.383526>
- ALLAN D. and WEISS M. [1980] *Accurate time and frequency transfer during common-view of a GPS satellite* Proc. of the 34th Annual Frequency Control Symposium 334–356.
- AUTOSAR [2020] Time Synchronization Protocol Specification, R20-11
- BEAUVOIR B. *et al.*, [1998] *Transmission of an Optical Frequency Through a 3 Km Long Optical Fibre*, The European Physical Journal D - Atomic, Molecular and Optical Physics 1 (2): 227–29. <https://doi.org/10.1007/s100530050085>
- BELOY K. *et al.*, [2021] *Frequency Ratio Measurements at 18-Digit Accuracy Using an Optical Clock Network*, Nature 591: 238–43. <https://doi.org/10.1038/s41586-021-03253-4>
- BERCY A, STEFANI F, LOPEZ O, CHARDONNET C, POTTIE P.-E and AMY-KLEIN A [2014] *Two-Way Optical Frequency Comparisons at 5×10^{-21} Relative Stability over 100-Km Telecommunication Network Fibres*, Physical Review A 90 (6). <https://doi.org/10.1103/PhysRevA.90.061802>
- BOVEN P [2018] *DWDM stabilized optics for White Rabbit* 2018 European Frequency and Time Forum (EFTF), Turin, Italy, 2018, pp. 213–216, doi: 10.1109/EFTF.2018.8409035.
- BIPM [2021] BIPM guidelines for GNSS calibration, URL: <https://webtai.bipm.org/ftp/pub/tai/publication/gnss-calibration/guidelines/>
- BIPM [2026] BIPM Time Department Data Base: Time dissemination services list and Time signals emissions, URL accessed Jan 2026: <https://webtai.bipm.org/database/dissemination.html>
- CALDWELL E.D, TRIANO T.M, SINCLAIR L.C. [2025] *High-Precision Optical Time and Frequency Transfer*. Advances in Optics and Photonics 17 (2): 375–75. <https://doi.org/10.1364/AOP.545290>
- CALOSSO C.E *et al.*, [2015] *Doppler-Stabilized Fibre Link with 6 dB Noise Improvement Below the Classical Limit* Optics Letters 40 (2): 131. <https://doi.org/10.1364/OL.40.000131>
- CANTIN E, TØNNES M, TARGAT R.L, AMY-KLEIN A, LOPEZ O and POTTIE P.-E [2021] *An Accurate and Robust Metrological Network for Coherent Optical Frequency Dissemination*, New Journal of Physics 23 (5): 053027. <https://doi.org/10.1088/1367-2630/abe79e>
- CARVER C.J and ZHOU X. [2024] *Polarization Sensing of Network Health and Seismic Activity over a Live Terrestrial Fibre-Optic Cable*, Communications Engineering 3 (1): 91. <https://doi.org/10.1038/s44172-024-00237-w>
- CHIODO N. *et al.*, [2015] *Cascaded Optical Fibre Link Using the Internet Network for Remote Clocks Comparison*, Optics Express 23 (26): 33927–37. <https://doi.org/10.1364/OE.23.033927>
- CLIVATI C., BOLOGNINI G., CALONICO D., FARALLI S., MURA A. and LEVI F. [2015] *In-Field Raman Amplification on Coherent Optical Fibre Links for Frequency Metrology*, Optics Express 23 (8): 10604–15. <https://doi.org/10.1364/OE.23.010604>
- CLIVATI C. *et al.*, [2018] *Optical Frequency Transfer over Submarine Fibre Links*. Optica 5 (8): 893–901. <https://doi.org/10.1364/OPTICA.5.000893>
- CLIVATI C. *et al.*, [2022] *Coherent Optical-Fibre Link Across Italy and France*, Physical Review Applied 18 (5): 054009. <https://doi.org/10.1103/PhysRevApplied.18.054009>
- CLONETS Consortium [2018] *Deliverable D1.2 – Current and Emerging Commercial/Research Technologies*. https://www.clonets.eu/CLONETS_Deliverable_D1.2_V1.0c029.pdf
- COHEN L.G. and FLEMING J.W [1979] *Effect of Temperature on Transmission in Lightguides* Bell System Technical Journal, ahead of print. <https://doi.org/10.1002/j.1538-7305.1979.tb03328.x>
- COTA E. G., LIPINSKI M., WLOSTOWSKI T., VAN DER BIJ E, SERRANO J. [2011] *White Rabbit Specification: Draft for comments*, version 2.0 (06-07-2011)

- DANIELSON M. [2014] “Time transfer capabilities in the DTM transmission system”. 2014 European Frequency and Time Forum (EFTF), Neuchatel, Switzerland, 2014, pp. 227-230, doi: 10.1109/EFTF.2014.7331472.
- DANILUK G. [2015] *White Rabbit Calibration Procedure*, version 1.1, November 9, 2015
- Defraigne P., ACHKAR J., COLEMAN M. J., GERTSVOLF M., ICHIKAWA R., LEVINE J., UHRICH, WHIBBERLEY P., WOUTERS M, AND BAUCH A. [2022] *Achieving traceability to UTC through GNSS measurements*. Metrologia 59, 064001.
- DESURVIRE E. [1994] *Erbium-Doped Fibre Amplifiers: Principles and Applications*. John Wiley & Sons.
- DIERIKX E. F. *et al.*, [2016] *White Rabbit Precision Time Protocol on Long-Distance Fibre Links* in IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, vol. 63, no. 7, pp. 945-952, July 2016, doi: 10.1109/TUFFC.2016.2518122.
- DONADELLO S. *et al.*, [2024] *Seismic Monitoring Using the Telecom Fibre Network*, Communications Earth & Environment 5 (1): 1–9. <https://doi.org/10.1038/s43247-024-01338-2>
- EIDSON J [2006]. *Measurement, Control, and Communication Using IEEE 1588* 10.1007/1-84628-251-9.
- FROGGATT M and MOORE J. [1998] *High-Spatial-Resolution Distributed Strain Measurement in Optical Fibre with Rayleigh Scatter*. Applied Optics, ahead of print. <https://doi.org/10.1364/AO.37.001735>
- FUJIEDA M, KUMAGAI M and NAGANO S. [2010] *Coherent Microwave Transfer over a 204-Km Telecom Fibre Link by a Cascaded System*, IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control 57 (1): 168–74. <https://doi.org/10.1109/TUFFC.2010.1394>
- GUILLOU-CAMARGO F. *et al.*, [2018] *First Industrial-Grade Coherent Fibre Link for Optical Frequency Standard Dissemination*, Applied Optics 57 (25): 7203. <https://doi.org/10.1364/AO.57.007203>
- GROSCHÉ G. *et al.*, [2009] *Optical Frequency Transfer via 146 Km Fibre Link with 10^{-19} Relative Accuracy*, Optics Letters 34 (15): 2270. <https://doi.org/10.1364/OL.34.002270>
- GUERRIER S. *et al.*, [2025] *Distributed Coherent Sensing over Deployed Fibres for Network as a Sensor Applications*, Journal of Lightwave Technology 43 (4): 1736–45. <https://doi.org/10.1109/JLT.2024.3498070>
- HANSON D.W. [1989] *Fundamentals of Two-Way Time Transfers by Satellite*, Proceedings of the 43rd Annual Symposium on Frequency Control. <https://doi.org/10.1109/FREQ.1989.68861>
- HARA T. *et al.*, [1999] *Development of Long Range Optical Fibre Sensors for Composite Submarine Power Cable Maintenance*, IEEE Transactions on Power Delivery 14 (1): 23–30. <https://doi.org/10.1109/61.736674>
- HAZELL P, BARON S, GOUSIOU E, HAGMANN G, GILL J, LIPINSKI M, GENOUDQ, GINGOLD T, WLOSTOWSKI T [2024] *Phase stability compliancy testing of a White Rabbit based solution for the LHC RF and Timing Distribution Backbone upgrade* Published 23 December 2024 in Journal of Instrumentation, Volume 19
- HOGHOOGHI N *et al.*, [2025] *Ultrastable Optical Frequency Transfer and Attosecond Timing in Deployed Multicore Fibre*, Optica 12 (6): 894–99. <https://doi.org/https://opg.optica.org/optica/abstract.cfm?uri=optica-12-6-894>
- HUSMANN D *et al.*, [2021] *SI-Traceable Frequency Dissemination at 157206 Nm in a Stabilized Fibre Network with Ring Topology*, Optics Express 29 (16): 24592. <https://doi.org/10.1364/OE.427921>
- HUSMANN D, ZAUGG S and MOREL J [2024] *Detection and Analysis of Anthropogenic Patterns in a Phase-Stabilized Optical Fibre Network*, 2024 European Frequency and Time Forum (EFTF), June, 3–6. <https://doi.org/10.1109/EFTF61992.2024.10722402>
- IEC [2021] IEC 62439-3 Industrial communication networks – High availability automation networks – Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR)

- IEEE [2002] *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, in IEEE Std 1588-2002, vol., no., pp.1-154, 31 Oct. 2002, doi: 10.1109/IEEESTD.2002.94144.
- IEEE [2008] *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, in IEEE Std 1588-2008 (Revision of IEEE Std 1588-2002) - Redline, vol., no., pp. 1-300, 24 July 2008.
- IEEE [2017] C37.238-2017 – IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications
- IEEE [2019] *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems* in IEEE Std 1588-2019 (Revision of IEEE Std 1588-2008), vol., no., pp.1-499, 16 June 2020, doi: 10.1109/IEEESTD.2020.9120376.
- IEEE [2026] *IEEE Standard for Local and Metropolitan Area Networks – Timing and Synchronization for Time-Sensitive Applications*,
- IETF [2010] RFC 5905: *Network Time Protocol Version 4: Protocol and Algorithms Specification*, Internet Engineering Task Force (IETF), June 2010, ISSN: 2070-1721.
- IETF [2020] *RFC 8915: Network Time Security for the Network Time Protocol*, Internet Engineering Task Force (IETF), September 2020, ISSN: 2070-1721
- IMAOKA A and KIHARA M [1992] *Long-Term Propagation Delay Characteristics of Telecommunication Lines*, IEEE Transactions on Instrumentation and Measurement, ahead of print. <https://doi.org/10.1109/19.177337>
- JIANG H, KEFELIAN F, LEMONDE P *et al.*, [2008] *Long-Distance Frequency Transfer over an Urban Fibre Link Using Optical Phase Stabilization*, Journal of the Optical Society of America B 25 (12): 2029–35. <https://doi.org/10.1364/JOSAB.25.002029>
- JOBERT S, RUFFINI S, RODRIGUES S, OUELLETTE M, MONTINI L, MAYER M, GILSON M, FERRANT J-L [2013] *Synchronous Ethernet and IEEE-1588 in Telecoms*, John Wiley & Sons, ISBN 9781848214439, 2013
- JOHNSTON G, RIDELL A, HAUSLER G. [2017] *The International GNSS Service*. In Teunissen, Peter J.G., & Montenbruck, O. (Eds.), Springer, Handbook of Global Navigation Satellite Systems (1st ed., pp 967-982), Cham, Switzerland: Springer International Publishing.
- KAUR N., FRANK F., PINTO J., TURKEY P., POTTIE, P-E. [2022] *A 500-km Cascaded White Rabbit Link for High-Performance Frequency Dissemination* in IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, vol. 69, no. 2, pp. 892-901, Feb. 2022, doi: 10.1109/TUFFC.2021.3134163.
- KAUR N., FRANK F., POTTIE P-E. TURKEY P. [2017] *Time and frequency transfer over a 500 km cascaded White Rabbit network* 2017 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS), Besancon, France, 2017, pp. 86-90, doi: 10.1109/FCS.2017.8088808.
- KEISER G. [2015] *Fibre Optic Communications*. 5th ed. Springer. <https://doi.org/10.1007/978-3-319-42367-8>.
- KOKE S., BENKLER E., KUHL A. and GROSCHE G. [2021] *Validating Frequency Transfer via Interferometric Fibre Links for Optical Clock Comparisons*, New Journal of Physics 23 (9): 093024. <https://doi.org/10.1088/1367-2630/ac21a0>
- KOKE S *et al.*, [2019] *Combining Fibre Brillouin Amplification with a Repeater Laser Station for Fibre-Based Optical Frequency Dissemination over 1400 Km*, New Journal of Physics 21 (12): 123017. <https://doi.org/10.1088/1367-2630/ab5d95>
- KREHLIK P, ŚLIWZCYNski Ł, BUCZEK Ł, KOŁODZIEJ J and LIPINSKI M [2016] *ELSTAB-fibre optic time and frequency distribution technology – a general characterization and fundamental limits in*

IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, vol. 63, no. 7, pp. 993-1004, 2016

- LINDVALL T. *et al.*, [2025] *Coordinated International Comparisons Between Optical Clocks Connected via Fibre and Satellite Links*, *Optica* 12 (6): 843–52. <https://doi.org/10.1364/OPTICA.561754>
- LIPINSKI M. [2022] *The inclusion of White Rabbit into the global industry standard IEEE1588*, Published in: JACoW ICALEPCS2021 (2022) 126-131
- LIPINSKI M., WLOSTOWSKI T., SERRANO J. and ALVAREZ P. [2011] *White rabbit: a PTP application for robust sub-nanosecond synchronization* 2011 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control and Communication, Munich, Germany, 2011, pp. 25-30, doi: 0.1109/ISPCS.2011.6070148.
- LISDAT C *et al.*, [2016] *A Clock Network for Geodesy and Fundamental Science*, *Nature Communications* 7 (12443). <https://doi.org/10.1038/ncomms12443>.
- LOMBARDI M.A [2016] *Evaluating the Frequency and Time Uncertainty of GPS Disciplined Oscillators and Clocks*, *NCSLI Measure J. Meas. Sci. Vol. 11*(3-4), 30-44.
- LOMBARDI M.A, HANSON D.W. [2005] *The GOES Time Code Service, 1974-2004: A Retrospective*. *Journal of Research of the National Institute of Standards and Technology*, Vol. 11, 2, (<http://nvl.nist.gov/pub/nistpubs/jres/110/2/j110-2lom.pdf>).
- LOPEZ O, HABOUCHE A, CHANTEAU B, CHARDONNET C, AMY-KLEIN A, and SANTARELLI G [2010] *Cascaded Multiplexed Optical Link on a Telecommunication Network for Frequency Dissemination*, *Optics Express* 18 (16): 16849–57. <https://doi.org/10.1364/OE.18.016849>.
- LXI [2008] *LXI IEEE-1588 Profile 1.0*, December 1, 2008
- MARRA G. *et al.*, [2018] *Ultrastable Laser Interferometry for Earthquake Detection with Terrestrial and Submarine Cables*, *Science* 361 (6401): 486–90. <https://doi.org/10.1126/science.aat4458>.
- MARRA G. *et al.*, [2022] *Optical Interferometry–Based Array of Seafloor Environmental Sensors Using a Transoceanic Submarine Cable*, *Science* 376 (6595): 874–79. <https://doi.org/10.1126/science.abl939>
- MATSAKIS D.N.M, LEVINE J, LOMBARDI, M. A. [2018] *Metrological and legal traceability of time signals*. *Proc. ION PTTI 2018*, Institute of Navigation, Reston, Virginia, January 29 – February 1, 59-71.
- MAZUR M *et al.*, [2022] *Transoceanic Phase and Polarization Fibre Sensing Using Real-Time Coherent Transceiver*, *Optical Fibre Communication Conference (OFC) 2022*, March, M2F.2. <https://doi.org/10.1364/OFC.2022.M2F.2>
- MCKENZIE W. *et al.*, [2024] *Clock synchronization characterization of the Washington DC metropolitan quantum network (DC-QNet)*, *Applied Physics Letters*, [online], <https://doi.org/10.1063/5.0225082>
- MILLS D. [1991] *Internet Time Synchronization: The Network Time Protocol*, *IEEE Transactions on Communications*. 39 (10): 1482–1493. Bibcode:1991ITCom. 39.1482M. doi:10.1109/26.103043.
- MOREIRA P., ALVAREZ P., SERRANO J., DARWEZEH I., WLOSTOWSKI T. [2010] *Digital dual mixer time difference for sub-nanosecond time synchronization in Ethernet* 2010 IEEE International Frequency Control Symposium, Newport Beach, CA, USA, 2010, pp. 449-453, doi: 10.1109/FREQ.2010.5556289.
- MOREIRA P, ALVAREZ P, SERRANO J, DARWEZEH I AND WLOSTOWSKI T [2010] *Digital Dual Mixer Time Difference for Sub-Nanosecond Time Synchronization in Ethernet*, 2010 IEEE International Frequency Control Symposium, FCS 2010. 449-453. 10.1109/FREQ.2010.5556289.
- NEWBURY N.R, WILLIAMS P.A and SWANN W.C. [2007] *Coherent Transfer of an Optical Carrier over 251 Km*, *Optics Letters* 32 (21): 3056. <https://doi.org/10.1364/OL.32.003056>.

- NOE S, HUSMANN D, MULLER N., MOREL J and FICHTNER A. [2023] *Long-Range Fibre-Optic Earthquake Sensing by Active Phase Noise Cancellation*, Scientific Reports 13: 13983. <https://doi.org/10.1038/s41598-023-41161-x>
- NOVICK, A., JOHNSON, P. [2024] *The Long-Term Timing Performance of Satellite Time and Location Receivers Utilizing Signals from Low Earth Orbit Satellites*, ION 2024 Precise Time and Time Interval Systems and Applications Meeting Proceedings, Long Beach, CA, US.
- NTP ORG [2025] Network Time Protocol (NTP) Project, URL: ntp.org
- P1588 [2020] *IEEE 1588-2019 Evolves to Better Serve its Wide Variety of Applications*, P1588 Working Group, URL accessed Dec 2025: <https://sagroups.ieee.org/1588/news/ieee-1588-2019-evolves-to-better-serve-its-wide-variety-of-applications/>
- PETIT G [2021] *Sub-10 16 accuracy GNSS frequency transfer with IPFP*, GPS Solutions. 25. 10.1007/s10291-020-01062-2.
- PETIT G and JIANG Z [2008a] *GPS All in View time transfer for TAI computation*, Metrologia **45** 35–45.
- PETIT G and JIANG Z [2008b] *GPS Precise point positioning for TAI computation*, Int J Navig Obs Article ID 562878
- PETIT G, KANJ A, LOYER S, DELPORTE J, MERCIER F, PEROSANZ F [2015] *1 x 10⁻¹⁶ frequency transfer by GPS PPP with integer ambiguity resolution*. Metrologia 52(2):301–309.
- PINKERT T.J. *et al.*, [2015] *Effect of Soil Temperature on Optical Frequency Transfer Through Unidirectional Dense-Wavelength-Division-Multiplexing Fibre-Optic Links*, Applied Optics 54 (4): 728–38. <https://doi.org/10.1364/AO.54.000728>.
- PISTONI N.C and MARTINELLI M. [1991] *Polarization Noise Suppression in Retracing Optical Fibre Circuits*, Optics Letters 16 (10): 711–13. <https://doi.org/10.1364/OL.16.000711>.
- PREDEHL K *et al.*, [2012] *A 920-Kilometer Optical Fibre Link for Frequency Metrology at the 19th Decimal Place*, Science 336 (6080): 441–44. <https://doi.org/10.1126/science.1218442>.
- REID T, NEISH A, WALTER T and ENGE P [2016] *Leveraging Commercial Broadband LEO Constellations for Navigation*, Proceedings of the 29th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2016), 15 p., September 2016.
- RIZZI M, LIPINSKI M, FERRARI P, RINALDI S, FLAMMINI [2018] *White Rabbit Clock Synchronization: Ultimate Limits on Close-In Phase Noise and Short-Term Stability Due to FPGA Implementation in IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 65, no. 9, pp. 1726–1737, Sept. 2018, doi: 10.1109/TUFFC.2018.2851842.
- SCHIOPPA M. *et al.*, [2022] *Comparing Ultrastable Lasers at 7×10^{-17} Fractional Frequency Instability Through a 2220 Km Optical Fibre Network*, Nature Communications 13 (1): 212. <https://doi.org/10.1038/s41467-021-27884-3>.
- SHEN Q. *et al.*, [2022] *Free-Space Dissemination of Time and Frequency with 10^{-19} Instability over 113 Km*, Nature 610 (7933): 661–66. <https://doi.org/10.1038/s41586-022-05228-5>.
- SHERMANN J A *et al.*, [2021] *A Resilient Architecture for the Realization and Distribution of Coordinated Universal Time to Critical Infrastructure Systems in the United States*. NIST Technical Note 2187
- SMPTE [2021] ST 2059-2 - SMPTE Standard - SMPTE Profile for Use of IEEE-1588 Precision Time Protocol in Professional Broadcast Applications
- SLIW CZYNSKI Ł, KREHLIK P, BUCZEK Ł., [2024] *Calibrated Fibre-Optic Time Distribution Using Dispersion-Insensitive Configuration*, 2024 European Frequency and Time Forum (EFTF), Neuchâtel, Switzerland, 2024, pp. 94-95
- SLIW CZYNSKI Ł, KREHLIK P, IMLAU H, ENDER H, SCHNATZ H, PIESTER D, and BAUCH A., [2020] *Fibre-Based UTC Dissemination Supporting 5G Telecommunications Networks*, in IEEE Communications Magazine 0163-6804/20/, DOI 10.1109/MCOM.001.1900599.

- SLIWICZYNSKI Ł, KREHLIK P, KOŁODZIEJ J, SCHNATZ H, PIESTER D, BAUCH A, IMLAU H, ENDER H [2019] *Calibrated optical time transfer of UTC(k) for supervision of telecom networks*, *Metrologia* **56** 015006, 2019
- STOWE D, MOORE D and PRIEST R [1982] *Polarization Fading in Fibre Interferometric Sensors*, IEEE Journal of Quantum Electronics, ahead of print. <https://doi.org/10.1109/JQE.1982.1071402>.
- TECKNEXUS [2024] *Precision TimeNet – how it works*, Tecknexus 5G Magazine, Feb 2024, pp. 19
- TAMPELLINI A, CLIVATI C, LEVI F, MURA A and CALONICO D [2017] *Effect of a Timebase Mismatch in Two-Way Optical Frequency Transfer*, *Metrologia* **54** (6): 805–9. <https://doi.org/10.1088/1681-7575/aa8a41>.
- TER-MIKIRTYCHEV V [2014] *Fundamentals of Fibre Lasers and Fibre Amplifiers*, Springer. <https://doi.org/10.1007/978-3-319-02338-0>.
- TERRA O, GROSCHE G and SCHNATZ H [2010] *Brillouin Amplification in Phase Coherent Transfer of Optical Frequencies over 480 Km Fibre*, *Optics Express* **18** (15): 16102–11. <https://doi.org/10.1364/OE.18.016102>.
- TOLMAN J, PTACEK V, SOUCEK A and STECHER R [1967] *Microsecond Clock Comparison by Means of TV Synchronizing Pulses*. IEEE Transactions on Instrumentation and Measurement vol. 16(3), 247–254.
- TØNNES M.B.K [2022] *Understanding and Exploiting Metrological Fibre Networks for Sagnac Sensing, Frequency Dissemination and Geosensing*, PhD thesis, Université Paris Sciences et Lettres (PSL). <https://hal.science/tel-03984045>.
- TØNNES, M.B.K. *et al.*, [2022] *Coherent Fibre Links Operated for Years: Effect of Missing Data*. *Metrologia* **59** (6): 065004. <https://doi.org/10.1088/1681-7575/ac938e>.
- WILKINSON M, DUMONTIER M, AALBERSBERG I, *et al.*, [2016] *The FAIR Guiding Principles for scientific data management and stewardship*, *Sci Data* **3**, 160018 (2016). <https://doi.org/10.1038/sdata.2016.18>
- WŁOSTOWSKI T. [2011] *Precise time and frequency transfer in a White Rabbit network*, MSc thesis, Warsaw University of Technology, Warsaw, Poland.
- WR [2021] *The White Rabbit Project*, URL accessed Dec 2025: <https://www.cern.ch/white-rabbit>
- WR [2025] *The White Rabbit Technology*, URL accessed Dec 2025: www.white-rabbit.tech/wr-technology
- XU D, LEE W.-K, STEFANI F, LOPEZ O, AMY-KLEIN A and POTTIE P.-E [2018] *Studying the Fundamental Limit of Optical Fibre Links to the 10^{-21} Level*, *Optics Express* **26** (8): 9515–27. <https://doi.org/10.1364/OE.26.009515>
- XU D, LOPEZ O, AMY-KLEIN A and POTTIE P.-E [2021] *Non-Reciprocity in Optical Fibre Links: Experimental Evidence*, *Optics Express* **29** (11): 17476–90. <https://doi.org/10.1364/OE.420661>
- YAN Y, CHANGJIAN G, XIONG W *et al.*, [2019] *Distributed Vibration Sensing Based on Forward Transmission and Coherent Detection*, arXiv Preprint.
- ZHOU Q, XIANG Z, QI Z *et al.*, [2025] *Characterization on the Phase Noise Power Spectral Densities of Various Coherent Fibre Links Established in China*, *Optical Fibre Technology* **93** (September): 104184. <https://doi.org/10.1016/j.yofte.2025.104184>

ITU Texts

Handbook “Satellite time and frequency transfer and dissemination” [ITU, 2010]

ITU-T SG15 – Networks, technologies and infrastructures for transport, access and home, URL accessed Dec 2025: <https://www.itu.int/en/ITU-T/studygroups/2025-2028/15/Pages/default.aspx>

- Recommendation ITU-R P.368-9 – Ground-wave propagation curves for frequencies between 10 kHz and 30 MHz
- Recommendation ITU-R TF.374-6 – Precise frequency and time-signal transmissions
- Recommendation ITU-R TF.460-6 – Standard-frequency and time-signal emissions
- Recommendation ITU-R TF.583 – Time codes
- Recommendation ITU-R TF.686 – Glossary and definitions of time and frequency terms
- Recommendation ITU-R TF.768-7 – Standard frequencies and time signals
- Recommendation ITU-R TF.768-7 – Annex – Characteristics of standard frequency and time signal emissions in allocated bands and characteristics of stations emitting with regular schedules with stabilized frequencies outside of allocated bands
- Recommendation ITU-R TF.1153-4 – The operational use of two-way satellite time and frequency transfer employing PN codes Recommendation
- Recommendation ITU-R P.684-7 – Prediction of field strength at frequencies below about 150 kHz
- Recommendation ITU-T G.781 – Synchronization layer functions for frequency synchronization based on the physical layer
- Recommendation ITU-T G.781.1 – Synchronization layer functions for packet-based synchronization
- Recommendation ITU-T G.810 – Definitions and terminology for synchronization networks
- Recommendation ITU-T G.811 – Timing characteristics of primary reference clocks
- Recommendation ITU-T G.811.1 – Timing characteristics of enhanced primary reference clocks
- Recommendation ITU-T G.812 – Timing requirements of slave clocks suitable for use as node clocks in synchronization networks
- Recommendation ITU-T G.8260 – Definitions and terminology for synchronization in packet networks
- Recommendation ITU-T G.8261 – Timing and synchronization aspects in packet networks
- Recommendation ITU-T G.8261.1 – Packet delay variation network limits applicable to packet-based methods (Frequency synchronization)
- Recommendation ITU-T G.8262 – Timing characteristics of synchronous equipment clocks
- Recommendation ITU-T G.8262.1 – Timing characteristics of enhanced synchronous equipment slave clock
- Recommendation ITU-T G.8263 – Timing characteristics of packet-based equipment clocks
- Recommendation ITU-T G.8264 – Distribution of timing information through packet networks
- Recommendation ITU-T G.8265 – Architecture and requirements for packet-based frequency delivery
- Recommendation ITU-T G.8265.1 – Precision time protocol telecom profile for frequency synchronization
- Recommendation ITU-T G.8266 – Timing characteristics of telecom grandmaster clocks for frequency synchronization
- Recommendation ITU-T G.8271 – Time and phase synchronization aspects of telecommunication networks
- Recommendation ITU-T G.8271.1 – Network limits for time synchronization in packet networks with full timing support from the network
- Recommendation ITU-T G.8271.2 – Network limits for time synchronization in packet networks with partial timing support from the network
- Recommendation ITU-T G.8272 – Timing characteristics of primary reference time clocks
- Recommendation ITU-T G.8272.1 – Timing characteristics of enhanced primary reference time clocks
- Recommendation ITU-T G.8273 – Framework of phase and time clocks

Recommendation ITU-T G.8273.2 — Timing characteristics of telecom boundary clocks and telecom time synchronous clocks for use with full timing support from the network

Recommendation ITU-T G.8273.3 – Timing characteristics of telecom transparent clocks for use with full timing support from the network

Recommendation ITU-T G.8273.4 – Timing characteristics of telecom boundary clocks and telecom time synchronous clocks for use with partial timing support from the network

Recommendation ITU-T G.8275 – Architecture and requirements for packet-based time and phase distribution

Recommendation ITU-T G.8275.1 – Precision time protocol telecom profile for frequency synchronization

Recommendation ITU-T G.8275.2 – Precision time protocol telecom profile for frequency synchronization

CHAPTER 7

Use of frequency sources

TABLE OF CONTENTS

	<i>Page</i>
7.1 Use of frequency sources in science and metrology	188
7.2 Definition of the standard frequency	188
7.3 Fundamental physics	188
7.4 Applied physics	189
7.4.1 Local measurement of g , the acceleration of gravity	189
7.4.2 Global measurements of the gravitational constant, G , and the gravitational potential, g	189
7.4.3 Very Long Baseline Interferometry (VLBI)	189
7.5 Positioning and navigation	190
7.5.1 System design	190
7.5.2 Doppler solutions	191
7.5.3 Differential solutions.....	191
7.5.4 Spherical solutions	191
7.6 Telecommunications.....	192
7.6.1 Channel capacity	192
7.6.2 Frequency division multiplexing.....	193
7.6.3 Time division multiplexing.....	193
7.6.4 The ITU-T standardization work.....	193
7.6.5 Trellis coding	194
7.7 Provision of electrical power	194
7.8 Concluding comment.....	194
References	194
ITU Texts	195

7.1 Use of frequency sources in science and metrology

A frequency source (FS) is a device that provides an electrical signal whose output frequency satisfies the requirements, such as the accuracy, stability and immunity from environmental perturbations, needed for a particular application. The frequency of the output signal may be derived from the mechanical properties of a quartz crystal or from the frequency associated with a transition in an atom or molecule. Many devices combine the two reference techniques and use a quartz crystal oscillator whose frequency is monitored or controlled by a frequency derived from the properties of an atom or molecule. Although the frequency associated with any transition in any atom or molecule could be used as the reference for a frequency source, there are technical advantages to using a hyperfine transition in the ground states of caesium, rubidium, or hydrogen, and most commercial frequency sources use the hyperfine transition frequency in one of these atoms as the reference. Details on the various devices have been provided in Chapter 3.

At a given level of accuracy, frequency sources and frequency measurement techniques are much simpler and cheaper than measurement techniques for other quantities such as length or mass. It is usually convenient to convert the quantity of interest – voltage, pressure, humidity etc. – into a frequency or a time interval via a transducer. Very often this improves the accuracy and reduces the cost of making the measurement and processing the results.

7.2 Definition of the standard frequency

Until May 2019 the International System of Units (the *Système International d'Unités* or SI) was defined in terms of seven base units and derived units defined as products of powers of the base units. From 20 May 2019 the SI units are defined in terms of seven defining constants. The complete system of units can be derived from the fixed values of these defining constants, expressed in the units of the SI. Despite of all progress made over the years, the standard frequency remains defined as the frequency associated with a hyperfine transition in the ground state of Cs^{133} . The practical realization of the definitions of the units is described in the SI Brochure (BIPM 2019), Appendix 2, which is accessible at the BIPM website (www.bipm.org). It is likely that some other atomic transition, or some combination of atomic transitions, will replace the definition based on caesium in the next 10 years (see Chapter 9 for a forecast). Devices that will realize the new definition will be complicated and expensive and will probably have to be operated in a controlled, laboratory environment. This is also true of caesium fountains, the devices that realize the current primary definition of the SI second. In addition, primary frequency devices generally sacrifice short-term stability to realize long-term accuracy and are not suitable for applications that depend on a real-time or near-real-time frequency reference. The re-definition of the SI second will be unlikely to affect applications that depend on frequency or time interval, since many of these applications require a real-time or a near-real-time reference signal or are limited by other constraints, and the importance of commercial frequency sources based on caesium, rubidium, or quartz crystals, is likely to continue for the foreseeable future. Hydrogen masers will continue to be important frequency sources because they have very good short-term stability, but they are too expensive and cumbersome to be used in most applications that do not require excellent short-term stability.

7.3 Fundamental physics

The stability and accuracy of the output frequency of any device is derived from the fundamental physics of the reference system, so that comparing the output frequencies of devices stabilized in different ways or to different atomic transitions can provide important information on the underlying physics of the atomic reference systems. These applications are particularly powerful for studying simpler atomic systems because these frequencies can be accurately calculated by using well-known quantum-mechanical principles.

The principles of Relativity affect the frequency measured by an observer who is moving with respect to an oscillator or who is at a different gravitational potential than the frequency standard. The “gravitational red shift” can be exploited to study the spatial variation of the gravitational potential on Earth. Comparing the frequencies of oscillators at very different gravitational potentials (on Earth and on a satellite, for example) provide an important test of Relativity. Up to now, the most accurate test of the so-called relativistic redshift was made using atomic clocks on two satellites of the European satellite navigation system Galileo [Delva *et al.*, 2018; Herrmann *et al.*, 2018] when a partial launch failure brought the satellites into an elliptical rather

than circular orbit. The eccentricity of the orbit resulted in a large periodic variation in the frequency of the transmissions observed by the tracking stations on Earth.

The principles of both Special and General Relativity are important in understanding the differences in time and in rate between oscillators on the surface of Earth and those located on the surfaces of the Moon or Mars, or at one of the stable Lagrange points. The calculation of these effects is quite complicated and includes a number of approximations, so that a comparison between theory and experiment will be very useful. Maintaining a communication link between Earth and a spaceship traveling to one of these remote locations will be challenging. The operation of LISA, the space-based interferometer to detect gravitational waves, will have similar difficulties.

7.4 Applied physics

Frequency references, usually quartz-crystal oscillators, are very widely used, and it is impossible to list all of the possible applications. Devices such as radio and television receivers depend on a local frequency reference in a fundamental way, but the “internet of things” has added requirements for a local frequency reference to devices such as washing machines and kitchen ovens. Most of these devices have only modest accuracy and stability requirements for the reference frequency, and oscillators that use a quartz crystal or a MEMS (Micro-Electro-Mechanical Systems) device as a frequency reference are often adequate. The considerations of SWAP (size, weight, and power) are often more important than the accuracy or even the stability of the oscillators in these systems. Monitoring the operation and validating the accuracy of a large number of these “dumb” devices will be challenging.

7.4.1 Local measurement of g , the acceleration of gravity

The measurement of the acceleration of gravity is obtained by observing the free-fall of an object in vacuum. The object is a corner-cube mirror that is part of a Michelson interferometer. The position of the corner-cube is measured in terms of the wavelength of the light that illuminates the interferometer, and its velocity and acceleration are determined by combining these position measurements with timing signals from a local frequency reference. The local value of the g can be determined with a fractional accuracy of about 10^{-9} . The measurement is challenging but is generally not limited by the accuracy or the stability of the wavelength and time reference systems.

7.4.2 Global measurements of the gravitational constant G , and the gravitational potential g

The gravitational constant, G , is difficult to determine because gravitational effects depend on GM , the product of G with the mass of the source (or on GMm , the product of the G , the mass of the source and the mass of the test object in the Eotvos experiment). The product GM , where M is the mass of Earth, is determined by using laser ranging between a station on Earth and a satellite with a retro-reflector. (The measurement does not depend on the mass of the satellite.). The measurement has a fractional accuracy of about 10^{-9} . The product is known with greater accuracy than the accuracy of either G or M .

The spatial variation of the gravitation potential of Earth can also be determined by measuring the variation in the distance between two satellites that are relatively close together in the same orbit. Successful examples are the Gravity Recovery and Climate Experiment, GRACE, and the follow-on measurements, GRACE-FO. The spatial variation is expressed in a spherical-harmonic expansion with a degree and order of about 300, so that the angular resolution on Earth is of order 1° , which corresponds to about 130 km in spatial resolution at the equator. The variation of the distance between the two satellites is measured interferometrically and is combined with orbital parameters observed from ground tracking stations.

7.4.3 Very long baseline interferometry (VLBI)

A Very Long Baseline Interferometer is based on a number of radio telescopes located at widely separated locations. Each site records the transmissions received from a distant radio source using a local hydrogen maser as the time and frequency reference. The recordings from the different sites are cross-correlated to solve for the three-dimensional orientation of Earth in space with respect to the distant radio sources. The data provide a measurement of UT1, the rotation angle of Earth, and polar motion – the precession, nutation, and irregular motion of the axis of rotation. The cross-correlation process depends on the stability of the frequency reference

at each site, so that hydrogen masers are typically used because they have very good fractional-frequency stability of order 10^{-15} over the duration of the measurements, which are often as long as a few hours. In addition to providing information on the position of Earth as a whole, the data can also be used to study continental drift and similar geophysical effects.

The method of cross-correlating radio data received at different sites can also be used to process the signals from artificial satellites. This technique is called quasi-VLBI and is useful in geodesy and space navigation. The motion of the satellites that are the sources of the signals cannot be neglected, so that this method is more complicated than VLBI, which uses signals from sources that are so far away that they can be treated as not moving.

7.5 Positioning and navigation

Determining the position of a receiver, and navigation – measuring the evolution of the position in time, are related applications, and the same techniques can be used to address both problems. However, the statistics are different. The navigation application must operate in real-time, and cannot benefit from the post-processing, averaging, and outlier detection that can be used to improve a position solution. The timing of electromagnetic signals is a fundamental aspect of both applications, and the methods that a receiver can use to determine its position (or the evolution of its position) by processing electromagnetic signals from a single transmitter or an ensemble of them will be discussed. The statistics of the clocks in the transmitters and the receiver are important, but an equally important problem is relating the physical position of the receiver to the reception time of the signal from the transmitters. The refractivity of the atmosphere complicates the solution because the geometrical distance between the transmitters and the receivers is less than the distance deduced from a measurement of the time of flight of the signals. A number of different methods have been developed to address this problem. Following general considerations on system design, examples are given. GNSS is currently the most widely used system for positioning, navigation and timing (PNT).

7.5.1 System design

The methods that will be described in the next sections depend on receiving the signals from multiple sources whose positions (or orbital parameters) are known and whose clocks are synchronized to a common reference, and realizing these requirements is outside of the scope of a user's application. The orbital parameters of each satellite and the offset of its clock with respect to the common reference (which is usually the ensemble-average system time) are estimated by a ground control station by combining the data received at a number of tracking stations with statistical models of the clocks. These estimates are uploaded periodically into each satellite and are transmitted to users via the navigation message. The system time may be steered to UTC based on a prediction of UTC from a timing laboratory, but the positioning and navigation applications do not depend on this steering. The accuracy of the solution obtained by a receiver is compromised both by the stability (i.e. predictability) of the clock in the satellite and by the accuracy of the broadcast ephemeris that is used for estimating the geometrical time of flight. The impacts of these contributions can be addressed by post-processing using precise orbits and clock solutions, but this advantage may not be possible in all configurations.

The time of flight of the radio signals is modified by the refractivity of the medium along the path. For signals from global navigation system satellites, the largest contribution to the refractivity comes from the ionosphere. This refractivity can be estimated by measuring its dispersion – the dependence of the measured time of flight on the frequency of the signal, and navigation satellites transmit information on multiple frequencies for this reason.

This method cannot be used to estimate the refractivity of the troposphere, since this refractivity has no dispersion at radio frequencies. The refractivity of the troposphere adds about 2 m (or 6 ns) to the geometrical path delay for a satellite at the zenith, and a larger value for satellites at lower elevations. This contribution must be estimated in other ways and is often ignored.

The position of the receiver (or the evolution of its position) can be determined by measuring the times of flight between the antenna and four satellites (see § 7.5.4.). The contribution of the troposphere increases the length of every path so that it is correlated with the determination of the vertical position of the receiver, and

the determination of this coordinate often has a larger error than the error in the North and East position determinations.

7.5.2 Doppler solutions

If the transmitters are moving (on satellites, for example), or if the receiver is moving with respect to stationary transmitters, then the frequencies of the received signals will be modified by the component of the motion along the line of sight between each transmitter and the receiver in first order. These Doppler shifts can be used in a positioning or navigation application instead of (or in addition to) the time-of-flight method described in the previous section. The component of the velocity perpendicular to the line of sight is responsible for a second-order Doppler shift, which is smaller than the first-order shift by v/c , the ratio of the speed of the source to the speed of light. If the receiver can determine this Doppler shift, then each determination locates the receiver on a conical surface whose apex is at the transmitter. The receiver can determine its position by repeating this measurement using the signals from multiple sources. If the clocks in all the sources transmit at the same frequency, then the receiver can solve for its position and for the offset of its clock with respect to the system time of the transmitters. If the orbit of a satellite transmitter is accurately known, then the receiver may be able to solve for its position by measuring the evolution of a single Doppler shift as the position of the satellite changes. This method places a more stringent requirement on the stability of the oscillator in the receiver. The principal advantage of this technique is that the Doppler shift is not sensitive to the refractivity of the atmosphere, although the measurement of the frequency offset will be made more difficult by fluctuations in the atmospheric refractivity.

For most applications, Doppler navigation has been replaced by methods that use the transit times of signals from GNSS satellites. The Doppler method is still used to determine the ground-speed of low-flying aircraft (helicopters, for example) or when the signals from GNSS satellites are not available or are not reliable. The method is also incorporated into the tracking loops of many GNSS receivers. The data from receivers at different sites can be useful in detecting “spoofing” – signals from stationary rogue transmitters that attempt to mimic the signals from the real satellites but cannot reproduce the Doppler shifts each receiver expects

7.5.3 Differential solutions

A second method for attenuating the impact of the atmospheric refractivity in positioning and navigation applications is for a receiver to measure the *difference* in the arrival times of signals from two (or more) transmitters whose signals are synchronized to a common reference time. The impact of the refractivity is attenuated to the extent that its contribution is the same to both paths and cancels in the difference. A measurement of the difference in arrival times locates the receiver on a hyperbola with the transmitters at the foci. If this measurement is repeated with another pair of transmitters, the position of the receiver is the point where the various hyperbolas intersect. This technique was widely used as the basis for the LORAN-C system. Its accuracy depends on the equality of the refractivity along the paths from the receiver to the two transmitters and was often most useful over water where this requirement was usually satisfied. A weakness of the ground-based transmitters used in the LORAN-C system was that the path difference was not sensitive to the altitude of the receiver, so that the method was not useful for airplane navigation.

7.5.4 Spherical solutions

Spherical navigation is currently the most widely used technique for both navigation and positioning applications, and this technique is based on receiving the transmissions from global navigational satellite constellations such as GPS, GLONASS, Galileo, or BeiDou. All of the constellations use the same basic principles. They comprise a space segment in MEO (Medium Earth Orbit) at about 20 000 km altitude on several orbital planes and transmit radio signals on at least two L-band frequencies. The position of the receiver (or its motion) is determined by measuring the times of flight between the receiver and several satellites in known orbits. Each satellite transmits information about its orbit, and an estimate of the offset of the clock on the satellite with respect to the system time of the constellation. The measurement of the transit time from each satellite locates the receiver on a sphere whose radius is defined by the transit time, and the position of the receiver is determined by the intersection point of the spheres determined by observing several satellites simultaneously. The receiver solution has four unknowns – the three coordinates of the position and the offset of the receiver clock with respect to the system time of the constellation, and these four unknowns can be

determined by observing the signals from four satellites from the same constellation. This enables a multitude of applications, such as incorporating receivers in modern wristwatches and smartphones: only a low-performance oscillator is required in the receiver. It is possible to use data from fewer satellites if the time of the clock in the receiver is known from other data or if its position is constrained, to be on the surface of Earth, for example. A receiver at most locations on Earth will receive signals from many satellites, so that solutions based on a redundant array of independent measurements (RAIM) can be used to detect outliers and errors. The Handbook “Satellite time and frequency transfer and dissemination” [ITU, 2010] provides a comprehensive description of the GPS and GLONASS satellite systems (Chapters 4 and 5 respectively), as well as the principles of time transfer via GPS (Chapters 12 and 13). At the time of publication of this older Handbook, the European GNSS Galileo had not been completely developed. The interested reader is referred to the web site of the European GNSS service centre (<https://www.gsc-europa.eu>) for information and documentation of the Galileo system performance.

7.5.4.1 Timing applications based on spherical solutions

The importance of GNSS signal reception for the realization of TAI and UTC was explained in Chapter 5. GNSS signals have evolved as the most important means for global time dissemination as explained in Chapter 6.

The signals from global navigation satellites have a central role in the telecommunications sector, and this will be discussed in the next section.

7.6 Telecommunications

Early communication methods used a low-rate digital method such as the Morse code or analogue systems such as point-to-point telephone circuits. The methods did not have strong requirements for time accuracy or frequency stability. They were typically realized by using a physical connection between the sender and the receiver. They had a stable transmission delay and could also be used to transmit time and frequency signals with minimal degradation. These simple channels were not adequate to support the increase in the demand for long-distance communication, and various methods were developed to support an increase in the capacity of communications circuits to handle this increase in the demand.

The increase in capacity was realized by sharing the physical circuit among multiple logical circuits, and the details will be discussed in the following sections. The sharing increased the capacity of the physical circuit, but the stability of the circuit delay was degraded by the multiplexing and de-multiplexing operations. Dedicated point-to-point circuits, especially circuits that use optical fibres as the transmission medium, have seen a renaissance to transmit time and frequency data when the highest level of delay stability is needed. Examples are given in Chapter 6.

7.6.1 Channel capacity

The Nyquist theorem relates the capacity of a digital circuit to its physical bandwidth. For example, if the physical bandwidth of a telephone circuit is 4 kHz, then the circuit can support a digital signal of 8 000 samples per second. An important advance in telecommunications methods was the realization that the capacity of a physical channel also depended on the signal to noise ratio (S/N) of the connection. The Shannon-Hartley theorem relates the maximum rate at which information can be transmitted over a communications channel, C , in bits per second, to the bandwidth of the channel, B in Hz and the S/N :

$$C = B \ln\left(1 + \frac{S}{N}\right) \quad (7-1)$$

where the noise is generated by a Gaussian random process. Although the Nyquist theorem limits the *signalling* rate to twice the physical bandwidth, the *information* rate can be increased by transmitting more than one bit per symbol, where the amplitude of the signal also carries information, for example. The limit is set by how many discrete amplitudes the receiver can distinguish, and this depends on the S/N of the system. The theorem asserts that an encoding method exists that can realize this capacity but does not define the method. Three methods have been developed to increase the rate at which information is transferred of a channel.

7.6.2 Frequency division multiplexing

One of the first methods that was used to increase the capacity of a channel was frequency-division multiplexing. Each transmission on the physical circuit is assigned a unique frequency. The data for each link are encoded on a unique frequency at the transmitting end and decoded at the receiving end. The number of distinct messages that can be combined on the circuit depends on the frequency accuracy of the encoding hardware and the bandwidth of the frequency filters used in the receiver.

The method depends on a “channel bank” – an encoder/decoder pair for each frequency. It was too expensive to be used for copper circuits, but the method has seen a renaissance as wavelength-division multiplexing in circuits that use optical fibres.

7.6.3 Time division multiplexing

Another method to increase the capacity of a channel is time-division multiplexing, in which each transmission is assigned a time-slice in a round-robin fashion. For example, in the early T1 system, 24 voice conversations were multiplexed using the signalling rate of 8 000 samples per second. The transmission is divided into packets, and each packet has one bit of each conversation. The method depended on the fact that the actual information rate of a voice conversation is less than the signalling rate of the circuit so that each sample of the voice transmission is broken up into 8 packets. In a later version, the signalling rate was increased to 1.544×10^6 samples per second, giving 24 channels each of which can support 64k bits per second.

The usefulness of the method depends on the synchronization of the encoding and decoding oscillators at the ends of the circuit. The system can handle small discrepancies in the frequencies at the end points with “slip buffers.” This technique cannot cope with a systematic frequency offset between the clocks at the endpoints. White frequency noise, which integrates to a random-walk in time can also be a problem.

The method of time division multiplexing becomes more complicated when the end-point of a circuit is a switching centre where the various messages must be separated and then combined with other messages from another circuit into a multiplexed transmission to another switching centre. Various “slip buffer” and FIFO (first-in-first-out) buffers are used to compensate for short-term fluctuations in the rates of the different channels.

The frequencies of the clocks in a complex network that is based on time-division multiplexing must be nearly synchronous (“plesiochronous”) and cannot have long-term systematic frequency offsets that will eventually saturate the largest slip buffer. This requirement is usually realized by an ensemble of primary oscillators, which are often caesium-based devices that can be monitored and validated by using data from global navigation satellites. The frequency output of this primary ensemble is used to control a larger ensemble of oscillators that have poorer stability. The devices in this pyramid scheme are identified by an integer “stratum number,” which specifies the number of links from that device back to the primary reference. Clocks that operate at each stratum number typically have poorer accuracy and stability than the devices at a lower stratum.

(In networks that support the network time protocol (NTP), the stratum number of a system specifies only the number of links back to a primary reference source and does not necessarily have any relationship to the stability or accuracy of the device.)

7.6.4 The ITU-T standardization work

International standards (ITU-T Recommendations), which detail technical specifications regarding the elements of the global communication infrastructure, have been developed by ITU-T Study Group 15. The group administers many Recommendations, including Recommendation ITU-T G.8260, which contains *Definitions and terminology for synchronization in packet networks*. This document specifies that UTC is needed as a common time reference for mobile time division duplex (TDD) circuits.

There are three recommendations that specify the performance of Primary Reference Clocks, which differ in the prescribed performance of the output signals and the architecture of their installation. The PRTC device is specified in Recommendation ITU-T G.8272, the enhanced device, ePRTC in Recommendation ITU-T G.8272.1, and the coherent network device, cnPRTC in Recommendation ITU-T G.8272.2. The reception of GNSS signals is common to all PRTCs, either as the primary or as the back-up timing source. UTC information can be retrieved from the GNSS navigation message, which contain the prediction of the difference between

the system time and a prediction of UTC (see Chapter 5). The ITU-T website of Study Group 15, Question 13 describes the current performance specifications.

7.6.5 Trellis coding

A Trellis-coding method encodes several bits of information onto one signal bit. For example, the amplitude of a single bit could be modulated to transmit more than just one bit of information. Or the phase of the transmitted frequency could be modulated by a multiple of $\pi/2$ to transmit 2 bits of information on each phase shift. (This technique is equivalent to transmitting different signals on the in-phase and in-quadrature components of the carrier.) In each case, the limit depends on the S/N of the connection – what is the smallest change in the amplitude or the phase of the received signal that can be detected by the receiver. This method can be combined with other methods such as time division multiplexing.

7.7 Provision of electrical power

Most activities of a society depend on the reliable and stable availability of electrical power. A long-lasting interruption of electrical supply would compromise the distribution of drinking water, the refrigeration and cooking systems that are required to store and prepare food, the operation of the heating or air-conditioning systems, and the numerous commercial, industrial, and residential systems that depend on electrical power. Electrical power is distributed from multiple sources to consumers in grids. Grid operators use precision timing for grid monitoring, protection and operation. The wide-area monitoring and situational awareness of grid performance rely on synchrophasor technology which monitors the amplitude and phase of the electrical power at many points on the grid. The proper operation of these units depends on tight timing requirements. Phase measurement units (PMU) are employed to measure the voltage and current waveforms at many points in the grid with high accuracy and with high temporal resolution. Each measurement is time-stamped, with reference to UTC. PMU data is collected at 30 to 120 samples per second. The measurements are analysed to compute the magnitude and the phase angle of the electrical power relative to a cosine function at the nominal transport frequency (50 Hz in Europe, 60 Hz in the US, to give examples) that is derived from UTC. The IEEE standard C37.118-1-2011 specifies the measurement rate and accuracy of these devices and covers a wide range of uses of the timing data, such as the detection of anomalous events and the location of line faults. The most demanding requirements depend on a synchronization to UTC with 1 μ s uncertainty. The signals received from GPS (or GNSS in general) are widely available and can satisfy these requirements. These signals have become the de facto standard in many deployed installations for these reasons.

7.8 Concluding comment

In the two preceding chapters, the technical requirements for timing in networks of the telecommunications and electrical power sectors have been presented in brief. In addition to the technical requirements, some countries and the European union also have legal requirements. [Lapuh, 2011] For example, the registration and billing of telephone calls must be documented with time stamps in legal time. Time information is needed only with 1 s resolution in this application, but it requires considerable effort to be compliant with such requirements and to be able to verify it in repeated audits. There are similar requirements for the billing of electrical energy, so that smart grid elements and electricity meters must be connected to a source of legal time. This brings us to the discussion of time dissemination traceable to UTC in the following Chapter.

References

- DELVA P *et al.*, (2018) Gravitational Redshift Test Using Eccentric Galileo Satellites, *Physical Review Letters* 121, 231101 (2018)
- HERRMANN S *et al.*, (2018) Test of the Gravitational Redshift with Galileo Satellites in an Eccentric Orbit, *Physical Review Letters* 121, 231102 (2018)
- IEEE (2011) *C37.118.1-2011 – IEEE Standard for Synchrophasor Measurements for Power Systems*, IEEE, DOI: 10.1109/IEEESTD.2011.6111219

LAPUHR (2011) EURAMET Countries' Legal Time Regulations and Practices, URL: <https://www.euramet.org/technical-committees/tc-projects/>

ITU texts

Handbook “Satellite time and frequency transfer and dissemination” [ITU, 2010]

Recommendation ITU-T G.8260 – Definitions and terminology for synchronization in packet networks

Recommendation ITU-T G.8271 – Time and phase synchronization aspects of telecommunication networks

Recommendation ITU-T G.8271.1 – Timing characteristics of enhanced primary reference time clocks

Recommendation ITU-T G.8271.2 – Timing characteristics of coherent network primary reference time clocks

CHAPTER 8

Time as a Service

TABLE OF CONTENTS

	<i>Page</i>
8.1 Overarching issues regarding the use of frequency sources	198
8.2 Reliability, good practice of installation.....	198
8.3 Resilience.....	199
8.4 Traceability	200

8.1 Overarching issues regarding the use of frequency sources

Time is ubiquitous and as has been seen, many sectors of modern technology are strongly dependent on the availability of local physical signals as well as of steering information so that a common time reference can be distributed in networks. In the previous chapters the different potential sources, steering references and dissemination services have been presented. The focus was in most cases on the performance, the achievable frequency accuracy and stability.

The term *Time as a Service* has become somewhat popular, as many deployed systems do not only include the necessary elements for their operation but provide options to distribute the embedded time information to third parties, potentially beyond the borders of the system operator. In this case the operator needs to assure the origin of time information and its unaltered transport through the network. Ideally, all the steps from the agreed international time reference UTC or its realizations at an NMI are transparent up to the signal received by the end user, including the uncertainty of each step of conversion and routing. This is by far not an easy undertaking, and often the term *traceability* is employed without adherence to strict metrological traceability. In this chapter, some material has therefore been compiled on the topics' traceability and, before that, on reliability and resilience. Each of the topics has been treated by numerous organizations and agencies in many documents in a verbose and often very general style. In order not to add one more contribution of that kind, the following content concentrates on some "simple" rules drawn from many years of practical work in the time and frequency business.

8.2 Reliability, good practice of installation

What is the need for a frequency reference and/or time reference? Does it need to be traceable to the SI second and/or to UTC? Does it need to be kept within some boundary frequency limit or time limit as part of a network? Does the signal at its full accuracy need to be available nominally 24/7/365? Which deviation from the ideal situation can be tolerated?

Any equipment, from the source to the end user terminal can – inevitably - fail. The choice of products must be guided a priori by the reliability specification of the vendor. If highest availability is requested, duplication or triplication of key elements are necessary. This comprises signal sources, receiving channels for timing information and equipment along the dissemination path.

The sophistication of the electrical installation depends on the local situation in the region or country. Feeding a time and frequency source with AC and DC in parallel should be standard. A double DC input can help when the equipment is supposed to be transported in operating conditions. In case that equipment is operated in duplicate at a given site, it should be powered by different circuits (different fuses) or even located in different rooms or buildings. Some frequency sources are sensitive to power line voltage changes. In addition, after a power outage, the frequency source may require several days or longer to reach a stable, accurate frequency again. For these reasons it is advisable to operate the frequency source from an uninterruptible power supply (UPS). A number of good commercial units are available. It is important that the size of the battery in the UPS be large enough to last for the longest expected outage and that the batteries be changed periodically since they have a limited lifetime.

When the optimum performance is expected, the location of the installation must match the specifications of environmental sensitivity of the devices. At some level, all electric equipment is sensitive to temperature, humidity, ambient magnetic field, and vibrations. If the sensitivity is not specified by the manufacturer, it is safer to choose another product, unless one has the resources to determine the sensitivity by own means. It is unphysical to expect perfect insensitivity.

Looking on an installed system from a higher-level position one should try to identify potential single points of failure and modify the installation to avoid them. So-called automatic switches may indeed change from a defect input channel to another, but the output may suffer a short gap and a discontinuity in phase. Whether this is tolerable will depend on the application.

Much has been said on the weakness of GNSS signals and the ease of disturbing its reception, intentionally or unintentionally. In GNSS-disciplined oscillators, the hold-over capacity needs to be known, and the product adapted to the needs. Good examples are in the ITU-T recommendations which specify various types of Primary Reference Clocks, mentioned in Chapter 6.

The reliability of devices and services will also likely be improved when following some ideas mentioned in the next section on resilience.

This section concludes with a short bullet list of potential sources of insufficient performance, inherited from the 1996 ITU-Handbook with the same title.

– *Ground loops*

When multiple instruments are interconnected, there is a problem of having multiple ground connections through which rather large currents can flow because of the low loop resistance. These currents cause extraneous voltages at the inputs and outputs of the various instruments. These voltages can degrade the performance of the frequency source, the distribution system and the measurement instrumentation. The best solution to this problem is to have all the ground connections made at a single point. This may require isolation transformers.

– *Power-line related magnetic fields*

Electronic instruments contain power supplies which often have stray magnetic fields at the power line frequency (50 or 60 Hz). Depending on the quality of the supplies and the position in the electronics rack, this field has been measured to be as great as 50 μT (500 milligauss). Careful positioning of the various instruments or magnetic shielding may be necessary to reduce the field enough so that it does not degrade the performance of the frequency source.

– *VSWR, cables, connectors*

The commonly used BNC connectors are notorious for mismatch and for delay variations and step changes. As an example of a cable length problem, the wavelength at 5 MHz is about 60 m. Unterminated cables with $1/4 \lambda$ (15 m) electrical length, which is not uncommon for general use, can act like a short. Cables themselves can cause problems because of phase changes and jumps when they are moved or when the temperature is changed. Even if currently the use of electrical signals exceeding 18 GHz in frequency is quite rare, it is advised to use the following connectors, in increasing order of quality and suitability: BNC, TNC, N-type, SMA (up to 18 GHz). If higher, frequency signals shall be handled, then PC-3.5 mm (26.5 GHz), PC-2.92 mm (40 GHz), PC-2.4 mm (50 GHz), PC-1.85 mm (67 GHz) and PC-1 mm (110 GHz).

8.3 Resilience

The Cambridge Dictionary defines resilience as “the ability to be happy, successful, etc. again after something difficult or bad has happened”, and in another context as “the ability of a substance to return to its usual shape after being bent, stretched, or pressed”. In a technical context, a system may be called resilient when it operates irrespective of changing conditions and recovers rapidly after a disruption, an attack from outside or an incident. One may distinguish between autonomous recovery or recovery involving human interaction. The most trivial action that may be needed during recovery from a power failure is the manual re-boot of the IT system that monitors and controls the local installation. IT-security in general, however, is beyond the scope of this Handbook. A wider interpretation of resilience has been proposed by Erik Hollnagel [Hollnagel *et al.*, 2013] as “The intrinsic ability of a system to adjust its functioning prior to, during, or following changes and disturbances, so that it can sustain required operations under both expected and unexpected conditions.” This implies not only resistance to disturbances but the capacity to adapt under disturbances.

The following features of a system architecture help to speed up or even enable recovery after a major disruption:

- Documentation of the system architecture (e.g. signal routing, software driving measurement equipment) are updated and available easily at a well-known place.
- During normal operation, all elements are monitored, and internal parameters are logged and stored on a server out of the laboratory in parallel.
- Local frequency standards are compared using duplicate measurement systems via PPS signals and/or frequency signals. Again, availability of the last recorded data should be easily accessible.
- Data sets in storage should indicate the time of data taking, the measurement quantity, and the unit of measurement results.

- In case that data are exported to third parties, like GNSS data files or clock measurement data etc., the IT routine to affect the transfer is understandable and allows to check which data export has eventually failed. Logging all transactions is very helpful.
- Key devices should ideally be kept in storage to be ready for replacement of a defect device. Some care is needed that the configuration of ports and power supplies is sufficiently similar so that cabling need not be changed.

An option to prepare for potential disruptions is the set-up of a watchdog system. It should record relevant internal parameters, but also, e. g. the size of a file that nominally should grow during a day and compares to the expected values or behaviour. In case of mismatch alarm is distributed in a suitable way to competent staff. Timely detection of anomalous behaviour is a key preventive measure. Another measure is keeping all software and firmware updated as far as possible and monitoring available storage capacity routinely.

One could add many detailed suggestions, but they would be very much tailored to specific applications or equipment. On the other hand, a very competent document was released by NIST [Sherman *et al.*, 2021] which is recommended for further reading.

An attempt at a common framework for classification into graded resilience levels is the scope of the ongoing IEEE Standards Association project IEEE P1952: *Standard for Resilient Positioning, Navigation, and Timing (PNT) User Equipment*. Note that the scope for its classification is limited to user equipment, explicitly excluding PNT sources, such as GNSS satellites or ground stations. At the same time, the text feeds its classification from assumptions about those PNT sources and their workings. Note also that the scope goes beyond timing and not only includes spatial positioning, but puts its emphasis on it. Proposed graded resilience levels have been the following:

- 1 Detect (and alert): The user equipment can detect any adversity that might prevent equipment from maintaining acceptable PNT performance
- 2 Recover: The user equipment can automatically regain acceptable PNT performance, once the adversities are no longer present
- 3 Resist: The user equipment can maintain acceptable PNT performance even under adversity, for some period of time
- 4 Withstand: The user equipment can maintain acceptable PNT performance indefinitely, even under adversity
- 5 Verify: The user equipment can determine with certainty whether any PNT information it receives is correct

Note that the resilience levels and their wording are based on an ongoing project as of March 2026, so the language and proposed classification are subject to change.

8.4 Traceability

This final section of the chapter was motivated by the following observations. As noted previously on several occasions, access to accurate time is indispensable for the functioning of modern infrastructure world-wide in virtually all sectors, including electrical power supply, telecommunications, financial institutions, time and frequency metrology, and (quite naturally) positioning and navigation. In some of the above sectors, the demand for demonstrating traceability to national or international standards has been imposed by legislation or regulation. This becomes particularly relevant when organizations advertise *Time as a Service*. Reference to UTC is made in many normative documents and international standards issued by and applicable in such user groups that share an interest in timing. Often the word “traceability” is found in such documents but not used with the strict meaning applying in the metrology community.

The reception of signals from Global Navigation Satellite Systems (GNSS) has developed as a common standard in equipment in use. In addition to the broad use of GNSS for navigation, positioning and scientific applications, the GNSS also fulfil an effective time dissemination function. This is inherent in the basic technical functionality of a GNSS, as explained in Chapters 6 and 7 of this Handbook. The impression is often raised that such GNSS reliance would already provide the required metrological traceability to UTC of disseminated time signals, based on the GNSS input. In reply to this, the BIPM had created a Task Group to

elaborate on the subject, and its findings have been compiled in [Defraigne *et al.*, 2022]. A more detailed version (the White Paper) can be obtained from BIPM with a request to tai@bipm.org.

In the following, the main thoughts from these two publications are given.

The International vocabulary of metrology (VIM) reference document provides a definition of “metrological traceability” to a given reference [JCGM, 2012]: it is the “property of a measurement result whereby the result can be related to a reference through a documented unbroken chain of calibrations, each contributing to the measurement uncertainty”. The ITU-R adopted an almost identical definition in its Glossary [Recommendation ITU-R TF.686 – Glossary]: “The property of the result of a measurement or the value of a standard whereby it can be related to stated references, usually national or international standards, through an unbroken chain of comparisons all having stated uncertainties”. In simple words this means: From the source of timing information to the end user the individual steps of signal generation, processing, and transmission should be transparent to the end user. Each step potentially entails a delay or alteration of the forwarded signal which need to be determined by a calibration, including its uncertainty. A good example is the GNSS delay calibration scheme practised by the BIPM as explained in Chapter 5.

Three use categories can be distinguished:

- 1) Use of GNSS signal-based information and data, including the control of a local oscillator, for internal applications;
- 2) Same as 1) but explicit reference is made to “traceability to UTC”;
- 3) Calibration laboratories using GNSS signals as source of reference signals for their services offered to third parties.

Uses in the realm of ITU will fall in categories 1 and 2, and proposals have been included in [Defraigne *et al.*, 2022], suggesting how to achieve the traceability strived for. Validating legal traceability generally requires some form of log file system and may require monitoring by an external regulatory authority. The details are defined by local statutes or regulations.

In closing the Chapter on *Time as a Service*, it is recalled that legal prescriptions or regulations issued by many user communities ask for traceability to national or international standards when measurements are made or when time stamps are issued. This trend is in line with general metrological requirements. International consistency and comparability of measurements are essential to ensure that measurement results can be universally accepted. To achieve this, measurement results need to be metrologically traceable to internationally recognized references. This required level of traceability is not attainable by blindly trusting the output of any GNSS device. For other externally supplied time references, similar considerations apply.

CHAPTER 9

Future prospects

TABLE OF CONTENTS

	<i>Page</i>
9.1 Introduction	204
9.2 Optical frequency standards	204
9.2.1 Single ion clocks	205
9.2.2 Neutral atom clocks.....	205
9.2.3 Specialized optical clocks	205
9.2.4 Optical local oscillators and frequency combs.....	205
9.2.5 Comparison of optical clocks.....	206
9.3 Commercial Optical Clocks.....	206
9.3.1 Optical Atomic Clocks Based on the Rb Two-Photon Transition	207
9.3.2 Optical Frequency Standards Based on I ₂ Molecular Transitions.....	207
9.3.3 Optical Clocks Based on the laser-cooled single ions or atoms.....	208
9.4 Applications.....	208
9.4.1 Optically steered time scale	208
9.4.2 Chronometric Geodesy.....	208
9.4.3 Information and communication systems	209
9.5 The redefinition of the SI second.....	210
References	210

9.1 Introduction

This Chapter describes the fields of standard time and frequency that have made rapid progress since the 2000s and are expected to develop further in the future – namely, optical frequency standards, optical clocks and their applications in time scale, geodesy and information and communication technologies.

Atomic clocks based on microwave transitions contributed greatly to the establishment of time and frequency standards in the latter half of the 20th century. However, atoms also possess optical transitions at much higher frequencies, and optical frequency standards based on these transitions have achieved dramatic progress in the 21st century. Since optical transitions are typically four to five orders of magnitude higher in frequency than microwave transitions, their fractional frequency uncertainty can be significantly reduced for the same level of absolute uncertainty.

Such high precision also reveals relativistic frequency shifts that were previously negligible, giving rise to new applications, such as using a stationary frequency source as a sensor for measuring gravitational potential.

Another important domain where time and frequency standards play a crucial role is information and communication technology (ICT). In wireless communications, same frequency between the transmitter and the receiver is a fundamental prerequisite for communication, and the importance of frequency standards has long been recognized while the need for phase/time alignment has become important relatively recently in the context of mobile networks. In wired communication, particularly in optical fibre systems, the unambiguous definition of optical frequency is also indispensable. In recent years, the importance of time synchronization within communication networks has been increasing rapidly.

9.2 Optical frequency standards

Optical frequency standards [Ludlow *et al.*, 2015; Fortier *et al.*, 2026] are based on electronic transitions of atoms or ions in the optical domain, in contrast to microwave clocks probing hyperfine transitions in the ground states. These transitions are chosen for their narrow natural linewidth, and their low sensitivity to environmental perturbations. Such narrow transitions can typically be inter-combination lines between the lowest lying single and triplet electronic spin states in divalent atoms (such as alkaline-earth like atoms e.g. Sr, Yb, Hg, Mg, Cd,...), or ionic species such as Al⁺, In⁺, or multi-polar transition (such as quadrupole transitions in Sr⁺, Ca⁺, Hg⁺, Lu⁺, or octupole transitions in Yb⁺).

These optical transitions, which span the visible spectrum from red to ultra-violet with sub-Hz linewidth, exhibit extreme quality factors on the order of 10^{15} . These large quality factors yield record fractional frequency instabilities from 10^{-15} to sub 10^{-16} at 1 s, offering the ability to resolve and control systematic effects in the 10^{-19} range after only hours or days of integration.

The clock transitions of optical clocks are probed by narrow linewidth, coherent laser sources, typically obtained by pre-stabilizing the frequency of the laser on an ultra-stable Fabry-Perot cavity. The frequency noise of these laser sources plays a crucial role in the performances of optical clocks: the coherence time of the laser is a factor limiting the maximum probing time of the clock transitions; the stroboscopic sampling of this noise by the operation cycle of the clock (known as the Dick effect) is a significant contribution to the final frequency noise of optical clocks. Because these two effects set the current instability of optical clocks, the improvement of ultra-stable cavity is an active domain of research (see § 9.2.3).

The Doppler effect is one of the main effects that limit the systematic uncertainty of state-of-the-art microwave frequency standards (atomic fountain clocks). Because the relative frequency shift arising from this effect only depends on the velocity of atoms, exploiting the fine accuracy permitted by optical transitions requires departing from the fountain geometry. Indeed, optical frequency standard rather make use of trapped ions or atoms, tightly confined over a size less than the optical wavelength, in the so-called Lamb-Dicke regime [Dicke, 1953]. This confinement ensures that the motion of atoms or ions is quantized, effectively changing the Doppler-broadened spectrum of free atoms with discrete resonances, the central one being unaffected by the motion of the atom [Wineland, 1979]. The methods used for this trapping classify optical clocks in two categories: single ion optical clocks in electromagnetic traps, or neutral atom clocks in optical lattices.

9.2.1 Single ion clocks

Ion clocks typically probe a single ion in a Paul trap, confining an ion using and oscillating RF electric potential. This tight confinement, associated with laser cooling results in trapping lifetimes reaching up to several months. The typical systematic effects that limit the accuracy of these clocks are the residual micro-motion in the oscillating trap-field, as well as the blackbody radiation frequency shift. But optimized trap design allows to finely control these effects. The excellent control of systematic effects in ion clocks, together with the availability of ion species with intrinsically low sensitivity to systematic effects such as Al^+ , In^+ or Lu^+ , results in the recent demonstration of ion clocks with total systematic uncertainty in the low 10^{-18} , or even the 10^{-19} range [Huntemann *et al.*, 2018; Marshall *et al.*, 2025]. However, ion clocks probe a single ion at a time because of the Coulomb interaction between ions. Consequently, their fractional frequency stability is limited by the fundamental Quantum Projection Noise associated in the 10^{-16} range at 1 s. To improve on this limit, new research topics include probing ions for extended durations, performing correlation spectroscopy using several clocks, or designing micro-fabricated ion traps able to trap tens of ions at a time. These multi-ion traps can for instance be achieved with In^+ , benefiting from their low quadrupole shift [Hausser *et al.*, 2025].

9.2.2 Neutral atom clocks

Using neutral atoms allows to escape the limitations arising from the Coulomb interaction between ions, hence opening the possibility to probe thousands or more atoms simultaneously, thereby decreasing the Quantum Projection Noise by several orders of magnitude. The counterpart is the difficulty to trap atoms in the Lamb Dicke Regime without perturbations to the transition frequency. In optical lattice clocks, this is achieved by confining laser cooled atoms in a powerful standing laser beam, the atoms being attracted to the intensity maxima of the standing wave by the dipole force. The light shift on the clock frequency induced by this powerful trapping laser is effectively cancelled at the “magic wavelength” [Ido, 2003], [Devianko and Katori, 2011; Ludlow and Ye, 2015] a particular wavelength of the trapping laser for which the polarisabilities of the two clock states are identical. Prominent systematic effects in optical lattice clocks are the residual lattice light shift, cold collisions between trapped atoms, and the black-body radiation shift. These effects are currently well characterized in the low 10^{-18} range. A key advantage of optical lattice clocks is the low Quantum Projection Noise, scaling down as the square root of the number of atoms simultaneously trapped. As such, the Dick effect remains an important limitation for optical lattice clocks, with demonstrated instabilities below 10^{-16} at 1 s. Nevertheless, optical lattice clocks offer interesting perspective to further improve their stability and accuracy, such as probing up to 10^6 atoms in shallow multidimensional lattices.

9.2.3 Specialized optical clocks

Other optical transitions are under consideration for optical clocks, mainly the low energy nuclear transition in the nucleus of ^{229}Th ions, lying in the UV (148 nm), or electronic transitions in various highly charged ions. These transitions offer numerous advantages for fundamental research, owing to their different nature from more conventional optical clocks, e.g. for the search for variation of fundamental constants [Peik *et al.*, 2021]. Also, nuclear transitions are more immune to environmental perturbations, allowing for probing a large number of ions in crystalline matrices without experiencing a prohibitive broadening of the transition. However, their practical applications remain limited by the cumbersome technology required to produce them and to probe their clock transition in the UV domain [Tiedau *et al.*, 2024].

9.2.4 Optical local oscillators and frequency combs

Ultra-stable lasers are used to probe the narrow transition of optical clocks. Most of them are obtained by locking the laser frequency on a high finesse, ultra-stable optical resonator composed of a pair of highly reflective mirrors contacted to a hollow spacer. In the locking process, the relative frequency fluctuations of the laser become limited by the relative length fluctuations of the cavity. An important effort has thus been dedicated to cancelling effects affecting the length of optical cavities: mechanical vibrations are reduced by a proper design of the cavity shape and support points, together with the use of vibration-insulating platforms; deformations resulting from thermal fluctuations are mitigated by the use of spacers made of Ultra-Low-Expansion Glass, of compensation rings placed at the back of the cavity mirrors, and by in-vacuum thermal shielding; the impact of electronic noise in the lock is reduced by using high finesse cavities ($> 10^5$). Once

these technical noise sources are made negligible, one faces the thermal noise of the material, resulting from mechanical dissipation processes. The thin reflective coating layer at the surface of the mirror is the dominant source of thermal noise, typically in the 10^{-16} range. Strategies to reduce this effect include the design of longer cavities, the use of structured reflective material with lower mechanical damping (semi-conductor based crystalline coatings), or large area cavity modes. Ultimately, silicon or sapphire cavities in cryogenic environments are an active research topic to further reduce the thermal noise, with instabilities now approaching 10^{-17} .

Although optical cavities have very low short term (< 100 s) instability, they remain highly sensitive to environmental fluctuations in the mid to long term. For this reason, a suitable drift compensation method must be implemented when probing optical transitions with cavity-stabilized lasers, e.g. based on comparisons with H-masers, or with high up-time operational optical clocks. Nevertheless, single crystal cryogenic silicon cavities can feature a low and predictable drift, well below the MHz/s mark. These cavities could be used in the future as a key ingredient for all-optical time scales, replacing H-masers.

State-of-the-art optical cavities are often built at telecom wavelength, taking advantage of ubiquitous laser sources and components available in this domain. Transferring the cavity's frequency stability is then ensured by an optical frequency comb, a pulse laser whose spectrum is composed of phase coherent lines with a regular spacing in the RF domain. Frequency combs have now matured into commercially available devices. A typical usage of frequency comb is to directly lock them on an ultra-stable laser: this effectively makes every single tooth of the comb an ultra-stable frequency reference, possibly extending over the whole visible spectrum (so-called super-continuum references). Optical frequency combs are thus used to compare optical clocks with different clock frequencies with negligible instability and inaccuracy (below 10^{-19}); but they can also be used to transfer the stability of optical sources (cavity-stabilized ultra-stable lasers, and optical clock signals) down to the RF domain.

9.2.5 Comparison of optical clocks

Comparing optical clocks is a crucial test to measure their frequency stability and to confirm their systematic evaluation. For the last decade, optical clocks have been compared with total uncertainties surpassing the realisation of the SI second. These comparisons showed agreements in the 10^{-17} to 10^{-18} range, but also sporadically exhibited disagreement up to 10^{-16} . This illustrates that regular, renewed clock comparison are a stringent test of their performances.

The first generations of optical clocks were remotely compared via micro-wave oscillators, either by comparison with primary frequency standards, or through GNSS signals. These comparison methods are nevertheless limited at the 10^{-16} level, hence they are not adapted to the current performances of optical clocks. To compare remote optical clocks at the continental scale, phase compensated optical fibre links can be deployed over thousands of kilometres. The bidirectional probing of these links allows to measure and compensate for phase perturbation occurring in optical fibres. The attenuation in optical fibre is a significant challenge for long distance links, which is circumvented by installing amplifiers and repeating stations. Phase compensated fibre networks can be operated over dark fibre, or over a dark channel in existing fibres used for telecommunications. The latter offer the advantage of having a better accessibility but require bypassing any existing unidirectional amplifier within the network [Schioppo *et al.*, 2021], [Koke *et al.*, 2021].

Transportable optical clocks allow the comparison of optical clocks that are not connected by optical fibre links. Several such clocks have been demonstrated in Asia, Europe, and America, with total uncertainties from 10^{-17} to 10^{-18} . The engineering effort conducted in several projects led to the recent availability of high performance commercial optical clocks.

9.3 Commercial optical clocks

With continued progress in research on optical frequency standards, the development of commercial optical clocks has begun in recent years (i.e. 2020s) [Schmittberger, Marlow and Scherer, 2021]. At present, however, their deployment remains more limited than that of microwave atomic clocks such as caesium clocks and hydrogen masers, since a certain level of specialized knowledge is still required for their operation and maintenance. As a result, widespread use in large-scale infrastructures such as telecommunication networks and data centres is yet to come.

Nevertheless, optical clocks based on room-temperature atoms stabilized to atomic or molecular transitions can achieve a level of frequency stability comparable to that of general-purpose hydrogen masers, while being implemented in compact enclosures, for example within a 19-inch rack occupying approximately 4U. It should be noted that, although these systems rely on optical atomic transitions, almost all practical users require microwave-frequency outputs. Therefore, a commercial optical clock is of limited practical value unless it provides a microwave-frequency output derived from the optical reference.

For this reason, commercial development of optical clocks was very limited during the early years of this century. The emergence of fibre-based optical frequency combs has fundamentally changed this situation, enabling compact and efficient conversion from optical to microwave frequencies, and thereby promoting the commercialization of optical clocks.

In addition, commercial clocks based on laser-cooled atoms or ions can achieve an unprecedented level of accuracy for commercially available devices, reaching uncertainties below 1×10^{-15} . At present, four types of optical atomic clocks operating in the optical domain have reached the stage of commercialization and are available for procurement. These systems exhibit performance equal to or exceeding that of caesium clocks and hydrogen masers. An overview of each type is provided below.

9.3.1 Optical atomic clocks based on the Rb two-photon transition

Alkali atoms exhibit a two-photon transition from the ground state $S_{1/2}$ to the $D_{5/2}$ state. Energetically, an intermediate $P_{3/2}$ state lies approximately halfway between these two levels and is strongly dipole-allowed with respect to both states. As a result, the two-photon transition can be excited relatively efficiently.

When the two-photon transition is driven by two counter-propagating laser beams, atomic recoil associated with photon absorption is cancelled, enabling Doppler-free spectroscopy. In this configuration, absorption resonances with linewidths of a few megahertz or less can be obtained. Furthermore, because the spectroscopy is Doppler-free, atoms with all velocity classes contribute to the signal, which is advantageous for signal strength. A review of optical atomic clocks based on the Rb two-photon transition is available in the literature [Obaze-Adeleke *et al.*, 2025].

Although the transition is described as relatively easy to excite, the laser frequency is detuned by approximately 1 THz from the $S_{1/2}$ – $P_{3/2}$ single-photon resonance. Consequently, when a single laser source is used to drive the two-photon transition, a high optical intensity is required. This leads to a significant AC Stark shift, and the long-term frequency stability is therefore strongly dependent on the stability of the optical power.

An alternative approach employs two optical fields resonant with the $S_{1/2}$ – $P_{3/2}$ (780 nm) and $P_{3/2}$ – $D_{5/2}$ (776 nm) transitions, respectively, with the reference frequency derived from their sum frequency. In this case, near-resonant spectroscopy can be performed with lower optical intensities. However, the spectroscopy is no longer fully Doppler-free, resulting in broader linewidths and a dependence on atomic temperature.

Both approaches are currently implemented in commercially available Rb optical frequency standards. The adopted scheme can generally be identified from the wavelengths of the optical outputs provided by the system.

In terms of performance, both types typically achieve a short-term stability of approximately 1×10^{-13} at 1 s averaging time, and a long-term stability in the 10^{-15} range. This level of performance is comparable to that of hydrogen masers, and Rb two-photon optical standards are therefore considered promising candidates as potential alternatives to hydrogen masers.

9.3.2 Optical frequency standards based on I_2 molecular transitions

Molecular iodine (I_2) exhibits a large number of absorption lines, some of which coincide with the emission wavelengths of gas lasers and solid-state lasers. In the 20th century, iodine-stabilized He–Ne lasers operating at suitable transitions were developed and commercialized.

In the final decade of the 20th century, solid-state Nd:YAG lasers with low phase noise became available. By stabilizing the output frequency of such lasers to iodine molecular transitions, systems with improved short-term stability were developed. Initial developments were carried out at research institutes in the early 2000s, and in 2020s, compact commercial systems fitting within a 19-inch rack occupying approximately 3U have become available in some countries [Roslund *et al.*, 2024].

With respect to performance, the long-term stability (on the order of one day) is comparable to that of Rb two-photon optical standards. Owing to the use of a solid-state laser as the primary oscillator, the short-term stability is superior, typically reaching approximately 3×10^{-14} .

9.3.3 Optical clocks based on the laser-cooled single ions or atoms

Two kinds of optical clocks based on laser-cooled ions or atoms have been commercialized around 2025. The electric quadrupole (E2) transition in Yb^+ ions was among the first systems developed for optical single-ion clocks. Through public–private collaboration in Germany, this technology has been successfully commercialized and is now available for procurement. The achievable systematic uncertainty is at the level of 10^{-17} [Stuhler *et al.*, 2021]. Optical lattice clocks generally involve more complex system architectures than single-ion clocks. Nevertheless, through public–private collaboration in Japan, the commercialization of optical lattice clocks based on strontium atoms has been achieved, and such systems are now available on a commercial basis [Muramatsu *et al.*, 2024].

9.4 Applications

9.4.1 Optically steered time scale

The maturation of optical frequency standards opens the possibility of using them for frequency steering of local time scale signals [Yao *et al.*, 2018]. Although UTC is a paper clock, NMIs realize and disseminate a real-time signal, UTC(k), to society, which is steered to remain consistent with UTC. Because the offset between UTC and UTC(k) is reported with a latency of about one month in Circular T issued by the BIPM, steering based solely on Circular T is inherently inefficient.

International Atomic Time (TAI) is generated so that its scale interval is identical to the SI second within the uncertainty of clocks used as reference (see Chapter 5). Therefore, with respect to the scale interval, an NMI possessing a high-accuracy local frequency reference can generate an accurate time scale by steering its scale interval to match that of the local reference. While a Cs fountain clock operating in high-accuracy mode requires nearly one day of averaging to evaluate the fractional uncertainty of a hydrogen maser (HM) at the mid- 10^{-16} level, an optical clock can achieve comparable stability within only a few hours of operation, limited primarily by the HM. This enables rapid calibration of the UTC(k) scale interval, making it possible to perform steering through regular intermittent operation of relatively short runs [Hachisu *et al.*, 2018]. Using an optical oscillator instead of a microwave flywheel would ideally allow the realization of a purely optical time scale. An attempt using a cavity-stabilized laser has already been reported [Milner *et al.*, 2019].

Consequently, UTC(k) realizations in which the scale interval is steered using optical clocks rather than traditional Cs fountain standards are beginning to emerge. However, continuous 24/7/365 operation of a single optical clock remains technically challenging, and achieving redundancy with multiple optical clocks to ensure uninterrupted signal delivery is not yet straightforward. Nevertheless, the 24/7 availability of a standard frequency derived from an optically generated stable time scale would be highly beneficial.

9.4.2 Chronometric geodesy

In relativity, two distinct concepts of time exist: proper time and coordinate time. The one-second interval realized by an atomic clock corresponds to proper time, whereas International Atomic Time (TAI) is defined within the framework of coordinate time. Therefore, the timescale intervals produced by national metrology institutes' atomic clocks are converted into coordinate-time intervals by taking relativistic effects into account before being used for the calibration of the TAI scale rate.

The conversion factor is determined by two primary effects: the special relativistic effect due to the velocity of the clock and the general relativistic effect due to the gravitational potential at the clock's position. The latter has a particularly large influence. A height difference produces a fractional frequency shift of gh/c^2 , corresponding to a fractional frequency shift of approximately 10^{-16} per one meter of height difference – comparable to the present uncertainty level of primary frequency standards.

Optical frequency standards achieve uncertainties two orders of magnitude smaller, reaching the 10^{-18} level, and thus offer the possibility of serving as sensors for directly measuring gravitational potential. This new application, known as chronometric geodesy, was once a theoretical idea but is now a realistic research field [Wu *et al.*, 2019].

The advantages of optical clocks as sensors arise not only from their extremely small uncertainty but also from their superior short-term stability, which enables detection of dynamic variations in the gravitational environment that were beyond the capability of microwave standards.

While gravimeters measure the gradient of the gravitational potential – that is, the local acceleration due to gravity – optical clocks measure the potential itself. Furthermore, optical clocks can only observe frequency differences between two clocks; therefore, by connecting two distant clocks via an optical fibre link and measuring their frequency difference precisely, the gravitational potential difference between the two sites can be determined. This approach can be applied not only to evaluate static height differences but also to detect temporal variations in gravitational potential through long-term observations.

Such measurements are expected to reduce cumulative errors in height determination. When heights measured by optical clocks are used as reference points (pivots) for levelling networks, the accumulated errors in conventional levelling can be greatly mitigated. Moreover, this technique allows determination of potential differences between locations separated by oceans, something impossible by conventional levelling.

In addition, monitoring time variations of potential differences with optical clocks could enable the detection of mass redistributions within the Earth, such as magma movement or tectonic plate motion. The gravitational effect of a mass varies inversely with the square of the distance from that mass, while its effect on gravitational potential varies inversely with the first power of distance. Consequently, gravimeters are more susceptible to noise from near-surface mass changes (e.g. groundwater), whereas optical clocks, acting as potential meters, are more sensitive to deeper mass variations within the Earth.

9.4.3 Information and communication systems

Synchronization in the context of information and communication systems is typically provided within a localized network, in which an approximation of UTC is obtained from a GNSS-disciplined oscillator and distributed using such protocols as NTP (Network Time Protocol) and PTP (Precision Time Protocol), see Chapter 6. Such a localized network could be a segment of a telecommunication backbone (constrained to a building, a region, or entire country), a power substation, a data centre, or TV broadcasting studio. The synchronization requirements in such systems are divided into two parts:

- 1) Accuracy of synchronization within the network, discussed in Chapter 6.
- 2) Accuracy of approximation of UTC, discussed in Chapters 5 and 6, as well as referred to in Chapter 8.

For example, the most stringent requirements defined as of writing (2025) for telecommunication networks require enhanced Primary Reference Time Clock (ePRTC) to provide UTC approximation to within 30 ns. A widespread reliance on GNSS for providing such UTC approximation is reaching its performance limits and has been identified as vulnerability. Large-scale space weather events such as solar flares, or intentional interference such as jamming or spoofing, can cause GNSS reception failures and disrupt communications. This vulnerability is not limited to telecommunication networks but also affects critical infrastructures such as broadcasting, power grids, and finance, leading to an international recognition of the need for terrestrial time distribution systems independent of GNSS [Divis, 2019], [US 2020]. Therefore, alternative and/or complementary solutions to GNSS are being implemented. These include fibre-based direct connections to sources of UTC(k) as well as enhanced holdover capabilities.

In this context, exceptional stability of optical clock, once available for industry applications, can address future requirements in terms of accuracy of UTC approximation and duration of holdover when GNSS is not available.

9.5 The redefinition of the SI second

With the advent of optical clocks, now surpassing their microwave counterpart by at least a solid order of magnitude, the possibility to redefine the SI second, replacing the Cs transition by optical transitions, is currently being addressed by the time and frequency community. However, raw performance is not the only step to make such a redefinition possible. Indeed, optical clocks have to further confirm their systematic uncertainty by repeated inter-comparisons between independent clocks, and to show that they can replace Cs primary standard in an operational way. A significant step toward this latter goal has already been achieved by the first calibration of the scale interval of international atomic time (TAI) with optical clocks, dating back to 2014 for the first calibration and to 2019 for the first steering calibration. In initial stage, intermittent operations of optical clocks provide a calibration of a local flywheel oscillator such as hydrogen maser, but some optical clocks have achieved uptime ratios of more than 90%, similar to Cs fountain standards. Since then, as much as eleven optical frequency standards have produced TAI calibration with time series up to one year.

Building on these progresses, the working group of the Consultative Committee for Time and Frequency has crafted a roadmap towards the redefinition of the SI second [Dimarcq *et al.*, 2024]. This roadmap introduced quantitative criteria that optical clocks must achieve before deemed ready for the redefinition, in terms of performances, cross comparisons, link to the current SI second, and steering of TAI. If the reader is interested to follow the development in time, it is suggested to enter “BIPM Redefinition of the second” in the browser. The webpage is continuously updated. As of 2026, the above objectives are not complete, but advanced enough to envision a redefinition as early as 2030 or 2034. For this a consensus will have to be reached not only about the atomic species that will be chosen for the redefinition, but also about either single or multiple transitions [Lodewyck, 2019] are suitable for the redefinition. It is not straightforward as of 2026 since many species and even technologies (ion or neutral atom clocks) are currently evenly matched.

References

- Derevianko A and H Katori H [2011] Colloquium: *Physics of optical lattice clocks*, Rev. Mod. Phys. Vol. 83, 331.
- Dicke R. H. [1953], *The effect of collisions upon the Doppler width of spectral lines*, Phys. Rev., Vol. 89, 472-473.
- Dimarcq *et al.*, [2024] *Roadmap towards the redefinition of the second* Metrologia Vol 61, 012001.
- Divis D. [2019] *Homeland Security Says PNT a ‘National Critical Function’*, Inside GNSS (2019).
- Fortier T. M., Luiten A. N., Margolis H. [2026] *Optical atomic clocks: defining the future of time and frequency metrology*, Optica Vol. 13, p. 143-163, <https://doi.org/10.1364/OPTICA.575770>.
- [Hachisu H., Nakagawa F., Hanado Y. and Ido T. [2018] *Months-long real-time generation of a time scale based on an optical clock*, (online) WWW.Nature.com/Scientificreports | (2018) 8:4243 | DOI:10.1038/s41598-018-22423-5.
- Hausser H. N. *et al.*, [2025] $^{115}\text{In}^+ - ^{172}\text{Yb}^+$ Coulomb Crystal Clock with 2.5×10^{-18} Systematic Uncertainty, Phys. Rev. Lett. Vol. 134, 023201.
- Huntemann *et al.*, [2016] *Single-Ion Atomic Clock with 3×10^{-18} Systematic Uncertainty*, Phys. Rev. Lett. Vol. 116, 063001.
- Ido T and Katori H [2003] *Recoil-Free Spectroscopy of Neutral Sr Atoms in the Lamb-Dicke Regime*, Phys. Rev. Lett. Vol. 91, 053001.
- Koke S. *et al.*, [2021] *Validating frequency transfer via interferometric fiber links for optical clock comparisons*, New Journal of Physics Vol. 23, 093024.
- Lodewyck J. [2019] *On a definition of the SI second with a set of optical clock transitions*, Metrologia Vol. 56, 055009.
- Ludlow A. D. *et al.*, [2015] *Optical atomic clocks*, Rev. Mod. Phys. Vol. 87, 637.

- Ludlow A. D. and Ye Y. [2015] *Progress on the optical lattice clock*, Comptes Rendus. Physique, Vol. 16 (5), pp. 499-505.
- Milner W.R. *et al.*, [2019] *Demonstration of a Time scale Based on a Stable Optical Carrier* Phys. Rev. Lett. Vol. 123, 173201.
- Muramatsu T *et al.*, [2024] *Compact and robust laser system for a transportable strontium optical lattice clock*, J. Phys. Conf. Series Vol. 2889, 012042.

CHAPTER 10

Summary and conclusions

TABLE OF CONTENTS

	<i>Page</i>
10.1 Introduction	214
10.2 Clocks and oscillators	214
10.2.1 Devices based on a mechanical resonance	214
10.2.2 Microwave atomic clocks	214
10.2.3 Optical atomic clocks	215
10.2.4 The future	215
10.3 Clock ensembles	215
10.3.1 Incorporating primary reference oscillators	215
10.3.2° The Internet of Things (IoT)	215
References	217

10.1 Introduction

The chapters in this Handbook discuss many different types of clocks and oscillators, and methods for describing the stability and accuracy of these devices. In addition, the Handbook includes a discussion of how different devices can be combined into clock ensembles, in which the performance of the ensemble is better (in a statistical sense) and more reliable than any one of its members. Finally, the Handbook presents techniques for distributing time and frequency information. The Handbook presents an introduction and a general overview of each of these subjects, and the readers are encouraged to consult the very extensive list of references for a more in-depth discussion of each topic. The next sections summarize the current state of each of these major subjects and include a prediction of the likely developments in the near-term future.

10.2 Clocks and oscillators

10.2.1 Devices based on a mechanical resonance

Oscillators and clocks that use the mechanical properties of quartz crystals or micro-electronic mechanical systems (MEMS) as the frequency reference are very widely used because they are small, cheap, and require a very modest power source. The oscillation frequency of the device is sensitive to fluctuations in the ambient temperature and other environmental effects, so that the temperature of the device may be stabilized (OCXO, oven-controlled crystal oscillator) or the devices can include additional components to minimize the impact of temperature variations (TCXO – temperature compensated crystal oscillator). Special-purpose TCXO-based devices have reduced sensitivity to accelerations (Low-g-sensitivity) and very low phase noise. The frequency of a voltage-controlled oscillator (VCO) can be adjusted over a relatively narrow range by applying a voltage to a voltage-sensitive reactive element in the oscillator circuit.

10.2.2 Microwave atomic clocks

Atomic clocks lock the frequency of an oscillator to the frequency associated with the transition between two states of the atom. The oscillator is typically a voltage-controlled device described in the previous section. Although any transition could be used for this purpose in principle, there were engineering advantages to using the hyperfine transition in the ground states of hydrogen, rubidium, or caesium, and devices based on these frequencies are still the most widely used atomic clocks. The output frequency is usually generated by the oscillator – the atoms act as a frequency discriminator to control the oscillator frequency in the passive devices that are based on a transition frequency in rubidium or caesium. Hydrogen masers are often configured as active devices – state-selected hydrogen atoms are injected into a high-Q resonant cavity, and the reference for the local oscillator is produced by the stimulated emission of the atoms in the cavity. The stability of the reference frequency is degraded if the frequency of the hyperfine transition is different from the resonant frequency of the cavity, and active masers include some form of “auto-tuning” to minimize this effect.

A number of different methods are used to lock the oscillator to the atomic transition. The method of controlling the oscillator has a finite bandwidth, so that the statistical characteristics of the devices are dominated by the free-running characteristics of the oscillator at very short times. The long-term stability is derived from the characteristics of the atomic transition, but this stability is often not realized in practical devices because of interactions between the atoms and the surroundings, which perturb the energies of the clock states. These interactions may change slowly in time, so that atomic clocks may exhibit long-term frequency drifts. These drifts are significant in hydrogen masers and most rubidium devices because the atoms are in a cell, and collisions with the walls are important. The cells in rubidium devices are usually filled with a buffer gas to minimize these collisions, and the cells in hydrogen masers have special coatings. The drifts are reduced by these measures, but are usually not removed completely, so that the drift in frequency is often the dominant limitation to the long-term stability of rubidium and hydrogen devices.

10.2.3 Optical atomic clocks

All of the atomic clocks described in the previous section produce microwave frequencies, and this was an important characteristic when they were designed because the output signals could be processed by the electronic components of that time. A higher frequency would be potentially more favourable, other things being equal, and this consideration was important in the choice of the hyperfine transition in the ground state of caesium as the reference frequency.

The development of optical combs (Fortier and Baumann, 2019) changed this calculation because they provided a way of linking an optical frequency to a much lower frequency, which could then be processed electronically. The increase in the frequency of the reference by about a factor of 10^5 would more than compensate for the increase in the complexity of the device, and results to date confirm this better stability.

The link between optical and microwave frequencies that is provided by an optical comb is not very sensitive to the value of the optical frequency, so that there is no clearly best transition for a re-definition of the standard of frequency in 2025 (Dimarcq *et al.*, 2024).

10.2.4 The future

The future developments of clocks and oscillators is likely to focus on two areas: (1) improvements in the stability and accuracy of devices that minimize size, weight and power (SWaP). The chip-scale atomic clock is an example of this development; and (2) improving the accuracy and reliability of frequency references based on optical transitions that will lead to a re-definition of the standard of frequency, probably by 2036.

10.3 Clock ensembles

Ensembles of clocks are widely used in National Metrology Institutes, timing laboratories, and other facilities to generate a time and frequency reference that has better statistical characteristics and is more reliable than any one of the contributing devices. The ensemble algorithm also provides the machinery for combining data from devices with very different statistical characteristics. In addition, the ensemble algorithm reduces the impact of the failure of a single device.

10.3.1 Incorporating primary reference oscillators

The input data to the current ensemble algorithms are generally equally-spaced measurements of the time differences between the member clocks. The ensemble algorithm is usually not well suited to incorporating data from a primary reference device directly into the ensemble calculation, because the primary device often does not run continuously as a clock and may be available only on an irregular basis. The design of hybrid algorithms that can accept periodic time-difference data and aperiodic frequency data will be an important subject for further study, because these combinations are likely to become more common in the future. One possible configuration would be an ensemble composed of hydrogen masers to provide short-term stability and a primary reference device to provide long-term frequency accuracy. The ensemble algorithm would have to support both periodic time-difference measurements and aperiodic frequency inputs in this configuration. An alternative would be to use a more traditional ensemble composed of just the masers and to steer the frequency of the devices to the aperiodic data from the frequency standard.

10.3.2 The Internet of Things (IoT)

Clocks and oscillators are now embedded in a very wide range of ‘smart’ devices such as consumer appliances and process-control equipment. The time-sensitive applications in these devices do not require the accuracy and stability of atomic clocks, and simpler quartz oscillators are usually adequate. These devices often must cooperate with each other, so that it is important that the clocks share a common ensemble time. An ensemble in which all the members are peers (rather than a server-client configuration) has no single point of failure, and it is easier to protect against outside attacks.

Unfortunately, the ensemble members are simple oscillators with relatively poor deterministic and statistical characteristics, so that models of the clocks are not very useful in predicting the evolution of the time and frequency of the devices. In addition, no device in the network may have adequate computing power to compute the ensemble average, and there may be no mechanism for steering the outputs of the individual members to this system time even if it could be computed. Traditional ensemble algorithms are not well suited to this application, and future developments in ensemble algorithms will be needed to address this type of configuration.

10.4 Networks of clocks

Determining the transmission delay is important for synchronizing the times of clocks that are not at the same location, and for transmitting timing information in general, and the fluctuations in the delay are important for exchanging frequency data between remote devices. The delay can be measured by transporting a calibrated device between the two stations, although this method is not practical for measuring its stability. The impact of the path delay between the systems that are being compared can be attenuated by the common-view method, which depends on the approximate equality of the delays between the two stations and a common transmitter. The magnitudes of these delays are not important to the extent that they are common to both paths and cancel in the difference. The time and frequency characteristics of the source are attenuated for the same reason. The two-way method estimates the one-way delay between the two systems as one-half of the measured round-trip value. The accuracy of the estimate of the one-way delay is limited by the symmetry of the transmission delays in the two directions and not by the magnitude of the delay.

The difficulties of comparing the times of the two clocks in these configurations is made more complicated when the members are not stationary – when they are operating on orbiting micro-satellites for example. The orbits are not as predictable as the orbits of satellites in higher orbits, and each satellite can be observed for only a limited time on each pass. The geometrical path delays between satellites are continuously changing, and the messages between them are affected by Doppler shifts and possibly by gravitational red shifts. In addition, the ensemble can have hundreds of members, and the membership is constantly changing. These configurations are likely to become more important in the near future, so that constructing ensembles of these clocks will be challenging. One possibility would be to construct sub-ensembles that are composed of only nearby members in similar orbits, but the standard ensemble algorithms do not have this flexibility.

In spite of the long-time interval that has elapsed since the release of the first issue of this ITU-R Handbook in 1997, the basic rules, statements and recipes have not changed. The increase in the importance of atomic clocks and high-performance network synchronization in everyday daily life, however, is astounding. The use of GNSS for navigation, location and time dissemination is widespread. Now many other global navigation satellite systems can be seen in daily use. On all systems, research and development are ongoing to improve resilience, safety and quality assurance. In telecommunications applications, the request for network synchronization for high-speed internet traffic as well as in mobile communications has led to ever tighter requirements on the deployed equipment. Last but not least, the safe and reliable distribution of electric energy generated from multiple individual installations challenges network management and synchronization, and this would not be possible without increased use of time and frequency signal dissemination and reception.

Whether the current version of this Handbook is capable to answer all questions and education needs of incoming engineers and decision makers? This needs to be seen and is to be judged by those who read the Handbook for such purpose. One further source of tutorial material is available: if the interested reader enters “BIPM capacity building e-learning” in the browser, the link gives access to abundant material, inter-alia from time and frequency experts.

On the other hand, the authors are confident that the next generation of scientists will add additional topics, which are overlooked today, to a future Handbook release.

References

- Dimarcq n, Gertsvolf m, Mileti g, Bize S, Oates C W, Peik E, Calonico D, Ido T, Tavella P, Meynadier F, Petit G, Panfilo G, Bartholomew J, Defraigne P, Donley E A, Hedekvist p o, Sesia i, Wouters M, Dubé P, Fang F, Levi f, Lodewyck j, Margolis h s, Newell D, Slyusarev S, Weyers s, Uzan j-p, Yasuda M, Yu d h, Rieck c, Schnatz h, Hanado y, Fujieda m, Pottie p-e, Hanssen j, Malimon a and Ashby n [2024], *Roadmap towards the redefinition of the second*, Metrologia, **61** 012001, Doi: 10.1088/1681-7575/ad17d2.
- FORTIER T and BAUMANN E [2019], *20 years of developments in optical frequency comb technology and applications*, Communications Physics, **2**, 153 (2019). Doi: 10.1038/s42005-019-0249-y.
- HÄNSCH T W [2006] *Nobel lecture: Passion for precision* Rev. Mod. Phys. 78 pp 1297 – 1309 (DOI 10.1103/RefModPhys.78.1297)
-

International
Telecommunication
Union

Place des Nations
CH-1211 Geneva 20
Switzerland

ISBN 978-92-61-43261-4



9 789261 432614

Published in Switzerland
Geneva, 2026

Photo credits: Adobe Stock