



Time Traceability for the Finance Sector

Fact Sheet

Version 2.0
14 August 2018

All correspondence related to this document should be addressed to:
Dr Leon Lobo
National Physical Laboratory
Hampton Road | Teddington | Middlesex | TW11 0LW
+44 (0) 20 8943 6383
leon.lobo@npl.co.uk

Contents

1. Introduction	4
2. Technical standards for business clock synchronisation	4
3. Global timekeeping	5
4. Traceability and uncertainty	6
5. Sources of time	8
6. Obtaining traceability using GPS	10
7. Improving resilience	12
8. Conclusions	13
Appendix 1: Terminology of time measurement	15
Appendix 2: Example of NPL monthly GPS monitoring bulletin	16

Abbreviations

BIPM	International Bureau of Weights and Measures
Cs	Caesium
DCF77	German standard-frequency and time service, transmitted on 77.5 kHz
ESMA	European Securities and Markets Authority
GNSS	Global navigation satellite system
GPS	Global Positioning Service
GPSDO	GPS-disciplined oscillator
LAN	Local area network
MiFID II	Markets in Financial Instruments Directive II
MSF	UK standard-frequency and time service, transmitted on 60 kHz
NIST	National Institute of Standards and Technology, Boulder, Colorado, USA
NPL	National Physical Laboratory, Teddington, UK
NTP	Network Time Protocol
OEXO	Oven-controlled crystal oscillator
PPS	Pulse per second; usually 1 PPS
PTP	Precision Time Protocol; PTP version 2 is defined in standard IEEE 1588-2008
Rb	Rubidium
RTS	Regulatory Technical Standard; issued by ESMA
SI	International System of measurement units
TCXO	Temperature-controlled crystal oscillator
UKAS	United Kingdom Accreditation Service
USNO	United States Naval Observatory, Washington DC, USA
UTC	Coordinated Universal Time
UTC(k)	Realisation of UTC maintained by the contributing institute with abbreviation <i>k</i>
UT1	Universal Time 1; time based on the Earth's rotation

1. Introduction

This document describes the key concepts that underpin the requirements for accurate time-stamping of transactions in the finance sector. The Regulatory Technical Standard for clock synchronisation is summarised, and subsequent topics covered include the international time system and the importance of UTC, the concepts of traceability and uncertainty for time measurements, sources of time, and methods to improve resilience. Key terms used in time measurement are summarised in Appendix 1.

Accurate and trusted timing plays a critical role in financial markets, underpinning the time stamping of trades, the synchronisation of computer systems and the measurement of network latency for process optimisation. The rapid expansion of computer-based trading has increased the need for synchronisation of trading systems and traceability to a common reference time scale, to help prevent trading irregularities and to aid forensics. To encourage this, a number of regulatory bodies now require high precision traceable time-stamping to help understand activity across trading venues, working toward a consolidated audit trail.

Time stamps created by different systems or networks can only be compared meaningfully if they are based on the same reference. For time, the global reference is Coordinated Universal Time (UTC): the time scale that underpins GPS, broadcast time signals and all other precise time services. UTC is generated by the International Bureau of Weights and Measures (BIPM) through an international collaboration involving around 75 timing institutes. Each of these institutes maintains a physical representation of UTC, called generically a UTC(*k*) time scale (where *k* is the abbreviation for the institute), which can act as the reference for national or regional time dissemination services.

Whatever time source and dissemination method are used, they must provide traceability back to UTC. Traceability requires a continuous chain of comparisons with known uncertainties, all of which must be documented. The GPS satellite signals alone do not readily provide traceability to UTC, but users can demonstrate traceability by obtaining GPS monitoring bulletins from one of the regional UTC(*k*) timing centres.

There are two other requirements for demonstrating time traceability. One is for the timing equipment to be calibrated, so that its unknown internal delays do not bias its time output. The other is for the equipment to be monitored continuously, so that any fault or anomalous behaviour can be detected and the time output not used until the equipment is working correctly again. The calibration evidence and monitoring results should be archived so that the status of the timing equipment at any point in time can be verified at a later date.

2. Technical standards for business clock synchronisation

The European Securities and Markets Authority (ESMA) has issued a series of technical standards in support of the MiFID II (Markets in Financial Instruments Directive II) regulations, which have been approved by the European Commission and came into effect on 3 January 2018. One of these supporting Regulatory Technical Standards, RTS 25, deals with clock synchronisation. Its technical regulations are described under four articles, which are summarised below.

Article 1 states that business clocks that give the timestamp for any reportable event should be synchronised to Coordinated Universal Time (UTC), using either a link to one of the timing centres

maintaining a UTC(*k*) realisation of UTC or the timed signals disseminated by GPS or another satellite system. The concept of traceability is mentioned earlier in RTS 25.

Article 2 describes the level of accuracy (maximum divergence from UTC) that should be achieved by the operators of trading venues, taking into account the gateway to gateway latency of their trading systems. These are given in table 1.

Latency time	Max. divergence from UTC	Timestamp granularity
>1 millisecond	1 millisecond	1 millisecond or better
=< 1 millisecond	100 microseconds	1 microsecond or better

Table 1: Level of accuracy for operators of trading venues.

Voice trading systems or those needing human intervention need only ensure 1 second maximum divergence from UTC, and record times with a 1 second granularity.

Article 3 defines the level of accuracy that apply to members or participants of trading venues, as summarised in table 2.

Type of trading activity	Max. divergence from UTC	Timestamp granularity
High frequency algorithmic	100 microseconds	1 microsecond or better
Voice trading systems	1 second	1 second or better
Human intervention; non-algorithmic	1 second	1 second or better
Concluding negotiated transactions	1 second	1 second or better
Other	1 millisecond	1 millisecond or better

Table 2: Level of accuracy for members or participants of a trading venue.

Article 4 specifies the need to demonstrate traceability to UTC by documenting system design, functioning and specifications, and to identify the exact point at which a timestamp is applied. The traceability system should be reviewed at least once a year to ensure compliance with the regulations.

3. Global timekeeping

Accurate timekeeping worldwide is based on a single reference time scale, known as Coordinated Universal Time and abbreviated in all languages as UTC. The International Bureau of Weights and Measures (BIPM), based in Paris, computes UTC in monthly blocks by combining data from around 400 continuously-operating atomic clocks located in national timing centres. The duration of the UTC second is fine-tuned using measurements from a small number of primary frequency standards to ensure that it remains as close as possible to the standard unit of time, the *second*, defined in the International System of Units (SI).

UTC is computed monthly, so does not exist in real-time. Each institute contributing clock data to the BIPM maintains its own physical realisation of UTC, such as UTC(NIST) in the USA and UTC(NPL) in the UK, which are known collectively as the UTC(*k*) time scales. These national time scales are adjusted so that they remain close to UTC, usually within 1 microsecond, and in some cases the difference is kept below 10 nanoseconds. They are traceable to UTC, and serve as the reference standards for all accurate time measurements globally.

UTC is based on atomic clocks, giving it great stability and accuracy. However, it slowly deviates from time based on the Earth's rotation, which fluctuates unpredictably over time and experiences a long-term slowing due to friction caused by the tides. The increasing difference between UTC and Earth rotation time (UT1, which can be thought of as a more precisely-defined version of Greenwich Mean Time) is occasionally corrected by a 1-second adjustment in UTC known as a leap second. A leap second is inserted into, or removed from, the final UTC minute of either 31 December or 30 June, so occurs at the same instant world-wide. A decision on the need for the next leap second is announced a little under 6 months in advance.

Between the introduction of the leap second mechanism on 1 January 1972 and 1 January 2019 there have been 27 leap second insertions into UTC. They have so far all been positive, so the final minute of the day has 61 seconds rather than 60.

4. Traceability and uncertainty

The concept of *traceability* for a measurement of time (sometimes referred to as *epoch* or *time-of-day*) requires a continuous chain of comparisons extending from the generation of a time stamp or synchronisation of a clock, back through the time distribution to one of the UTC(*k*) time scales, and so to the reference time scale UTC¹. Figure 1 provides a graphical illustration of the measurement steps linking a timestamp back to the reference time scale UTC, using the NPLTime[®] service as an example.

Each comparison in the chain will be carried out with some inherent level of inaccuracy, which can be quantified by an estimated standard deviation known as the *standard uncertainty* of the measurement. The standard uncertainties of every link in the dissemination solution can then be combined to give the total standard uncertainty of the time stamp or synchronisation at the endpoint, where it is available to the user.

Understanding the time dissemination chain and calculating the total uncertainty is an essential requirement for demonstrating traceability. Each link must be documented, along with the associated uncertainty evaluation, and the combined standard uncertainty of the timestamp or time output must be determined and recorded.

The standard uncertainty of a measurement will generally have 2 components. The random fluctuations in the measurement result can be expressed as a statistical value termed the Type A uncertainty, u_A . It is usually determined by taking repeated measurements to analyse their distribution. The likely range of systematic bias or offset in the measurement result is quantified by the Type B uncertainty, u_B . It is usually based on scientific judgement using all available evidence, including previous measurements, calibration results and manufacturers' specifications. If these two components are independent, a total uncertainty u for the measurement can be calculated as their root-sum-square combination:

$$u = \sqrt{u_A^2 + u_B^2}$$

¹ "Metrological and legal traceability of time signals", D. Matsakis, J. Levine and M. Lombardi, Proceedings of the 49th Annual Precise Time and Time Interval Systems and Applications Meeting, Reston, Virginia, January 2018: <https://tycho.usno.navy.mil/papers/ts-2018/Traceability-of-Time-Signals.pdf> or <https://tf.nist.gov/general/pdf/2941.pdf>

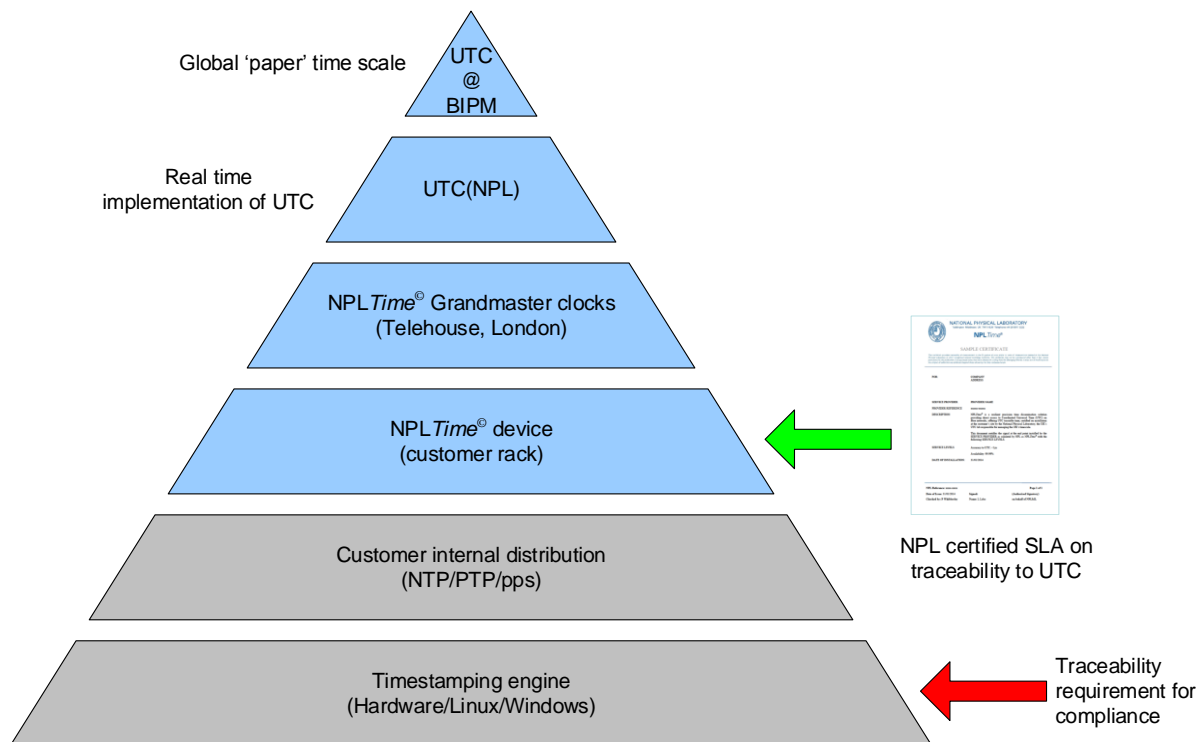


Figure 1: Example of a traceability pyramid showing the comparisons linking a timestamp to UTC, taking the NPLTime® service as an example.

As a simple example of uncertainty evaluation, a *uncertainty budget* is shown in table 3 for an NTP service provided by an institute maintaining a UTC(k) time scale. The service is assumed to feed into a local area network synchronised internally using NTP, and to culminate in a Windows-based time-stamping engine.

Link in distribution chain	Standard uncertainty
Sync of NTP server to UTC(k)	0.000001 s
NTP distribution over the internet	0.05 s
NTP distribution over LAN	0.01 s
Timestamping engine	0.1 s
Total uncertainty	0.11 s

Table 3: Example of an uncertainty budget for an NTP service. In this case there are no systematic offsets and all of the components of the uncertainty budget are statistical terms.

The combined standard uncertainty² of the traceability chain between the user and UTC can generally be calculated as the quadrature sum of the individual uncertainties of each link. Taking the values in table 3, the combined standard uncertainty can be calculated as follows:

$$\begin{aligned}u_c &= \sqrt{a^2 + b^2 + c^2 + \dots} \\ &= \sqrt{(0.1)^2 + (0.01)^2 + (0.05)^2 + (0.000001)^2} \text{ s} \\ &= 0.11 \text{ s}\end{aligned}$$

Note that to determine a combined uncertainty value, the uncertainty of each link in the traceability chain has to be measured or otherwise determined, for example in a calibration carried out by an accredited institute.

In addition, it is not sufficient simply to evaluate the uncertainty budget once, and assume it remains valid indefinitely. There must be continuous monitoring at key points of the distribution to demonstrate that the traceability chain remains valid and the time available at each point is within the calculated uncertainty at all times when measurements are being taken. The monitoring records must therefore be stored so that they can be checked at some later time to confirm that the equipment was operating correctly.

5. Sources of time

Time is readily available from a broad range of sources and over many different distribution protocols, with considerable variations in the resolution, stability and accuracy that they offer the user. All disseminated time signals and services are derived from UTC, but in practice it can be difficult or impossible to validate the entire distribution chain well enough to calculate the uncertainty of every link and establish traceability.

Commonly used sources of time include:

1. **Free-running clock.** The most commonly encountered types of clock are quartz crystal oscillators, usually temperature-controlled (TCXO) or oven-controlled (OCXO), rubidium (Rb) oscillators, and caesium (Cs) clocks. A clock on its own, even a caesium atomic clock, is generally not a viable source of time-of-day. Even if synchronised against a reliable UTC reference, its time and frequency offsets from UTC will tend to increase over time due to the presence of a range of noise processes within the clock. The characteristics of this noise vary considerably from one type of clock to another, and with the measurement time. Averaging an ensemble of several clocks can give some improvement, but does not eliminate the problem as the average will itself drift.
2. **GPS.** The Global Positioning System (GPS) and other global navigation satellite systems (collectively termed GNSS) are great sources of accurate time. In a typical installation a GPS-disciplined oscillator will disseminate time across a local-area network using NTP or PTP (discussed below). However, such solutions need to be implemented with care³. Potential

² A Beginner's Guide to Uncertainty of Measurement, Stephanie Bell, NPL, March 2001:
<http://www.npl.co.uk/publications/guides/a-beginners-guide-to-uncertainty-of-measurement>

³ Global Navigation Space Systems: reliance and vulnerabilities, The Royal Academy of Engineering, Mar 2011:
<http://www.raeng.org.uk/publications/reports/global-navigation-space-systems>

sources of error include multipath reflections of the satellite signals, space weather events⁴, RF interference at the antenna, and uncalibrated delays in antenna cables or receiver hardware. A particular concern is interference caused by GPS jammers, which are readily available at low cost and can block GPS signal reception over a radius of hundreds of metres. A GPS-disciplined oscillator feeding a network time distribution must be monitored or compared continuously with another, different time source (such as a free-running clock or a different model of GPS receiver) to verify that it remains locked correctly to the satellite signals.

3. **NTP servers.** The Network Time Protocol (NTP) is widely used to disseminate time over the internet and large numbers of servers can be found online. Servers in unknown locations should be avoided as many are based on GPS-disciplined oscillators and their performance may be affected by local factors such as interference and multipath effects. Most institutes maintaining a UTC(*k*) time scale operate NTP servers that are monitored and synchronised to their time scales. These servers can in principle deliver traceability to UTC, particularly if an NTP authentication method is employed. However, the NTP protocol can only provide synchronisation over wide-area networks with an accuracy of a few tens of milliseconds and will not meet the MiFID II requirements.
4. **Standard-frequency and time signals.** A number of countries operate radio signals that provide access to time based on a UTC(*k*) time scale. These services, which include MSF (60 kHz) in the UK and DCF77 (77.5 kHz) in Germany, are widely used to synchronise radio-controlled clocks. In general, though, the accuracy of the signals varies from tens of milliseconds down to tens of microseconds, depending on the form of the modulation and on variations in the signal propagation, and they are rarely used for network synchronisation.
5. **Time delivery over fibre.** NPL*Time*[®] and similar services disseminate UTC-traceable time over managed fibre links using methods such as the Precision Time Protocol (PTP version 2, defined in the standard IEEE 1588-2008). PTP is a dissemination method originally developed for local area networks, and is capable of achieving accuracy better than 100 nanoseconds over stable and symmetric links. PTP can deliver synchronisation over longer distances using telecoms fibre networks, employing dedicated channels or PTP-compatible switches to maintain accuracy and offering MiFID II compliance at the ingress point of the customer distribution system. As with NTP, the latency to each end point is continuously measured by the protocol (assuming out and back symmetry), enabling the end point time offset to be corrected. The major risk to this type of service is of the fibre link being severed, due to roadworks for example, but local holdover mechanisms can be put in place to maintain service provision whilst the repair is being effected.

In summary, it is not enough to have an accurate time source feeding a distribution network: it is critical to understand, manage and document how time is delivered, distributed and consumed. At each step in the process, knowledge of the traceability of the time signal to UTC is essential to ensure regulatory compliance at the time-stamp. Achieving MiFID II compliance at the time-stamp requires a complete understanding of the traceability chain and the uncertainty of every link, all the way from UTC to the time-stamp, together with an archive of the evidence needed to demonstrate this knowledge.

⁴ Extreme space weather: impacts on engineered systems and infrastructure, The Royal Academy of Engineering, Feb 2013:

<http://www.raeng.org.uk/publications/reports/space-weather-full-report>

6. Obtaining traceability using GPS

GPS receivers have become widely used as reference time standards, providing synchronisation of devices across local area networks. Most of these devices are GPS-disciplined oscillators (GPSDOs). Their 'self-adjusting' behaviour and ease of use make them an attractive choice for many applications. However, it is very difficult to demonstrate direct traceability to UTC using a GPSDO, and they should be installed and used with care.

A GPSDO contains an internal oscillator, usually a quartz crystal oscillator or a rubidium atomic frequency standard, that generates signals with good short-term frequency stability. Long-term accuracy is obtained by steering the oscillator to the time scale broadcast by the GPS satellites, known as GPS time, which is derived from the realisation of UTC at the United States Naval Observatory, UTC(USNO). GPS time does not implement leap seconds so has an integer-second offset from UTC(USNO), but the current offset and forthcoming changes are broadcast within the GPS navigation message. A GPSDO applies the offset to its time output so that it provides a representation of UTC rather than GPS time. This steering procedure enables a GPSDO to deliver a high level of performance, and to maintain that performance indefinitely.

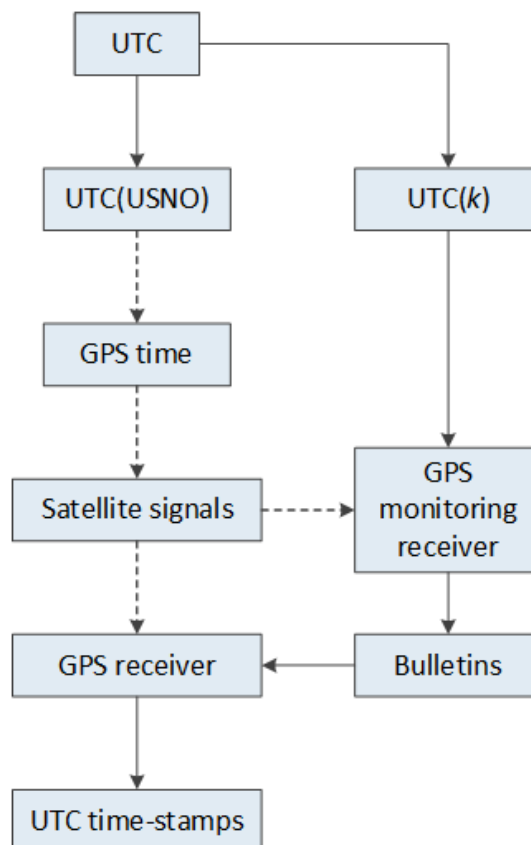


Figure 2: Chain of comparisons from UTC to the time-stamps generated by a GPS disciplined oscillator, and (on the right) the use of bulletins of GPS monitoring results from a UTC(k) institute to demonstrate traceability to UTC. The dashed lines indicate links where it can be difficult or impossible to demonstrate traceability. Note that although the pulses of GPS time are aligned closely with those of UTC, there is an integer-second difference between them which changes when a leap second is inserted into UTC (this difference is broadcast in the satellite signals).

The distribution chain from UTC to the time output of a GPSDO is shown on the left side of figure 2. The links represented by dashed lines cannot easily be evaluated and assigned an uncertainty by an external user, making direct traceability difficult to establish. For example, time obtained from satellite signals can be affected by a number of variables, including delays in the antenna and receiver, environmental effects such as propagation delays through the ionosphere and troposphere, and multipath signal reflections.

A solution often adopted by calibration laboratories is to subscribe to one of the GPS monitoring bulletins published daily, weekly or monthly by some national measurement institutes⁵. These bulletins give a measured value for the time difference between that institute’s UTC(*k*) time scale and GPS time, providing traceability between the satellite signals observed in that region and UTC. An example of the NPL monthly GPS bulletin is included as Appendix 2.

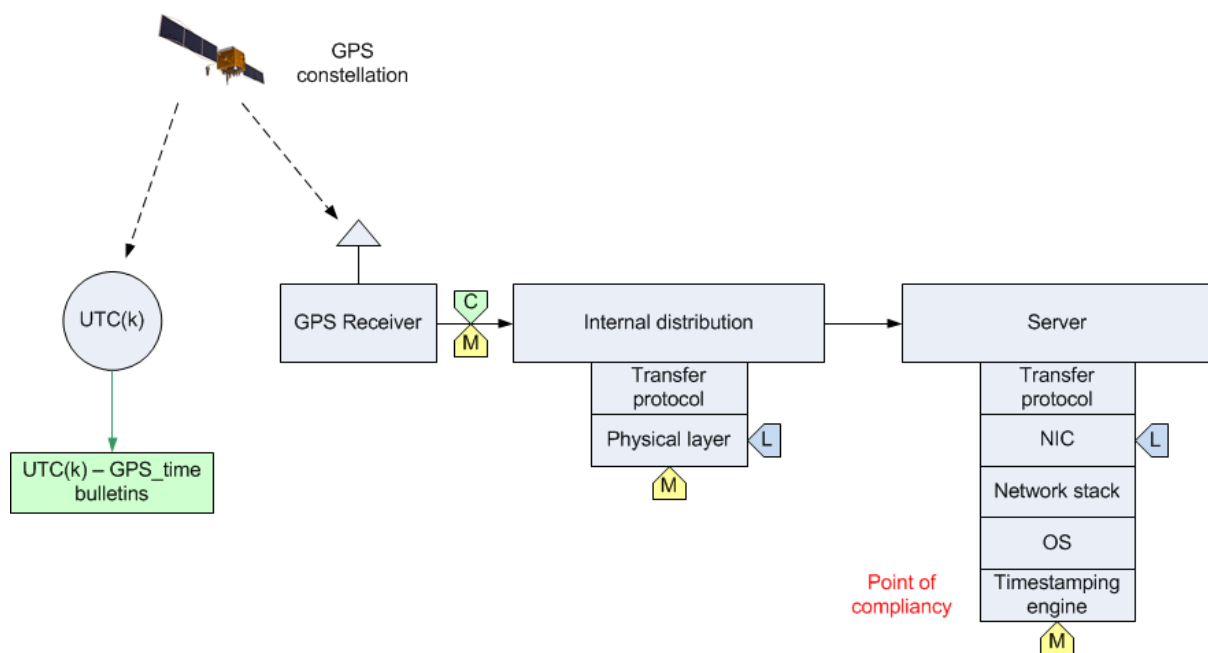


Figure 3: Diagram of a typical GPS-based distribution, indicating the traceability chain. ‘C’ indicates the point where a calibration may be carried out to determine the offset introduced by the GPS receiver, ‘L’ indicates points where latency (distribution offsets) can be determined, and ‘M’ indicates points where continuous monitoring of the time distribution can be carried out.

Even if monitoring bulletins or other methods are used to establish traceability through the GPS satellite signals, there are two other aspects that must be considered before UTC traceability can be claimed for the time output of a GPSDO. Figure 3 illustrates how these apply to a typical GPS-based distribution.

The first requirement is for calibration of the GPSDO. Although the time derived from the satellite signals should be aligned closely with UTC in the long-term, the output signal from the GPSDO may have a constant offset due to uncorrected internal delays, including the delay of the antenna cable.

⁵ Use of GPS disciplined oscillators for frequency or time traceability, EURAMET Technical Guide No.3, March 2016: <https://www.euramet.org/publications-media-centre/technical-guides/>

The total delay of the equipment can be determined in a calibration, and the appropriate correction applied to remove the offset. In some situations, for example where the claimed uncertainty of the UTC output is 1 microsecond or more, it may be possible to use the GPS receiver manufacturer's own evaluation of its time offset to satisfy the traceability requirement. However, in most cases a calibration will be needed.

Calibrations should be carried out by an organisation with the required technical competence. Many countries have an official agency responsible for assessing the capabilities and integrity of test and calibration laboratories. In the UK for example, the national accreditation authority is the United Kingdom Accreditation Service, UKAS.

The second requirement is to have some method of monitoring a GPSDO acting as a reference time source to verify that it remains correctly synchronised to the satellite signals. The most effective solution is to compare the output signal from the GPSDO continuously with the similar output from another device. A sudden change in the time difference between the two devices will indicate that there is a problem with one of them, and neither should be used as a reference standard until the fault has been identified and fixed. The second device can be either a free-running clock or a different model of GPS receiver, preferably from a different manufacturer (since two receivers of the same type may misbehave in the same way). If a second GPS receiver is used, it should be a completely separate installation with its own antenna, mounted well away from the antenna of the primary GPSDO to minimise the risk of both receivers being affected in the same way by factors such as multipath reflections or local interference.

If a second standard is not available, an alternative approach is to monitor the operating parameters of the GPSDO and to check for changes or discrepancies. The parameters observed might include the oscillator locking voltage, the number of satellites tracked, or the satellite signal strengths. Whatever method is adopted, the monitoring results should be recorded and archived so that the status of the GPS receiver at any point in time can be checked at a later date.

7. Improving resilience

A range of methods are available to improve the resilience and reliability of a time delivery solution. Two of the most common are discussed in this section: holdover clocks, and multiple sources of time.

Holdover clocks

A time delivery service is often combined with some form of clock or oscillator to maintain service provision for a period if synchronisation to the delivery source is lost. The clock is commonly referred to as a holdover clock, and the period after the loss of synchronisation is called holdover. Any of the clock types listed in section 5 can provide holdover, but the choice of clock will determine the length of time that the service can be maintained under holdover. The second factor to consider is the synchronisation process that keeps the holdover clock aligned in time, and often also frequency, with the time source for as long as that source is available.

The best holdover performance is obtained using a Cs clock as it does not display long-term frequency drift, unlike quartz or Rb oscillators that experience frequency changes over time due to ageing processes. The stability of a Cs clock, measured using statistical tools such as Allan deviation, is typically below 10 ns when averaged over 1 day. As a result, a caesium clock used for holdover need only be re-synchronised to the time source – a process known as steering or disciplining - daily or even less frequently, and can maintain service delivery to within a few tens of ns over many days. In contrast, a rubidium oscillator might only maintain delivery within 1 μ s over several hours.

To maintain UTC traceability when synchronisation is lost, the time offset between the source and the holdover clock – and its uncertainty – must be known at that instant. In addition, the frequency offset of the clock from the synchronisation source, and its uncertainty, has to be known so that the divergence of the clock from UTC can be predicted. If the clock type has frequency drift, this should also be taken into account, so that at any later time the clock offset from UTC and its uncertainty (or prediction error) can be estimated.

Holdover clocks are most commonly encountered in GPS receivers intended for use as time sources, a type of device known as a GPS-disciplined oscillator (GPSDO). The oscillator in this case is usually quartz or rubidium, and the disciplining process is a phase-locked loop with an integration time of a few minutes (quartz) to a few hours (Rb).

Multiple time sources

We saw in section 5 that a number of different sources of time and time dissemination methods are available, which can provide a traceable chain of measurements between a time stamp and UTC. The resilience of a time service against interruptions or other disruptions can be increased substantially by using two independent dissemination methods. For example, a GPS-based solution can be combined with fibre delivery of time, potentially enabling the two sources to be compared against each other to maintain time synchronisation to UTC in the event of loss of one or other of them.

The key to the effective use of two (or more) time sources is a method to compare them, and to switch from one to the other if necessary. This requires a specialised device able to monitor the difference between the two sources, for example by measuring the time offset between 1 pulse-per-second (1 PPS) signals derived from each. If one of the signals is lost, the comparison device will automatically switch to the other. More commonly, one of the time sources will develop a fault and begin to deviate from the correct time. An example of this would be a GPS receiver losing its lock to the satellite signals and free-running on its internal oscillator. An error of this type is harder to deal with if the comparison device has no means of detecting which source is correct.

To get around this problem, the comparison device can incorporate a holdover clock, loosely disciplined to one of the time sources as discussed above. In the event of the time sources starting to diverge rapidly, a 3-way comparison with the holdover clock will usually reveal which of the sources has developed the fault. The other, correct, source can then be selected as the reference.

8. Conclusions

The MiFID II regulations in force in Europe require time-stamping of automated transactions with an accuracy of 100 microseconds (μs). To have any validity, the time source and dissemination method used must be traceable to UTC, the international reference time scale. In practice, a time measurement or timestamp should relate back through a traceability chain of comparisons with known uncertainties to one of the timing institutes that maintain a physical representation of UTC, or UTC(k) time scale. The traceability chain and the uncertainty evaluation must be documented thoroughly.

Traceability is particularly difficult to demonstrate when using a GPS receiver as the time source. GPS monitoring bulletins issued by one of the regional UTC(k) timing centres can facilitate this by providing a record of the time offset between that UTC(k) time scale and GPS time.

The resilience of a time delivery service can be improved by incorporating holdover clocks to maintain the time at the endpoints if the distribution path is interrupted, or by using two independent distribution methods. In both cases, care is needed to understand the offsets and instabilities present, and their evolution over time, in order to maintain traceability through a disruption to the delivery service.

Whatever timing equipment is used, there are two other requirements for demonstrating traceability to UTC. The first is that the timing equipment should be calibrated to determine and if necessary correct the error in its time output due to its internal delays. The other is that the equipment should be monitored continuously to detect any fault that might affect its time output. Monitoring and calibration records must be kept to provide evidence of correct operation.

Appendix 1: Terminology of time measurement

Brief explanations of some of the more common terms used in time measurement are given here.

Accuracy of a measurement represents its closeness of agreement with the true value, and can be quantified by assigning an uncertainty to it. For time measurements, the 'true value' is provided by the international reference time scale UTC.

Precision is the level of agreement between repeated measurements of a quantity, and can often be calculated using a statistical tool such as a standard deviation or a variance.

Resolution is the smallest difference that can be observed in a measurement, and is dependent on the properties of the measuring device.

Stability (or equivalently **instability**) is a measure of the fluctuations in a measured quantity over a given time interval. The dominant noise processes observed in a clock or time distribution system depend on the measurement time, and the stability is quantified using statistics such as the Allan deviation that are a function of the sampling period.

Traceability of the time given by a clock requires an unbroken chain of measurements with known uncertainties that link it back to the agreed reference standard, which for accurate timekeeping worldwide is UTC. A timestamp can only be considered accurate if its traceability back to a source of UTC has been established.

Uncertainty of a measurement, or more correctly the standard uncertainty, is a numerical assessment of the level of doubt in the result. The standard uncertainty can often be separated into two components: a statistical term due to the random fluctuations in the measurement result, and a systematic term due to a bias or offset in the measurement result. If these terms are independent, the combined standard uncertainty can be calculated as their root-sum-square combination.

Coordinated Universal Time (UTC) is the international reference time scale. It is computed monthly by the International Bureau of Weights and Measures (BIPM) using more than 500 atomic clocks in around 75 timing centres world-wide. These timing centres maintain physical time scales that provide real-time realisations of UTC, referred to as UTC(*k*), where *k* designates the abbreviation allocated to an institute. The scale interval of UTC is the SI second.

The SI second (s) is the basic unit of time in the International System of Units (SI). It is defined as the duration of a specified number of cycles of the microwave radiation corresponding to a particular transition within caesium 133 atoms. The second can be subdivided into milliseconds (1 s = 1000 ms), microseconds (1 s = 1 000 000 μs) and nanoseconds (1 s = 1 000 000 000 ns).

Appendix 2: Example of NPL monthly GPS monitoring bulletin

Several National Measurement Institutes in Europe monitor continuously the GPS time broadcast in the GPS satellite signals against their respective UTC(k) time scales, and publish the a summary of the measurements in weekly or monthly bulletins. An NPL GPS monthly bulletin is included below as an example.

NATIONAL PHYSICAL LABORATORY
Time and Frequency Services
Teddington, Middlesex, TW11 0LW, UK

Web site: www.npl.co.uk

N P L G P S B u l l e t i n

No.2018-07 July 2018

MJD	Date	[UTC(NPL) - GPS_time] mod 1 s (ns)
58300	2018-07-01	2.5
58301	2018-07-02	2.5
58302	2018-07-03	2.6
58303	2018-07-04	2.5
58304	2018-07-05	3.0
58305	2018-07-06	1.9
58306	2018-07-07	2.2
58307	2018-07-08	1.1
58308	2018-07-09	1.0
58309	2018-07-10	1.1
58310	2018-07-11	1.6
58311	2018-07-12	2.9
58312	2018-07-13	3.6
58313	2018-07-14	4.0
58314	2018-07-15	3.3
58315	2018-07-16	3.4
58316	2018-07-17	2.9
58317	2018-07-18	3.2
58318	2018-07-19	3.7
58319	2018-07-20	3.9
58320	2018-07-21	4.8

58321	2018-07-22	5.8
58322	2018-07-23	6.6
58323	2018-07-24	6.1
58324	2018-07-25	6.5
58325	2018-07-26	6.8
58326	2018-07-27	6.0
58327	2018-07-28	4.6
58328	2018-07-29	4.6
58329	2018-07-30	4.4
58330	2018-07-31	4.7

NOTES:

1. This Bulletin contains daily average measurements of [UTC(NPL) - GPS_time].
2. ## indicates that NPL data are not available for that day.
3. The total 95% confidence interval on each daily value is +/-22 ns.
4. Due to accumulated leap seconds, [UTC(NPL) - GPS_time] div 1 s = -18 s.
5. $UTC(NPL)-GPS_time = [UTC(NPL)-GPS_time] \div 1\ s + [UTC(NPL)-GPS_time] \bmod 1\ s$.
6. This Bulletin has been compiled using GPSMONITOR201.EXE version 2.01.
7. The measurements in this Bulletin were taken by Mesit GTR51 GNSS timing receiver s/no. 1401406.
8. The measurements in this Bulletin are single-frequency C/A code observations with the ionospheric delay corrected using a P3 combination of the P1 and P2 code measurements.
9. No anomalous GPS measurements were detected during the period covered.