WHITE PAPER

Accurate timing for the IoT

Four questions to identify the timing technology that meets the requirements of your wireless application.

Abstract
In this white paper, we walk you through the process of selecting the right technology for your IoT application. We outline existing timing technologies for the IoT and present a hybrid timing approach that opens the door to new application indoors, outdoors, and even underground.

White paper authors:
Drazen Drinic, Principal Product Strategy, Product Center Cellular, u-blox AG
Chris Marshall, Senior Principal Engineer, u-blox AG

Contents

Timing is essential in IoT

Time synchronization is crucial in all network applications, from the Internet to industrial, financial, and scientific applications. This has led to the development of a variety of approaches to achieve highly accurate synchronization of devices across networks. These include timing protocols for connected devices such as NTP (Network Time Protocol) and 1588v2 PTP (Precision Time Protocol), as well as using GNSS receivers as a highly reliable source of time.
Today, the spread of the Internet of Things (IoT) is changing the timing game. An increasing number of distributed applications are wireless, so we face a big opportunity but also a challenge, how do we distribute highly synchronized timing wirelessly? It can be done with today's technology – if, from early on in the development, you ask the right questions.
In this white paper, we walk you through the process to help you select the right technology for your IoT application. We outline existing timing technologies for the IoT and present a hybrid timing approach that allows new wireless IoT applications indoors, outdoors, and even underground.

Timing continues to drive innovation

Scientific and technological progress have long evolved in lockstep with our ability to keep time. Today, nanosecond-accurate precision timing using atomic clocks underpins everything from the Internet and satellite-based positioning to high frequency trading and mobile telecommunication networks. The next frontier involves extending the reach of accurate timing to the next generation of distributed systems that are open, highly interconnected, and deeply embedded in the physical world.

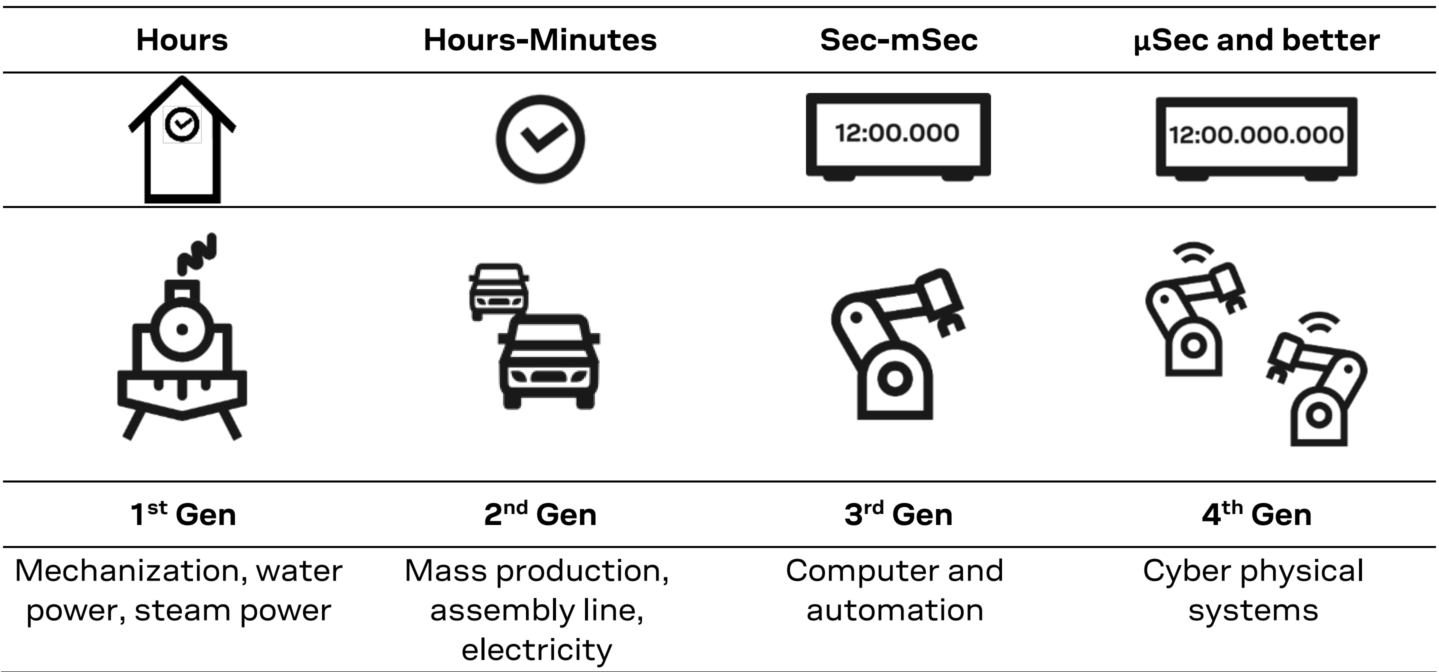

Figure 1: The relevance of timing in industrial revolutions.

At the heart of the ongoing digitalization of every aspect of our lives lies the Internet of Things (IoT), the sum total of devices – things – that are connected to the Internet and to each other. By leveraging wireless connectivity and cloud-based data storage and analytics, these connected devices are enabling countless new use cases, from smart cities to security, from the connected industry to transportation. In each of these, timing plays a central enabling role.
Ongoing developments in timing technology promise to further grow the IoT ecosystem and broaden its scope, paving the way for a wealth of new business opportunities, and potentially giving new players a foothold in the industry.

The lineup of timing technologies for IoT applications continues to grow. Successful applications will stand out, at least in part, by adopting the timing technology that is best suited to meet their specific needs. With this white paper, we aim to assist product developers in the early stages of development by outlining key questions that are critical to their decision-making process.

Question 1: Which time do you need?

Customers often tell us that they need a highly precise timing solution for their IoT application. When asked to explain exactly what it is they need, the answer is typically: "Highly accurate universal time." Timing modules based on a GNSS-receiver will deliver highly accurate universal time provided a sufficiently unobstructed view of the sky. Unfortunately, however, this isn't the case for all scenarios, in particular those involving indoor and underground applications.
The obvious follow-up question is: is highly accurate universal time really necessary? Digging deeper, we often learn that they are actually after a highly reliable and resilient timing reference that works out of the box in the broadest range of settings. That's a different challenge, but fortunately, it's also one that we can solve.

To understand the difference between the two, it helps to take a look at the four classes of time that serve the needs of most IoT applications. As we will explain later on, the type of timing required directly influences the technology options that are available to a specific application.

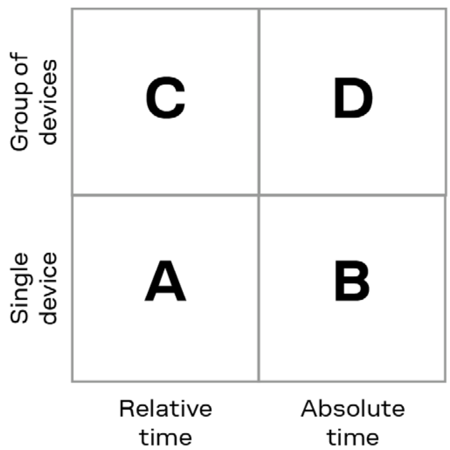

Figure 2: Identifying which of the four classes of time you need can help narrow down the range of available timing solutions.

Quadrant A: Timers

Many timing applications require no more than an accurate measure of the time elapsed between events on a single device, with no need to synchronize multiple devices. Examples are egg timers, stopwatches, and heart rate monitors. These applications require an accurate “tick” – in other words, an accurate measure of time intervals between isolated events where the alignment with absolute time is not critical.

Quadrant B: Astronomical reference clocks

Other applications, such as the global reference clock in Greenwich, London, require absolute time on a single device. Others include turret clocks, alarm clocks, and watches. Absolute time may be Coordinated Universal Time (UTC) or the local time in a given time zone.

Quadrant C: Signal synchronization

Applications that need a shared notion of time intervals fall into quadrant C. These include groups of distributed sensors that record and share sets of signals. Examples include Phasor Measurement Units in electricity smart grids, and wireless media equipment (such as microphones) in broadcast and auxiliary services.
These distributed applications need to be able to synchronize their “tick” and establish an agreed upon initial time, t0. In these applications, the requirement to align the time to absolute time is typically less stringent.

Quadrant D: Financial trading networks

The most demanding distributed applications require all devices to be highly synchronized to an absolute measure of time. A well-known example is that of financial trading networks, which need to be able to determine the order in which stocks are traded with an accuracy in the order of nanoseconds.
These applications need to maintain highly precise time intervals and synchronize their measurement scale with absolute time, for example, using GNSS-based timing modules that are synchronized with the atomic clocks on board GNSS satellites.

Question 2:
How accurate do you need it?

Now that we've sorted out the type of timing your application requires, it is time to turn to its demands in terms of accuracy. But first, keep in mind the difference between accuracy and precision: a high quality watch set to the wrong time can be extremely precise, but inaccurate.

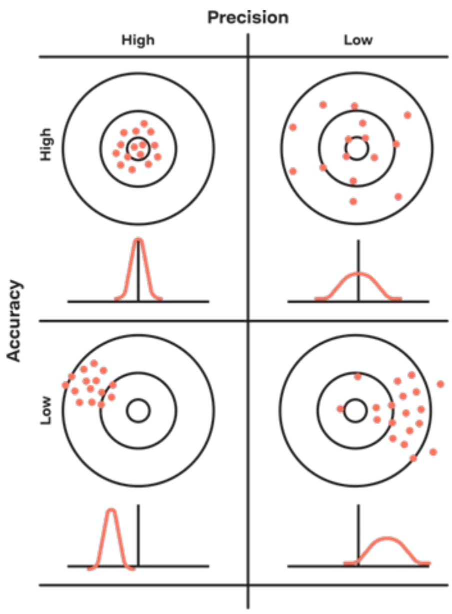

Figure 3: A high quality watch set to the wrong time may be precise, but inaccurate.

It might seem to be a smart strategy to err on the side of caution and equip any application with highly accurate absolute timing. After all, there's no harm in having your egg timer go off after 3 minutes +/- 5 nanoseconds and knowing the time to the same accuracy. That said, setting unnecessarily stringent requirements on timing accuracy can lead you to overlook timing technologies that are cost-effective and optimized for your application.

The tradeoff is inevitable, but a careful assessment of accuracy requirements will give you the broadest possible array of technologies to choose from, ultimately improving the size, the cost, and the performance of your device.
Some applications set different requirements on different parts of the process. Consider the example of an average speed trap that uses two cameras located one kilometer apart to track the passage of a vehicle.

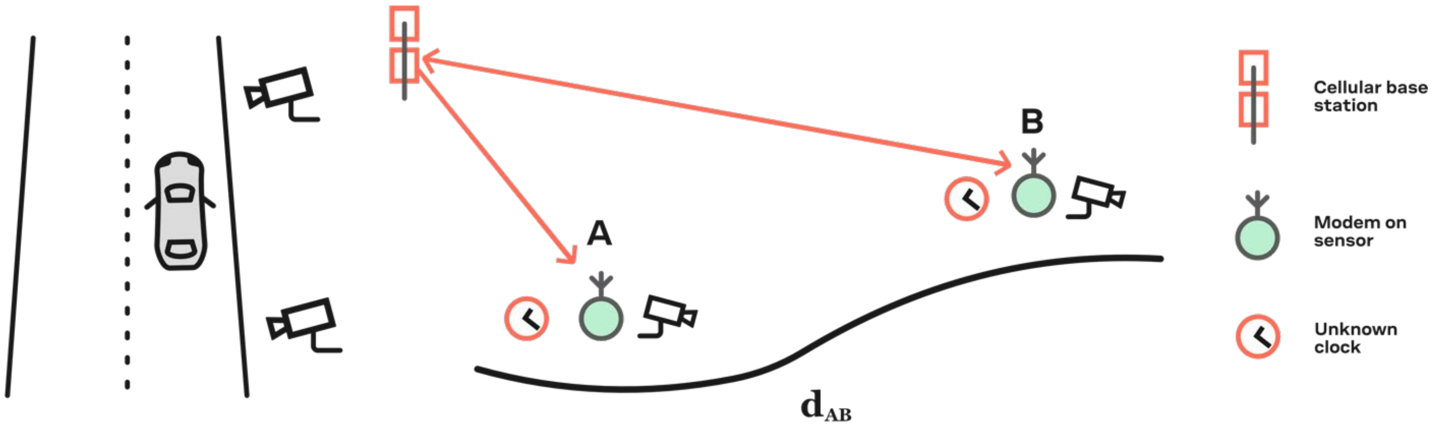

Figure 4: Camera-based speed trap setup along a road requiring different levels of timing accuracy for different parts of the process.

In this example, the relative timing accuracy between the cameras should be kept low, as a 0.1 second uncertainty gives an uncertainty in the speed measurement of 0.2 km/h. Meanwhile, accuracy requirements on the absolute time used to report the event are much more relaxed, with one second likely to be sufficient.

Smart grids, which involve devices distributed over wide areas that measure sampled values and record events, also have a range of time performance requirements, defined in IEC 61850.

Time Performance Class	Accuracy	Purpose
T1	±1 ms	Time tagging of events
T2	± 100 µs	Time tagging of zero crossings and of data for the distributed synchrocheck. Time tags to support point on wave switching.
T3	±25 µs	Instrument transformer synchronization (sampled values)
T4	±4 µs	
T5	±1 µs	

Table 1: Time performance requirements laid out in IEC 61850

Question 3: What are your constraints?

A clear understanding of the constraints on your application is essential to select the optimal timing technology, in particular in terms of coverage, mobility, power consumption, and cost.

Signal coverage

Where will your device operate? Wireless timing devices typically leverage some form of electromagnetic signal to track time. Because some signals penetrate deeper into buildings or underground than others, the quality, reliability, and continuity of the signals that the device receives can narrow down the list of technology options.

GNSS signals, for example, penetrate very poorly into buildings, while low power wide area cellular signals, such as LTE Cat 1, LTE-M, and NB-IoT have much deeper signal penetration.

Robustness

How sensitive is your application to interruptions? Cellular networks can fail, GNSS signals can be jammed, and internal clocks can drift. Critical applications need to be designed to operate reliably even under the most challenging circumstances.

One way for critical applications to increase their robustness is by drawing on multiple sources of timing information.

Mobility

Will your device stay put or will it move around? This question is particularly relevant when cellular 4G LTE signals are used to determine time. Moving mobile devices towards or away from cell towers leads to a “Doppler-like” effect that affects the length of each “tick,” impacting the accuracy of their alignment with the reference time.

Spatial distribution

How distributed will your application be? IoT applications often involve devices spread over an area, the size of which influences the timing precision that can be achieved by each candidate technology. While wireless transmission services such as those used in radio-controlled clocks can precisely synchronized devices within a small area, the synchronization becomes less precise as the devices move apart.

On a larger scale, devices that rely on cellular 4G LTE signals to track time might be allocated different or changing serving cells by the network, and so might not all be connected to and receiving the same base station. Such use cases require careful attention, as not all base stations are synchronized to the same reference time.

For wide area applications, GNSS is most appropriate for providing accurate synchronization, using the signals from multiple satellites to find the position and the time at the receiver.

Energy consumption

What level of battery autonomy are you aiming for? Keep in mind that update frequency and data transfer come at the cost of battery longevity, which is also strongly influenced by the technology choices you make in your designs.

Cost

What’s your budget? Traditionally, distributing time from the timing “grand master” down to the end device has been a complex and costly task. New solutions are offering simpler and more affordable ways to distribute accurate timing across systems.

Even in the simplest solutions, it makes sense to step back to consider the benefits of different technologies beyond accuracy. Returning to the example of the simple egg timer with no further connectivity requirements, a standard 32 kHz clock might be an excellent and cost-effective solution. But if you are developing a networked washing machine that uses a cellular communication module to connect to the cloud, the marginal cost of using an LTE-based timing solution might be so low that it makes sense, even more so if it turns out to be the simplest to implement and maintain.

Question 4: What are your options?

So, with a clear understanding of the type of time we need, the accuracy we are aiming for, as well as the external, design, and use constraints our application is subject to, it's time to look at the available technology options.

IoT applications can draw on a range of sources for timing, each with their own set of constraints. Below are the most common and, arguably, the best-suited timing technologies for wireless IoT applications.

Local timer

Local clocks are straightforward to implement and always available provided there is sufficient power supply. That said, they are limited both in precision and accuracy, which depends on the frequency of operation, the quality factor of the oscillator, the power consumption, the power source, and, importantly, environmental variability, particularly temperature. They further require initialization to set the time.

Low frequency clocks, such as the common 32 kHz clock, are excellent for keeping an approximate view of time, for example scheduling the duration of sleep intervals. High frequency, high quality, temperature-compensated oscillators may be used to provide a high precision time base to the equipment, with the minimum of timing jitter.

Wireless broadcast transmitters

Wireless transmitters have long been used to provide timing signals for devices over a wide area, notably the Rugby clock, latterly relocated and renamed MSF in the UK, WWVB in the US, and DCF77 in Germany.

In addition to a frequency reference that is not affected by a propagation delay between transmitter and receiver, they also broadcast various modulation signals to indicate seconds and absolute timing.

When used to synchronize multiple devices, the difference in propagation time from the transmitter to the devices, given by $\Delta_t = (d_A - d_B) / c$, where c is the speed of light, may need to be taken into account. For applications where the IoT devices are distributed over a small area (e.g. 100 meters apart) or which do not require high precision (e.g. <1 milliseconds), this error effect can safely be neglected.

LTE signals

Cellular wireless signals such as those from LTE base stations are widely available, even indoors, and, with the deployment of low power wide area networks, may increasingly also be received in previously unserved locations such as underground. It is clearly convenient for a cellular modem to use such wireless signals for timing applications, as well as for communications.

The cellular signals provide excellent short term and usually good long term stability, and can provide relative time-keeping for static, distributed IoT sensor networks connected to the same base station. Depending on the cellular system design, the cellular signal may be linked to absolute time more or less accurately. For example CDMA and TDMA systems tend to be accurately synchronized to absolute time using GNSS, while GSM networks are less tightly controlled. The absolute timing of LTE base station depends on the network configuration.

In common with broadcast wireless signals, the cellular base station may send messages indicating the absolute time of transmission, which when received by each device gives an absolute time that depends on the propagation distance of the signal from the base station.

Network timing

Once an internet connection is established, timing services and protocols such as NTP can be used to provide timing services to the modem device.

Timing is also provided as a part of some location services that provide assistance to GNSS receivers. Absolute timing information is provided over the cellular network as part of the CellLocate feature of u-blox's cellular modems.

GNSS signals

GNSS signals offer IoT applications atomic clock-level timing accuracies. Designed to determine their x, y, and z coordinates, as well as the absolute time based on the signals from GNSS satellites carrying atomic clocks, timing units that leverage GNSS signals have become the solution of choice to synchronize mobile communication network infrastructure, and are perfect for wide area applications. However, because of the low signal levels involved they are currently limited to applications that mainly operate outdoors.

Overview of timing options

We can summarize the applicability of each timing source for the different types of applications introduced here as follows:

Local timer	**LTE signal**	**Network timing**	**GNSS**
			
Elapsed time locally	**Relative time everywhere**	**Absolute time everywhere**	**Absolute and relative time outdoors**
Pros			
• Always available provided sufficient power	• Widely available, even underground • Excellent short term stability	• Widely available for connected devices	• Excellent accuracy • Excellent long term stability
Cons			
• Applicability limited by drift, temperature and aging	• Absolute timing may be undefined • Observed time depends on signal propagation distance	• Limited by variable latency, depending on the network, which affects timing accuracy	• Limited to the cases where signal reception is not obstructed (not indoors and not underground)

Figure 5: A comparison of time sources for cellular modem IoT applications

Hybrid Timing

On its own, each timing solution is subject to limitations that restrict its applicability. Local clocks drift. LTE-based solutions can fall out of sync when the devices move from cell to cell. GNSS-based solutions are limited by their need for an unobstructed sky view, as well as by signal jamming.

Combining the timing accuracy offered by GNSS signals with the reach, penetration, and reliability of LTE signals would be desirable to meet the performance specifications required by many IoT and distributed wireless applications. Such a hybrid approach offers several benefits:

- Increased resiliency: providing an accurate timing output beyond the reach of GNSS signals.
- Increased reliability: improving accuracy when GNSS reception is limited.
- Increased availability: providing absolute time when GNSS signals are obstructed.

Depending on the application, the timing sources are used effectively by the hybrid timing system to provide the measurement of the timing intervals within the device, estimate the timing between events at multiple IoT devices, and to report absolute time.

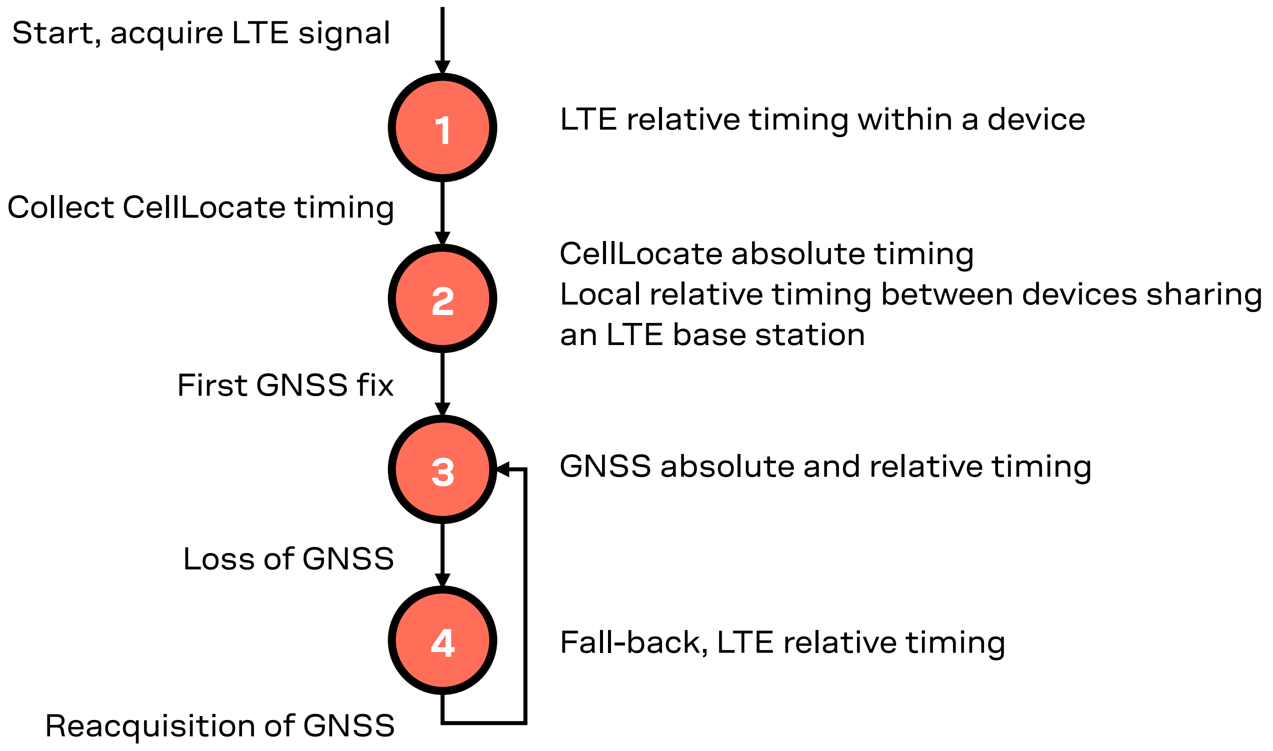

Figure 6: Illustrative sequence of events and the timing available for a modem providing accurate and robust hybrid timing to an application.

In such a system the limitations of one source of time may be compensated for by another:

- The relative time measurement provided by LTE can bridge outages in the GNSS.
- The approximate absolute time provided by network timing can resolve ambiguities in the LTE frame timing.
- The local clock can detect jumps and errors in LTE or GNSS timing.
- The position and time provided by GNSS over wide areas resolves the offsets arising from propagation distances (and signal visibility) of LTE signals.

Applications of hybrid timing

Some applications illustrate the use of hybrid timing to meet the relative and absolute timing requirements for IoT devices.

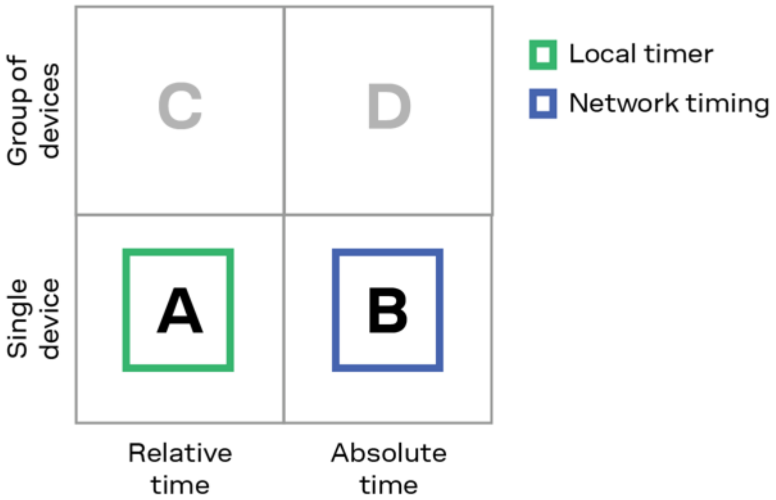

Figure 7: A meter application

The simplest example is for a meter, or even for a radio alarm clock. In the meter application, each device acts on its own, without interaction with its neighbors. Timing accuracy is not that critical, but cost and coverage in difficult environments are the priority. Time is initially set and periodically updated by a network timing service, and timing is maintained by the local clock that provides local autonomy.

If high precision is required, then the device can use clock LTE signals for timing the elapsed intervals of time.

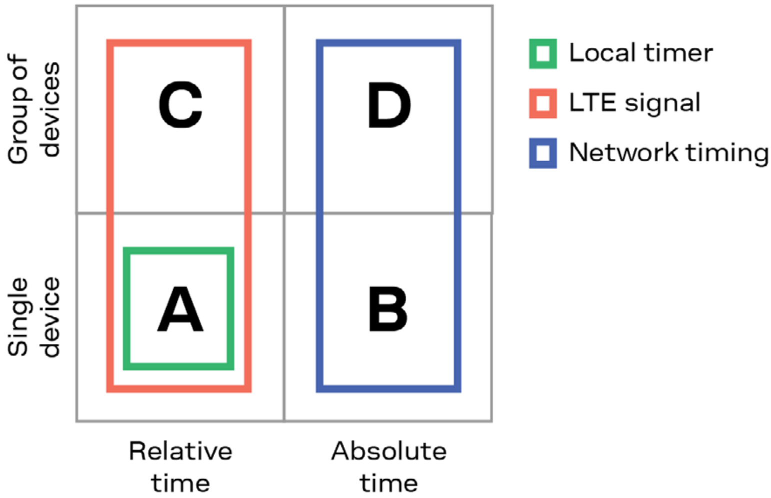

Figure 8: Factory of the future and signal handling applications

On the other hand, more demanding IoT applications require the synchronized processing of signals or information with multiple devices, indoors. The devices are often reasonably close together. Examples are:

- large machine control in factories of the future
- fault finding in smart grids
- audio streaming with wireless microphones and broadcast auxiliary services
- average speed measurement

In these applications a local clock has its tick supervised by the LTE base station, which also provides accurate relative timing between devices. Network timing is used to resolve any ambiguities in the frame timing of the LTE base station, and to time-stamp events if necessary for annotating measurement reports.

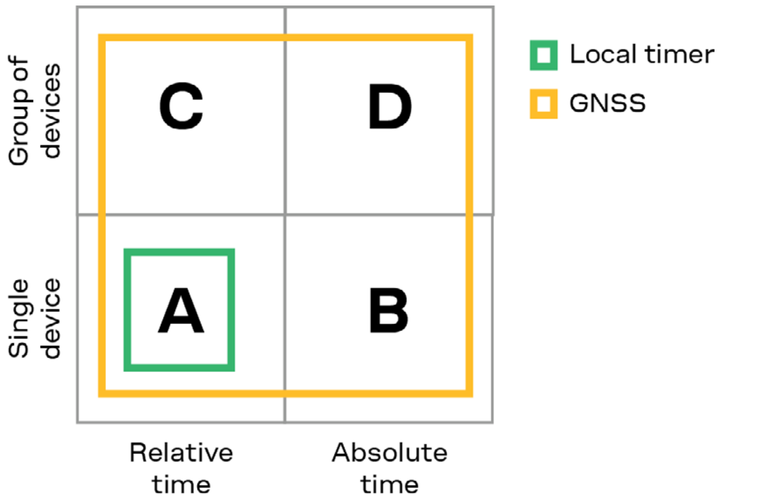

Figure 9: Wide area sensor applications

For wide area applications with multiple sensors outdoors, such as seismology, GNSS provides unrivalled performance.

A local clock may be used for general scheduling of tasks, and GNSS will typically be used to control a high quality local oscillator, which then drives the application system to obtain the best possible jitter performance.

Timing in marginal conditions

Indoor applications are typically beyond the reach of GNSS signals. There are, however, use cases in which devices deployed indoors featuring a GNSS receiver occasionally see a few GNSS satellites.

It is possible to estimate the local clock time from just a single satellite, if the position of the device is known. The accuracy of the time estimate that is then produced by the GNSS receiver depends on the accuracy of the position information with which it is provided, scaled by the speed of light.

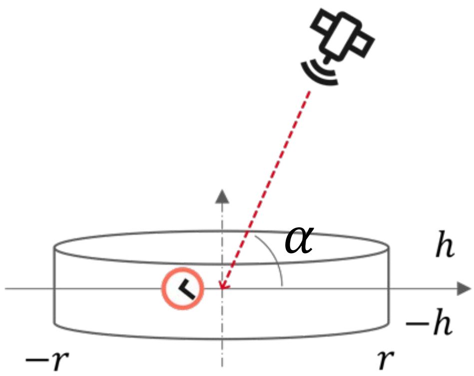

Figure 10: Determining the uncertainty in time estimate for a given accuracy in the position estimate.

With a horizontal uncertainty of radius ±*h* and a height uncertainty of ± *h* then for a satellite at an angle of elevation α the uncertainty σ in the time estimate is given by σ =(*r* cos *a* +*hsin* α) ⁄ *c*, where *c* is the speed of light. Even a rough position estimate, as provided by u-blox's CellLocate service, can be sufficient to achieve high time accuracy when operating in single satellite Time Mode.

For example, if the location of the IoT device is estimated to an accuracy of 150m, then measurement of a single satellite signal can give a timing accuracy of the order of 0.5 µs.

Coping with changes

In marginal environments, such as indoors, the GNSS fix may become lost. In these instances, a hybrid solution allows the modem to maintain timing by referring to the LTE signal as reference, as shown below. In the absence of a GNSS fix, timing may not be as precise, but is sufficient for many applications.

When the GNSS signal is reacquired, timing relative to GNSS and UTC time is re-established. At this point there is potentially a small discontinuity that depends on the configuration and application, as illustrated in the figure.

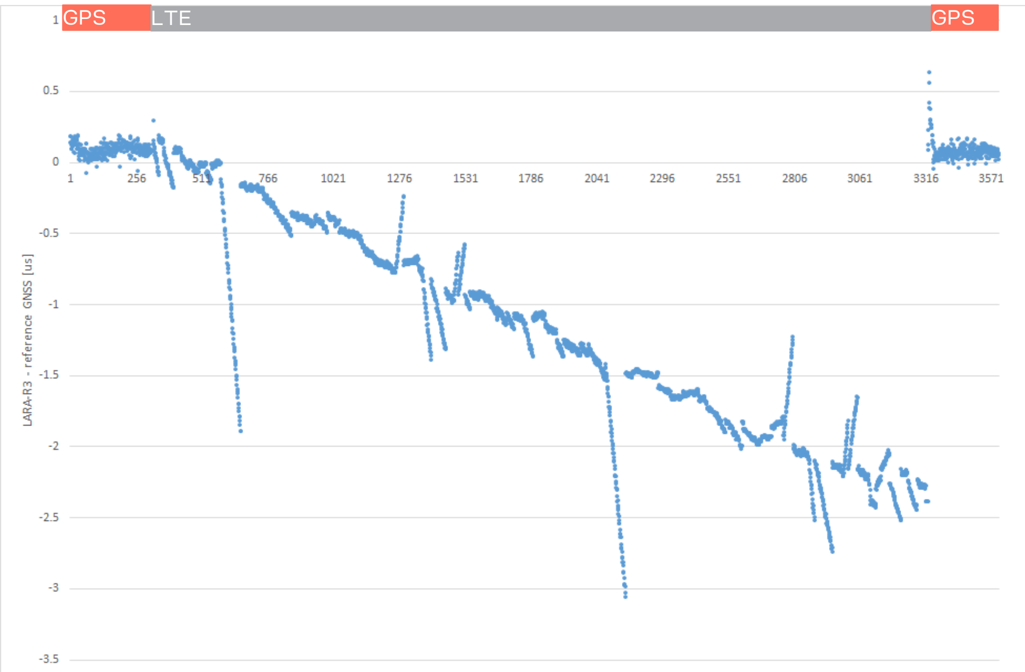

Figure 11: Hybrid timing output of 1pps, with GPS becoming unavailable and falling back to an LTE base station as a reference clock source.

A hybrid timing system should therefore:

- Indicate the current status and source of timing to the application.
- Measure and report any transitions and updates in its knowledge of time, for example as information such as GNSS becomes available.

Sharing absolute timing deep indoors

Using Hybrid timing, GNSS can be used to provide accurate absolute timing to a group of modems. This brings two important benefits:

- It provides timing deep indoors, with a reference modem with a GNSS receiver providing supporting measurements.
- It allows the application to synchronize separate groups of modems deployed over a wide area, with a reference modem with GNSS receiver supporting each group.

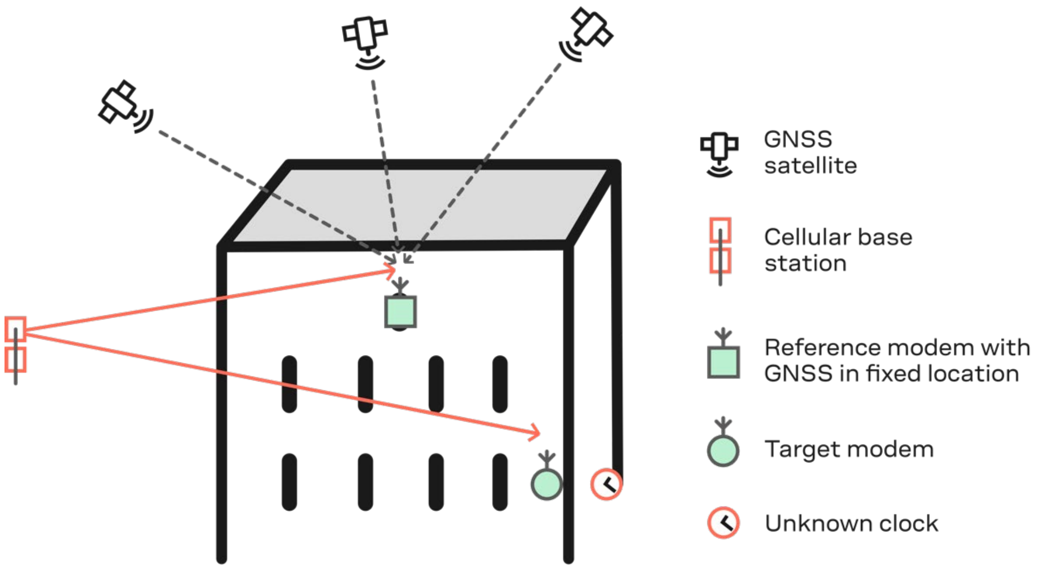

Figure 12: Using a local anchor with GNSS to calibrate the timing of a signal from an LTE base station.

In these cases, assistance information is provided by the reference modem with the GNSS receiver so that neighboring modems can set the absolute timing of the signal received from the LTE base station. The neighboring modems can then operate successfully and provide absolute timing, even if GNSS is not available.

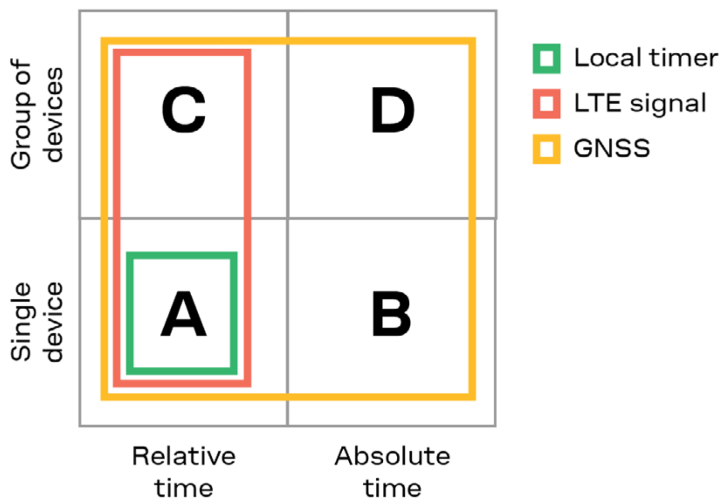

Figure 13: A reference modem with a GNSS receiver measuring the absolute time of an LTE wireless signal.

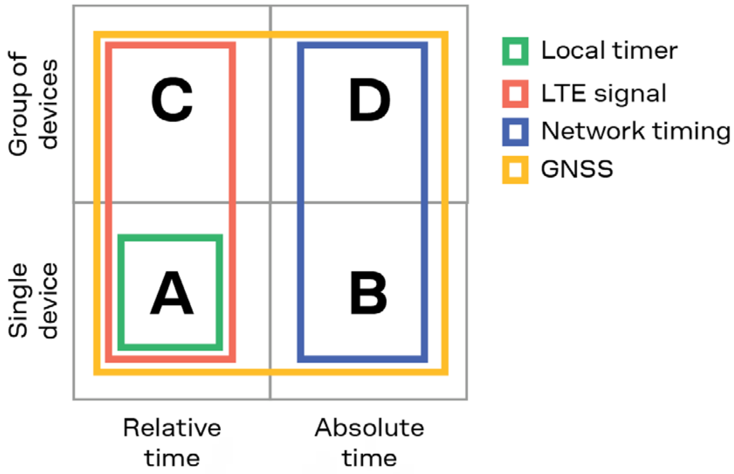

Figure 14: A modem using the timing of the LTE wireless signal, measured by the reference modem.

The accuracy is limited by the service area over which the neighbor modems in the group are distributed, as explained earlier.

In a wide area deployment with a number of groups of modems, the use of a reference modem with GNSS in each group then allows multiple separate groups to be accurately synchronized together, since they are each synchronized to absolute time by means of the GNSS.

Conclusions

From this overview of the use of time in IoT applications, it can be seen that there are two key aspects to time – its “tick”, or ability to measure intervals, and its “setting” compared to some time base. Together these provide the information to measure the timing of IoT events within a device, between groups of devices, and referred to absolute UTC time, according to the needs of the application.

Hybrid timing can use a variety of information sources, from a local clock and wireless signals – both cellular and GNSS, in order to improve performance. The wireless signals can be used to provide time information and to establish the timing “tick” and “setting” for the device time.

By this means we can provide timing to IoT devices and applications with good coverage, good accuracy, and good reliability.

About the authors

Drazen Drinic, Principal Product Strategy,
Product Center Cellular, u-blox

Drazen Drinic is an electrical engineer from the Technical University in Dortmund, Germany. He has over 20 years' experience working in various product management roles in Telecomm infrastructure and Energy and Automation technology.

Chris Marshall, Senior Principal Engineer, u-blox

Dr. Chris Marshall is an electronics engineer from Cambridge University and Imperial College, London, UK. He has enjoyed and led Research and Development projects and activities relating to wireless circuits, systems, and software, from radio paging receivers, through the digital cellular revolution, to software GPS.

About u-blox

u-blox (SIX:UBXN) is a global provider of leading positioning and wireless communication technologies for the automotive, industrial and consumer markets. Its solutions let people, vehicles and machines determine their precise position and communicate wirelessly over cellular and short-range networks.

With a broad portfolio of chips and modules, and a growing ecosystem of product supporting data services, u-blox is uniquely positioned to empower its customers to develop innovative solutions for the Internet of Things, quickly and cost-effectively.

With headquarters in Thalwil, Switzerland, the company is globally present with offices in Europe, Asia and the USA.

UBX-18051297

u-blox AG
Zuercherstrasse 68
8800 Thalwil
Switzerland

www.u-blox.com