Accuracy, Precision, Delays and Drift: How Good is Your Time Reference ?

Dr. Conrado Seibel
mergedK GmbH

April 2014

Abstract

Time is of paramount importance for the power industry: understanding and controlling the interactions between what happens in different parts of the power grid is only possible if a common time-reference is used. The required accuracy ranges from sub-second (multi-rate billing) up to 1 microsecond (power-flow monitoring with synchrophasors) and even better (traveling wave-based fault location).

Using GPS-based clocks only solves part of the problem since misconfiguration of clocks and/or devices, the wrong choice of time-distribution media and the occasional loss of lock to satellites can significantly affect the system as a whole.

In this paper we will describe accuracy and precision of master clocks and the errors caused not only by antenna cable delays but also by the use of serial protocols, optical fibers, and the unintended consequences of using redundant Ethernet switches.

We will also look at the performance of master clocks in the so-called "holdover mode", i.e. what happens after loss of lock to GPS satellites, the error estimate of a clock and the ways of reporting this estimate to different devices in the substation.

1 Time in the Power Industry

Time is of paramount importance for the power industry: understanding and controlling the interactions between what happens in different parts of the grid is only possible if a common time-reference is used.

Table 1 lists some applications along with the typically required accuracy.

Application	**Required Accuracy**
Multi-rate billing	0.5 to 1 s
Event reconstruction, sequence-of-events recording	1 to 4 ms
Lightning strike correlation	1 ms
Time-tagging of zero crossings, point-on-wave switching	100 μs
Synchrophasors, power flow monitoring	1 to 10 μs
Traveling wave-based fault location	better than 1 μs

Table 1: Sample applications and typically required accuracy

1.1 Signals and Protocols

There is a broad spectrum of signals and protocols used for time-synchronization in the power industry due to the fact that equipment in power substations remain operational over several decades.

The most common signals and protocols will be described in the following sections.

Time Strings

Very common in substation automation using SCADA (Supervisory Control and Data Acquisition) and RTUs (Remote terminal Units) due to the low-cost. Many formats are used, there is no common standard.

Since messages are exchanged over a serial port at low bit-rates (typically 9600 bps), accuracy is low (but still sub-second).

Timing Pulses

Most common one is 1PPS (or PPS, One Pulse Per Second). Sometimes 1PPM (or PPM, One Pulse Per Minute) is also used.

Can be very accurate (sub-microsecond) but is ambiguous since there is more than one possible interpretation for a pulse. Timing pulses are frequently used in serial ports to complement time strings.

Dedicated Timing Signals

The most common ones are IRIG-B and DCF77. In both cases, time and date information is modulated on top of timing pulses.

IRIG-B is very accurate (tens of nanoseconds for DC level shift and Manchester modulation) and widely used in North and South America. There are different types, the most common (and complete) one is IRIG-B004 with IEEE 1344 extensions.

DCF77 is slightly less accurate than IRIG-B but still good enough for all applications in the power industry. Very common in Europe and North Africa.

Ethernet Based

Ethernet-based time protocols have a cost-advantage since no additional wiring is needed in the substation for the distribution of time signals. Obviously, all receiving devices need an Ethernet port and, less obviously, good processing capabilities.

There are two Ethernet-based protocols in use in power substations: SNTP/NTP (very common, up to 1 millisecond accuracy) and PTP (up to 1 microsecond accuracy, part of the IEC61850 Standard, only newer and relatively complex devices support it).

Any master clock supporting Ethernet-based time distribution has to address cybersecurity concerns.

1.2 Distribution of Time and Frequency in Power Substations

For the same reason already mentioned in the previous section (equipment in power substations are a mix of new and old, sometimes spanning a few decades), several different media is used to distribute time and frequency in the power industry, as described below.

Serial Port

Voltage-level based interface (V24/RS232, sometimes RS485). Good for short distances (up to 5 meters, unless optical decouplers are used), more or less only usable inside the same cabinet. Very low-cost, used for time strings and PPS/PPM signals.

Optical Fiber

Point-to-point only communication, good for long distances (up to 1500 meters), very robust. Drivers, cables and connectors are expensive (but not that much if compared to 10 years ago).

Used for PPS, PPM, IRIG-B, DCF77 and, sometimes, time strings.

Twisted-pair cables

Point-to-point connection, short distances, very cheap connectors and cabling.

Used for PPS, PPM, IRIG-B and DCF77.

Hi-Power, galvanically isolated coaxial cable

Multipoint (each output can drive several devices), good for medium distances (50 to 100 meters). Requires coaxial cables, BNC connectors, adapters and terminators but is not as expensive as optical fiber.

Used for PPS, PPM, IRIG-B and DCF77.

Ethernet

Isolated, multi-point, very flexible, no additional wiring required. Only works with relatively new devices.

Used for SNTP/NTP and PTP.

2 Sources of Errors

2.1 Unchecked Delays and Time-Offsets

Time-offset is the difference between the reported time and true time. When using GPS-based master clocks, time-offset will be more or less constant and caused by propagation delays in the antenna cable and the distribution media or by misconfiguration of serial ports.

Complex network topologies might also introduce errors in the delay compensation algorithms used by some network-based time distribution protocols.

Antenna Cable

Propagation delay on the coaxial cable used for the antenna connection to a GPS receiver is around 4 ns/m (1.2 ns/foot) and will not really be a problem unless very long antenna cables are demanded by the application. Some GPS Master Clocks allow for this correction and, if available, it should be used to minimize this time-offset.

Distribution Media

For copper cables the propagation delay is 4 ns/m (1.2 ns/ft). Since copper cables are rarely used for longer distances, this propagation delay can be ignored most of the time.

Optical fiber has a propagation delay of around 6 ns/m (1.8 ns/foot). Care should be taken in this case since optical fibers are frequently used to cover long distances in a substation. A optical fiber cable of 150 m (500 ft) will introduce a time-offset of nearly one microsecond. This will make it completely unsuitable for a synchrophasor application or a traveling wave-based fault locator.

Data Rates and On-Time Marks in Serial Protocols

When using serial protocols (very commonly used by programmable logic controllers, SCADA systems and RTUs), care should be taken in correctly configuring the so-called OTM (on-time mark) character.

At a rate of 9600 bps around 1 ms is needed to transmit each single character. Since typical serial messages have around 10 to 20 characters, an error of nearly a full period can be introduced.

There are two ways around this delay:

1. Combine the serial port information with a PPS signal that will latch the time information into the device. In that case, make sure that the time reported corresponds to the expectations of the device being served (i.e., is the time information AHEAD or BEHIND the PPS pulse). See figures 3 and 2 below.

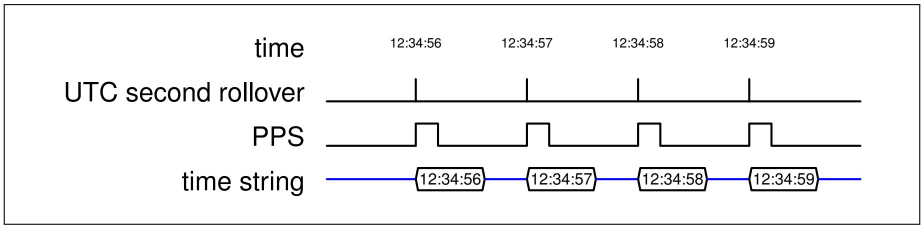

Figure 1: Reported time BEHIND PPS signal

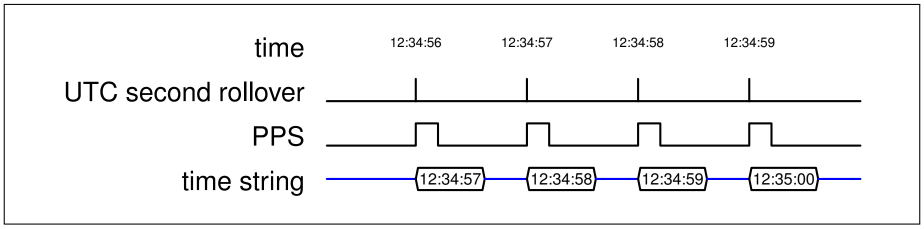

Figure 2: Reported time AHEAD of PPS signal

2. If a separate PPS signal isn't available, the time at which the serial message will be transmitted can be configured in some master clocks. As an example let us assume that a RTU requires a time string compatible with the Traconex 1020

messages. The time part of this message is 15 characters long and the on-time mark is the falling edge of the start bit of the ASCII carriage return at the end of the message. At 9600 bps, 8 data bits, 1 start, 1 stop and no parity bits, the message should be transmitted roughly 13 ms prior to the second rollover. This will make sure that the on-time mark is aligned within one millisecond with the UTC second rollover.

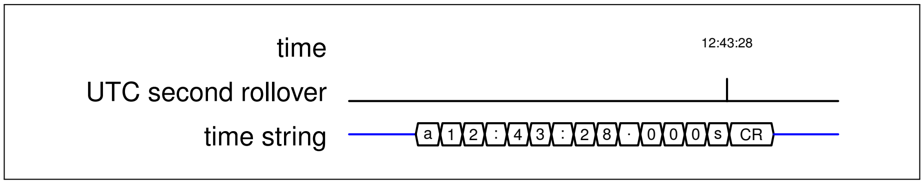

Figure 3: Correct position of OTM character in a Traconex 1020 message

Redundant Ethernet Switches

Every network-based time distribution protocol uses a delay measurement mechanism to correct the offset error introduced by the network layer itself.

In the Network Time Protocol (NTP) the measurement of this delay assumes symmetric times (sending data is delayed as much as receiving data). If redundant switches are used, care must be taken that packets routed in one direction travel along the same (or at least equivalent) route as packets traveling in the other direction. See figure 4 for a representation of the potential problem.

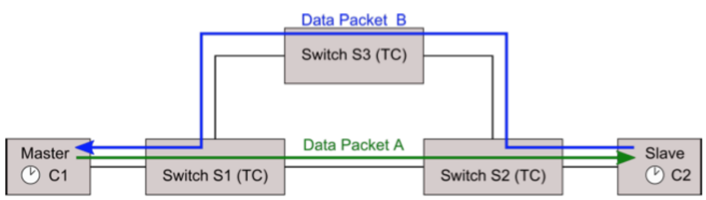

Figure 4: Different packet routing will cause wrong propagation delay estimates

PTP (the Precision-Time Protocol) is not affected by this since it keeps track of the propagation delays of (and compensates for) each network link.

2.2 Holdover mode and Drift

Under normal operation conditions, a GPS-based master clock is "locked" to the time-signals transmitted from the GPS satellites.

However, the GPS signal is a very weak one and this locking might be broken in different ways:

- blocking by nearby buildings or trees (good positioning of antenna is paramount)
- interference from other nearby radio-frequency sources
- jammers (for GPS or mobile networks)

In such cases a GPS-based master clock will switch to a "holdover" mode and try to keep a time estimate until the locking can be re-established.[1]

During holdover the master clock will use a local oscillator. Different types of oscillators might be used resulting in different maximum drift:

Oscillator Type	**Stability (in ppm)**	**1 s per**	**1 ms per**	**1 μs per**
XO (crystal oscillator)	10	1.16 days	100 s	0.1 s
TCXO (temperature compensated)	1	11.6 days	1000 s	1 s
OCXO (oven-controlled)	0.01	116 days	1.16 days	100 s
Rubidium oscillator	0.0001	317 years	116 days	2:46:40

Table 2: Local oscillators and their typical drift

A good master clock will report not only if it is locked to the GPS satellites, but also provide an estimate of the accumulated local oscillator drift (using, for example, the IRIG-B extension bits described in IEEE C37.118).

Once the GPS-based master clock is able to re-establish the locking state, there will a discrepancy between the time reported by the GPS Satellites and the time internally kept in the clock.

A good master clock should slew towards the correct time (make small adjustments to the time reported over several seconds or minutes) instead of making a single step adjustment.

Time is strictly monotonic (it only moves forward and never stands still) and synchronized devices might have trouble with backward jumps.

2.3 Leap Seconds

People around the world use either local time or UTC. Local time is UTC plus minus an offset and (if used) daylight saving time.

[1]Notice that other sat-based systems (GLONASS, Galileo) suffer from similar problems (low intensity signals, similar operating frequencies).

UTC is based on TAI (International Atomic Time), a time standard calculated using a weighted average of signals from atomic clocks located in nearly 70 national laboratories around the world.[2]

However, a "leap second" is occasionally added to UTC in order to keep civil time aligned with the perceived time-of-day. This is due to the Earths rotation speed being constantly slowed down by the tides.

The insertion of a leap second is decide at least six months in advance. In the 40 years (from the standard creation in 1961 up to 2011), a total of 25 leap seconds have been added; the most recent was added on 30 June 2012.

Figure 5: Leap second displayed on a Caesium Beam Tube frequency standard

A good master clock should correctly announce and report the insertion of leap seconds so that the synchronized devices are able to handle it accordingly.

3 Checklist for a sound time reference

Below is a checklist for a sound time reference application.

1. What accuracy is required by the application ?

   Knowing the required accuracy will help judging how important are the next questions.

2. Is the antenna cable delay significant ?

   Antenna cable delays are only significant if very long cables are used or if the required accuracy is 1 μs or better.

3. Is time distribution over optical fiber being considered ? Will it be significant ?

   Time distribution over optical fibers can introduce delays up to 10 μs over long distances.

4. Is a serial port being used with a dedicated time signal ?

   Make sure that the time reported corresponds to the time expected by the device (ahead or behind PPS). Otherwise an offset of 1 s will be noticed.

[2]Time reported by the GPS satellites is TAI-based, with a constant offset of 19 seconds.

5. Is a serial port being used without a dedicated time signal ?

   This will possibly introduce errors up to 20 ms. If possible, adjust the OTM settings of the master clock.

6. Is NTP being used ?

   If yes, check the network topology and make sure that the propagation delays are symmetric.

7. What are the requirements for holdover operation (drift limits) ?

   The requirements for holdover operation will help select the right local oscillator for the master clock.

8. Is an estimate of time error required ?

   If yes, make sure that it is correctly reported by the master clock.

9. Is clock slewing required ?

   If yes, make sure that the master clocks implements it.

10. How are leap seconds announced ?

    Leap seconds should be properly announced and handled, especially for new installations.

About the author

Conrado Seibel received his Bachelor's degree in Electrical Engineering and a Ph.D. in Information Systems from the Universidade Federal de Santa Catarina, Brazil. Conrado has been working in the development of data acquisition and data processing systems for over 25 years. He is currently the lead developer and managing director at mergedK GmbH in Berlin, Germany.