# Specifying Holdover Performance

Bill Dickerson, P.Eng., FIEEE

## Introduction

Though GNSS clocks are extremely reliable, and rarely lose lock, especially not for an indefinite period, modern GNSS clocks increasingly specify holdover performance. We have found that antenna system damage due to unrelated maintenance activities is the No. 1 cause of an unlocked condition. Failures, while rare, still happen occasionally also. Our customers would like to know what to expect in such cases.

## What is the definition of 'holdover?'

Holdover is, quite simply, the time error expected to accumulate over a specified interval when the GNSS clock is unlocked. It includes all sources of error that affect clock performance. Most of them are due to the imperfections of the holdover oscillator incorporated in the clock.

What though is meant by the word 'expected' in that definition? Customers say, 'all I want to know is  how far off the clock will be after a day…' which seems like a simple question. However, holdover error can't be predicted exactly: if it could, we would be able to correct for it. And, we do correct for as much of it as we can measure. Other clock manufacturers probably try to do this as well.

So what we have left, after correcting for the things we think we know about our holdover oscillator, is a combination of random noise and perhaps unknown systematic effects, that is, sensitivities to factors we can't or for practical reasons, don't measure. Such factors must be treated statistically. That is, we measure the actual behavior of the oscillator while it is locked; we assign as much of the variation as we can to 'known' mechanisms including aging and environmental effects (temperature, humidity, and/or barometric pressure for instance); and what remains is called the 'residual' – all of which appears as noise. Then, based on that, we can make a statistical estimate of how far off that clock might have drifted, if it had lost lock, considering the predicted trajectory and the residual errors.

This estimate varies for different holdover intervals. What we are doing is extrapolating all known factors into the future, and evaluating the accuracy of the resulting estimate. Extrapolation is most accurate for a short period; the longer the extrapolation interval, the greater the expected errors in the extrapolation. We report estimated holdover performance for time intervals from 15 minutes to 30 days (depending on clock model). Holdover estimates apply for a single interval after loss of lock; they are not ongoing or cumulative specifications or estimates. Particularly, the expected error for a 7-day period is not exactly seven times the 24-hour number (paradoxically, it can be better or worse).

Making these estimates requires enough data to have some confidence in the estimates. Our clocks generally require seven times the prediction interval; so, to estimate holdover uncertainty for a 24 hour period requires seven days of operating data. Up to a point, additional run time improves the estimate.

Since we measure environmental sensitivities along with aging of the resonator itself, in real time, we can estimate actual holdover uncertainty in that same environment. That is, if the temperature is changing with a ± 5 Kelvins diurnal (daily) period, the estimate of holdover uncertainty applies over that same range. Our holdover estimates are based on that specific oscillator, mounted in that specific clock, under



Arbiter Systems, Inc. · 1324 Vendels Circle, Suite 121 · Paso Robles, CA  93446 · USA
Tel: +1.805.237.3831 · Fax: +1.805.238.5717 · E-mail: sales@arbiter.com · Internet: http://www.arbiter.com

the actual operating conditions it has experienced in the recent past.

## What is the meaning of Arbiter's holdover specification?

Holdover specifications apply for the stated conditions. For our clocks, we specify holdover as if the unit was operated under the complete range of environmental parameters. Of course, this is a pessimistic (i.e., very conservative) assumption, but we do not want to mislead our customers. We would rather you be happy with a product that greatly exceeds its specifications, rather than disappointed and angry with a product that does not meet your expectations and needs.

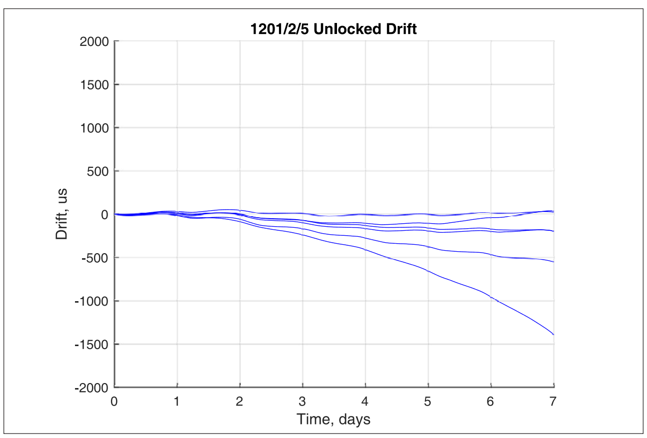


Figure 1: Measured holdover error, 7 days; Models 1201/1202/1205 (OCXO holdover oscillator)

## How about our competitors?

Many vendors specify holdover under unrealistically restricted conditions – including absurd assumptions such as a constant operating temperature. We believe that they do this for these reasons:

1. Holdover is a key specification; they might be able to use it as a lock-out
2. Improving performance costs money (often, quite a lot)
3. Vendors cannot know the actual operating environment
4. Specifying at a constant temperature provides the best holdover for a given clock
5. Customers have no practical way to verify actual performance, so they will believe the claim

The constant temperature condition is particularly troublesome, because holdover performance is in most cases dominated by temperature influences. So, vendors often quote an unrealistic holdover number under unrealistic conditions (fixed temperature), and then they might provide some information about the sensitivity to that condition. For instance, let's say they quote 7 µs / 24 hours at constant temperature, and then in the footnotes they mention 'one part in $10^8$ temperature sensitivity.'

How do you convert that to an expected total holdover number? The fact is, without a lot more information, you probably can't. The one part in $10^8$ is probably their holdover oscillator vendor's specification, and it probably applies over the entire operating range, though what the unit does within that range is also probably not linear or even consistent from unit to unit. Without knowing the actual performance under your conditions, you can't tell what that clock will do if it loses lock.

## Why is constant temperature a bad restriction? I have an air conditioner…

Quite simply, air conditioners have two significant issues: they cycle on and off; and they do not have infinite 'loop gain.' This means that the temperature is constantly changing, and that even the average temperature in the room varies from the warmest to the coolest time of the day. Let's look at what happens if we take that 'one part in $10^8$' and make a few simplifying assumptions.

The first assumption is that this is a linear drift: the oscillator is - 1 x $10^{-8}$ at its lower operating limit (let's say - 20 °C) and + 1 x $10^{-8}$ at its upper limit, 65 °C. That gives a temperature slope of 2.35 x $10^{-10}$ per degree. And for simplicity, let's just say that the temperature changes by 3 degrees and stays there. That alone results in a holdover error of 60 µs/day: 86 400 s/day * 3 K * 2.35 x $10^{-10}$ / K. Which is about nine times worse than the vendor's 'holdover' specification.

'Real' conditions are far more complicated to evaluate. Which is why we provide actual, real-time estimates of expected holdover uncertainty for our clocks under your actual operating conditions. We even allow you, the user, to specify how conservative the estimates are. We use a default estimator of two sigma, which means that the actual holdover error will be less than predicted ~ 95 % of the time. You can change it, for instance, to a three-sigma estimate, in which case the actual will be less than predicted ~ 99.7 % of the time.

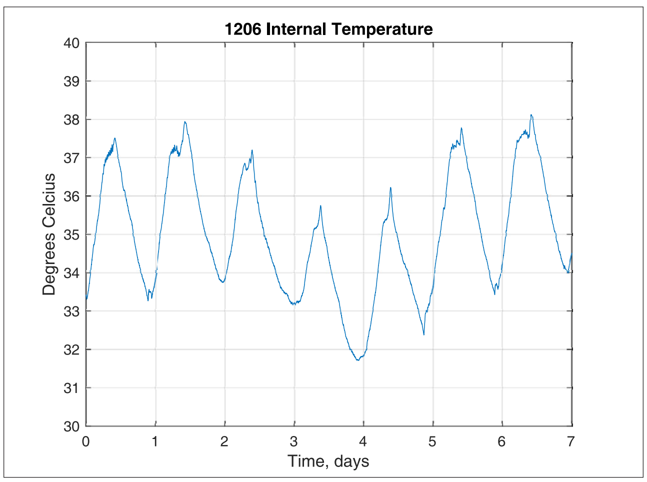


Figure 2: Measured temperature variations, 7 days

## Why is holdover performance important?

There is one obvious reason, and one not so obvious. The obvious reason is that this gives an indication of how far off the clock might be in case GNSS is lost. As I said before, this is normally due to damage rather than failure, but anything is possible.

The less obvious reason is that one of the indicators of spoofing is an unexpected time error from the GNSS receiver. So, the more stable the holdover oscillator, the more sensitive our spoofing detection becomes. You would not be happy with us if we raised a 'spoofing alarm' because the temperature of the clock changed... so, by actually measuring performance under actual operating conditions, we can get the best estimate of likely oscillator stability – and the best sensitivity for spoofing detection, without raising false alarms. Of course, spoofing detection is also based on numerous other factors; but the better the performance of the holdover oscillator, the more likely that we can correctly detect spoofing.

## What do we do to optimize our holdover performance?

Arbiter Systems, Inc. holds two patents (US 9 362 926 and US 9 979 406) for 'High-Reliability Holdover Method and Topologies' which offer several improvements to previous state of the art. One of the key improvements is achieved by using a fractional-N loop to control our system clock frequency, rather than an analog-tuned holdover oscillator.

Analog-tuned oscillators have a number of limitations: the tuning process is noisy and non-linear; errors in the tune voltage generator degrade performance; and even more importantly they are not necessarily repeatable. By allowing the holdover oscillator to 'free run' and locking a flywheel oscillator to it with the fractional-N loop, we eliminate all of the uncertainties related to analog tuning. This gives us exact digital control over the holdover process. Benefits include exact linearity and repeatability, essentially unlimited resolution, and freedom from added noise and errors due to analog limitations. When we command the fractional-N loop to adjust the frequency by 0.1 ppm, we can be confident that it will change by exactly that amount every time – not 0.1 ppm ± 25 %.

Having exact measurements of the holdover oscillator frequency as it changes under real operating conditions also enables accurate estimation of the predicted holdover performance. Using an analog tune signal with its ± 25 % non-linearity and often even more loosely-specified tune coefficient to estimate oscillator frequency degrades performance estimates accordingly.

The consequence of all of this is that we can offer exceptional real-world holdover performance for a reasonable price. The figures below show the results. For our standard OCXO, specified at 1 ms/24 hours, 'typical' performance is over forty times better (Figure 1). That's right, forty (40) times. This is measured in the lobby of our facility in Paso Robles, California (Figure 3) – a part of the building that 'inherits' its air conditioning from neighboring spaces, and is subject to occasional blasts of hot or cold outside air when the door is opened about 3 m (10 ft) away. Further, for energy conservation, the HVAC system is turned off when the building is unoccupied, so the temperature is allowed to vary significantly for more than 100 hours of the 168-hour week, which is about 2/3 of the time (Figure 2).

Figure 3: GNSS clock long-term performance evaluation test system (the 'Droid Ranch')

Also shown: typical performance for Model 1206C clock, which incorporates a rubidium holdover oscillator (Figure 4).

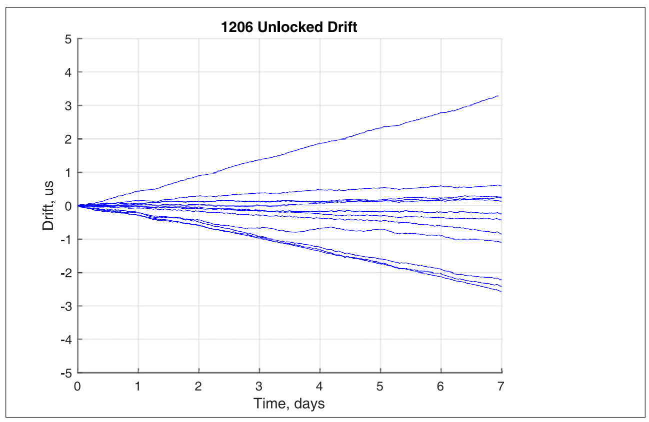


Figure 4: Measured holdover error, 7 days, Model 1206C (rubidium holdover oscillator)



Arbiter Systems, Inc. · 1324 Vendels Circle, Suite 121 · Paso Robles, CA 93446 · USA
Tel: +1.805.237.3831 · Fax: +1.805.238.5717 · E-mail: sales@arbiter.com · Internet: http://www.arbiter.com

## Annex: What are the consequences of holdover error?

Let's say we have a GNSS clock that has lost lock. A reasonable question is, what will be the effect on our ability to maintain normal operations under this condition? Arbiter's clocks report their estimated actual holdover error as part of the relevant time codes and protocols, so that each application can make an informed decision regarding what to do as the error accumulates.

Table 1, which follows, shows the effects of varying holdover error on some power system applications.

| **Application** | **Holdover Error** | | | | | |
|---|---|---|---|---|---|---|
| | 0.1 µs | 1 µs | 10 µs | 100 µs | 1 ms | 10 ms |
| Traveling-wave fault location (location error) | 15 m (50 ft) | 150 m (500 ft) | 1.5 km (1 mile) | *15 km (10 miles)* | *150 km (100 miles)* | |
| Synchrophasor measurement unit (PMU) (phase error) | 0.002 ° | 0.02 ° | 0.2 ° | 2 ° | *20 °* | *200 °* |
| Merging unit/ sampled values (MU / SV) | Absolute phase error: same as for PMU above<br>Relative phase error (between MUs connected to same clock): none<br>Recommendation: connect all MUs on a process bus to same clock | | | | | |
| Relay, event/fault recorder | Meets requirements of NERC PRC-002-2 requirement R10 | | | | | *Exceeds limits* |

Table 1: Effect of holdover errors on some power system applications. Normal font: usable for most applications; *italic font*: probably not usable.