SyncServer® S600/S650 Series Timing Architecture

Summary

In the world of up-to-date, high-accuracy time servers, the fundamental technique to keep a time server extremely accurate has changed from traditional varactor-steering of a physical oscillator to deploying superior multifactor software-based algorithms coupled with a free-running oscillator and frequency synthesizer. The SyncServer® S600 series implements this proven timing architecture for world-class time and frequency performance. Understanding the basics of how the SyncServer operates helps to answer many questions about how it behaves in different scenarios.

The Role of Internal Oscillator

At the core of time and frequency synchronization technology is the capability to accurately transfer time and/or frequency from a reference time source to the local timing outputs, like the outputs from a SyncServer time server. Central to the time server is the local oscillator, which provides a stable internal frequency reference that effectively provides a buffer between an external reference input and local output. Accurately frequency-aligning the internal oscillator to a reference source is traditionally referred to as "oscillator disciplining". This term derives from the common practice of directly steering the oscillator to match the selected input reference frequency. Because many types of oscillators inherently have significant frequency error and drift (as they are not perfect resonators), the disciplining process is the technique to carefully "steer" the local oscillator frequency to keep it as on-frequency as possible relative to the reference.

GPS/GNSS as the Reference Time Source

The most common time reference used to discipline a local oscillator in a stratum 1 server is the GPS satellite system, as well as other GNSS constellations. Each GPS satellite has an atomic clock on board whose accuracy is traceable to the many atomic clocks located at the U.S. Naval Observatory (USNO). The GPS satellite system offers broad signal coverage over the planet, from which a SyncServer locks to the available broadcast signals, providing a reliable and accurate time source from which to precisely synchronize the time and frequency. Within the SyncServer, GPS satellite data is converted to a 1PPS signal that provides the on-time mark used to derive that precise time and frequency.

Historical Varactor-Steering

A traditional way to steer an oscillator involves creating a closed-loop feedback system where the frequency output of the local oscillator is measured against the output of the GPS receiver. The local oscillator is actively varactor-steered to adjust its output frequency to closely follow the GPS 1PPS frequency (there is filtering involved). This technique is called varactor-steering and while it has its place and tradeoffs and has been used for many years, the technology deployed in the SyncServer substantially improves upon the varactor-steering technique.

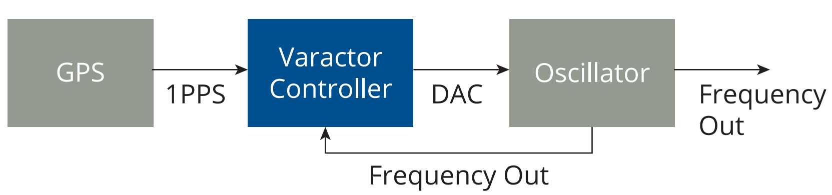

Figure 1: Traditional GPS-referenced clocks use varactor-steering to discipline the local oscillator.

Avoiding the Downsides of Varactor-Steering

The SyncServer S600 series uses a modified approach to system frequency-steering that avoids active varactor-steering adjustment of the internal frequency reference. Figure 2 shows the simplified architecture compared with the direct oscillator steering method shown in Figure 1. The internal oscillator is allowed to free run (no active adjustment) and becomes an input to a high-resolution frequency synthesizer. With this approach, frequency adjustments are applied directly to registers in the synthesizer, resulting in an "on-frequency" output from the synthesizer. This output is then used to drive the frequency of all outputs in the SyncServer. The benefits of this approach compared with traditional varactor-steering are numerous:
The synthesizer provides a linear frequency control (a given step-size change of rate-control results in a predictable frequency change). In contrast, direct varactor-based oscillator steering can have significant non-linearity in steering, leading to variable and non-optimized output behavior (depending on the portion of the steering curve being used).

- Because this control system exhibits improved linearity, it allows signal-processing algorithms to be more effective, unencumbered by the deleterious effects associated with unaccounted for non-linearity.
- The on-frequency output from the core is decoupled from the internal reference oscillator frequency, which allows for driving the internal system with a much higher initial on-frequency signal.
- The synthesizer provides high-resolution frequency control that is the same regardless of the selected reference oscillator. This is compared to direct oscillator steering where the rate control sensitivity will likely vary between different internal oscillators that can be installed in the unit (for example, standard quartz, OCXO, Rubidium).

For these reasons, the incorporation of a free-running internal oscillator followed by synthesizer provides an improved overall approach for the S600 SyncServer series core clock architecture.

Superior Softclock Technology

The synthesizer-based rate control is combined with an embedded software library (known at Microchip as Softclock) that takes full advantage of the overall system linearity and awareness of the stability characteristics of the specific internal reference oscillator to optimize system time and frequency. The Softclock software simultaneously observes all provided time and frequency references to ensure that smooth transitions between references will occur when needed.

Softclock Learns and Adapts to Multiple Inputs and Environmental Conditions

The benefit of the Softclock system is a more accurate and stable time server relative to UTC(USNO) through GPS/GNSS or whatever time reference is provided. The Softclock leverages multivariate algorithms to attain a form of artificial intelligence for synchronization. The system can monitor and adapt to multiple simultaneous inputs in addition to GPS, such as IRIG in, PTP in, NTP peering, 10 MHz in, and 1PPS in. Additionally, internal temperature is monitored and (in some cases) servo-controlled as a method to reduce potential temperature-associated timing errors. This includes anticipating frequency shifts in the free running local oscillator as a function of observed temperature change and air flow inside the SyncServer chassis. Through manufacturing calibration, the SyncServer achieves a specification of 15 ns RMS UTC(USNO).

The primary function of the SyncServer is to output accurate timing signals through the Ethernet ports or BNC connections. To accomplish this, the output of the frequency synthesizer serves as the rate reference basis for all SyncServer outputs ranging from the 1PPS to the IRIG signals, sine waves, programmable rates, NTP and PTP timestamping operations, and so on. All of these signals are connected to the common core that is initiated from the synthesizer output.

Demonstrated, Industry-Proven Softclock Performance

While the Softclock technology is new to SyncServers, the technology itself is patented and has been refined for over 10 years in Microchip telecommunication synchronization systems widely deployed around the world. The result is a robust and stable solution that is not only used in Microchip products, but also licensed to manufacturers that need precise timing in their products. The chances are high that the mobile backhaul of voice and data for your cellular phone is being synchronized by Softclock systems.

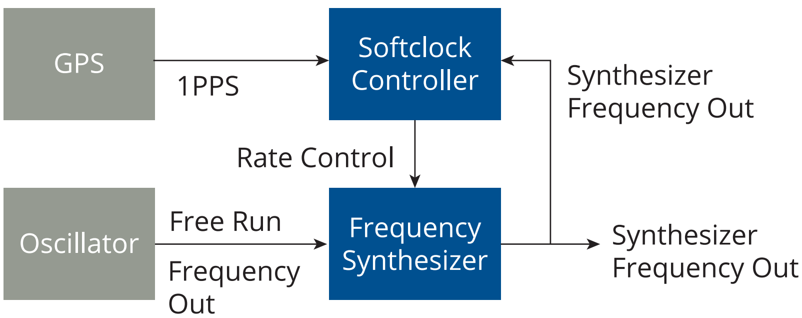

Figure 2: The patented Softclock controller with frequency synthesizer improves overall time and frequency accuracy and stability.

Oscillator Choice Matters

As good as the SyncServer timing architecture coupled with Softclock technology is, the choice of which oscillator is installed in the SyncServer is important as well. In general, the less expensive the oscillator, the less stable it is. When locked to a reference, the stability of the internal oscillator impacts two important areas:

- A less stable oscillator requires more frequent synthesizer adjustments in order to bound the negative impact of oscillator drift. This results in higher variance on timing outputs when a less stable oscillator is used, even with a highly stable input reference such as GPS.
- When potentially less stable input references are provided—such as PTP passing through multiple network hops—the ability to recover good synchronization can depend on the maximum time over which the input can be observed without oscillator instability becoming a factor. With more stable oscillators, the ability to observe for much longer before making control decisions can significantly improve the output timing performance.

When connection to input references is lost, the SyncServer depends on the Softclock historical modeling and calibration of the installed oscillator to maintain time and frequency performance. This condition is referred to as "holdover", and the performance achieved is strongly dependent on the selected oscillator stability. The SyncServer 600 series holdover performance by oscillator type is very good, as reflected in the commonly-cited industry comparison accuracy after 24 hours. This information is listed in the following table.

Oscillator	Holdover Drift (1st 24 hours)
Standard	400 microseconds
OCXO	25 microseconds
Rubidium	<1 microsecond
External Cesium*	<100 nanoseconds

Holdover drift is the anticipated accuracy of the SyncServer over a period of time if disconnected from all external time references.

*Cesium Standard with ~1 x 10^{-12} frequency stability.

Excellent Output Signal Coherency for Rates, Timecodes, and Standard Sine Waves

SyncServers are deployed in a variety of applications where timing output signal alignment to each other and to UTC becomes particularly critical. SyncServer Timing I/O modules that output time codes and rates are coherent with each output on the module, as well as with alignment to other modules and UTC when tracking GPS/GNSS. This includes sinewaves generated from the Timing I/O modules. Also, programmable pulse outputs set to periods that do not repeat every second will still maintain coherency with other similar outputs because they all begin together and remain aligned thereafter.

Low Phase Noise and Coherency

Low Phase Noise (LPN) applications are a situation where direct oscillator steering may be needed. This is because one of the only practical ways to accomplish close-in phase-noise specifications is directly from a quartz oscillator output that exhibits the needed close-in characteristics (which is a specially made oscillator for this purpose). The technique is then to position the oscillator as close to the output connections as possible, minimizing any signal degradation due to excessive signal conditioning in the path.

Despite the outstanding LPN properties of these special quartz oscillators, they are still not inherently accurate and will drift. The solution in SyncServer 600 series is to use the synthesizer-based steering method (described earlier) as the "on frequency" signal that then controls the LPN oscillator in order to also keep it on frequency. Care is taken to minimize the rate of steering of the LPN oscillator because adjusting it too fast can degrade the phase-noise performance. The relationship is "loosely-coupled." When using an LPN module in a SyncServer, the main reference oscillator should be a Microchip Rubidium Miniature Atomic Clock (MAC) because its inherent longer-term stability leads to less frequent adjustments at the LPN oscillator input. This provides excellent long-term stability of the 10MHz signal as well as the exceptionally good LPN characteristics.

All LPN 10 MHz outputs from the SyncServer are coherent. However, some applications also require exact coherency between the 10 MHz outputs and the UTC-aligned 1PPS output to be near perfect. To achieve this with little or no close-in phase noise degradation, the installed main oscillator in the SyncServer must be either an OCXO or the preferred Rubidium MAC.

Softclock Time and Frequency Technology is the Future

The Microchip timing experts have learned that with the inexpensive computational power and time and frequency hardware available today, a superior timing system no longer relies on the physics of vibrating quartz crystals or atomic oscillators alone to keep a time server extremely accurate. Software modeling and compensating for non-linearities in oscillator behavior, monitoring changing environmental conditions, and better overall electronic design and packaging delivers the outstanding time and frequency performance found in an advanced time server like the SyncServer S600/S650 series.