BY JIM OLSEN AND KISHAN SHENOI • SYMMETRICOM INC

Accurate clocks optimize network service

ATOMIC CLOCKS EXTEND SIGNAL HOLDOVER TO MAINTAIN ACCURATE NETWORK TIMING AND ENSURE CONSISTENT SERVICE QUALITY DURING A GPS OUTAGE.

Wire-line- and wireless-communication systems depend on continuous access to a reliable PRS (primary-reference-source) timing signal to ensure high-quality service, especially for real-time applications, such as voice, IPTV (Internet Protocol television), and online gaming. Accurate and stable timing plays a critical role in maintaining synchronization between connected nodes in the network. Synchronized nodes improve handshaking and packet transfers and minimize latency, errors, and dropped packets; they also maximize available-bandwidth usage.

The GPS (global positioning system) is an ideal reference signal because it is readily accessible, traceable to UCT (universal coordinated time), and cost-effective for many applications. However, continuous access to GPS is not always guaranteed. When the GPS-reference signal is unavailable, network elements begin to fall out of synchronization, making it more difficult to maintain a reliable network link. As a consequence, network efficiency and reliability significantly decrease, resulting in lost packets, increased network contention, and lower overall quality of service. These problems can directly impact revenue, as they do in the cellular market, in which the competitive spotlight has shifted from comprehensive coverage to network stability. Dropped calls remain one of the key reasons users cite when changing service providers.

When you lose the GPS reference, a holdover oscillator holds the last good frequency signal until the provider restores the PRS. The quality and performance of the holdover oscillator determine the service performance of the network, as well as how long it will continue to operate efficiently once it goes into holdover mode. Compact rubidium-based atomic oscillators, for example, maintain Stratum 1 quality synchronization and timing to maintain services during extended holdover periods (Reference 1).

Certain characteristics of a holdover oscillator are critical when timing goes into holdover operation, and it is important to select the most effective oscillator for a particular application and network element. Diligent system design can generally minimize initial frequency-offset errors, leaving the frequency-aging rate and temperature stability as the dominant performance factors determining holdover performance.

THE TIMING HIERARCHY

It is important to understand that reliable networks must maintain a timing hierarchy to avoid degrading clock-holdover mechanisms. There are many ways in which you can create loops between, for example, the CO (central office) and smaller remote offices using the POP (Post Office Protocol), because people manage and engineer provisioning rules. It is relatively easy to create a distribution-timing loop within a network, such as Network Element A providing a reference to Network Element B, which provides a reference to Network Element C, which provides a reference signal back to Network Element A (**Figure 1**). SONET (synchronous-optical-network) and SDH (synchronous-digital-hierarchy)-ring applications are particularly prone to timing loops because the architectures themselves create a loop through which timing must be distributed.

Timing loops can unexpectedly wreak havoc on the network when the network loses the primary GPS reference. The loop

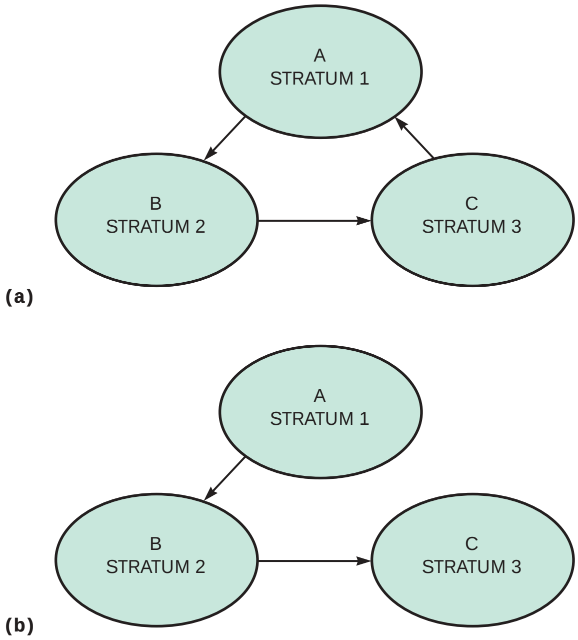

Figure 1 Distribution-timing loops (a), in which Network Element A provides a reference to Element B, which provides a reference to Element C, which provides a reference signal back to Element A, can quickly degrade holdover-oscillator timing signals unless a timing hierarchy exists (b). In a timing hierarchy, elements can receive a timing signal only from an element higher in the hierarchy.

drives oscillators to the end of their tuning range, which can create a large frequency offset that can isolate network elements and create traffic outages. The tuning range of an oscillator is a measure of how far the oscillator can move from its center frequency. The oscillator centers on the GPS-based PRS signal when it is available and maintains synchronization between network elements. When the GPS-based PRS-reference signal is unavailable, however, the oscillator pulls away from its center frequency, and the oscillators in network elements begin to pull from each other. The farther an oscillator is from its center frequency, the shorter the period of time it can provide accurate synchronization.

When a timing-reference signal passes around a loop, its accuracy quickly degrades. Consider, for example, a crystal that has a tuning range of one part in 10^8 that centers on an incoming PRS. Depending on the shared reference signal, the crystal can pull either above or below the center frequency. If the reference signal is bad or corrupted because of a timing loop, the crystal quickly pulls one way or the other to the edge of the tuning range. The narrower the tuning range of an oscillator, the less it can pull away from the center frequency. The tuning range and intrinsic accuracy of an oscillator are related entities. For example, rubidium oscillators have a tuning range (intrinsic accuracy) on the order of parts in 10^9, which means that they cannot pull far from their original PRS reference; even if you drive them to the end of their tuning range, rubidium oscillators can still maintain an accurate enough synchronization signal to avoid immediately impacting the network.

One method for distributing timing that reduces the possibility of timing loops is to propagate timing references according to a hierarchy. In other words, a timing-reference signal never follows a less accurate reference. For example, a Stratum 2 network component can feed a timing-reference signal to a Stratum 3 component, but Stratum 3 components cannot feed signals to Stratum 2 components. Thus, a signal from a less accurate holdover oscillator with a wide tuning range can never corrupt an accurate holdover oscillator with a narrow tuning range. Table 1 shows the four stratum levels, their respective accuracy, their typical placement in a network, and the type of holdover oscillator each network requires.

It is important to be aware of the difference between a timing-reference signal that is off center frequency and one that has drifted. When a PRS is available through a GPS-based reference source, the holdover oscillator locks to the PRS because it is constantly recalibrating to the PRS. Drift comes into play when the PRS is no longer available and the system must rely on the holdover oscillator. At this point, the holdover oscillator is on its own. The accuracy of the holdover oscillator determines how quickly the center frequency of the oscillator drifts. Thus, the ability of a holdover oscillator to maintain an accurate synchronization timing reference depends on both the tuning range and the drift of the oscillator (Figure 2). Table 2 compares the tuning range and drift of various oscillators.

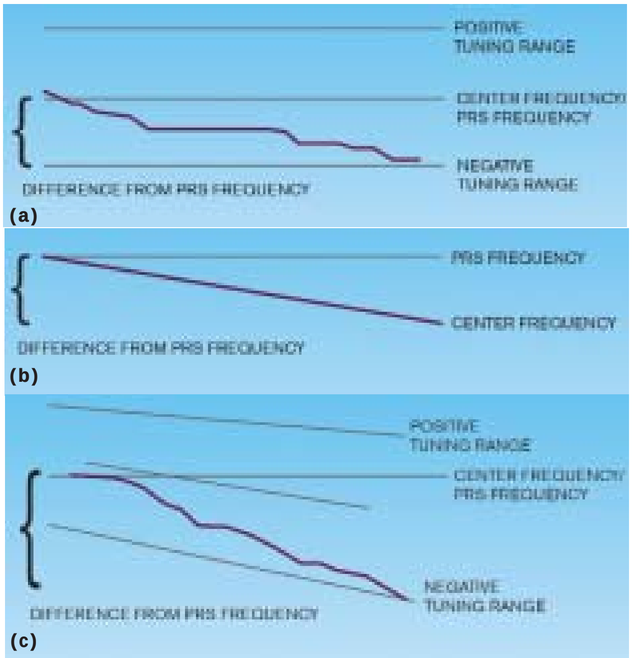

Figure 2 The tuning range of a holdover oscillator is a measure of its ability to stay close to its center frequency (a). Drift is the change in center frequency over time (b). Accurate synchronization depends on both the tuning range and the drift of an oscillator (c).

SETTLING TIME

A holdover oscillator can support a single network element or clock distribution throughout a network. This factor determines the accuracy of the oscillator's overall timing-reference signal. However, another aspect that you must consider is settling time. "Settling time" refers to how long it takes a system to adjust to the timing-reference signal from a holdover oscillator alone—that is, after it loses the PRS—as well as how long it takes a system to readjust to the PRS once the oscillator recovers the GPS-based reference signal.

Traditionally, oscillators remained on frequency through the use of a tuning voltage. As an oscillator drifted, the system would keep the oscillator on frequency by adjusting the tuning voltage. Sooner or later, however, the system would hit the end of the voltage-tuning range. Depending on the oscillator and application, this time would be two to 10 years. Now, the network element requires a manual recalibration to move the oscillator back to its center voltage.

Modern network-timing systems have eliminated the need for

(continued on page 78)

TABLE 1 STRATUM LEVELS, PRECISION, PLACEMENT IN A NETWORK, AND BEST SUITED OSCILLATOR

Level	Accuracy	Place in network	Holdover oscillator
Stratum 1	PRS: 1×10^{-11}	Synchronization feed	Local cesium clock or rubidium oscillator
Stratum 2	9×10^{-10}	Synchronization elements	Rubidium oscillator
Stratum 3	Parts in 10^{-8}	Network elements (typically end office)	Ovenized crystal oscillator
Stratum 4	Parts in 10^{-6}	Network elements (typically customer premises)	Crystal oscillator

TABLE 2 OSCILLATOR TYPES, ACCURACY, AND DRIFT

Technology	Intrinsic accuracy	Stability (1 sec)	Stability (floor)	Initial to ultimate aging (per day)	Hold Stratum 1	Cost
Hydrogen maser	$\sim 10^{-10}$	$\sim 10^{-13}$	$\sim 10^{-15}$	10^{-15} to 10^{-16}	Years	~200×
Cesium beam	$\sim 10^{-11}$	$\sim 10^{-11}$	$\sim 10^{-14}$	NA	NA	~20×
Rubidium-gas cell	$\sim 10^{-9}$	$\sim 10^{-11}$	$\sim 10^{-13}$	10^{-11} to 10^{-13}	Days	~×
High-quality quartz	10^{-6} to 10^{-8}	$\sim 10^{-12}$	$\sim 10^{-12}$	10^{-9} to 10^{-11}	Hours	~0.7×

and high expense of manual tuning through the use of DDS (direct digital synthesis). DDS circuitry allows the use of a free-running oscillator to create an internal reference signal that can lock to the incoming PRS feed. DDS circuits use well-known digital elements, such as accumulators. A digital word controls the output-frequency rate, simplifying microprocessor control. The precision of the system can generate a frequency value that directly links to the word length of the implementation. The range of a DDS circuit is so large that it can easily compensate for the drift in an oscillator over more than 20 years.

When the system loses the GPS-based PRS feed, DDS relies solely on the signal from the oscillator, which differs somewhat from the combined PRS-plus-oscillator reference. If the change in reference frequency is too large, it causes a phase transient in the timing signal that may cause network elements to unlock. As these elements attempt to relock to the reference signal, the network element will experience disruptions, errors, and dropped traffic. To avoid losing lock, the system must adjust slowly over time to the new reference signal.

MORE AT EDN.COM ▷

+ Go to www.edn.com/ms4195 and click on Feedback Loop to post a comment on this article.

Settling occurs whenever the reference-timing signal of a network element changes. Thus, even if the system only briefly loses the GPS-based reference signal, the system needs to settle to adjust to both the loss and the immediate regaining of the PRS. The longer the settling time, the greater the impact on system reliability.

In general, settling time depends on the overall difference between the PRS and the holdover frequency. The farther the frequency has drifted, the longer the settling time. Tuning range also comes into play, because the narrower the tuning range, the faster an oscillator can settle and restore network stability.

PUTTING IT ALL TOGETHER

As you can see, you must consider many variables when selecting a holdover oscillator to use at each level of a network. Typically, the more accurate and reliable an oscillator is, the more expensive it is. Arbitrarily using less accurate holdover oscillators directly affects network reliability. On the other hand, specifying too accurate a holdover oscillator unnecessarily increases system cost.

First, consider how long a system can afford to operate using a holdover frequency. If a technician can quickly address the loss of a PRS, then no pressing need exists for a long-term holdover capability. However, if signal loss could extend across days or weeks, more accurate holdover is the only way to maintain the synchronization to uphold network reliability during this time. A second relevant aspect to consider is accuracy. Crystal oscillators provide a sufficient timing signal for hours to days. Hence, after a few hours or days, timing accuracy for this network element decreases to the accuracy of the crystal rather than to the original reference or PRS that the oscillator was following. For mission-critical applications, a rubidium oscillator holds an accurate reference for more than 30 days, depending on the application.

Developers also need to consider the overall impact of a holdover oscillator on the system. For example, if a network element propagates timing signals down through the timing hierarchy, a less accurate holdover oscillator speeds the degradation of holdover oscillators further down the hierarchy. Finally, developers need to take into account the impact that settling time has on the network. For example, although a network element may require only the precision of an ovenized crystal oscillator, a developer may opt to implement a rubidium oscillator to minimize settling time and, as a result, reduce the impact of PRS loss and gain.

Synchronization is the foundation of reliable connectivity in networks. Using an accurate timing hierarchy maximizes bandwidth usage and uptime. Reliable networks require the presence of a holdover oscillator to maintain accurate network timing when something interrupts access to a PRS. By understanding the difference between oscillator precision, drift, tuning range, and settling time, engineers can select the most cost-effective and reliable holdover oscillator to reduce unnecessary system downtime and loss of service quality.EDN

REFERENCE

1 "Synchronization Interface Standards for Digital Networks," American National Standards Institute, ANSI/T1.101-1998, www.ansi.org.

AUTHORS' BIOGRAPHIES

Jim Olsen is a business-development director at Symmetricom (San Jose, CA), where he has worked since 1984. He has designed and implemented telecommunication-synchronization networks and works with telecom manufacturers on adopting and integrating new technologies. Before joining Symmetricom, he served for six years in the US Navy in electronic warfare. Olsen is a much-sought-after speaker at national and international conferences and has authored numerous white papers and articles about synchronization and next-generation networks.

Kishan Shenoi, PhD, is a chief technologist at Symmetricom (San Jose, CA), where he has worked since 1992. He works on the research and development of the company's synchronization and timing technologies from the conceptual phase to the prototype phase. The industry knows him as Professor Sync, the dean of Sync University, an online, industry-neutral forum that provides network engineers information ranging from synchronization basics to next-generation network issues. Shenoi has written numerous white papers and application briefs about issues that affect telecom-synchronization networks. He received a doctorate in electrical engineering from Stanford University (Stanford, CA).