In Sync with GPS:

GPS Clocks for the

Wireless Infrastructure

by

Peter Kuykendall and

Dr. Peter V. W. Loomis

Trimble Navigation

There is a subtle revolution taking place. Quietly wireless infrastructure is conspiring to overthrow the tried and true system of copper wire and land-line telephony. First pagers moved into our lives. Initially used by doctors, the pager has proliferated so that even moms use them to keep track of their kids. Next the cellular phone moved from a tool so expensive that only high-powered salespeople and important executives could afford it to a device so common it is starting to lose its standing as a status symbol.

The wireless revolution is not complete. One of the current frontiers is competition for local telephone service. The Baby Bells control the copper wire infrastructure and charge for its use. A new technology, wireless local loop, provides telephone service without using a hardwire connection to the user. This technology is expected to bring telephones to third world nations in Asia and Africa where wire-line infrastructure was never built. It is easy to see that the demand for wireless services will increase as these services weave their way into the fabric of our lives.

There is a catch to this revolution. Unlike wire-based systems where lines can be put up until they obliterate the sky, each wireless system requires its own unique slice of the limited radio spectrum. For practical use, this spectrum is physically limited from a few hundred kHz to more than 1000 MHz. In order to get the most out of its assigned slice of the radio spectrum, a wireless system must be carefully timed and synchronized. This is a difficult problem as the timing required can be sub-microsecond precision for base-stations located across very large geographical areas.

Fortunately there is an elegantly simple and cost effective solution in the form of a “GPS Clock.” A GPS Clock is a combination of a GPS receiver and a high-quality, stable oscillator. GPS is used to discipline (calibrate) the oscillator to remove small biases in the frequency. The GPS Clock can synchronize both system timing and transceiver frequency, and is virtually fail-safe. It generates timing signals whenever there is power, and never needs to be re-calibrated. Most importantly, the recent emphasis on mass-production of GPS technology for the automobile navigation market has lowered the price of GPS hardware drastically. This allows GPS integration into many installations requiring high-quality timing subsystems.

PAGING: A Simple Example

Paging provides a good example for understanding how timing synchronization is important to the wireless infrastructure. Most of us have sent or received a page at some point in our lives. It is a simple procedure. You call up the paging service and are prompted to enter a telephone number and then hang up. The page is sent and, with any luck, the “pagee” returns the call. Behind the scenes, every tower in the system is sending that same page simultaneously. The logic is that the paging signal will reach the receiver wherever it may be in the service area, even if the receiver is deep within a building. It is likely when you receive your page it will reach you from two or even three different towers. If the signals were not synchronized, the page might be received and announced two or three times.

The precision required by a paging tower is only about a microsecond, easily provided by GPS. A typical installation on a paging transmitter tower uses a "smart antenna", a GPS receiver and antenna in a single weather-tight unit. Paging infrastructure providers prefer smart antennas because they are easy to install and perform reliably in the hostile radio frequency environment of the paging tower.

The paging towers are able to synchronize off the Pulse Per Second (PPS) signal which is a standard output of many GPS receivers. It may seem the GPS $\pm$600 nanosecond (nsec) accuracy is overqualified for this application, but there is no other way to provide fail-safe timing as cost-effectively.

CDMA Mobile Telephones: A Challenging Application

The cell phone revolution started with a technology called Advanced Mobile Phone System or AMPS. Most of the phones and cell sites in the United States are still based on this analog technology. The weakness of AMPS is that it can support only a limited number of users.

To support more users, cell phone providers must move to a digital system. There are two principal types of digital systems, Time Division Multiple Access (TDMA) and Code Division Multiple Access (CDMA). TDMA shares its allocated slice of the frequency spectrum among a number of users and gives each access based on a unique time slot. CDMA shares the frequency slice based on spread spectrum technology, in which all users broadcast and transmit simultaneously. Each user's transmission is "spread" digitally across the allowable frequency band with a unique code that distinguishes it from other users' transmissions at the same frequency. The unique codes are used to separate users' transmissions at the base station. CDMA offers the better sound quality and allows the most number of users on the system.

CDMA is a challenge to implement. The spread spectrum signal requires sophisticated broadcast power management and "soft hand-offs" between base stations. This requires that the base stations be precisely timed. With the upcoming CDMA wireless telephony systems, each transmitter must be maintain its frequency to within one part in 10^{10} — no more than 7 microseconds in a day. This requirement, known as "holdover", assures the high reliability of the telephone system. Should anything happen to an antenna or cable the phone companies want every chance for the system to continue working until a service crew can reach it. Early CDMA systems used atomic frequency standards to maintain this accuracy. The current generation uses a GPS Clock. Interestingly, neither of the GPS Clock's two components, a GPS receiver and a high-quality crystal oscillator, can meet the demanding specification alone. The long-term stability of GPS complements the short-term accuracy of the crystal oscillator to complete the task.

Enhanced 911 Location

One sign (or symptom) of the success of the wireless revolution is that there are a large number of emergency calls coming in from cell phones. Motorists are calling 911 to report drunk drivers, accidents, and car fires. Unfortunately, when a 911 distress call is placed from a cell phone, the dispatcher has no access to the caller's location. The locations of landline phones are stored in a sophisticated 911 data base and used to direct the rescuer to the scene, but there is currently no comparable method of locating a mobile phone. Consequently, the dispatcher must rely on information from the cell phone user, who often cannot give his location accurately, and delaying resolution of the crisis. A recent example: a woman was stranded in a blizzard for hours while a rescue team searched for her car.

In response to this problem Congress has mandated an *enhanced* 911 or E911 system that is able to locate callers within 125 meters be implemented within the next few years. There are a number of ways to do this, one of the suggestions being to include GPS receivers in all mobile phones. The leading candidate, however, is to locate the caller's handset directly with the cell towers, using a method called Time Difference Of Arrival (TDOA).

TDOA works by measuring the difference in arrival time of the signal between two cell sites. This is also known as "hyperbolic navigation", and is the basis for many radio-navigation systems, including GPS and LORAN. The technique is simple: a call transmitted from a cell phone arrives at two separate cell sites, but the closer site receives the signal a bit earlier. The difference in signal arrival time is converted to difference in distance using the speed of light. The knowledge that the caller is 1200 meters closer to one cell site than the other cell site places his position on a hyperbolic curve on the ground. Adding a third cell site, and another hyperbolic curve, pinpoints the caller's position at

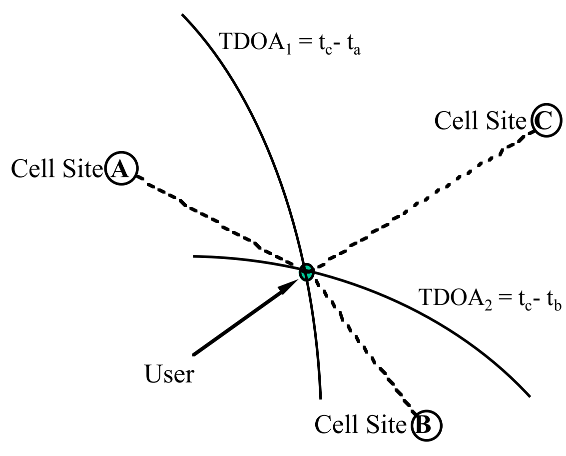

Figure 1. Two hyperbolas fix user's location

the intersection of the two curves (Figure 1).

The two cell sites must be time-synchronized very accurately and reliably for this technique to work. At the speed of light, each nsec of error in timing translates into a foot or more of error in position. As inter-tower synchronization degrades, the TDOA measurements grow inaccurate, the hyperbolas become "fuzzy", and the position error increases proportionately. A GPS Clock at each cell tower easily synchronizes the cell towers to within 100 nsec —100 foot TDOA accuracy — or better.

The GPS Clock

Most users think of GPS as a means of determining position, but the constellation of 24 satellites is also an excellent time keeper. Each satellite contains two Rubidium and two Cesium atomic clocks. These are monitored against atomic clocks on the ground, and the whole system is continuously calibrated against the world wide time standard, Universal Coordinated Time (UTC). Radionavigation signals such as GPS are naturally constructed as a time signal, so it is fairly straightforward to use GPS as our "atomic clock in the sky."

The signal from each satellite is very accurate. The precision of measurement is better than a nsec. Atmospheric modeling errors can account for 50 nsecs of error. By far the largest source of error is Selective Availability (SA). The Department of Defense intentionally degrades GPS accuracy by forcing the GPS signal clock to drift slowly in frequency. With SA, each satellite's signal experiences a timing error of about 100 nsec (1σ) and a frequency error of about one part in 10^8. Without SA, the timing errors would be at best about 10 nsec and the frequency error almost one part in 10^{10}. A multi-channel GPS receiver can average the SA errors over seven or eight satellites if clear view of the sky is available, reducing the effect of SA by almost a factor of three (Figure 2).

Figure 2

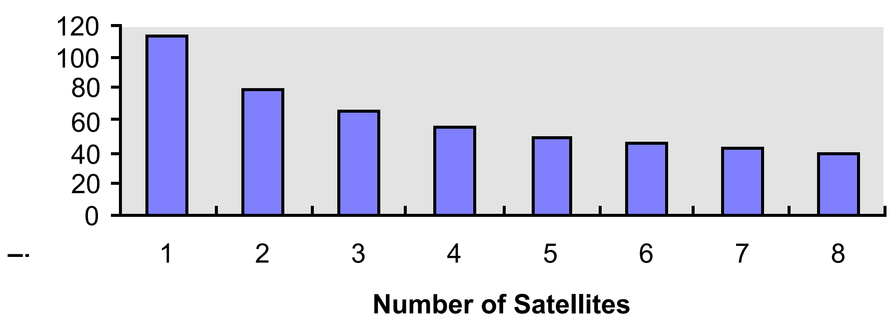

The above accuracy meets the microsecond timing requirements of the paging industry, using a standard GPS receiver with a PPS output. The CDMA “holdover” requirement, on the other hand, is a much more difficult accuracy specification. It requires not only a high degree of timing accuracy, but also a very accurate frequency calibration that minimizes oscillator drift over extended periods of GPS outage. The first generation of GPS Clocks used expensive Rubidium oscillators to accomplish this “holdover” specification. Cost-competitive Rubidium oscillators sell for around $3000 and require regular maintenance and replacement. Since CDMA is battling both TDMA and AMPS for market share, every part that goes into a CDMA base station must be as cost competitive as possible.

With SA, the GPS signal is a poor direct provider of frequency. Even without SA, GPS by itself is not stable enough to provide the reference frequency for CDMA. Atomic oscillators, and even many quartz oscillators, provide a more stable frequency in the short term than GPS. However, all oscillators drift, some more slowly than others. Eventually even a sophisticated cesium oscillator will drift away from UTC. Conversely, GPS is kept on a short leash and, while it wanders back and forth in the short term, it always is calibrated to within a few hundred nanoseconds of UTC.

The GPS Clock exploits GPS’ inherent long term stability and combines it with the good short term stability of a quality quartz (XO) or rubidium (Rb) oscillator. As mentioned before, these oscillators will have a very high quality signal over a short period of time, but will tend to wander over longer periods of time. The wander can be kept in check by gently steering the oscillator with GPS (Figure 3). The average frequency of the oscillator over a certain time interval can be measured by the GPS. The accuracy of this measurement is equal to the GPS timing accuracy (SA, 100 nsec 1σ) divided by the time interval. Over periods of a few minutes, this measurement is so inaccurate that it cannot calibrate the oscillator accurately. Over longer averaging times of perhaps 1000 seconds, calibration accuracy approaches one part in 10^{10} (100 nsec/1000 s). For the calibration to be accurate, the oscillator must be stable, i.e. the frequency offset must be relatively constant over the measurement interval.

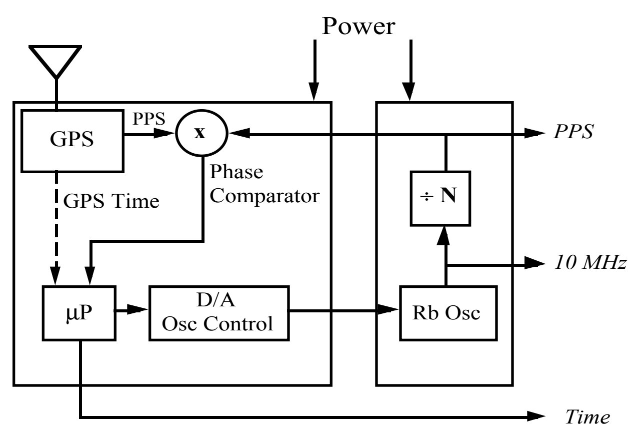

Figure 3. First-Generation GPS Clock

With the GPS receiver keeping long-term wander in check, a relatively inexpensive ovenized oscillator (OCXO) achieves a high level of both short-term and long-term stability. Recently quartz oscillator manufacturers have responded to the challenge of the CDMA holdover requirement by developing ultra-stable double-oven OCXO technology. This oscillators, once it is disciplined by GPS, is able to maintain its frequency for extended periods of time even during a GPS outage. It sells for a fraction of the price of a Rubidium oscillator, has a much wider temperature range, and a significantly longer service life.

This lower-cost oscillator created a second generation of simpler, less expensive GPS Clocks. Bringing down the cost and improving the reliability is essential to the success of CDMA systems in the marketplace. Without an economical way to synchronize towers across the country and the world, CDMA would have remained just an experimental technology.

Like the first-generation GPS Clock, the second-generation clock uses an autonomous GPS receiver that outputs a PPS signal. This PPS is compared with a PPS derived from the output of the OCXO and the difference from the phase comparator is used by a microprocessor to steer the OCXO (Figure 4). This is a competent design, but there is a better way to build this mouse trap.

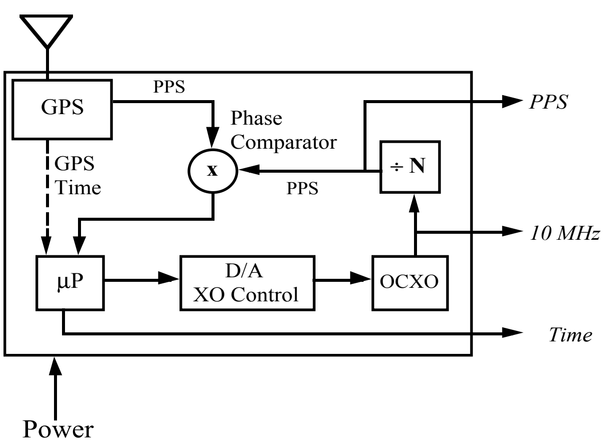

Figure 4. Second-Generation GPS Clock

The third generation of GPS Clocks continues the cost reduction profile. In the third -generation design shown in Figure 5, the OCXO and GPS receiver are now tightly integrated onto a single printed circuit board.

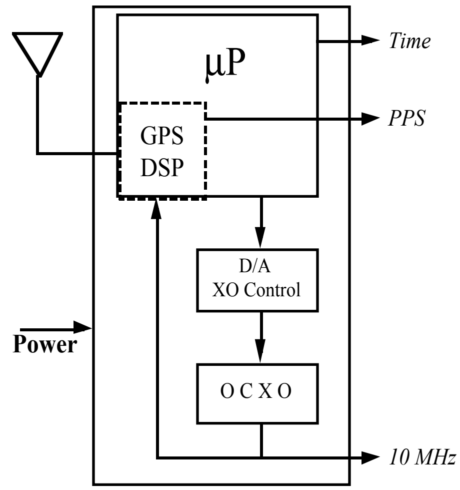

Figure 5. Third-Generation GPS Clock

The 10 MHz frequency from the double-oven quartz oscillator is used directly as the timing source for the GPS' digital signal processor and RF front end chip. Moreover, a single microprocessor runs both the GPS receiver and the clock steering functions. In the first- and second-generation designs, the precision of the PPS output signal (40 to 100 nsec) and put limits on the GPS Clock accuracy. In third-generation designs, the oscillator is directly compared to the GPS signal without using the PPS, allowing accuracy to approach theoretical limits.

The third-generation GPS Clock is also ideal for E911 positioning. There is no holdover requirement as in the CDMA application, so a less expensive single-oven oscillator can be used. System reliability is very important, however. The key benefit of the third-generation GPS Clock is the high-reliability design made possible by its high degree of integration.

In Sync with the Future

Twenty years ago, had one suggested that position accurate to a few meters would be delivered gratis via satellite to mass-produced credit-card-sized radios, it would have been called a believable dream. But unlike most believable dreams, this one has come true. Just like the personal pocket-size radiotelephone.

Although these consumer items are beguiling simple, behind them lies a wireless revolution that is voraciously demanding more. More from its limited slice of the frequency spectrum: more reliability, more clarity, more capacity, more coverage, more features. Time and frequency synchronization enable technologies such as digital CDMA telephony, Enhanced 911 location, and efficient radio-paging to turn these demands into economical realities. And the GPS Clock, the child of that credit-card position radio, is keeping time.