A NEW CLASS OF PRECISION UTC AND FREQUENCY REFERENCE USING IS-95 CDMA BASE STATION TRANSMISSIONS

Bruce M. Penrod
EndRun Technologies

ABSTRACT

A new class of precision timing and frequency reference is introduced that indirectly receives GPS timing and frequency information via TIA/EIA Standard IS-95 Code Division Multiple Access (CDMA) mobile telecommunications base station transmissions. Like cell phones, these products operate indoors without external antennas and provide accuracy, low cost and ease of installation. The technology fits particularly well in IP network synchronization and quality-of-service monitoring applications where rooftop antenna installation is often impossible. The salient characteristics of the IS-95 CDMA signals that make them suitable for this purpose and a general CDMA timing receiver architecture are described. Performance data versus similar references that use conventional GPS reception are also presented.

KEYWORDS

IS-95 CDMA, indirect GPS, UTC, time dissemination, frequency control

INTRODUCTION

This year, products based on a new way of obtaining precision, UTC traceable time and frequency will be on display in the exhibit area of the ITC 2004 meeting. Though the exhibit area might be deeply buried within the hotel in a windowless room, these new products will receive GPS-derived, UTC-traceable time and frequency indoors, using cell phone antennas.

Indirect GPS is arguably the best description for this new technology that has been made possible by the rapidly expanding, global deployment of IS-95 compliant mobile telecommunications systems. Since IS-95 system time is by definition equal to GPS time [1], base stations operating in these systems act as *repeaters* of the GPS timing information they receive from the satellites in order to synchronize themselves. The spread-spectrum modulation scheme employed in IS-95 systems has striking similarities to that of the GPS, and allows the underlying GPS time reference to be extracted from the base station transmissions with a high degree of precision. This new approach eliminates the cost and hassle of installing a rooftop antenna for timing users that are within range of one of these base stations.

Why IS-95 Systems Are Synchronized to GPS Time

Like the GPS, IS-95 is a spread-spectrum system that uses CDMA techniques to differentiate multiple, simultaneous users of the same frequency channel. In the GPS, all of the satellites transmit simultaneously on the same frequency, but each satellite uses a different pseudonoise (PN) spreading code to distinguish its transmissions from those of the other satellites. In the IS-95 system, all base stations transmit on the same frequency and they also transmit the same PN spreading code. *How does this work?*

Each base station transmits the code with a different time delay offset relative to the on-time code. Since the base stations are not moving and the range of their coverage is intentionally very short, this relative time offset can be made large enough that a mobile unit could never be close enough to one base station, and far enough from another, that the codes received from the two base stations would line up and interfere with each other. The typical light speed propagation delay from a base station to the edge of its coverage cell is about 10 microseconds. In an IS-95 system base stations transmit with code offsets that are multiples of about 50 microseconds, allowing them to have a timekeeping error budget on the order of 10 microseconds and still maintain adequate separation from the PN codes transmitted from neighboring cells.

It should be evident that each base station must have a way of knowing when to begin transmitting its replica of the PN code relative to all of the other base stations in the system. There really is only one globally available option for maintaining this level of synchronization. For this reason, the IS-95 standard defines CDMA system time to be GPS time and requires that base stations be synchronized to GPS time to less than 10 microseconds, even during periods of GPS satellite unavailability lasting up to 8 hours [2].

IS-95 systems takes advantage of this inherent level of synchronization to implement an efficient, *soft hand-off* call transfer strategy as the mobile user moves through the cells. This feature is based on the ability of the mobile handset to calculate the correlator offset to the neighboring base stations with sufficient accuracy to allow a much abbreviated search when changing over. This results in many fewer *dropped* calls. It turns out that dropped call statistics kept for each base station provide a key indicator of timing impairments in the CDMA system [3].

How IS-95 Systems Are Synchronized to GPS Time

The timing and frequency signals transmitted by the typical IS-95 base station are sourced from either one or two high performance GPS timing receivers equipped with state-of-the-art disciplined oscillators. Since it cannot be assumed that the typical base station will have a climate control system, the GPS timing receiver(s) at each base station is(are) equipped with either a rubidium vapor local oscillator or an ultra-stable, ovenized quartz local oscillator with software temperature compensation. Such elaborate means are needed to ensure that the base station will maintain the required timing accuracy during GPS signal outages. During such outages, the base station timing must *flywheel* on the free-running stability of the local oscillator.

IS-95 SYSTEM OVERVIEW

Figure 1 depicts the general architecture of an IS-95 CDMA network [4]. One or more base stations are connected via dedicated wire lines to a Mobile Telephone Switching Office (MTSO). The MTSO is responsible for coordinating the base stations that are connected to it and interfacing their voice and data traffic to the Public Switched Telephone Network (PSTN). The MTSO is also responsible for maintaining system time and monitoring the timing status of its connected base stations. This function requires the co-location of at least one GPS time and frequency reference receiver. Each base station must also independently maintain system time and frequency, so a co-located GPS time and frequency receiver is required at those locations as well.

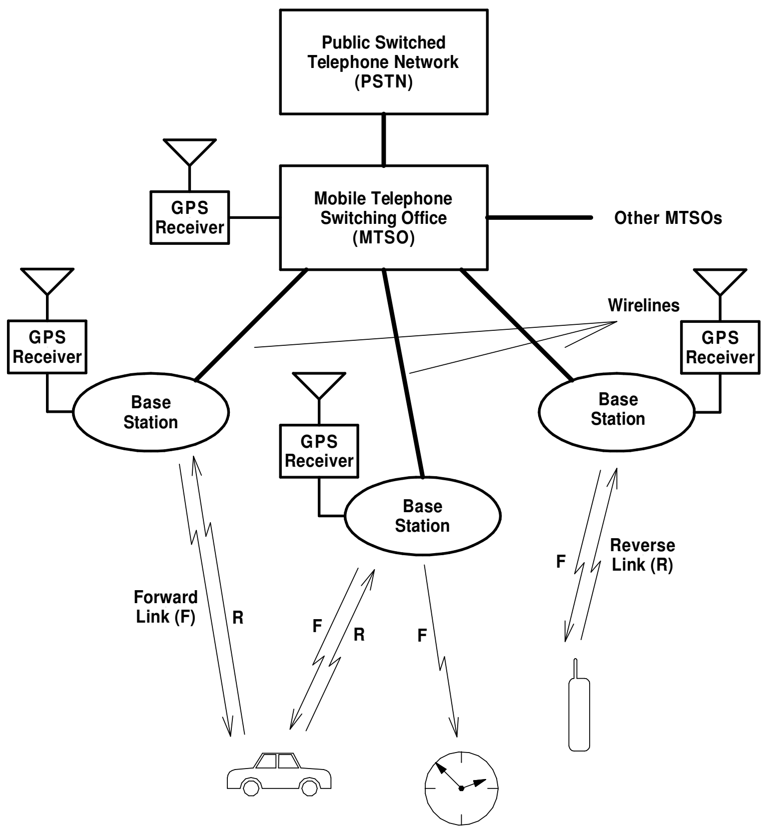

Figure 1 – CDMA Mobile Telecommunications System

Worldwide, IS-95 CDMA cellular systems generally operate in two carrier frequency bands. The original analog system, known as the Advanced Mobile Phone System (AMPS), occupies the 869 to 894 MHz band, commonly referred to as the *cellular* band. The largest CDMA provider using the AMPS cellular band is Verizon Wireless [5]. More recently, the Personal Communications System (PCS) was introduced, operating at 1930 to 1990 MHz—the *PCS* band. U. S. Sprint is the largest CDMA provider using the PCS band.

Due to more rapid signal attenuation at the PCS carrier frequency, PCS offers providers the ability to space the coverage cells more densely in heavily populated urban environments and thereby handle more traffic. However, the *cellular* carrier frequency offers better performance in terms of range and building penetration [4]. For the purpose of transferring time and frequency to stationary users inside of buildings or at the fringes of a coverage area, this is a definite advantage.

In a mobile telecommunications system, transmissions from the base station to the mobile user are on the *forward link*. Transmissions from the mobile user to the base station are on the *reverse link*. As with the GPS, time and frequency transfer via the IS-95 CDMA technology is passive, or receive-only. The precise time and frequency information is contained in the transmissions on the forward link, and it is not necessary to handshake with, or be a subscriber in order to just *listen* to these transmissions. Consequently, the reverse link will not be discussed here.

Forward Link Structure

Mobile telephones when first energized must receive and decode the system timing information on the forward link in order to initiate or receive calls. The forward link is composed of multiple, simultaneous channels. Multiplexing is by a set of orthogonal *covering codes* that bi-phase modulate the data stream in each channel priot to PN code spreading. These covering codes are chosen from a set of order 64 Walsh functions. The channels consist of the pilot channel, sync channel, up to 7 paging channels and up to 55 traffic channels.

Rejection of narrowband signal interference as well as neighboring base station interference is via direct-sequence, PN code bi-phase modulation. The Walsh-covered channel data is spread with two different PN codes designated I and Q. The resulting pair of bit streams are used to modulate the I (cosine) and Q (sine) radio frequency (RF) carriers to generate a form of quaternary phase-shift keying (QPSK). The PN codes are 2^{15} chips long, and the chipping rate is 1.2288 Mcps, so the code repeats every 26.666… milliseconds. Isolation between base stations is accomplished by assigning each a PN code offset that is a multiple of 64 chips, which equates to multiples of 52.0833… microseconds. 512 such offsets span the full length of the PN code.

For time and frequency transfer, only two of the channels must be demodulated. Figure 2 shows the details of the generation of the pilot and sync channels at the base station. The pilot channel is broadcast at the highest power level of the forward link channels and is unique in that its data is all zeros. Since the pilot data is constant, the receiver's initial PN correlation search may use longer integration times to bring the signal up out of the noise. After pilot phase lock has been accomplished, it provides the reference phase for coherent demodulation of the remaining data channels.

Sync channel data is a fixed-length *message* sent at 1200 bits per second (bps) that the receiver uses to establish GPS, UTC and local time. Prior to transmission it is convolutionally encoded using a rate ½, constraint length 9 encoder, effectively creating two *symbols* for each bit of data to allow error correction at the receiver. Each symbol is then repeated once, and a *block* of 128 of these is *interleaved* to provide temporal diversity for burst error mitigation. This block of symbols is transmitted at 4800 bps as a *frame* with each frame being aligned with, and having the duration of, one repetition of the PN code. Three contiguous frames compose a *superframe*.

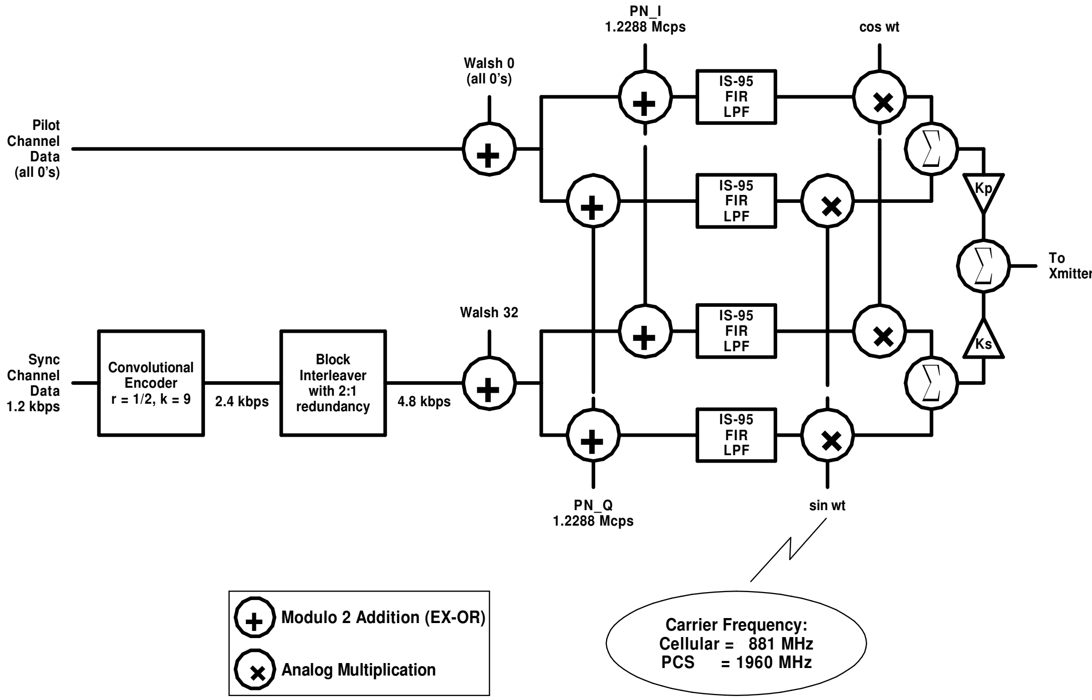

Figure 2 – Pilot and Sync Channel Signal Structure

The time-of-day information contained in the message is valid four superframes after the end of the last superframe containing the message. Since each base station transmits its PN code with an offset delay relative to the zero offset PN code, the PN code offset of the specific base station is also contained in the sync channel message so that this fixed delay can be corrected.

The waveforms that compose the transmitted forward link signal are coherent with the co-located GPS time and frequency reference. The PN codes, Walsh functions, frames and superframes all began synchronously with second 0 of GPS time, and are again aligned with every even GPS second thereafter, except that superframes are coincident only every third even second [1]. During normal operation, this even second alignment to GPS time is at the 1 microsecond level [2]. Although the IS-95 standard does not require base station RF carrier frequency control to better than 5×10^{-8} [2], in practice they are precisely phase locked to the GPS reference frequency.

IS-95 CDMA SIGNAL RECEPTION

Figure 3 depicts the general architecture of a receiver specifically designed to demodulate and recover the underlying precision time and frequency signals present in the forward link transmissions of an IS-95 CDMA network.

Front End

The desired signal band is pre-selected using a surface acoustic wave (SAW) filter appropriate for the particular system carrier frequency. Following pre-selection filtering, the signal is amplified with automatic gain control due to the wide dynamic range of the received signals. The signal is then downconverted in quadrature to a low intermediate frequency (IF) suitable for digitizing with high speed analog-to-digital converters. Prior to digitizing, the I and Q signals are bandpass filtered to match the IS-95 bandwidth of 1.25 MHz.

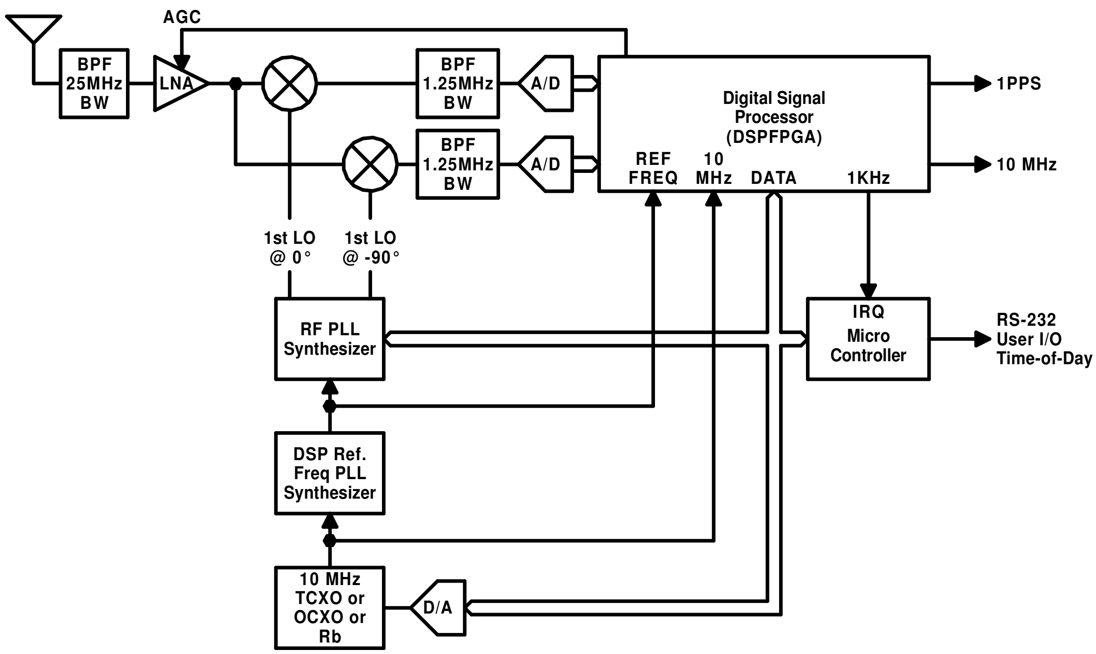

Figure 3 – CDMA Time and Frequency Receiver

Frequency Plan

The receiver operates with a 10 MHz reference oscillator. Signals required for IS-95 demodulation are synthesized from this reference. Since the clocking frequencies needed for digital signal processing (DSP) of the IS-95 signal are not direct divisions of 10 MHz, a DSP reference frequency is synthesized via a separate crystal oscillator phase locked loop (PLL). This DSP reference oscillator frequency is divided as needed in the DSP field programmable gate array (DSPFPGA) to perform the final downconversion and baseband processing. It also provides the reference for the 1st local oscillator (LO) RF PLL synthesizer.

Digital Signal Processing

The DSPFPGA performs high speed multiply-and-accumulate (MAC) arithmetic for the final conversion to baseband of the digitized I and Q signals. It synthesizes I and Q replica waveforms by modulo-2 additions of the various divisions of the DSP reference frequency, and generates the I and Q PN codes via linear feedback shift registers (LFSR). The I and Q PN codes then spread the replica waveforms, which are then multiplied with the received signal and accumulated for some integer number of symbol intervals. These accumulations are then passed to the microcontroller for interpretation.

The high-performance microcontroller receives the raw baseband data from the DSPFPGA and controls the stepping of the phase of the DSPFPGA correlators to implement the pilot channel

PN code search. After pilot detection, the early-late correlator powers are used to phase-step the correlator so as to maintain PN code phase lock. While PN code locking, the I and Q correlator sums are used to phase adjust the DSPFPGA-generated 2nd LO so as to align it with the pilot carrier phase. These code and carrier phase corrections are used to implement an outer PLL using the digital-to-analog converter to phase lock the 10 MHz reference oscillator to the received signal. This approach optimizes the short term stability of the time and frequency outputs of the receiver by taking advantage of the higher resolution and lower noise characteristics of the carrier phase measurements to smooth the code phase measurements.

Sync Channel Data Decode and Timekeeping

Once pilot code and carrier phase lock has been attained, the sync channel correlator sums contain symbol data. The microcontroller *de-interleaves* these and processes them using a Viterbi decoder to recover the transmitted sync channel data. Real time is maintained in the microcontroller using the DSPFPGA-generated 1 kHz interrupt that is coherent with the 10 MHz reference oscillator. The location of the superframe boundaries relative to the 1 kHz interrupt is measured in the DSPFPGA. After successful data decode, as validated by the CRC, the microcontroller commands the DSPFPGA to synchronize the 1 kHz interrupt to the appropriate superframe boundary, advanced by the decoded base station PN offset. During the millisecond prior to the next second, the DSPFPGA is commanded to output the 1PPS synchronous with the next millisecond. After the intial synchronization of the 1PPS, all phase alignment is maintained via frequency control of the 10 MHz oscillator, so that the outputs have very low phase jitter.

PERFORMANCE CHARACTERISTICS

Spread spectrum systems are ideal for high precision time transfer, so the performance of the described CDMA receiver in terms of its precision and short term stability is very similar to that of a well designed GPS receiver. Since the CDMA time scale is in fact phase locked to the GPS time scale, the long term frequency accuracy and stability of the technology is essentially equal to that of GPS. Its limitations are in its absolute time accuracy and medium term frequency stability. The limitations are due to several factors which are discussed in decreasing order of magnitude.

Propagation Delay

Depending upon the density of the cells in an area, which is generally a function of the expected number of simultaneous users and/or the density of signal-obstructing buildings, the time transfer uncertainty due to propagation delay ranges from less than 5 microseconds (urban, dense) to as much as 25 microseconds (suburban, sparse). Since it is always present and has the largest potential variation, this is the dominant factor in limiting the absolute time accuracy. Since both the base station and the receiver are stationary, it has no affect on the stability or absolute accuracy of the frequency that is transferred.

IS-95 Timing System Requirements

The system is not intended to disseminate time to a level of accuracy beyond the stated requirements of the IS-95 standard, or 10 microseconds. Though the equipment that is being used to maintain that level of accuracy is normally much better than that [2], service providers have no responsibility or motivation to ensure any higher level of accuracy continuously. Although an alarm is generated immediately when a base station GPS receiver is operating in flywheel mode, it does not necessarily cause the base station to be turned off, and the system does not provide an indication of the GPS timing status in the forward link data.

A visit to the base station by a technician might not occur until the base station experiences a statistically significant upturn in its dropped calls [3]. Service providers are motivated to maintain an acceptable level of service, and dropped calls represent loss of revenue. The 10 microsecond timing sub-system error budget allows flywheel operation to continue without interruption of service while the severity of the problem is being evaluated.

What this means to time and frequency dissemination over the IS-95 airwaves is that a very small, but not negligible, probability exists that the time being transferred is from a high stability oscillator that has been free-running for some fraction of a day. This is the dominant factor limiting the medium term stability of the realized time and frequency transfer.

Base Station Switching

From time to time, base stations are taken down for various reasons. When this occurs, the signals from neighboring base stations become receivable and the receiver will begin to track the strongest one of those. Due to the change in propagation delay from the new base station, this switchover can cause a step transition in the timing output. Depending upon the drift rate of the local oscillator in the CDMA receiver, a transient in the frequency output can occur while it is being brought back on frequency. The frequency transient can be controlled by using a higher stability local oscillator, just as is done in conventional direct GPS receivers.

Multipath

Due to its specific utility as an indoor antenna alternative to rooftop antenna direct GPS, multipath effects are significantly greater than they are with rooftop mounted GPS antennas that have clear visibility to the horizon. Rearrangement of the furniture or relatively small movements of the antenna can, by altering the strength of the direct signal relative to a long delayed reflection, cause larger timing variations than might be expected. In *urban canyon* installations of direct GPS, where only one or two satellites may be visible and there are multiple signal paths due to reflections off of taller adjacent buildings, similar instabilities are seen as the satellites move across the sky.

Due to the correlation properties of the CDMA PN codes, timing shifts due to variations in the multipath reflections inside of a building are bounded by the width of a PN code chip to about 800 nanoseconds. In most situations, multipath induced instabilities are secondary in magnitude to those that could arise from GPS flywheeling intervals at the local base station, but multipath phenomena will be experienced more often and by more users.

PERFORMANCE DATA

Figures 4-6 represent data gathered during November 14 – 24 of 2001 at the EndRun Technologies facility in downtown Santa Rosa, CA. The reference for the time interval measurements was a high performance GPS timing receiver disciplining a rubidium local oscillator, with rooftop mounted antenna, tracking a full constellation of satellites. The phase residuals of the digital phase lock loop controlling the rubidium oscillator were continuously monitored during the data logging. Figure 4 shows these residuals plotted along with ambient temperature. With SA being off, the phase error was maintained at less than 10 nanoseconds, with the measurements exhibiting a peak time deviation (TDEV) of less than 3 nanoseconds.

The cellular band CDMA receiver was equipped with a miniature, dual-inline packaged (DIP) oven controlled crystal oscillator (OCXO). This OCXO exhibits a temperature stability of about 2×10^{-9}/°C and a short term stability of about 2×10^{-10} at one second. The antenna was a ¼ wave monopole with magnetic base attached upside down to the metal framework that supports the acoustic tiles in the suspended ceiling. This is a recommended configuration as it places the antenna above most obstructions inside of a typical room and provides some ground plane.

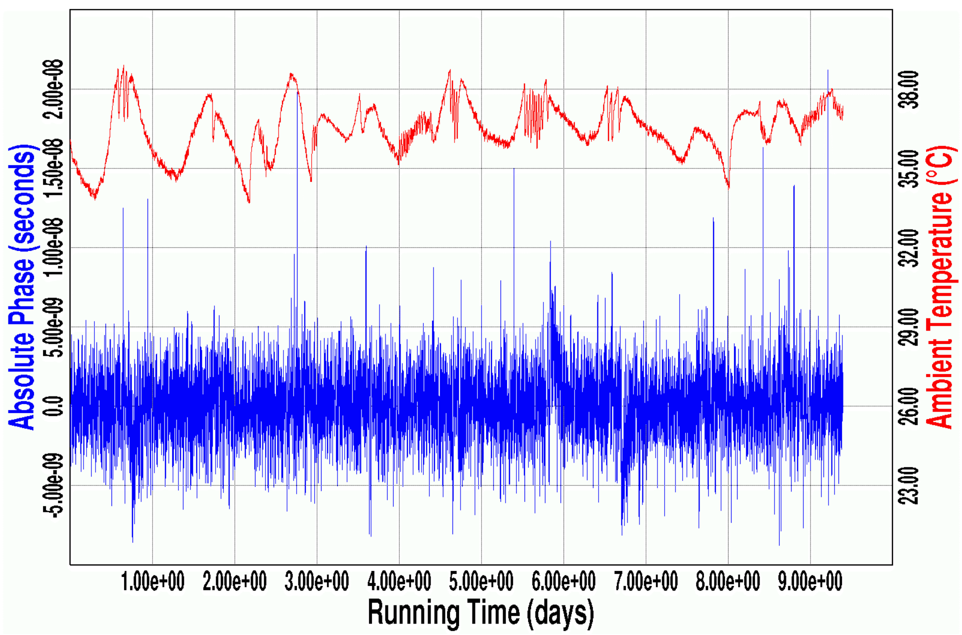

Figure 4 – GPS Disciplined Rubidium Residuals Nov. 14-24, 2001

Figure 5 shows the time interval measurements of the CDMA 1PPS versus the direct GPS 1PPS plotted along with the ambient temperature. Data was collected at 10 second intervals and the plot was smoothed over ten samples. The propagation delay is about 2.5 microseconds. The peak-to-peak phase "wander" of about 600 nanoseconds, not diurnal in nature, seems to be typical of this base station, the timing sub-system of which is believed to have been supplied by Lucent [5].

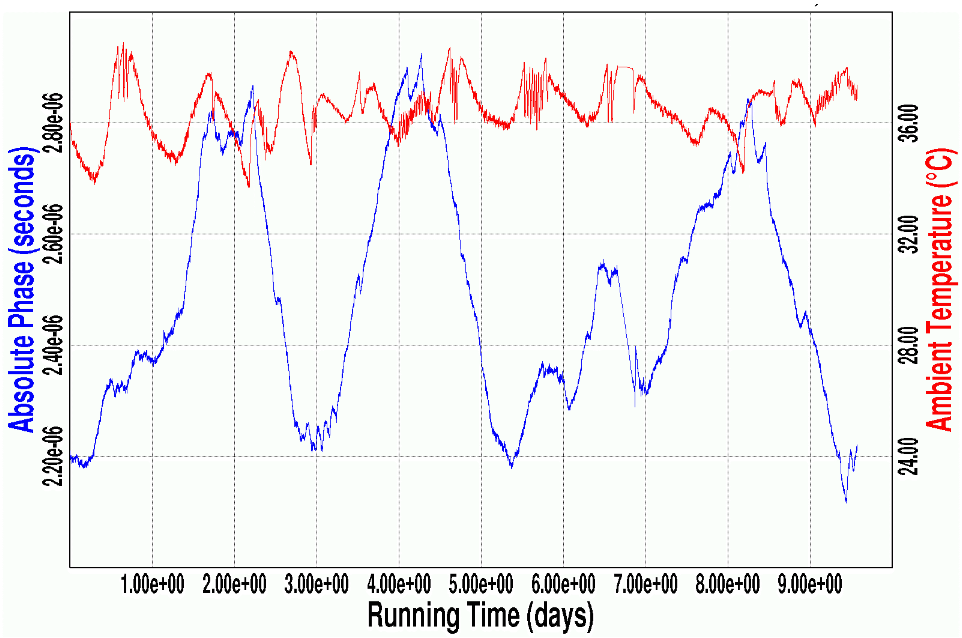

Figure 5 - CDMA 1PPS vs GPS 1PPS Phase Nov. 14-24, 2001

Figure 6 shows the Allan deviation of these time interval measurements, again with 10 sample smoothing applied. The short term stability is similar to that of a conventional direct GPS receiver with an equivalent oscillator, while the longer term stability is degraded relative to direct GPS due to the characteristics of this particular base station. It is not known if this performance is typical of Lucent supplied equipment.

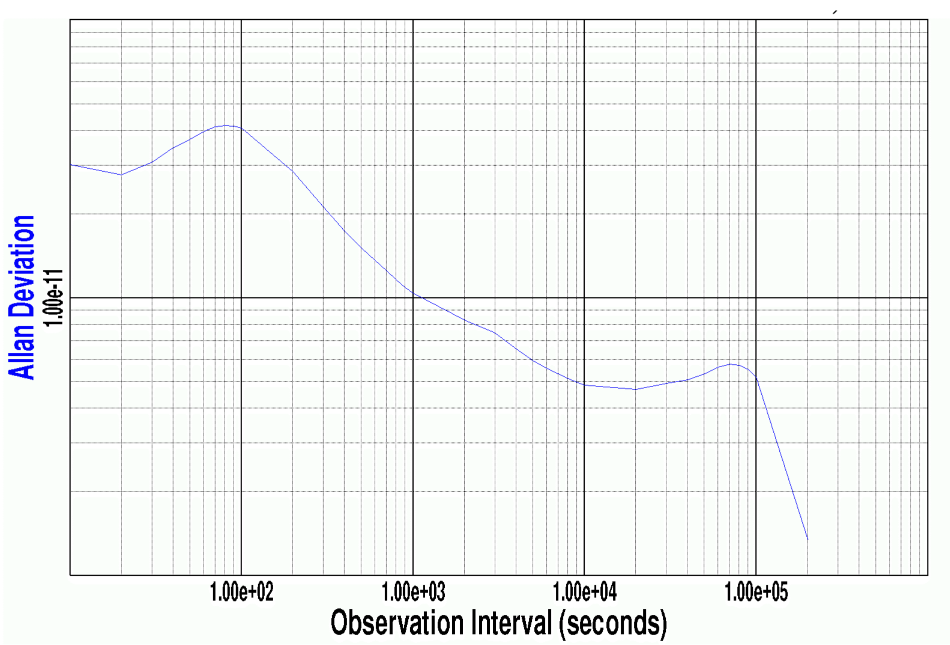

Figure 6 – CDMA 1PPS vs GPS 1PPS Allan Deviation Nov. 14-24, 2001

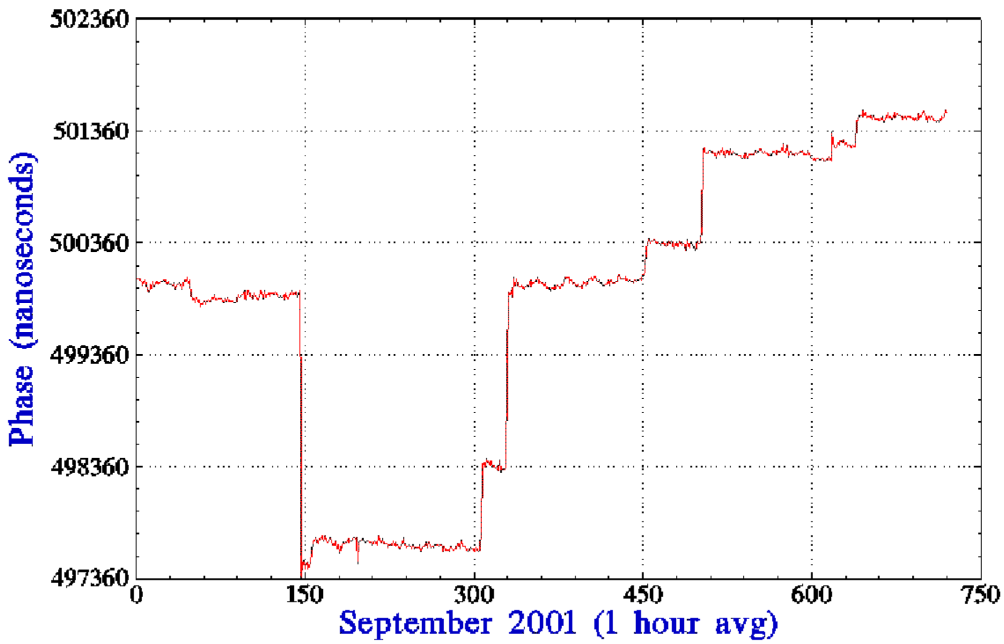

Figure 7 – CDMA Receiver TCXO 10 MHz Phase vs UTC(NIST) September 2001

Figure 7 was provided by Michael Lombardi of the National Institute of Standards and Technology (NIST) in Boulder, CO [6]. It shows data collected over the month of September 2001 on one of these CDMA receivers equipped with the standard temperature compensated crystal oscillator (TCXO). The 10 MHz output was divided down to 1PPS and compared against the UTC (NIST) timescale. As such it provides information only about the frequency accuracy and stability of the received CDMA signal.

Status logging was not performed during the logging period, so the exact cause of the phase steps is not known. They could be either from base station switching or short signal outages, during which the 10 MHz TCXO frequency would accumulate phase error fairly quickly. Since, the CDMA receiver algorithms currently do not attempt to maintain strict coherence between the phase of the 10 MHz output and the 1PPS output following periods of signal outage, these accumulated phase errors would persist exactly as shown by Figure 7. Phase shifts of the magnitudes shown in Figure 7 would not be seen with higher stability disciplined oscillators.

Of particular interest in this data is the much better phase stability during the long, continuous tracking intervals. The deviation appears to be diurnal with about 100 nanoseconds of peak-to-peak movement. It is believed that the timing sub-system for this base station was provided by Motorola [5]. It is not known if this performance is typical of Motorola-supplied GPS timing sub-systems.

CONCLUSIONS

A new technology has been introduced that offers a promising alternative for terrestrial based time and frequency dissemination using modern wireless techniques. CDMA mobile telecommunications background information, including system standards and signal characteristics, and a timing and frequency optimized CDMA receiver architecture have been presented. Performance limitations have been discussed, and long term data taken by two independent sources in two geographically remote locations has been presented. These indicate the reliability, and performance capabilities of the new technology. The economic and practical advantages of indoor operation are gained with a marginal tradeoff in absolute timing accuracy relative to conventional, direct GPS.

REFERENCES

[1] "Mobile Station-Base Station Compatibility Standard for Wideband Spread Spectrum Cellular Systems," TIA/EIA Standard 95-B, Arlington, VA: Telecommunications Industry Association, March 1999.

[2] "Recommended Minimum Performance Standards for Base Stations Supporting Dual Mode Spread Spectrum Systems," TIA/EIA Standard 97-D, Arlington, VA: Telecommunications Industry Association, June 2001.

[3] James Cigler, Sr. RF Engineer – System Performance, Verizon Wireless, teleconversations during November 2001

[4] J. S. Lee, L. E. Miller, *CDMA Systems Engineering Handbook,* Artech House Publishers, Boston, 1998, ISBN 0-89006-990-5

[5] CDMA Development Group, CDMA Regional Worldwide Database, web document *http://www.cdg.org/world/region.asp?region=North+America*, November 2001

[6] Michael Lombardi, NIST Time & Frequency Division, e-mailed phase data plotted during November 2001