GNSS Solutions:

Signal acquisition and search, and antenna polarization

"GNSS Solutions" is a regular column featuring questions and answers about technical aspects of GNSS. Readers are invited to send their questions to the columnists, **Professor Gérard Lachapelle** and **Dr. Mark Petovello**, Department of Geomatics Engineering, University of Calgary, who will find experts to answer them. Their e-mail addresses can be found with their biographies at the conclusion of the column.

How does signal acquisition with batch processing (i.e., FFTs) work? How is it different than traditional serial search techniques?

In order to answer these questions, let us briefly review the serial-search acquisition process. The core architecture of a serial-search receiver is depicted in **Figure 1**.

Numerically controlled oscillators (NCOs) drive the so-called "local" code and carrier generators (recall that in order to track the signal, the receiver must generate local copies of the code and carrier and must synchronize them with the received signal). The locally generated code and carrier signals are multiplied with the received signal (which itself has been amplified, filtered, down-converted in frequency and digitized with an analog-to-digital or A/D converter), and the results are then accumulated (discrete-time equivalent of integration).

In order to account for the as yet unknown phase of the received signal, the locally generated carrier is broken into two components, which are phase-shifted by 90 degrees with respect to each other. These two orthogonal signals, after multiplication with the received signal and local code, and subsequent accumulation, are referred to as the in-phase (I) and quadrature or quadraphase (Q) components. The code-lock detector then is given simply by the sum of the squares of the I and Q components:

$$Code\_Lock\_Detector = I^2 + Q^2$$

In order to detect a satellite, the locally generated code must first be the same PRN code as the one the satellite is broadcasting and must also be aligned within half a PRN code bit (known as a "chip") of the received signal. Furthermore, for relatively strong signals such as those received in open-sky environments, the locally generated carrier frequency must be generated within about 1 kHz of the down-converted carrier of the received signal.

Consequently, the serial search process involves picking a PRN code and then searching all possible combinations of local code offset and carrier frequency. In the GPS C/A code, for example, the PRN code is 1023 chips long. In order to get within half a chip of the received signal, the receiver must search 2046 half-chips. Frequency uncertainty is on the order of +/- 20 kHz for high dynamic applications and/or for receivers with poor oscillators and is searched in 1 kHz bins, thus yielding 40 frequency bins that must be searched.

In practice the receiver will pick a frequency setting for its local car-

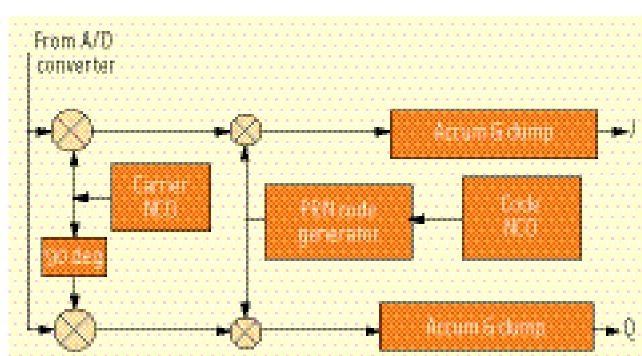

FIGURE 1 Serial-Search Receiver Architecture

rier and then search through all possible code offsets. If no satellite signal is detected, the frequency setting is changed and the process is repeated. This continues until either a) the satellite signal is acquired, or b) all possible combinations of local code and carrier (known as the search grid) have been searched.

If the entire search grid yields no satellite, then another PRN code is selected and the entire process starts over. An example of code-lock detector output for an entire search grid (for a visible satellite) is given in **Figure 2**. As one can see, the detector output rises above the noise only when the locally generated code and carrier are close to that of the received signal.

(Note: The example plots given here show the results of processing GPS C/A-code signals that have been bandlimited to 2 MHz in the receiver front-end. The final intermediate frequency was 1.27 MHz and the sampling rate was 5 MHz. The figures here show only a portion, +/–5 kHz, of the frequency search space. Also note that with a sampling rate of 5 kHz, there are 5000 samples per C/A-code.)

Although this so-called serial search process works well, it is fairly slow. For the GPS C/A code, the minimum accumulation interval (also known as the predetection integration interval) is one whole C/A code or, equivalently, one millisecond. Given 40 bins of frequency uncertainty and 2,046 half-chips of code uncertainty, there are thus 40×2,046 = 81,840 code/frequency combinations in the search grid. Even though each combination takes only 1 millisecond to process, searching the entire grid would take well more than one minute.

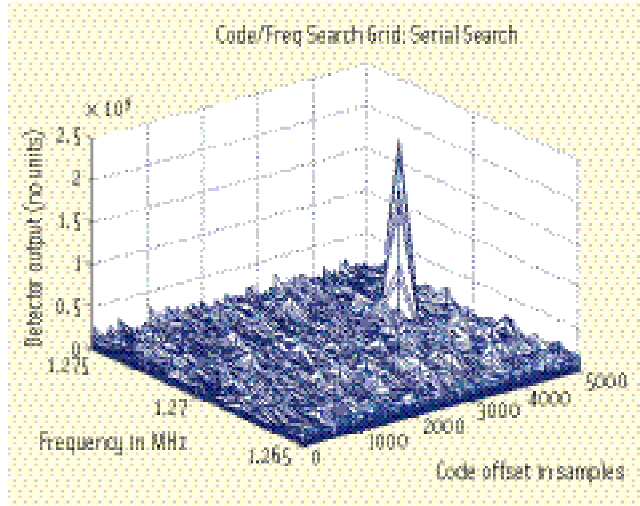

FIGURE 2 Sample Code-Lock Detector Output for a Single Visible Satellite

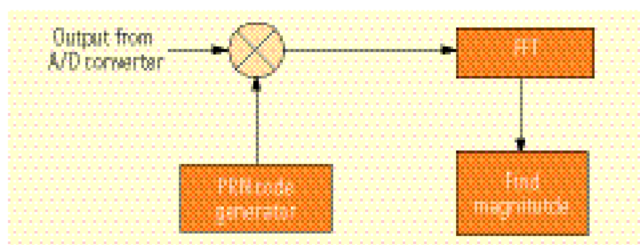

FIGURE 3 Parallel Frequency Search Receiver Architecture

Almost since the beginning of GNSS itself, researchers have been looking for ways to speed up the acquisition process. If multiple channels are available (i.e., multiple implementations of Figure 1), the search grid can be divided up among them and, in the extreme case, application-specific integrated circuit (ASIC) chips have been

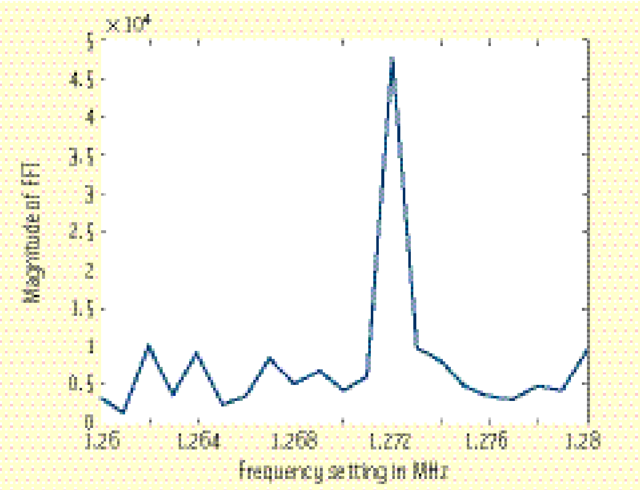

FIGURE 4 **Doppler Frequency Determination using Parallel Frequency Search Approach**

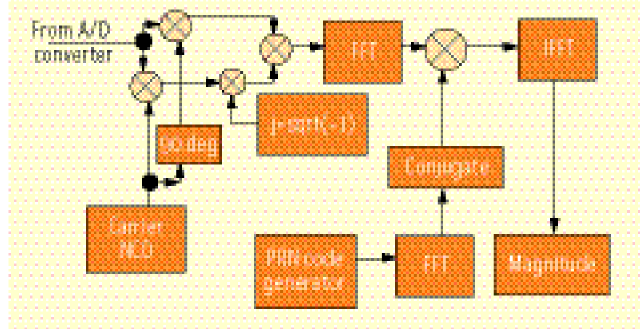

FIGURE 5 **Parallel Code Search Receiver Architecture**

developed to search several thousand grid spaces simultaneously.

(We should mention, however, that these chips have been developed for the encrypted military signals that have at least an order of magnitude greater search grid size). All of these techniques, however, are still restricted to time-domain processing. That is, the accumulation/integration is carried out over time and must be done once for each search grid space.

The motivation behind the modern fast Fourier transform or FFT-based search techniques was the desire to parallelize the search space by exploiting frequency-domain processing. The acquisition process could be accelerated if one could search, say, all possible code offsets for a given frequency setting or, vice versa, all possible frequency settings for a given code offset.

To see how this works, first consider the *parallel frequency search* architecture given in **Figure 3**. The theory behind this technique is amazingly straightforward. The input signal consists of the PRN code, Doppler-shifted carrier and noise. If the locally generated PRN code aligns with the code in the received signal, then they cancel each other out in the multiplication process and all that is left is the Doppler-shifted carrier plus noise. The frequency of this signal can then be determined by computing the FFT and looking for a spike in the spectrum. An example of this is shown in **Figure 4**, which was obtained using the same data as in Figure 2.

Of course, if the locally generated PRN code is not aligned with the received signal, then the FFT output will only show noise and no discernable peak in the spectrum. Thus, in practice the search requires the receiver to generate all possible PRN code offsets (2,046 for the case of the GPS C/A code).

However, for each code offset, all frequency bins are searched simultaneously and thus the search grid is reduced from 81,840 bins down to 2,046. This obviously can be performed significantly faster than the serial search so long as the computational burden associated with the FFTs can be accommodated.

The other FFT-based search technique is the *parallel code search* architecture shown in **Figure 5**. To understand how it works, we need to review something we probably hated in school: convolution. We recall that the convolution integral was quite unpleasant to evaluate. However, one may also recall that convolution in the time-domain corresponded to multiplication (an easy task) in the frequency-domain.

In normal GNSS acquisition and tracking we do not utilize convolution, but we do utilize a closely related process known as correlation. Although we did not refer to it that way in Figure 1, multiplying the received signal with the locally generated signal and then integrating the result is known as correlation. If we could perform this via multiplication in the frequency domain, then all possible code offsets could be searched simultaneously.

To see how this is done, first consider the convolution of two periodic signals (with period T):

$$x(t) * v(t) \overset{\Delta}{=} \int_0^T x(\lambda) v(t-\lambda)\, d\lambda$$

Correlation of these two signals is very similar:

$$R(\tau) \overset{\Delta}{=} \int_0^T x(t) v(t+\tau) dt$$

The key difference between the two operations is that one of the signals (v in the equations here) is time-reversed in the convolution integral, whereas it is not in the correlation integral. If we compute the Fourier Transform of the convolution integral, we get (as stated earlier) multiplication:

$$F\{x(t) * v(t)\} = X(\omega) \cdot V(\omega)$$

We can now extend this by recognizing that convolution is equivalent to correlation when one of the signals has been time reversed. From Fourier theory we know the Fourier Transform of a time-reversed signal is given by the complex conjugate of its frequency domain representation:

$$F\{v(-t)\} = V^*(\omega)$$

We can thus see the Fourier Transform of the correlation process is given by:

$$F\left\{\int_0^T x(t)v(t+\tau)dt\right\} = F\{x(t)*v(-t)\} = X(\omega)\cdot V^*(\omega)$$

This process is implemented in the right half of the diagram in Figure 5. Although it appears complicated, the left half of the diagram simply corresponds to multiplication of the received signal by a complex exponential at the generated carrier frequency, ω:

$$e^{j\omega t} = \cos\omega t + j\sin\omega t$$

The search process involves setting a local carrier frequency and multiplying the local carrier with the received signal and taking the FFT of the result. This is multiplied with the complex conjugate of the FFT of the locally generated PRN code. The result is inverse FFT'd. The magnitude of the inversed FFT (IFFT) is then equivalent to the code lock detector shown previously.

If the local carrier frequency is close to that of the received signal, then the IFFT will have a spike at the code offset (the offset between the local and received code). An example is given in **Figure 6**, computed from the same data set as above. Thus, for a given local frequency setting, all possible code offsets are searched simultaneously.

The receiver must still repeat this process for all possible frequency bins but the search grid size has been reduced from 81,840 down to 40. This allows for a tremendous increase in search speed over the traditional serial search — again, assuming the computational burden associated with the FFTs and IFFTs can be accommodated.

Before leaving this particular technique, we should also note the IFFT output is a measure of the entire PRN code

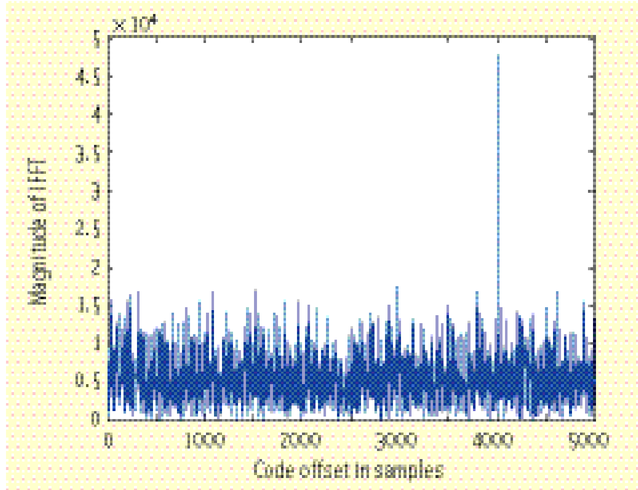

FIGURE 6 Code Offset Determination using Parallel Code Search Approach

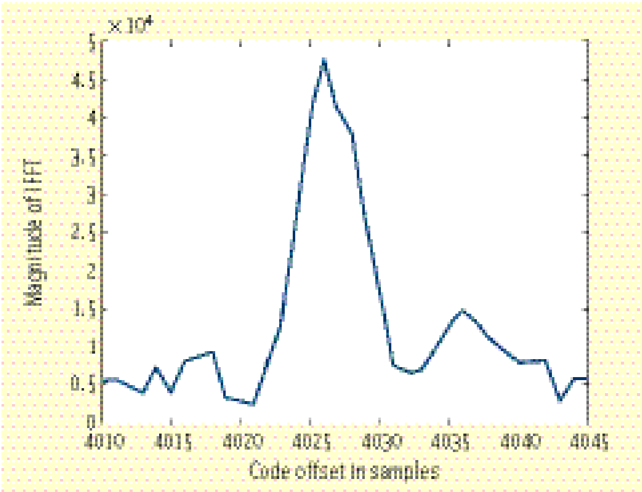

FIGURE 7 **Detail of Correlation Peak Obtained using Parallel Processing**

correlation function and thus can be used for signal quality monitoring purposes. As an example, consider the zoomed-in view of the correlation peak (from Figure 6) given in **Figure 7**. The right side of the correlation peak is clearly deformed from the ideal because the peak is supposed to be symmetrical.

In this particular case the deformation is likely due to multipath, but further consideration is beyond the scope of the topic being considered. The point is that a batch search technique yields insight into the signal that was not possible with the traditional serial search.

To be fair, parallel search techniques have some negative characteristics. We have already mentioned the increased computational power required (compared to the serial search). In addition, delays are associated with the accumulation of the samples along with the computation of the FFTs and IFFTs.

The delay associated with the accumulation of the samples is not of concern, however, because this is the same delay as would be experienced during a normal predetection integration interval in traditional processing. The computational delay, however, is an additional delay not encountered in traditional processing, and it complicates the hand-off from the search algorithm to the tracking algorithm. Given the speed of modern digital signal processing and field programmable gate array chips, however, the delay can be reduced to a manageable value.

When first proposed more than 15 years ago, these parallel search techniques required computational power that did not exist. Now, however, these techniques have been implemented in real-time and have found their way into commercial products. Research is now being conducted to extend the benefits of frequency-domain processing to tracking, measurement generation, signal quality monitoring, and more.

MICHAEL S. BRAASCH, PH.D., P.E.

Michael S. Braasch, Ph.D., P.E., is the Thomas Professor of Engineering in the School of Electrical Engineering and Computer Science at Ohio University. He is also a principal investigator with the Avionics Engineering Center also at Ohio University.

Why do GNSS Systems Use Circular Polarization Antennas?

GNSS systems generally use *right hand circular polarization* (RHCP) antennas. Why not vertical polarization or horizontal polarization? After all, linear polarized antennas are easier to build.

Perhaps the most basic motivation is that circular polarizations are more tolerant of physical orientation mismatches. As an example: if we transmit using a vertically oriented whip antenna (vertical polarization) and use a similar antenna for reception, we would see a strong signal if the antenna is vertically oriented. But take the receive antenna and lay it on its side (90 degree rotation), and the signal strength will go down by 20 decibels or more, because we are now trying to receive a vertically polarized signal with a horizontally polarized antenna.

A circular polarized antenna can be built using two antennas: a horizontal polarization antenna and a vertical polarization antenna with the outputs from each added together with a 90 degree phase shift. The relative phasing of the addition determines whether we have constructed an RHCP antenna or a *left hand circular polarization* (LHCP) antenna.

Repeating the physical rotation experiment but with two RHCP antennas pointed at each other, the signal strength stays the same independent of rotation angle. The only effect seen is that the carrier phase advances 1 degree for each 1 degree of physical rotation. This is the so-called *phase windup* or *phase wrap-up* effect. The effect is frequency-independent, and we would see the same carrier phase advance on both L1 and L2 signals in units of cycles.

In a GNSS system, the advantages of circular polarization are clear; the user doesn't have to worry about the orientation of his antenna other than to make sure it is pointed in the general direction of the satellites. With a linear polarization; the user's receive polarization would be dependent on his/her direction of travel and performance would be unacceptable.

As a further motivation, the ionosphere also affects signal polarization. Free electrons in the ionosphere rotate the signal's polarization. Called Faraday rotation, this rotation is inversely proportional to frequency squared; so, the rotation at L2 (1227.6 MHz) is 1.65 times larger than that seen at L1 (1575.42 MHz) for the same total electron count (TEC). The effect yields up to about 90 degrees of rotation at L2 and 50 degrees of rotation at L1.

Using linear polarizations, receivers could see significant signal attenuations even if the physical orientations are correct. With circular polarization, signal strength is constant, but we would see a frequency-dependent phase windup effect due to Faraday rotation.

Circular polarization is also advantageous in fighting the effects of multipath. When an RHCP signal reflects off a surface at a large angle of incidence, it tends to reverse its polarization sense to LHCP. A good RHCP antenna will attenuate LHCP signals by 10 to 20 dB relative to an RHCP signal. For smaller angles of incidence below about 30 degrees, the reflected multipath signal will tend towards linear or RHCP and, so, the receive antenna's multipath rejection will not be as strong (See the Editor's note for a reference on multipath polarization).

Expanding on this notion, several researchers are working on active antenna designs that exploit polarization differences between direct-path and multipath signals. The basic idea is to strongly attenuate multipath by developing precise opposite sense polarization in the direction of multipath sources whilst passing the direct-path signal which has different polarization. A key problem is seeing the multipath signals as they usually have lower signal-to-noise ratio than the direct path signals.

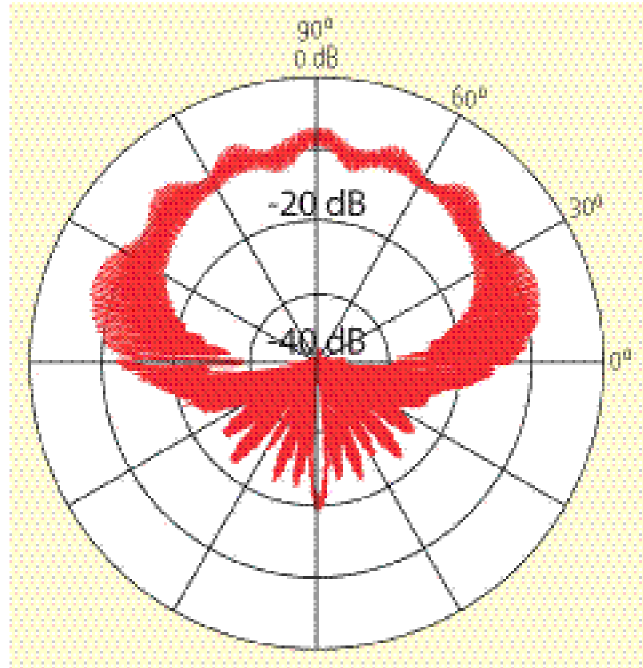

FIGURE 1 Patch Antenna Elevation Gain Pattern (-1 dBiC Peak)

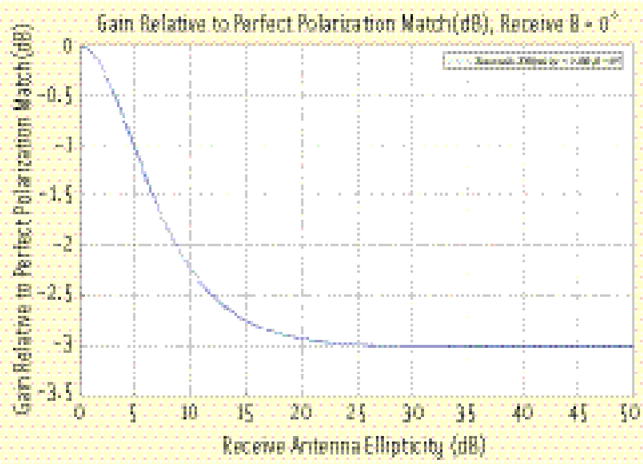

FIGURE 2 **Losses Due To Ellipticity Assuming A Perfect Satellite Antenna**

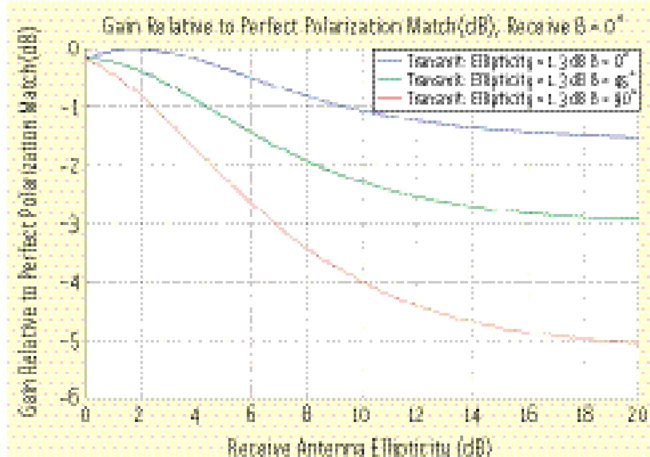

FIGURE 3 **Polarization Losses Accounting for Satellite Antenna Imperfections**

An ideal RHCP antenna would have the same response to any linear polarization stimulus; namely -3 dB relative to its response to an RHCP stimulus. In practice, the antenna gain in two orthogonal polarizations (e.g., horizontal & vertical) is not perfectly balanced. As a result, the antenna will actually have elliptical polarization.

Referring to **Figure 1**, RHCP antenna performance is often characterized using a spinning dipole. The dipole has a linear polarization, and, as it spins, it transitions from vertical to horizontal polarization and so on.

Up near the zenith (elevation angle of 90 degrees), relatively little variation occurs in gain (i.e., "thickness" of the line in Figure 1) as the dipole spins indicating near circular polarization. At lower elevation angles, the variation is much greater (i.e., "thicker" line) indicating close to linear polarization. In fact, the predominant polarization of patch antennas is horizontal at low elevation angles — a result of their small apparent vertical extent.

These results are fairly typical for a patch antenna, but much better performance is available from other antenna types; for example a NovAtel Pinwheel antenna or the quadrafilar helix antenna used in many Garmin products. The main catch with these designs is that they are not flat like a patch antenna, and so they occupy a greater volume.

The ratio of maximum gain to minimum gain is often referred to as *axial ratio* or *ellipticity* and is usually quoted in units of decibels (dB) with smaller magnitudes being preferred for GNSS antennas. (More formally, ellipticity is the inverse of the axial ratio, but by convention the sign is often dropped and the terms are used interchangeably.)

Assuming the transmit antenna at the satellite has perfect circular polarization, then ellipticity in the receive antenna yields the losses shown in **Figure 2**. High ellipticity, which corresponds to essentially linear polarization, leads to a 3 dB loss. In effect, half of the signal energy is lost due to polarization mismatch.

Referring to the GPS signal specification, ICD-GPS-200D, the reader will discover that the GPS satellite antennas have polarization errors themselves. Specifically:

> *3.3.1.9 Signal Polarization. The transmitted signal shall be right-hand circularly polarized (RHCP). For the angular range of ±14.3 degrees from boresight, L1 ellipticity shall be no worse than 1.2 dB for Block II/IIA and be no worse than 1.8 dB for Block IIR/IIR-M/IIF SVs. L2 ellipticity shall be no worse than 3.2 dB for Block II/IIA SVs and shall be no worse than 2.2 dB for Block IIR/IIR-M/IIF over the angular range of ±14.3 degrees from boresight.*

The ±14.3 degrees aspect of the specification ensures total earth coverage, since the edge of earth nadir angle is 13.9 degrees (as seen from the satellites). Taking the GPS Block IIR satellite's L1 maximum ellipticity of 1.8 dB as an example, we obtain the polarization loss results shown in **Figure 3**.

Now, because both the transmit and the receive antennas are presumed to have elliptical polarizations, we also need to account for the relative physical orientations of the two antennas' major elliptic axes of polarization, expressed as the rotation angle β. The

Mark Petovello is a Senior Research Engineer in the Department of Geomatics Engineering at the University of Calgary. He has been actively involved in many aspects of positioning and navigation since 1997 including GNSS algorithm development, inertial navigation, sensor integration, and software development.

Email: mpetovello@geomatics.ucalgary.ca

Professor Gérard Lachapelle holds a CRC/iCORE Chair in Wireless Location in the Department of Geomatics Engineering at the University of Calgary. He has been involved with GNSS since 1980 and has received numerous awards for his contributions in the area of differential kinematic GPS and indoor location.

Email: lachapel@geomatics.ucalgary.ca

β = 0 degrees case corresponds to the situation where the transmit and the receive antennas have their peak gain polarizations oriented the same way, producing the interesting result that, if both antennas have the same ellipticity, no polarization loss occurs (See blue line in Figure 3 when the receive antenna ellipticity is 1.8 dB). Assuming both antennas have a 1.8 dB ellipticity, physically rotating one of the antennas by 90 degrees about the boresight axis (to β = 90 degrees), we would see a 0.7 dB polarization mismatch loss.

At higher receive antenna ellipticities, the losses can be much greater. This situation tends to prevail for low elevation satellites where patch antenna ellipticities are much degraded. The patch antenna of Figure 1 has an ellipticity of about 20 dB at low elevation angles. So, we could see anywhere from 1.5 to 5 dB of polarization mismatch loss depending on polarization axis orientations.

In conclusion, circular polarization is preferred for GNSS systems because it is less sensitive to specific antenna orientations. Circular polarization also has the advantage of discriminating against multipath because reflections tend to have opposite sense polarizations that the antenna rejects. Finally, we note that practical antennas are not perfectly circular in their polarizations and that polarization mismatches can still be significant.

Editors' Note. An excellent discussion of multipath polarization can be found in Hannah, Bruce M., Modelling and Simulation of GPS Multipath Propagation, Ph. D. Thesis, Queensland University of Technology, 2001. Available online at <http://adt.library.qut.edu.au/adt-qut/public/adt-QUT20020326.160949/index.html>.

Manufacturers. The patch antenna used to illustrate an elevation gain pattern was a Novatel 521 from **NovAtel, Inc.**, Calgary, Alberta, Canada. The figure graphic was used by permission from NovAtel.

LOGAN SCOTT

Logan Scott, based in Breckenridge, Colorado USA, is a consultant specializing in radio frequency signal processing and waveform design for communications, navigation, radar, and emitter location. He has more than 27 years of military and civil GPS systems engineering experience. As a senior member of the technical staff at Texas Instruments, he pioneered approaches for building high-performance, jamming-resistant digital receivers. He is currently active in location-based encryption and authentication, high performance/low bias adaptive array technologies, and RFID applications. He holds 29 U.S. patents. IG