Lockheed photo

The Soft Approach

A Recipe for a Multi-System, Multi-Frequency GNSS Receiver

PC-based software GNSS receivers, which have served as powerful research tools for many years, are beginning to find applications outside the laboratory — in such areas as atmospheric modeling and synthetic aperture radar. The article describes the design of a multi-channel software receiver developed by the University of Calgary — challenges, trade-offs, results, and plans for further capabilities. A sidebar describes how the PLAN software receiver worked as a signal-monitoring tool to analyze the L5 signal broadcast by GPS satellite SVN 49.

CILLIAN O'DRISCOLL, DANIELE BORIO, MARK PETOVELLO, TOM WILLIAMS AND GÉRARD LACHAPELLE
UNIVERSITY OF CALGARY

On the morning of April 9, 2009, researchers in the Position, Location And Navigation (PLAN) group at the University of Calgary received an e-mail from a collaborator in the United States asking for assistance in assessing the performance of the new L5 test signal.

The signal was due to begin transmission from GPS space vehicle number (SVN) 49 the following morning at 6 a.m. local time in Calgary. We were informed that the test signal would be transmitting the pilot component of pseudorandom noise (PRN) code 63, contrary to the L1 and L2C signals broadcast from the same satellite, which would broadcast PRN 1 signals.

At that time the finishing touches had recently been made to the latest generation of the PLAN group's GNSS software Navigation receiver, a multi-system, multi-frequency software receiver. This receiver was already L5-capable, having been tested on signals from a hardware simulator; however, some modifications were still required to enable processing of the test signal. Within 15 minutes of receiving the SVN 49 signal specifications, the necessary modifications had been made and the receiver was ready to be tested.

Less than 24 hours later one of the authors (Dr. Borio) arrived in the PLAN laboratory, ready to begin recording the raw IF samples. Using a triple-frequency front-end and a RAID array, over one hour's worth of synchronous, high-rate data was collected in the L1, L2, and L5 frequency bands, both before and after the 6 a.m. switch-on of the L5 signal. (See the sidebar, "A Signal Monitoring Tool," for a discussion of the results from these observations of SVN 49.)

As shown in **Figure 1**, the receiver modifications worked perfectly and all signals from SVN 49 (shown as GPS.1 in the figure) were tracked successfully.

Although neither the L5 nor L2 signals contain any navigation message, we were still able to generate pseudorange measurements by sharing the navigation information from the L1 signal. Unfortunately, the ephemeris information provided by the satellite is unreliable and so these extra pseudorange measurements are not used in the navigation solution.

Notably, SVN 49's L2 carrier-to-noise (C/N$_0$) levels are considerably lower than the L1 levels. The variation in this difference corresponding to changes in satel-

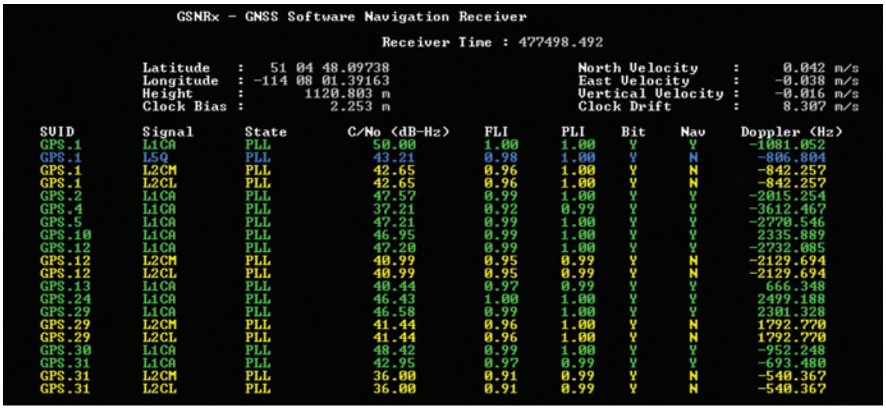

```
                    GSNRx - GNSS Software Navigation Receiver

                                 Receiver Time : 477498.492

                Latitude    :   51 04 48.09738                        North Velocity    :     0.042 m/s
                Longitude   : -114 08 01.39163                        East Velocity     :    -0.038 m/s
                Height      :       1120.803 m                        Vertical Velocity :    -0.016 m/s
                Clock Bias  :          2.253 m                        Clock Drift       :     8.307 m/s

   SVID         Signal      State        C/No (dB-Hz)     FLI        PLI      Bit     Nav     Doppler (Hz)
   GPS.1        L1CA        PLL              50.00        1.00       1.00      Y       Y      -1081.052
   GPS.1        L5Q         PLL              43.21        0.98       1.00      Y       N       -806.804
   GPS.1        L2CM        PLL              42.65        0.96       1.00      Y       N       -842.257
   GPS.1        L2CL        PLL              42.65        0.96       1.00      Y       N       -842.257
   GPS.2        L1CA        PLL              47.57        0.99       1.00      Y       Y      -2015.254
   GPS.4        L1CA        PLL              37.21        0.92       0.99      Y       Y      -3612.467
   GPS.5        L1CA        PLL              47.21        0.99       1.00      Y       Y      -2770.546
   GPS.10       L1CA        PLL              46.95        0.99       1.00      Y       Y       2335.889
   GPS.12       L1CA        PLL              47.20        0.99       1.00      Y       Y      -2732.085
   GPS.12       L2CM        PLL              40.99        0.95       0.99      Y       N      -2129.694
   GPS.12       L2CL        PLL              40.99        0.95       0.99      Y       N      -2129.694
   GPS.13       L1CA        PLL              40.44        0.97       0.99      Y       Y        666.348
   GPS.24       L1CA        PLL              46.43        1.00       1.00      Y       Y       2499.188
   GPS.29       L1CA        PLL              46.58        0.99       1.00      Y       Y       2301.328
   GPS.29       L2CM        PLL              41.44        0.96       1.00      Y       N       1792.770
   GPS.29       L2CL        PLL              41.44        0.96       1.00      Y       N       1792.770
   GPS.30       L1CA        PLL              48.42        0.99       1.00      Y       Y       -952.248
   GPS.31       L1CA        PLL              42.95        0.97       0.99      Y       Y       -693.480
   GPS.31       L2CM        PLL              36.00        0.91       0.99      Y       N       -540.367
   GPS.31       L2CL        PLL              36.00        0.91       0.99      Y       N       -540.367
```

FIGURE 1 Screen shot showing simultaneous tracking of L1, L2, and L5 signals from SVN 49. For clarity GPS L1 C/A signals are indicated in green, L2C signals in yellow, and the L5 signal in blue.

A Signal Monitoring Tool

Software receivers are not only research tools for the development and test of new navigation and positioning algorithms. The flexibility of software architectures enables them to record several pieces of information that are not limited to position and velocity. Correlator and discriminator outputs, frequency and phase lock indicators and several synchronization messages are just a few examples of the parameters that a software receiver makes available to users and researchers.

Hardware receivers are unable to provide the same level of accessibility to the GNSS signal and its parameters. In this respect, a software receiver can become a sophisticated signal-monitoring tool.

A clear example is provided by the PLAN group's GNSS software receiver, which has been successfully used for the analysis of the L5 signal broadcast by SVN 49. Moreover, the ability of this receiver to process simultaneously and synchronously data sets from different frequencies — L1, L2, and L5 in this case — has revealed consistent timing relationships between the L5 signal anomalies and the start of the L2CL epochs.

The complex correlator outputs (1 millisecond coherent integration) generated by the software receiver processing a data set recorded a few minutes after SVN 49 started broadcasting the L5 signal (April 9, 2009, 6 a.m., Calgary local time) are shown in **Figure 1**. Spikes are clearly present in the quadrature component of the L5 signal and are caused by phase jumps in the transmitted signal.

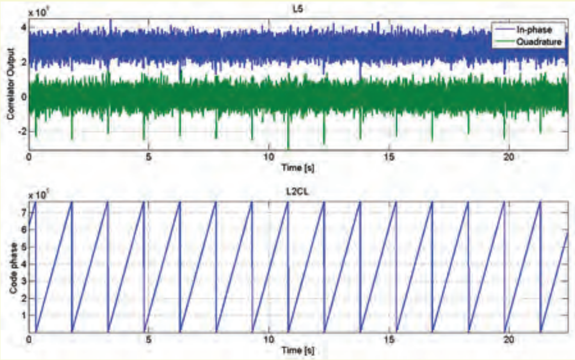

FIGURE 1 Comparison between the L5 correlator outputs (1 ms integration) with the code phase of the L2CL signal. Phase jumps in the transmitted L5 signal induce residual power in the quadra-phase component of the L5 signal. These phase jumps are synchronous with the beginning of the L2CL code periods.

In Figure 1, the L5 complex correlator outputs are compared against the code phase of the L2CL signal. The jumps in the L5 quadrature components are perfectly aligned with the start of the L2CL epochs (which in turn are aligned to the X1 epochs in the transmitter).

Since the phase jumps and the anomalies in the quadrature components repeat periodically every 1.5 seconds, the approach described in **Figure 2** was adopted to extract the average behavior of the correlator outputs. The correlator outputs were grouped in blocks of 1.5 seconds duration and averaged in order to remove the noise effects. Thus, the one-millisecond correlator outputs were divided into 1,500 bins, and averaging was conducted over all outputs, which are separated by 1.5 s. The average was performed over a total duration of 20 minutes.

Figure 3 shows the phase of the averaged correlator outputs: a phase jump of about -55 degrees clearly emerges. Moreover, other secondary peaks, regularly spaced by 80 milliseconds, can also be observed. Note that the "zero" correlator bin corresponds to the first millisecond after transmit time was extracted on the L1 signal (i.e., that time at which the first Z-count was encountered). Similar results have been obtained by processing different datasets, showing the consistency of the L5 anomalies.

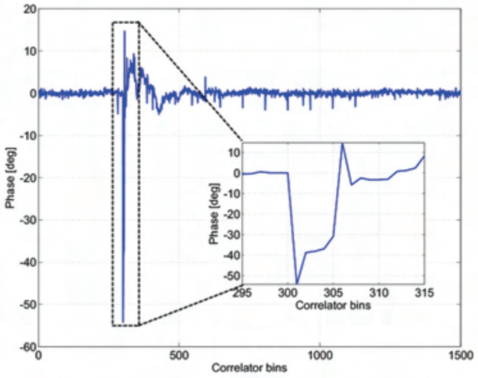

FIGURE 3 Mean phase anomaly observed on the L5 data transmitted by the SVN 49. A phase jump of approximately -55 degrees occurs every 1.5 seconds at the beginning of each L2C/X1 epoch. Secondary phase jumps, regularly spaced by 80 milliseconds, can also be observed.

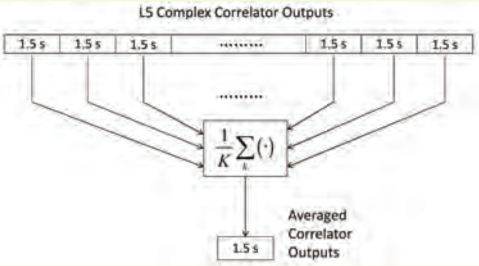

FIGURE 2 Processing scheme adopted for the analysis of the L5 signal anomalies. The L5 complex correlator outputs are grouped in blocks of 1.5 seconds (1,500 bins) and averaged in order to reduce the noise effect. The average is performed over a total duration of 20 minutes.

lite elevation leads us to conclude that the variation results from three factors: 1) the difference in transmitted power (nominally 1.5 decibels), 2) the difference in receiver antenna gain pattern in the L1 and L2 frequency bands, and 3) the three-decibel loss from processing the L2CM and L2CL codes separately.

The key to the success of the L5 experiment is, in a word, design. Software receiver implementation in the PLAN group has gone through a number of design cycles since it began in 2004, with work on the latest incarnation, beginning in mid-2006.

This article will describe the objectives, design challenges, and engineering solutions that led to creation of PLAN's GNSS software receiver. It will also discuss the features and applications of software receivers in general and the PLAN receiver in particular.

The Secret to This Success

The original version was a flexible single-frequency receiver, capable of processing GPS and Galileo signals in the L1 band using a variety of processing strategies, including standard, vector-based, and ultra-tight integration of GNSS with inertial navigation systems (INS).

By May 2008 it was clear that a redesign was necessary in order to expand the receiver in a number of key ways:

- to more easily introduce multi-frequency support
- to simplify the addition of new systems and signals
- to separate the high rate signal processing tasks to ease the integration of co-processors and thus improve efficiency.

A design process involving many months of discussion, testing, and retesting led to the current receiver architecture, which is illustrated schematically in **Figure 2**.

The key components of a software receiver that we identified during this process are 1) a front-end, 2) a processing manager, 3) a list of satellites and 4) a navigation solution. With this in mind, an object-oriented software architecture was developed in C++ to parallel these components.

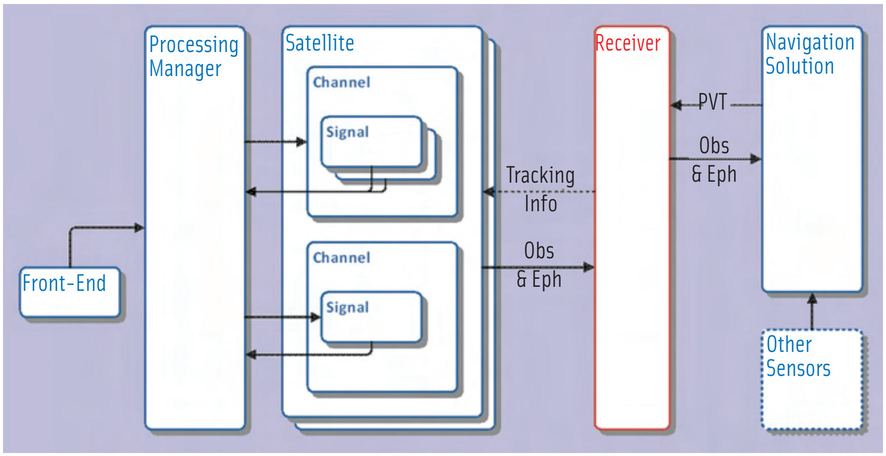

FIGURE 2 Structural overview of the PLAN software receiver (dotted lines represent optional items or data flow)

We made a particular effort to define simple, yet comprehensive interfaces for related objects (e.g., all front-end objects). In so doing, a GNSS receiver can be "composed" (to use programming jargon) using any objects that conform to the defined interfaces. In turn, this provides a very high level of flexibility and code re-use.

In Figure 2, front-end objects act as a source of synchronized samples from one or multiple frequency bands. The processing manager object is responsible for performing all high-rate processing on a per-signal basis. Usually this consists of performing the Doppler removal and correlation (DRC) — also called baseband mixing and despreading, respectively — but other operations are also possible, such as implementation of multi-access interference mitigation techniques.

Each satellite object encapsulates a set of signals which are acquired and tracked in channels (either individually or in groups), and from which measurements are obtained. These objects can use any strategy they choose to perform the receiver operations, as long as they conform to the interfaces as defined by the base class object.

The receiver uses a "factory" creational pattern (again, more jargon) to create new satellite objects. Finally, the navigation solution object accepts satellite observations (pseudorange, pseudorange rate, and carrier phase) and ephemeris objects and then solves for position, velocity, and time. The receiver is responsible for managing the interaction between these various objects.

Under this design, the receiver accepts a front-end, a processing manager, a satellite factory, and a navigation solution at start-up, and does not need to be aware of the particular algorithms used by each object to achieve its tasks. Such a general structure greatly speeds up development time.

A Multitude of Multiplicities

We are entering a golden age in GNSS research. With multiple GNSS systems, each transmitting multiple signals on multiple frequencies, the opportunities for research and development are manifold. As has been well documented elsewhere, software receivers provide an excellent platform for exploiting these opportunities. However, we believe that the structured approach undertaken by the PLAN group represents a distinct way forward in receiver development.

The L5 anecdote related at the start of this article reveals one advantage of the structured approach: adaptability. We have found this structure to be incredibly useful on numerous occasions. For example, in late March 2009 we decided to add GLONASS capability to the PLAN receiver. With only a few hours (about three) of design and programming, the ability to acquire and track GLONASS C/A signals on both L1 and L2 had been added to the receiver.

The addition of navigation message decoding and satellite trajectory calcu-

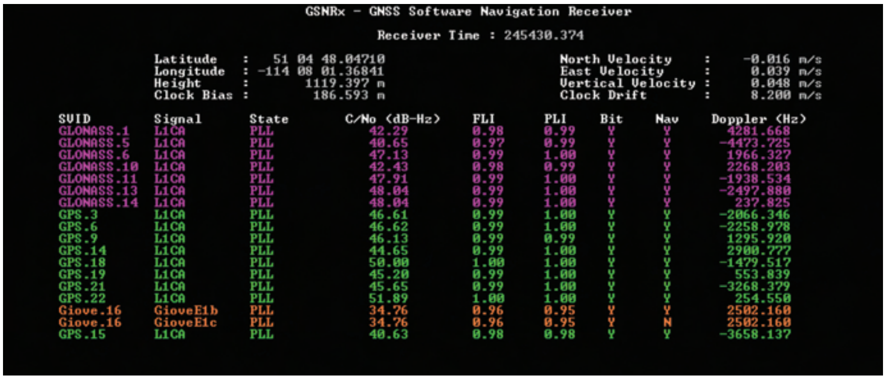

```
                     GSNRx - GNSS Software Navigation Receiver

                           Receiver Time : 245430.374

          Latitude    :   51 04 48.04710              North Velocity     :    -0.016 m/s
          Longitude   : -114 08 01.36841              East Velocity      :     0.039 m/s
          Height      :       1119.397 m              Vertical Velocity  :     0.048 m/s
          Clock Bias  :        186.593 m              Clock Drift        :     8.200 m/s

 SVID         Signal      State      C/No (dB-Hz)    FLI      PLI    Bit    Nav    Doppler (Hz)
 GLONASS.1    L1CA        PLL            42.29       0.98     0.99    Y      Y      4281.668
 GLONASS.5    L1CA        PLL            40.65       0.97     0.99    Y      Y     -4473.725
 GLONASS.6    L1CA        PLL            47.13       0.99     1.00    Y      Y      1966.327
 GLONASS.10   L1CA        PLL            42.43       0.98     0.99    Y      Y      2268.203
 GLONASS.11   L1CA        PLL            47.91       0.99     1.00    Y      Y     -1938.534
 GLONASS.13   L1CA        PLL            48.04       0.99     1.00    Y      Y     -2497.880
 GLONASS.14   L1CA        PLL            48.04       0.99     1.00    Y      Y       237.825
 GPS.3        L1CA        PLL            46.61       0.99     1.00    Y      Y     -2066.346
 GPS.6        L1CA        PLL            46.62       0.99     1.00    Y      Y     -2258.978
 GPS.9        L1CA        PLL            46.13       0.99     0.99    Y      Y      1295.920
 GPS.14       L1CA        PLL            44.65       0.99     1.00    Y      Y      2900.777
 GPS.18       L1CA        PLL            50.00       1.00     1.00    Y      Y     -1479.517
 GPS.19       L1CA        PLL            45.20       0.99     1.00    Y      Y       553.839
 GPS.21       L1CA        PLL            45.65       0.99     1.00    Y      Y     -3268.379
 GPS.22       L1CA        PLL            51.89       1.00     1.00    Y      Y       254.550
 Giove.16     GioveE1b    PLL            34.76       0.96     0.95    Y      Y      2502.160
 Giove.16     GioveE1c    PLL            34.76       0.96     0.95    Y      N      2502.160
 GPS.15       L1CA        PLL            40.63       0.98     0.98    Y      Y     -3658.137
```

FIGURE 3 Screenshot of the software receiver tracking signals from three constellations (GPS, GLONASS, GIOVE). Note that the SV Giove.16 corresponds to the GIOVE -B satellite, as per the GIOVE ICD. For clarity GPS signals are indicated in green, GLONASS signals in magenta and GIOVE signals in orange.

lations required some additional work (not least of which was deciphering the GLONASS interface control document, or ICD). Nonetheless, the generic structure of the receiver meant that no receiver level modifications were required.

The receiver still handles generic "satellite," "ephemeris," and "observation" objects, and remains agnostic as to the particulars of these objects. **Figure 3** presents one of the first GPS/GLONASS/GIOVE L1 data sets that we processed, in which seven GLONASS satellites, nine GPS satellites, and the Giove-B satellite are all clearly present.

In addition to the ability to track multiple systems at multiple frequencies, the PLAN receiver is also able to apply a variety of higher level integration algorithms. At present, the receiver has the capability for stand-alone, vector-based, and ultra-tight GNSS/INS integration.

We have tested the receiver using GPS L1/L2C measurements in the ultra-tight configuration mode and, although too few L2C-capable satellites are operating at present to compute a stand-alone navigation solution, the results are promising. The receiver has also been successfully used in high-accuracy carrier-phase data-processing software packages.

More esoteric processing strategies are also possible using this structure. For example, the Galileo E1a signal consists of two side-lobes separated by a little over 30 megahertz. The software receiver is capable of tracking either the wideband signal or the narrow band side-lobes (upper and lower). **Figure 4** shows the Doppler measurements obtained from each of the three E1a signal components.

Independently processing the full E1a signal and its side-band components allows comparative analysis of the measurements obtained from the different channels. In addition, the Figure 4 inset shows our receiver simultaneously processing all five components of the GIOVE-B E1 signal.

Challenges

The software receiver structure described here yields many benefits. However, a number of challenges associated with the software receiver approach remain. Primary amongst these is the issue of data throughput.

As with any design process, many trade-offs have been made in developing the PLAN receiver. In particular, receiver designers commonly expect that flexibility comes at the cost of speed. Indeed, the use of an object-oriented approach is often advised against on the grounds that it is inherently less efficient than a purely procedural approach.

In the face of this expectation, one of our primary goals while designing the receiver was to enable *more* efficient computation of the high-rate processes in the receiver, while maintaining a high degree of flexibility and adaptability. The identification of the processing manager object referred to earlier was the key to achieving this goal.

In our experience, much greater computational efficiencies can be achieved by bringing all the high-rate computations together, such that all

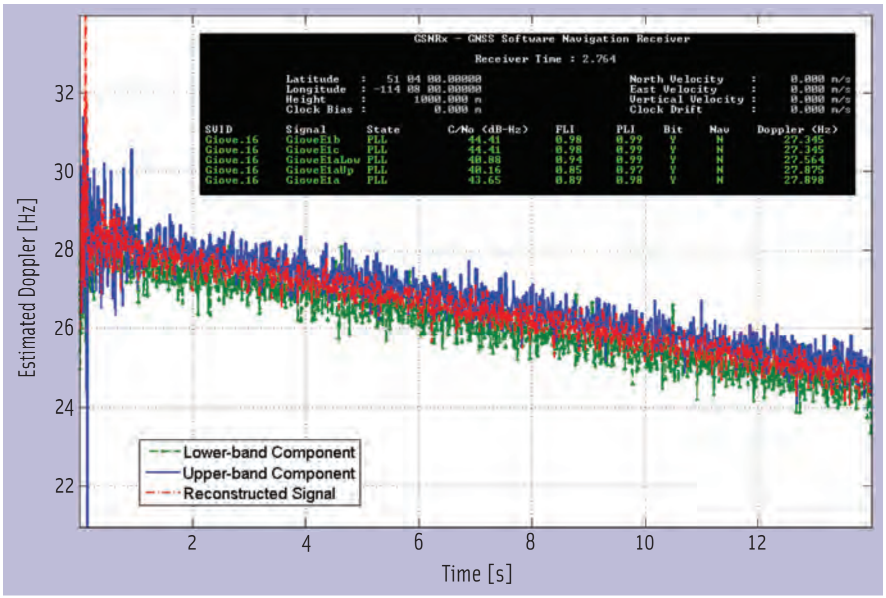

```
                    GSNRx - GNSS Software Navigation Receiver

                            Receiver Time : 2.764

           Latitude   :   51 04 00.00000                North Velocity    :      0.000 m/s
           Longitude  : -114 08 00.00000                East Velocity     :      0.000 m/s
           Height     :       1000.000 m                Vertical Velocity :      0.000 m/s
           Clock Bias :          0.000 m                Clock Drift       :      0.000 m/s

SVID       Signal      State      C/No (dB-Hz)     FLI      PLI     Bit    Nav    Doppler (Hz)
Giove.16   GioveE1b    PLL           44.41         0.98     0.99     Y      N        27.345
Giove.16   GioveE1c    PLL           44.41         0.98     0.99     Y      N        27.345
Giove.16   GioveE1aLow PLL           40.88         0.94     0.99     Y      N        27.564
Giove.16   GioveE1aUp  PLL           40.16         0.85     0.97     Y      N        27.875
Giove.16   GioveE1a    PLL           43.65         0.89     0.98     Y      N        27.898
```

FIGURE 4 Estimated Doppler for GIOVE-B E1a upper, lower and combined signals. Inset shows a screen shot of the receiver simultaneously tracking E1a, E1b, E1c as well as the upper and lower sidelobes of E1a from the same satellite.

the information required to perform the Doppler removal and correlation operations for *all* signals is available in one object at the same time. This structure has the added advantage of enabling the integration of co-processors, such as field programmable gate arrays (FPGAs) and graphics processing units (GPUs) with relative ease.

A clear demonstration of the ability of the PLAN software receiver to handle large computational burdens in an efficient manner appeared in late March of this year, when the software was successfully interfaced to a commercially available, 12-channel portable (being in a USB flash drive form factor), L1-only front-end. This front-end has a sampling rate of slightly greater than 16 Msps and an IF bandwidth of approximately four megahertz.

Running on a laptop PC with two gigabytes of RAM and dual-processors, the PLAN receiver is capable of acquiring, tracking, and navigating in real time, using up to 12 channels simultaneously. Note that this includes both GPS and Galileo/GIOVE constellations. This real-time receiver uses only the CPU cores on the PC and does not take advantage of any GPU co-processors that may be present.

The remaining key challenge in GNSS software receiver design is the availability of hardware for down-converting and digitizing the variety of RF spectral bands that are now of interest to GNSS researchers. The commercially available L1 front-end is one of only a handful of complete designs available on the market at present that easily interface with PC hardware. Very few currently offer more than single-frequency operation, and none that we are aware of currently offer the flexibility to acquire signals from all bands of interest.

At present, the PLAN group uses a high-fidelity vector signal analyzer (VSA) in conjunction with a RAID array to acquire high rate, synchronous multi-frequency data on up to three frequency bands at a time.

Summary

The modular design adapted for PLAN's software receiver greatly facilitates research in many areas of GNSS signal and navigation processing. The ability to quickly add new signals, new systems, and new signal processing techniques was a fundamental goal of the design process, and the success of this approach has been demonstrated through a number of examples.

A common misconception suggests that a software receiver designed entirely in C++ would be too cumbersome and slow to operate in real time. On the contrary, the PLAN receiver can operate in real-time with up to 12 channels at sampling rates of over 16 Msps.

Much research and development work remains to be done, even within the constraints of the software receiver as defined in Figure 2, and certainly a whole world of possibility exists beyond this specific definition. Signal quality monitoring, high-sensitivity operation, atmospheric modeling, and synthetic aperture radar applications are only a few of the potential further uses of the software receiver.

We also recently commenced development of a graphical user interface (GUI) to replace the current console interface. Although this development does not contribute directly to the research goals of the PLAN group, the ability to visualize receiver data at run-time can be very useful for quick identification of interesting phenomena, such as RF interference.

Manufacturers

The software receiver developed by the University of Calgary PLAN is the GNSS Software Navigation Receiver (GSNRx), which has been licensed to a number of third parties. The L1-only front-end used for real-time development of the GSNRx is the SiGe GN3S Sampler available from **Sparkfun** <www.sparkfun.com>, Boulder, Colorado, USA. (The GN3S Sampler was co-developed by the **University of Colorado Aerospace Department** and **SiGe Semiconductor**, Ottawa, Canada, to make the software for SiGe's SE4110L GPS front-end ASIC work with GNU/Linux. The Sampler incorporates the ANT-555 L1 patch antenna from **Onshine Enterprise Company Ltd.**, Chung Ho City, Taipei, Taiwan. However, most of PLAN's post-mission testing uses the NovAtel 702 GG antenna from **NovAtel, Inc.**, Calgary, Alberta, Canada, and, for the L1/L2/L5 work, the NovAtel 704-X.) The vector signal analyzer (VSA) is the NI PXI-5661 2.7 GHz RF Vector Signal Analyzer from **National Instruments**, Austin, Texas, USA. The hardware simulator used to test the GSNRx is the GSS7700 from **Spirent Communica-**

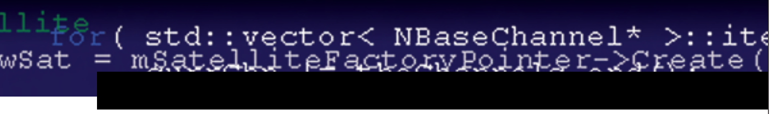

tions, Paignton, England. The GSNRx runs on a PC computer with an Intel Core 2 Duo processor from the **Intel Corporation,** Santa Clara, California, USA.

Additional Resources

[1] Capua, R. and M. Antonini (2008), "A Soft Touch: Sogei's GPS/Galileo Software Receiver and Institutional GNSS Applications in Italy", *Inside GNSS*, Fall 2008, pp. 33-39

[2] Petovello, M. G., and C. O'Driscoll, G. Lachapelle, D. Borio and H. Murtaza (2008), "Architecture and Benefits of an Advanced GNSS Software Receiver," *Journal of Global Positioning Systems*, vol. 7, no. 2, pp. 156-168.

[3] Won, J.-H., and T. Pany and G. W. Hein, "GNSS Software Defined Radio: Real Receiver or Just a Tool for Experts?", *Inside GNSS*, July/August 2006, pp. 48-56

Authors

Cillian O'Driscoll Ph.D. is a senior research associate in the Position, Location And Navigation (PLAN) group at the University of Calgary. He received his Ph.D. in electrical and electronic engineering from University College Cork, Ireland in 2007. His research interests are in all areas of GNSS signal processing, with emphasis on weak signal acquisition.

Daniele Borio received the M.S. degree in communication engineering from Politecnico di Torino, Italy, the M.S. degree in electronics engineering from ENSERG/INPG de Grenoble, France, and the doctoral degree in electrical engineering from Politecnico di Torino. Dr. Borio has been a senior research associate at the PLAN group of the University of Calgary, Canada, since January 2008. His research interests include the fields of digital and wireless communication, location, and navigation.

Mark Petovello is an assistant professor 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.

Tom Williams has B.E.(Hons) and M.E. degrees from the University of Canterbury, New Zealand, and a PhD from the University of Alberta, all in mechanical engineering. He is a senior research associate at the PLAN group in the Department of Geomatics Engineering at the University of Calgary. His research interests include inertial navigation and multi-sensor integration with particular focus on pedestrian applications.

Gerárd Lachapelle is a professor of Geomatics Engineering at the University of Calgary where he is responsible for teaching and research related to location, positioning, and navigation. He has been involved with GPS developments and applications since 1980. He has held a Canada Research Chair/iCORE Chair in wireless location since 2001. IG