The GPS Assimilator: a Method for Upgrading Existing GPS User Equipment to Improve Accuracy, Robustness, and Resistance to Spoofing

Todd E. Humphreys, Jahshan A. Bhatti, *The University of Texas at Austin, Austin, TX*
Brent M. Ledvina, *Coherent Navigation, San Mateo, CA*

BIOGRAPHIES

Todd E. Humphreys is an assistant professor in the department of Aerospace Engineering and Engineering Mechanics at the University of Texas at Austin. He received a B.S. and M.S. in Electrical and Computer Engineering from Utah State University and a Ph.D. in Aerospace Engineering from Cornell University. His research interests are in estimation and filtering, GNSS technology, GNSS-based study of the ionosphere and neutral atmosphere, and GNSS security and integrity.

Jahshan A. Bhatti is pursuing a Ph.D. in the Department of Aerospace Engineering and Engineering Mechanics at the University of Texas at Austin, where he also received his B.S. His research interests are in development of small satellites, software-defined radio applications, and GNSS technologies.

Brent M. Ledvina is Director of New Business and Technology at Coherent Navigation in San Mateo, CA. He is also an adjunct faculty member of the Bradley Department of Electrical and Computer Engineering at Virginia Tech. He received a B.S. in Electrical and Computer Engineering from the University of Wisconsin at Madison and a Ph.D. in Electrical and Computer Engineering from Cornell University. His research interests are in the areas of ionospheric physics, space weather, estimation and filtering, and GNSS technology and applications.

ABSTRACT

A conceptual method is presented for upgrading existing GPS user equipment, without requiring hardware or software modifications to the equipment, to improve the equipment's position, velocity, and time (PVT) accuracy, to increase its PVT robustness in weak-signal or jammed environments, and to protect the equipment from counterfeit GPS signals (GPS spoofing). The method is embodied in a device called the GPS Assimilator that couples to the radio frequency (RF) input of an existing GPS receiver. The Assimilator extracts navigation and timing information from RF signals in its environment—including non-GNSS signals—and from direct baseband aiding provided, for example, by an inertial navigation system, a frequency reference, or the GPS user. The Assimilator optimally fuses the collective navigation and timing information to produce a PVT solution which, by virtue of the diverse navigation and timing sources on which it is based, is highly accurate and inherently robust to GPS signal obstruction and jamming. The Assimilator embeds the PVT solution in a synthesized set of GPS signals and injects these into the RF input of a target GPS receiver for which an accurate and robust PVT solution is desired. A prototype software-defined Assimilator device is presented with three example applications.

INTRODUCTION

What will GNSS receivers look like five years from now? The answer, of course, depends on the application. Mass-market receivers used in applications that do not require precision positioning and timing (e.g., hand-held units for hikers) will likely remain simple single-frequency L1-C/A-based GPS devices. On the other hand, a growing segment of military and civilian GNSS users will demand greater accuracy and reliability from their receivers than can be offered by single-frequency GPS. They will want their GNSS devices to be multi-frequency to combat ranging errors due to ionospheric delay, and multi-system to improve satellite availability and robustness against signal interference.

Major commercial GNSS receiver manufacturers already have product roadmaps in place that anticipate these demands. Manufacturers realize that they will be at a competitive disadvantage relative to their peers if they only offer a subset of available GNSS signals to sophisticated users. "Why should I have to choose between signals?" their customers will complain, "I'd like all of them!"

Then there is the issue of GNSS security. There was a time, perhaps 20 years ago or more, when computer users were largely unconcerned with the security of their personal computers. That time has passed. As any victim of a computer virus knows, firewalls, anti-virus software, and protocols for secure data transfer are no longer optional, but required. Likewise, the deepening dependence of the civil infrastructure on GNSS—especially for timing synchronization—and the potential for financial gain or high-profile mischief make civil GNSS jamming and spoofing a gathering threat. Since the publication of the U.S. Department of Transportation's Volpe Report on GPS dependence nearly a decade ago [1], GNSS security researchers have repeatedly warned that civil GPS is not yet

Preprint of the 2010 ION GNSS Conference Portland, OR, September 21–24, 2010

secure, and that users trust its signals at their peril. As Professor David Last commented at a recent conference on GNSS security, "Navigation is no longer about how to measure where you are accurately. That's easy. Now it's how to do so reliably, safely, robustly."

Secure positioning, navigation, and timing (PNT) will require use of all available means: inertial navigation systems, stable frequency sources, multiple antennas [2], cryptographic authentication [3], and all radio frequency signals from which PNT information can be extracted—including non-GNSS signals and signals never intended to be used for PNT.

In short, PNT devices in critical applications five years from now will likely be remarkably capable and secure devices that adhere to an all-signals-in-view, all-available-means philosophy.

Meanwhile, however, the overwhelming majority of GNSS receivers—even those in critical applications—are simple L1 C/A-based devices that fail when signals are blocked or jammed, complaining "Need clear view of sky." What is more, no commercially-available civil GNSS receiver, as far as the authors are aware, incorporates even rudimentary defenses against spoofing. Are these receivers to be considered obsolete? Perhaps. And perhaps the prudent course of action is to replace them with secure and reliable modern devices.

A decision to replace existing receivers, however, cannot be made lightly. The hundreds of thousands of deployed GNSS receivers across the globe today represent an enormous investment in equipment and training. Moreover, in many cases the GNSS receiver is only an embedded subcomponent of a larger PNT-reliant system. It may be inconvenient, unsafe, or expensive to replace these embedded devices with modern counterparts. Nonetheless, the vulnerability of existing receivers, embedded and otherwise, to signal obstruction, jamming, and spoofing, and their inability to make use of modernized GNSS signals and other signals of opportunity, leaves much to be desired.

As an alternative to replacement of existing equipment, this paper proposes augmentation. A technique has been developed for upgrading existing GNSS user equipment to address their shortcomings without requiring hardware or software modifications to the equipment. The technique re-purposes the portable civil GPS spoofer described in [4] to generate "friendly" spoofing signals whose implied navigation solution is derived from a fusion of GPS and other observables. The technique is embodied in a device, called the GPS Assimilator, whose output is injected directly into the radio frequency (RF) input of existing GPS equipment to immediately robustify the equipment against GPS outages and interference.

Consider three examples of existing devices for which Assimilator augmentation is potentially preferable to replacement:

Time Reference Receivers These devices, which typically cost several thousand dollars apiece, couple a GPS receiver to a stable OCXO or atomic frequency reference. They are used extensively in telecom networks; in particular, the IS-95 CDMA-based digital cellular standard and its progeny require each base station to be synchronized with a GPS receiver so that the timing of transmissions can be controlled to better than 10 μs. Modern time reference receivers easily meet this requirement. They are also capable of extended (24-36 hours) "holdover mode" operation in case of signal blockage or jamming. Nonetheless, recent experiments with a popular time reference receiver model at the University of Texas Radionavigation Laboratory have revealed that these receivers are easily spoofed. Within the span of one hour, the pulse-per-second output of a receiver originally locked to authentic GPS signals can be driven to a larger than 10-μs error without raising any of the receiver's internal alarms. The obvious implication—that it would take a malefactor less than one hour to render a cell-phone base station inoperable via spoofing—is cause for concern. Unfortunately, due to the tight coupling between the GPS receiver and the oscillator it disciplines, the GPS receiver component cannot be easily swapped with a modern, secure version.

Phasor Measurement Units (PMUs) Also known as synchrophasors, these devices couple a GPS receiver to power measurement equipment in order to simultaneously obtain the phasor values of voltages and currents at particular instants of time [5]. Although currently used primarily for monitoring, it is anticipated that these devices will see widespread future use in closed-loop control systems designed to increase the carrying capacity of the power distribution grid. The installed base of PMUs represents a considerable investment. Unfortunately, like time reference receivers, the GPS receivers coupled to PMUs are easily spoofed. Timing errors lead to phasor angle errors, and, when these reach several degrees ($\sim$100 μs), they can destabilize closed-loop control of transient swings.

Embedded Military GPS Receivers Unsurprisingly, GPS receivers find widespread use in armed forces worldwide. Several hundred thousand devices have been procured by the armed forces of the United States and foreign military sales customers over the last five years—a substantial collective investment. A large fraction of military-grade GPS receivers are used as embedded receivers, being coupled to targeting, tracking, and communications equipment via well-defined and field-tested interfaces. These critical downstream systems can be rendered inoperable when incoming GPS signals are blocked or jammed. Counterintuitive as it may seem that a vehicle-to-vehicle communications system would be disabled by GPS blockage, such is the nature of modern dependence on precision nav-

igation and timing.

Each of these example devices would see benefits in accuracy, robustness, or security from coupling to an Assimilator. The following section describes the Assimilator as a conceptual device. Thereafter, a working prototype is introduced and initial experimental results are given.

I. CONCEPTUAL ASSIMILATOR

The Assimilator concept is based on the principle that from virtually any modern environment one can extract a wealth of navigation and timing-related information. Thus, the Assimilator behaves opportunistically, scanning ambient radio waves for PNT information while also accepting baseband data from an inertial navigation system (INS), an external time source, or directly from the user. All extracted PNT information is fused to yield an optimal navigation and timing solution. Up to this point, the Assimilator is no different from other proposed systems for robust navigation and timing that employ an "all available means" philosophy. For these proposed systems, as for the Assimilator, GPS is but one of several potential sources of PNT data.

Having obtained a fused PNT solution, however, the Assimilator takes an unusual additional step: it embeds the PNT solution in a consistent set of synthesized GPS L1 C/A signals, the common-denominator of all existing GNSS equipment. By casting its solution into this output format, the Assimilator can deliver the additional accuracy, robustness, and security of its solution to any GNSS device by simply injecting its output into the RF input of the target device. Thus, the Assimilator acts as a conduit for funneling ambient PNT information to existing GNSS equipment, without requiring hardware or software changes to the equipment. Despite the inefficiency of regenerating GPS RF signals after already having obtained a PNT solution, the assimilative approach is warranted in cases where, due to tight embedded coupling with expensive downstream equipment or due to user familiarity, it becomes more cost-effective or safer to augment existing equipment than to replace it.

A. Assimilator Components

Figure 1 introduces the Assimilator concept in block diagram form. Each subcomponent will be treated in turn.

A.1 Front-End Bank, Multi-System Receiver Module, and Navigation and Timing Fusion Module

A bank of RF front ends digitizes segments of the RF spectrum containing signals that potentially bear PNT information. Naturally, GPS, Galileo, and other GNSS signals are among these, but so are terrestrial radionavigation signals such as LORAN/eLORAN (now only outside the U.S.) as well as communication signals from cell phone towers or Iridium satellites—signals not originally designed for PNT.

The digitized data exiting the RF front-end bank are routed to a software-defined multi-system receiver module implemented on a digital signal processor (DSP). Here, each target signal is tracked, either independently or as part of a vector tracking loop. The observables extracted from each signal, including pseudorange, carrier phase, carrier Doppler, signal strength, and navigation data, are sent to a central navigation and timing fusion module, which also accepts baseband PNT inputs. For each input signal there exists in the fusion module an *a priori* noise and signal dynamics model, enabling the data to be optimally fused in a Bayesian estimation framework.

A.2 Anti-Spoofing Modules

Both the multi-system receiver module and the navigation and timing fusion module are equipped with anti-spoofing software. Module A, which resides within the multi-system receiver module, continuously scans the RF data streams entering the Assimilator for spoofing signatures. Module B, which resides within the navigation and timing fusion module, watches for inconsistencies between observables in the centralized solution. Further details on these modules will be given later on.

A.3 Embedded GPS Signal Simulator

The output of the navigation and timing fusion module feeds an embedded GPS signal simulator. The embedded signal simulator is depicted in Fig. 1 as residing partly outside the DSP because it includes an external RF upconversion component. Several options are possible for signal simulation:

Impaired GPS L1 CA If reception of GPS L1 C/A signals is impaired, then the embedded simulator can generate signal replicas in two ways:

Matched The simulator generates the same GPS L1 C/A signals that would be visible to the Assimilator in clear sky conditions, corrected to be consistent with the fusion module's PNT solution.

Fictitious The simulator generates a partially or entirely fictitious signal set whose implied solution is consistent with the fusion module's PNT solution.

Unimpaired GPS L1 CA If several sufficiently strong GPS L1 C/A signals are available in the Assimilator's environment, then the embedded simulator can generate clean replicas of these, passing on to the target receiver the additional accuracy of the fused PNT solution by continually applying corrections to the code start times and carrier phases of the GPS L1 C/A signals. Again, there are two options for generation:

Non-aligned No attempt is made to align the code phase of the simulated and authentic signals.

Code-aligned The simulator generates GPS signals that are approximately code-phase-aligned with the corre-

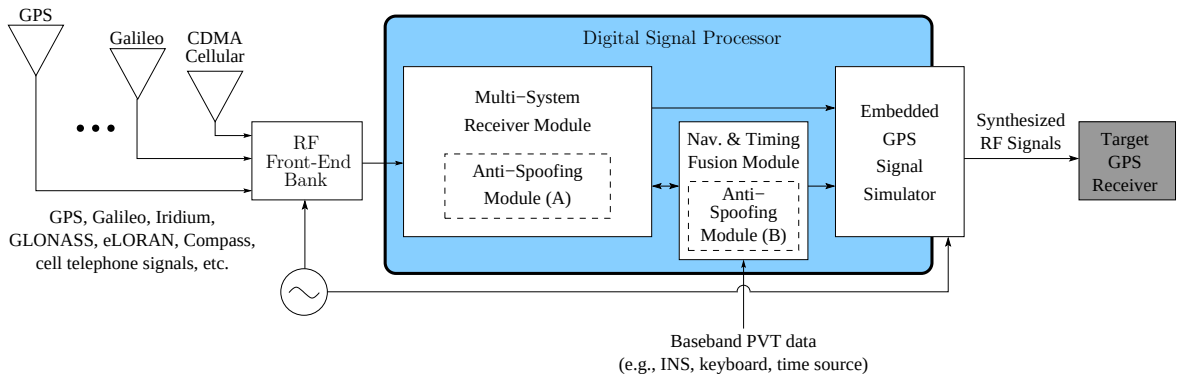

Fig. 1. Block diagram of the conceptual Assimilator.

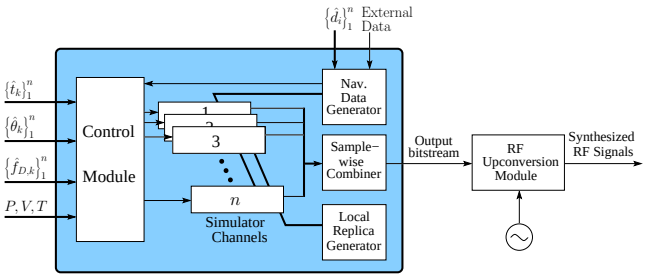

Fig. 2. Expanded view of the embedded GPS signal simulator.

sponding authentic GPS L1 C/A signals at the RF input to the target GNSS receiver. (The simulated signals are not in general exactly code-phase aligned due to corrections applied to match the simulated signals' implied solution with the fusion module's PNT solution.) The advantage of such alignment is that the synthesized signals appear approximately as the authentic signals to the target receiver, which enables a user to "hot plug" the Assimilator into a target receiver with no interruption in PNT.

An expanded view of the embedded GPS signal simulator is shown in Fig. 2. Its subcomponents will be treated in turn.

Control Module The control module coordinates the generation of the synthesized GPS signals by directing the carrier phase, carrier frequency, and code phase in each of n simulator channels. It accepts the following inputs from the multi-system receiver module's L1 C/A tracking channels: the estimates $\{\hat{t}_k\}_1^n$ of the start times of the kth ranging code interval on receiver channels 1-n; the estimates $\{\hat{\theta}_k\}_1^n$ of the beat carrier phase on receiver channels 1-n at times $\{\hat{t}_k\}_1^n$; and the estimates $\{\hat{f}_{D,k}\}_1^n$ of the Doppler frequency shift on receiver channels 1-n at times $\{\hat{t}_k\}_1^n$. In addition, the control module accepts the following inputs from the fusion module: the estimated current time T and the estimated position P and velocity V of the Assimilator's antenna. The control module configures each simulator channel to generate a single GPS L1 C/A signal with a carrier frequency and code phase that are consistent with the fusion module's position P, velocity V and time T solution.

Simulator Channels Each of the n simulator channels can be configured to generate a unique GPS C/A signal, modeled as

$$x_n(\tau_i) = A_n(\tau_i) d_n(\tau_i) C_n(\tau_i - t_{n,k}) \times Q\left\{\sin\left[2\pi f_{IF}\tau_i + \theta_n(\tau_i)\right]\right\} \quad (1)$$

$$\dot{\theta}_n(\tau_i = t_{n,k}) = f_{D,n,k} \quad (2)$$

where $x_n(\tau_i)$ is the ith sample of the signal, τ_i is the time of the ith sample, $A_n(\tau_i)$ is the amplitude at τ_i, $d_n(\tau_i)$ is the navigation data bit value that applies at τ_i, $C_n(\tau_i - t_{n,k})$ is the ranging code chip value that applies at τ_i, $t_{n,k}$ is the start time of the kth ranging code interval, $Q\{\cdot\}$ is a quantization function, f_{IF} is the intermediate frequency, $\theta_n(\tau_i)$ is the beat carrier phase at τ_i, and $f_{D,n,k}$ is the Doppler frequency shift at time $t_{n,k}$. The ranging code function $C_n(\tau)$ can be expressed as

$$C_n(\tau) = \sum_{j=-\infty}^{\infty} c_{n,j} \Pi_{T_c}(\tau - jT_c) \quad (3)$$

and the navigation data bit function $d_n(\tau)$ as

$$d_n(\tau) = \sum_{j=-\infty}^{\infty} d_{n,j} \Pi_{T_d}(\tau - jT_d) \quad (4)$$

where $\{c_{n,j}, c_{n,j+1}, ...\}$ and $\{d_{n,j}, d_{n,j+1}, ...\}$ are the unique ranging code chip sequence and navigation data bit sequence corresponding to the GPS satellite whose signal is being emulated on the nth simulator channel, T_c and T_d are the duration of one ranging code chip and one navigation data bit, and $\Pi_T(\tau)$ is the usual rectangular support function equal to unity over $0 \le \tau < T$ and zero otherwise.

Local Replica Generator The local replica generator generates the ranging code samples $\{C_n(\tau_i)\}, i = 1, 2, ...,$ and

the quantized carrier replica samples

$$Q\left\{\sin\left[2\pi f_{IF}\tau_i + \theta_n(\tau_i)\right]\right\}, \quad i = 1, 2, 3, ... \tag{5}$$

The generation of carrier replicas is described in [6]. A command and data bus conveys phase and frequency information from the simulator channels to the local replica generator, and returns local carrier and code replicas from the local replica generator to the simulator channels.

Navigation Data Generator The navigation data bit sequence $\{d_{n,j}, d_{n,j+1}, ...\}$ required by the nth simulator channel is generated in one of two ways. When GPS L1 C/A signals are available, a steady stream of navigation data bits is taken from the GPS L1 C/A channels of the multi-system receiver module. Data bits extracted from the authentic signals are fed to the navigation data generator and compiled into a signal-specific databit library. If the receiver module loses lock on an authentic signal corresponding to one being simulated, then the navigation data generator continues to populate the simulated bitstream with navigation data from the data bit library, but the library no longer gets updated as before. This continues until the control module reconfigures the simulator channel to generate another signal.

If a simulator channel is configured to generate a signal for which there is no recent databit library, then the navigation data generator produces data bits consistent with a standard ephemeris for the corresponding satellite. The generator can also accept user-supplied satellite data in the form of a databit library or a satellite ephemeris.

Sample-wise Combiner Combination of the signals generated in each of the simulator channels is performed digitally sample-by-sample in the sample-wise combiner. For typical Assimilator operation, all output signal are weighted equally so that the target GPS receiver sees a set of received signals with equal carrier-to-noise (C/N_0) ratios. However, the Assimilator is also capable of matching ambient C/N_0 ratios in case this information is useful for downstream systems. In this case, the ith sample from the nth simulator channel is weighted by the simulated amplitude $A_n(\tau_i)$ and summed with the corresponding samples from the other simulator channels, each weighted appropriately. The combined signal is then re-quantized to produce an output bitstream.

RF Upconversion Module The output bitstream of the sample-wise combiner is routed to an RF upconversion module comprising a digital-to-analog converter, frequency mixers, filters, and a signal attenuator. The upconversion module converts the digital signal into a set of synthesized GPS signals at RF. The reference oscillator that drives the RF upconversion module must be the same oscillator that drives the Assimilator's RF front-end bank.

B. Capabilities and Limitations

B.1 Accuracy

For maximum compatibility with legacy GNSS receivers, the Assimilator outputs only GPS L1 C/A signals. The narrow bandwidth (∼2 MHz) of the C/A ranging code limits the accuracy of the pseudorange-based position and time solution that the Assimilator can deliver to the target receiver. The Assimilator compensates for this limitation by generating synthetic signals with high C/N_0 and by selecting a constellation geometry that minimizes the geometric dilution of precision (GDOP).

Suppose elevation masking in the target receiver is disabled so that for any number n of synthesized signals the optimal constellation geometry, which includes below-the-horizon satellites, can be employed. The inverse square-root relationship between minimum GDOP values implies that the minimum attainable GDOP for n satellites is

$$\mathrm{GDOP}(n) = \mathrm{GDOP}(4)\sqrt{\frac{4}{n}} \tag{6}$$

where $\mathrm{GDOP}(4) = \sqrt{5/2}$ is the minimum GDOP for 4 satellites [7].

Gains in precision due to increasing n are, unfortunately, vitiated by a decrease in C/N_0 for all signals as a consequence of the interference contributed by each additional simulated signal. Assume that signals are simulated with equal power and that the combined signal is quantized to one bit in the sample-wise combiner (one bit quantization is a practical choice for signal generation on a compact, low-power platform). Then one can show empirically that the C/N_0 value for each simulated signal drops from roughly 54 dB-Hz for $n = 4$ to 51 dB-Hz for $n = 9$.

These C/N_0 values can be incorporated into a ranging precision model and combined with the GDOP values from Eq. (6) to bound the position and time precision the Assimilator can deliver to the target receiver. Figure 3 shows this precision bound as a function of the number of simulated signals. The plot indicates that under the assumptions presented here a precision of approximately 13 cm is attainable. The plot further indicates that there is no advantage in simulating more than 6 signals.

Like the precision of code phase measurements derived from Assimilator-generated signals, carrier phase precision is determined by the C/N_0 values of the simulated signals. For $n = 6$ simulated signals of equal power, the rms carrier phase errors are less than 0.4 degrees for each signal.

Another aspect of accuracy has to do with how well the Assimilator can eliminate the effects of ionospheric delay by exploiting multiple-frequency transionospheric signals. The broadcast ionospheric model used in legacy

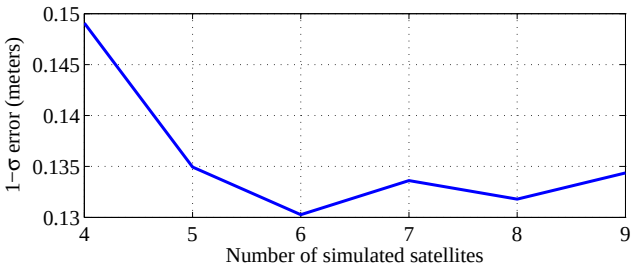

Fig. 3. Minimum 1-σ position and range-equivalent timing error deliverable to a target receiver by a 1-bit quantized Assimilator output stream, as a function of the number of simulated signals. For the underlying ranging precision model, a code tracking loop averaging time of 10 seconds and an early-minus-late correlator spacing of 1 chip is assumed.

single-frequency GPS receivers is known to eliminate only about 50% of the rms ranging error [8], leaving day-side zenith delay errors up to 10 m for a quiet ionosphere and much worse for an active ionosphere. By contrast, standard dual-frequency L1 C/A & L2C GPS measurements at $C/N_0 = 45$ dB-Hz can be used to reliably estimate the ionospheric delay to within 0.66 meters. Thus, the multi-frequency Assimilator can pass on to a single-frequency target receiver significant benefits in position and timing accuracy.

B.2 Robustness

The Assimilator's PNT solution is, by virtue of the diverse navigation and timing data that feed it, inherently robust against GNSS signal obstruction and jamming. Signals from cell phone base stations, Iridium satellites, and LORAN transmitters are tens of dB stronger than those from GNSS satellites. Thus, not only is the Assimilator robust to GNSS outages, it can also withstand substantial blockage, jamming, or other interference in the cell phone (1.9 GHz), Iridium (1.6 GHz), and LORAN (100 kHz) frequency bands. Naturally, in a complete GNSS signal blackout, the PNT solution that the Assimilator feeds to the target receiver will be degraded, but by leveraging non-GNSS navigation and timing sources, the Assimilator limits this degradation substantially. Baseband aiding from an INS or stable frequency reference lowers the Assimilator's tracking threshold for GNSS signals and permits the Assimilator to "coast" through periods of complete RF blackout.

Ionospheric scintillation poses another challenge for GNSS receiver robustness. The deep power fades and accompanying fast phase transitions induced by equatorial ionospheric scintillation stress a receiver's carrier tracking loops, and, as severity increases, can lead to navigation bit errors, cycle slipping, and complete loss of carrier lock. The Assimilator makes best use of incoming GNSS signals by incorporating carrier phase tracking loops that are specially designed for scintillation robustness (for maximum navigation accuracy, all carrier tracking loops within the Assimilator track carrier phase, not just frequency). One simple technique for extending the mean time between cycle slips (and decreasing the chances of frequency unlock) is to wipe off the navigation data bits from data-bearing channels so that a traditional full-cycle carrier tracking loop can be employed instead of a half-cycle Costas loop [9]. The navigation data generator within the Assimilator's embedded signal simulator stores a signal-specific data bit library for each GPS L1 C/A signal. Because the C/A navigation message repeats every 12.5 minutes, this library can be used to predict the value of data bits that are received during scintillation-induced power fades. A network connection on the Assimilator permits data bit libraries to be downloaded from a remote server. Also, the Assimilator benefits from access to modernized GNSS signals whose pilot (data-free) channels are by design more scintillation-robust than the legacy GPS C/A signal.

B.3 Security

As with robustness, the Assimilator is inherently more secure against signal spoofing than legacy civil GNSS receivers because of the diverse data from which its PNT solution is derived. In Ref. [4] it was demonstrated that developing a portable single-frequency civil GPS spoofer is relatively straightforward, but carrying out a successful spoofing attack becomes much more difficult when multiple signal types must be spoofed simultaneously. The Assimilator's anti-spoofing module B, which resides within the navigation and timing fusion module, exploits this inherent security by comparing the observables output by each of the receivers in the multi-system receiver module. The module looks for outliers in pseudorange measurements in an extension of the popular receiver-autonomous integrity monitoring (RAIM) algorithm.

At a lower level in the processing chain, the Assimilator's anti-spoofing module A, which resides within the multi-system receiver module, monitors the individual incoming data streams for irregular behavior, employing, for example, the data bit latency and vestigial signal defenses introduced in [4], the multi-antenna defense introduced in [2], or the signal quality monitoring technique introduced in [10].

Stronger civil spoofing defenses than these require cryptographic signal authentication. Over the last decade, GNSS security researchers have proposed techniques for navigation message and spreading code authentication [3,11], and authentication based on exploiting the known relationship between GPS C/A signals and the encrypted GPS Y code on L1 [12]. Of these, navigation message authentication (NMA) appears to be the most practical while still providing a strong defense against spoofing.

The primary objection to civil NMA as applied to GPS is that it would require a change to the GPS signal-in-space specification, otherwise known as the GPS interface control document (ICD). However, NMA could in fact be introduced in a way that is fully compatible with the current ICD by exploiting the flexible L2 or L5 navigation message structure. In this approach, one of the presently undefined navigation messages would be allocated for transmission of a digital signature. The GPS Control Segment would generate a public and private key pair, publishing the public key and keeping the private key secret. Users would download the public key from a trusted repository. With the private key, the Control Segment would digitally sign all navigation data messages transmitted between the kth and $(k+1)$th digital signature message. The signature, at least 128 bits long for adequate security, would be transmitted to the user in the payload of the $(k+1)$th signature message.

With this scheme, users could periodically authenticate all navigation data bits modulated on a given signal. Furthermore, only a slight extension to the GPS ICD would be required: an additional paragraph would be added to define the new digital-signature-bearing message

The Assimilator could be adapted to implement the algorithm that verifies the received digital signature. Thus, by combining the Assimilator with navigation message authentication on the GPS L2C or L5 signals, legacy GNSS receivers could be cryptographically secured against GPS spoofing attacks.

Whether by cryptographic means or otherwise, if the Assimilator detects a spoofing attack it alerts the user and excludes the spoofed signals from its internal PNT estimate. The synthesized GPS signals that the Assimilator continuously sends to the target receiver are accordingly spoof-free, and the target receiver is protected from the spoofing attack.

II. PROTOTYPE ASSIMILATOR

A prototype Assimilator has been built to prove the basic feasibility of the conceptual Assimilator introduced above. The prototype is an extension of the GRID software-defined GNSS receiver introduced in [13] and [14] and the GPS spoofer introduced in [4]. The following two subsections will describe the prototype and offer initial experimental results.

A. Description

The prototype Assimilator is a dual-frequency device with a rudimentary spoofing defense. Its embedded signal simulator generates output signals in the code-aligned, unimpaired GPS L1 C/A simulation mode. The device receives L1 C/A and L2C GPS signals and outputs L1 C/A signals with code phases corrected for ionospheric delay. Figure 4 shows a picture of the prototype assimilator coupled to a GPS time reference receiver. The lower right panel of Fig. 4 is a close-up of the DSP chip that simultaneously handles dual-frequency signal acquisition and tracking, and single-frequency signal simulation.

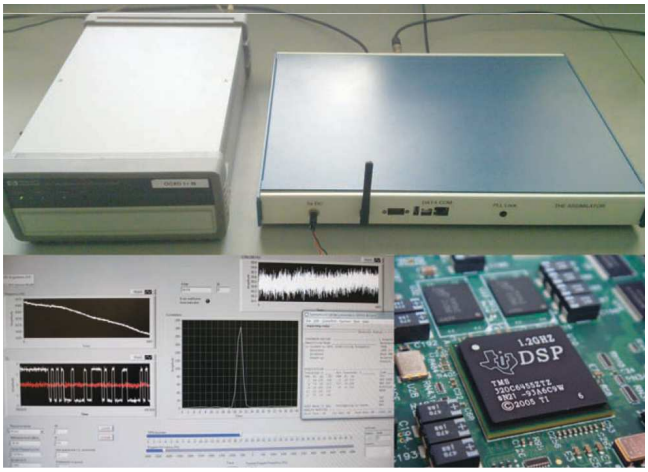

Fig. 4. Top: The prototype Assimilator device connected to a GPS time reference receiver. Lower left: Screen shot of a software GPS receiver tracking the Assimilator output and of the GPS time reference receiver's status page showing a timing lock. Lower right: The Assimilator's DSP core. This single chip performs all processing required for both signal tracking and signal generation.

A.1 Multi-System Receiver Module

Though it currently only tracks GPS L1 C/A and L2C signals, the prototype Assimilator's multi-system receiver module is designed for expansion. Written in object-oriented C++, the module's principal feature is an extensible array of so-called Bank objects, each of which acts as an independent software receiver. Class polymorphism is exploited so that all Bank objects share a common structure.

Figure 5 illustrates the principal components of the Bank object. Data from each RF front end are stored in separate data buffers, with sub-buffers assigned to each of the front ends' quantization bits. The prototype uses two-bit quantization but its processing algorithms can be adapted for N-bit quantization. Each data buffer feeds a separate Bank object, of which the prototype Assimilator has two: one for GPS L1 C/A and one for GPS L2C. Each Bank object, in turn, houses an array of Channel objects, one for each independent signal to be tracked. Channel objects are made up of phase, frequency, and code tracking loop objects and other supporting members. All Channel objects within a Bank feed into a single Observables object that encapsulates the observables extracted from the signals to which the Bank is devoted. Each Bank object feeds a sequence of Observables objects to the navigation and timing fusion module for fusion into a single PNT solution.

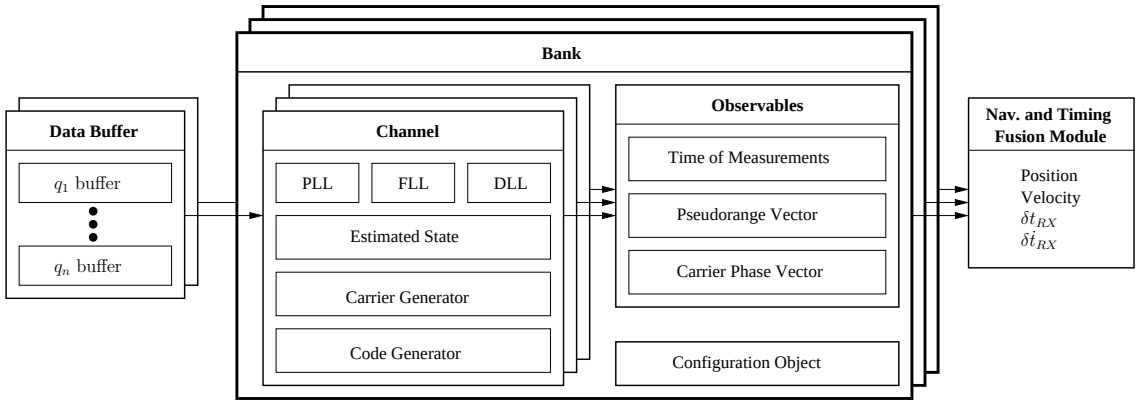

Fig. 5. Software objects within the prototype implementation of the multi-system receiver and the navigation and timing fusion modules.

A.2 Navigation and Timing Fusion Module

The prototype Assimilator fuses dual-frequency GPS measurements into a single-frequency simulated GPS RF output by re-generating clean versions of the C/A signals that it tracks and by compensating for ionospheric delay on each of the simulated signals.

Only 8 L2C-capable GPS satellite are currently in orbit, which implies that direct measurements of the L1 ionospheric delay are not available for all satellite-to-receiver paths. To compensate for ionospheric delay in all simulated signals, the prototype Assimilator uses whatever dual-frequency measurements are available to adjust the parameters of a simple single-layer ionospheric model. Delay estimates drawn from this model are then applied to adjust the code phase of simulated signals. No attempt is made at present to compensate for ionospheric-induced advances in the simulated signals' carrier phases.

A.3 Anti-Spoofing

For anti-spoofing, the prototype Assimilator implements the data bit latency defense [4,10] and a phase trauma indicator. The simple data bit latency defense is premised on the difficulty of (1) predicting or synthesizing a consistent stream of navigation data bits for each signal, or (2) re-transmitting the broadcast GPS data bits with an undetectable latency. In the prototype Assimilator's implementation of the defense, Channel objects continuously monitor each navigation data bit stream and flag suspicious shifts in bit synchronization. The defense is far from foolproof, but it is simple to implement and substantially raises the bar for a successful spoofing attack.

The Assimilator's phase trauma indicator is simply an implementation of the phase lock indicator introduced in [15]. Like the data bit latency defense, the phase trauma indicator is simple to implement yet difficult to avoid triggering during a spoofing attack. A triggering of the phase trauma indicator during an interval of high nominal carrier-to-noise ratio raises suspicion of a spoofing attack.

A.4 Embedded Signal Simulator

The prototype Assimilator's embedded signal simulator functions just as described earlier for the conceptual Assimilator except that the prototype makes no attempt to choose the best possible combination of signals to simulate; it simply selects the strongest n tracked C/A signals (usually 8) for simulation.

B. Preliminary Performance Results

B.1 Processing Demands

When tuned for efficiency, the prototype Assimilator meets real-time deadlines with computational resources to spare. The processing power of the prototype's DSP is such that it can run the equivalent of 135 parallel GPS C/A channels. Because their longer ranging codes cannot be stored in on-chip memory, L2C channels require the equivalent processing of 4 C/A channels. More expensive still are the simulator channels, each of which requires an equivalent of 5.4 C/A channels. At full capability, the Assimilator can track 14 GPS L1 C/A signals and 14 GPS L2C signals while simultaneously generating 8 simulation signals, in addition to performing a 1-Hz navigation solution and periodic background acquisition.

B.2 Power Requirements

Of the prototype's hardware subsystems, the DSP and its peripherals, unsurprisingly, draw the most power, requiring 6 W. The dual-frequency RF front end draws 2 W. A separate single-board computer used for network communications, remote reprogrammability, and housekeeping draws 1.3 W. The RF upconversion module draws less than 700 mW. Thus, the total power draw of the prototype Assimilator is less than 10 W.

III. USAGE EXAMPLES

A. Protecting a Time Reference Receiver from an Unsophisticated Spoofing Attack

As a first usage example, consider the Assimilator as an in-line anti-spoofing module. The top panel of Fig. 4 shows the Assimilator in this role protecting a time reference GPS receiver. The Assimilator monitors incoming GPS signals and raises a flag upon detection of irregularities in the start time of the navigation data bits or upon detection of phase trauma not explained by low C/N_0 values. As long as this flag remains unasserted, the time reference receiver has a reasonable assurance that its pulse-per-second output is properly synchronized with true GPS time.

B. Reducing Ionospheric Errors in Single-Frequency Target Receivers

A dual-frequency Assimilator can reduce ionospheric errors in a single-frequency target receiver to which it is attached. An experiment involving ionospheric delay modeling errors was conducted to illustrate the utility of the prototype Assimilator in this role. The experiment was carried out in after-the-fact processing on a desktop PC running the Assimilator code. A dual-frequency front end identical to that of the prototype Assimilator was used to digitize 300 seconds of dual-frequency data from a high-quality GPS signal simulator. The signal simulator's scenario profile was set to generate signals consistent with a Klobuchar ionospheric model with low electron content (∼2 meters zenith delay at L1). Data from the signal simulator were used to ensure a well-defined truth position. After-the-fact processing was required since the GPS signal simulator and the prototype Assimilator were not co-located.

The Assimilator ingested the dual-frequency digital data and generated a combined set of output signals. These were routed to a software-defined single-frequency GPS receiver (the target receiver) whose internal ionospheric model had been disabled. Midway through the run, the Assimilator began compensating for ionospheric delay by adjusting the code phase of its simulated signals. Prior to this moment, the target receiver showed positioning errors of approximately 3 meters, mostly in the altitude component. To the target receiver it appeared as though there was a mismatch between the broadcast ionospheric model and the true ionosphere. However, when the Assimilator activated its dual-frequency-based ionospheric compensation, the target receiver's altitude errors were immediately reduced. Figure 6 plots the altitude error time history. Clearly, even with this experiment's low-electron-content ionosphere, the target receiver benefited from its coupling to the Assimilator.

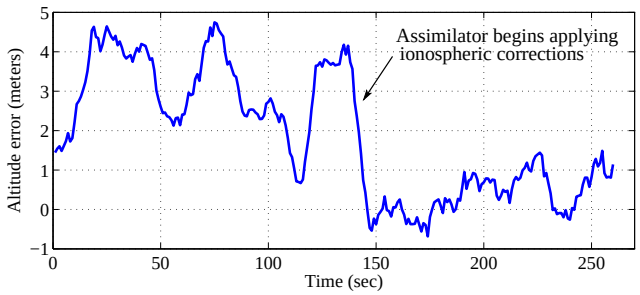

Fig. 6. Time history of altitude errors within an Assimilator-aided single-frequency receiver. Midway through the processing run the Assimilator activates ionospheric compensation, reducing the target receiver's altitude errors by approximately 2.5 meters.

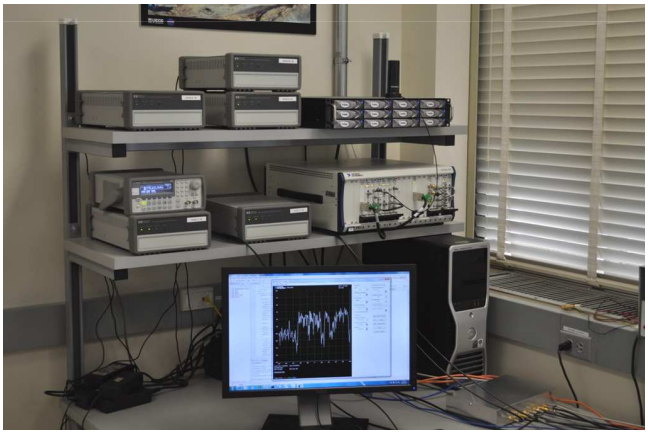

Fig. 7. The alternate prototype Assimilator setup. A general-purpose wideband dual-frequency front-end feeds data to a desktop processor in which the Assimilator processing takes place. The CDMA cellular antenna can be seen atop the RAID storage array.

C. Weak-Signal Tracking via CDMA Cellular Signal Aiding

In a third example application, the Assimilator was configured to receive code division multiple access (CDMA) cellular telephone signals from nearby towers. This application required an alternate prototype Assimilator, pictured in Fig. 7, with a front-end capable of simultaneously receiving the GPS L1 and 1900-2000 MHz cell telephone bands. This setup allows the Assimilator to exploit the native frequency stability of the CDMA cell telephone signals to compensate for frequency instability within its own (inexpensive) local oscillator. Adequate compensation for frequency instability permits the long coherent integration intervals required to track GNSS signals with low C/N_0. Thus, the Assimilator enables GNSS use deep indoors or in environments where GNSS may be subject to interference.

Figure 8 illustrates in block diagram form the processing required to exploit the cellular aiding signal. Weak GNSS signals are despread and down-converted based on approximate knowledge of receiver position and time. The comparably much stronger cell telephone signals are tracked

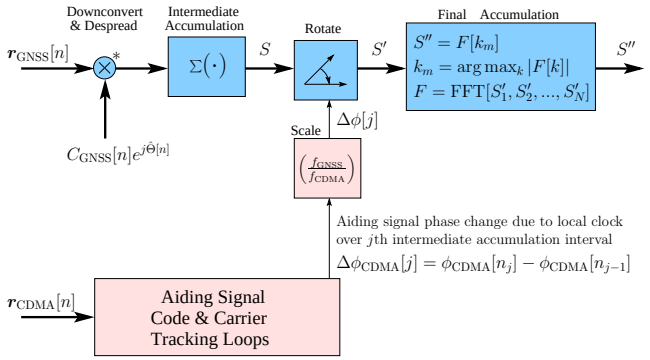

Fig. 8. Block diagram illustrating the technique of exploiting stable CDMA cellular signals of opportunity to extend the coherence time of a low-cost local oscillator.

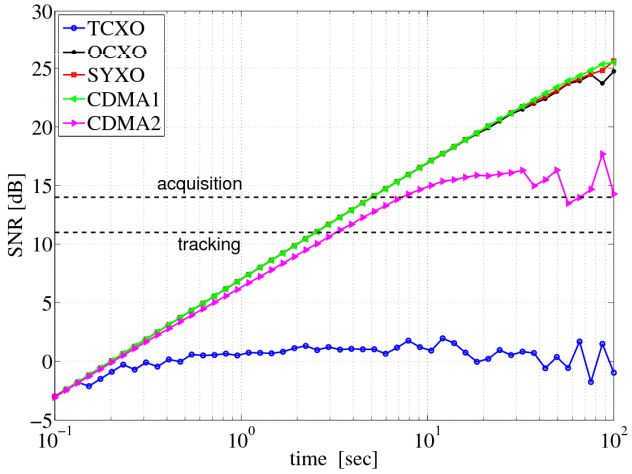

Fig. 9. Pre-detection SNR as a function of integration time for a GNSS receiver driven by several different local oscillators. The two dashed horizontal lines show pre-detection SNR thresholds for reliable acquisition and tracking; see [16] for details on the threshold values. CDMA1 and CDMA2 refer to the local oscillator as "fixed up" by stable CDMA cellular signals. SYXO refers to single-differenced GPS carrier phase playing the role of the local oscillator and represents a limiting case for coherent integration: atmospheric variations and GPS satellite clock instability prevent pre-detection SNR from rising any faster than this trace.

with an independent set of code and carrier tracking loops to extract estimates of the shift in beat carrier phase that occurs over each GNSS intermediate accumulation interval. The phase shift estimates, after proper scaling, drive a rotation of the complex intermediate GNSS signal accumulations, thereby compensating for local-clock-induced phase shifts over the intermediate accumulation intervals.

The benefit of this frequency stability transfer technique is evident in Fig. 9, which shows the pre-detection signal-to-noise ratio (SNR) as a function of integration time for a GPS signal with nominal $C/N_0 = 7$ dB-Hz as tracked by a receiver driven by several different local oscillators. The short coherence time of the low-cost temperature-compensated crystal oscillator (TCXO) prevents the pre-detection SNR from ever reaching thresholds required for acquisition and tracking. By contrast, the black trace, which corresponds to an oven-controlled crystal oscillator (OCXO) with a coherence time exceeding 100 seconds, easily meets the thresholds. Likewise, a CDMA-cellular-signal-aided low-cost local oscillator can sustain long coherence times, as indicated by the CDMA1 and CDMA2 traces. Interestingly, the relatively poor performance of CDMA2 compared with CDMA1 suggests that not all base-station-transmitted cellular signals are of equivalent stability.

By extending the coherence time of its local clock, the Assimilator is able to track severely attenuated GNSS signals in indoor environments and, in turn, pass on to the target receiver a set of strong GNSS signals whose implied timing and positioning solution is of useful accuracy. This technique of opportunistic frequency stability transfer for extending the coherence time of inexpensive GNSS receiver clocks is described in further detail in [16].

IV. CONCLUSIONS

A technique has been presented for upgrading existing GNSS user equipment, without requiring hardware or software modifications to the equipment, to improve its accuracy, to increase its robustness in weak-signal or jammed environments, and to secure it against counterfeit GNSS signals. The technique is embodied in a device called the GPS Assimilator that acts opportunistically to extract navigation and timing information from its environment. The Assimilator encodes this information into a set of standard GPS L1 C/A signals with which all legacy GNSS receivers are natively compatible. A dual-frequency prototype Assimilator with a rudimentary spoofing defense has been presented. Initial experimental results show the prototype successfully correcting ionospheric errors in a single-frequency target receiver and suggest that, by exploiting the native frequency stability of cell telephone signals, the Assimilator can enable indoor GNSS reception.

Efforts are underway to develop the next-generation Assimilator prototype: a compact device equipped with a robust cryptographic defense against spoofing and capable of tracking dual-frequency GPS and CDMA cell telephone signals. Eventually, as board sizes are reduced, the Assimilator's processing core can be housed within its antenna enclosure, offering GNSS users the possibility of upgrading their current receivers with a simple change of antenna.

V. END NOTES

The assimilator concept and early hardware were developed at Coherent Navigation, Inc., a startup company of which Dr. Ledvina and Dr. Humphreys are co-founders, along with four others. Coherent Navigation Inc. has filed a patent covering the assimilator concept and related tech-

nologies.

References

[1] "Vulnerability assessment of the transportation infrastructure relying on the Global Positioning System," Tech. rep., John A. Volpe National Transportation Systems Center, 2001.

[2] Montgomery, P. Y., Humphreys, T. E., and Ledvina, B. M., "A Multi-Antenna Defense: Receiver-Autonomous GPS Spoofing Detection," *Inside GNSS*, Vol. 4, No. 2, April 2009, pp. 40–46.

[3] Scott, L., "Anti-spoofing and authenticated signal architectures for civil navigation systems," *Proc. ION GPS/GNSS 2003*, Institute of Navigation, Portland, Oregon, 2003, pp. 1542–1552.

[4] Humphreys, T. E., Ledvina, B. M., Psiaki, M. L., O'Hanlon, B. W., and Kintner, Jr., P. M., "Assessing the spoofing threat: development of a portable GPS civilian spoofer," *Proceedings of ION GNSS 2008*, Institute of Navigation, Savannah, GA, 2008.

[5] Phadke, A., Pickett, B., Adamiak, M., Begovic, M., Benmouyal, G., Burnett Jr, R., Cease, T., Goossens, J., Hansen, D., Kezunovic, M., et al., "Synchronized sampling and phasor measurements for relaying andcontrol," *IEEE Transactions on Power Delivery*, Vol. 9, No. 1, 1994, pp. 442–452.

[6] Ledvina, B., "Real-Time Generation of Bit-Wise Parallel Carrier Replicas Applied to a GPS/GNSS Software Receiver," *IEEE Transactions on Aerospace and Electronic Systems*, 2010, to be published.

[7] Spilker, J. J., *Global Positioning System: Theory and Applications*, chap. 5: Satellite Constellation and Geometric Dilution of Precision, American Institute of Aeronautics and Astronautics, Washington, D.C., 1996, pp. 177–208.

[8] Klobuchar, J. A., *Global Positioning System: Theory and Applications*, chap. 12: Ionospheric Effects on GPS, American Institute of Aeronautics and Astronautics, Washington, DC, 1996, pp. 485–515.

[9] Humphreys, T. E., Psiaki, M. L., and Kintner, Jr., P. M., "Modeling the effects of ionospheric scintillation on GPS carrier phase tracking," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 46, No. 4, Oct. 2010.

[10] Ledvina, B. M., Bencze, W. J., Galusha, B., and Miller, I., "An In-Line Anti-Spoofing Module for Legacy Civil GPS Receivers," *Proceedings of the ION ITM*, Institute of Navigation, San Diego, CA, Jan. 2010.

[11] Hein, G., Kneissi, F., Avila-Rodriguez, J.-A., and Wallner, S., "Authenticating GNSS: Proofs against spoofs, Part 2," *Inside GNSS*, September/October 2007, pp. 71–78.

[12] Lo, S., DeLorenzo, D., Enge, P., Akos, D., and Bradley, P., "Signal Authentication," *Inside GNSS*, Vol. 0, No. 0, Sept. 2009, pp. 30–39.

[13] Humphreys, T. E., Ledvina, B. M., Psiaki, M. L., and Kintner, Jr., P. M., "GNSS Receiver Implementation on a DSP: Status, Challenges, and Prospects," *Proceedings of ION GNSS 2006*, Institute of Navigation, Fort Worth, TX, 2006.

[14] O'Hanlon, B. W., Psiaki, M. L., Kintner, Jr., P. M., and Humphreys, T. E., "Development and Field Testing of a DSP-Based Dual-Frequency Software GPS Receiver," *Proceedings of ION GNSS 2009*, Institute of Navigation, Savannah, GA, 2009.

[15] Van Dierendonck, A. J., *Global Positioning System: Theory and Applications*, chap. 8: GPS Receivers, American Institute of Aeronautics and Astronautics, Washington, D.C., 1996, pp. 329–407.

[16] Wesson, K., Pesyna, K., Bhatti, J., and Humphreys, T., "Opportunistic Frequency Stability Transfer for Extending the Coherence Time of GNSS Receiver Clocks," *Proceedings of the ION GNSS Conference*, Institute of Navigation, Portland, Oregon, 2010.