The WAAS/L5 Signal for Robust Time Transfer: Adaptive Beamsteering Antennas for Satellite Time Synchronization

David S. De Lorenzo[1], Sherman C. Lo[1], Jiwon Seo[1], Yu-Hsuan Chen[2], Per K. Enge[1]
[1] *Stanford University, USA*
[2] *National Cheng Kung University, Taiwan*

BIOGRAPHY

David De Lorenzo is a Research Associate in the Stanford University GPS Research Laboratory and an Engineer in the Research Group at Polaris Wireless.

Sherman Lo is a Senior Research Engineer in the Stanford University GPS Research Laboratory, where he is Associate Investigator for Stanford University efforts on the FAA Alternate Position Navigation & Time (APNT) study.

Jiwon Seo is a Postdoctoral Scholar with joint appointments in the Stanford University GPS Research Laboratory and the Space Environment and Satellite Systems Laboratory.

Yu-Hsuan Chen is a Ph.D. candidate in electrical engineering at National Cheng Kung University, Taiwan, and a visiting scholar in the Stanford University GPS Research Laboratory.

Per Enge is Professor of Aeronautics and Astronautics at Stanford University, where he is the Kleiner-Perkins, Mayfield, Sequoia Capital Professor in the School of Engineering. Dr. Enge has received the Kepler, Thurlow, and Burka Awards from the Institute of Navigation and is a Fellow of the Institute of Navigation and of the Institute of Electrical and Electronics Engineers.

ABSTRACT

Geostationary satellites of the Wide Area Augmentation System (WAAS) offer a novel, robust, and cost-effective means of synchronizing time at widely-separated ground facilities, to levels of ~50ns, without the need for dedicated long-distance wired communication networks. However, reliance on satellite-based signals for time synchronization in high-reliability applications is problematic without explicit hardening against radio frequency interference (RFI). The primary innovations to be discussed in this paper are: (1) adaptive electronically-steered multi-element antenna arrays and signal processing strategies for RFI mitigation, (2) live signal, synthetic interference, and hardware-in-the-loop testing of jammer cancellation algorithms, and (3) preparations for over-the-air interference tests which will probe the effects of front-end saturation on digital beamsteering performance. We describe a hardware system assembled from readily-available commercial building blocks (data acquisition system, antennas, etc.), and a critical goal of this research is to realize significant GPS anti-jam performance in an open-architecture (or non-defense-related) platform. Therefore, the central innovations in this research enable adaptive electronic beamsteering with high-dynamic-range signals (14-bit I/Q digitization), employing commercial off-the-shelf (COTS) hardware and computer systems, and targeted for a civilian high-reliability, high-volume (many hundreds of deployed systems) GPS timing application.

INTRODUCTION

GPS increasingly provides the basis for aircraft navigation, surveillance, and air traffic control systems. Furthermore, many new aviation operations which improve efficiency and capacity will depend on the capabilities provided by GPS. However, when GPS is unavailable, these services no longer function as initially anticipated or planned. Therefore, a high-performance yet cost-effective means of aviation alternate navigation is critical to the continued safe and efficient utilization of the nation's airspace in the event of primary guidance system unavailability.

One of the most likely threat scenarios leading to loss of primary GPS guidance is interference, whether unintentional radiation or deliberate jamming. In order for an aviation alternate navigation system to prevail in such conditions, it must be resistant to that failure mode.

This was one of the beneficial factors previously cited in the consideration of ground-based radionavigation systems for aviation alternate navigation use.

Based on the above considerations, our research team is working with the FAA and other collaborators to define an aviation alternate navigation system which meets requirements for high-performance, cost-effectiveness, and resistance to jamming that can deny access to the primary GPS utility [1]. Candidate architectures include systems which utilize extant or planned aviation infrastructure such as distance measuring equipment (DME) augmented with other ground-based transmitters (GBTs). Using such infrastructure is particularly attractive from a cost perspective. However, system design trade studies indicate that precise time synchronization between these ground stations will have many performance benefits, primarily in terms of capacity, coverage, and availability. The time synchronization budget, on the order of 50ns between widely-separated ground facilities, either requires dedicated and expensive ground-based wired network infrastructure or would make use of satellite-based time transfer methods.

Geostationary satellites of the Wide Area Augmentation System (WAAS) offer a novel, robust, and cost-effective means of synchronizing time at widely-separated ground facilities, to levels of ~50ns, without the need for dedicated long-distance wired communication networks (Figure 1). WAAS satellites have been utilized in previous studies for highly-accurate two-way time transfer and synchronization [2, 3]. However, reliance on satellite-based signals for aviation alternate navigation, without explicit hardening against the interference threat mechanisms which could disrupt or deny GPS, would leave the alternate system susceptible to the same vulnerabilities as the primary service.

High-gain mechanically-steered directional dish antennas have been employed in previous research [4], but these systems are undesirable in the present context for a number of reasons. First, a directional dish requires precise orientation and installation in order to point to the desired geosynchronous location, adding expense to the siting and build-out plan. Second, a directional dish needs a steering mechanism to account for geo-satellite slot changes or satellite change-over, with the steering sub-system potentially impacting reliability. Third, to track more than one geosynchronous satellite requires additional directional antennas (one per tracked satellite signal), which limits scalability and reduces geometric diversity. And fourth, while the side-lobes of a directional dish have much lower gain than the main

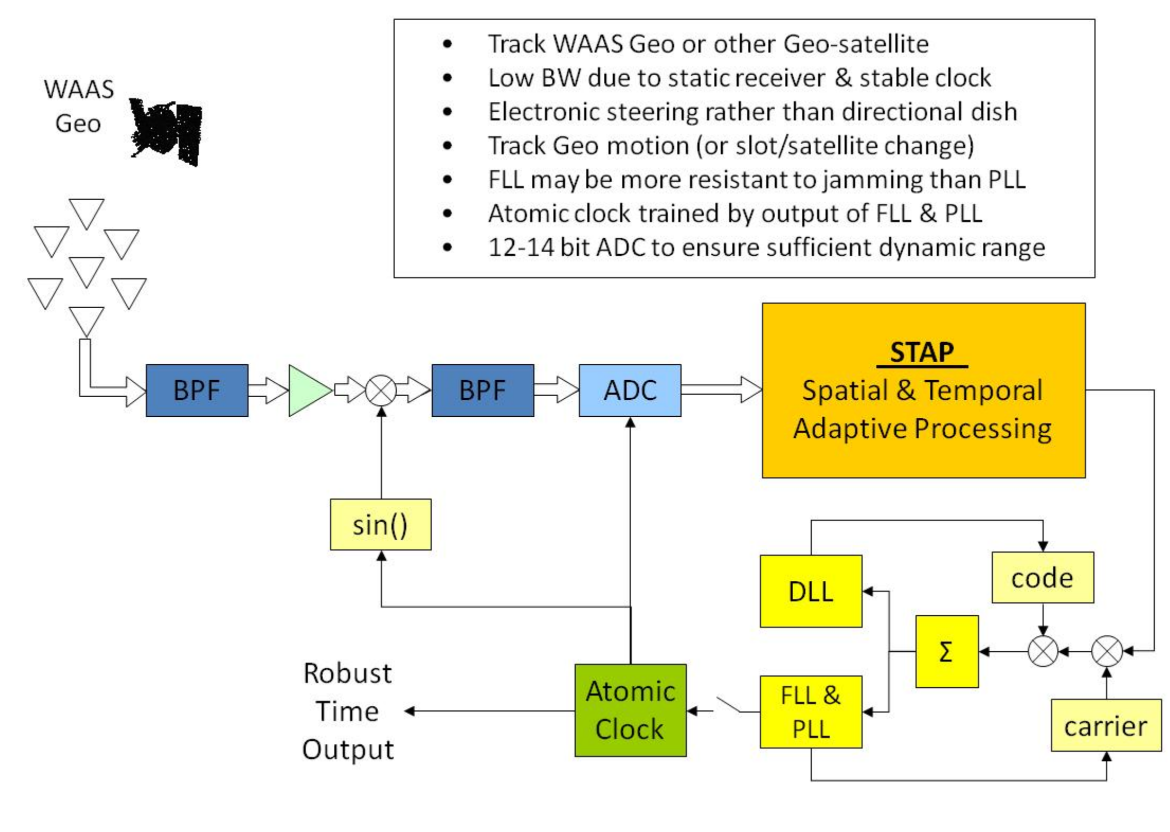

Figure 1. Time Transfer & Synchronization with the WAAS/L5 Signal.

beam, side-lobe directions are uncontrolled and thus potentially point in an undesirable direction vis-à-vis a terrestrial jamming source.

The primary benefits to using WAAS for time synchronization in the manner described by this research include balancing the requirements for: (1) robustness against inadvertent or deliberate radio-frequency interference (RFI), (2) leveraging existing and planned aviation ground infrastructure facilities, (3) reasonable cost targets for deployment and on-going maintenance, and (4) time synchronization accuracy sufficient to support aviation alternate navigation requirements. The primary innovations to be discussed in this paper are: (1) adaptive electronically-steered multi-element antenna arrays and signal processing strategies for RFI mitigation, (2) live signal, synthetic interference, and hardware-in-the-loop testing of jammer cancellation algorithms, and (3) preparations for over-the-air interference tests which will probe the effects of front-end saturation on digital beamsteering performance.

It should be emphasized that we describe a hardware system assembled from readily-available commercial building blocks (data acquisition system, antennas, etc.), and a critical goal of this research is to realize significant GPS anti-jam performance in an open-architecture (or non-defense-related) platform. Therefore, the central innovations in this research enable adaptive electronic beamsteering with high-dynamic-range signals (14-bit I/Q digitization), employing commercial off-the-shelf (COTS) hardware and computer systems, and targeted for a civilian high-reliability, high-volume (many hundreds of deployed systems) GPS timing application.

The remainder of this paper is organized as follows. We first define the interference threat scenarios that motivate this research, and briefly review adaptive processing for interference cancellation. Next we discuss the multi-element antenna array and the high-fidelity data acquisition testbed utilized for live signal testing. Then we describe system architecture and signal processing algorithms used to cancel incident RFI, which allows a tracking receiver to remain locked on the desired satellite signal(s). Finally we preview the live over-the-air jamming tests which will allow quantification of actual anti-jam performance improvement, and introduce our migration path to a real-time signal processing capability.

INTERFERENCE THREAT SCENARIOS

Interference threats to the continued aviation use of GPS generally can be classed into several categories:

- Scheduled outages, such Department of Defense testing as described and promulgated in NOTAMs (notices to airmen)
- Unintentional outages due to anomalous events, such as equipment malfunction or spurious out-of-band emissions which overlap onto aeronautical-protected GPS/GNSS spectrum
- Short-range jamming from low-power portable GPS denial devices (the so-called “personal privacy” market)
- Intentional jamming that deliberately seeks to deny GPS/GNSS services over a large geographic area, whether targeted against the aviation user or for which aviation is a collateral victim

Furthermore, the interference can be classified as tone jamming (CW (continuous-wave) interference), pulsed interference, or wideband (e.g., white-noise interference, jamming containing a code-division multiple-access (CDMA) modulation overlay, or more sophisticated jamming/spoofing waveforms). In general, a frequency-selective notch filter or frequency-domain adaptive processing (FDAP) can be an effective remedy against CW jamming. Pulsed interference is a likely candidate for time-domain blanking; for example, this can be an effective remedy against DME interference in the L5 band. Therefore, the greatest interference challenge to time synchronization in the aviation alternate navigation service is likely to be wideband or CDMA jamming, which motivates our consideration of adaptive electronically-steered antenna arrays for mitigation.

With this threat scenario defined, the goals for the interference mitigation sub-system become:

1. Remain impervious to low-power or inadvertent RFI (the first three bullets in the above list);
2. Maximize resistance to deliberate in-band jamming (the fourth bulleted item); thereby
3. Make the jammer radiate enough energy to enable localization and mitigation; or
4. Require that the jammer turns off (ceases radiating) in order to remain covert, in which case the time synchronization system re-acquires the satellite signal

These operational modes are summarized in Figure 2.

ADAPTIVE PROCESSING OVERVIEW

There is a broad range of interference mitigation techniques described in the GPS literature. These techniques include filtering and signal processing to reduce out-of-band and in-band interference power, automatic gain control (AGC) to exploit quantization effects of the analog-to-digital conversion process, antenna designs that suppress low-elevation signals or

Condition	Response	Duration
Primary GPS Nav. Service	Robust time system in standby mode Training of clock drift parameters Steady-state antenna operation	Indefinite
Primary GPS Nav. Service Denied (e.g., inadvertent or deliberate jamming)	Robust time system rejects interference Maintains stable time sync (e.g. WAAS L5) Provides time to Alternate Nav. Service	Indefinite
APNT WAAS/L5 Time Sync Denied (e.g., regional service disruption or high-power deliberate jamming)	Robust time system flywheels on atomic timekeeping standard Continues to provide time to Alternate Nav. Service Awaits opportunity to re-sync to WAAS/L5	Limited e.g., 30 min. with timing accuracy better than 50ns

Figure 2. Time Transfer Operational Modes.

greatly enhance pattern directionality, and changes to the phase-lock and delay-lock loops to improve tracking robustness, including augmentation by inertial measurement units and vector processing in the delay-lock loops [5]. Electronically-steered controlled-reception pattern antenna (CRPA) arrays are among the most aggressive anti-jam technologies, and CRPA arrays implementing space-time adaptive processing (STAP) can be extremely effective in mitigating radio frequency interference (RFI) [6, 7]. CRPA arrays can increase the signal to interference plus noise ratio (SINR) in two ways:

1. *Beamsteering*: enhancing the array gain in the direction of GPS satellites, even as the antenna and satellites move (reorient) relative to each other
2. *Nullsteering*: attenuating interference signals that arrive from directions other than those of the desired GPS satellite signals

The characteristics of the CRPA beam pattern are determined electronically by a set of weighting coefficients that may be computed either deterministically or adaptively. A deterministic CRPA increases gain in a desired look direction (or directions, if multiple signals are being tracked by the same beamformer); however the sidelobes and nulls of the array are uncontrolled and may not adequately suppress interference. An additional step of interference detection and localization can produce nullsteering constraints which significantly reduce the array gain in the direction of undesired signals. In practice, this method of nullsteering has difficulties because performance falls off dramatically with only small errors in interference localization [8].

Adaptive algorithms control sidelobes and steer nulls without the same sensitivity to small interference localization errors. Adaptive array processing increases the SINR by using feedback to optimize some characteristic of the array output. Suppression of narrowband or continuous-wave (CW) interference can be achieved by adaptive spatial filtering. Greater interference rejection (particularly of multiple, high-power, or wideband sources) can be realized by incorporating temporal filtering as well, for example with a tapped-delay-line antenna array.

"Adaptive" in this context means that the array gain pattern *automatically* adjusts to the signal and noise environment, subject to user-specified constraints. The constraint or optimization criteria broadly can be classified either as maximizing the signal to interference plus noise ratio (SINR) at the array output [9, 10], or as minimizing the mean-square error (MSE) between the actual array output and the ideal array output [11]. In both of these cases, the array weights adapt to maximize the desired signal and to reject interference.

For this investigation, the Applebaum beamformer [10], or minimum-variance distortionless-response (MVDR) array, is studied; the MVDR array is in the SINR class of methods. This algorithm constrains to unity the array gain in a particular look direction (it also may have side

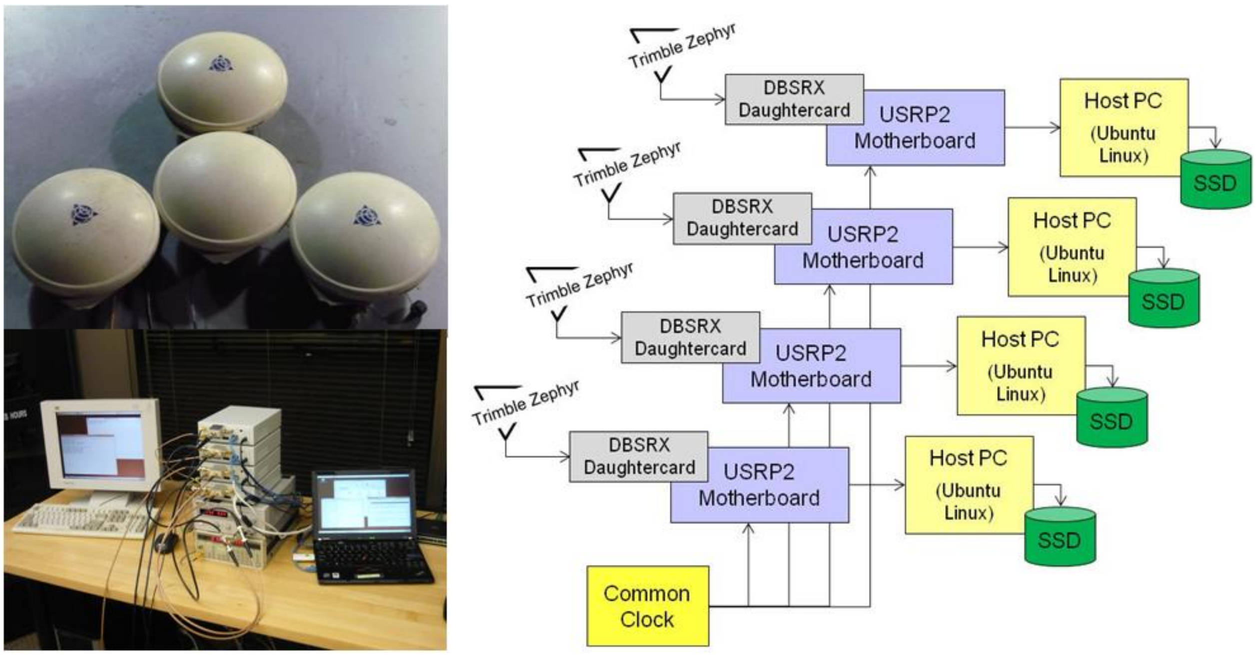

Figure 3. Four-Antenna Data Collection System for L1 and L5 Signals.

constraints for nullsteering), while rejecting coherent interference down to the noise floor. In the interference-free limit, the MVDR beamsteering constraint is equivalent to the weight vector calculated for the deterministic CRPA.

TEST HARDWARE DESCRIPTION

There are four identical signal conditioning and data acquisition systems, comprising in aggregate the 4-element antenna array testbed (Figure 3). The individual data acquisition systems are frequency synchronized with a common 10 MHz external clock source; this is the only physical interconnect between antenna channels.

Each data acquisition system begins with a Trimble L-band Zephyr antenna [12], whose signal passes to a USRP2 software-configurable FGPA-based data acquisition board equipped with a DBSRX programmable mixing and down-conversion daughtercard [13]. (It should be mentioned that mixing to baseband (i.e., 0 Hz intermediate frequency), is undesirable, as there is minor instability in USRP2 clock steering when the mixing frequency is equal to the programmable filter center frequency.) Samples are digitized at 14-bit resolution with complex I/Q samples, and these are sent as 16-bit data words across a gigabit ethernet cable to a host computer. The host computer runs the Ubuntu 9.04 distribution of Linux [14], and the open-source GNU Radio software-defined radio module is used as the configuration and data logging interface [15].

Data are streamed either to system memory (for short data captures) or to a solid-state hard-drive. For lower sampling frequencies, such as for GPS L1-C/A signals sampled at (approximately) 2-6+ Msamples/second, a traditional rotating hard-drive is acceptable. But for higher sampling frequencies required for modernized signals with greater signal bandwidths, sampling at 20-25 Msamples/second overloads the write speed of a rotating hard-drive, and solid-state storage is necessary. Due to load-balancing firmware in current-generation solid-state drives, write speeds can be variable (and degrade) over drive lifetime, and re-initialization to the as-new configuration (i.e., zero-hour state) can help to maintain initial manufacturer-specified write performance.

Note that the operating system is installed and boots from an external USB "thumb-drive", so that the data capture disk is not tasked with operating system file access or cache read/writes. Also, configuration testing indicated that using an on-board gigabit ethernet adapter, rather than a USB-based or PCI-based add-in adapter, was necessary for driver installation compatibility with the GNU Radio software.

The downconversion frequency synthesizers on each USRP2/DBSRX system are synchronized to a common external 10 MHz clock. This ensures precise frequency stability between data acquisition systems. However, while sample clock *rates* are synchronized, the actual *time instant* of sample digitization is not synchronized. Thus, an initial data pre-processing step is to reduce the sample

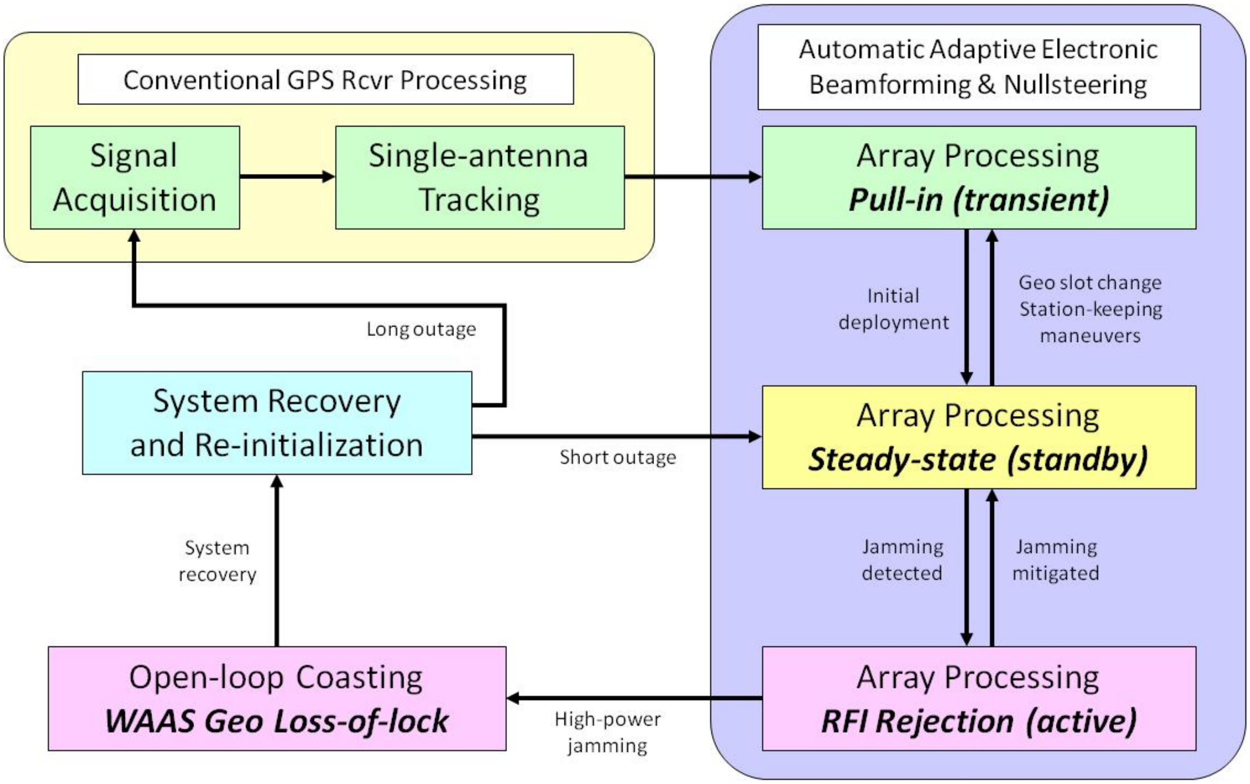

Figure 4. Robust Timekeeping Receiver: Operational Modes & State Transitions.

misalignment between captured signal files to sub-sample level. This pre-processing step first aligns to millisecond level by matching navigation data messages across signal recordings, and next aligns to sub-sample level by precise code-phase alignment within the CDMA sequence. For this short-baseline antenna system, code-phase skew between array elements due to satellite and array baseline geometry is negligible and can be completely disregarded.

The outcome of the data acquisition hardware system and the data pre-processing steps is four data files (for this four-antenna system), one for each antenna element, with data record alignment to sub-sample level. Note that additional antenna elements can easily be incorporated into this hardware platform, simply by splitting the 10 MHz clock synchronization signal to additional data acquisition systems. The primary reason for limiting this feasibility study to a four-element antenna array was for cost concerns.

SIGNAL PROCESSING ARCHITECTURE

The signal processing architectural paradigm (Figure 4) is to perform initial satellite acquisition (i.e., detection) and tracking in single-antenna mode, and then to transition to array processing mode. It should be noted that the interference threat scenario for satellite-based time synchronization means that operationally this system would expect to spend the majority of its mission duration in un-jammed (i.e., interference-free) conditions. Array processing mode begins with phase calibration between the multiple antennas in the array for each tracked satellite of interest. In the case of a geostationary satellite, this phase calibration is static, to the phase stability of the mixers and sampling sub-systems in the analog hardware. For a low-earth orbiting (LEO) or medium-earth orbiting (MEO) satellite, the phase calibration relationship can be measured in real-time for as long as the several single-antenna channels maintain carrier-phase lock on a satellite; if interference causes the single-antenna channels to lose carrier-phase lock, then the carrier-phase alignment can be derived from a stored database based on previous tracking results (since the satellite ground tracks repeat on a known/scheduled basis).

An alternative method of phase alignment is to utilize array synthesis techniques that consider baseline geometry, array orientation, satellite constellation, line biases, filter delays, etc. This is considerably more difficult, in our experience, than using spare receiver tracking channels to derive the phase differences between the satellite signals across antenna elements. Furthermore, antenna anisotropy and mutual electronic coupling will impact array response and cause the optimal weight coefficients to differ from those calculated by purely geometrical considerations [16, 17, 18].

Once the phase relationships between the antenna elements for a particular satellite are known, then these phases can be used as the steering constraint in either deterministic or (preferably) adaptive beamforming. (For the entire extent of un-jammed operation, the phase steering constraints can be observed and managed in real-time according to the single-antenna signal tracking channels.) It is expected that the system will spend the vast majority of its operational life in un-jammed conditions, with occasional transition to active anti-jam mode. When interference occurs, then there is seamless transition to active interference cancellation. When jamming ceases, then the system transitions back into the standby state, ready again for the next interference event. In the event that jamming is so powerful that the adaptive antenna system cannot attenuate the interference and the satellite tracking channel loses phase lock, then the system transitions to open-loop operation of the high-precision system clock (meaning that clock-steering parameters are no longer updated in real-time). At the cessation of jamming, the system enters a recovery and re-initialization mode. If the loss-of-lock was of short duration (i.e., tracking phase continuity persists), then the system can go back to steady-state (standby) mode; for longer outages, the system may have to re-acquire code-phase and/or carrier-phase/Doppler-frequency alignment to the satellite (i.e., re-entry of system initialization mode).

Actual mechanization of initialization mode, or really of any single-antenna operational tracking of a satellite, is achieved by setting to zero the antenna weighting vector entries for all antennas but the antenna of interest, i.e., tracking for antenna i is accomplished by the following:

$$\vec{W} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \\ \vdots \\ w_N \end{bmatrix},\ with\ w_{n,n\in 1,\ldots N} = \begin{bmatrix} 1\ for\ n = i \\ 0\ for\ n \neq i \end{bmatrix}$$

Phase alignment across antenna elements for a particular satellite is found by comparing the carrier-phase difference between tracking channels (Figure 5). The carrier-phase difference gives the phase rotation required to align the signals such that constructive interference will enhance signal power for this satellite, e.g., for antenna #2 with respect to antenna #1:

$$\Delta\phi_2 = \phi_2 - \phi_1$$

This carrier-phase difference forms the beamsteering constraint for deterministic or adaptive array processing, with no requirement for resolving array geometry, satellite constellation, or other array synthesis quantities.

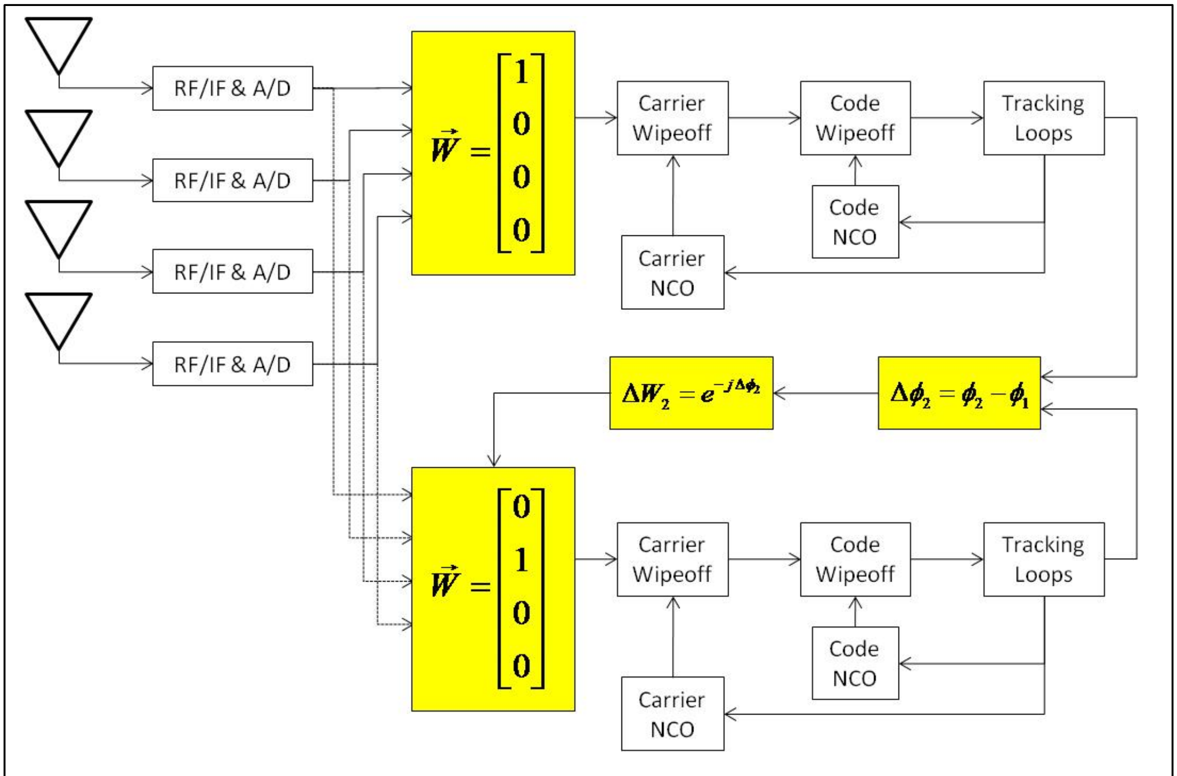

Figure 5. Acquisition, Initial Tracking, and Inter-channel Phase Alignment.

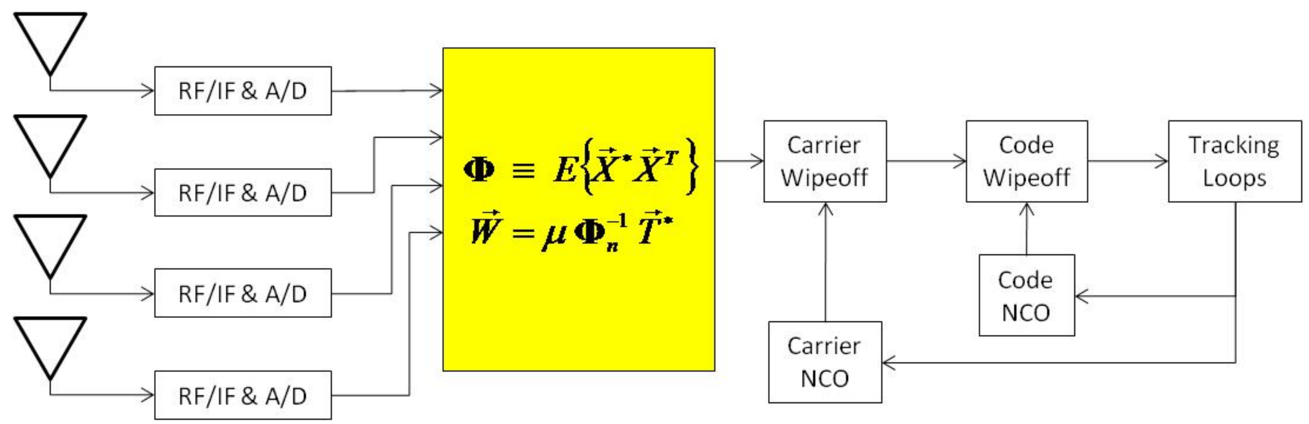

Figure 6. MVDR Adaptive Beamsteering for RFI Rejection.

Adaptive beamsteering is currently mechanized as minimum-variance distortionless-response (MVDR) processing [10]. To improve implementation efficiency with a software-defined radio (and future GPU-based baseband processing architectures), the beamsteering weight vector is calculated using the inverse of the signal covariance matrix, the Sample Matrix Inverse (SMI) Method (Figure 6). (Previous implementations have utilized an iterative approach to weight vector update [19, 20] – however, this will not be covered here.) The signal covariance matrix is defined as the expected value of $\vec{X}^{*}\vec{X}^{T}$, and is estimated accordingly on one-millisecond buffers of input data. (Note: the '*' operator here indicates conjugation.) Then, the optimal weight vector is estimated as follows:

$$\vec{W} = \mu\mathbf{\Phi}^{-1}\vec{T}^{*}$$

Since the covariance matrix is only an estimate of the true value, likewise the weight vector is only an estimate of the ideal filter for cancelling interference. Therefore, smoothing of the weight vector estimate is highly desirable. In general, the speed of response of the adaptive weight vector should be faster than the bandwidth of the carrier tracking loop, but no faster than is absolutely necessary. Specifically, one wants to reject interference before it can contaminate the carrier tracking loop and drive the carrier numerically-controlled oscillator (NCO) off frequency – this leads to an undesirable cycle slip; but one does not want too much noise in the weight vector – remember, the weight vector is only an estimate of the ideal Weiner filter response based on imperfect input data (i.e., the inverse of the signal covariance matrix), not the actual (hypothetical) filter that would perfectly cancel RFI. In short, smoothing is a tradeoff between response speed and limiting "gradient noise" in adaptive weight vector convergence behavior [7].

RESULTS WITH LIVE SIGNALS & SYNTHETIC RFI

The 4-antenna data collection system was installed on the roof of the Durand Building on the Stanford University campus, and data in the L1-band and the L5-band were sampled. (In-field stationary tests also were conducted, but those are not discussed here.) Several minutes of data were captured in each frequency band. The signals were tracked for each visible satellite, and in this way phase calibration alignment was achieved for these satellites.

Interference was injected synthetically by replaying the stored data records and combining with digital interference files (Figure 7). A reference satellite direction was selected for the phase alignment of the interference signal, and CDMA jamming signals synthesized accordingly; for L1 with a chipping rate of 1.023 Mchips/sec utilizing unused PRN code 168 and for L5 with a chipping rate of 10.23 Mchips/sec utilizing unused PRN code 122.

CDMA Interference Parameters	L1-band	L5-band
Center frequency	1575.42 MHz	1176.45 MHz
Code chipping rate	1.023 Mchips/sec	10.23 Mchips/sec
PRN code	168	122

The interference signal is amplified to achieve a desired jamming-to-signal (J/S) power ratio, and combined with the stored L1 or L5 signal records. Adaptive beamsteering is implemented as a stand-along module, or appliqué, following which the composite signal goes to the satellite tracking channels in a software receiver.

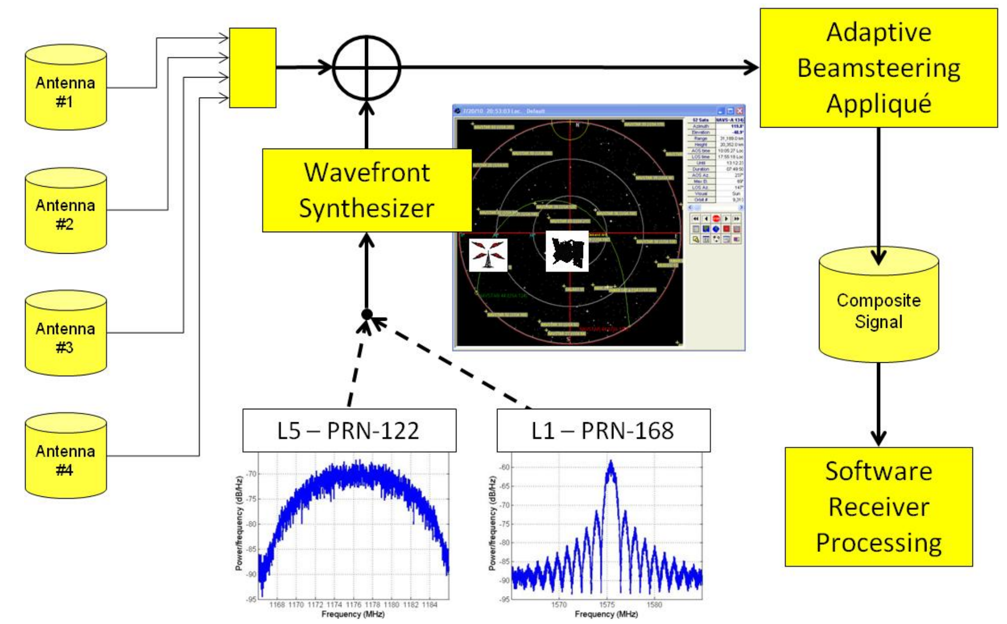

Figure 7. Archived Signals, Synthetic Jamming Overlay, and Beamsteering Appliqué.

At a low J/S power ratio, even a non-hardened single-antenna GPS receiver can reject interference – e.g., a J/S power ratio of 25 dB is rejected easily by such a receiver. As jamming power increases, at some point the non-hardened single-antenna GPS receiver loses lock on the satellite. Even a multiple antenna beamsteering receiver will lose lock, depending on the interference signal direction-of-arrival, since the sidelobes of the antenna pattern are uncontrolled. Our live-signal testing with synthetic interference shows that the multi-element adaptive antenna system rejects moderate to high levels of interference. As shown in Figure 8, a modest 45 dB J/S power ratio is easily tolerated, and preliminary tests (live signals and synthetic interference) indicate that considerably greater J/S power ratios are achievable. The issue remains to what extent saturation effects in the analog front-end impact and degrade anti-jam beamsteering performance.

CONCLUSIONS AND NEXT STEPS

We have described a research testbed and signal processing algorithms for GPS anti-jam in support of aviation alternate navigation. Specifically, this system enables satellite time synchronization between widely-separated ground stations in the event of GPS unavailability due to radio frequency interference (RFI). This system has several features which distinguish it from other GPS anti-jam applications, which we briefly summarize below:

- Stationary receiver and virtually stationary transmitter. This means that tracking loops only need to track receiver clock dynamics. As the system will employ a high-quality atomic clock with good long-term stability characteristics (for RFI ride-through), this allows selection of extremely low tracking loop bandwidths
- Aviation alternate navigation requires a time synchronization error budget of approximately 15m (50ns) code-phase ranging accuracy. This means that the antenna array system is essentially insensitive to array biases and mutual coupling effects that dominate at the sub-meter level required for high-precision applications (e.g., carrier-phase differential navigation).
- The current generation of software receivers supports dozens of real-time tracking channels. Our current implementation allocates spare channels to calibrate the beamsteering constraint vector, rather than utilizing array synthesis techniques. This improvement in algorithmic simplicity greatly enhances robustness and performance validation/verification.
- The data acquisition system employs 14-bit analog-to-digital converter to maximize signal dynamic range. This means that saturation and analog effects, which are going to be limiting factors in system anti-jam performance, can be isolated and studied in detail.

Only brief attention is devoted in this paper to anti-jam results, as *synthetically-injected* RFI bypasses one of the fundamental limiting mechanisms in adaptive array interference rejection performance. Namely, saturation effects in the analog front-end will constrain real-world

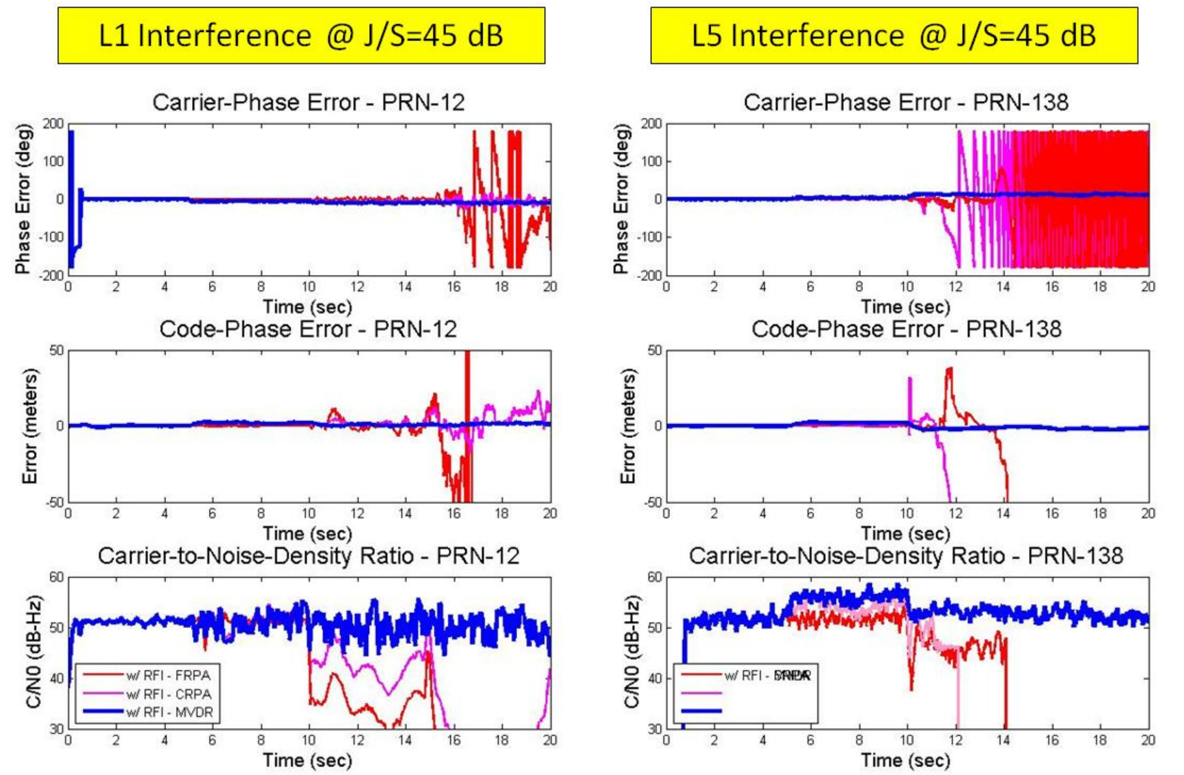

Figure 8. High-power Simulated RFI: Adaptive Beamsteering Predicted to Reject RFI. When do predictions diverge from actual system performance??

effectiveness, such that analytical and numerical studies neglecting these effects can be of somewhat modest utility. This motivates the current research program, which has completed to-date the definition and construction of an appropriate hardware data acquisition testbed and development and implementation of signal processing algorithms, including adaptive beamsteering mechanization. All hardware and computer systems are readily available from commercial vendors, which enhances the scalability and cost-effectiveness of the final solution. At the current time, our focus is planning for a live over-the-air interference test campaign (NAVFEST/2011) and migration to real-time signal processing capability. We expect to report on this research in future articles.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the FAA on cooperative agreement 2008-G-007, and of Stuart Riley and Trimble Navigation for the loan of GPS antennas. However, the opinions and findings expressed herein are solely those of the authors, and do not necessarily represent those of the Federal Aviation Administration, Trimble Navigation, or of other affiliated agencies and companies.

REFERENCES

1. M. Harrison, R. Kenagy, R. Lilly, L. Eldredge, M. Narins, D. Chapman, R. Loh, V. Massimini, R. Niles, P. Enge, and S. Lo, 2010, "Alternative Positioning, Navigation, and Timing (APNT) Services Research Development Program," *Proc. ION GNSS 2010*, in press.

2. W.J. Klepczynski, 1998, "The Time Distribution System for the Wide Area Augmentation System (WAAS)," *Proc. 30th Annual Precise Time and Time Interval (PTTI) Meeting*, pgs. 439-444.

3. P. Fenton, B. Klepczynski, E. Powers, and R. Douglas, 1999, "Time Transfer Using WAAS: An Initial Attempt," *Proc. 31st Annual Precise Time and Time Interval (PTTI) Meeting*, pgs. 191-202.

4. P. Fenton, E. Powers, B. Klepczynski, W. Lewandowski, and R. Douglas, 2001, "Capabilities of the WAAS and EGNOS using High Gain Antennas for Time Distribution and Frequency Transfer," *Proc. ION GPS 2001*, pgs. 2214-2219.

5. J.J. Spilker and F.D. Natali, 1996, "Interference Effects and Mitigation Techniques," in Global Positioning System: Theory and Applications I, B.W. Parkinson and J.J. Spilker (Eds.), AIAA, Reston, VA, pgs. 717-771.

6. R.T. Compton, Jr. 1988, Adaptive Antennas: Concepts and Performance, Prentice-Hall, Inc., Englewood Cliffs, N.J.

7. B. Widrow and S.D. Stearns, 1985, Adaptive Signal Processing, Prentice-Hall, Inc, Englewood Cliffs, N.J.

8. K.G. Gromov, 2002, "GIDL: Generalized Interference Detection and Localization System," PhD Dissertation, Stanford University.

9. O.L. Frost, 1972, "An Algorithm for Linearly Constrained Adaptive Array Processing," *Proc. IEEE*, vol. 60, no. 8, pgs. 926-935.

10. S.P. Applebaum, 1976, "Adaptive Arrays," *IEEE Trans. on Antennas and Propagation*, vol. 24, no. 5, pgs. 585-598.

11. B. Widrow, P.E. Mantey, L.J. Griffiths, and B.B. Goode, 1967, "Adaptive Antenna Systems," *Proc. IEEE*, vol. 55, no. 12, pgs. 2143-2159

12. Zephyr Geodetic Antennas, Trimble Navigation Inc., reachable on the web at http://www.trimble.com.

13. USRP2 motherboard and DBSRX programmable daughtercard, Ettus Research LLC, reachable on the web at http://www.ettus.com.

14. Ubuntu Linux, © Canonical Ltd., reachable on the web at http://www.ubuntu.com.

15. GNU Radio, © Free Software Foundation, reachable on the web at http://gnuradio.org.

16. G.F. Hatke, 1998, "Adaptive Array Processing for Wideband Nulling in GPS Systems," *Proc. 32nd Asilomar Conference on Signals, Systems, and Computers*, vol. 2, pgs. 1332-1336.

17. K.F. McDonald, R. Raghavan, and R. Fante, 2004, "Lessons Learned Through the Implementation of Space-Time Adaptive Processing Algorithms for GPS Reception in Jammed Environments," *Proc. IEEE PLANS*, pgs. 418-428.

18. U.S. Kim, D.S. De Lorenzo, J. Gautier, P. Enge, and J.A. Orr, 2004, "Phase Effects Analysis of Patch Antenna CRPAs for JPALS," Proc. ION GNSS 2004, pgs. 1531-1538.

19. D.S. De Lorenzo, J. Gautier, J. Rife, P.K. Enge, and D.M. Akos, 2005, "Adaptive Array Processing for GPS Interference Rejection," *Proc. ION GNSS 2005*, pgs. 618-627.

20. D.S. De Lorenzo, F. Antreich, H. Denks, A. Hornbostel, C. Weber, and P.K. Enge, "Testing of Adaptive Beamsteering for Interference Rejection in GNSS Receivers," *Proc. ENC 2007*, pgs. 1277-1287.