

GPS WORLD

WWW.GPSWORLD.COM

GNSS
POSITION
NAVIGATION
TIMING

Innovation

Galileo
Cycle-Slip
Detection

**LOW-COST
CDGNSS FOR
AUTOMATED
VEHICLES**



HOME
22min

**2016 STATE OF THE
GNSS INDUSTRY**

INTERGEO PREVIEW

OLYMPIC GPS



LOW-COST PRECISE POSITIONING FOR AUTOMATED VEHICLES

Carrier-Phase Differential GNSS Produces Mass-Market Centimeter Accuracy

A DENSE REFERENCE NETWORK FACILITATES LOW-COST carrier-phase differential GNSS positioning with rapid integer-ambiguity resolution. This could enable precise lane-keeping for automated vehicles in all weather conditions.

BY Matthew J. Murrian, Collin W. Gonzalez, Todd E. Humphreys, Kenneth M. Pesyna, Jr, Daniel P. Shepard and Andrew J. Kerns

Strong demand for low-cost precise positioning exists in the mass market. Carrier-phase differential GNSS (CDGNSS) positioning, accurate to within a few centimeters even on a moving platform, would satisfy this demand were its cost significantly reduced. Low-cost

CDGNSS would be a key enabler for many demanding consumer applications.

Centimeter-accurate positioning by CDGNSS has been perfected over the past two decades for applications in geodesy, precision agriculture, surveying and machine control. But mass-market adoption of this technology will demand much lower user cost — by a factor of 10



to 100 — yet still require rapid and accurate position fixing. To reduce cost, mass-market CDGNSS-capable receivers will have to make do with inexpensive, low-quality antennas whose multipath rejection and phase center stability are inferior to those of antennas typically used for CDGNSS. Moreover, there will be a strong incentive to use single-frequency receivers, whereas almost all receivers used for CDGNSS in surveying and similar applications are multi-frequency. Despite these user-side disadvantages, mass-market precise positioning will be expected to demonstrate convergence and accuracy performance rivaling that of the most demanding current precise positioning applications: Users will be dissatisfied with techniques requiring more than a few tens of seconds to converge to a reliable sub-decimeter solution.

Meeting this challenge calls for innovation targeting both the rover (user) equipment and the reference network. Here we examine the challenge from the point of view of the reference network and offer demonstration results for a low-cost end-to-end system.

The recent trend in precise satellite-based positioning has been toward precise point positioning (PPP), whose primary virtue is the sparsity of its reference network. But standard PPP requires several tens of minutes or more to converge to a sub-10-centimeter 95 percent horizontal accuracy. Faster convergence can be achieved by recasting the PPP problem as one of relative positioning, thereby exposing integer ambiguities to the end user. This technique, known as PPP-RTK or PPP-AR, is mathematically similar to traditional network real-time kinematic (NRTK) positioning. As the network density is increased, sub-minute or even

instantaneous convergence is possible with dual-frequency high-quality receivers. Even single-frequency PPP-RTK is possible, with convergence times of approximately 5 minutes for a 40-kilometer network spacing.

For PPP-RTK and NRTK, convergence time is synonymous with the time required to resolve the integer ambiguities that arise in double-difference (DD) carrier-phase measurements, referred to here as time to ambiguity resolution, or TAR. As reference networks become denser, they can better compensate for spatially-correlated variations in signal delay introduced by irregularities in the ionosphere and, to a lesser extent, in the neutral atmosphere. Improvement is manifest as reduced uncertainty in the atmospheric corrections that the network sends to the user. Reduced uncertainty in the atmospheric corrections is key to reducing TAR.

Prior work has established an analytical connection between uncertainty in the ionospheric corrections (denoted σ_i) and TAR. The existing literature does not, however, offer a satisfactory model for the dependence of σ_i on network density. The prevailing model is based on single-baseline CDGNSS, which is inapt for PPP-RTK and NRTK. Moreover, prior work does not address the effect of network-side multipath on the accuracy of the corrections data, which becomes increasingly important as low-cost and poorly-sited reference stations are used to densify the network.

Here, we examine the relationship between ionospheric uncertainty and probability of correct ambiguity resolution, and present the results of an empirical investigation of the relationship between network density and the total uncertainty in network correction data. The online version of this article (gpsworld.com/murrian)

develops a simple analytical model relating error variance in network corrections to network density. Our analysis and experiments indicate that for rapid TAR in challenging urban environments with low-cost receivers, network density must be significantly increased. We report on the design and deployment of a dense network in Austin, Texas, and demonstrate a new system that taps into the network to provide reliable vehicle lane-departure warning.

AMBIGUITY RESOLUTION

Reducing the ionospheric uncertainty σ_i allows a strong prior constraint to be applied in the ionosphere-weighted model, thereby increasing $P(\hat{z} = z)$, the probability that the estimated and true integer ambiguity vectors are equivalent. It is instructive to consider single-epoch ambiguity resolution (AR), for two reasons. First, for stationary users with low-cost equipment, multipath errors dominate in the carrier-phase measurement and are strongly correlated over 100 seconds or more. Thus, if single-epoch AR fails then a static user may have to wait an unacceptably long time for multipath errors to decorrelate enough to permit AR. In any case, single-epoch performance is a strong predictor of multi-epoch performance over an interval short enough (a few tens of seconds) to satisfy impatient mass-market users.

Second, a convenient and accurate analytical model (by Dennis Odijk and PJG Teunissen) for single-epoch AR reveals the dependency of $P(\hat{z} = z)$ on scenario parameters of practical interest: the standard deviation of ionospheric correction errors, the number of visible satellites, the standard deviation of undifferenced carrier- and code-phase measurement errors (including multipath-induced errors), a satellite geometry factor,




FIGURE 1 Single-epoch single-frequency ambiguity fixing. Blue traces (left axis) indicate the probability $P(\hat{z}=z)$ of correctly resolving all integer ambiguities with a single epoch of data as a function of the number of satellites m . Each trace represents $P(\hat{z}=z)$ for a different value of ionospheric uncertainty σ_i . Green bars (right axis) represent the probability mass function $P(m)$ for the number of satellites above an elevation mask angle of 15 degrees, assuming 31 GPS, 14 Galileo, and 3 WAAS satellites (projected mid 2017). Each blue trace is marked with the total probability of correct integer resolution P_T , a function of both the trace itself and $P(m)$. Other parameters of the scenario: geometry factor $f = 2.5$, standard deviation of undifferenced phase measurements $\sigma_\phi = 3\text{mm}$, standard deviation of undifferenced pseudorange measurements $\sigma_\rho = 50\text{cm}$, and number of estimated parameters $p=3$.




FIGURE 2 Total probability of a correct fix for the scenario of Figure 1 as a function of ionospheric uncertainty σ_i .

the number p of free parameters to be estimated ($p=3$ for negligible tropospheric error, $p=4$ to estimate a single additional tropospheric parameter), and the number of carrier frequencies broadcast by each of the satellites (1, 2, or 3) along with each carrier's wavelength. The model is highly accurate for single-epoch AR, but only approximate for multiple epochs, with accuracy degrading as the data interval lengthens. The model's inaccuracy results from its assumption that overhead satellites remain static from epoch to epoch, which yields pessimistic results for even fairly short data capture intervals (for example, 30 seconds). Fully accounting for satellite motion in an analytical model for $P(\hat{z}=z)$ is an open problem, which is why studies that wish to account for satellite motion resort to simulation.

FIGURES 1 and 2 show single-epoch, single-frequency results from the analytical $P(\hat{z}=z)$ model for parameters




FIGURE 3 As Figure 1 except for dual-frequency (L1-L2) measurements and the probability mass function $P(m)$ corresponds only to a constellation of 31 GPS satellites. The elevation mask angle is again taken to be 15 degrees. It is assumed that dual-frequency measurements can be obtained from every GPS satellite.




FIGURE 4 Total probability of a correct fix for the scenario of Figure 3 as a function of ionospheric uncertainty σ_i .

approximately reflecting the mass-market use case. The most important conclusion to draw from these figures is that for single-epoch, single-frequency AR to be even moderately reliable ($P_T \geq 0.9$) over the next few years, the ionospheric uncertainty σ_i must be held under 2 millimeters. This will relax somewhat as more Galileo and MEO BeiDou satellites come online, but signal blockage in built-up areas will raise the effective elevation mask angle significantly above the 15 degrees assumed here, reducing the number of available satellites. Thus, sub-2-mm ionospheric uncertainty remains desirable for urban environments even as GNSS constellations become fully populated.

FIGURES 3 and 4 offer results for a dual-frequency (L1-L2) single-epoch scenario. All other scenario parameters are held as for the single-frequency scenario except that, in an attempt to be somewhat more pessimistic, $P(m)$ is based only on GPS satellites. It is assumed that from each satellite the user can extract dual-frequency measurements. As with the single-frequency case, it is evident that dual-frequency P_T is strongly dependent on σ_i . The dual-frequency case is more forgiving, but substantial performance improvement can still be had by reducing σ_i to under 2 mm.

Corrections Uncertainty and Network Density. This section is posted in the online version, gpsworld.com/murrian.




FIGURE 5 Standard deviation of uncorrected DD carrier-phase residuals versus baseline distance between each of the 22 reference stations and the master reference station.

ANALYSIS OF A DENSE REFERENCE NETWORK

We examined σ_i as a function of network density using data from several organizations providing GNSS reference station observations: National Geodetic Survey Continuously Operating Reference Stations, UNAVCO, and the California Real Time Network. This combination allowed analysis of a hypothetical reference network of 23 high-quality GNSS receivers with an overall network density of approximately 8 nodes/1,000 km², or an average inter-station spacing of 14 km. The sites for this network are located between Los Angeles and Pomona, CA.

DD carrier-phase observations from GPS L1 C/A signals spanning GPS weeks 1850 through 1859 were used for the analysis. A minimum satellite elevation mask was enforced at 20 degrees. Any satellite not above the elevation mask and providing carrier-phase observations at both the beginning and end of each processing window was excluded. A step size of 10 minutes was used. The longest available sub-window, meeting the above requirements and providing a minimum of 6 satellite vehicles (1 pivot satellite and 5 others), was selected for processing. To facilitate batch processing, integer ambiguities were assumed to be resolved correctly when the mean standard deviation of carrier-phase residuals for that solution was less than one quarter wavelength of the GPS L1 frequency. In application, this constraint resulted in rejecting only 0.6 percent of all solutions.

Network Corrections Estimation. See the online version for equations and discussion from this section.

Uncorrected Carrier-Phase Residuals. FIGURE 5 shows the expected values for DD carrier-phase residual standard deviations for all secondary stations, based on observations that were not corrected for atmospheric

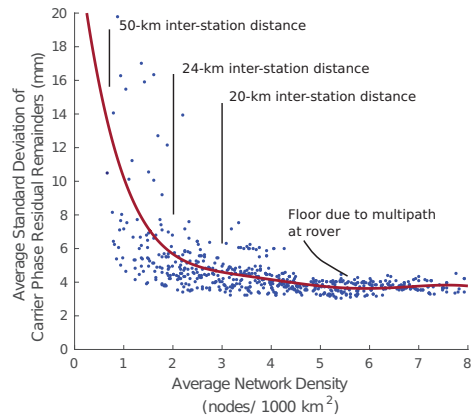


FIGURE 6 Standard deviation of carrier-phase residual remainders (the carrier-phase residuals which remain after application of network corrections) versus average network density. The fitted curve is simply a polynomial fit of these data; it is not based on any theoretically anticipated behavior.

delay. These data were produced by averaging the standard deviation of the DD carrier-phase residuals calculated at each epoch across all satellites present in the solution. The fitted curve indicates a linear growth of DD carrier-phase residuals with distance to the

i80

GNSS Receiver

E4.001@Hall A4
INTER GEO

Combining quality, support and leading-edge GNSS Performance

- The user-friendly LCD interface allows for common workflows without any external device.
- The ultra-rugged alloy cased i80 is the smallest and lightest receiver with dual hot swappable batteries, allowing continuous work and reduced user fatigue.
- Internal UHF, Cellular data modems, Wi-Fi, Bluetooth, Serial, and USB communication supports varied processes and any peripheral device.
- GPS, GLONASS, Galileo, Beidou, SBAS, QZSS; L1-L2-L2C-L5: every satellite and every signal tracked for the most reliable and accurate RTK solutions.

www.chcnave.com



Figure 6 shows that carrier-phase residuals after application of network corrections are considerably reduced



NETWORK DEPLOYMENT

FIGURE 11 GNSS antenna configuration. A single-baseline precise position solution between the primary antenna and the master reference station provides precise vehicle position. A constrained-baseline 2D attitude solution between the primary and secondary antennas provides heading.



Live Vehicle Demonstration. In partnership with Radiosense, an




FIGURE 9 Reference station components.

Austin-based precise positioning startup, we have developed and demonstrated a low-cost vehicle lane departure warning system that receives corrections from our dense reference network. The system takes in lane widths from an external database and infers a safe driving corridor within each lane by analyzing the behavior of human drivers on the same road. A vehicle's proximity to the lane boundary is displayed in real time to the driver and passengers.

For robustness against cycle slips and to provide a baseline against which to compare future improvements, the system currently employs single-epoch CDGNSS positioning without aiding from additional sensors. In choosing a single-epoch approach, the system naively discards information regarding the underlying integer ambiguities at the beginning of each measurement epoch. Still, the system performs well with the typical number of overhead signals in a light urban environment: correct and internally-validated solutions were available in over 92 percent of measurement epochs. When a second rover antenna is included to combat multipath with spatial diversity, this percentage improves to 96. Such good single-epoch performance suggests that, when armed with additional sensor aiding and proper integer ambiguity persistence, reliable and accurate vehicle positioning can be maintained in more challenging environments.

Demonstration Setup. The live demonstration followed a predetermined route in the vicinity of the University of Texas campus. The 1-mile route (**FIGURE 10**) passed through both open-sky and partially-blocked environments.

Prior to the demonstration, the vehicle was driven several times on the same route collecting GNSS measurements to precisely map typical driving trajectories on the route. The ensemble of trajectories was used to build a centimeter-accurate model of the lane center along the route. The sensing equipment employed during this mapping phase is no different than that used during the demonstration,




FIGURE 13 Vehicle position relative to lane edge (left) synchronized in time with video still frame (right), centered safely within the lane, as depicted by green rectangle.



FIGURE 14 Vehicle nearing lane edge, as depicted by yellow rectangle.




FIGURE 15 Vehicle crossing lane edge, as depicted by red rectangle.

making feasible eventual crowd-sourcing, wherein end-user vehicles generate and update the centerline models.

The demonstration vehicle was outfitted with two dual-frequency GNSS antennas mounted with magnetic bases onto the roof. The first antenna, designated primary, operated as the rover in a single-baseline CDGNSS solution against the master reference station of the Longhorn Reference Network, as illustrated in **FIGURE 11**. This baseline provided the geo-referenced, centimeter-accurate vehicle position. The other antenna, designated secondary, was paired with the primary antenna to produce a constrained-baseline CDGNSS solution providing sub-degree-accurate vehicle heading. The secondary antenna also served as a backup when the primary antenna produced a result that did not pass the precise positioning engine's internal validity testing.



The GNSS antennas were connected to a low-cost, dual-frequency front-end in the trunk of the vehicle (FIGURE 12) which downconverted and digitized the incoming signals and subsequently fed them to a low-cost single-board computer running the precise positioning engine. A cellular modem received real-time measurements from the master reference station, while a WiFi router streamed real-time solutions to several Android devices in the vehicle for real-time visualization of precise within-lane position.

Demonstration Results. FIGURES 13, 14 and 15 show snapshots of the Android application and a still frame of the side of the vehicle in three different scenarios. The large rectangle indicates vehicle position with respect to the modeled lane center, changing color from green, when the vehicle is within the safe driving corridor, to yellow as the vehicle nears the edge of the lane, and finally to red if the vehicle breaks the lane boundary. One could imagine wrapping a control loop around these signals to enable last-moment lane-keeping.

FIGURE 16 reveals the precision with which the positioning engine was able to locate the vehicle's driver-side antenna in four repeated passes along the test route. The variation between the four yellow traces is primarily due to driver non-repeatability; actual measurement precision is at the centimeter scale. A small bias in the traces' registration to the picture is present because Google Earth imagery is only registered to the International Terrestrial Reference Frame with meter-level accuracy.

FIGURE 17 shows a time history of the vertical deviation from the route mean, in meters. The zoomed view of the vertical deviation shown in FIGURE 18 allows one to appreciate the precision of the positioning engine: the vertical trajectory is smooth at the centimeter level. FIGURE 19 shows the DD residuals in carrier phase



FIGURE 16 Four repeated traces of driver's side antenna as vehicle made a turn.

and pseudorange for GPS PRN 30 during the four loops in Figure 17. One-sigma undifferenced phase and pseudorange deviations are 3.4 mm and 42 cm, respectively.

The figures demonstrate that the precise positioning engine fed by reference data from the Longhorn Reference Network maintained centimeter-accurate knowledge of the vehicle's position during almost the entire trajectory, despite passing between a large football stadium and parking garage, each of which introduced significant signal blockage and multipath.

For the data shown in Figure 17, 96 percent of the 3.3-Hz measurement epochs resulted in a correct and internally-validated positioning solution. The majority of the remaining solutions were correct but did not pass internal validation. For only 0.6 percent of solutions were the carrier-phase integer ambiguities resolved incorrectly, but all of these incorrect solutions were caught and excluded by the validation algorithm. Furthermore, the number of overhead signals during the time in which this particular dataset (set A) was taken was average, as seen in the upper plot of FIGURE 20. 16 signals above 15 degrees elevation were available during this time. In contrast, the number of overhead signals for a second dataset taken 8 days prior (set B) was much worse, with only 12 signals above 15 degrees elevation, as seen in the lower plot.




FIGURE 17 Time history of the vertical deviation from the route mean, in meters.




FIGURE 18 Zoomed view of the time history of the vertical deviation from the route mean, showing the centimeter-level precision in the 3.3 Hz positioning data.




FIGURE 19 Double-difference carrier phase (top) and pseudorange (bottom) residuals for GPS satellite 30 at frequency L1 over the full time interval shown in Figure 17.

For insight into the performance of the positioning engine as a function of the number of overhead satellites, TABLE I details the performance of these two datasets (as well as a third dataset) in terms of the percentage of epochs that passed the positioning engine's internal validation testing, based on a ratio test with a fixed threshold of 2.0. Results are shown for single- and dual-antenna positioning solutions and for dual-antenna vehicle heading solutions.




FIGURE 20 The number of signals above a 15-degree elevation mask. Each plot spans an entire day. The black arrows denote the time of day in which two different datasets, A and B, were taken. The dashed red line represents the mean number of signals above the mask over both days. Dataset A was taken during a nominal time when 16 signals were available, while dataset B was taken during a worst-case time when only 12 signals were available.

VALID SOLUTION PERCENTAGE PER DATASET				
Dataset	# Signals	Single Antenna	Dual Antenna	Heading
A	16	92%	96%	96%
B	12	41%	54%	94%
C	15	60%	71%	99%

TABLE 1 The performance of each dataset in terms of the percentage of solutions that passed validation testing.

A large dropoff in positioning performance occurs when the number of overhead signals is reduced below 16, while the constrained-baseline heading determination performance remains good throughout. Fortunately, it will not be long until even more signals are available. Within the next 8 months, the Galileo constellation will add six fully operational satellites. These will bring the number of GPS L1, GPS L2C, Galileo E1, and SBAS signals that are above 15 degrees elevation to 16 or more 95 percent of the time, enabling high-reliability single-epoch CDGNSS positioning.

CONCLUSIONS

For a sufficiently dense reference network, linear least squares estimation can be applied to the task of reducing uncertainties due to tropospheric and ionospheric

delays for the purposes of providing improved positioning accuracy as well as faster time to ambiguity resolution for carrier-phase differential positioning. High network density allows use of a strong linear model for atmospheric delays, which has the virtue of suppressing network-side multipath errors in the provided corrections.

A network of 23 high-quality reference stations in the vicinity of Los Angeles, California was studied to determine what network density is sufficient to make all network-side error sources negligible compared to rover receiver multipath. A density of three stations per 1,000km², or an average inter-station spacing of 20 km, was found to drive network-side ionospheric, tropospheric, and multipath errors well below rover receiver multipath.

These findings motivate a significant densification of permanent reference networks, at least in built-up areas where signal blockage and multipath are common, to support mass-market applications for which low user (rover receiver) cost and rapid convergence to a reliable sub-decimeter position are a priority. In a light urban setting, and with the kind of satellite coverage that will soon become the norm, we demonstrated vehicle lane departure warning in a field test that produced highly reliable instantaneous sub-decimeter positioning.

ACKNOWLEDGMENTS

This work was supported in part by Samsung Research America, by the Data-Supported Transportation Operations and Planning Center (D-STOP), a Tier 1 USDOT University Transportation Center, and by the Texas Department of Transportation under the Connected Vehicle Problems, Challenges and Major Technologies project. 🌐

MATTHEW J. MURRIAN served as a U.S. Navy submarine officer before joining the Radionavigation Laboratory at the University of Texas at Austin, where he obtained a Masters in aerospace engineering and is now pursuing a Ph.D.

COLLIN W. GONZALEZ received his B.S. in aerospace engineering from The University of Texas at Austin.

TODD E. HUMPHREYS is an associate professor at The University of Texas at Austin, Director of the UT Radionavigation Laboratory, and scientific advisor to Radiosense.

KENNETH M. PESYNA, JR. received his PhD from The University of Texas and is co-founder of Radiosense, which develops precise positioning systems for automated vehicles.

DANIEL P. SHEPARD is a Ph.D. candidate in the Radionavigation Laboratory at The University of Texas at Austin and is co-founder of Radiosense.

ANDREW J. KERNS received his M.S. in electrical engineering from The University of Texas at Austin and is co-founder of Radiosense.