



Technical Report 115

High-precision GPS Vehicle Tracking to Improve Safety

Research Supervisor
Todd E. Humphreys
Wireless Networking and Communications Group

September 2016

Data-Supported Transportation Operations & Planning Center (D-STOP)

A Tier 1 USDOT University Transportation Center at The University of Texas at Austin



D-STOP is a collaborative initiative by researchers at the Center for Transportation Research and the Wireless Networking and Communications Group at The University of Texas at Austin.

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated under the sponsorship of the U.S. Department of Transportation's University Transportation Centers Program, in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof.

Technical Report Documentation Page

1. Report No. D-STOP/2016/115		2. Government Accession No.		3. Recipient's Catalog No.	
4. Title and Subtitle High-precision GPS Vehicle Tracking to Improve Safety				5. Report Date September 2016	
				6. Performing Organization Code	
7. Author(s) Matthew J. Murrian, Collin W. Gonzalez, Todd E. Humphreys, Kenneth M. Pesyna, Jr., Daniel P. Shepard, and Andrew J. Kerns				8. Performing Organization Report No. Report 115	
9. Performing Organization Name and Address Data-Supported Transportation Operations & Planning Center (D-STOP) The University of Texas at Austin 1616 Guadalupe Street, Suite 4.202 Austin, Texas 78701				10. Work Unit No. (TRAIS)	
				11. Contract or Grant No. DTRT13-G-UTC58	
12. Sponsoring Agency Name and Address Data-Supported Transportation Operations & Planning Center (D-STOP) The University of Texas at Austin 1616 Guadalupe Street, Suite 4.202 Austin, Texas 78701				13. Type of Report and Period Covered	
				14. Sponsoring Agency Code	
15. Supplementary Notes Supported by a grant from the U.S. Department of Transportation, University Transportation Centers Program.					
16. Abstract Commercial Global Positioning System (GPS) devices are being used in transportation for applications including vehicle navigation, traffic monitoring, and tracking commercial and public transit vehicles. The current state-of-practice technology in GPS devices typically has 10-meter accuracy and can properly answer the needs in the above applications. When it comes to safety, where driver behavior is important, new technologies for high-precision (i.e., centimeter-level) mobility detection are required. It is with high-precision GPS traces that safety can be evaluated by identifying when drivers are drowsy or distracted, and anticipating problems before they occur. In this project, the researchers exploring the use of carrier-phase differential GNSS to produce 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.					
17. Key Words GPS, safety, GPS devices, centimeter-level GPS			18. Distribution Statement No restrictions. This document is available to the public through NTIS (http://www.ntis.gov): National Technical Information Service 5285 Port Royal Road Springfield, Virginia 22161		
19. Security Classif.(of this report) Unclassified		20. Security Classif.(of this page) Unclassified		21. No. of Pages 29	
				22. Price	

Form DOT F 1700.7 (8-72)

Reproduction of completed page authorized

Disclaimer

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. Mention of trade names or commercial products does not constitute endorsement or recommendation for use.

Acknowledgements

The authors recognize that support for this research was provided by a grant from the U.S. Department of Transportation, University Transportation Centers, as well as by Samsung Research America and the Texas Department of Transportation under the Connected Vehicle Problems, Challenges, and Major Technologies project.

Table of Contents

1. Introduction	1
2. Ambiguity Resolution	2
2.1 Corrections Uncertainty and Network Density	4
3. Analysis of a Dense Reference Network.....	6
3.1 Network Corrections Estimation	7
3.2 Linear Least-Squares compared to Quadratic-Least Squares Estimation	9
3.3 Uncorrected Carrier-Phase Residuals.....	13
3.4 Network-Corrected Residuals	13
4. Network Deployment	15
4.1 Live Vehicle Demonstration	15
4.2 Demonstration Setup	16
4.3 Demonstration Results	17
5. Conclusions	22

List of Illustrations

Figure 1. Single-epoch single-frequency ambiguity fixing.....	3
Figure 2. Total probability of a correct fix for the scenario of Figure 1 as a function of ionospheric uncertainty σ_i	3
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.....	4
Figure 4. Total probability of a correct fix for the scenario of Figure 3 as a function of ionospheric uncertainty σ_i	4
Figure 5. Map showing trends in σ_v across a simulated reference network assuming a linear model for combined DD ionospheric and tropospheric delays and independent errors due to multipath at each station.....	5
Figure 6. Depiction of the placement of the 23 GNSS reference stations.	7
Figure 7. Probability densities of the terms estimated at the station location for SPMS of UNAVCO.	10
Figure 8. Probability densities of the terms calculated at the station location for SPMS of UNAVCO.	10
Figure 9. Probability densities of the terms for every GPS satellite observed, calculated at the station location for SPMS of UNAVCO, where each plot line represents a different GPS satellite.	11
Figure 10. Probability densities of the terms for every GPS satellite observed, calculated at the station location for SPMS of UNAVCO, where each plot line represents a different GPS satellite.	11
Figure 11. Probability densities of the difference between linear least-squares and quadratic least-squares network correction estimates for representative reference stations.	12
Figure 12. Standard deviation of uncorrected DD carrier-phase residuals versus baseline distance between each of the 22 reference stations and the master reference station.	13
Figure 13. Standard deviation of carrier-phase residual remainders (the carrier-phase residuals that remain after application of network corrections) versus average network density.	14
Figure 14. Overview of the planned Austin area reference network (Google Maps).	15
Figure 15. Low-cost reference station in the Longhorn Reference Network.....	15
Figure 16. Reference station components.	15
Figure 17. Demonstration route.	16
Figure 18. GNSS antenna configuration.	17

Figure 19. Low-cost, dual-frequency rover system in the trunk of the vehicle.	17
Figure 20. 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.....	18
Figure 21. Vehicle nearing lane edge, as depicted by yellow rectangle.	18
Figure 22. Vehicle crossing lane edge, as depicted by red rectangle.....	18
Figure 23. Four repeated traces of driver’s side antenna as vehicle made a turn.	18
Figure 24. Time history of the vertical deviation from the route mean, in meters.	19
Figure 25. 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.	19
Figure 26. Double-difference carrier phase (top) and pseudorange (bottom) residuals for GPS satellite 30 at frequency L1 over the full time interval shown in Figure 24.....	20
Figure 27. The number of signals above a 15-degree elevation mask.	21

1. Introduction

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. This report 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.

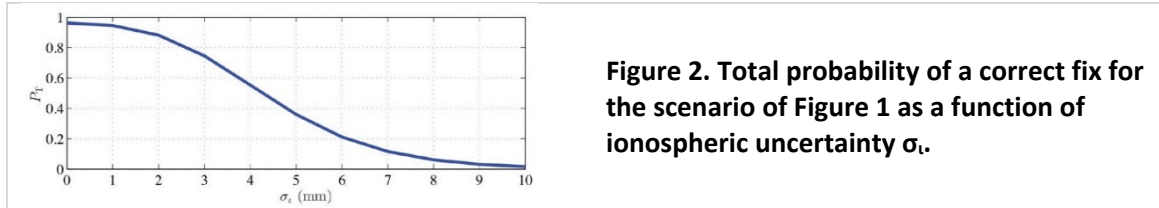
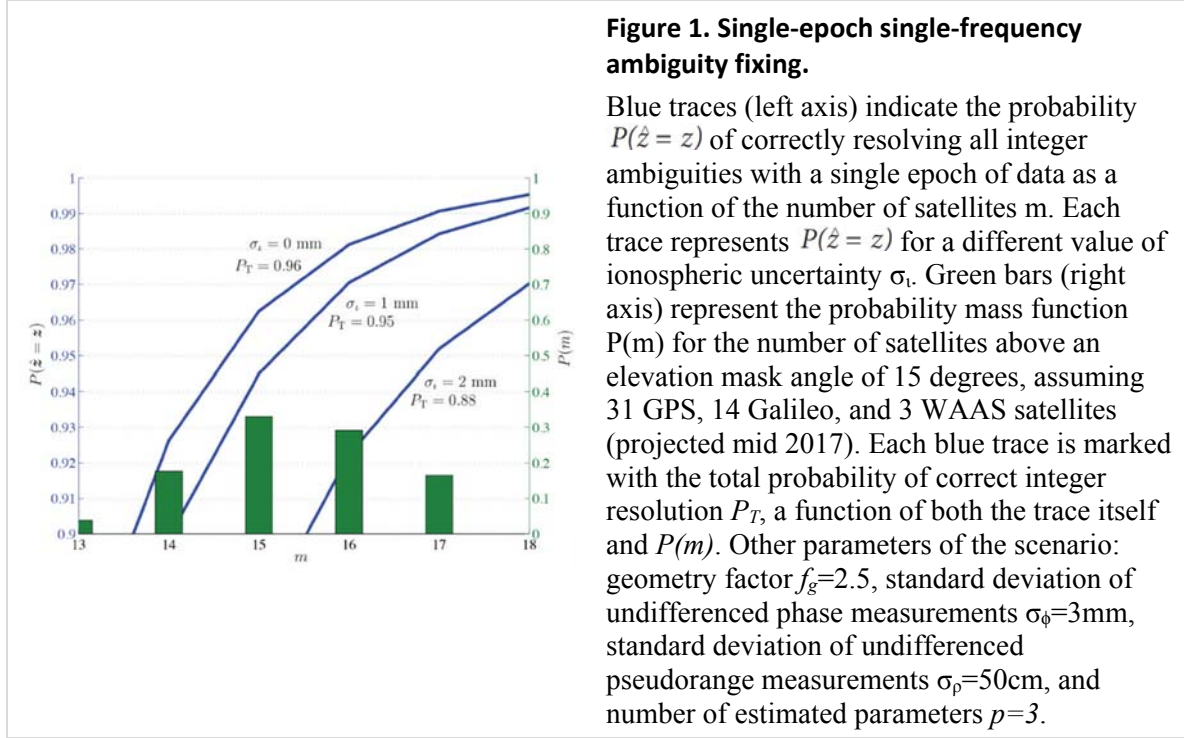
2. Ambiguity Resolution

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, 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 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.

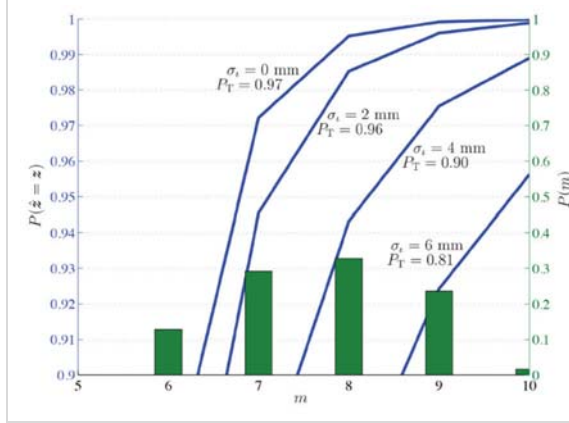


Figure 3. As Figure 1 except for dual-frequency (L1-L2) measurements and the probability mass function $P(m)$ corresponds only to a constellation of 31 GPS satellites.

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.

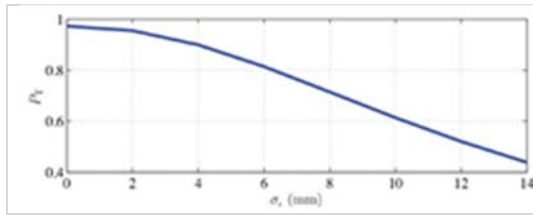


Figure 4. Total probability of a correct fix for the scenario of Figure 3 as a function of ionospheric uncertainty σ_i .

2.1 Corrections Uncertainty and Network Density

A key question arises in connection with σ_i : How is it related to reference network density? One expects to decrease with increased network density, but what is the exact relationship?

Dennis Odijk's work adopts a linear relationship between σ_i and the distance l between the user and the nearest reference station: $\sigma_i = \beta l$, $0.3 \leq \beta \leq 3 \text{ mm/km}$

Parameter β depends on ionospheric activity; Odijk recommends determining β empirically. Similarly, his other work adopts a linear relation, with $\beta = \text{mm/km}$. But there appears to be no justification for applying this linear model to ionospheric corrections provided to a user by a network of reference receivers. The linear trend corresponds to individual single-baseline solutions involving a single master reference station without network aiding; it is not representative of how σ_i varies for a rover within a reference network.

Instead of determining how σ_i varies throughout a reference network, it will be more useful to consider the spatial variation in the variance of aggregate error in network-provided corrections. The aggregate error variance, denoted σ_v^2 , can be modeled as the sum of variances associated with (1) residual ionospheric delay error, (2) residual neutral atmospheric (hereafter tropospheric) error, and (3) error due to carrier-phase multipath at the reference network stations:

$$\sigma_v^2 = \sigma_t^2 + \sigma_t^2 + \sigma_m^2 \quad (1)$$

This model assumes that precise orbital ephemerides are used to eliminate spatially-correlated errors due to satellite ephemeris errors and that the contribution σ_v^2 from reference station carrier-phase thermal noise is negligible compared to reference station carrier-phase multipath error.

Focusing therefore on σ_v , consider its relationship to reference network density γ , expressed in stations per unit area. This relationship depends on the assumed model for the DD ionospheric and tropospheric delays. Let a denote the master reference station and let $S = \{s1, s2, \dots, sN\}$ denote the set of all secondary stations available in the network. Then, for pivot satellite i and alternate satellite j , suppose that the true combined DD atmospheric delay at secondary station $s \in S$ can be accurately modeled as follows, where x_s , y_s , and z_s represent the secondary station's east, north, and up displacement from the master:

$$v_{as}^{ij} = c_x^{ij}x_s + c_y^{ij}y_s + c_z^{ij}z_s + c_0^{ij} \quad (2)$$

Dai et al. refer to this model as a linear interpolation model or first-order surface model.

The quantities $c_x^{ij}, c_y^{ij}, c_z^{ij}$ and c_0^{ij} are the model parameters for the satellite pair i, j . Figure 5 illustrates trends in σ_v across a simulated reference network.

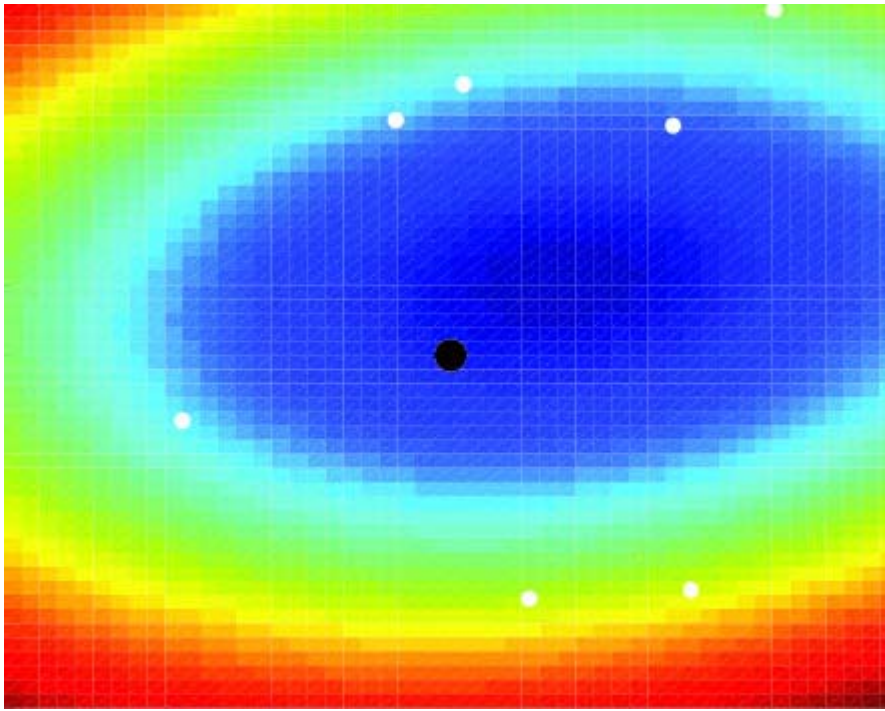


Figure 5. Map showing trends in σ_v across a simulated reference network assuming a linear model for combined DD ionospheric and tropospheric delays and independent errors due to multipath at each station.

The master station is marked in black; secondary reference stations are marked in white. Blue denotes low σ_v ; red denotes high σ_v .

For the linear model in (1), one can show that if stations are sufficiently uniformly distributed (i.e., no station clumping), then the average value of σ_v across a network, denoted $\bar{\sigma}_v$, is approximately related to the network density γ by

$$\bar{\sigma}_v = \frac{q}{\sqrt{\gamma}} \quad (3)$$

where q is a parameter related to the variance of the uncorrelated errors w_{as}^{ij} , $s \in S$. This approximation becomes highly accurate as γ increases.

It is clear from (3) that, for the linear model (1), $\bar{\sigma}_v$ can be driven to an arbitrarily small value by increasing the network density γ , and this is true despite the presence of multipath in the reference station carrier-phase measurements. Whether (3) applies in practice depends on whether (1) can be considered an accurate model for v_{as}^{ij} , at least over a compact region. The following section examines this question empirically. It further seeks to identify, for an example dense reference network, the density γ beyond which further reduction $\bar{\sigma}_v$ no longer matters (would no longer improve $P(\hat{z} = z)$) because rover multipath dominates.

3. Analysis of a Dense Reference Network

We examined σ_v 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 relative positions of the sites selected to comprise this reference network, located between Los Angeles and Pomona, California, are depicted graphically in Figure 6.

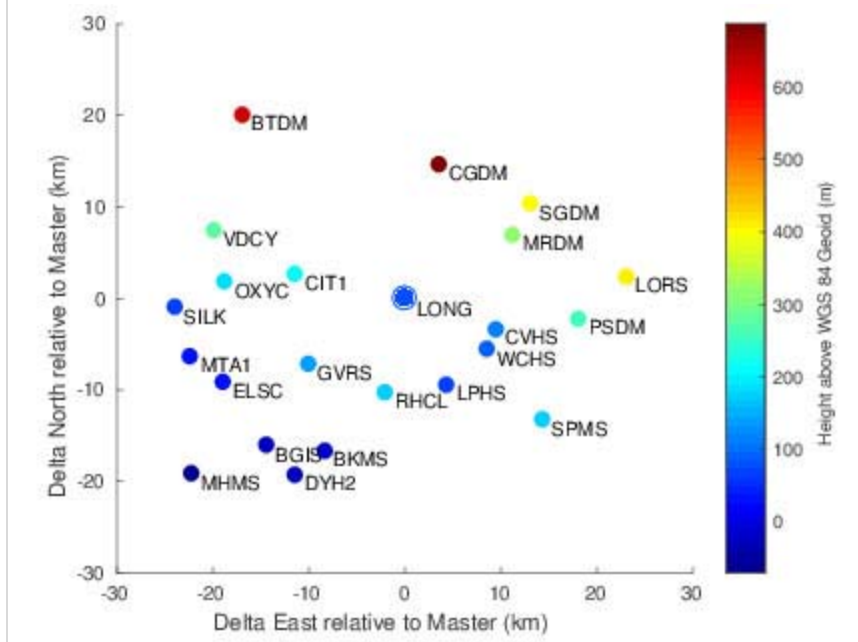


Figure 6. Depiction of the placement of the 23 GNSS reference stations.

Horizontal positions are relative to the master station, LONG of CRTN, in kilometers. The color map indicates the height of each station above the WGS 84 geoid in meters.

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.

3.1 Network Corrections Estimation

Estimation of network corrections made use of least-squares estimation applied to carrier-phase residuals measured between master station LONG, denoted a hereafter, and secondary reference stations $s \in \mathbf{S}$, where \mathbf{S} is now taken to be the set of all stations other than LONG. Consider the following model for the DD carrier-phase measurement, expressed in meters, between master station a , secondary station $s \in \mathbf{S}$, pivot satellite i , and alternate satellite j :

$$\lambda \phi_{as}^{ij} = r_{as}^{ij} + v_{as}^{ij} + \lambda N_{as}^{ij} + w_{as}^{ij} \quad (4)$$

Here, λ is the carrier wavelength; ϕ_{as}^{ij} is the DD carrier-phase measurement, in cycles; r_{as}^{ij} is the DD range; $N_{as}^{ij} \in \mathbb{Z}$ is the DD integer ambiguity; v_{as}^{ij} is the DD combined atmospheric delay, which includes tropospheric and ionospheric delays; and w_{as}^{ij} is the DD carrier-phase measurement error, which is dominated by carrier-phase multipath error at a and s .

Experimental analysis of $\bar{\sigma}_v$ as a function of network density proceeded as follows. A subset of secondary stations $\mathbf{S}_k \subset \mathbf{S}$ was chosen, together with a , to act as the k th test network. A large number K of subsets \mathbf{S}_k of various geographic size and density were analyzed. Let $\{\mathbf{S} \setminus \mathbf{S}_k\}$ denote the set of secondary stations not in the k th test network. For each \mathbf{S}_k , $k = 1, 2, \dots, K$, all secondary stations in $\{\mathbf{S} \setminus \mathbf{S}_k\}$ were designated, one at a time, to act as a test station, or rover. Atmospheric delays estimated by the k th network for test station $s \in \{\mathbf{S} \setminus \mathbf{S}_k\}$ were then differenced from actual delays measured by s to evaluate the quality of the atmospheric delay estimates.

Details of the atmospheric delay estimation procedure for the k th test network are as follows. For each $s \in \mathbf{S}_k$, a DD measurement residual was formed for each pivot satellite i and alternate satellite j as

$$\tilde{v}_{as}^{ij} = \lambda \phi_{as}^{ij} - r_{as}^{ij} - \lambda N_{as}^{ij} \quad (5)$$

Where r_{as}^{ij} was assumed known to sub-millimeter accuracy and was assumed to have been resolved correctly. The true DD atmospheric error v_{as}^{ij} contributing to (5) was assumed to vary linearly with geometry over sufficiently short baselines as modeled in (2). The DD multipath error term w_{as}^{ij} was assumed to be zero mean, and the component w_{as}^{ij} due solely to s was assumed to be uncorrelated with all corresponding components w_{au}^{ij} , $u \in \{\mathbf{S}_k \setminus s\}$.

Under these assumptions, v_{as}^{ij} can readily be estimated via least squares. Let $\tilde{\mathbf{v}}^{ij}$ be the vector containing the residuals for $|\mathbf{S}_k| \times 1$. This residuals vector can be modeled as

$$\tilde{\mathbf{v}}^{ij} = H \mathbf{c}_1^{ij} + \mathbf{w}^{ij} \quad (6)$$

where H is an $|\mathbf{S}_k| \times 4$ matrix whose rows are of the form $[x_s \ y_s \ z_s \ 1]$. The 4×1 vector contains the parameters of the hyper-plane to be estimated at each epoch. The $|\mathbf{S}_k| \times 1$ vector \mathbf{w}^{ij} contains DD measurement errors.

An estimate $\hat{\mathbf{c}}_1^{ij}$ from a least-squares solution of (6) was used to produce a network correction $\hat{v}_{as,l}^{ij}$ for a test secondary station $s \in \{\mathbf{S} \setminus \mathbf{S}_k\}$, acting as rover, at location x_s, y_s, z_s :

$$\hat{v}_{as,l}^{ij} = \hat{c}_{x1}^{ij} x_s + \hat{c}_{y1}^{ij} y_s + \hat{c}_{z1}^{ij} z_s + \hat{c}_0^{ij} \quad (7)$$

The subscript l on the atmospheric correction $\hat{v}_{as,l}^{ij}$ indicates that the correction is based on a linear model for DD atmospheric errors; it is used to distinguish the correction from those produced by a quadratic model later on. The correction $\hat{v}_{as,l}^{ij}$ was applied at test station $s \in \{S \setminus S_k\}$ to produce a corrected DD phase measurement: $\lambda \tilde{\phi}_{as}^{ij} = \lambda \phi_{as}^{ij} - \hat{v}_{as,l}^{ij}$.

This procedure was repeated at each epoch for each satellite pair i, j visible to each test station $s \in \{S \setminus S_k\}$ of the k th test network, $k = 1, 2, \dots, K$.

If the assumed models hold, then in the limit as the network density increases, $\lambda \tilde{\phi}_{as}^{ij}$ can be modeled as

$$\lambda \tilde{\phi}_{as}^{ij} = r_{as}^{ij} + \lambda N_{as}^{ij} + w_{as}^{ij} \quad (8)$$

where w_{as}^{ij} is DD phase measurement error due only to multipath at s . In other words, as network density increases, application of the network correction $\hat{v}_{as,l}^{ij}$ eliminates not only v_{as}^{ij} but also w_{as}^{ij} , the component of the DD phase measurement error due to multipath at the master.

3.2 Linear Least-Squares compared to Quadratic-Least Squares Estimation

To evaluate the assumption that DD tropospheric and ionospheric errors vary proportional to relative position, c_1 was estimated with the full set of secondary stations S for single epochs at 300-second intervals. The probability distributions of the contributions of those parameters (e.g., $c_{xl}x_s$ and not simply c_{xl}) are shown below. For comparison, equivalent values are calculated for a quadratic least-squares estimate of the following form:

$$\begin{aligned} \hat{v}_{as,q}^{ij} = & \hat{c}_{x2}^{ij} x_s^2 + \hat{c}_{y2}^{ij} y_s^2 + \hat{c}_{z2}^{ij} z_s^2 \\ & + \hat{c}_{x1}^{ij} x_s + \hat{c}_{y1}^{ij} y_s + \hat{c}_{z1}^{ij} z_s + \hat{c}_0^{ij} \end{aligned} \quad (9)$$

Here, the subscript q of $\hat{v}_{as,q}^{ij}$ denotes a quadratic model for DD atmospheric delays. The distributions of comparable terms from (9) are also shown in Figures 7 and 8. These data represent the collection of all satellites above the elevation mask angle. It is noted that when all satellites are considered together, the expected value of these terms is very near zero.

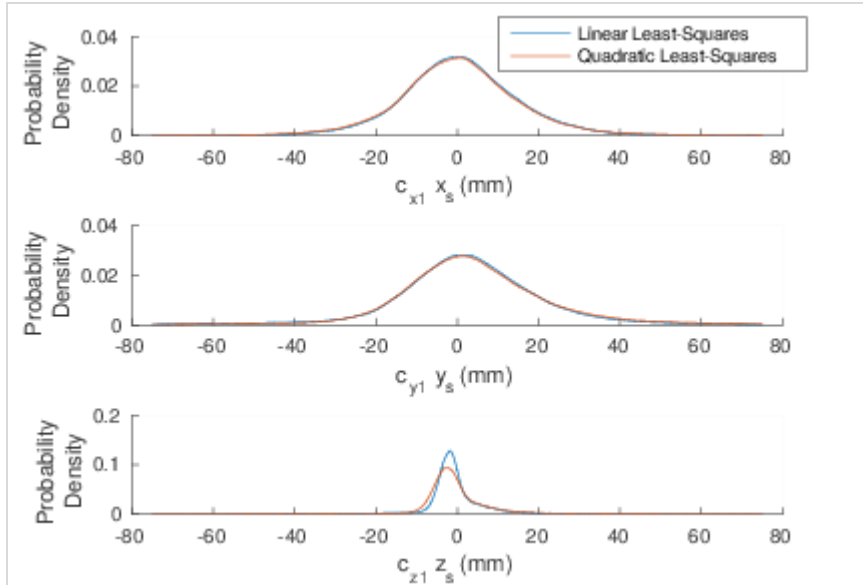


Figure 7. Probability densities of the terms estimated at the station location for SPMS of UNAVCO.

As indicated by the legend, the linear components are shown for a linear least-squares estimation as well as the linear components for a quadratic least-squares estimation. These data represent the probability densities for all GPS satellites combined.

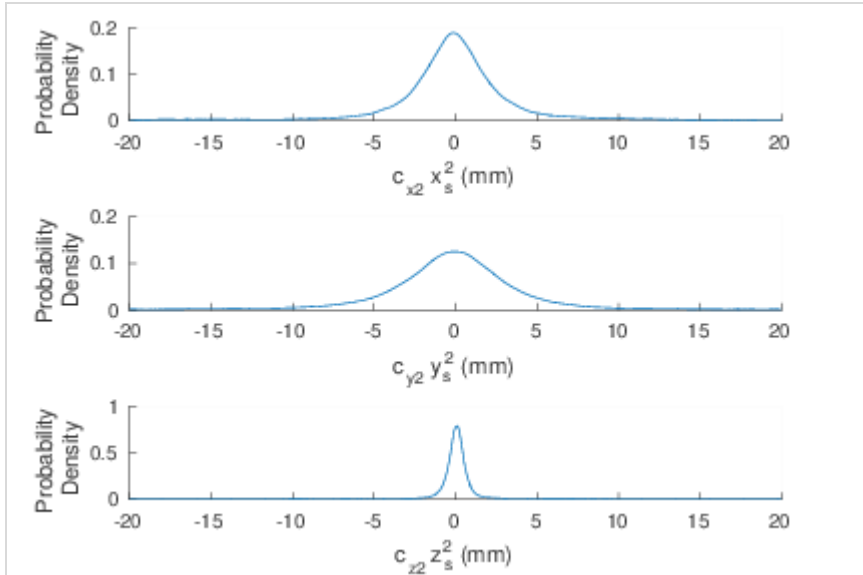


Figure 8. Probability densities of the terms calculated at the station location for SPMS of UNAVCO.

Figures 9 and 10 show the same data as the two above, but with each GPS satellite plotted separately. It is noted that the linear parameters, when considering only a particular satellite, are not necessarily zero-mean. This is hypothesized to be a manifestation of the satellite orbit reflected in the tropospheric and ionospheric errors. It

is interesting to note that the quadratic terms shown in Figure 10 largely exhibit zero mean behavior despite non-zero mean for the associated linear terms.

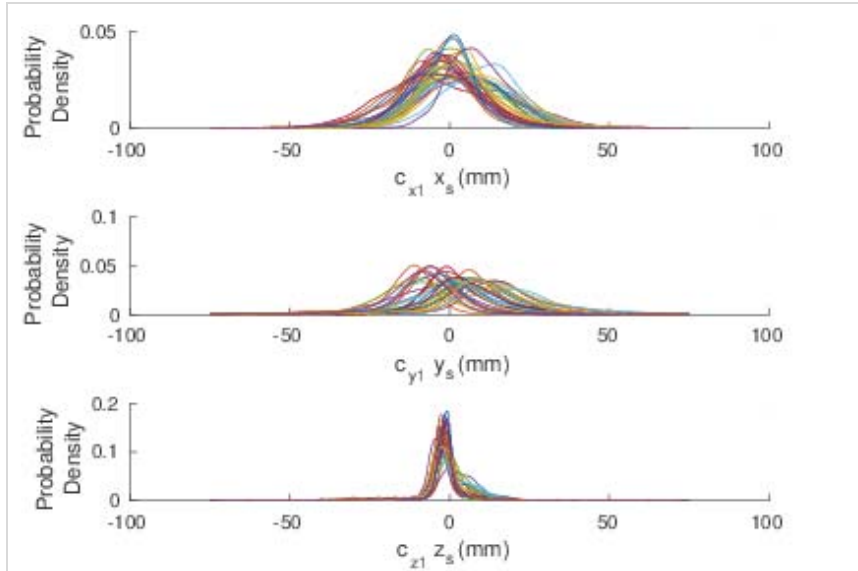


Figure 9. Probability densities of the terms for every GPS satellite observed, calculated at the station location for SPMS of UNAVCO, where each plot line represents a different GPS satellite.

This figure is intended to qualitatively illustrate the non-zero mean nature of these linear terms when considered for individual satellites.

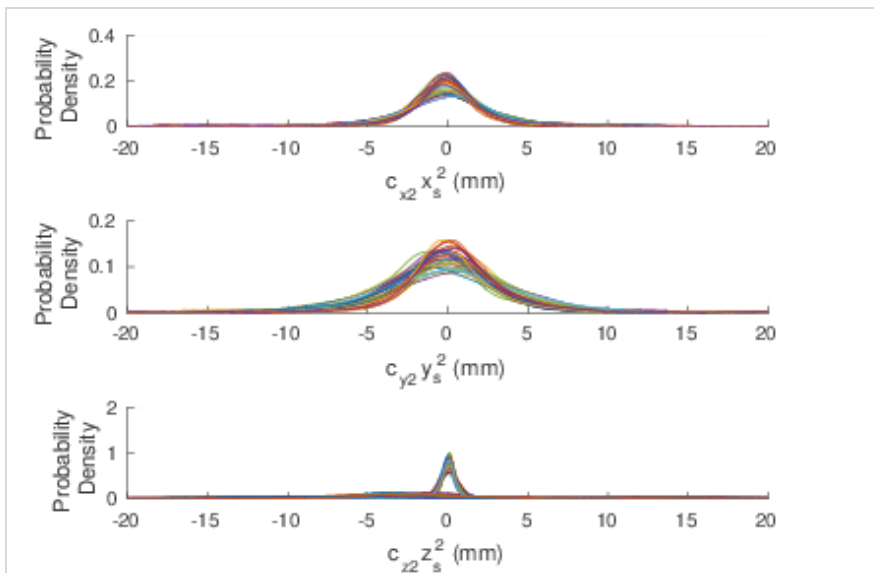


Figure 10. Probability densities of the terms for every GPS satellite observed, calculated at the station location for SPMS of UNAVCO, where each plot line represents a different GPS satellite.

This figure is included to qualitatively illustrate the largely zero mean nature of these quadratic terms when considered for individual satellites.

Figure 11 shows the probability distributions of the difference between (7) and (9) (i.e., $\hat{v}_{as,l}^{ij} - \hat{v}_{as,q}^{ij}$) at three representative reference station positions. It can be noticed that despite the increasing baseline distance of LORS and BGIS as compared to CGDM, there is no apparent correlation in these estimation errors. Notice that CGDM and LORS have very similar distributions despite their difference in baselines. BGIS and LORS, with similar baselines, exhibit very different distributions. There is no apparent correlation found between reference station positions and these error terms. Additionally, these distributions are zero-mean for all $s \in \mathbf{S}$ (to within 0.5 mm in each case) with 68.27% boundaries positioned between 1.5–5.5 mm. Because these errors appear indistinguishable from multipath, it is concluded, for this specific network and time period, that linear least-squares estimation is sufficient for estimating tropospheric and ionospheric errors. This is fortunate, because the linear model for atmospheric DD delays provides an averaging effect on multipath present at the reference stations, which minimizes the introduction of multipath errors into the estimates produced.

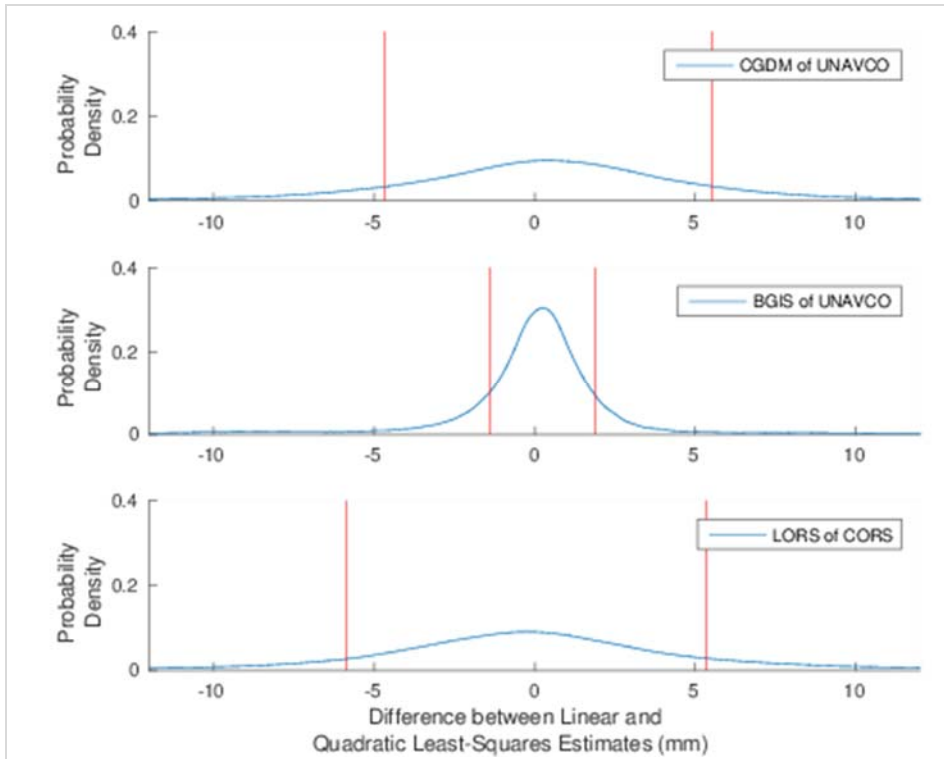
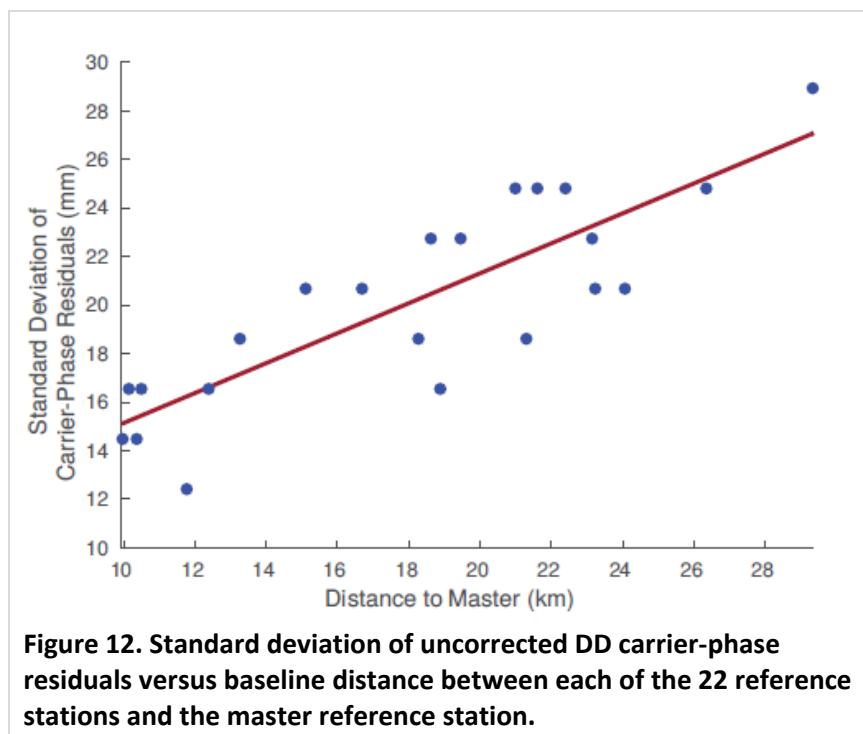


Figure 11. Probability densities of the difference between linear least-squares and quadratic least-squares network correction estimates for representative reference stations.

The red vertical lines denote boundaries between which 68.27% of the probability distribution is contained, displayed as a comparative proxy to 1σ of the Gaussian-distribution (these distributions are non-Gaussian). Recall that CGDM has a distance to the master station of 15.1km, BGIS is at 21.6km, and LORS is at 23.1km.

3.3 Uncorrected Carrier-Phase Residuals

Figure 12 shows the expected values for DD carrier-phase residual standard deviations for all secondary stations ($s \in S$) through use of uncorrected observations. 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 master, $\beta = 0.62$ mm/km. Additionally, the mm-level scatter of these data points suggests that position biases of the resolved reference station positions are also mm-level. If the linear fit is shifted down by approximately 4 mm (e.g., taking the minimum data points as those with very little position bias) and extrapolated to 0 km, one can consider this as providing a rough estimate of DD multipath at the reference stations; 4.7 mm (DD) or 3.3 mm (single-difference equivalent).



3.4 Network-Corrected Residuals

Figure 13 displays similar data to Figure 12, except that the carrier-phase residuals are those that remain after network corrections are applied. Each data point corresponds to a particular subset of secondary stations together with the master, and a particular rover selected at random from the remaining stations. Both the size and specific selection of secondary stations comprising each subset were randomly selected. In all, 70 different network configurations and more than 3.67 million NRTK solutions were analyzed.

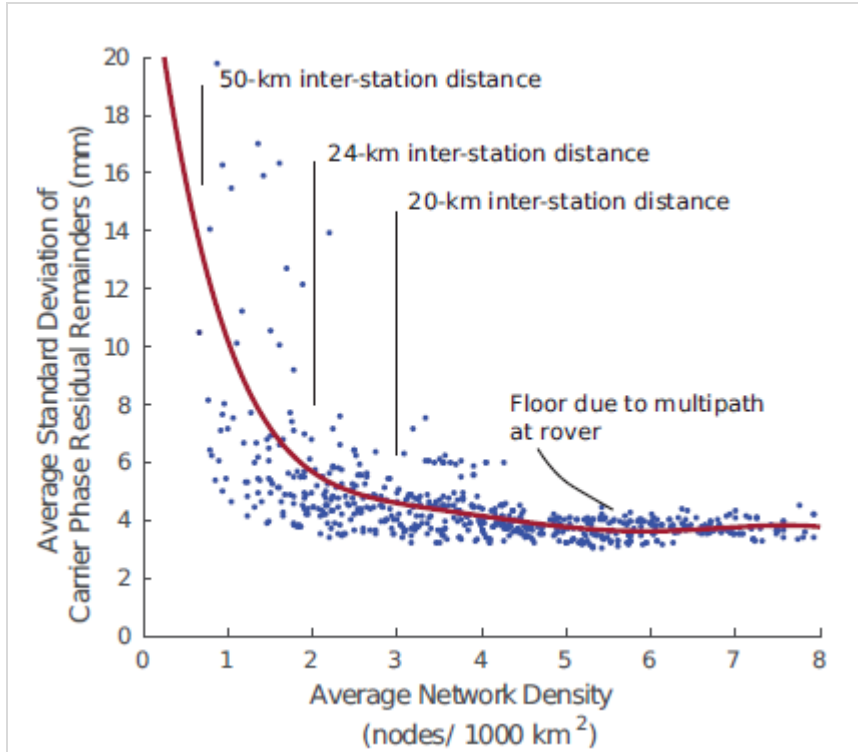


Figure 13. Standard deviation of carrier-phase residual remainders (the carrier-phase residuals that remain after application of network corrections) versus average network density.

The fitted curve is simply a polynomial fit of these data; it is not based on any theoretically anticipated behavior.

Figure 13 shows that carrier-phase residuals after application of network corrections are considerably reduced compared to those original magnitudes seen in Figure 12. With increasing network density, the DD residuals' deviation asymptotically approaches a minimum value of about 4 mm, which corresponds to an undifferenced deviation of 2 mm. This floor is due to multipath at the rover. Deviations in excess of this floor are caused by residual ionospheric errors and, to a lesser extent, neutral atmospheric errors. Attributing the excess deviation entirely to residual ionospheric errors, and assuming these are uncorrelated with multipath, one can estimate from Figure 13 the undifferenced ionospheric uncertainty. For example, for a 50-km inter-station distance, $\sigma_i = (\sqrt{(14^2 - 4^2)})/2 = 6.7$ mm. To achieve the $\sigma_i < 2$ mm recommended earlier for fast and reliable AR, station separation should be no more than 22 km, which we round down to a recommended value of 20 km to provide a margin of station redundancy.

4. Network Deployment

We have developed and deployed a low-cost reference network testbed in Austin, Texas, with site hosting courtesy of the Texas Department of Transportation. The Longhorn Reference Network boasts a dozen stations, with plans for 20 (Figure 14). The network's average inter-station spacing is far shorter than the 20-km spacing recommended earlier. The tighter spacing provides redundancy and flexibility of experimentation. The low-cost reference stations are deployed in environments with greater multipath and signal blockage than those of the high-quality stations studied earlier. Such non-ideal signal environments are to be expected in a dense low-cost reference network, for which choice of station siting is driven largely by opportunity.

The reference station design, pictured in Figure 15 and diagrammed in Figure 16, is novel. Each station is a self-contained, solar-powered node supporting a software-defined dual-frequency, dual-antenna GNSS receiver with an always-on cellular connection to university servers for data collection and software maintenance.

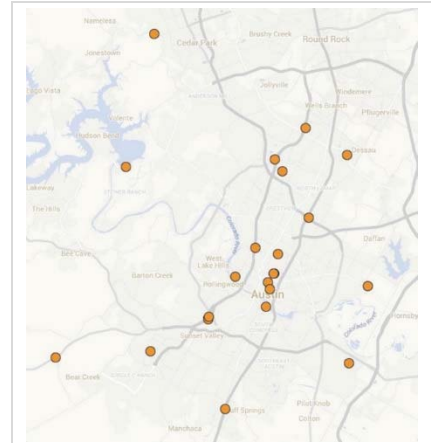


Figure 14. Overview of the planned Austin area reference network (Google Maps).



Figure 15. Low-cost reference station in the Longhorn Reference Network.

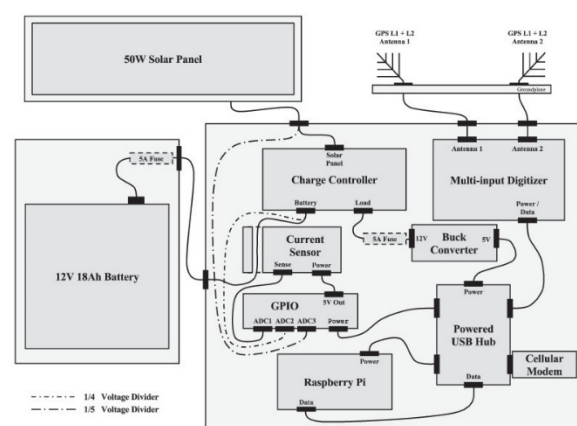


Figure 16. Reference station components.

4.1 Live Vehicle Demonstration

In partnership with Radiosense, an 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.

4.2 Demonstration Setup

The live demonstration followed a predetermined route in the vicinity of the University of Texas campus. The 1-mile route (Figure 17) passed through both open-sky and partially-blocked environments.

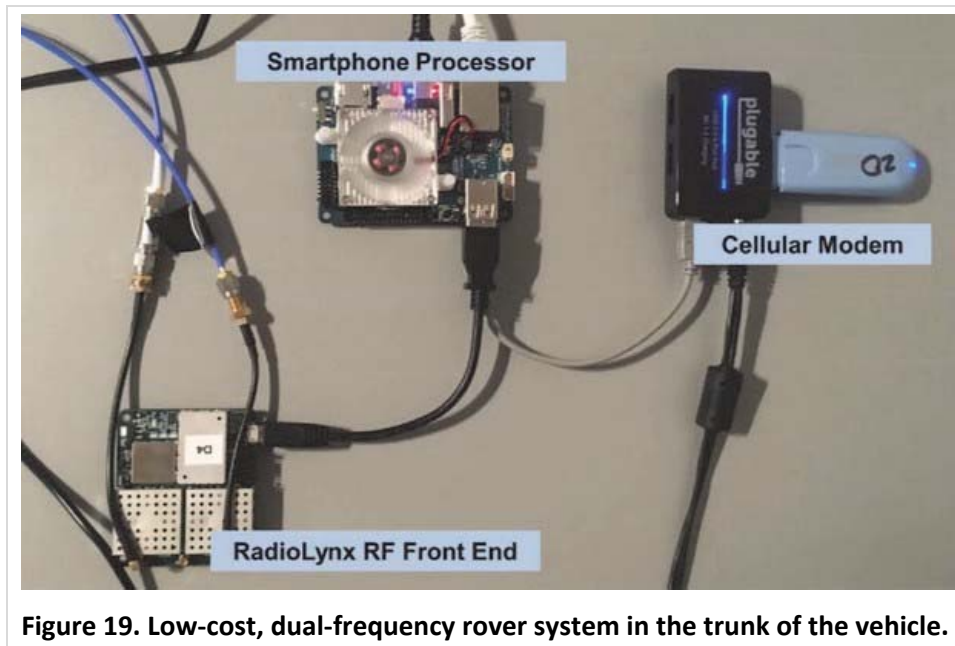
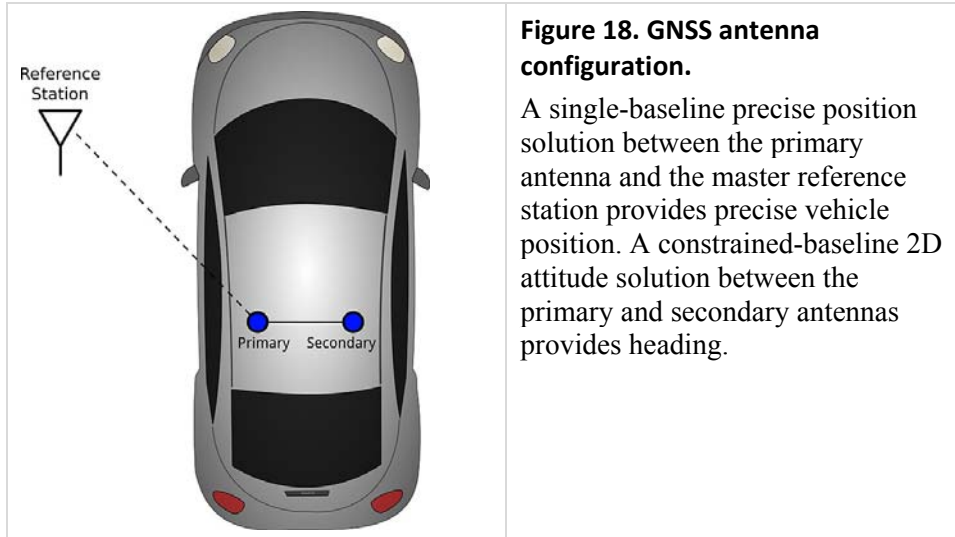
Prior to the demonstration, the vehicle was driven several times on the same route collecting GNSS measurements to precisely map typical driving trajectories on the route. The ensemble of trajectories was used to build a centimeter-accurate model of the lane center along the route. The sensing equipment employed during this mapping phase is no different than that used during the demonstration, making feasible eventual crowd-sourcing, wherein end-user vehicles generate and update the centerline models.



Figure 17. Demonstration route.

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 18. 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 19), 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.



4.3 Demonstration Results

Figures 20, 21, and 22 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 20. 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 21. Vehicle nearing lane edge, as depicted by yellow rectangle.



Figure 22. Vehicle crossing lane edge, as depicted by red rectangle.

Figure 23 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 23. Four repeated traces of driver's side antenna as vehicle made a turn.

Figure 24 shows a time history of the vertical deviation from the route mean, in meters. The zoomed view of the vertical deviation shown in Figure 25 allows one to appreciate the precision of the positioning engine: the vertical trajectory is smooth at the centimeter level. Figure 26 shows the DD residuals in carrier phase and pseudorange for GPS PRN

30 during the four loops in Figure 27. One-sigma undifferenced phase and pseudorange deviations are 3.4 mm and 42 cm, respectively.

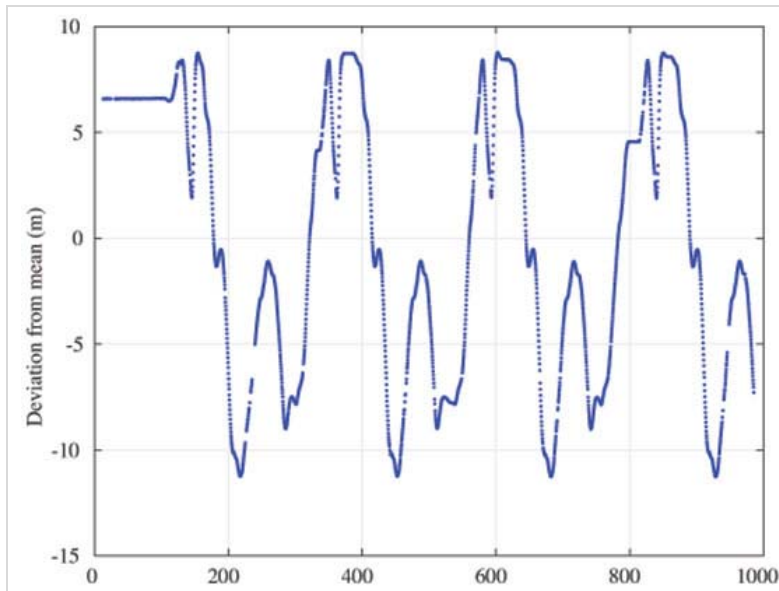


Figure 24. Time history of the vertical deviation from the route mean, in meters.

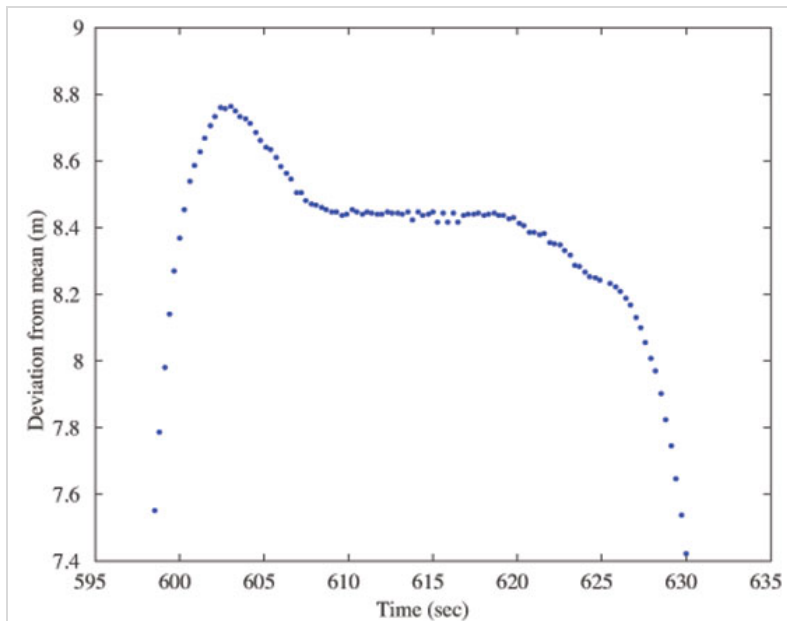


Figure 25. Zoomed view of the time history of the vertical deviation from the route mean, showing the centimeter-level precision in the 3.3 Hz positioning data.

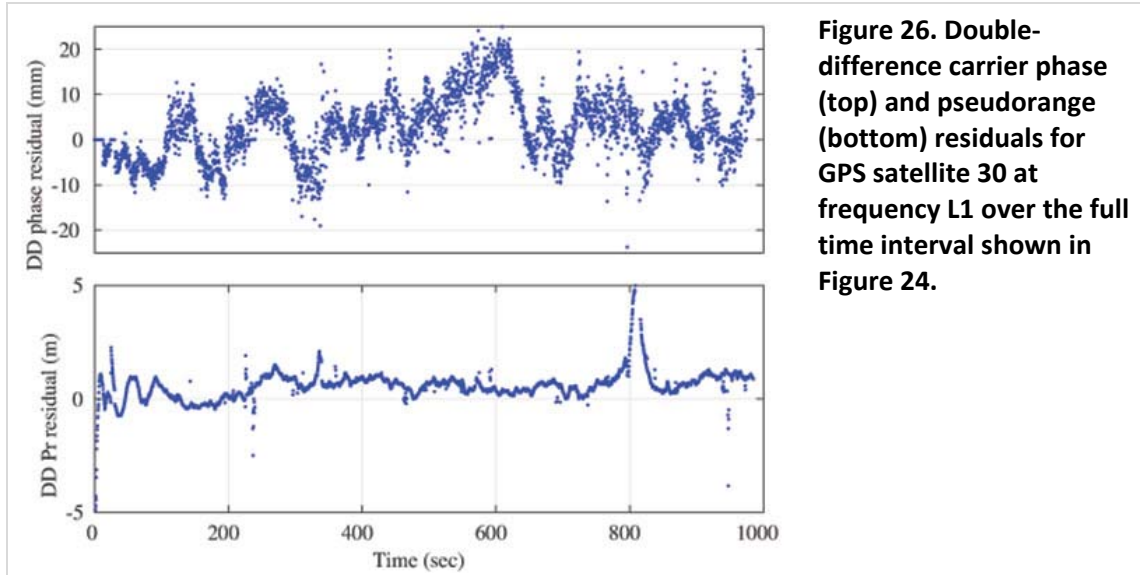


Figure 26. Double-difference carrier phase (top) and pseudorange (bottom) residuals for GPS satellite 30 at frequency L1 over the full time interval shown in Figure 24.

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 24, 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 27. Sixteen 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.

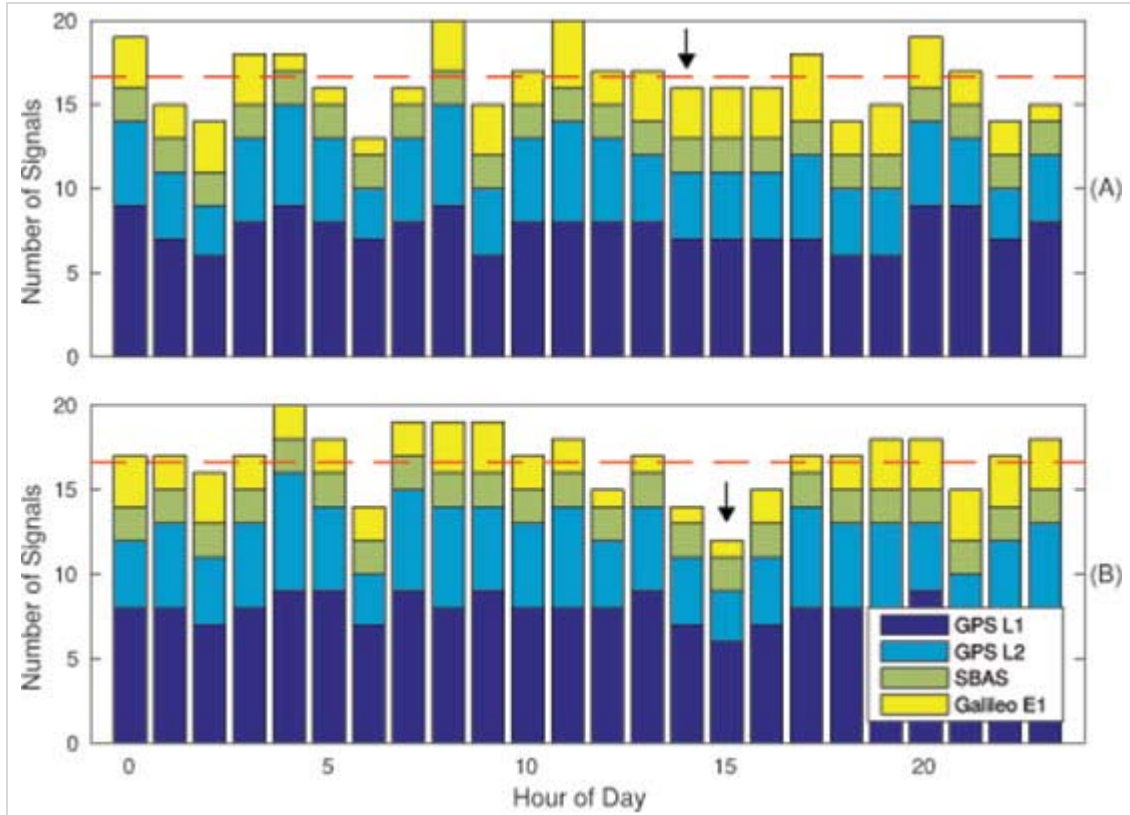


Figure 27. The number of signals above a 15-degree elevation mask.

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.

For insight into the performance of the positioning engine as a function of the number of overhead satellites, Table 1 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.

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

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%

A large drop-off 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.

5. 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.