Estimating the Residual Tropospheric Delay for Airborne Differential GPS Positioning[‡]

J. Paul Collins and Richard B. Langley

Geodetic Research Laboratory, Department of Geodesy and Geomatics Engineering, University of New Brunswick, Fredericton, N.B., Canada.

BIOGRAPHIES

Paul Collins graduated from the University of East London in 1993 with a B.Sc. (Hons) degree in Surveying and Mapping Sciences. He is currently enrolled in the M.Sc.E. degree program in the Department of Geodesy and Geomatics Engineering at the University of New Brunswick (UNB), where he is investigating the effects of the troposphere on kinematic GPS positioning.

Richard Langley is a professor in the Department of Geodesy and Geomatics Engineering at UNB, where he has been teaching since 1981. He has a B.Sc. in applied physics from the University of Waterloo and a Ph.D. in experimental space science from York University, Toronto. Prof. Langley has been active in the development of GPS error models since the early 1980s and is a contributing editor and columnist for GPS World magazine.

ABSTRACT

In post-processing dual frequency GPS carrier phase data, the residual tropospheric delay can easily be the largest remaining error source. This error can contribute a bias in height of several centimetres even if simultaneously recorded meteorological data are used. This shortcoming is primarily due to the poor representation of the water vapour profile in the tropospheric delay models. In addition, a lack of real-time meteorological data would force the scaling of either surface values or standard atmosphere values, neither of which are likely to accurately represent the ambient atmosphere.

To obtain the highest precision in kinematic GPS some advantage may be obtained by estimating the residual tropospheric delay along with the position of the moving platform. The simple tests reported in this paper removed biases of upto ten centimetres in height when estimating the residual tropospheric delay from GPS data recorded at an aircraft in flight. However, important limitations exist in the geometry of the satellite coverage which must be considered before the full reliability of the technique can be assessed.

INTRODUCTION

This paper describes an investigation into the estimation of the residual tropospheric delay from GPS signals. This parameter is the remaining part of the tropospheric delay not predicted by empirical models. In post-processed dual frequency carrier phase data, it can easily be the largest remaining error source.

Unlike most applications of this technique, where the receiver is static, we have used data recorded at an aircraft in flight. This idea was motivated by the fact that highly accurate aircraft positions are required for gravimetric, altimetric and photogrammetric surveying purposes. Increasingly, GPS is being used to provide the decimetre-level accuracy required for some of these techniques. This level of precision can be achieved using carrier phase observables, but we will show that unmodelled tropospheric effects could potentially contribute a bias of a similar magnitude.

The Tropospheric Delay

An electromagnetic signal propagating through the neutral atmosphere is affected by the constituent gases. The fact that the refractive index is slightly greater than unity gives rise to a decrease in the signal's velocity. This increases the time taken for the signal to reach a GPS receiver's antenna, increasing the equivalent path length (both effects are often referred to as the "delay"). Refraction also bends the raypath and thereby lengthens it, further increasing the delay. Because the bulk of the delay occurs within the troposphere, the whole delay is often referred to solely as the "tropospheric delay".

[‡] *Poster presentation at ION GPS '97, 10th International Technical Meeting of the Satellite Division of The Institute of Navigation, Kansas City, Mo., September 16-19, 1997.*

By assuming that the neutral atmosphere is both horizontally stratified and azimuthally symmetric, the tropospheric delay can be modelled in two parts: the delay experienced in the zenith direction and the scaling of that delay to the delay experienced at the zenith angle of the raypath. The functions that undertake the scaling are usually termed mapping functions, although obliquity factor is sometimes used. This leads to the common formulation of zenith delays and mapping functions seen in the space geodetic literature. The typical formulation of the tropospheric delay is described as:

$$T_i^k = t_i^z(hyd)\cdot m_i^k(hyd) + t_i^z(wet)\cdot m_i^k(wet), \qquad (1)$$

where at the antenna of receiver *i*, the delay on the signal from satellite *k* is a function of the delays in the zenith direction caused by the atmospheric gases in hydrostatic equilibrium and by those gases not in hydrostatic equilibrium (primarily water vapour), $t_i^z(hyd)$ and $t_i^z(wet)$ respectively; and their mapping functions, $m_i^\kappa(hyd)$ and $m_i^\kappa(wet)$ respectively. The mapping functions are usually described as functions of the satellite elevation angle – the complement of the zenith angle.

For simplicity, we will consider equation (1) as:

$$T_i^k = t_i^z(total)\cdot m_i^k(combined), \qquad (2)$$

where

$$m_i^k(combined) = \frac{t_i^z(hyd)\cdot m_i^k(hyd) + t_i^z(wet)\cdot m_i^k(wet)}{t_i^k(hyd) + t_i^k(wet)}. \qquad (3)$$

When processing GPS observations, a value for the tropospheric delay is predicted using empirical models which in general must be provided with values of the ambient temperature, pressure and relative humidity. Unfortunately, even with accurate values, these models rarely predict the true delay with a high degree of accuracy. In theory, the hydrostatic component of the delay can be predicted in the zenith to the millimetre level, however the highly variable nature of atmospheric water vapour means that the accuracy of the non-hydrostatic delay is at the centimetre, or even decimetre level.

In addition, when recording GPS data at an aircraft, it is often the case that no meteorological information is recorded at the same time. When processing this data, assumed meteorological values must be used, and in addition to the poorly modelled wet component, there could also be a bias contributed by the hydrostatic component. The recovery of these errors is the aim of the techniques presented here.

The GPS Observables

Because of its high precision and low noise characteristics we used the GPS carrier phase observations. In general, the pseudorange measurements are too noisy to allow for the accurate estimation of the tropospheric delay. To remove the bias of the satellite and receiver clocks we have used the double-difference observable. Ignoring multipath and noise, we have

$$\Phi_{ij}^{kl} = \rho_{ij}^{kl} + T_{ij}^{kl} - I_{ij}^{kl} + \lambda N_{ij}^{kl}, \qquad (4)$$

where ρ represents the differential geometric range between satellites *k* and *l* and stations *i* and *j*; *T* is the differential delay caused by the troposphere; *I* is the differential delay caused by the ionosphere; λ is the carrier frequency wavelength; and *N* is the differential integer cycle ambiguity.

The differential ionospheric delay can be removed from equation (4) by using dual frequency data with the standard inter-frequency combination. The double-difference integer ambiguity term must also be resolved by some suitable method. It is important that the ambiguities then remain fixed for the solution to be consistent. For this reason, it is important that the carrier phase data be free from cycle-slips and data gaps.

THE RESIDUAL TROPOSPHERIC DELAY

There are two ways of estimating the residual tropospheric delay: either as a scale factor, *s*; or as a residual zenith component, r^z:

$$T_i^k = (1 + s_i)\cdot t_i^z \cdot m_i^k, \qquad (5)$$

$$T_i^k = (t_i^z + r_i^z)\cdot m_i^k. \qquad (6)$$

where we have dropped the zenith delay and mapping function parenthetical labels for clarity.

By restricting the residual error to the zenith delay, we are assuming that there are no errors in the mapping function. This is obviously untrue, however recent mapping functions such as those of *Niell* [1996] have been shown to be very accurate and any remaining error will likely come from unmodelled atmospheric gradients and azimuthal asymmetry. In theory these gradients can also be modelled, however, our data may not have the sensitivity to detect them.

Modelling Considerations

The differential tropospheric delay is given by:

$$T_{ij}^{kl} = T_j^l - T_j^k - T_i^l + T_i^k , \qquad (7)$$

which, because the zenith delay at a particular station will be the same for satellites *l* and *k*, can be written as:

$$T_{ij}^{kl} = t_j^z\left(m_j^l - m_j^k\right) - t_i^z\left(m_i^l - m_i^k\right). \qquad (8)$$

Estimating a station dependent scale factor, *s*, gives:

$$T_{ij}^{kl} = \left(1+s_j\right)t_j^z\left(m_j^l - m_j^k\right) - \left(1+s_i\right)t_i^z\left(m_i^l - m_i^k\right), \qquad (9)$$

with partial derivatives for a least-squares adjustment:

$$\frac{\partial T_{ij}^{kl}}{\partial s_j} = t_j^z\left(m_j^l - m_j^k\right) \text{ and } \frac{\partial T_{ij}^{kl}}{\partial s_i} = t_i^z\left(m_i^k - m_i^l\right),$$

which are the between-satellite single difference tropospheric delays. Estimating a residual zenith delay, r^z gives:

$$T_{ij}^{kl} = \left(t_j^z + r_j^z\right)\left(m_j^l - m_j^k\right) - \left(t_i^z + r_i^z\right)\left(m_i^l - m_i^k\right), \qquad (10)$$

with partial derivatives:

$$\frac{\partial T_{ij}^{kl}}{\partial r_j} = m_j^l - m_j^k \text{ and } \frac{\partial T_{ij}^{kl}}{\partial r_i} = m_i^k - m_i^l ,$$

which are the differential mapping functions.

Condition of the Normal Equations

Previous studies of estimating the tropospheric delay from GPS data (e.g. *Van Hove et al.* [1993]) have highlighted the problem of using differenced data over short baselines. For this situation, there exists a strong mathematical correlation between the partial derivatives of the tropospheric delay at the two stations. For baseline lengths up to several hundred kilometres the elevation angles to a particular satellite will be similar and hence so will the partial derivatives (but with opposing signs). Even if the meteorological conditions are drastically different at the ends of such a baseline, it is difficult for a least-squares model to separate the two contributions.

The usual technique to overcome this problem is known as *levering* [*Rocken et al.,* 1995] and works by simply fixing the tropospheric delay at the reference station and estimating the *relative* delay at the secondary station. Our use of real-time meteorological data at the reference station will help minimise the error in the estimated residual delay. While some error will be present, it will be constant for all solutions computed with different tropospheric delay models at the aircraft.

This problem of estimating the tropospheric delay over short baselines is compounded by the fact that the determination of the height component of position is sensitive to the existence of any unmodelled tropospheric zenith delay and vice-versa. To explain this we can consider the observation equation for a single station/satellite measurement contributing to the double-difference observable. Ignoring the ionospheric and integer ambiguity terms (which we can remove), we have:

$$\Phi_i^k = \rho_i^k + T_i^k = \left[\left(E^k - E_i\right)^2 + \left(N^k - N_i\right)^2 + \left(H^k - H_i\right)^2\right]^{\frac{1}{2}} + t_i^z \cdot m_i^k \qquad (11)$$

where the receiver and satellite coordinates are defined in the local geodetic coordinate system (easting, northing and height). Linearising around approximate values and subtracting from the observation gives:

$$\Delta\rho = \Phi - \rho^0 - T^0 \doteq \Delta E_i \cdot \cos(e) \cdot \sin(a) + \Delta N_i \cdot \cos(e) \cdot \cos(a) + \Delta H_i \cdot \sin(e) + \Delta t_i^z \cdot \csc(e) , \qquad (12)$$

where the zero superscripts indicate predicted values; Δ represents the small corrections to the a-priori estimates; and *a* and *e* represent the azimuth and elevation angle of the satellite. The mapping function is approximated with the cosecant of the elevation angle. If we consider that the corrections to the horizontal position components ΔE and ΔN are largely decoupled from the height and tropospheric delay components, then we can see how the presence of height and zenith delay errors affect the retrieval of each other.

Figure 1 shows the effect on the modelled range difference of a zenith delay error of 2 mm, a height error of 200 mm and their combined effect ($\Delta\rho$) from the zenith down to an elevation angle of five degrees. Additionally, we show what happens when we try to recover the height and zenith delay components assuming (incorrectly) that the contribution of the other component is zero. We are able to accurately recover the height component down to an elevation angle of

approximately thirty degrees, beyond which the accuracy increasingly degrades. At the same time we are only able to recover well the zenith delay error at the low elevation angles. At very high elevation angles an error in the tropospheric zenith delay is almost indistinguishable from the unmodelled height component.

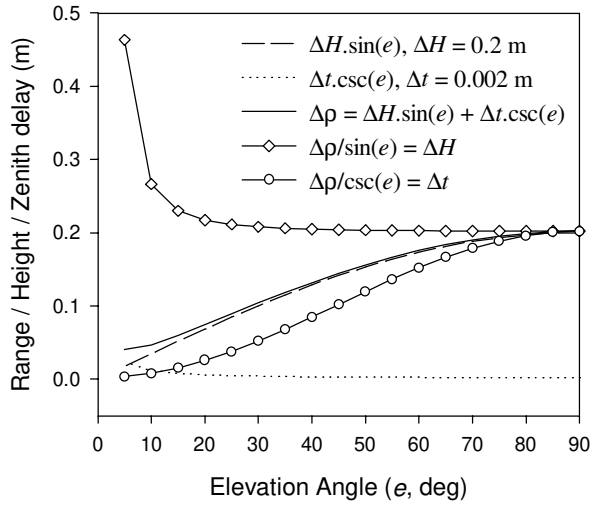

Figure 1. Simulation of the impact and recovery of height and zenith delay components on a GPS range.

What this means is that an *unmodelled* tropospheric zenith delay error causes an error in height determination, which increases with the inclusion of lower elevation data. This is a well known fact in GPS, but importantly for us we can see that attempting to solve for the zenith delay is hindered without low elevation angle data. These results will be modified if tight constraints are placed on the station height components in the least-squares adjustment. By closely constraining the height to its known value, the least-squares model is better able to estimate the tropospheric zenith delay.

IMPLEMENTATION AND DATA PROCESSING

A least-squares positioning model using double-differenced, dual-frequency, GPS carrier phase observations is implemented in the KARS processing software [*Mader*, 1996]. The code has been modified at UNB to allow for the estimation of the tropospheric delay as either a scale factor or a zenith delay residual at either the secondary roving receiver or at both the rover and the reference receiver.

To test the estimation of the residual delay, some of the flight data from the St. John's, Newfoundland–based Frizzle '95 campaign was used (see *Collins and Langley* [1997] for more details). Meteorological data consisting of pressure, temperature and relative humidity were simultaneously recorded along with dual frequency GPS data at an aircraft and at a ground reference station. Both data types are available at a two second sampling interval.

It has been possible to use the carrier phase data down to an elevation angle of five degrees on one of the data sets. Of the data recorded on the other days of this experiment, there are too many cycle slips and data gaps for the software to adequately process the data. Figure 2 shows the flight path over which the data was recorded on March 3rd 1995. The maximum distance reached from the reference station was 210 kilometres. As Figure 2 indicates, data collection was halted in mid-flight. This was due to a lack of memory in the aircraft receiver.

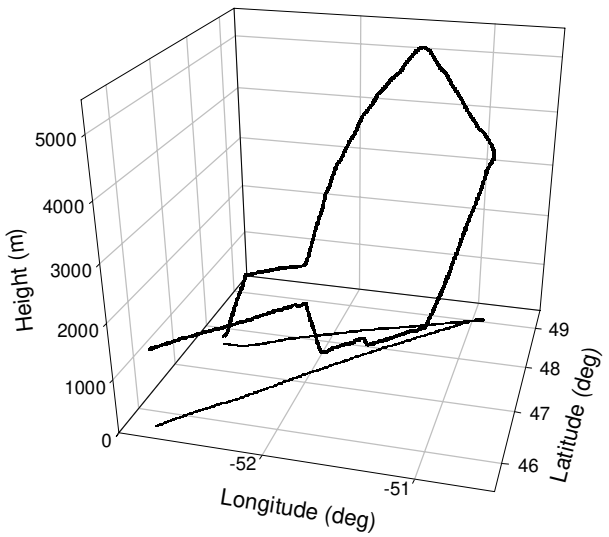

Figure 2. Flight path of aircraft for data used in this study.

A set of fixed carrier phase integer ambiguities for all satellites on both frequencies was derived. This was done by processing the flight data at various elevation cut-off angles while resolving the ambiguities "on-the-fly". Comparing the ambiguities from these solutions with ambiguities computed for the short static period on the ground before the flight, has enabled stable sets of integers to be selected. While confident that these are the correct values, without actually estimating these values in flight (in a similar manner to *Sonntag et al.* [1995]), we can only confirm this by examining the residuals of the position solutions to see if they diverge over time.

Of the remaining error sources, the primary one is the satellite position error. To minimise it as much as possible, International GPS Service for Geodynamics (IGS) precise orbits were used. This leaves multipath and noise as remaining unmodelled errors which should be of the order of centimetres or less for the carrier phase observable.

RESULTS

One tropospheric zenith delay and mapping function combination was adopted for the reference station for all the solutions. These were the *Saastamoinen* [1973] zenith delays using the simultaneously recorded meteorological data and the mapping functions of *Niell* [1996], which only require position and day-of-year information.

This combination was also used at the aircraft and provided with the simultaneously recorded meteorological data. This model is denoted as SAANf ('f' for full-met. input). The model denoted as SAANx ('x' for extrapolated) used the reference station meteorological data scaled to the height of the aircraft using standard atmospheric scaling equations (see *Collins and Langley* [1997]); and the SAANh model ('h' for hydrostatic only) which predicted only the hydrostatic delay at the aircraft from the real-time data while the wet delay was set to zero.

Three other models were also tested for modelling the delay at the aircraft: UNB4, which supplies meteorological data based on the 1966 U.S. Standard Atmosphere Supplements to the Saastamoinen and Niell algorithms; the initially proposed WAAS model; and the NATO recommended model (see *Collins and Langley* [1997] for more details).

One solution was computed for each of these models estimating the three-dimensional Cartesian positions of the aircraft along with the residual tropospheric delay as a scale factor at each epoch. No filtering was applied and no a-priori constraints were placed on any of the parameters. Each epoch provided an independent solution.

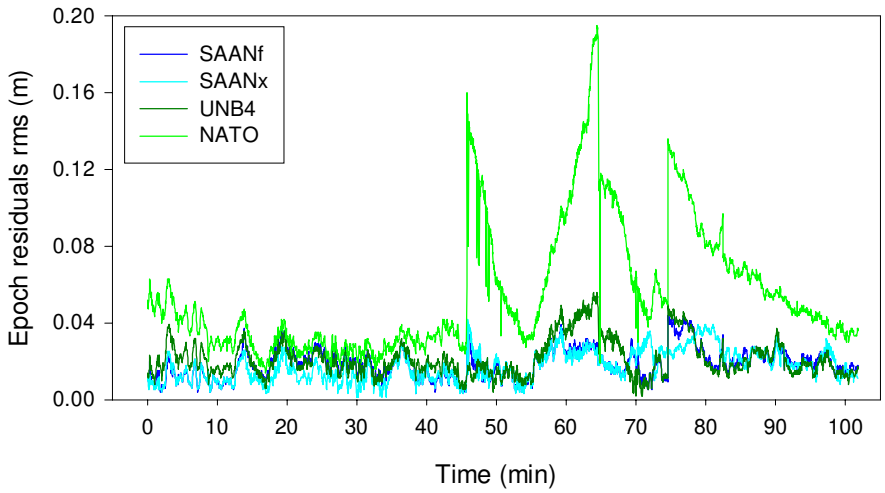

Figure 3. Root-mean-square double-difference carrier phase residuals without residual tropospheric delay estimation.

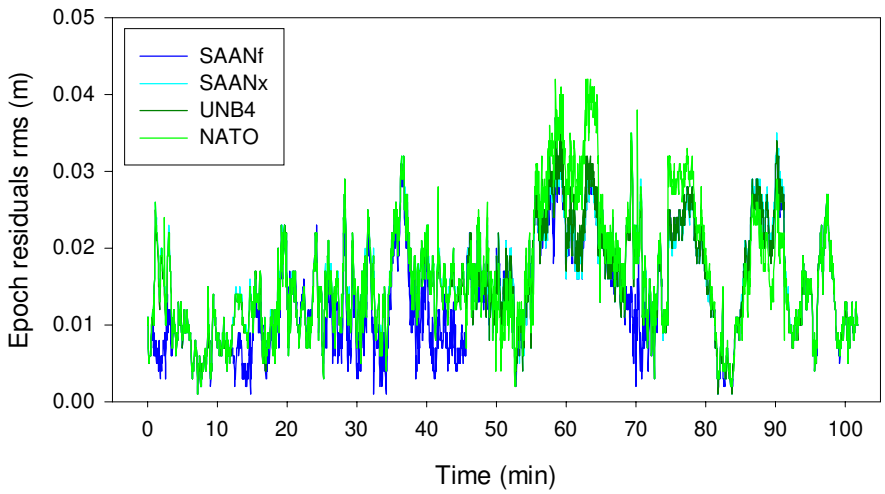

Figure 4. Root-mean-square double-difference carrier phase residuals with residual tropospheric delay estimation.

Adjustment Residuals

Considering first of all the double-differenced carrier phase residuals after the least-squares adjustment, comparison of Figure 3 and Figure 4 shows the general improvement gained by estimating a residual tropospheric delay parameter. Only four of the models are plotted for clarity. Figure 3 represents the root-mean-square (rms) of the residuals computed without estimating the residual tropospheric delay. Almost every plot has a distinct trace indicating the impact of each tropospheric delay model. The exception is for the traces of the SAANf and UNB4 solutions, which closely follow each other. In Figure 4 however, all the traces have merged to become almost indistinguishable, indicating that estimation of the residual tropospheric delay has largely removed the impact of the choice of a particular model on the solution (note also the change of scale).

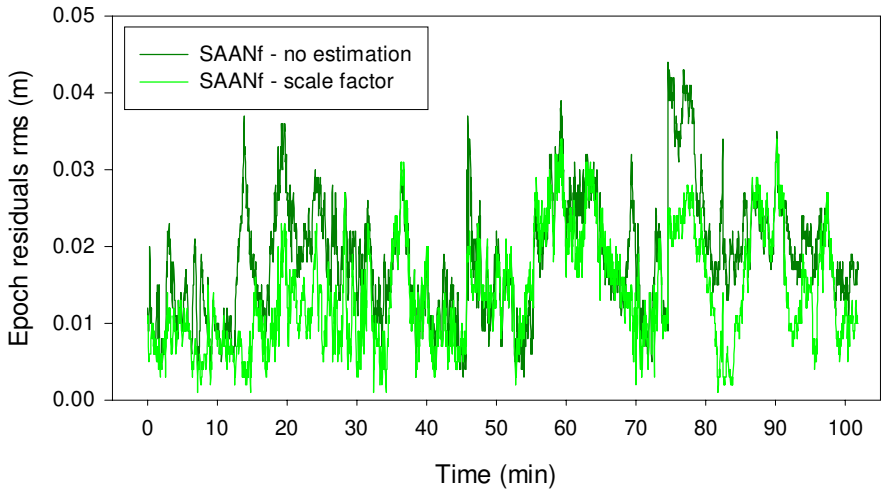

Figure 5. Root-mean-square double-difference carrier phase residuals with and without residual tropospheric delay estimation — SAANf model used with aircraft data.

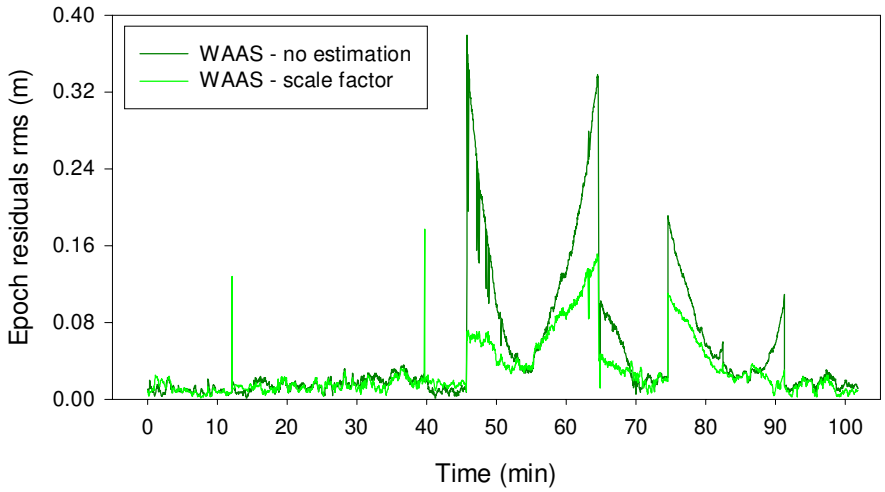

Figure 6. Root-mean-square double-difference carrier phase residuals with and without residual tropospheric delay estimation — WAAS model used with aircraft data.

Examining two models more closely, Figure 5 shows the rms residual for each epoch for the SAANf model. In general, only a small improvement has been made in estimating a residual correction. This is to be expected because the tropospheric delay prediction of this model is fairly good due to the use of meteorological measurements recorded at the aircraft. A better indication of the improvement possible with estimating the residual delay is gained from examining the impact of the initially-proposed WAAS model on the solution. This model was left out of Figure 3 and Figure 4 because of its large impact. This can be seen in Figure 6 where large jumps correspond to the rising and setting of satellites at 5 degrees elevation angle. This model uses only a modified cosecant of the elevation angle mapping function and consequently large errors would be expected

with these satellites. As this plot shows, some, but not all, of the error has been absorbed by estimating the residual tropospheric zenith delay. In addition, the spikes visible in this plot at approximately 12 and 40 minutes into the flight are an artifact of a discontinuity in the formulation of this model. This occurs whenever the aircraft crosses the 1500m height level (see *Collins and Langley* [1997] for further discussion of this feature).

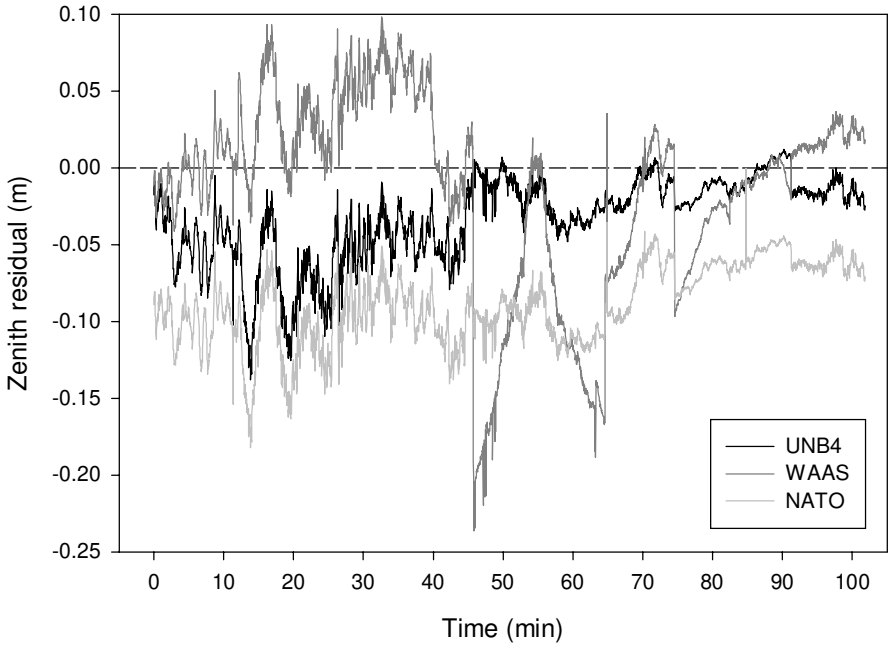

Figure 7. Residual tropospheric delay estimates for models with no real-time meteorological data.

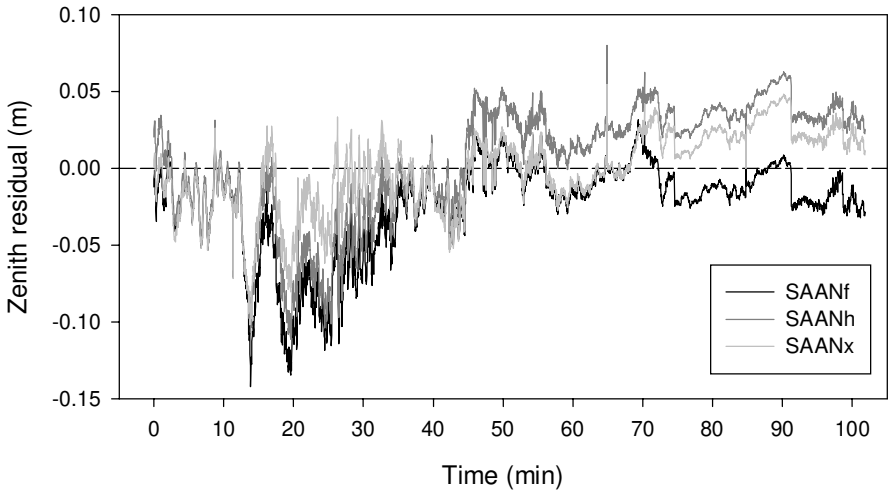

Figure 8. Residual tropospheric delay estimates for models with real-time meteorological data.

Residual Delay Estimates

Turning to the actual residual delays estimated, Figure 7 shows the values for the models without real-time meteorological input and Figure 8 for those with real-time input. Both plots can be considered in two halves — before and after the 45 minute epoch. Consideration of Figure 9 shows that before this point in the flight there are no satellites at low elevation angles (< 10 degrees). As indicated previously, this limits the potential for adequately estimating the tropospheric delay. As an example, given that the SAANh model predicts only the hydrostatic delay we would expect positive zenith residuals to represent the remaining wet delay. This is generally true in the second half of Figure 8, but the wide variation in the first half, coupled with the large negative values could mean that the residual estimates for this time span are unreliable.

At the same time however, it is interesting to note from Figure 7 that the residual estimates for the NATO model solution have an almost constant bias component over the

whole time span. This trend is what we would expect given that the NATO model is formulated with a constant value of surface refractivity which, to a first order consideration, is biased from the real surface refractivity experienced over the flight path. Both of the other models used in the solutions shown in Figure 7 change their atmospheric parameter inputs primarily as a function of latitude, hence there is no constant bias in their solutions.

Considering the results for the other models, Figure 7 shows that UNB4 gives the smallest residual tropospheric delay for the models without real-time meteorological input; and that the initially-proposed WAAS model is greatly influenced by low elevation satellites in the solution. Given the correct zenith delay, the initially-proposed WAAS model will over-predict the delay at low elevation angles. This can be seen at ~45 minutes when a new satellite appears. The residual estimate jumps to a large negative value to try to account for the over-prediction. The residual then increases toward zero as the satellite rises in the sky.

Figure 8 shows that the full-meteorological model, SAANf, has the smallest residual delay, which is what we would hope for. The extrapolation of surface values introduces some bias (model SAANx), but the largest residual delay is observed when estimating the whole wet zenith delay (model SAANh).

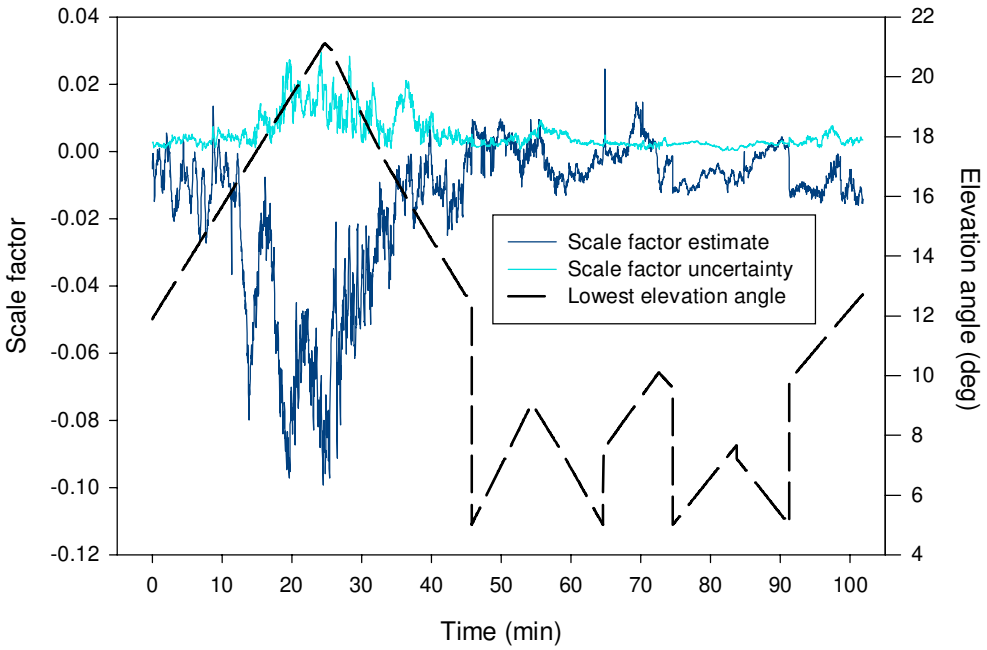

Figure 9. Comparison of SAANf residual tropospheric delay estimates and their uncertainty with the lowest elevation angle used in the solution.

Position Differences

Turning to the impact of the residual tropospheric delay estimation on the position determination, we can first of all consider the residual delay that remains when using real-time meteorological data. Without residual delay estimation, we would consider the SAANf solution to be the "best" obtainable because of its realistically-modelled zenith delays and mapping functions driven by real-time meteorological data. By estimating the residual delay we would hope to model any deviations from the average atmospheric structure implied by these models. Figure 10 shows the difference in the position components for solutions computed using the SAANf model with and without residual tropospheric delay estimation. The difference in the height component is considerable: of the two sets of statistics for this data, even when considering only the "good" estimates after the 45 minute epoch, there is a mean bias in height of ~5 cm with an rms of ~9 cm.

In addition, using the solution with the SAANf model and residual tropospheric delay estimation as a benchmark, we can compare the impact of estimation with other models and confirm that estimation of the residual delay helps to remove the impact of less accurate models. Figure 11 shows the position differences of the solution computed using UNB4 at the aircraft and Figure 12 shows the influence of using the NATO model at the aircraft. The biases in these two plots are predominantly a function of the lowest elevation angle (cf. Figure 9).

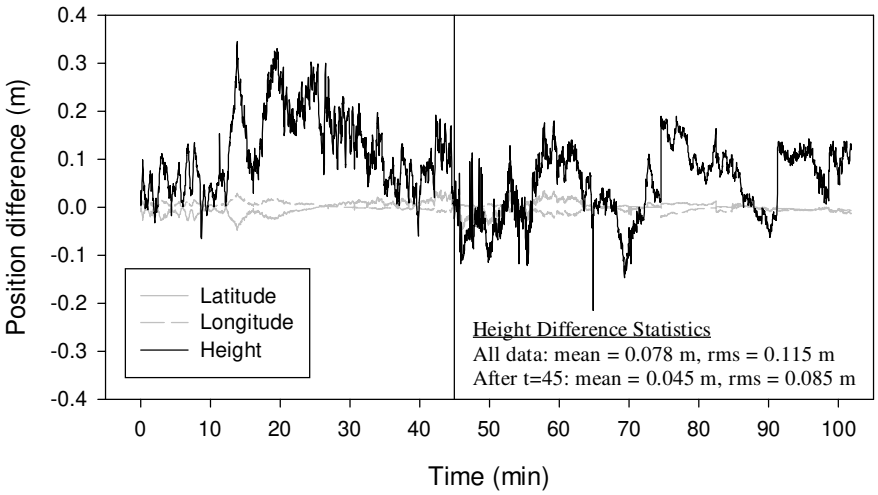

Figure 10. Difference in position solutions with and without residual delay estimation from predictions with real-time meteorological data (SAANf model).

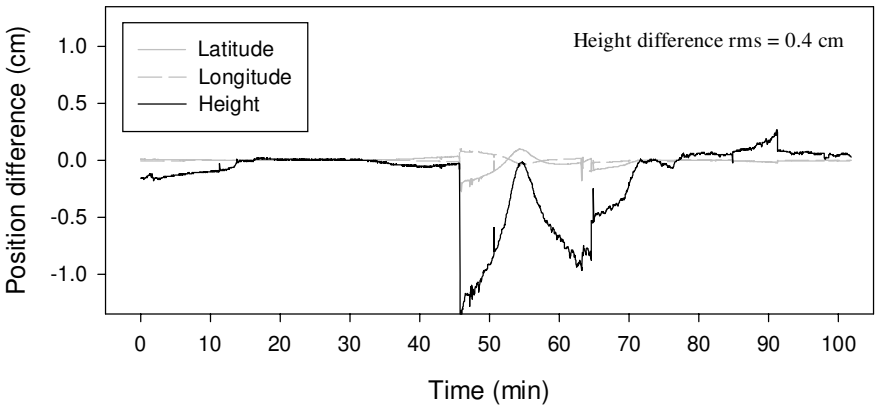

Figure 11. Position differences between the UNB4 solution and the SAANf solution. (Residual delay estimated in both solutions.)

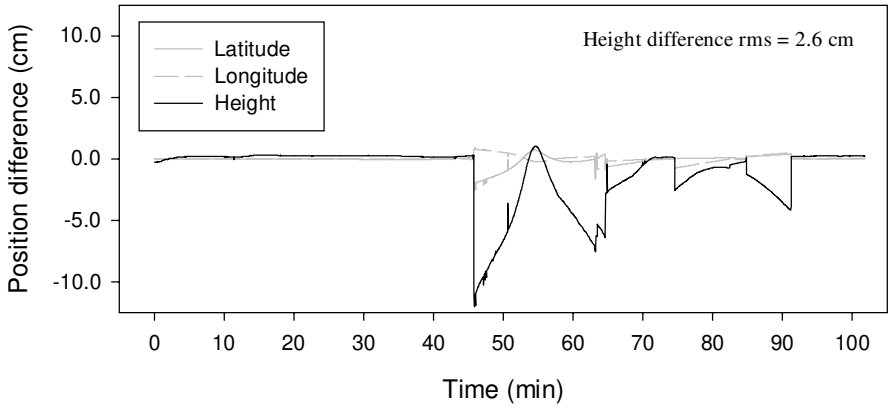

Figure 12. Position differences between the NATO solution and the SAANf solution. (Residual delay estimated in both solutions.)

Where there are no low elevation satellites to magnify the errors, all three solutions compare at the millimetre level in all three position components. Even so, the impact of using the UNB4 model instead of real-time meteorological data is still only of the order of 1 cm in height at maximum. The impact of the NATO model is one order of magnitude larger.

CONCLUSIONS

We have attempted to show in this paper the effect of implementing residual tropospheric delay estimation from GPS data recorded at an aircraft in flight. The aim was to remove any unmodelled effects of the troposphere that cannot be predicted by empirical models, even when using meteorological measurements of the ambient atmosphere.

Problems were encountered with the general stability of the least-squares solution when attempting to estimate the residual delay at the reference station as well as at the aircraft. The similarity of the normal equation coefficients prevents the separation of the contribution of the atmosphere at the two stations in the double-difference observable. By estimating only the residual delay at the aircraft we are assuming that there is no residual effect at the reference station, or that this effect is absorbed by the estimate for the aircraft.

Estimating the residual delay appeared to almost wholly remove the impact of a particular tropospheric delay model. However, the accuracy of the mapping function and the impact of the satellite geometry are important. It appears crucial that there exists data at low elevation angles (less than 10 degrees) for the tropospheric residual estimate to be meaningful. With this condition, plus an accurate mapping function, estimating the residual delay not only removes the impact of one particular tropospheric delay model but also the biases from using non-real-time meteorological parameter values at the aircraft. Therefore, if the highest possible precision is required for aircraft positioning the estimation of a residual delay should be considered, otherwise biases of the order of ten centimetres may be present in the solution.

The work reported in this paper has been only a preliminary study and further investigations are required to study the condition of the least-squares normal equations and the overall reliability of the technique. New investigations could include the impact of antenna phase centre corrections, as the data is particularly sensitive to these at low elevation angles. Additionally, the implementation of a Kalman or other type of constraining least-squares filter could significantly enhance the technique by providing some a-priori constraints to the estimates.

ACKNOWLEDGMENTS

The support of Nav Canada (formerly Transport Canada Aviation), the Natural Sciences and Engineering Research Council of Canada, the Atmospheric Environment Service and the U.S. Federal Aviation Administration are gratefully acknowledged.

REFERENCES

Collins, J.P. and R.B. Langley (1997). *A Tropospheric Delay Model for the User of the Wide Area Augmentation System*. Final contract report prepared for Nav Canada, Department of Geodesy and Geomatics Engineering Technical Report No. 187, University of New Brunswick, Fredericton, N.B., Canada.

Mader, G.L. (1996). "Kinematic and rapid static (KARS) GPS positioning: Techniques and recent experiences". *IAG Symposia No. 115*, Eds. G. Beutler, G.W. Hein, W.G. Melbourne and G. Seeber. IUGG/IAG, Boulder, Colo., 3-4 July. Springer-Verlag, Berlin, pp. 170-174.

Niell, A.E. (1996). "Global mapping functions for the atmosphere delay at radio wavelengths." *Journal of Geophysical Research*, Vol. 101, No. B2, pp 3227-3246.

Rocken, C., T. Van Hove, J. Johnson, F. Solheim and R. Ware (1995). "GPS/STORM — GPS sensing of atmospheric water vapour for meteorology." *J. Atmos. Oceanic Technology*, Vol. 12, pp. 468-478.

Saastamoinen, J. (1973). "Contributions to the theory of atmospheric refraction." In three parts. *Bulletin Géodésique*, No. 105, pp. 279-298; No. 106, pp. 383-397; No. 107, pp. 13-34.

Sonntag, J.G., C.F. Martin and W.B. Krabill (1995). "Ambiguity resolution over long baselines for airborne differential GPS positioning." *Proceedings of ION GPS-95*, Palm Springs, Calif., September 12-15, pp. 1117-1125.

Van Hove, T.M., C. Alber and J.M. Johnson (1993). "Atmospheric water vapour as noise and signal for Global Positioning System applications." *Proceedings of ION GPS-93*, Salt Lake City, Utah, September 22-24, pp. 797-804.