

INTEGRATION OF RAW GPS MEASUREMENTS INTO A BUNDLE ADJUSTMENT

Cameron ELLUM

Mobile Multi-sensor Systems Research Group
Department of Geomatics Engineering, University of Calgary, Calgary, AB, Canada
cmellum@ucalgary.ca

KEY WORDS: Photogrammetry, Mapping, Bundle Adjustment, GPS Integration, Software

ABSTRACT

GPS data is typically included in photogrammetric adjustments as externally processed position observations. This implementation has obvious benefits in its simplicity; however, a more fundamental fusion of the GPS data into the bundle adjustment is possible. In this paper, an investigation is made into the inclusion of GPS pseudoranges directly into a photogrammetric bundle adjustment. The advantages of the technique include improved accuracy and reliability, and the ability to use GPS data when less than four satellites are available. Notes are made regarding pseudorange errors and their mitigation using atmospheric models, linear-combinations, and precise orbits and clock corrections. Using the new technique, tests are performed with aerial GPS/photogrammetric data that demonstrate that the method provides accuracies that are superior to those obtained when exposure station position observations are used in the adjustment. The paper concludes with some notes regarding the design and implementation of the combined GPS/photogrammetric adjustment, with an eye towards maintainability, extensibility, and performance. The hierarchical structure of the program is described, and the benefits of an object-oriented design using inheritance and polymorphism are outlined.

1 INTRODUCTION

Kinematic GPS controlled aerial photogrammetry has become an omnipresent technology in both the scientific and commercial mapping communities. Virtually all airborne mapping systems now integrate a GPS receiver with their camera. This integration is done at the hardware level, as the GPS receiver and camera must communicate, either for the GPS to trigger the camera or for the camera to record the exposure time. Unfortunately, on the software side, the integration of GPS and photogrammetry is not as close. Typically, the GPS data is included in the photogrammetric bundle adjustment only as processed positions (Ackermann, 1992; Greening et al., 1994; Mikhail et al., 2001). In effect, the GPS and photogrammetric processing engines operate largely in isolation. This implementation has obvious benefits in its simplicity; however, a more fundamental fusion of the GPS data into the bundle adjustment may provide improvements in both accuracy and reliability.

This paper outlines a tighter coupling of the GPS and photogrammetric processing engines where the GPS code pseudoranges are directly included in the bundle adjustment. The goal of this integration is to improve the accuracy and reliability when compared to the naïve inclusion of GPS positions.

2 THEORETICAL FOUNDATIONS

In the following section, a brief theoretical background is given on both the current technique for including GPS data in a photogrammetric bundle adjustment, and on the altered technique being pursued as part of this project.

2.1 Existing Technique for Including GPS Data

With some rare exceptions (for example, Kruck et al., 1996) GPS data is almost always included in photogrammetric adjustments as processed positions. In other words, the raw GPS measurements are first processed using an external processing program that provides position and covariance estimates. These positions are then included in the adjustment using parameter observation equations. The nominal form of these equations is (Mikhail et al., 2001)

$$\mathbf{r}_{GPS}^M(t) = \mathbf{r}_c^M(t) + \mathbf{R}_c^M(t)\mathbf{r}_{GPS}^c + \left(\mathbf{b}_{GPS}^M + \mathbf{d}_{GPS}^M(t - t_0)\right), \quad (1)$$

where $\mathbf{r}_{GPS}^M(t)$ is the position of the GPS antenna, $\mathbf{r}_c^M(t)$ is the position of the camera perspective centre, $\mathbf{R}_c^M(t)$ is the rotation matrix that aligns the camera axes to the mapping space axes, and \mathbf{r}_{GPS}^c is the offset between the GPS antenna and camera perspective centre. The bias and drift terms $-\mathbf{b}_{GPS}^M$ and \mathbf{d}_{GPS}^M respectively – are included as unknown parameters in the adjustment and are intended to account for the errors caused by incorrect GPS ambiguity resolution. These terms can also account for datum inconsistencies.

2.2 Modification of the Collinearity Equations

In Ellum (2001), an alternative technique was investigated for including exposure station position observations in the photogrammetric adjustment. Derivation of the relevant equations begins with the forward conformal transformation that relates the GPS positions with the image co-ordinates,

$$\mathbf{r}_P^M = \mathbf{r}_{GPS}^M(t) - \mathbf{R}_c^M(t)\mathbf{r}_{GPS}^c + \mu\mathbf{R}_c^M\mathbf{r}_p^c. \quad (2)$$

As in equation (1), this transformation requires the $\mathbf{R}_c^M(t)$ rotation matrix that aligns the camera axes to the mapping space axes and the \mathbf{r}_{GPS}^c offset vector between the GPS antenna and camera perspective centre.

By rearranging Equation (2), the reverse transformation is found to be

$$\mathbf{r}_p^c = \mu^{-1} [\mathbf{R}_M^c (\mathbf{r}_P^M - \mathbf{r}_{GPS}^M) + \mathbf{r}_{GPS}^c]. \quad (3)$$

Elimination of the third equation yields a pair of modified collinearity equations,

$$x_p = -c \frac{r_{11}(X_P - X_{GPS}) + r_{12}(Y_P - Y_{GPS}) + r_{13}(Z_P - Z_{GPS}) + x_{GPS}}{r_{31}(X_P - X_{GPS}) + r_{32}(Y_P - Y_{GPS}) + r_{33}(Z_P - Z_{GPS}) + z_{GPS}} \quad (4a)$$

$$y_p = -c \frac{r_{21}(X_P - X_{GPS}) + r_{22}(Y_P - Y_{GPS}) + r_{23}(Z_P - Z_{GPS}) + y_{GPS}}{r_{31}(X_P - X_{GPS}) + r_{32}(Y_P - Y_{GPS}) + r_{33}(Z_P - Z_{GPS}) + z_{GPS}}. \quad (4b)$$

By examining Equation (4), it can be seen that the exposure station positions are no longer explicitly present in the collinearity equations, and that essentially, the GPS positions form the ‘base’ of the equations. This has a number of advantages. First, the GPS positions can be directly used as the initial approximates in the linearised collinearity equations. Second, because the GPS positions are one of the quantities being adjusted, the position measurements can be directly used as parameter observations. In this case, the parameter observation equation is

$$\mathbf{0} = \mathbf{r}_{GPS}^M - \hat{\mathbf{r}}_{GPS}^M, \quad (5)$$

where $\hat{\mathbf{r}}_{GPS}^M$ represents the current estimate of the position during the adjustment. Adjusting the GPS positions directly also means that they are one of the quantities output by the adjustment. This allows for easy comparison with the input positions, which in turn simplifies the analysis of the results. Finally, expressing the collinearity equations as a function of the GPS positions means that the inclusion of the raw GPS pseudorange and phase measurements in the adjustment can be done with greater ease than would otherwise be possible, as such measurements are also functions of the GPS positions. This last point provides the motivation for the project under investigation in this paper.

2.3 Inclusion of the Pseudorange Measurements in the Photogrammetric Adjustment

With the collinearity equations expressed by Equation (4), inclusion of the GPS pseudorange measurements in the photogrammetric adjustment is straightforward. Essentially, simple observation equations with the form

$$p = |\mathbf{r}_{GPS/SV}| + c\Delta t_{rx} \quad (6)$$

are added to the adjustment, where p is the GPS pseudorange measurement, c is the speed of light, and Δt_{rx} is the receiver clock bias. This last term is added to the adjustment as an unknown parameter.

Including the pseudorange measurement in the adjustment should improve the mapping accuracy and, more importantly, reliability. Furthermore, it enables GPS data to be used when less than four satellites are visible. While this is not generally an issue for aerial mapping platforms, it could be beneficial for terrestrial mobile mapping systems.

3 GPS ERRORS

Equation (6) assumes that the only error in the code pseudorange measurement is random noise. In reality, of course, this is not the case. There are a number of systematic errors present in the observations, and when the largest of these are accounted for the pseudorange observation equation more completely resembles

$$p = |\mathbf{r}_{GPS/SV} + \delta\mathbf{r}_{SV}| \quad (7)$$

$$- c(\Delta t_{sv} - \Delta t_{rx}) + d_{iono} + d_{tropo}. \quad (8)$$

In the above equation, $\delta\mathbf{r}_{SV}$ is the error in the satellite coordinates, Δt_{sv} is the satellite clock bias, d_{iono} is the ionospheric delay, and d_{tropo} is the tropospheric (or neutral atmosphere) delay. These error sources and the techniques to mitigate them are well documented – see, for example, Hofmann-Wellenhof et al. (1994). However, for completeness these errors and the specific steps taken in this project to mitigate them are detailed below.

3.1 Satellite Position Errors

The GPS satellite position errors are commonly divided into along-track, across-track, and radial components. Due to the great distance to the satellites, the former two error components do not significantly project onto the measured ranges. Thus, they can, effectively, be disregarded. The radial error, however, directly impacts ranges observed from the satellites and must therefore be examined. Figure 1 shows the radial satellite position errors for the entire GPS constellation during a one-week period in May of 2003. For this period, the root-mean-square (RMS) radial error is just over 1.1m and the maximum error is close to 5m.

The satellite position error, radial or otherwise, is easily corrected for using precise ephemerides. Precise ephemerides are observed or predicted orbits that are freely available from a number of organisations that include the United States’ National Imagery and Mapping Agency (NIMA) and the International GPS Service (IGS). In the case of the latter, several products are available, with accuracies ranging from 25cm for predicted orbits to better than 5 cm for observed orbits with a two-week latency. Either accuracy is well below the expected accuracy of the pseudorange measurements. Precise ephemerides typically have a sample interval of 15 minutes. To determine satellite positions between samples, polynomial interpolation is normally used (Hofmann-Wellenhof et al., 1994). A seventh-order interpolator is typically sufficient.

It should be noted that it is also possible to view the GPS satellite position errors as errors in the control points defining the photogrammetric network datum, instead of grouping them with the other range errors as was done here. This

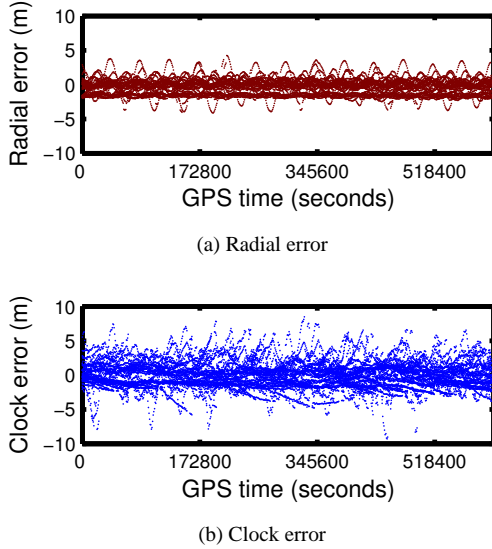


Figure 1: Satellite radial position and clock errors (Entire constellation, GPS week 1217)

implies that the satellite positions could be weighted in the adjustment instead of being fixed, just as the GPS exposure station position observations are weighted in contemporary GPS-controlled aerial photogrammetry.

3.2 Satellite Clock Biases

A satellite clock bias will manifest itself entirely as a range error. Most of the satellite clock biases can be removed using correction coefficients broadcast as part of the satellite ephemeris. The residual error that remains, however, can still be significant. Figure 1(b) shows the difference between the broadcast and precise satellite clock corrections. When compared with Figure 1(a), it can be seen that the residual satellite clock error is larger than the radial satellite orbit errors. For the one week period shown in the figure, the RMS clock error for the entire constellation was just under 2m, and the maximum error was close to 10m.

To correct for the residual satellite clock errors, precise clock corrections can be used. Such corrections are normally included with precise ephemerides, and, just as with precise ephemerides, polynomial interpolation can be used to determine the correction between sample epochs. Because the satellite clocks are very stable, a lower-order interpolator can be used than that for the positions. However, regardless of the interpolator order, with 15-minute clock corrections maximum errors of close to a metre may still occur. Fortunately, clock corrections at a 5-minute sample interval are also available from the IGS, and when these higher-rate corrections are used with a third-order interpolator, the maximum errors can be reduced to under half a metre.

3.3 Ionospheric Delays

Essentially, the only option for dealing with the ionospheric delays is to use the ionospheric-free linear combination

to eliminate the first-order effects. The only other convenient option is to use the broadcast ionospheric prediction model to estimate its effect. However, the broadcast model only removes 50%-60% of the error, and can leave maximum errors of some tens of meters.

Figure 2 shows the noise in the ionospheric-free linear combination, as determined by differencing the code and phase measured ionospheric error and removing the mean difference. When such a difference is performed, all common errors are eliminated, and what remains is predominantly code multipath and receiver noise. For the satellite depicted, the noise for the near-zenith measurement was about 0.3m. As the elevation of the satellite decreases, the noise increases, following a relationship that – until about 15° elevation – can roughly be described by

$$\varepsilon(e) = \frac{\varepsilon(90^\circ)}{\sin(e)}, \quad (9)$$

where $\varepsilon(90^\circ)$ is noise at the zenith angle, and e is the satellite's elevation angle. This simple cosecant relationship was used in this project to estimate the variances of the ionospheric-free pseudorange observations.




Figure 2: Measured code ionospheric delay noise (SV 1, day 2 of GPS Week 1217)

3.4 Tropospheric Delays

The errors due to the tropospheric delay are typically mitigated using a combination of zenith-delay models and mapping functions. The tropospheric models may use surface measurements of temperature, pressure, etc., or they may use standard empirical values. In this project, the UNB2 tropospheric model was used in conjunction with the Niell mapping function (Collins et al., 1996; Niell, 1996). When the tropospheric delay is corrected for using a model, a residual tropospheric delay (which is generally elevation-dependant) will still remain. Any residual delay common to all satellites will, however, be compensated for in the estimate of the receiver clock offset.

4 TESTING AND RESULTS

The direct inclusion of the GPS pseudoranges in the bundle adjustment was tested using a block of imagery captured using a medium resolution digital camera with an image size of 4096×4096 pixels. The block consisted of 42 images collected from 7 parallel flight lines at a flying height of roughly 900m. Fifty-three well-distributed check points

were available for comparing with the adjustment output. Dual-frequency GPS observations at 2Hz were collected along with the imagery. Additionally, a dual-frequency base station in the centre of the block collected GPS observations at 1Hz.

There were several problems with the data set that complicated the generation and analysis of results. Foremost among these was that only orthometric heights were available for the check points. Because an accurate geoid model for the test region was unavailable, these heights could not be converted into ellipsoidal heights compatible with the GPS heights determined in the adjustment. In an admittedly imperfect solution, the vertical datum shift was solved for in an adjustment that treated all the check points as control points and used exposure station position observations generated from the best possible dual-frequency carrier-phase GPS solution (to solve for the datum shift it is necessary to constrain both datums). In addition to the large vertical datum shift, it was felt that there may also have been small horizontal datum shifts. These were not solved for, and, if present, contribute to the mean errors seen in the results presented below. An additional problem with the data set was that the lens distortion available for the camera was not in a format compatible with the adjustment software used. Consequently, the lens distortion was calibrated for using the same adjustment that solved for the vertical datum shift. This may mean that the standard deviations in the results are somewhat optimistic as the camera may ‘fit’ the data better than it should.

Before looking at the results available when the GPS measurements are included in the adjustment, it is worthwhile to get some idea of the noise within the network. Table 1 shows the results from a conventionally controlled adjustment where approximately one-third of the check points were used as control points. The remaining check points were used to calculate the statistics in the table. These results should be an indication of the best possible accuracy available from the network.

Table 1: Check Point Error Statistics (m): Control Points

	Horizontal	Vertical
Mean	0.18	-0.19
Std. dev.	0.09	0.45
RMSE	0.20	0.49
Absolute maximum (mean removed)	0.27	1.00

The comparison of results will primarily be done using the standard deviations of the check point errors. This in acknowledgement of the fact that a mean error – primarily due to unmodelled tropospheric delays – will almost certainly be present in the networks determined using the undifferenced GPS pseudoranges. It may be tempting to believe that the GPS errors would ‘average out’ over the entire block. Unfortunately, because of the relatively short timespan in which the imagery was captured, the errors at the individual GPS stations will be highly correlated (during this time period, the troposphere and satellite positions do

not change significantly). The common errors among GPS stations will cause the entire network to translate.

Finally, it should be emphasised that in the tests that follow, no ground control points are used. The networks are controlled entirely by the GPS measurements.

4.1 Broadcast Orbits and Clocks

The first tests were performed using the broadcast satellite orbits. Table 2 contains the results for when the pseudoranges are included directly in the adjustment. Notably, the standard deviations of the check point errors are only slightly worse than those in Table 1. In other words, directly including the GPS pseudoranges in the adjustment yields object space accuracies that are comparable to those obtained from the same network controlled via well-distributed ground control points. This is a promising first result; however, it must be restated that the efforts made to overcome difficulties with the data may mean that this result is somewhat optimistic.

Table 2: Check Point Error Statistics (m): Pseudorange observations, Broadcast Orbits

	Horizontal	Vertical
Mean	0.98	3.06
Std. dev.	0.21	0.47
RMSE	1.00	3.09
Absolute maximum (mean removed)	0.46	1.25

Of course, rather than being directly integrated into the bundle adjustment, the pseudorange measurements can also be used to generate single-point exposure station positions. These positions could then be added to the adjustment as position observations in the typical fashion (see 2.1). Table 3 shows the results for when the network is controlled using such positions. By comparing the results in this table with those in Table 2, it can be seen that directly including the pseudoranges in the adjustment yields object space accuracies that are about 30% better than when single-point position observations are used. Both approaches use exactly the same data, but the closer integration that comes from directly including the pseudoranges in the adjustment leads to a significant improvement in accuracy.

Table 3: Check Point Error Statistics (m): Single-point position observations, Broadcast Orbits

	Horizontal	Vertical
Mean	1.09	3.07
Std. dev.	0.35	0.69
RMSE	1.15	3.14
Absolute maximum (mean removed)	0.63	1.61

In spite of the favourable standard deviations, it should be noted that, as predicted, large mean errors exist in both tests shown above. In these tests, the mean error also reflects the residual satellite clock error (and, to a lesser extent, satellite position error) in addition to the unmodelled tropospheric delay spoke of above.

4.2 Precise Orbits and Clock

The second set of tests were performed using the precise satellite orbits and clocks. As alluded to above, this should reduce the mean errors in addition to improving overall accuracy. Both trends are visible in Table 4, where both the standard deviations and mean errors are substantially less than those in Table 2. In the case of the standard deviations, an improvement of approximately 25% was observed. Surprisingly, the standard deviations are even better than those from the ground-controlled network in Table 1. This is an auspicious result; however in light of the problems with the data, it is also one that needs further study.

Table 4: Check Point Error Statistics (m): Pseudorange observations, Precise Orbits

	Horizontal	Vertical
Mean	0.37	2.88
Std. dev.	0.16	0.35
RMSE	0.40	2.90
Absolute maximum (mean removed)	0.47	0.97

As shows in Table 5, the use of precise ephemeris also improves results when single-point exposure station position observations are used to control the network. As with the broadcast orbits, however, the position observations approach is not as accurate as when the pseudoranges are directly included in the adjustment.

Table 5: Check Point Error Statistics (m): Single-point positions observations, Precise Orbits

	Horizontal	Vertical
Mean	0.54	2.67
Std. dev.	0.15	0.52
RMSE	0.56	2.72
Absolute maximum (mean removed)	0.61	1.34

5 ADJUSTMENT SOFTWARE

To perform the combined adjustment of photogrammetric and GPS data, the original plan was to use an existing bundle adjustment software package developed as part of a prior research project (see Ellum, 2001, for details). All that was required was the addition of the equations developed in section 2.3. Most of the GPS-specific calculations, such as orbit calculation and application of atmospheric corrections, could have been done in stand-alone programs prior to commencement of the bundle adjustment. However, rather than follow this path, it was decided that the start of a new research project presented a good opportunity to rework the current software, making it easier to maintain and extend.

To satisfy the two goals of maintainability and extensibility, the adjustment program has been divided into individual adjustment modules as shown in Figure 3. In this scheme, the overall adjustment is divided into sub-adjustments that are connected in a hierarchical fashion.

Each sub-adjustment need only make a few generic routines available to the parent adjustment. The parent adjustment then only has to call the routines in the appropriate order. The strategy is further simplified through inheritance and polymorphism – all the adjustments can inherit a generic behaviour from a common base, or, when necessary, implement their own custom behaviour. A program following this design is more maintainable than a single monolithic design because the individual adjustments (and adjustment quantities) can be tested and debugged in isolation. Also, the inheritance and polymorphism results in less code, further improving maintainability.

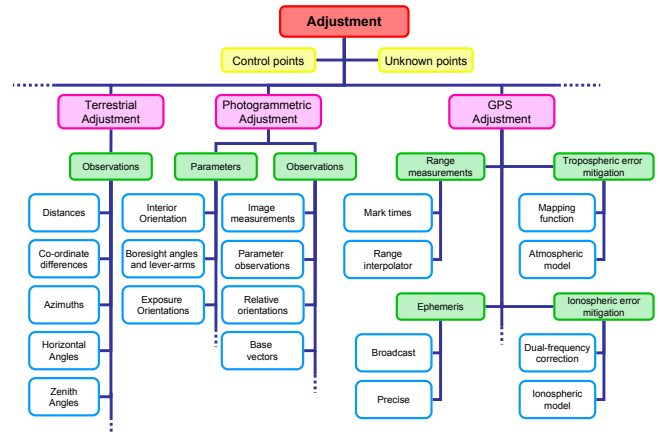


Figure 3: Design of combined adjustment program

A disadvantage of organising the adjustment in the manner described above is that it makes it more difficult to use the reduced normal equations when solving the system of equations. This is because the primary adjustment module, which is responsible for solving the system of equations, has no knowledge of the structure of the normal matrix. For it to maintain such knowledge would be contrary to the goal of genericity. As a consequence, no special techniques – such as the method of reduced normals – are used when solving the system of equations, and the system is solved using Cholesky decomposition only. Naturally, this results in degraded performance. At first, this was a concern; however, at the same time as the adjustment was being re-implemented a move was also being made towards the use of machine-specific tuned BLAS (Basic Linear Algebra Subprograms) and LAPACK (Linear Algebra PACKage) libraries (Anderson et al., 1999). These libraries contain high-performance routines for performing matrix operations and for solving linear systems, and these routines replaced the naïve (but optimised) ‘C’ and ‘C++’ routines that had been previously been responsible for such operations. BLAS and LAPACK libraries are freely available for most computing platforms – in this case, the ATLAS (Automatically Tuned Linear Algebra Software) library was used (ATLAS, 2003). As it turned out, the improvement from using these libraries far outweighed the costs of not using the reduced normal equations. As shown in Table 6, the time required for a moderately sized adjustment using Cholesky decomposition and the tuned libraries was less than one-third the time required when the reduced normal equations were used with the existing naïve routines.

It should be emphasised that naïve does not mean non-optimised, as the existing routines had been compared with a number of other implementations and had been found to be superior.

Table 6: Timings (Pentium IV 1.7GHz, 58 Photos, 1207 Redundancy)

Solution Technique	Solution Implementation	
	Naïve 'C'	BLAS/LAPACK
Cholesky decomposition	92.3s	10.9s
Reduced normals	33.3s	5.2s

6 CONCLUSIONS AND OUTLOOK

Difficulties with the data used during testing mean that it is not possible to state any conclusions definitively. However, the results appear to indicate the following:

- Directly including GPS pseudoranges in a photogrammetric bundle adjustment yields relative mapping accuracies close to those available from a network controlled using ground control points.
- Despite using exactly the same data, directly including GPS pseudoranges in a photogrammetric bundle adjustment provides better results than when single-point exposure station position observations are used.
- Regardless of whether the pseudoranges are directly included in the adjustment or single-point exposure station position observations are used, precise orbits and clocks can substantially improve mapping accuracy.

With regards to the software implementation of the combined GPS/photogrammetric adjustment, an object-oriented program design can improve both maintainability and extensibility of the software. Also, use of tuned linear algebra libraries can dramatically improve performance of numerically intensive adjustment computations. Their use is advised even for moderately sized adjustment problems.

To the best of the author's knowledge, neither the technique of including the GPS pseudoranges in a bundle adjustment, nor a hierarchical combined photogrammetric/GPS/network adjustment have been discussed in the literature before. Both are, however, straightforward extensions of existing practices.

A number of investigations closely related to this paper have been performed but could not be included here because of length restrictions. Foremost among these are tests with fewer than three satellites. Also, the technique of including raw GPS measurements has been extended to include double-difference code ranges. Both of these investigations will be discussed in a forthcoming paper.

Finally, it should be noted that the ultimate goal of this research is the creation of an integrated GPS/photogrammetric processing package. In such an arrangement, the photogrammetric adjustment would feed position updates into

a GPS Kalman filter, aiding the GPS ambiguity resolution. The GPS filter would, in turn, feed highly accurate ambiguity resolved carrier-phase ranges into the photogrammetric adjustment.

ACKNOWLEDGEMENTS

As the author's supervisor, Dr. Naser El-Sheimy is thanked for his ongoing personal and financial support. Joe Hutton and Dr. Mohammed Mostafa at Applanix, Inc. are thanked for providing the data used in this study. Funding for this research was provided by the Killam Trusts and the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

- Ackermann, F., 1992. Kinematic GPS control for photogrammetry. *Photogrammetric Record* 14(80), pp. 261–276.
- Anderson, E., Bai, Z., Bischof, C., Blackford, S., Demmel, J., Dongarra, J., Croz, J. D., Greenbaum, A., Hammarling, S., McKenney, A. and Sorensen, D., 1999. *LAPACK Users' Guide*. 3rd edn, Society for Industrial and Applied Mathematics (SIAM), Philadelphia.
- ATLAS, 2003. Automatically tuned linear algebra software (ATLAS). URL <http://math-atlas.sourceforge.net>, accessed 28 May, 2003.
- Collins, P., Langley, R. and LaMance, J., 1996. Limiting factors in tropospheric propagation delay error modelling for GPS airborne navigation. In: *The Institute of Navigation 52nd Annual Meeting*, The Institute of Navigation (ION), Cambridge, MA, U.S.A., pp. 519–528.
- Ellum, C., 2001. The development of a backpack mobile mapping system. Master's thesis, University of Calgary, Calgary, Canada.
- Greening, W., Chaplin, B., Sutherland, D. and DesRoche, D., 1994. Commercial applications of GPS-assisted photogrammetry. In: *GIS/LIS 1994*, pp. 390–401.
- Hofmann-Wellenhof, B., Lichtenegger, H. and Collins, J., 1994. *GPS: Theory and Practice*. 3rd edn, Springer-Verlag, New York.
- Kruck, E., Wübbena, G. and Bagge, A., 1996. Advanced combined bundle block adjustment with kinematic GPS data. In: *Proceedings of the 18th ISPRS Congress, International Archives of Photogrammetry and Remote Sensing, Volume 31, PartB3, International Society of Photogrammetry and Remote Sensing (ISPRS)*, Vienna, pp. 394–398.
- Mikhail, E. M., Bethel, J. S. and McGlone, J. C., 2001. *Introduction to Modern Photogrammetry*. John Wiley and Sons, Inc., New York.
- Niell, A., 1996. Global mapping functions for the atmosphere delay at radio wavelengths. *Journal of Geophysical Research* 11(B2), pp. 3227–3246.