

Increasing GNSS RTK availability with a new single-epoch batch partial ambiguity resolution algorithm

Alex Parkins

University College London, Gower Street, London, WC1E 6BT, UK
Email: a.parkins@ucl.ac.uk

Abstract GPS Single-epoch Real-Time Kinematic positioning is immune to cycle slips and can be immediately reinitialized after loss-of-lock, providing high availability. This technique requires reliable ambiguity resolution: incorrect ambiguities can cause position errors of several meters, and failed ambiguity resolution reduces availability. However, a bias or inaccuracy in a single phase observation can prevent successful resolution of the whole set of ambiguities. Partial ambiguity resolution allows a subset of ambiguities to be resolved with greater probability of success than the full set. A new algorithm for resolving a subset of ambiguities with validation from previous epochs is described. If normal ambiguity resolution fails, all ambiguity subsets are generated and ordered with the best subsets first. Each subset is then resolved in turn. Fixed subsets are validated against values from previous epochs; this validation procedure greatly reduces the proportion of epochs with incorrect ambiguities. An additional algorithm is described that uses the fixed ambiguities as precise ranges to resolve the remaining unfixed ambiguities. In order to test these new algorithms, GPS data were collected from static and ship-based GPS receivers around Harwich harbor and processed from reference stations at distances up to 111 km. In the static tests the distance over which a 90% ambiguity resolution success rate for dual-frequency data was achieved was increased from 15 km to 76 km. However, in some cases the processing time was too long for this algorithm to be practical without a time-based cutoff. There is also a risk of incorrect ambiguities being propagated, particularly for single-frequency processing. In a ship-based test, the distance over which sufficient availability to support harbor navigation was achieved using single-epoch dual-frequency RTK was increased from 1 km to 66 km.

Keywords

Partial ambiguity resolution, single-epoch, RTK, phase biases

Introduction

The use of Global Navigation Satellite Systems (GNSS) for navigation in safety-critical environments requires robust, reliable positioning with high availability. The most precise positioning requirements necessitate the use of Real Time Kinematic (RTK) positioning, which uses the phase observations as precise ranges to obtain centimetric accuracy. A GNSS receiver measures the difference between the internally-generated and received carrier phases, which leaves the observation biased by an unknown number of integer cycles, called the integer ambiguity. The phase observations alone do not contain enough information to resolve this ambiguity in a single epoch, so it is necessary either to collect phase data over several epochs or to use the code observations for a single epoch. The single-epoch technique is very useful in safety-critical applications because it is immune to cycle slips and can be immediately re-initialized after loss-of-lock, making it more robust than the multi-epoch approach in difficult environments. This results in improved availability of the positioning system.

The ambiguities are resolved in a two-step procedure. A real-valued float estimate of the ambiguities is first obtained by either the single- or multi-epoch technique; this is then fixed to integer values using some ambiguity resolution procedure. If the fixed values pass a validation test, a positioning solution is determined using the ambiguity-fixed phase observations as precise range measurements. It is not necessary to solve for the receiver position until the final step, which leaves two possible parameterizations in the initial step to find the float ambiguities. In the geometry-free technique each receiver-satellite range is solved for separately, whereas in the geometry-based technique the ranges are parameterized in terms of the receiver position. The geometry-based technique solves for fewer parameters and takes advantage of the additional satellite geometry and so gives the greatest probability of successful ambiguity resolution, as well as benefiting from additional satellites. However, a bias or inaccuracy in a single observation can prevent the whole set of ambiguities from being successfully resolved, and incorrect ambiguity resolution affects the whole set of ambiguities,

producing position errors of several meters. The geometry-free technique is more robust to biased observations because a biased observation only affects a single ambiguity, but is a less powerful approach due to the increased number of parameters and loss of geometric information (Eissfeller et al. 2001). The overall geometry-free success rate is reduced by additional satellites, as they are effectively solved independently and the extra geometric information is ignored. However, the geometry-free technique is commonly used for ambiguity resolution research as it is simpler and easier to analyze.

It is well established that the optimal method of ambiguity resolution for the full set of unbiased ambiguities is integer least squares (Teunissen 1999) as automated by the LAMBDA technique (Teunissen 1993). However, it is not always necessary to resolve all ambiguities in order to achieve the desired precision in the final solution, and when using integer bootstrapping (Blewitt 1989) to fix the ambiguities, or when using the geometry-free parameterization, then attempting to fix all ambiguities reduces the overall success rate (Cao et al. 2007). This leads to the problem of partial ambiguity resolution, where a subset of ambiguities may be fixed with an increased probability of success. However, the problem of identifying which subset to fix is still an open one when using integer least squares (Teunissen and Verhagen 2007).

The probability of successful ambiguity resolution is affected by the presence of biases on the phase observations. Unlike biases in the code or ambiguity-fixed phase observations, phase biases cannot be detected and excluded a priori in a single-epoch because they are directly absorbed into the corresponding float ambiguity parameter. This can cause incorrect ambiguity resolution, or prevent the ambiguity validation test from being passed. Biased observations in the geometry-free model only affect the ambiguity from a single satellite by up to half a cycle, except when the observations have been differenced; a bias in the reference satellite affects all the observations. However, in the geometry-based model a biased observation can cause incorrect resolution of the whole set of ambiguities, leading to position errors of several meters. It is therefore desirable to develop a partial ambiguity resolution technique that will enable, at a given epoch, the exclusion of any biased or inaccurate observations that are preventing successful ambiguity resolution.

This paper describes a new single-epoch technique to resolve a subset of the ambiguities based upon batch least-squares ambiguity fixing and the errors in the observations for the given epoch. This algorithm is only applied if the full set of ambiguities is not resolved (i.e. fails a validation test). All subsets

of the full set of ambiguities are generated and ordered according to some criterion, and ambiguity resolution and validation are then attempted on each of these subsets in turn. If the validation test is passed for a subset then the fixed ambiguities are compared to the “weighted mode” of the values for each ambiguity from previous epochs. If all ambiguity values match then the subset is accepted; if they do not it is rejected and another subset is tried. This algorithm results in an increase in the number of epochs with successfully fixed ambiguities when using the single-epoch approach, but maintains the advantage over multi-epoch positioning that it is robust against cycle slips and can re-acquire lost satellites in a single epoch, even with full loss-of-lock.

Partial ambiguity resolution techniques

It is often not necessary to resolve all the ambiguities to achieve the required accuracy for the final position, particularly when there are many satellites and frequencies available. There have been several techniques proposed that will allow a subset of the full set of ambiguities to be fixed with an increased probability of successful ambiguity resolution.

The most widely-discussed existing partial ambiguity resolution technique, introduced in (Teunissen et al. 1999), is based on the geometry-free model and the idea that additional satellites reduce the ambiguity resolution success rate, as discussed in Mowlam and Collier (2004). This idea is illustrated through the equation for the bootstrapped success rate (Teunissen 1998):

$$P_s = \prod_{i=1}^n \left(2\Phi\left(\frac{1}{2\sigma_{\hat{a}_{i|I}}}\right) - 1 \right) \quad (1)$$

where n is the number of ambiguities, $\sigma_{\hat{a}_{i|I}}$ is the i th ambiguity estimate conditioned on the previous $I = \{1, \dots, (i-1)\}$ sequentially rounded ambiguities, and the cumulative normal distribution

$$\Phi(x) = \int_{-\infty}^x \frac{1}{\sqrt{2\pi}} e^{-\frac{v^2}{2}} dv \quad (2)$$

It is clear that, because the overall success rate in (1) is the product of the success rate for each individual ambiguity, fixing more ambiguities reduces the overall success rate. The first step in the partial ambiguity resolution algorithm is to apply a linear transformation to decorrelate the ambiguities: this improves the success rate for bootstrapping and the search speed for LAMBDA (Verhagen 2005). A decorrelating transformation is generated as part of the LAMBDA algorithm. The decorrelated ambi-

guities are ordered by decreasing precision, and a minimum required success rate is chosen. The ambiguities are then fixed in order until fixing another would reduce the total success rate to below the chosen threshold, at which point the algorithm terminates. (1) is a lower bound for the integer least squares success rate; this bound is sharper the more decorrelated the ambiguities are (Verhagen 2003).

The decorrelation step complicates the interpretation of (1). In the geometry-free case the linear combinations are formed between different frequencies to the same satellite, and often correspond to e.g. fixing the widelane ambiguities to all satellites, then fixing the narrowlane ambiguities. In this case, applying the partial ambiguity resolution algorithm might be equivalent to e.g. fixing widelane only, or fixing widelane for all and narrowlane for some satellites. In the geometry-based case the decorrelating combinations are formed across all satellites and frequencies, so any unfixed decorrelated ambiguities will be linear combinations of the original ambiguities across all satellites and frequencies.

(1) may appear to suggest that having more visible satellites reduces the overall success rate. This is indeed the case for geometry-free fixing. However, for the geometry-based approach, additional satellites improve the decorrelation, which results in a higher success rate for each step in (1) and a higher overall success rate for integer bootstrapping; the additional observations also in general increase the success rate for integer least squares. Therefore geometry-based batch-mode fixing will not necessarily benefit from excluding the observations suggested by this technique, although it will be correct for bootstrapping (because the success rate equation is exact) and for geometry-free (because the additional observations from new satellites do not help to fix the existing ones). More frequencies improve the decorrelation for both techniques, and therefore increase the success rate.

The problem with this technique is that it relies on the theoretical precision of the ambiguities to decide which ones to try and fix, and this does not necessarily correspond to which ambiguities can be fixed at the current epoch because it only gives a general precision rather than the specific errors at a given epoch. This technique also only considers the optimum subset to fix based on sequential fixing (bootstrapping) rather than the superior batch-fixing (least squares) approach.

Dai et al. (2007) discuss two partial ambiguity resolution techniques that are applied only if the ambiguity validation test for a given epoch has been failed. The first technique involves trying to

fix all sets of ambiguities from 5 or more satellites; ambiguity resolution is only failed for a given epoch if all these sets have been tried. However, this technique is only briefly mentioned and is not implemented, and would produce a large proportion of epochs with incorrect ambiguities, as shown in the results section of this paper. In the second technique it is attempted to fix the wide- and narrow-lane ambiguities using LAMBDA. If the validation test is failed then the subset of ambiguities that are the same in the most-likely and second-most-likely sets of ambiguities, as obtained from the LAMBDA search technique, are fixed. In Dai et al. (2007), this technique is combined with several other new algorithms, so the results of the partial ambiguity resolution are not clear. However, there are a relatively high proportion of epochs with incorrect ambiguities, so this technique may not be suitable for a safety-critical environment. This may be a problem with the algorithm in that the ambiguities that are the same in the best and second-best sets of ambiguities may not necessarily be the correct ones, since both these sets are based on the same set of float ambiguities, which may contain biases.

Henkel et al. (2009) propose a modification of the LAMBDA technique for multi-epoch satellite-satellite single difference Galileo partial ambiguity resolution in the presence of biased observations. Like the partial ambiguity resolution technique based on the bootstrapping success rate, this technique relies on the theoretical precision of the observations to determine which subset of ambiguities to fix.

New subset ambiguity resolution technique

A new subset ambiguity resolution technique has been developed that uses batch analysis of real observations in a single epoch to resolve a subset of the ambiguities. This method uses the real observations rather than their formal precision, so can take account of the specific biases on each observation at the current epoch rather than relying on the formal precision of the observations. The ambiguities are analyzed in batch mode, which is how they are resolved when using integer least squares, rather than through the sequential bootstrapping success rate. The algorithm is designed to have a low probability of incorrectly-fixed epochs. An overview of the algorithm is shown in Figure 1.

The subset ambiguity resolution algorithm is only used if normal ambiguity resolution, usually LAMBDA followed by the ratio test, fails, i.e. the validation test rejects the best set of ambiguities as not being sufficiently distinguishable from the second-best to be confident that they are correct.

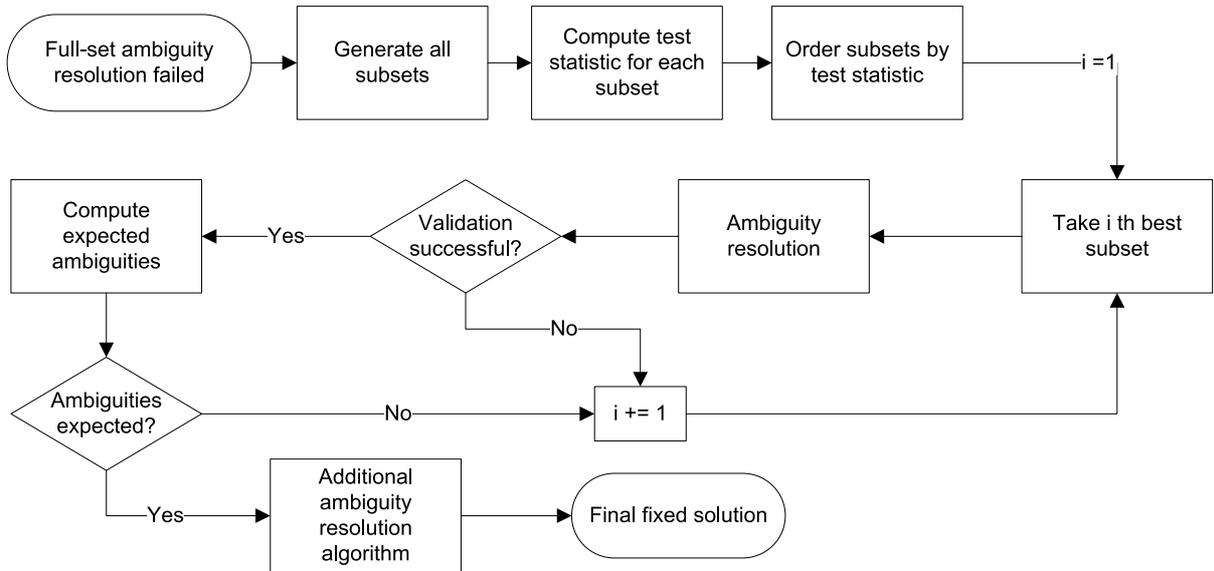


Figure 1: Subset ambiguity resolution algorithm

All subsets of the full set of ambiguities are generated and ordered according to some criterion. Each subset is then analyzed in turn: the corresponding float ambiguity values and covariance matrix are extracted and used to attempt to fix the ambiguities using LAMBDA. The point of the ordering is that the subsets that are most likely to be fixed correctly are tried first, in order to reduce both the probability of an incorrect fix and the processing time. The two ordering methods that have been analyzed are Ambiguity Dilution of Precision (ADOP), a computed measure, and mean Signal to Noise Ratio (SNR), an observed measure.

ADOP (Teunissen et al. 2000) quantifies the a priori precision and geometry of the ambiguities and is defined as:

$$ADOP = \sqrt{\det Q_a}^{\frac{1}{n}} \quad (3)$$

where Q_a is the covariance matrix of the ambiguities and n is the number of ambiguities. Ambiguity sets with low ADOP are more likely to be correctly resolved if all observations are unbiased, so it is attempted to fix these first. Although ADOP does not contain any information about the presence of biases, ambiguity sets with more precise observations and better geometry are more robust, so it is better to try to fix these first in a difficult environment.

SNR shows the strength of the signal, an observed rather than theoretical measure. The idea behind using SNR as an ordering metric is that, as an observed quantity, SNR could contain information about the presence of biases in an observation: signals with a low SNR are more likely to be contaminated by multipath and other errors. Subsets with the lowest mean SNR are tried first. This

favors smaller sets of ambiguities that exclude noisy observations, but in general larger sets are more likely to be fixed correctly using the geometry-based technique, as discussed in the introduction. Therefore the subsets are first ordered by size and then by SNR. A potential problem with ordering by SNR is that it may disproportionately favor high-elevation satellites (with high SNR), to the detriment of the geometry of the system

It is possible to devise other ordering schemes, for example some combination of SNR and ADOP in order to take account of the measured signal strength, theoretical precision and overall system geometry, but it is not expected that this will have a significant effect on the results due to the similarity of the SNR and ADOP ordering schemes as shown in the results section.

The problem with this technique of partial fixing is that any ambiguity validation procedure, such as the ratio test, has a probability of accepting the incorrect set of ambiguities due to the statistical nature of the test; the true ambiguity values are not known. Therefore with each subset attempted there is some probability of accepting an incorrect set of ambiguities, which may result in position errors of several meters. If many subsets are tried then there is a high probability that eventually one will be fixed but, as shown in the results section, not necessarily to the correct subset. This means that the technique described above results in a very high fail rate, making it unsuitable for use; an extra step must be added in order to reduce the probability of incorrect ambiguity resolution.

The additional validation step uses information from previous epochs where the ambiguities were fixed normally in order to validate the fixed ambiguity subset for the current epoch. For a given satellite and frequency, the “weighted mode” of the fixed ambiguity value is taken over the previous epochs, with more weighting given to more recent

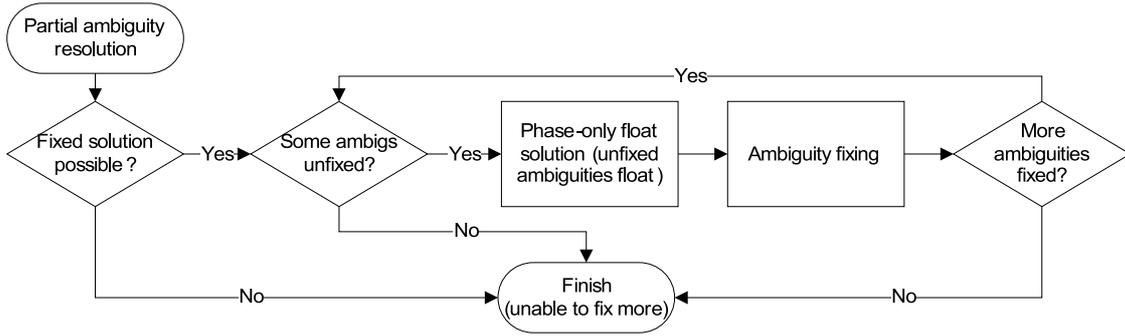


Figure 2: Additional ambiguity resolution algorithm

values, in order to determine the predicted value of each ambiguity. Different weighting schemes affect how many epochs must have their ambiguities resolved to the same values before these values are accepted for use in the algorithm; this is a trade-off between availability and reliability.

The weight applied to the value at each epoch for the processing is $w(\Delta N) = 1/\Delta N$, where ΔN is the time offset in epochs. For a given ambiguity, the sum of the weights for each different value in the previous 20 epochs is computed, and the value with the highest sum is taken as the expected one for that ambiguity for the current epoch. This means that if the three most recently fixed full sets of ambiguities agree on the value of an ambiguity, then the subset ambiguity resolution algorithm will accept this value. If the receiver loses track of a satellite, the previously stored ambiguity values are deleted, as the true ambiguity may be different when the satellite is re-acquired. For the analysis in the results section, the partial ambiguity resolution algorithm is applied if there is at least one stored set of previous values; requiring more previous values would decrease availability but improve reliability.

Only values where the whole set was fixed normally (i.e. not using this algorithm) are used to compute the weighted mode in order to prevent the situation where incorrect values are propagated indefinitely. When a subset is fixed in the current epoch then the values of all ambiguities are compared to the expected values from the previous epochs; if any ambiguities do not match then the fixed subset is discarded.

The effect of this algorithm is as follows: if the true ambiguity values are the same as expected from the previous epochs and the subset is fixed to these values, then it is correctly kept; if it is incorrectly fixed then it is rejected. If the true ambiguity values have changed from the previous epochs and the ambiguities are fixed correctly, then they will incorrectly be rejected; however, this situation will

only last until the whole set has been correctly fixed for a few epochs, and unless all ambiguities have changed (complete loss-of-lock) then it will still be possible to fix subsets that do not contain changed ambiguities. The only dangerous situations are when the ambiguities change but a subset accidentally incorrectly fixes upon the old values, or when the ambiguities have been fixed to the same incorrect value over several previous epochs, so the algorithm treats these incorrect values as true and propagates the error. The latter of these seems the more likely, but could be controlled by a good validation algorithm with appropriate critical value. It may be beneficial to use a stricter validation algorithm when using this technique, as epochs where the validation test fails can still have a subset resolved, but full sets of ambiguities that have been incorrectly resolved can cause problems for the subset ambiguity resolution algorithm in subsequent epochs.

After a subset of ambiguities has been fixed, these ambiguity-fixed observations can act as precise ranges and allow the fixing of further ambiguities, with the aim of improving the precision of the final position. This is done through a least squares adjustment in a similar manner to the single-epoch positioning where the code provides the position, except that ambiguity-fixed observations are used in place of the code observations. The ambiguity-fixed phase observations are much more accurate than the code, and so provide better estimates of the remaining unfixed ambiguities and therefore increase the likelihood of successful ambiguity resolution. If fewer than half of the ambiguities have been fixed then it may be necessary to also include the code observations in this adjustment, although this was not necessary for the data analyzed here. These additional ambiguities can be fixed with the subset ambiguity resolution algorithm, and the whole procedure iterated until it is not possible to resolve any more ambiguities, as shown in Figure 2. This completes the subset ambiguity resolution algorithm for the current epoch.

The subset ambiguity resolution algorithm in this form is a hybrid of the single- and multi-epoch approaches. The ambiguity resolution is based on the observations from the current epoch only, so it is not necessary to worry about cycle slips or re-initialization after loss-of-lock. In particular, cycle slips will only result in the rejection for a few epochs of subsets that include the slipped ambiguities (assuming that normal single-epoch ambiguity resolution does not fix the ambiguities), but subsets that do not include the cycle-slipped ambiguities will still be resolved. In contrast, an un-detected cycle slip in multi-epoch ambiguity resolution will prevent successful resolution. The benefit of the subset ambiguity resolution algorithm over normal single-epoch ambiguity resolution is that there is a possibility to resolve more ambiguities after the normal algorithm has failed, resulting in increased positioning availability and continuity.

Another advantage of this algorithm over some other partial ambiguity resolution techniques is that the subset of ambiguities that are fixed are “real” ambiguities, corresponding to real measurements. In partial ambiguity resolution algorithms where the ambiguities are decorrelated before deciding which to fix, the fixed ambiguities are linear combinations of the true ambiguities. For geometry-free fixing then they are a linear combination of ambiguities from a single satellite (e.g. widelane), and for geometry-based fixing they are a linear combination of ambiguities from all satellites and frequencies. The effect of this is that the partially-fixed observations are less precise: only fixing widelane ambiguities results in an imprecise ambiguity-fixed position. However, the subset ambiguity resolution technique only degrades the position accuracy through fewer observations, which will have less of an effect, particularly when using multiple GNSSs with many satellites visible.

The main disadvantage of this algorithm is the computation time, which is significantly increased due to the necessity of running the LAMBDA algorithm for each subset attempted. This will be analyzed further in the results section.

Performance with real data

A data collection exercise was performed in Harwich harbor, as described in (Parkins 2008, Parkins 2009). A reference station was set up on a roof near the harbor (BASE) and another in a field 15 km away (WIX), and a receiver positioned onboard the ship THV Alert (SHIP) as it performed maneuvers in the harbor. Data from nearby Ordnance Survey (OS) stations were also used, as shown in Figure 3, and the ship was tracked by two total stations to provide a truth model. The data were processed

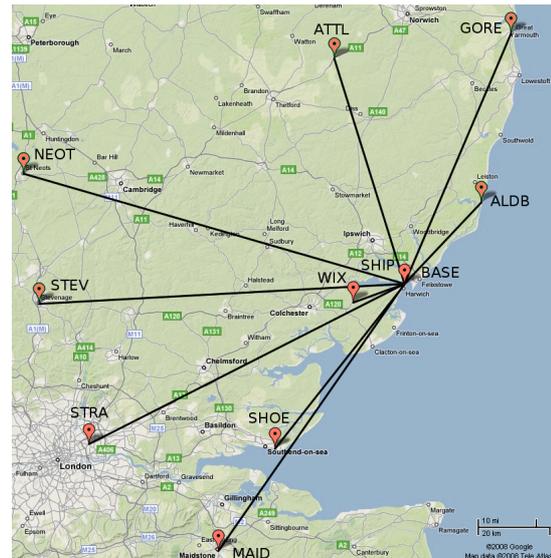


Figure 3: Map of reference stations

using software written in C++ and co-developed by the author at UCL; this software is described in greater detail in Parkins (2009).

90 minutes of 1 Hz dual-frequency data were processed using a single-epoch geometry-based double-difference RTK technique. The Klobuchar model was used to model the ionosphere, and the troposphere was modeled by the ESA model with the Global Mapping Function (GMF). The data were processed over all baseline lengths with a variety of different processing techniques, and the ambiguity resolution success and fail rates determined as the proportion of epochs with correct or incorrect ambiguities respectively. The ratio test with a critical value of 2.5 was used for ambiguity validation: if this failed then the ambiguity was either left unfixed (LAMBDA only) or some variant of the subset ambiguity resolution algorithm applied. The ambiguity-fixed positions are compared to the known positions to determine whether the ambiguities have been fixed correctly, as the true ambiguity values are unknown. Although this technique is not completely reliable it will suffice for comparing the different ambiguity resolution procedures, as the same criterion is applied to both. Future analyses could use a batch-processing technique to more rigorously determine the success rates.

Two different analyses are presented here. The first uses BASE as the rover in order to demonstrate the development of the subset ambiguity resolution algorithm; BASE is static and therefore it is much easier to determine if the ambiguities are correct, and to examine the accuracy achieved. The second analysis uses SHIP as the rover in order to demonstrate the effect of the algorithm in a safety-critical environment; the total station observations are used to validate the results.

Table 1: Static experiment success rate percentage (fail rate percentage in brackets)

Reference station	Baseline length (km)	LAMBDA only	Basic subset fixing	Full algorithm SNR	Full algorithm ADOP
WIX	15	92 (0)	100 (0)	100 (0)	100 (0)
ALDB	31	84 (0)	100 (0)	100 (0)	100 (0)
SHOE	60	56 (0)	95 (0.53)	93 (0)	93 (0)
ATTL	66	39 (0)	89 (5.94)	87 (0)	87 (0)
GORE	76	35 (0)	94 (4.02)	85 (0)	92 (0)
MAID	92	20 (0)	79 (6.50)	65 (0)	65 (0)
STRA	100	2 (0)	50 (26.34)	38 (0)	40 (0)
STEV	102	5 (0.06)	58 (24.85)	51 (2.16)	45 (2.16)
NEOT	111	10 (0)	59 (25.09)	51 (0)	57 (0)

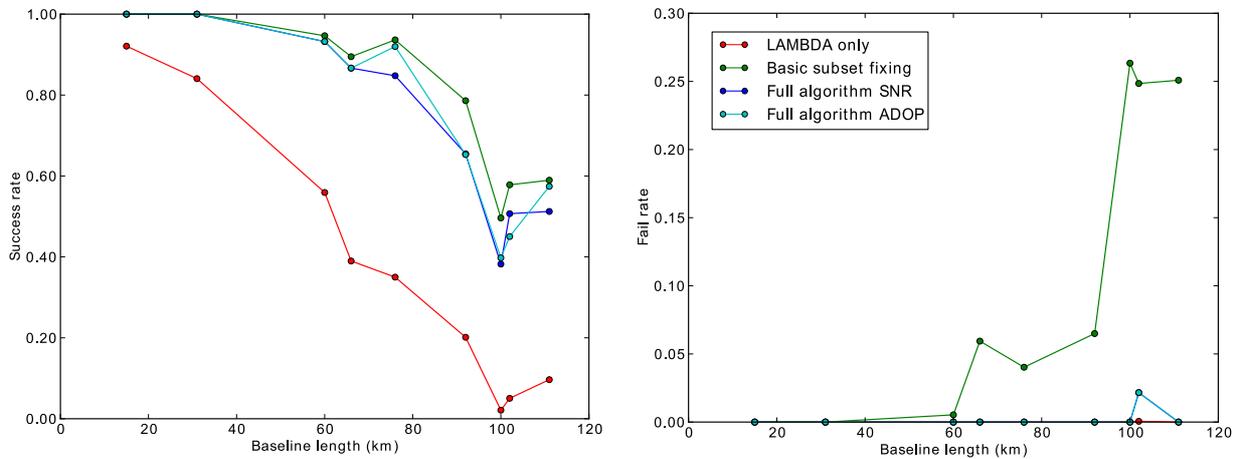


Figure 4: Static experiment results

Static experiment

This experiment uses the temporary reference station set up near Harwich harbor (BASE) as a rover in order to illustrate the development of the partial ambiguity resolution technique. The “true” position of BASE, required to compute the accuracy of the different techniques, was determined by averaging the single-epoch ambiguity-fixed positions over all baselines and epochs. Five different processing techniques were used; the results are shown in Figure 4 and Table 1.

The first processing technique is the conventional procedure, using LAMBDA and the ratio test with a critical value of 2.5. It can be seen that the success rate is significantly lower than for the other techniques, all of which employ some form of additional ambiguity fixing in the case that the ratio test fails using LAMBDA alone. However, the fail rate is also low, with incorrect ambiguities from the 111 km baseline only. The success rate decreases with reference station distance due to the reduced cancellation of spatially-correlated errors.

The second technique used is basic subset fixing, where the subsets are ordered according to ADOP and each is attempted to be fixed in turn, using LAMBDA and the ratio test; no further verification tests are applied. Although this technique results in the highest success rate for many baseline lengths, it also has a high fail rate, which increases with baseline length to over 25%. This is because each time the ratio test is used to validate the fixing for a subset, there is a small probability that an incorrect set of integer values is accepted. With no additional validation and a large number of subsets to try, it is likely that the ambiguities will be fixed incorrectly eventually if they are not fixed correctly first. This means that very few epochs have unfixed ambiguities, because this only occurs when all subsets have been attempted.

The last two techniques employ the full algorithm, with verification of the chosen subset by computing the expected values from previous epochs; the difference between these two is whether the ambiguity subsets are ordered according to ADOP or SNR. These two ordering schemes produce slightly different results, but it is not clear which one is superior. Compared to basic subset fixing, the

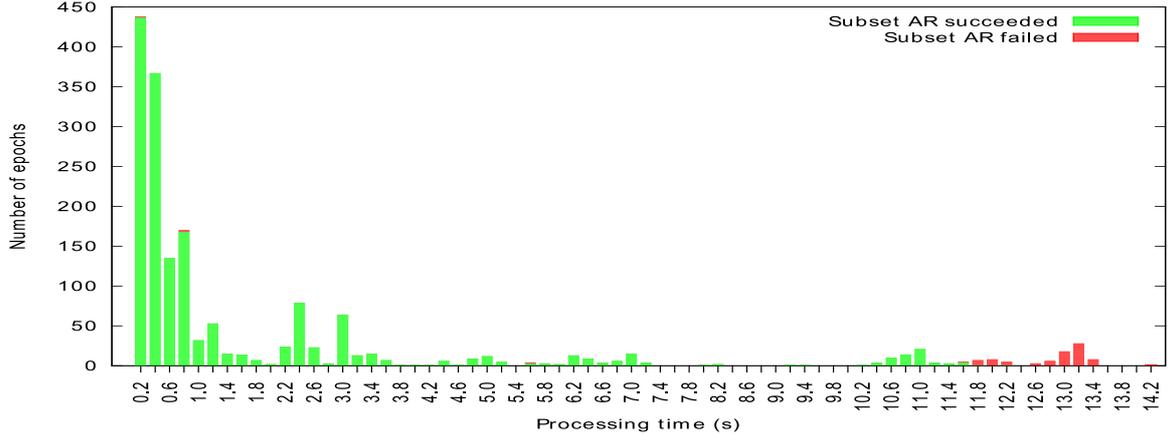


Figure 5: Subset ambiguity resolution algorithm processing time from GORE

success rate is slightly lower, but the fail rate over all baselines except 100 km (STRA) has been reduced to zero. The fail rate for both techniques is still higher than for the full ambiguity resolution to STRA, which is a cause for concern for safety-critical applications. This is due to the propagation of incorrect ambiguities from LAMBDA and ratio test. This could be reduced by either increasing the ratio test critical value over longer baselines to guard against incorrect ambiguity resolution, or by changing the subset acceptance weighting scheme so that more previous epochs are required to be fixed to a value before it is accepted. However, this will result in a longer period of time where this algorithm cannot be used after loss-of-lock.

All techniques show a substantial improvement in ambiguity resolution success rates. However, it is not feasible to attempt to fix subsets without some additional validation, due to the high fail rate. The verification test chosen, comparing the fixed ambiguities for a given subset with those from previous epochs, greatly reduces the fail rate, although over one baseline it is still higher than for the full set due to propagation of incorrect ambiguities.

Table 2 shows the effect the additional ambiguity fixing algorithm on the accuracy of the final position. Out of 9 reference stations, the precision was improved in 6 and degraded in 3, although the accuracy is very poor for one of these latter stations due to incorrect ambiguity resolution. The precision can only be affected in epochs where the subset ambiguity resolution algorithm is used, so the difference is less for the shorter baselines where there is a high success rate with the full ambiguity set. The results show that this technique can improve the precision of the final solution, but this is not guaranteed. This may be because the least pre-

cise observations are excluded by the partial ambiguity resolution algorithm, so subsequently fixing them can introduce poor-quality observations into the solution. However, overall, this data set supports the use of this algorithm to improve the precision of the final ambiguity-fixed solution. The accuracy for both subset ambiguity resolution techniques is worse than for LAMBDA alone, as the additional epochs use fewer ambiguity-fixed observations to obtain the final position. The main benefit of partial ambiguity resolution is in improved success rate and therefore availability, rather than final positioning accuracy.

Figure 5 shows the ambiguity resolution processing time from GORE (66 km baseline) with a 1.8 GHz processor for those epochs where the subset ambiguity resolution algorithm was used. It can be seen that in the majority of epochs the ambiguity subset was resolved within 1 s, and over half of the benefit of the algorithm was obtained within 0.5 s. Many of the epochs that took a long time to process resulted in failed ambiguity resolution, so it may be beneficial to implement a time-based search cutoff; this would be necessary for real-time applications.

Table 2: Static test plan accuracy (2σ , mm)

Reference station	Baseline length (km)	LAMBDA only	Subset ambiguity resolution	Fixing additional ambiguities
WIX	15	9.1	10.2	10.0
ALDB	31	12.4	13.2	12.4
SHOE	60	14.8	22.2	17.3
ATTI	66	21.7	21.4	25.0
GORE	76	19.8	30.1	21.9
MAID	92	14.0	26.9	19.0
STRA	100	18.0	138.1	29.6
STEV	102	31.4	496.7	526.5
NEOT	111	64.8	40.2	43.9

Table 3: Kinematic experiment success rate percentage (fail rate percentage in brackets)

Reference station	Baseline length (km)	Dual-frequency LAMBDA	Dual-frequency subset	Single-frequency LAMBDA	Single-frequency subset
BASE	1	99.6 (0)	100 (0)	40 (0.3)	55 (0.6)
WIX	15	92 (0)	100 (0)	21 (2.1)	46 (3.2)
ALDB	31	93 (0)	99.9 (0)	6 (1.2)	19 (1.7)
SHOE	60	68 (0)	98 (0)	5 (1.1)	9 (1.1)
ATTL	66	41 (0)	95 (0)	10 (1.0)	22 (1.0)
GORE	76	50 (0)	92 (0)	1 (2.5)	4 (3.4)
MAID	92	33 (0)	81 (0)	5 (3.4)	12 (5.5)
STRA	100	5 (0)	33 (0)	1 (3.2)	4 (4.3)
STEV	102	10 (0.4)	41 (2.5)	2 (2.0)	4 (2.4)
NEOT	111	11 (0)	53 (0)	3 (1.3)	6 (0)

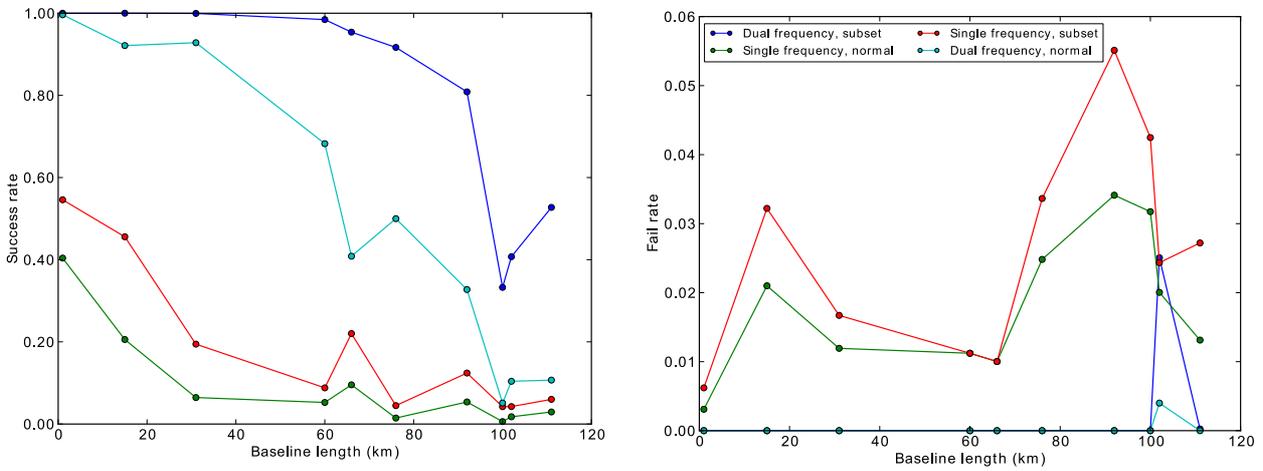


Figure 6: Kinematic experiment results

Kinematic experiment

This experiment uses the receiver onboard the ship (SHIP) as the rover, in order to investigate the performance of the subset ambiguity resolution technique in a real safety-critical environment. The ambiguities are validated by the positions derived from total stations tracking a prism underneath the antenna. Due to the poor accuracy of the total station truth model, it is not possible to assess the accuracy of the GNSS positions in the same way as for BASE. However, the total stations were sufficiently accurate to validate the ambiguities, as described in Parkins (2009). The ship was very close to the reference station at BASE and moved less than 1 km, so the baseline lengths are similar to those for the static experiment.

Figure 6 and Table 3 show the ambiguity success and fail rates over all baselines to SHIP. It is necessary to have a success rate near 100% in order to achieve the positioning requirements for this safety-critical application, as discussed in Parkins (2009). This is only achieved with normal ambigui-

ty resolution over the 1 km baseline using dual-frequency positioning; the ambiguity resolution success rate drops off rapidly with baseline length due to increasingly decorrelated atmospheric and orbit errors at the two receivers. Although this effect is still evident with the subset ambiguity resolution algorithm, the ambiguity resolution success rate is maintained at a sufficiently high level to support safety-critical positioning at baseline lengths up to 66 km. The single-frequency ambiguity resolution success rate is very low for both techniques, but is also improved by the subset ambiguity resolution algorithm, although the fail rate is also increased.

With dual-frequency ambiguity resolution, the ambiguities are only resolved incorrectly over the 102 km baseline (STEV) when using either technique. However, over this baseline, the few epochs of incorrect ambiguity resolution with the full set are propagated with the subset ambiguity resolution algorithm, and the proportion of epochs where the ambiguities are resolved incorrectly is increased from 0.4% to 2.5%. This shows a vulnerability of

the algorithm: if the ambiguities from several previous epochs are fixed to the same incorrect values, then the algorithm will treat these as correct and search for subsets that can be fixed to these incorrect values. This effect could be reduced by making sure that the validation test for the initial ambiguity resolution is relatively conservative; in this scenario it may be that a ratio test critical value of 2.5 is appropriate for the shorter baselines, but a higher value is more appropriate for the longer baselines. The same error propagation effect can be seen with the single-frequency data over most baselines, but there is already a high fail rate with L1 only, which shows that the ratio test critical value is too low. However, the L1 success rate is already very low, and reducing the critical value will result in very few epochs with fixed ambiguities; this shows the advantage of two frequencies for ambiguity resolution.

Performance with simulated data

In order to analyze the effectiveness of the algorithm in a controlled environment, dual-frequency GPS data were simulated for sites BASE and GORE for the whole day of the trial. The data interval was 1 s, which resulted in 86400 epochs of data being processed. The error sources simulated include the ionosphere from IONEX files, solid earth tide and ocean tide loading. The orbit positions were generated from the IGS final SP3 files, which are relatively close to truth compared to the broadcast orbits used in the processing. The phase noise standard deviation was 0.3 m, and the code noise standard deviation was 0.003 m.

The results are shown in Table 4. The ambiguity resolution success rate is increased from 58.4% to 80.5% with the subset ambiguity resolution algorithm, while the fail rate remains the same. This clearly shows the benefits of the subset ambiguity resolution algorithm in improving the availability of single epoch RTK positions. The 2σ accuracy is poor for both techniques due to the effect of the incorrect ambiguities; when these are disregarded

Table 4: Results from simulated data

	LAMBDA only	Subset ambiguity resolution
Ambiguity resolution success rate (%)	58.4	80.5
Ambiguity resolution fail rate (%)	0.2	0.2
95% accuracy all fixed (m)	0.427	0.368
95% accuracy correctly fixed	0.029	0.031

the position accuracies are at the 3 cm level for both techniques.

Conclusion

This paper describes a technique for partial ambiguity resolution in the presence of biased observations, based upon a batch-mode analysis of the observations at the specific epoch rather than their theoretical precisions. All subsets of the full set of ambiguities are generated, ordered according to some criterion (SNR or ADOP), and fixed in turn. If a subset is successfully fixed then it is compared to the “weighted mode” of the values from previous epochs where the full set was fixed: if the values are the same then the subset is kept; if not then it is discarded. This has the effect of validating the subset against previous epochs, and greatly reduces the probability of the algorithm fixing the ambiguities incorrectly, while still achieving the benefits of the single-epoch approach in that it is not necessary to worry about cycle slips and an instantaneous position can be achieved.

Results from real data show that the algorithm improves the ambiguity resolution success rate over baseline lengths from 1 km to 111 km to both a static station and to a rover onboard a ship in a safety-critical environment. In the static tests, an ambiguity resolution success rate greater than 90% was achieved with baselines up to 76 km using subset ambiguity resolution, compared to 15 km using full ambiguity resolution. When the receiver was mounted on a ship, high-availability safety-critical positioning was achieved up to a 66 km baseline using subset ambiguity resolution, compared to 1 km baseline for the normal technique. This shows that the algorithm would be useful for positioning scenarios where high availability must be achieved. Over one long baseline, incorrectly-resolved ambiguities were propagated by the subset ambiguity resolution technique, resulting in a higher fail rate than with the normal processing technique. This is a risk that is present when the normal fail rate is high, and should be guarded against by choosing a sufficiently conservative validation test for the full set of ambiguities. The single-frequency success rate was still low and the fail rate was increased, which demonstrates the necessity of two frequencies for reliable RTK positioning.

The most significant disadvantage to this technique is the increased processing time, as it is necessary to run the LAMBDA algorithm for each subset, until a subset is fixed. The processing was performed with non-optimized code on a 1.8 GHz processor. In the majority of epochs where the subset ambiguity resolution algorithm was applied, ambiguities were resolved within 1 s, and over half

of the benefit of the algorithm was obtained within 0.5 s. However, the lengthy processing time would be a problem for real-time applications where only limited computing resources were available.

A second algorithm was described where the subset of phase observations with fixed ambiguities was used to provide precise ranges in order to fix additional ambiguities. It was shown that this technique can increase the accuracy of the final ambiguity-fixed position, although the improvement is not large and there is the possibility of decreasing the final precision by incorporating less accurate observations.

Future work

Future work should concentrate on the problem of processing time. In particular, it would be useful to introduce a time-based cutoff, where ambiguity resolution for a given epoch is aborted if the processing takes too long; in this case it may be better to ignore the previous epoch and process the next one. In addition, there may be further techniques available to improve the processing efficiency. The decorrelation matrix Z computed as part of the LAMBDA algorithm depends on the geometry of the satellites, which changes slowly. Therefore this matrix could be computed for each subset and stored between epochs for some period of time, being used whenever that particular subset is processed. Efficiency could instead be increased by using bootstrapping rather than LAMBDA to resolve the ambiguities, as this algorithm is computationally cheaper. This would result in a lower success rate for each subset attempted, but it would be possible to attempt more subsets in a given period of time.

With the addition of a time-based cutoff, the ordering of the subsets might become more important, as only the top few would be tried. It would be useful to perform a more comprehensive analysis on the effect of the metric chosen on the ordering of the subsets: as shown in the results section, ordering by ADOP or SNR gives similar results with regards to ambiguity resolution success rate, but this might not be the case if not all subsets are processed. Some combination of SNR and ADOP, or another metric, could be tried. In addition, different decay functions for the weighted mean could be investigated: these affect how many previous epochs need to be fixed to the same value before it is accepted as correct.

Acknowledgements

This work was funded by an Engineering and Physical Sciences Research Council (EPSRC) studentship with the General Lighthouse Authorities of the United Kingdom and Ireland (GLA). The author would like to thank the GLA for their financial support and assistance during the project, and the crew of THV Alert for their patience and cooperation during the data collection exercise. The author would also like to thank Topcon UK for their loan of equipment and assistance with data collection. The author gratefully acknowledges the help and support of Paul Cross during his supervision of this project.

References

- Blewitt, G. (1989). Carrier phase ambiguity resolution for the Global Positioning System applied to geodetic baselines up to 2000 km. *Journal of Geophysical Research*, 94, 10–187.
- Cao, W., O’Keefe, K. & Cannon, M.E. (2007). Partial Ambiguity Fixing within Multiple Frequencies and Systems. In *proc. ION GNSS 2007*.
- Dai, L., Eslinger, D. & Sharpe, T. (2007). Innovative Algorithms to Improve Long Range RTK Reliability and Availability. In *proc. ION GNSS 2007*.
- Eissfeller, B., Tiberius, C., Pany, T., Bibberger, R. & Heinrichs, G. (2001). Real-Time Kinematic in the Light of GPS Modernization and Galileo. In *Proc. of ION GPS 2001*.
- Henkel, P., Gomez, V. & Gunther, C. (2009). Modified LAMBDA for Absolute Carrier Phase Positioning in the presence of biases. In *Proc. ION ITM, 2009*.
- Mowlam, A. & Collier, P. (2004). Fast Ambiguity Resolution Performance using Partially-Fixed Multi-GNSS Phase Observations. In *The 2004 International Symposium on GNSS/GPS, Sydney, Australia*.
- Parkins, A. (2009). Performance of precise marine positioning using future modernised global satellite positioning systems and a novel partial ambiguity resolution technique. Ph.D. thesis, UCL.
- Parkins, A., Grant, A. & Cross, P. (2008). The impact of new signals on precise marine navigation - initial results from an experiment in Harwich harbour. In *proc. NAV08, London*.
- Teunissen, P. (1993). Least-squares estimation of the integer GPS ambiguities. In *Section IV Theory and Methodology, IAG General Meeting, Beijing, China*.

- Teunissen, P. (1998). Success probability of integer GPS ambiguity rounding and bootstrapping. *Journal of Geodesy*, 72, 606–612.
- Teunissen, P. (1999). An optimality property of the integer least-squares estimator. *Journal of Geodesy*, 73, 587–593.
- Teunissen, P. & Verhagen, S. (2007). GNSS Carrier Phase Ambiguity Resolution: Challenges and Open Problems. In *Proceedings of the Scientific meetings of the IAG General Assembly 2007*, Perugia, Italy.
- Teunissen, P., Joosten, P. & Tiberius, C. (1999). Geometry-free Ambiguity Success Rates in Case of Partial Fixing. In *Proceeding of National Technical Meeting & 19th Biennial Guidance Test Symposium*, ION 1999, 201–207
- Teunissen, P., Odijk, D. & Jong, C. (2000). Ambiguity dilution of precision : an additional tool for GPS quality control. In *LGR-Series*, 21, 261–270, Delft Geodetic Computing Centre, Delft.
- Verhagen, S. (2003). On the approximation of the integer least-squares success rate: which lower or upper bound to use? *Journal of Global Positioning Systems*, 2, 117–124.
- Verhagen, S. (2005). The GNSS integer ambiguities: estimation and validation. Ph.D. thesis, Delft University of Technology.

Author Biography

Alex Parkins is a Research Associate in the GNSS, Geodesy and Navigation Research Group at University College London (UCL). He received his PhD from UCL in 2009, for work on GNSS for high-precision, safety-critical applications. His current area of research interest is the use of GNSS for seismometry.

