Enhancements of the Range Consensus Algorithm (RANCO)

Georg Schroth[1], Markus Rippl[1], Alexandru Ene[2], Juan Blanch[2], Boubeker Belabbas[1], Todd Walter[2], Per Enge[2], Michael Meurer[1]

[1]German Aerospace Center
[2]Stanford University

BIOGRAPHY

Georg Schroth is currently working on his Master's Thesis in Electrical Engineering and Information Technology at the Technical University of Munich (TUM) and he is a working student at the Institute of Communications and Navigation (IKN) of the German Aerospace Center (DLR). He also participates in the graduate program on Technology Management of the Center for Digital Technology and Management (a joint venture of both Munich Universities), which is part of the Elite Network Bavaria.

Markus Rippl received his Diploma in Electrical Engineering and Information Technology from TUM in 2007. Since then, he has been a research fellow in the IKN at DLR in Oberpfaffenhofen near Munich. His field of work is the integrity of Global Navigation Satellite System (GNSS) based navigation using receiver-side algorithms.

Alexandru Ene is a Ph.D. candidate in Aeronautics and Astronautics working in the GPS Laboratory at Stanford University. His research focus is on software simulation in the area of combined GPS/Galileo signals, positioning error threat space and navigation integrity.

Juan Blanch received a Ph.D. in Aeronautics and Astronautics in December 2003 from Stanford University, for his work on ionospheric estimation for the Wide Area Augmentation System (WAAS). He continued as a Research Associate in the Stanford GPS laboratory, where he is currently developing algorithms for civil aviation navigation integrity.

Boubeker Belabbas is a research fellow in the IKN at DLR in Oberpfaffenhofen. He is a PhD student at ENPC (Ecole Nationale des Ponts et Chaussées) – Paris and an external PhD student at the TUM. His field of work is that of GNSS integrity, with its applications in augmentation systems and for Safety of Life receivers.

Todd Walter is a Senior Research Engineer and the WAAS Lab Director in the Department of Aeronautics and Astronautics at Stanford University. He received his Ph.D. from Stanford in 1993 and is currently developing WAAS integrity algorithms and analyzing the availability of the WAAS signal. He is a fellow of the ION.

Per Enge is a Professor of Aeronautics and Astronautics in the School of Engineering at Stanford University. He directs the GPS Research Laboratory, which develops satellite navigation systems based on the Global Positioning System (GPS). He has been involved in the development of WAAS and LAAS for the FAA. Per is a Fellow of the ION and the IEEE, and has received the Kepler, Thurlow, and Burka Awards from the Institute of Navigation (ION) for his work.

Michael Meurer received a Ph.D. degree in Electrical Engineering and in 2005 became an Associate Professor (PD) at the University of Kaiserslautern, Germany. Since 2006, Dr. Meurer has been with the IKN at DLR, where he is currently the director of the Department of Navigation.

ABSTRACT

In anticipation of the future GNSS constellations becoming operational, it will no longer be possible to assume that the probability of failure for more than one satellite within a certain timeframe is negligible. Further, it is questionable whether it is always reasonable to compute a position estimate based on all satellites in view, rather than selecting the "best" subset.

The Range Consensus (RANCO) algorithm is not only capable of detecting multiple satellite failures at a time, but it also allows the determination of good estimates of the current ranging biases. RANCO calculates position solutions based on subsets of four satellites and compares

this estimate with the pseudoranges of all the satellites not contributing to this solution. The residuals of this estimate are then used as a measure of statistical consensus.

The scope of this work is the optimization of the performance of RANCO by restricting it to the detection of a certain number of failed satellites at a time and by finding an optimal subset selection process for this constraint. Furthermore, the computation of the subset quality was reconsidered and significantly improved by the use of the Weighted Dilution of Precision (WDOP). In this paper, the physical model for determining the threshold for the separation between correct and faulty satellite signals has been extended. The RANCO algorithm was also verified with respect to its capability of detecting and identifying satellites with a bias higher than a given threshold. Throughout the paper those satellites are defined to fail.

The abilities of RANCO, to exclude multiple simultaneous ranging faults and low biases, paves the way for safety critical applications by combining receiver autonomous algorithms with the integrity channel information from future GNSS systems.

INTRODUCTION

With the rise of enhanced GNSS services over the next decade (i.e. the modernized GPS, Galileo, GLONASS, and Compass constellations) [11], the number of ranging sources (satellites) available for a positioning will significantly increase to more than double the current value. Simulations show that with full Galileo and GPS constellations an average of 18 satellites and a minimum of 13 will be in view for most users. Hence, the applicability of RAIM techniques for the purpose of monitoring position integrity will increase. Additionally, the use of dual frequency receivers will eliminate almost completely the largest magnitude errors for unaided GPS, those caused by the ionospheric delay [7,8].

Unfortunately, one cannot assume that GNSS services, other than GPS, will have the same satellite failure probabilities. A failure probability of 10^{-3} might be proven and realized by the control segment much more easily than the currently accepted probability of 10^{-5}. Altogether, it will no longer be possible to assume that the probability of failure for more than one satellite within a timeframe, significant for safety critical applications, is negligible.

Selecting only a subset of the “best” satellites might be superior to computing a position estimate based on all satellites in view if it can be guaranteed that this subset includes only small biases in the range measurements. In Augmented GPS scenarios like the Local Area Augmentation System (LAAS), it could be necessary to consider and correct only a subset of the current constellation, for reasons related to the available signal bandwidth or due to large propagation errors affecting a number of satellite signals.

The Range Consensus (RANCO) algorithm [1], developed at the Stanford GPS Lab and presented at the European Navigation Conference 2008 in Toulouse, is not only supposed to be capable of detecting multiple satellite failures at a time but also to allow the determination of good estimates of the current ranging biases. Throughout the simulations in this paper, we will consider the particular case where all ranging sources are satellites of the GNSS constellations (i.e. no SBAS, GEOs, or pseudolites).

RANCO calculates position solutions based on subsets of four satellites and compares this estimate with the pseudoranges of all the satellites not contributing to this solution. The residuals of this comparison are then used as a measure of statistical consensus. The satellites that have a higher estimated range error than a certain threshold are identified as outliers, as their range measurements disagree with the expected pseudoranges by a significant amount given the position estimate. All subsets of four satellites that have an acceptable geometric conditioning with respect to orthogonality will be considered.

The chances are very high that a subset of four ranging sources that is consistent with all the other “healthy” range measurements will be found. The subset with the most inliers is consequently utilized for identification of the outliers in the combined constellation.

This approach allows the identification of as many outliers as the number of satellites in view minus five: four for the estimation, and one additional satellite that confirms this estimation. As long as more than five measured pseudoranges are consistent with respect to each other, one can exclude the ones that have a bias higher than the threshold. This approach is similar to the Random Sample Consensus Algorithm (RANSAC), which is applied for computer vision tasks [3], as well as previous Range Comparison RAIM algorithms [5]. As this algorithm is designed for the use with a combined GPS and Galileo constellation, it is performing significantly better in this environment. However, it can also be used with GPS only and still detect failed satellites at low biases.

The first section is a review of the basic ideas of the RANCO algorithm, which form the basis for the enhancements presented in the second section. Here, the procedure for the determination of the fault-free subset is improved and a physical model for the determination of the threshold for the separation between correct and faulty satellite signals is extended.

In the next section, the performance improvement of RANCO by restricting it to the detection of a certain number of failed satellites at a time and using a new optimal subset selection process for this constraint is presented. Further, the computation of the subset quality is reconsidered and significantly enhanced with the aid of the so called Weighted Dilution of Precision.

The fourth section contains a comparison of the results with the initial and enhanced versions of the algorithm, followed by a conclusions section

RANGE CONSENSUS APPROACH

The algorithm developed and investigated in this work is based on the elementary idea of the Random Sample Consensus (RANSAC) algorithm, which is used in the field of graphics and image processing. The algorithm is capable of interpreting data containing a significant percentage of gross errors [3].

Usually, by computing a Least Squares (LS) solution based on multiple measurement samples that correspond to a noise distribution, a single biased sample will influence the result to a considerable extent. Therefore, it is very important to detect and identify outliers and remove them from the final solution. Figure 1 shows a two-dimensional abstraction of this problem. The blue noisy measurement points correspond to the green line that represents the true model behind the samples. One of them has a large bias and causes an inaccurate estimate (the red line) of the true model when computing a LS solution over all measurements.

The RANSAC approach calculates an estimate based on the minimal necessary subset of sample points, in order to minimize the amount of corrupted measurements employed in the estimation.

In the two-dimensional example, an estimate is directly computed based on only two samples. This stands in contrast to the most RAIM algorithms like the Multiple Hypotheses Solution Separation (MHSS) algorithm, where the maximum number of samples is used [9].

As displayed in Figure 2, this may result in many bad estimates (e.g. the blue line), depending on the selected sample pair. In order to find the best pair, the algorithm iterates through all possible combinations of subsets and counts the number of samples that lie within a box surrounding the model (the box is defined by a threshold distance from the model solution space, represented in this two-dimensional case by a line).

If the count of the "inliers" is high, this indicates a high consensus of our current solution with the remaining samples. The measurement points that lie outside of the box are called "outliers". With a threshold that corresponds to the distribution of the noise, it can be assumed, that there is a subset, which corresponds to all other unbiased samples, see Figure 3. Therefore, this approach is applicable to detect multiple biased samples.

Now this approach will be transferred to the satellite navigation case with four-dimensional estimates. Here, the pseudorange measurements are used as sample points and the minimum subset position estimation is based on a combination of four satellites.

These position estimations are compared with the pseudoranges of all satellites. If the residuals of this comparison are higher than the threshold, the corresponding satellites are called outliers. Again, the algorithm iterates through all subsets that are acceptable with respect to the conditioning of their geometry matrix and skips the weak geometries that lead to a higher position Dilution of Precision (DOP) and worse estimates.

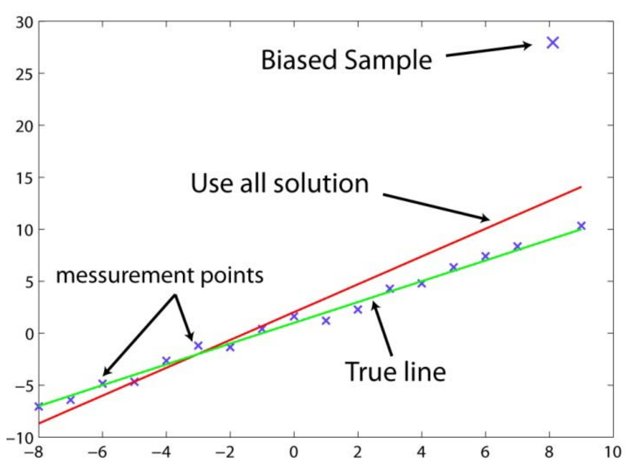

Figure 1: All in view solution

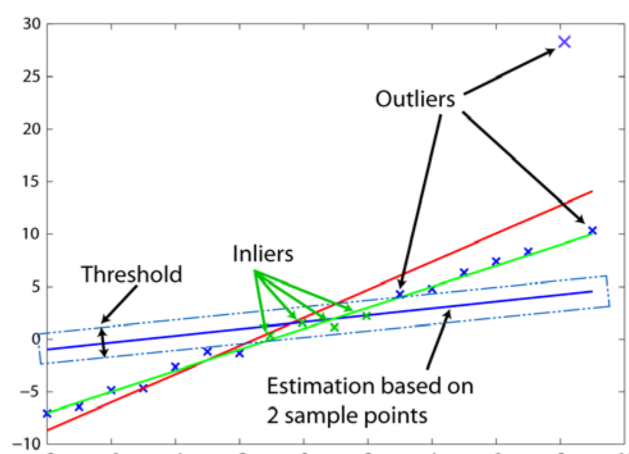

Figure 2: Minimum subset solution

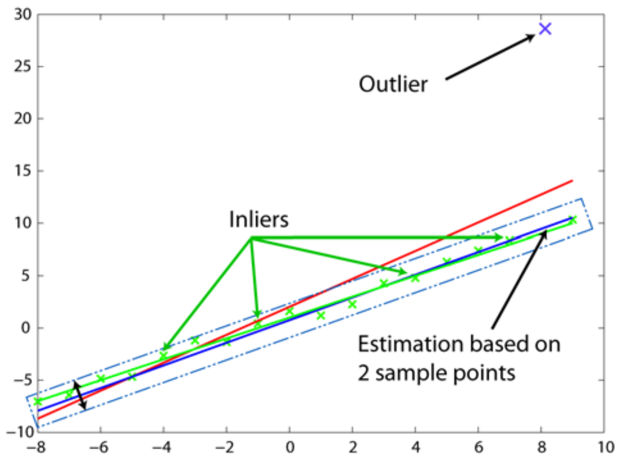

Figure 3: Best subset solution

The best position estimate is based on the subset of four satellites, which leads to the highest consensus with the other pseudoranges and therefore has the highest inlier count. It also defines which satellites are believed to have a bias higher than acceptable. Those biased satellites are referred as outliers relative to this final estimate.

To simulate and evaluate this approach it is not necessary to use the real pseudoranges or to calculate the real position solution. As we are interested in the degree of consensus between the ranges, we rather look at the distributions and errors to avoid many unnecessary computations. The well-known position determination in equation (1) shows the true position vector x, the geometry matrix G, the pseudorange vector y, and the noise vector n:

$$y = Gx + n \ (1)$$

This equation also holds for a single satellite, where $\tilde{y}$ and $\tilde{n}$ are the pseudorange and noise scalars and g^T is the corresponding line in the geometry matrix, where the first three columns are the components of the normal vectors between the true position and the individual satellites:

$$\tilde{y} = g^T x + \tilde{n} \ (2)$$

The LS estimation for the position is obtained by inverting the G matrix. As only subsets of four are considered, the linear system is not over determined and therefore it is not necessary to build the so called Moore-Penrose pseudoinverse:

$$\hat{x} = Hy = HGx + Hn \ \ (3) \qquad H = G^{-1} \ (4)$$

Now, the consensus between the position estimate that was derived by a subset of four satellites and the remaining satellites has to be evaluated. Therefore, equation (2) is remodeled and stated for the noise free case:

$$g^T \hat{x} - \tilde{y} = 0 \ \ (5)$$

This is the main relation, which has to be evaluated for all satellites and with every reasonable subset of four. As already mentioned, it is not necessary to calculate the true position estimates but only to investigate the errors. This fact stems from inserting equations (2) and (3) into equation (5) and reducing $HG = I$:

$$g^T Hn - \tilde{n} = 0 \ (6)$$

The final equation (6) can now be used for the simulations of the RANCO approach.

RANCO AND ITS NOVEL EXTENSIONS

This section gives a detailed illustration of RANCO and its novel extensions. According to equation (6), the normal vectors and consequently the geometry matrix of all satellites in view as well as the error vectors are necessary inputs; see Figure 4. Additionally, the sigma values of the expected error distributions that result by modeling the effects of the troposphere and the ionosphere are required. Those will be used to define appropriate thresholds.

As described above, the algorithm is identifying biased ranging sources by analyzing the agreement of all satellites with all possible subsets of four. As the number of possible subsets is rather high and many of them have a weak geometry, which means that some of the satellites are close to each other in the sky, it is reasonable to consider only the best subsets. The process of the subset selection is described in the following section.

We can assume at this point that there is at least one subset that includes no faulty satellites and has a reasonable geometry.

The position estimations that are based on these subsets are then compared with the pseudoranges of all satellites in view, indicated as the range comparison block in Figure 4. As mentioned, this process is accomplished based on equation (6), in order to reduce the computational complexity.

The expected deviation of a pseudorange from an assumed model is generally related to the individual measurements and the geometry of the subset. Therefore, the error tolerance should be different for each satellite. To allow a separation of errors that lie within the expected error distribution, and biases of unacceptable magnitude, the thresholds are individual and multiples of the expected standard deviation of the residual of the range comparison.

This standard deviation is a function of the expected measurement error standard deviations σ and the standard deviation of the position estimate.

With W, the inverse of the covariance matrix, the term $(G^TWG)^{-1}$ computes a four-dimensional measure of confidence of the current position estimate. This allows us to determine the confidence in the direction of the satellite that has to be investigated by projecting the four-dimensional measure on the corresponding line-of-sight vector g. This results in the expected variance of the position in the direction of the satellites to be compared. Due to the fact that at each range comparison the set of satellites in the subset and those to be compared do not overlap, there is no dependency between them. Thus, the variance of the residual is given by the sum of the variances of the position estimations and the pseudoranges (equation 7).

$$\sigma_{residual} = \sqrt{g^T(G^TWG)^{-1}g + \sigma^2} \quad (7)$$

Hence, the threshold of the range comparison can now be defined to be a multiple of the expected residual standard deviation. The satellites, whose residuals of the comparison are smaller than the corresponding threshold, are defined as inliers of the current subset. Thus, the degree of discrepancy corresponds to the expected noise deviation and the geometry of the subset.

As shown in Figure 4, RANCO iterates through all selected subsets and performs the above described comparison. This allows the determination of the consensus set for each subset and thus the number of inliers. The subset with the highest inlier count can be assumed to have the highest correspondence with all other satellites. The count of inliers, k, has to be large enough to ensure that a correct estimate of the true position was detected. To avoid the possibility that the final consensus is compatible with faulty ranging sources (and assuming that z is the probability that any given measurement is within the error bounds of an incorrect position estimate), z^{k-4} must be very small. While there is no general way of precisely determining z, it is reasonable to assume that it is less than the a priori probability that a given measurement is within the error bounds of the correct model. Assuming $z < 0.5$, a value of k-4 equal to seven will provide a probability of better than 99 percent that compatibility with an incorrect position estimate will not occur. On the other hand, in the absence of faults, the algorithm can be stopped as soon as a subset that defines all satellites as inliers has been found. In this case, RANCO identified no satellites to have a bias higher than the threshold.

As indicated above, a failed satellite within a subset directly impacts the resulting position estimate since four satellites are used to compute four unknowns. Thus, a subset that includes a failed satellite will result in a small

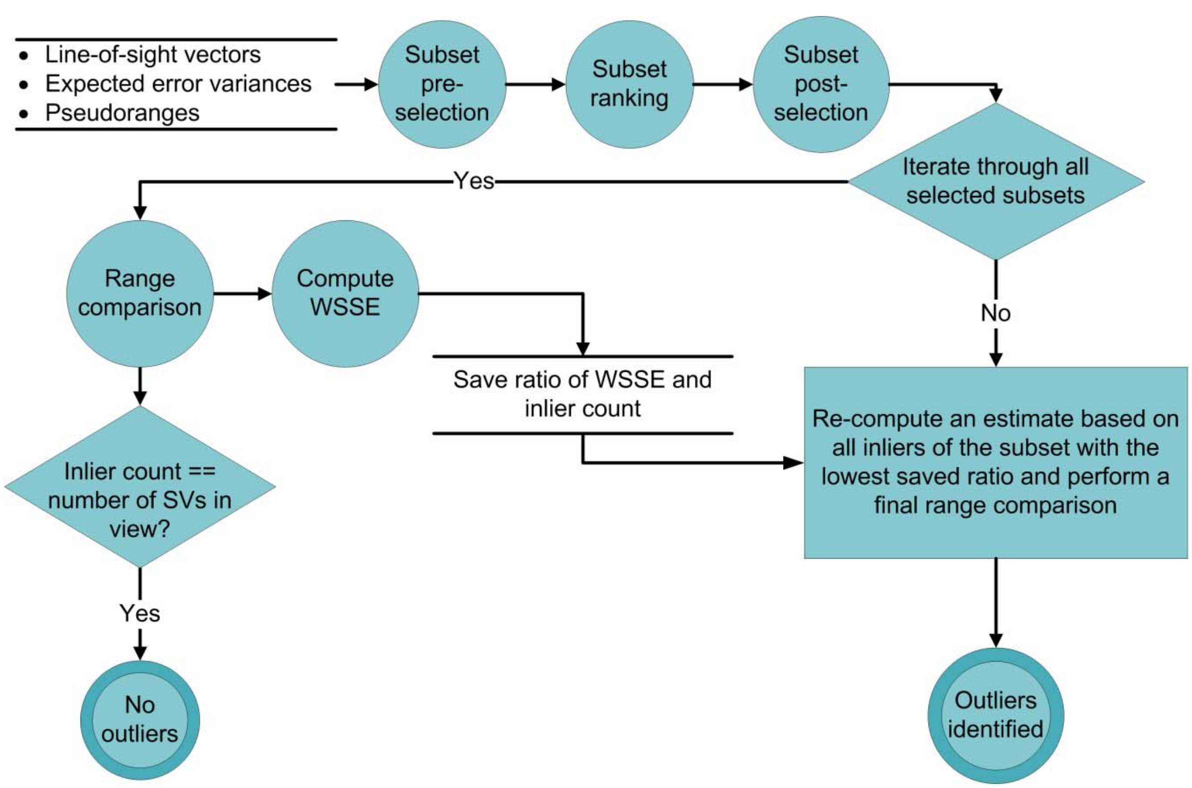

Figure 4: Data flow diagram for the RANCO algorithm

consensus set. This allows the detection of small biases. However, we can even improve this by the additional use of the Weighted RAIM approach [2].

Thus, we compute the Weighted Sum of the Squared Errors (WSSE) of each relevant consensus set, determined within the range comparison. As the consensus set includes more than four satellites, a LS solution can be computed. With the WSSE, the overall consistency of the solution is examined by the error residual of the fit. An estimate of the ranging errors from the LS fit can be obtained from the basic measurement equation

$$\hat{\varepsilon} = y - G\hat{x} = (I - GK)y = (I - P)y, \quad (8)$$

where K is the weighted pseudo-inverse of G

$$K = (G^T W G)^{-1} G^T W \quad (9)$$

and

$$P = GK. \quad (10)$$

Based on these error estimates ($\hat{\varepsilon}$), the scalar WSSE measure can be defined:

$$WSSE = \hat{\varepsilon}^T W \hat{\varepsilon} = [(I - P)y]^T W [(I - P)y] = y^T W (I - P)y. \quad (11)$$

Thus, a low WSSE indicates that the fit was good and the error in the position is most likely small.

We will combine the two approaches described, by building a quotient of the WSSE and the inlier count. If a fault free and a faulty subset result in the same or nearly the same number of inliers, due to the fact that the bias is very small, the WSSE is capable of indicating which of the subsets has the highest probability of being fault-free [2]. Thus, the detection rate can be further improved. If the best subset does not correspond to all the satellites in view, the outliers of this subset are then likely to have a bias higher than the threshold. Then, a final position estimate is computed with a Weighted Least Squares (WLS) solution based on all inliers of the best subset. As this solution is expected to be closer to the true position than the estimate based on four satellites, the residuals of the comparison between this position estimate and the pseudoranges of all satellites in view are determined once more. Thus, a very good guess of the true ranging errors for the satellites is obtained. In turn, this permits the ranking of the satellites with respect to their quality and the exclusion of satellites that have an unacceptable bias. It also allows detecting and removing of a specific bias that is common to multiple satellites, which is useful for reducing multipath effects [10]. This is equivalent to removing the information of one satellite from the final solution; an affordable tradeoff given a high number of satellites in view. As a final remark, with the knowledge about the position of the satellites, it is possible to detect geometric correlations with respect to the ranging errors, which can be used to detect ionospheric fronts [4].

SUBSET SELECTION

The selection of the useful subsets out of $\binom{n}{s}$ possible subsets is of central importance for the performance of the algorithm. Here, n is the number of satellites in view and s the number of satellites in a subset, which is set to 4. Only subsets having excellent satellite geometry shall be considered, as they are less sensitive to errors. Those subsets where line-of-sight vectors are far from orthogonality will generally be skipped.

On the other hand, we have to ensure that for every relevant combination of faulty measurements there is at least one subset that excludes it. This means, that the algorithm has to be aware whether it can guarantee that there is one fault-free subset with a good geometry. Otherwise, the algorithm has to give a warning, if the current constellation does not allow inspecting all possible combinations of failed satellites, called failure modes.

Hence, the subsets with acceptable quality have to be determined first. With a combined GPS and Galileo constellation the number of satellites in view and thus the number of subsets is quite high. Hence, the ranking of the satellite quality has to be very computationally efficient. We propose to use an iterative approach to first eliminate the subsets that do not have a robust geometry and thus are unlikely to result in a high consensus set.

This is done with the subset pre-selection algorithm shown in Figure 5.

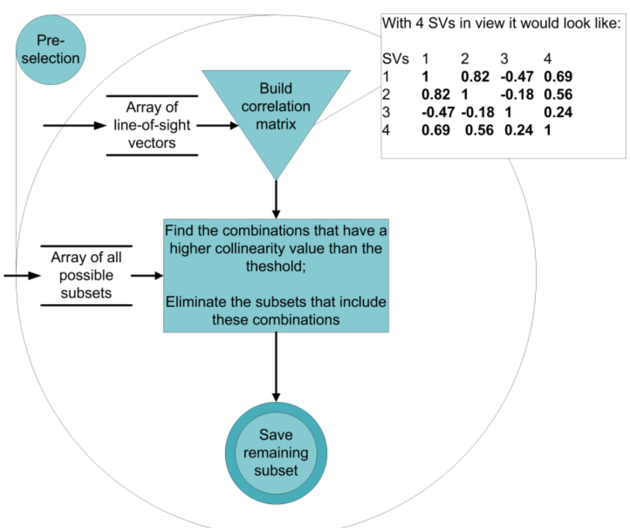

Figure 5: Subset pre-selection algorithm

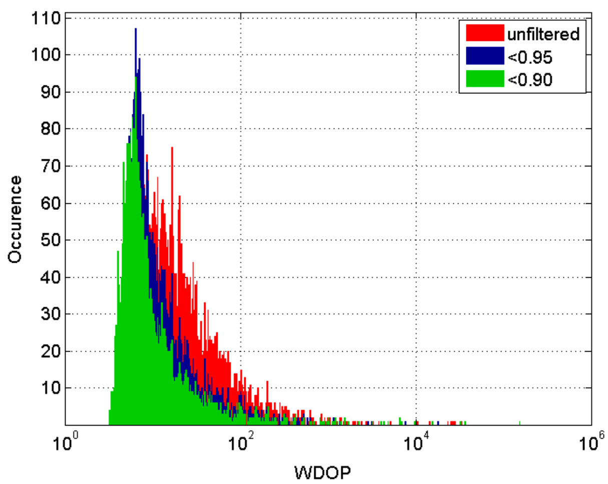

Figure 6: Distribution of the WDOP filtered with the pre-selection algorithm

Here, a two dimensional correlation matrix of all line-of-sight vectors is generated by computing the inner product of all possible pairs of line-of-sight vector combinations.

This allows the identification of the satellites whose lines of sight are nearly aligned. A good measurement geometry entails satellites well-dispersed across the visible sky. Therefore, subsets should not include satellites that are in about the same relative direction. Hence, we exclude the subsets that comprise pairs of satellites that are detected to be more collinear than a certain threshold.

The red histogram in Figure 6 shows a typical distribution of the WDOP of the available subsets for a single measurement. The WDOP is a measure of the expected positioning confidence and of the overall quality of the satellite geometry. It is determined as $WDOP = \sqrt{(G^TWG)^{-1}}$.

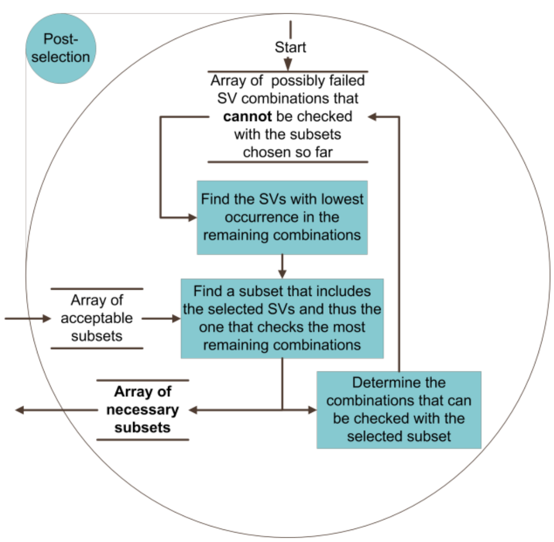

Figure 7: Subset post-selection algorithm

The application of a threshold for the allowed collinearity of a pair of line-of-sight vectors included in a subset results in the filtered blue and green distributions. It was empirically assessed that a considerable part of the unacceptable subsets can be excluded at thresholds of 0.95 and 0.90 respectively. Unfortunately, there are subsets with a very high WDOP and thus a bad geometry that cannot be identified by inspecting the geometry matrix from a two-dimensional perspective only. However, the advantage of applying this algorithm is that it is extremely fast in reducing the number of subsets to be considered in the next selection step.

In the initial version of the RANCO algorithm, the subsets have been evaluated by the computation of the condition number, which proved to be a reliable indicator for the geometry of the subsets. As the results of Singular Value Decomposition (SVD) are used for the computation of the geometry matrix inverse and of the conditioning number, this seemed to be an appropriate approach. However, with the implementation of the new threshold that is based on the standard deviations of the expected signal errors and the geometry, the WDOP turned out to be the optimal measure for the subset quality. Here, both the geometry and the expected signal errors of the satellites in the subset are considered. Hence, the subsets with the lowest WDOP values can be assumed to result in the best position estimates if they are fault-free. Experiments showed, that a WVDOP, that investigates the vertical quality of the position estimate can be computed faster but is an insufficient estimate of the WDOP. The threshold for the subset selection has been set to a WDOP value of 8 m. The change of this parameter results in an alteration of the availability of RANCO. If the threshold is set to a very low WDOP value, the chances are high that there are not sufficient subsets to investigate all possible failure modes. On the other hand, a high threshold increases the Missed Detection and False Alarm Rates.

After having determined which subsets have the highest probability to produce optimal position estimations, it is now important to find the subsets essential in investigating all possible failure modes. Whereas the previous version of RANCO was capable of identifying more than half of the satellites in view to be faulty, this ability is generally not necessary in integrity scenarios since the probability for this event can be neglected. The probability for a satellite fault is assumed to be lower than 10^{-3}, which allows limiting the assumed maximum number of satellites likely to be failed at a time.

This additional information allows RANCO to significantly reduce the computational complexity by reducing the number of subsets that have to be considered. By limiting the maximum number of failed satellites to a certain value, all possible satellite failure combinations can be determined. The post-selection algorithm, illustrated in Figure 7, starts with the identification of the $\binom{n}{f}$ failure modes, where n corresponds to the number of satellites in view and f to the maximum failed satellites. Subsequently, the goal of the algorithm is to find the smallest combination of subsets that can investigate all identified failure modes. Additionally, this combination of subsets should be ideally the one with the lowest WDOP values.

A subset can investigate all ranges not contained in it with respect to their biases, and thus all failure modes that are composed of these not contained measurements. For a maximum number of possibly failed satellites f, each satellite is part of $\binom{n-1}{f-1}$ failure modes. With every investigated failure mode, this count is reduced for the satellites included in these modes. Hence, the algorithm can identify the satellites that are in the least failure modes. The subset, which can investigate the most remaining failure modes, includes the satellites with the lowest count of being part of a failure mode.

The algorithm naturally begins with the subset, which has the lowest WDOP value and determines the failure modes that can be investigated with it. Subsequently, these can be deleted from the list of remaining failure modes. Hence, the count for the satellites of being part of the remaining modes can also be updated. By calculating the sum of these counts for every subset, the remaining subsets can be ranked by their ability to investigate the remaining failure combinations. Thus, the subset with the lowest sum comprises of the satellites which are included in the least failure modes. This subset can investigate the satellites that are not part of the subset and thus the ones that have the highest occurrence in the remaining failure combinations. Additionally, the subset ranking includes also the WDOP value to give preference to the subsets resulting in optimal position estimates.

The search for the next best subset continues until all failure modes can be investigated. For the case in which not all failure modes can be investigated, a warning has to be given. Experiments show that for an average of 19 satellites in view, about 5-6 subsets are required to cover all possible failure modes, assuming a maximum of four failed satellites at a time. Compared to the average of 2000 previously evaluated subsets, this poses an enormous advantage with respect to the computation time.

As the evaluation of the subsets will start with the best, from a WDOP point of view, it is therefore likely that the first comparison already identifies all satellites to be inliers and stops the algorithm in the error-free case.

SIMULATION RESULTS

The improvements and modifications to the original RANCO algorithm have been verified by simulating a worst case scenario and a realistic simulation over 78% of the earth.

I. Worst-case scenario

The worst case scenario is used to directly compare the previous RANCO algorithm with the enhanced version. In the worst case scenario, a geometry with 13 visible satellites providing a maximum GDOP is used, where two satellites failed at a time. As described above, RANCO needs at least one subset of four satellites with a good geometry that does not include any failed satellite. Hence, this very challenging scenario should be ideal to illustrate the strengths and weaknesses of this approach.

First, both versions of the algorithm are evaluated in this scenario with respect to the False Alarm Rate. Here, all 13 satellites were unbiased and a normally distributed noise has been applied to each of them.

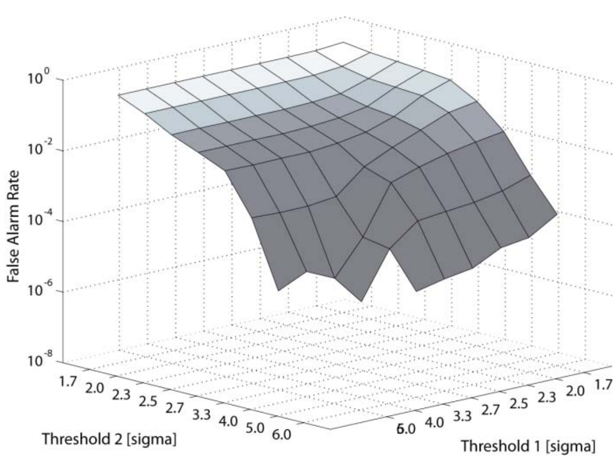

Figure 8: Previous RANCO worst case False Alarm Rate

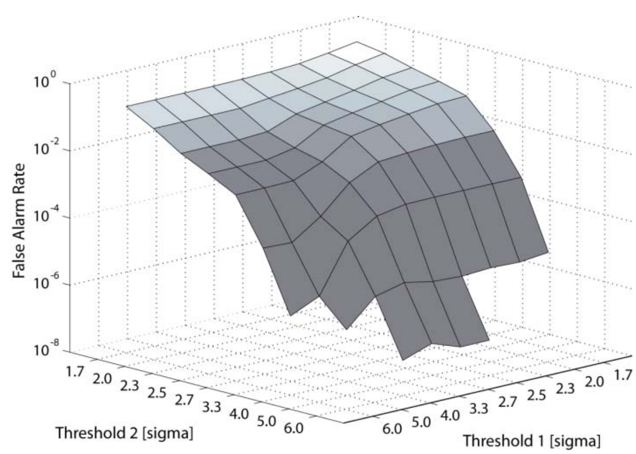

Figure 9: Enhanced RANCO worst case False Alarm Rate

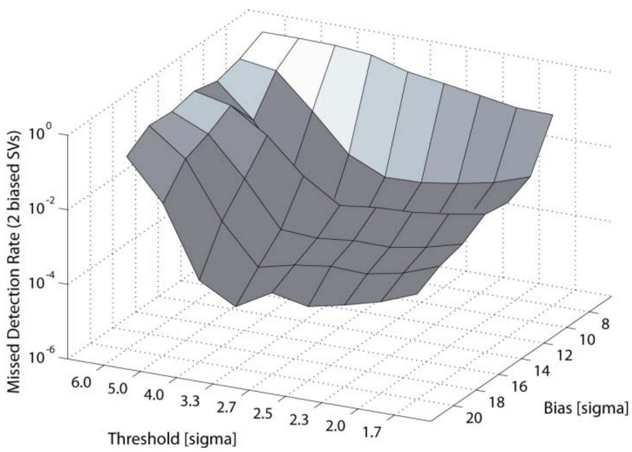

Figure 10: Previous RANCO worst case Missed Detection Rate for 2 biased SVs

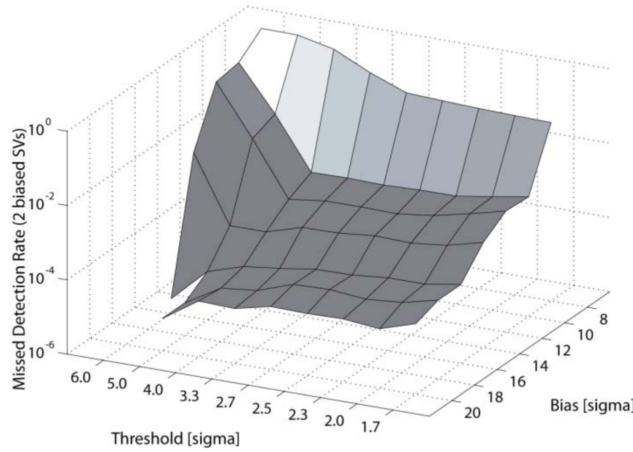

Figure 11: Enhanced RANCO worst case Missed Detection Rate for 2 biased SVs

In Figures 8, and 9 the False Alarm Rate (FAR) is plotted against the two thresholds, where the first is used to identify the subset with the highest probability to be fault free and the second to differentiate between biased and unbiased satellites.

Here, the thresholds range from 1.7 times to 6.0 times the expected noise standard deviation (sigma) for each individual satellite. As assumed, the FAR is mainly depending on the second threshold. However, if the first threshold is set to very low values, the identified fault free subset is small and the differentiation based on this subset is less reliable. In comparison, due to the new definition of the threshold, the enhanced version results in FARs of two orders of magnitude lower than the previous version.

In the second part of the worst-case scenario, the Missed Detection Rate (MDR) for the two failed space vehicles (SVs) is determined. Here, both thresholds are set to the same value to simplify the evaluation. The bias of the two satellites is applied in multiples of the expected noise standard deviation (sigma) and has a random sign. It ranges from 8 times to 20 times the standard deviation.

A Missed Detection is recognized if one of the two biased satellites has not been identified. The improvements with the new version of RANCO achieve a MDR that is one order of magnitude lower than the previous version. The results with the two versions are illustrated in Figures 10 and 11. It is important to note that at very low thresholds a further reduction of the threshold does not necessarily result in a lower MDR.

In these cases, the subset identified to be fault free is small due to the fact that also some correct satellites are excluded. Thus, as described above, the probability that this subset is indeed fault free decreases.

II. Realistic simulations

To provide the reader with more realistic statistics on the performance of the enhanced version of RANCO, this algorithm has been evaluated with MAAST, a simulation tool developed at Stanford [6]. The following statistics are based on constellations observed by 240 users, equally distributed over the earth surface between latitudes 70° S -70° N at 102 time points over 24 hours. A combined GPS and Galileo constellation with 54 satellites altogether has been simulated.

In these simulations, the first threshold has been set to 2.5 times sigma and the second to 3.5 times sigma as those values resulted in a good balance between MDR and FAR. Under these conditions, the FAR has been determined to be lower than 10^{-4}. The subset post-selection algorithm has been adjusted to assume a maximum of four failed satellites at a time.

Figure 12 illustrates the Missed Detection Rate relative to the bias applied to a single satellite. To show the impact of the subset post-selection algorithm, the blue curve represents the performance without the post-selection algorithm and the red line the performance with enabled post-selection.

The comparison of both clearly shows that the selection of only necessary and optimal subsets is superior to the selection of all subsets with good geometry. Hence, this enhancement not only significantly reduces the computational complexity but also improves the performance of RANCO.

In Figure 13, the same scenario with four biased satellites is illustrated.

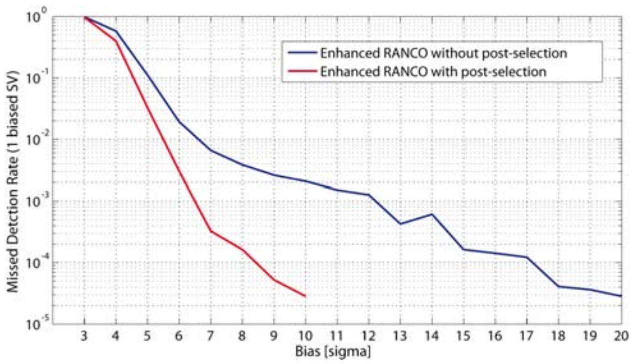

Figure 12: Enhanced RANCO Missed Detection Rate for 1 biased SVs simulated over 78% of the world

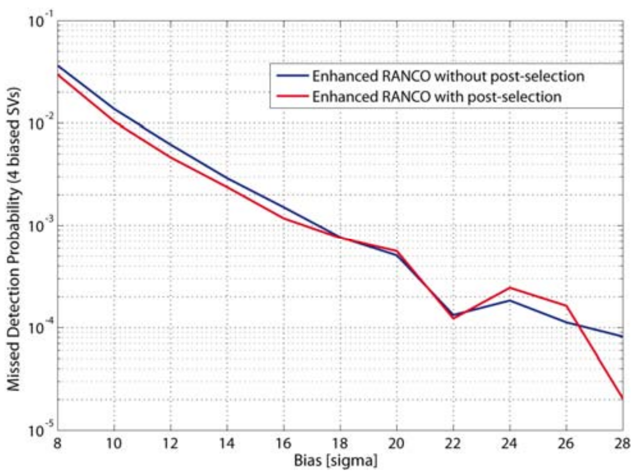

Figure 13: Enhanced RANCO Missed Detection Rate for 4 biased SVs simulated over 78% of the world

As the number of failed satellites now corresponds to the assumed maximum, the post-selection cannot outperform the version without the post-selection with respect to the Missed Detection Rate.

The post-selection algorithm gives warnings in the cases where it was not able to find sufficient subsets. This information is not included in the displayed MDR values.

Altogether, the enhancements to the previous version of RANCO resulted in Missed Detection Rates that are lowered by more than two orders of magnitude.

CONCLUSIONS

The reliable and fast detection of faulty range measurements is a central challenge in satellite navigation, especially with respect to safety of life applications. This fact is becoming more important with the upcoming GNSS and regional satellite navigation systems.

This novel algorithm, called RANCO (Range Consensus Algorithm), developed at the Stanford GPS Lab and the German Aerospace Center, addresses this problem by identifying faulty satellites in the range domain at very low biases.

In general, by knowing the error in the range domain, one can easily calculate the effect of measurement faults in the position domain and decide whether it is reasonable to exclude a satellite or not. RANCO calculates a position solution based on subsets of four satellites and compares this estimate with the pseudoranges of all the satellites not contributing to this solution. The residuals of this comparison are then used as a measure of statistical consensus.

This approach allows identifying as many outliers as the number of satellites in view minus five, four for the estimation, and one additional satellite that confirms this estimation. As long as more than at least five satellites in view are consistent with respect to the pseudoranges, one can reliably exclude the ones that have a bias higher than the threshold. This threshold allows balancing between the Missed Detection Rate and the False Alarm Rate at the same time.

The enhancements presented in this paper, such as the subset evaluation, the subset selection algorithm, and the modified threshold definition, resulted in a significant reduction in both the Missed Detection Rate and the False Alarm Rate. In parallel, the number of subsets that have to be evaluated in this approach could be reduced from approximately 2000 to less than ten in most scenarios.

It is important to note that the computational complexity of this approach decreases with the increasing number of available satellites because less subsets have to be evaluated to cover all possible failure modes. Altogether, this marks an important step in the ongoing development of this algorithm.

For future development it will be of importance to analyze the properties of the algorithm in more detail with respect to the optimal adjustment of the thresholds under more complex scenarios. Replacing the binary thresholds for the identification of the fault free subset by continuous measures, like the sum of the residuals, might reveal an additional performance gain. Additionally, an analytical derivation of the distribution for the missed detection and false detection is necessary.

RANCO's abilities to exclude multiple faulty ranging sources at a time and at low biases paves the way for safety critical and mass market applications by allowing reliable and accurate estimations of position, velocity, and time even during erroneous satellite constellations.

For RANCO, a European patent application has already been filed [12].

ACKNOWLEDGEMENTS

The authors would like to thank Prof. Dr. Christoph Günther, Director of the DLR Institute of Communications and Navigation, for his support in the Range Consensus Algorithms and the provision of the required resources.

REFERENCES

[1] Schroth, G., Ene, A., Blanch, J., Walter, T., Enge, P., "Failure Detection and Exclusion via Range Consensus", *European Navigation Conference*, Toulouse, France, 2008.

[2] Walter, T., and Enge, P., "Weighted RAIM for Precision Approach", *ION Institute of Navigation*, Palm Springs, 1995.

[3] Fischler, M., Bolles, R., "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography", *Communications of the ACM*, June 1981, Vol. 24, pp. 381-395.

[4] Konno, H., "Dual-Frequency Smoothing for CAT III LAAS: Performance Assessment Considering Ionosphere Anomalies", *ION Institute of Navigations Global Navigation Satellite Systems Conference*, Fort Worth, TX, 2007.

[5] Lee, Y., "Analysis of Range and Position Comparison Methods as a Means to Provide GPS Integrity in the User Receiver", *Annual Meeting of the Institute of Navigation*, Seattle, 1986, pp. 1-4.

[6] MAAST, *Matlab Algorithm Availability Simulation Tool*, Stanford, December 7, 2007, http://waas.stanford.edu/~wwu/maast/maast.html.

[7] Misra P., Enge, P., "Global Positioning System: Signals, Measurements, and Performance", Lincoln, Ganga-Jamuna Press, 2005, 2nd Edition.

[8] Parkinson, B., Spilker, J. Jr., "Global Positioning System: Theory and Applications", Washington, American Institute of Aeronautics and Astronautics, Inc., 1996, Vol. II.

[9] Ene, A., "Multiple Hypothesis RAIM with Real-Time FDE and Forecasted Availability for Combined Galileo-GPS Vertical Guidance", in Proceedings of the European Navigation Conference - GNSS/TimeNav, Geneva, Switzerland, 28-31 May 2007.

[10] Phelts, R.E., Enge, P., "The Multipath Invariance Approach for Code Multipath Mitigation", *Institute of Navigations GPS Conference*, Salt Lake City, UT, 2000.

[11] Revnivykh S., Falcone, M., Kogure, S., "GNSS All Over the World", *47th Civil Global Positioning System Service Interface Committee meeting*, Texas, 2007, Vol. 18, pp. 16-18.

[12] Schroth, G., "Failure Detection and Exclusion via Range Consensus", (file number 2008/048), European Patent, 2008.