Failure Detection and Exclusion via Range Consensus

Georg Schroth, *Stanford University and German Aerospace Center (DLR)*
Alexandru Ene, *Stanford University*
Juan Blanch, *Stanford University*
Todd Walter, *Stanford University*
Per Enge, *Stanford University*

BIOGRAPHY

Georg Schroth studies Electrical Engineering and Information Technology at the Technical University of Munich (TUM). Presently, he is working on his Master Thesis. He received the Bachelor of Science Degree in January 2007. 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. The work presented in this paper was performed during his stay as a Visiting Researcher at the GPS Lab of the Department of Aeronautics and Astronautics of the Stanford University.

Alexandru Ene is a Ph.D. candidate in Aeronautics and Astronautics working in the Global Positioning System (GPS) Laboratory at Stanford University. His research focus is software simulation in the area of combined GPS/Galileo signals, positioning error threat space and integrity. He holds a Bachelor in Astronomy and Astrophysics and a citation in German Language from Harvard University, and a Masters degree in Management Science and Engineering from Stanford University.

Juan Blanch graduated from the Ecole Polytechnique in France in 1999, majoring in applied mathematics and physics. In 2000, he obtained a Master of Science in aeronautics and astronautics from Stanford University and joined its GPS laboratory, working on ionospheric estimation for the Wide Area Augmentation System (WAAS). In January 2004, he graduated with a Master's degree in electrical engineering a PhD in aeronautics with the thesis "Using Kriging to Bound Satellite Ranging Errors due to the Ionosphere". He then joined the Stanford GPS laboratory as a Research Associate, where he is currently developing integrity algorithms.

Todd Walter is a Senior Research Engineer in the Department of Aeronautics and Astronautics at Stanford University. Dr. Walter received his PhD. in 1993 from Stanford 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 at Stanford University, where he is the Kleiner-Perkins, Mayfield, Sequoia Capital Professor in the School of Engineering. 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 has received the Kepler, Thurlow, and Burka Awards from the ION for his work. He is also a Fellow of the ION and the Institute of Electrical and Electronics Engineers (IEEE). He received his Ph.D. from the University of Illinois in 1983.

ABSTRACT

With the rise of enhanced GNSS services over the next decade (i.e. the modernized GPS, Galileo, GLONASS, and Compass constellations), the number of ranging sources (satellites) available for a positioning will significantly increase to more than double the current value. One can no longer assume that the probability of failure for more than one satellite within a certain timeframe is negligible. To ensure that satellite failures are detected at the receiver is of high importance for the integrity of the satellite navigation system. With a large number of satellites, it will be possible to reduce multipath effects by excluding satellites with a pseudorange bias above a certain threshold. The scope of this work is the development of an algorithm that is capable of detecting and identifying all such satellites with a bias higher than a given threshold.

The Multiple Hypothesis Solution Separation (MHSS) RAIM Algorithm (Ene, 2007; Pervan, et al., 1998) is one of the existing approaches to identify faulty satellites by calculating the Vertical Protection Level (VPL) for subsets of the constellation that omit one or more satellites. With the aid of the subset showing the best (or minimum) VPL, one can expect to detect satellite faults if both the ranging error and its influence on the position solution are significant enough. At the same time, there are geometries and range error distributions where a different satellite, other than the faulty one, can be excluded to minimize the VPL. Nevertheless, with multiple constellations present, one might want to exclude the failed satellite, even if this does not always result in the minimum VPL value, as long as the protection level stays below the Vertical Alert Limit (VAL).

The Range Consensus (RANCO) algorithm, which is developed in this work, calculates a position solution based on four satellites and compares this estimate with the pseudoranges of all the satellites that did not contribute 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. Hence, the chances are very high that a subset of four satellites that is consistent with all the other "healthy" satellites will be found. The subset with the most inliers is consequently utilized for identification of the outliers in the combined constellation.

This approach allows one to identify as many outliers as the number of satellites in view minus four satellites for the estimation, and minus at least one additional satellite, that confirms this estimation. As long as more than four plus at least one 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 approach is similar to the Random Sample Consensus Algorithm (RANSAC), which is applied for computer vision tasks (Fischler, et al., 1981), as well as previous Range Comparison RAIM algorithms (Lee, 1986).

The minimum necessary bias in the pseudorange that allows RANCO to separate between outliers and inliers is smaller than six times the variance of the expected error. However, it can be made even smaller with a second variant of the algorithm proposed in this work, called Suggestion Range Consensus (S-RANCO). In S-RANCO, the number of times when a satellite is not an inlier of a set of four different satellites is computed. This approach allows the identification of a possibly faulty satellite even when only lower ranging biases are introduced as an effect of the fault.

The batch of satellite subsets to be examined is preselected by a very fast algorithm that considers the alignment of the normal vectors between the receiver and the satellite (first 3 columns of the geometry matrix). Concerning the computational complexity, only 4 by 4 matrices are being inverted as part of both algorithms. With the reliable detection and identification of multiple satellites producing very low ranging biases, the resulting information will also be very useful for existing RAIM Fault Detection and Elimination (FDE) algorithms (Ene, et al., 2007; Walter, et al., 1995).

1. INTRODUCTION

In anticipation of the future GNSS constellations like GPS IIF/III, Galileo, GLONASS, and Compass becoming operational (Revnivykh, et al., 2007), a multitude of questions on the use of these numerous ranging sources will arise. 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, with the given threat models, the applicability of RAIM techniques for the purpose of monitoring position integrity will be increased. Additionally, the use of dual frequency receivers will eliminate almost completely the largest magnitude errors for unaided GPS, those caused by the ionospheric delay (Misra, et al., 2005; Parkinson, et al., 1996). Unfortunately, one cannot assume that GNSS services different from 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 certain timeframe is negligible.

The MHSS algorithm (Ene, 2007; Pervan, et al., 1998) is one of the existing approaches to identify faulty satellites by observing their influences on the VPL. This RAIM algorithm separates the computation of the VPL in multiple hypotheses, which include the cases where single and multiple satellites or even whole constellations have failed. By determining the individual VPL values under each of the hypotheses, weighted by the probability of their occurrence, one can determine the overall VPL. In order to identify faulty satellites, the algorithm builds subsets of the current geometry by excluding one or multiple satellites at a time. An overall VPL is computed for each subset and, as the VPL should increase with a decreasing number of correct satellites, one can expect that the VPL values for the subsets are all higher than for the full geometry. Nevertheless, if a satellite bias influenced the position estimation by a considerable extent, the computed VPL will decrease when excluding this faulty satellite. Therefore, the satellite that was excluded in the corresponding subset, which results in the lowest VPL, is assumed the faulty one.

By minimizing the VPL, satellites with a high ranging bias which does not translate in a large position domain error may not be excluded, as their contribution still reduces the VPL, even though to a small extent. Nevertheless, with multiple constellations present, one might want to exclude the failed satellite, even if this does not always result in the minimum VPL value, as long as the protection level stays below the VAL.

Further, it is questionable if it is always reasonable to compute a position estimate based on all satellites in view rather than selecting only a subset of the "best". 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. Hence, there is a need for a novel algorithm, which is not only capable of detecting multiple satellite failures at a time but also allows determining good estimates of the current ranging biases. This enables a system to deselect the satellites that have a bias higher than a given threshold. With a good estimate of the current ranging bias of each individual satellite, it might be possible to reduce multipath effects by excluding satellites with a pseudorange bias above a certain threshold.

The remainder of this paper is organized as follows: Section 2 discusses the main idea of the RANCO algorithm, which is designed to cope with the challenges and requirements discussed above. Section 2.1 is devoted to the thorough elaboration on its underlying methodology, while section 2.2 comprises a detailed presentation and comparison of the two major subset selection processes, which are part of the RANCO algorithm. The S-RANCO algorithm, a variation of RANCO that allows the suggestion of possibly failed satellites at very low biases, is introduced in section 2.3. Then, section 3 gives an overview on the simulation results of the algorithm and illustrates the differences with respect to the MHSS algorithm. Section 4 concludes the work with a brief summary and an outlook on future work.

2. A NOVEL, RANGE-CONSENSUS-DRIVEN APPROACH

The algorithm developed and investigated in this work is based on the elementary idea of the Random Sample Consensus (RANSAC) algorithm, which is well known in the field of graphics and image processing. The algorithm is capable of interpreting/smoothing data containing a significant percentage of gross errors (Fischler, et al., 1981). 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 at 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 a very bad estimate (the red line) of the true model when computing a LS solution over all measurements.

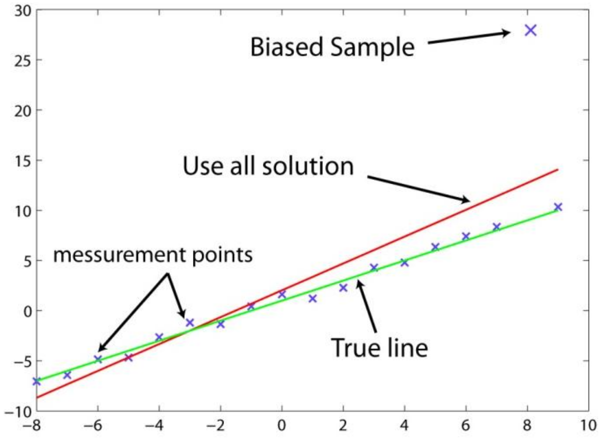

Figure 1: All in view solution

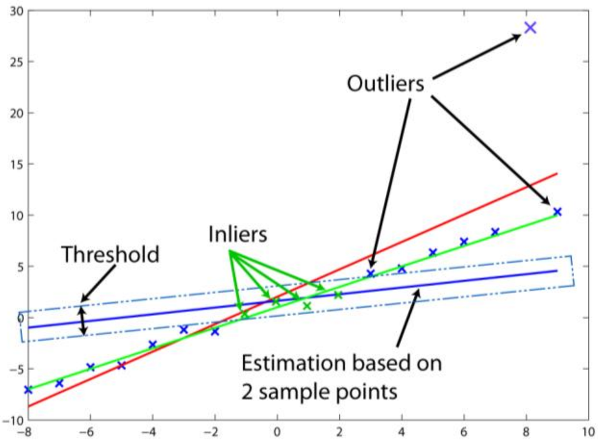

Figure 2: Minimum subset solution

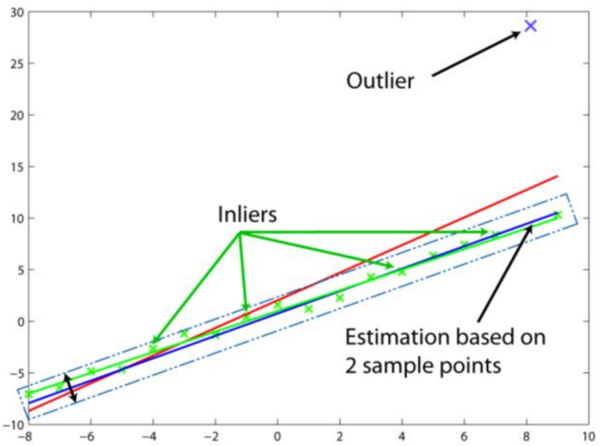

Figure 3: Best subset solution

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.

As displayed in Figure 2, this may result in many bad estimates (e.g. the blue line), depending on the sample pair we select. 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 value). If the count of the "inliers" is high, this indicates a high consensus of our current solution with the remaining samples. The ones 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 we want to transfer this approach to the satellite navigation case where one makes 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 their geometry matrix conditioning and skips the weak geometries as those lead to a higher position Dilution Of Precision (DOP) and worse estimates, which will be discussed in section 2.2. 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:

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

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:

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

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 Moore-Penrose pseudoinverse:

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

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:

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

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. Thus, equations (2) and (3) are inserted into equation (5) and $HG = I$ is eliminated.

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

The final equation (6) can now be used for the simulations of the RANCO approach, which will be explained in the following section.

2.1 A DETAILED ILLUSTRATION OF RANCO

After the discussion of the basic ideas behind the RANCO algorithm, this section will take a more detailed look at it. According to equation (6), the normal vectors and consequently the geometry matrix of all satellites in view, and also the error vectors are necessary inputs. 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 section 2.2. We can assume at this point that the subsets are sorted with respect to the robustness to errors, that every satellite will be included in at least one subset, and that no satellite is within all subsets.

The position estimations that are based on these subsets are then compared with the pseudoranges of all satellites in view. As mentioned, this process is accomplished based on equation (6), in order to reduce the computational complexity. The deviation of the residuals of the comparison is a function of the measurement error variances σ and the geometries of the subsets. The variances of the residuals are given by the sum of the variances of the position estimations and the pseudoranges (equation 7). Here, W is the inverse of the covariance matrix.

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

The expected deviation of a pseudorange from an assumed model is generally related to the individual measurements, and therefore, the error tolerance should be different for each satellite. Hence, the thresholds are individual and are multiples of the expected noise deviation. The satellites, whose residuals of the comparison are smaller than the corresponding threshold, are defined as inliers of the current subset. Here, the degree of discrepancy corresponds to the expected noise deviation.

As shown in Figure 4 the number of inliers is counted for each subset to find the one with the most inliers and thus 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 incorrect 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.

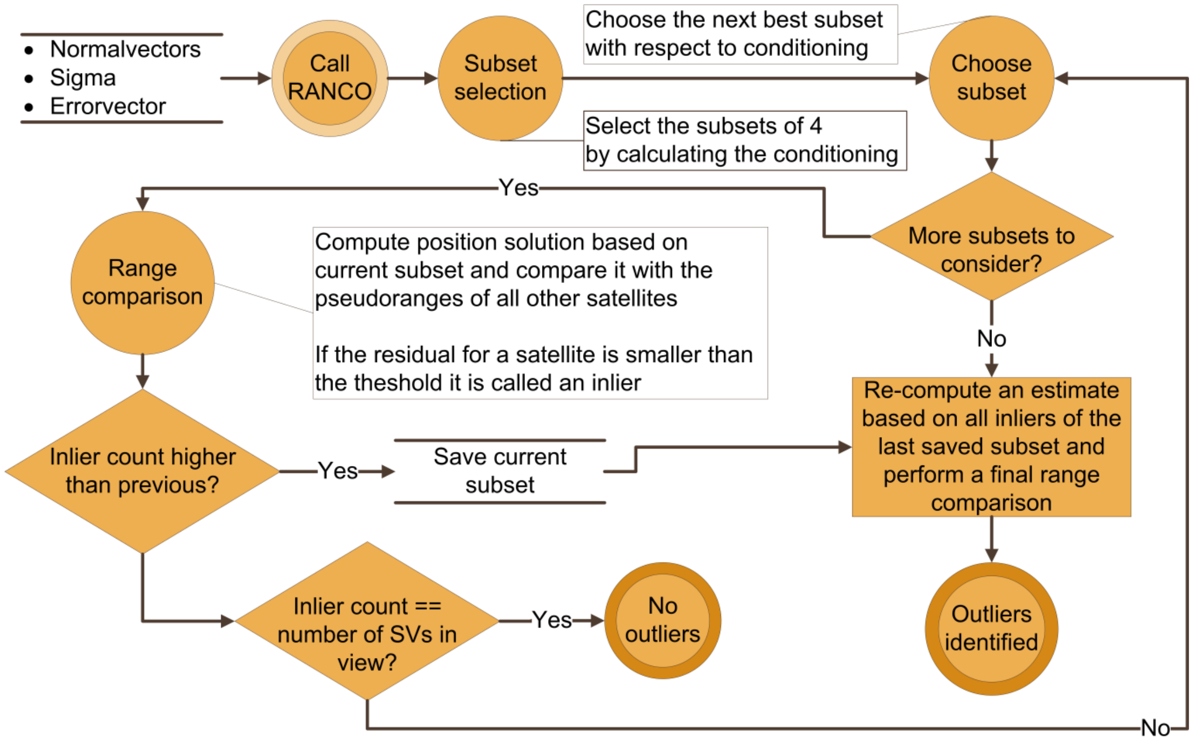

Figure 4: Data flow diagram for the RANCO algorithm

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.

Naturally, 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. However, 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. As this solution is expected to be closer to the true position than the estimate based on four satellites, once more the residuals of the comparison between this position estimate and the pseudoranges of all satellites in view are determined.

Thus, a very good guess of the true ranging errors for the satellites is obtained. This, in turn, allows the ranking of the satellites with respect to their quality and the exclusion of satellites that have an unacceptable bias.

This allows detecting and removing of a specific bias that is common to multiple satellites, which is useful for reducing multipath effects (Phelts, et al., 2000). It is equivalent to removing the information of one satellite from the final solution; nevertheless, this is easily affordable given a high number of satellites in view. Further, 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 (Konno, 2007).

2.2 THE SUBSET SELECTION

The selection of the useful subsets out of $\binom{k}{n}$ possible subsets is of central importance for the performance of the algorithm. Only subsets that have strong satellite geometry, as they are less sensitive to errors, shall be considered and those where satellite lines of sight are far from orthogonal will generally be skipped. A good measurement is the condition number of the geometry matrix.

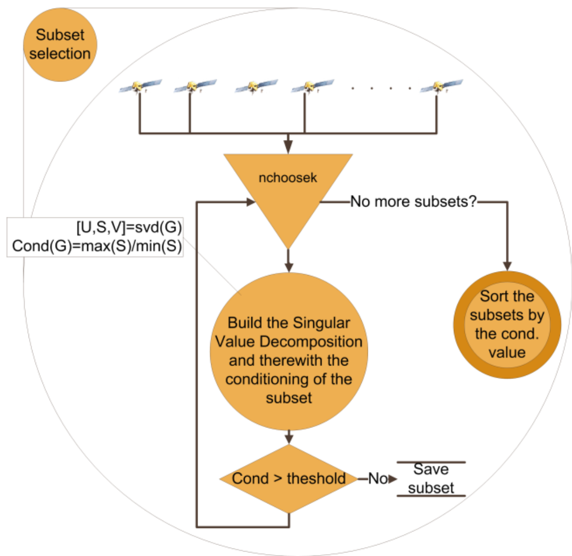

Figure 5: Subset selection algorithm #1

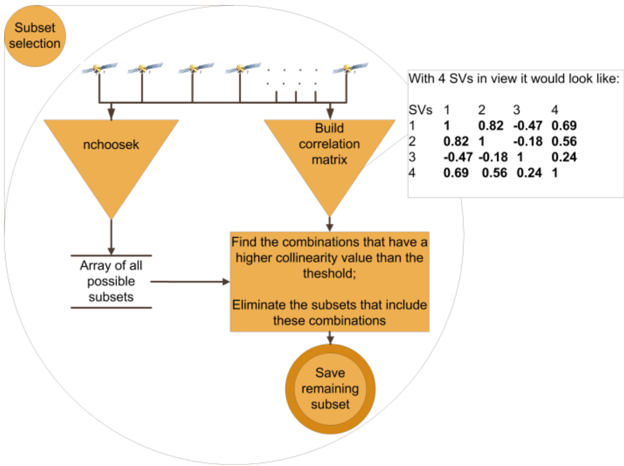

Figure 6: Subset selection algorithm #2

As the results of Singular Value Decomposition (SVD) are used for the computation of the inverse of the geometry matrices and of the conditioning number, this is an appropriate approach.

On the other hand, sufficient subsets are needed to ensure that there is at least one subset excluding any given satellite. If this is not the case, the given satellite cannot be identified to be failed, as it cannot be compared against an independent subset. In the case where this satellite is biased, all subsets containing it, are affected by the bias and consequently erroneous. The probability that at least one of our subsets is an error-free set of four satellites rises with the number of considered subsets. For the case where we have a huge amount of measurements, which is usually the case for RANSAC applications, the relation is given by equation (8).

Here f is the probability that a selected satellite is within the error bounds, u is the probability that a subset does contain no faulty satellite, and p is the probability that we have at least one fault free subset by selecting c independent subsets:

$$(1-u)^c = (1-p)\,; \quad u = f^4 \qquad (8)$$

$$c = [\log(1-p)]/[\log(1-u)] \qquad (9)$$

As in the satellite navigation case a maximum of five independent subsets are available, we cannot apply this relation directly. However, within a combined GPS and Galileo constellation finding sufficient subsets that have a conditioning number below a reasonable threshold is fortunately usually not a problem.

As subsets with a good conditioning are less sensitive to errors, they are sorted to allow the algorithm to start with the best subset as shown in Figure 5. In the error-free case, it is therefore likely that the first comparison already identifies all satellites to be inliers and stops the algorithm. The number of subsets that are finally considered is a tradeoff between computation time and performance of the algorithm. This approach is already fast by building the Singular Value Decomposition (SVD) of four-by-four matrices only and reusing the results in the further computations. It can nevertheless be improved by the selection process in Figure 6: all possible subsets of four are determined and saved to an array.

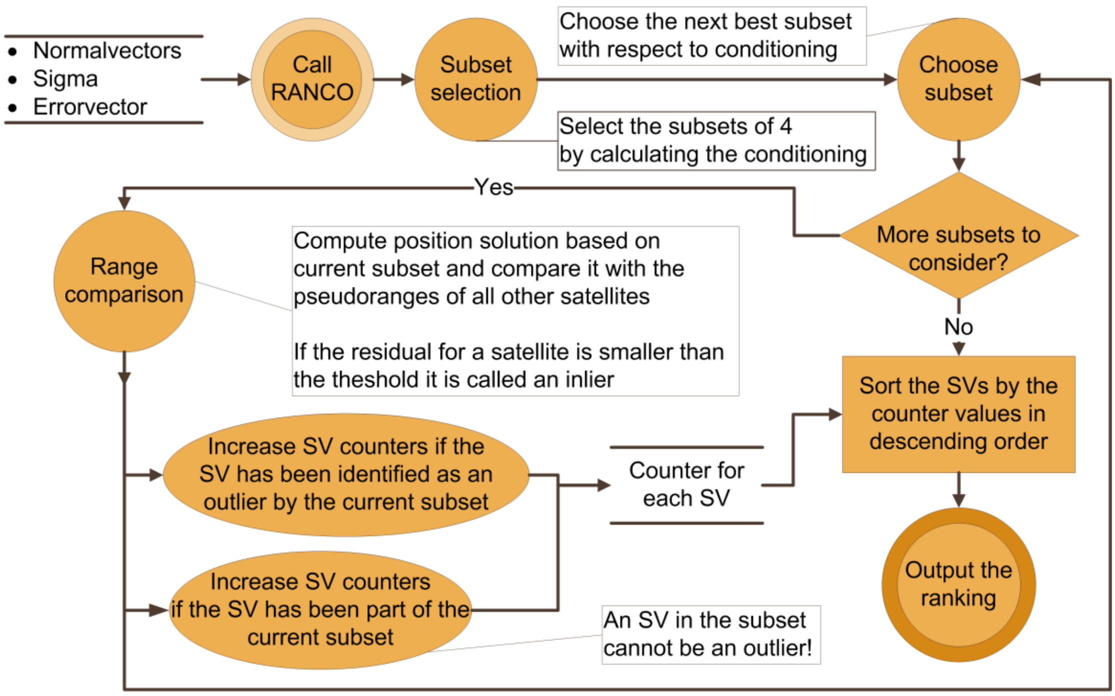

Figure 7: Data flow diagram for the S-RANCO algorithm

In parallel, the two-dimensional correlation matrix of the normal vectors between the satellites and the receiver position is computed. As the algorithm wants to consider subsets with satellites whose line-of-sight vectors are close to orthogonal, the scalar product of all possible combinations of normal vectors is computed. These products indicate the collinearity of the vectors. If an entry in the symmetric correlation matrix is high, the two corresponding satellites are in the same relative direction.

Based on this knowledge, the algorithm can exclude the subsets that comprise satellite combinations that are detected to be more collinear than a certain threshold. The computation of the collinearity matrix as well as the index search is a lot faster than the SVD computation. However, as this approach is restricted to two-dimensional combinations, it cannot evaluate the overall orthogonality.

This means that it excludes subsets that would have been accepted by the first approach. Nonetheless, this effect is relatively small at high thresholds and therefore the tradeoff is acceptable. The two approaches can also be combined in a way that the second one preselects subsets with a very high threshold to filter out certainly not acceptable subsets and forward the remaining ones to the first approach. Then, the original algorithm sorts the subsets again by the conditioning value and excludes the remaining unacceptable subsets. As the second algorithm has a negligible computation time compared to the first one, it immediately allows a reduction in the number of subsets to be inspected by the former.

2.3 S-RANCO, A VARIATION OF RANCO

After a close look at the subset selection procedure, this section will take a look at a second algorithm proposed in this paper, which is very closely related to RANCO.
S-RANCO is also capable of detecting satellite failures but its strength can be found in the suggestion of possibly failed satellites at very low biases. Therefore, the results with S-RANCO can serve as an input for additional algorithms. The major differentiator of this algorithm is that it does not search for the subset with the least outliers but counts the number of times for each satellite being an outlier, as shown in Figure 7. Every time a satellite is determined not to be an outlier, the counter for that satellite is increased.

As it is not guaranteed that every satellite is included in exactly the same amount of subsets, the times the satellite is part of the current subset are also counted. This is necessary, as a satellite that is part of the position solution cannot be an outlier. The addition of the counters normalizes these different initial conditions. The satellite with the highest counter value is most likely to be faulty. It should be investigated by a subsequently executed algorithm. If 1000 subsets are considered for instance and the value for a specific satellite reaches also 1000 or values close to it, is clear that this satellite has been an outlier for all or almost all subsets it was not part of. In this case, the algorithm could also detect a satellite to be failed depending on the threshold.

3. SIMULATION RESULTS

The analytically derived results are now supposed to be verified by simulations with the Matlab Algorithm Availability Simulation Tool (MAAST). This simulator has been developed at the GPS Lab in the Department of Aeronautics and Astronautics at Stanford University. It is a publicly available, customizable MATLAB toolset for simulating confidence estimation algorithms and evaluating their effects on service availability (MAAST, 2007). The RANCO algorithm is implemented within MAAST as shown in equation (6), Figure 4 and Figure 7. The following simulation results are based on a combined GPS and Galileo constellation. Users within 70% of the earth surface (by excluding the earth poles) in the vertices of a longitude and latitude grid with separations of 30 degrees are considered. The duration of the simulations is 48 hours, with measurements every 2.4 hours. In this way, each simulation run results in 1200 samples.

Figure 8 shows the distribution of the number of satellites in view during the simulation. As expected, there is no geometry of less than 13 satellites in a combined constellation. An average of 18 satellites in view can be fairly accepted. The following two graphs show comparisons between the two RANCO algorithms and the MHSS algorithm. In this experiment, a single satellite failure at a time was simulated. The threshold for the RANCO and S-RANCO was set to 2.5 times the sigma of the individual satellites to achieve a low missed detection rate. The applied biases were multiples of the sigma as well and were added to the random noise. Therefore, the failure bias and the random noise can add up constructively or destructively. If the failure bias is equal to the Gaussian noise variance (in the following referred to as sigma) then it is very likely that the overall error is about zero.

Therefore, it is evident that the algorithms can hardly detect any failures below the addition of the threshold and the noise of the satellites. Nevertheless, S-RANCO can still suggest the satellite with the highest posterior probability of being failed. This early knowledge is very useful for further algorithms in a snapshot approach as well as in following analyses. RANCO needs a bias of

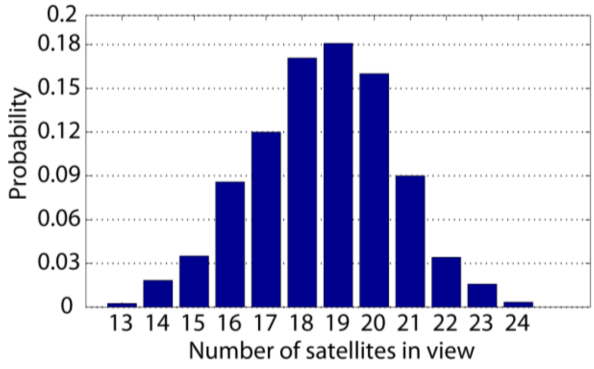

Figure 8: Distribution of the number of satellites in view

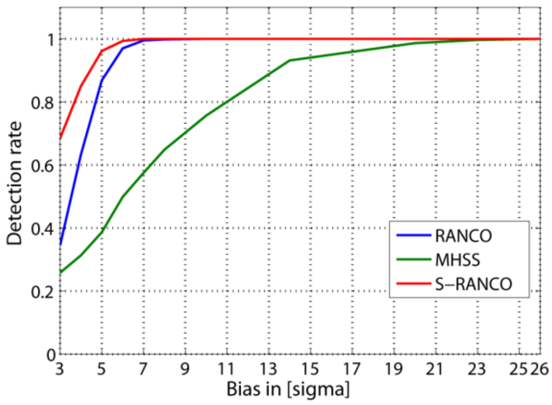

Figure 9: Detection rate by selecting the most critical satellite to be failed

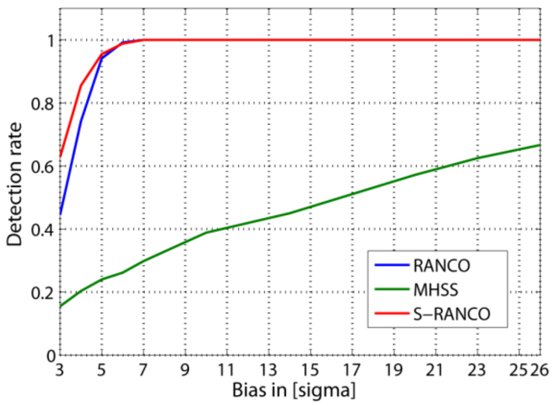

Figure 10: Detection rate by selecting the least critical satellite to be failed

about two times sigma higher than for the suggestion approach, but can identify the faulty satellite. Depending on the threshold and the variance of the Gaussian noise, there are also false detections, which will be discussed later.

The MHSS is identifying faulty satellites by searching for the subset of the current geometry that minimizes the VPL. Not only the ranging error, but also its influence to the position solution and the probability of the hypothesis that a satellite has failed is evaluated here. In Figure 9, the most critical satellite was selected to be the failed, and therefore the one with the highest influence to the position solution, whereas in Figure 10 it was the most unimportant one. RANCO is hardly influenced by this but the MHSS needs much higher biases to identify which satellite has failed. In this comparison, a not optimized version of MHSS was used and one can expect that MHSS can perform significantly better if it is adapted to the FDE application. An implementation of a Weighted RAIM (Walter, et al., 1995) approach would enhance the performance significantly.

Besides the ability to detect biases that are hardly above the noise, the main advantage of RANCO is the detection of multiple biases at a time. Figure 11 shows the results of an experiment with different numbers of satellites failed. There are cases where the algorithm detects only partially the failed satellites.

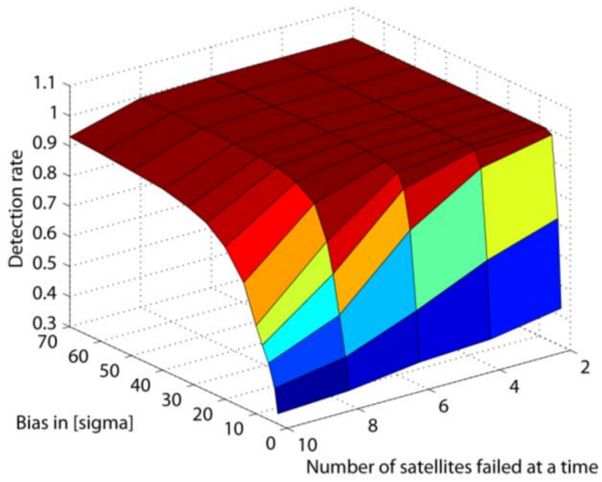

Figure 11: Multiple satellite failure detection

To quantify detection, a correct and complete detection is weighted as 100% and a partial detection of the failed satellites in the corresponding percentage. The average detection score is visualized in this figure. With an increasing number of satellite failures, the necessary bias for a correct detection is increasing by about one sigma for each additional failed satellite. This behavior changes when we encounter more than seven failures.

With the distribution of the number of visible satellites in mind, we see that at least four satellites plus one additional correct satellite are necessary to identify a subset that does not include biased ranging sources and consequently results in an acceptable position estimate. This is necessary in order to be able to correctly detect the remaining satellites as outliers (and thus faulty). Independent of the bias, the constellations where this constraint is not fulfilled cannot be correctly analyzed. For the case where ten satellites have failed, at least 15 satellites in view are necessary to identify all outliers.

The distribution in Figure 8 shows that 5.6% of the geometries considered have 15 satellites or fewer in view. At a bias of 70 times sigma, the detection rate is determined to be 93% with 10 failed satellites, which matches very well the theoretical limit. Further, it is important that the errors are not correlated. If there are more correlated faulty satellites than correct ones, the algorithm will also not be able to detect them. Altogether, the algorithm is able to identify at most the “number of satellites in view – (4+1)” faulty satellites. Besides the detection rate, the false detection probability is of high importance for using RANCO as a RAIM algorithm. To obtain results of statistical significance, the following simulations were based on a single geometry with only 13 satellites in view and one million samples were recorded. As described before, there are two thresholds where the first one is meant to identify the inliers within the run through all subsets and therefore to identify the best subset. The second threshold is applied after calculating a WLS solution based on all previously identified inliers.

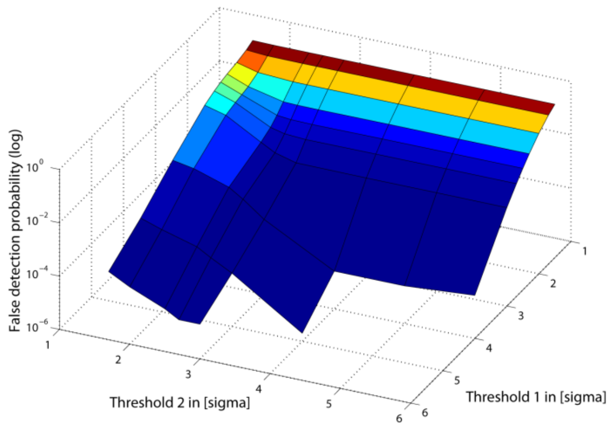

Figure 12: False detection probability as a function of the thresholds (RANCO)

As shown in Figure 12, in the most cases it is convenient to set both thresholds to the same values. However, in geometries with very few satellites it is reasonable to reduce the first threshold in order to increase the dynamic of the ranking of the subsets via the inliers. The final outliers can then be identified based upon the weighted and smoothed solution using all previous inliers and the second threshold. The blank areas in the logarithmic graph show that not a single false detection could be recognized within one million samples.

Thus, it could be shown that a threshold of five times sigma is sufficient to comply with the requirements for the False Alarm Probability. However, the requirements for the false detection rate can be reduced by excluding satellites only if this results in a reduction of the VPL. Further, the Missed Detection Probability for different biases and thresholds has to be analyzed. Figure 13 shows the relationship between the applied threshold and the necessary bias.

For these simulations, the applied thresholds were set solely relative to the expected variances of the pseudoranges for reasons of computational complexity. At high thresholds, the recalculation of the position estimation based on all detected inliers is of high importance. Here, the algorithm is rarely able to find the best subset, as it will identify all correct satellites already in one of the first subsets that are considered. By basing our decisions on all inliers rather than on a subset, we can reduce the necessary bias significantly. The missed detection probability at a threshold of five times sigma and a bias of 12 times sigma is therefore lower than 10^{-4}.

4. CONCLUSIONS

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

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

In general, knowing the pseudorange error in the range domain, one can easily calculate the effect of biases 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.

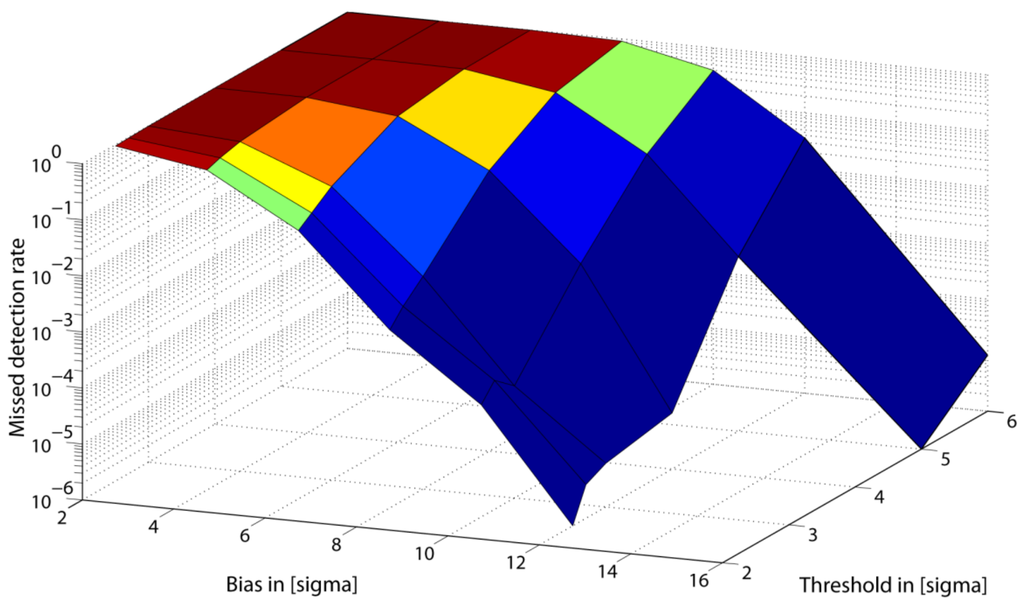

Figure 13: Missed detection probability as a function of the threshold (RANCO)

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.
As this algorithm is designed for the use at a combined GPS and Galileo constellation, it is performing significantly better in this environment. However, it can also be used with GPS only where it still detects failed satellites at lower biases than MHSS.

The minimum bias that allows the effective separation between outliers and inliers can be improved in the S-RANCO algorithm. With multiple constellations present, one might want to exclude a faulted satellite, even if this does not always result in the minimum VPL value, as long as the protection level stays below the VAL.

It will be of importance to analyze the properties of the algorithm in more detail with respect to the optimal adjustment of the thresholds concerning complex scenarios. In addition, the number of necessary subsets that are considered by the algorithm should be evaluated. These values strongly depend on the exact requirements of the purposes of RANCO. Further, an analytical determination of the distribution for the missed detection and false detection needs to be derived.

An additional improvement will be to evaluate a variation of RANCO where position solutions will be computed based on each individual subset, and test each position solution with respect to its consensus to all other position solutions. The subset that results in a position estimation corresponding to the majority of other estimations will be the optimal one. By determining which satellites are in none of the corresponding subsets, one will find the outliers of the current geometry.

RANCO’s abilities to exclude multiple faulty ranging sources at a time and 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 (Schroth, 2008).

REFERENCES

Ene Alexandru Multiple Hypothesis RAIM with Real-Time FDE and Forecast Availability for Combined Galileo-GPS Vertical Guidance, European Navigation Conference GNSS/TimeNav, Geneva, 2007.

Ene Alexandru, Blanch Juan and Powell J. David Fault Detection and Elimination for Galileo-GPS Vertical Guidance, ION Institute of Navigation, National Technical Meeting, San Diego, 2007.

Fischler Martin A. and Bolles Robert C. 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.

Konno Hiroyuki Dual-Frequency Smoothing for CAT III LAAS: Performance Assessment Considering Ionosphere Anomalies, ION Institute of Navigation, Global Navigation Satellite Systems Conference, Fort Worth, TX, 2007.

Lee Young C. Analysis of Range and Position Comparison Methods as a Means to Provide GPS Integrity in the User Reciever, Annual Meeting of the Insitute of Navigation, Seattle, 1986, pp. 1-4.

MAAST Matlab Algorithm Avalability Simulation Tool, Stanford, December 7, 2007, http://waas.stanford.edu/~wwu/maast/maast.html.

Misra Pratap and Enge Per Global Positioning System: Signals, Measurements, and Performance, Lincoln, Ganga-Jamuna Press, 2005, 2nd Edition.

Parkinson Bradford W. and Spilker James J. Jr. Global Positioning System: Theory and Applications, Washington, American Insitute of Aeronautics and Astronautics, Inc., 1996, Vol. II.

Pervan Boris S., Pullen Samuel P. and Christie Jock R. A Multiple Hypothesis Approach to Satellite Navigation Integrity, Journal of The Intitute of Navigation, 1998, Vol. 45, pp. 61-71.

Phelts R.E. and Enge P. The Multipath Invariance Approach for Code Multipath Mitigation, Institute of Navigation, GPS Conference, Salt Lake City, UT, 2000.

Revnivykh Sergey, Falcone Marco, and Kogure Satoshi GNSS All Over the World, 47th Civil Global Positioning System Service Interface Committee meeting, Texas, 2007, Vol. 18, pp. 16-18.

Schroth Georg Failure Detection and Exclusion via Range Consensus (file number 2008/048), European Patent, 2008.

Walter Todd and Enge Per Weighted RAIM for Precision Approach, ION Insitute of Navigation, Palm Springs, 1995.