Sensor Uncertainty Mitigation for Detect and Avoid Systems

E. Theunissen[1]

[1]Netherlands Defence Academy

Abstract

Using RTCA DO-365B test-vectors and Monte-Carlo simulations, the impact of position- and velocity uncertainty is illustrated in the TCPA-DCPA domain for three different sensor-types: ADSB, Mode C and ATAR. Following this, Sensor Uncertainty Management (SUM) strategies are presented and the required decisions concerning the quantification of the relevant design parameters are discussed. To illustrate the performance of a prototype SUM function, baseline alerting performance statistics are provided for two alerting algorithm configurations without SUM that use the bounds of the available trade-space of a DO-365B Detect and Avoid alerting function. This data provides insight into the possible trade-off between missed alerts and incorrect alerts without SUM. The performance statistics of the prototype SUM function demonstrate a significant reduction in missed alerts both for the use of ADSB and Mode C tracker data with no increase in incorrect alerts for ADSB and less than a 1% increase for Mode C.

Keywords: Sensor Uncertainty Management, Detect and Avoid

1. Introduction

To safely integrate remotely piloted aircraft systems (RPAS) into the U.S. National Airspace (NAS), Detect and Avoid (DAA) systems are required [1][1]. A DAA system is an electronic means to support the pilot with avoiding conflicting traffic before the situation has turned into a collision hazard. To achieve this, a DAA system provides the pilot with alerting and guidance to remain DAA Well Clear (DWC)[2]. To provide timely alerting and guidance, it needs to be assessed whether the future separation with other traffic is predicted to decrease below the DWC boundary within a specified alert time[3]. Such an assessment is typically performed using an extrapolation of the measured ownship- and traffic state.

1.1 The need for Sensor Uncertainty Management (SUM)

The measurements of ownship and traffic position and velocity will contain errors, the magnitude of which will be a function of sensor type but can furthermore depend on conflict geometry. In [3], the impact of position- and velocity uncertainty on the computation on the boundaries of conflict space defined by a spatial threshold is discussed. It is illustrated how the inaccuracy in velocity measurements constrains the effective look-ahead time that can be used by a conflict prediction function. In Figure 1 (from [3]), the solid red contour represents the space where horizontal separation is predicted to decrease below 2 NM based on the reference (true) state of ownship and intruder. The dashed, dash-dotted and dotted red contours represent the same threshold, but now for three different combinations of velocity and track uncertainty. The solid red line arc at the top of the figure represents the headings which must be avoided to prevent a decrease in separation below 2 NM. The dotted arc indicates these headings for the highest velocity and track uncertainty used in the example.

[1] Comparable requirements are under development for access to the European Airspace.
[2] For UAS access to the NAS, the DWC boundary is mathematically defined and quantified in [2].
[3] DO-365B specifies the boundaries of time (i.e. ‘no earlier than’ and ‘no later than’) within which an alert is to be provided and a minimum average alert time.

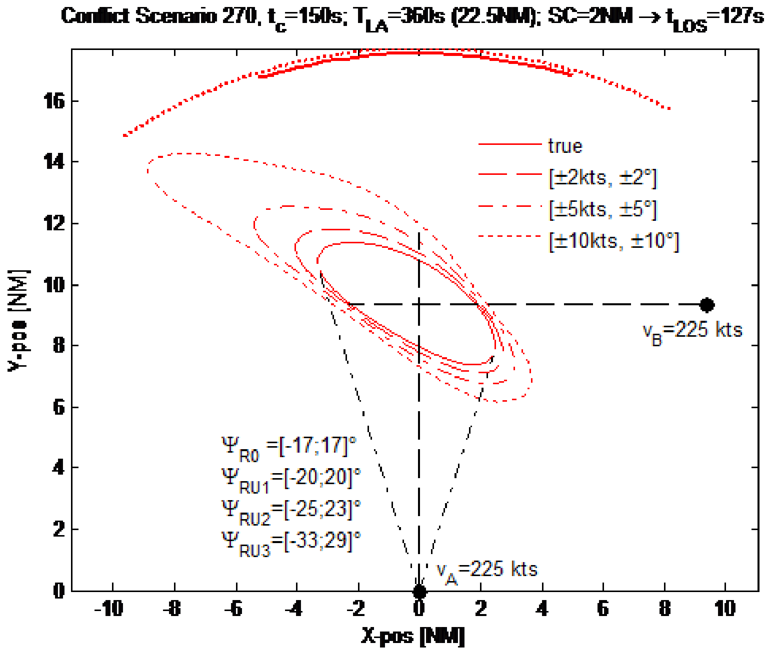

Figure 1 – Impact of velocity uncertainty on conflict space [3]

As can be seen by comparing the solid heading arc with the dotted arc, the increase in velocity uncertainty requires a significantly larger change in heading to avoid the predicted separation to decrease below 2 NM.

In [2] minimum sensor performance requirements are provided for the DAA alerting and guidance functions, and it is indicated that "*Implementing an alerting and guidance scheme using simple predictions and the HAZ values directly from the requirements will result in unacceptable performance*". This "*unacceptable performance*" is characterized by a too high percentage of late-and/or missed alerts, a too high percentage of nuisance alerts and too many intermittent alerts. Sensor Uncertainty Management (SUM) is intended to achieve the required performance by addressing all three issues. To test the performance of an alerting function and generate alerting performance statistics for a particular implementation, so-called test-vectors are provided in [2]. These test-vectors comprise ownship and intruder trajectories for hundredths of geometries and are accompanied by so-called Encounter Characterization Files (ECF) that provide a reference for the required alert. The availability of such a standard test-vector set allows comparisons between different implementations and/or configurations of DAA alerting functions while ensuring that performance differences are not caused by the input data. All data used in this paper to analyze encounters and generate alerting performance statistics is from the test-vector archive provided with [2].

1.1.1 TCPA, DCPA and alerting thresholds

Two key parameters in a DAA alerting algorithm are the predicted Distance at the Closest Point of Approach (DCPA) and the Time to Closest Point of Approach (TCPA). Due to accuracy limitations regarding the measured position and velocity of ownship and traffic, errors in the predicted DCPA and TCPA will result. For a given spatial alerting threshold, this will cause a false alert if the DCPA based on 'perfect data' is above the alerting threshold while the DCPA based on measured data is below it. The opposite situation will cause a missed alert.

In [4] it is described how the use of time to co-altitude, data filtering, alert filtering, alert hysteresis and sensor specific adaptive thresholds have been applied to increase the robustness of the alerting function. The TCPA-DCPA depiction (Figure 2) was used to illustrate the trade-space that exist in terms of spatial- and temporal alerting thresholds. In this figure, the origin (TCPA=0) is on the right side and TCPA increases from right to left. A converging encounter would yield a TCPA that moves from left to right.

A designer can position the alerting threshold in the may-Alert zone, which is the space between the Non-Hazard Zone (NHZ) and the Hazard Zone (HAZ). The NHZ is characterized by a DMOD of 1.5 [NM] for the Preventive and Corrective Alerts and 1.2 [NM] for the Warning Alert. If (A) represents the true state (lying in the NHZ) and (B) represents the computed state (lying in the Corrective may-Alert zone), the fact that (B) lies above the alert threshold (blue line) prevents an alert from being generated.

Only if the difference between true and computed state would be so large that the computed state lies below the blue line, a False Alert would be declared.
On the other hand, if the true state is represented by (C) and the measured state by either (B) or (A), no alert will be generated because both measured states lie above the blue line. This is still acceptable since the true state does not lie in the Must Alert zone represented by the white box. On the other hand, if the observed state in this case is represented by (D), an alert is generated. This is acceptable since the true state lies in the may-Alert zone. Finally, if the true state is represented by (D) and the computed state by (C), an alert will be generated because (C) lies below the blue line. However, if the computed state is represented by (B), a Missed Alert condition occurs.

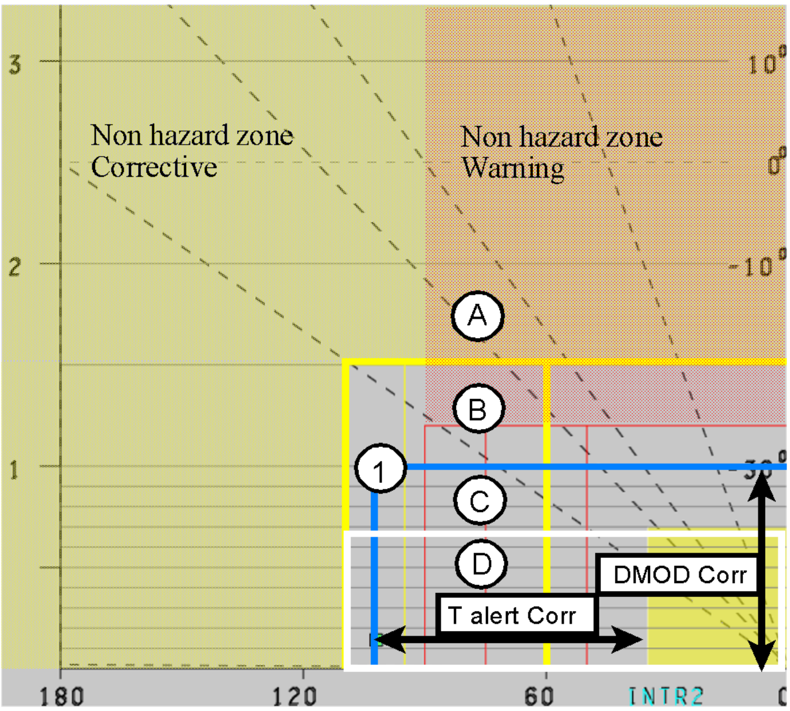

Figure 2 – TCPA-DCPA Chart with Example Alert Threshold [4]

The selection of the alerting threshold should follow from (statistical) information on sensor performance. If real-time information on the measurement uncertainty is available, a more robust approach for reducing missed alerts is to use this information to estimate the 95%[4] TCPA-DCPA containment for the received position and velocity data and test whether that containment intersects with an alerting threshold. In such a configuration, the alerting thresholds can be kept closer to the actual HAZ boundary.
In the DO-365B reference implementation [2], a SUM approach that is further detailed in [5] is applied. This SUM approach uses a 'wrapper' that besides the measured state also uses information about the standard deviations of measured position- and velocity data to feed the alerting algorithm with variations to the state to determine the worst-case in terms of future separation. To reduce the amount of intermittent alerts, m-out-of-n filtering is used.

2. Impact of sensor noise

To allow a designer to assess the alerting performance in the presence of sensor noise, Appendix P of [2] contains a total of 444 en-route test-vectors. Of these, 302 represent 'must-alert' situations and 84 'must not alert'. The remaining 58 represent 'may alert' situations. These are conflict geometries in which the intruder leaves the NHZ but never enters the HAZ. Each test-vector is available as a truth-reference, as the simulated output of an ADSB, mode C, mode S or ATAR sensor, and as the simulated output of a tracker that uses one of the previous four simulated sensor-outputs as input. Figure 3 provides an overview of the relation between these test-vectors.

[4] The 95% is an example. The actual percentage should follow from alerting performance requirements. Increasing the percentage will also increase the likelihood of false alerts.

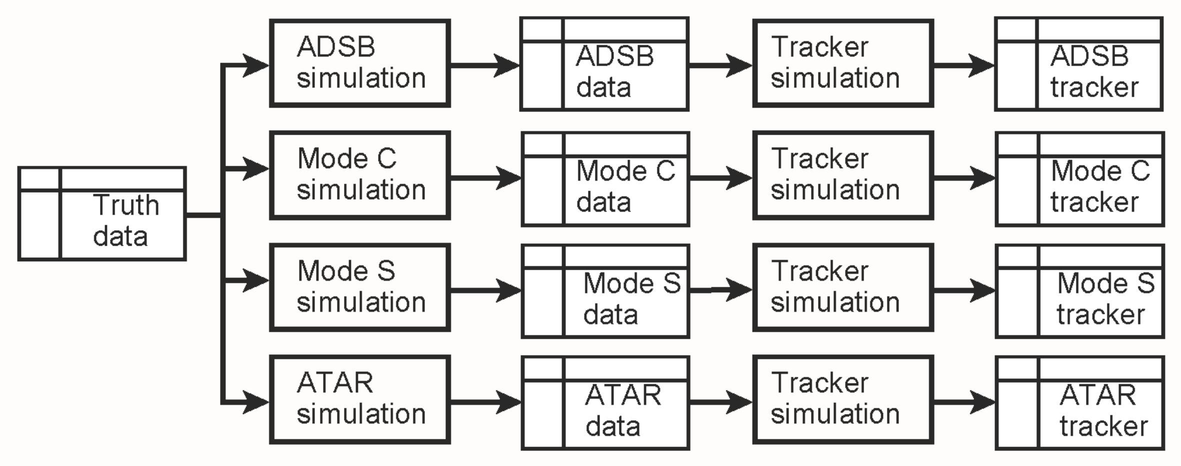

Figure 3 – Relation between test-vector truth-, sensor- and tracker representations from [2]

Figure 4 shows the plan-view of a truth-vector, an ADSB tracker output and a mode C tracker output. In these plan-views, the blue line represents ownship trajectory, with the start location indicated by a circle. In the ADSB- and mode C tracker plots, the red-dotted line represents the true intruder trajectory and the red-dashed line the estimated intruder trajectory. The estimated start location is indicated by a circle.

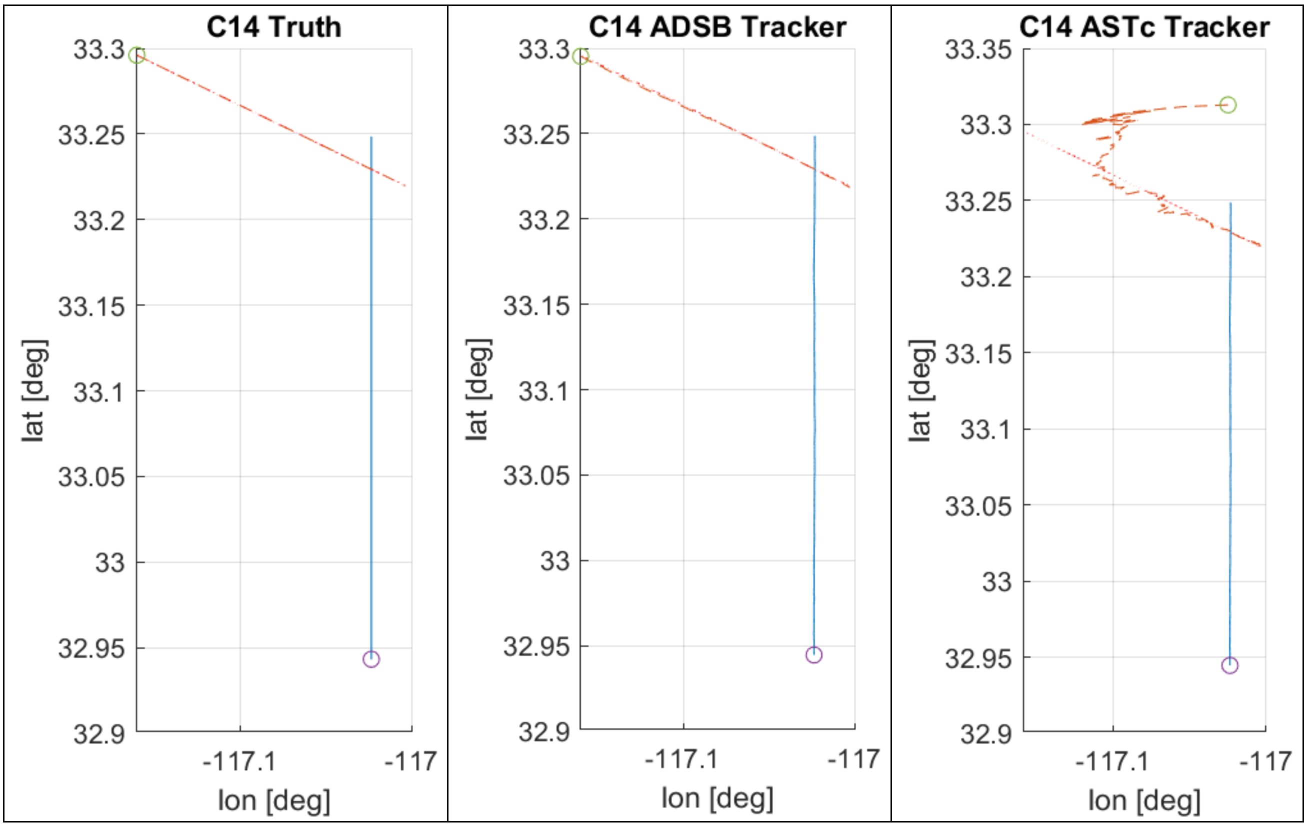

Figure 4 – Plan view of trajectory from test-vector C14 using truth, ADSB tracker and ASTc tracker data

As can be seen, the plan-view based on the ADSB tracker data is very similar to the plan-view based on the truth data. However, the plan view based on the Mode C tracker data illustrates that initially the estimated position is rather different from the true position. It takes some time before the estimates of the intruder position converge around the true trajectory. This can have an impact on alerting performance.

2.1 Sensor-specific and geometry specific accuracy performance

An ADSB report contains both the position and velocity as measured by the reporting traffic. As a result, the position- and velocity accuracy for ADSB data does not depend on distance to ownship or conflict geometry. This is different for the other sensors; Mode C and Mode S estimate range to traffic from the time between the transmitted interrogations and the received replies. Bearing is estimated using a direction-finding antenna. The accuracy of the resulting position estimate is inversely proportional to range. A tracker requires multiple range-bearing measurements to estimate the relative velocity vector. Because the position error in the direction of the range measurement is typically (much) smaller than the position error perpendicular to it, the accuracy of the relative velocity is a function of conflict geometry (i.e. head-on, crossing, overtake). In the DO-365B tracker output, the sensor uncertainty is specified as 1-sigma values for the North-South direction and the East-West direction, and the error-covariance. Figure 5 depicts the position- and velocity uncertainty as a function of TCPA for tracker output of three different sensors for test-vector C14.

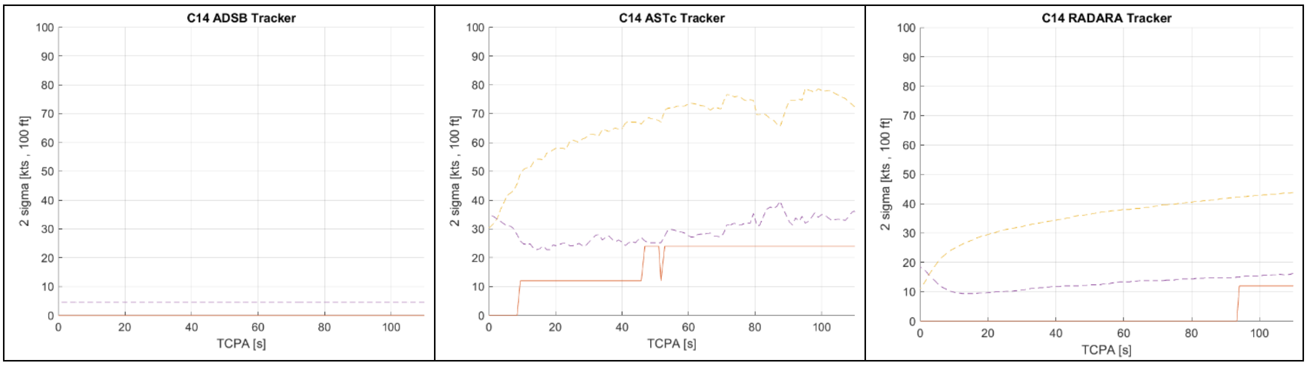

Figure 5 – 2 sigma value of position and velocity uncertainty for ADSB, mode C and RADAR tracker

In the plots in Figure 5, the East-West velocity uncertainty is represented by the dashed orange line. The North-South velocity uncertainty is represented by the dashed purple line. The scale on the vertical axis is in [kts] for these velocity uncertainties. For the 1-sigma position uncertainty, the resolution of the source data is 0.1 NM, which is about 600 ft. The red solid line in the above figures depicts the 2-sigma position uncertainty. The numbers of the vertical scale represent the uncertainty divided by 100, e.g. 12 represents 1200 ft. Given the magnitude of the velocity uncertainty in case of the mode C and RADAR tracker, this will be the determining factor for missed- and incorrect alerts. The next section will discuss a method to illustrate how such uncertainties impact the alerting decision. Whereas Figure 2 depicted the alerting thresholds in the TCPA-DCPA domain, the next section will present the impact of sensor noise in the TCPA-DCPA domain.

3. Depicting the impact of sensor noise in the TCPA-DCPA domain

In Figure 2 it was illustrated how the depiction of the computed TCPA and DCPA of an intruder relative to the thresholds depicted in the TCPA-DCPA space can be used to determine the horizontal alert condition. The transformation of subsequent trajectory points of ownship and intruder to the TCPA-DCPA space can be used to plot a TCPA-DCPA profile as presented on the right in Figure 6. Note that compared to Figure 2, the origin of the plot is now on the left side, and TCPA increases from left to right. As a result, subsequent TCPA-DCPA points for a converging encounter will move from right to left.

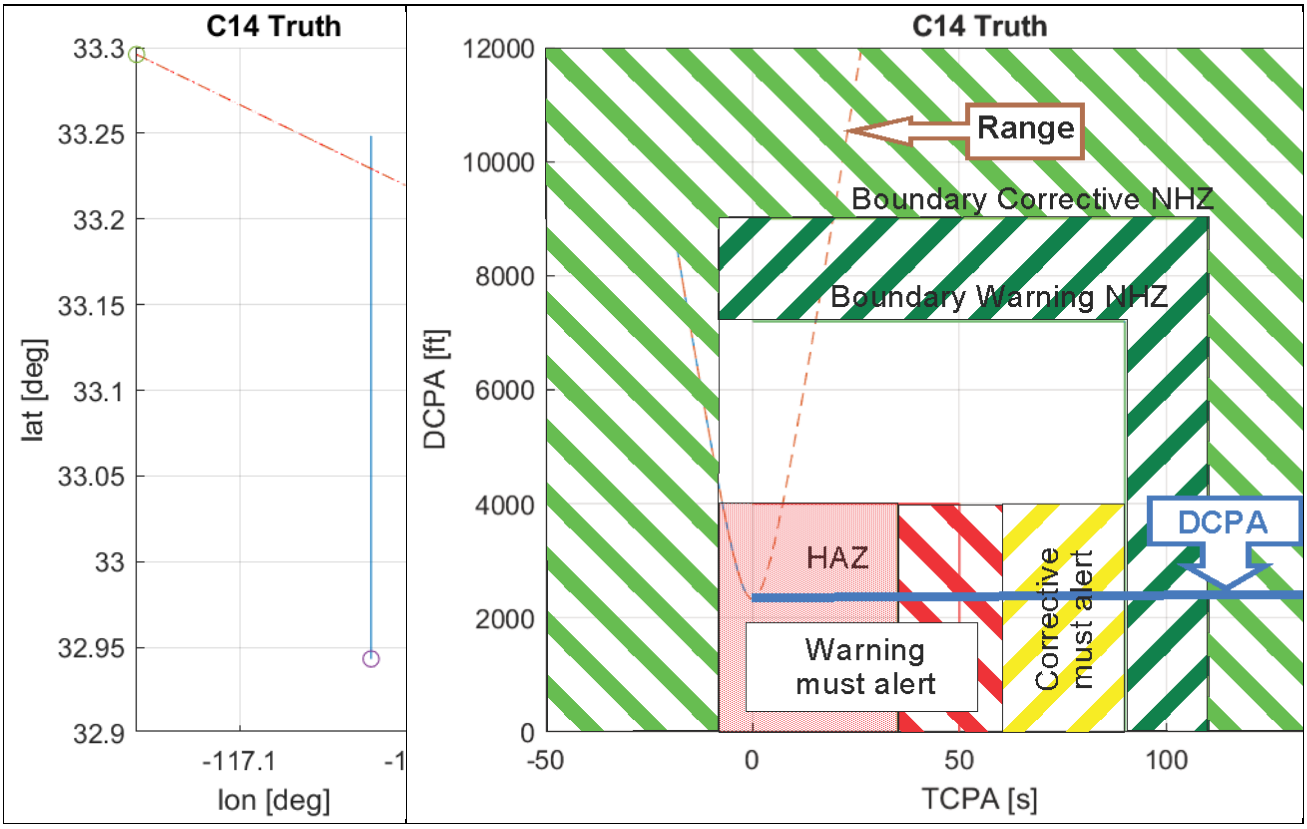

Figure 6 – Trajectory plan-view and associated TCPA-DCPA profile for test-vector C14 truth data

In the test-vectors, for ADSB data the position and velocity noise is based on a Gauss-Markov model with a 5 minute decorrelation [2]. For Mode C, per test-vector the range noise is based on a single sample from a uniform distribution for the bias and Gaussian noise to represent the additional jitter. The bearing noise is Gaussian [2].

In case the accuracy of velocity is independent on range to the intruder and direction of the relative velocity (as is the case with ADSB), the 2-sigma value of the velocity noise can be used to compute and depict the 95% TCPA-DCPA containment space (due to velocity uncertainty) relative to the truth-line. Figure 7 shows the results for two test-vectors using a 2-sigma value of 20 kts in the North-South direction and in the East-West direction.

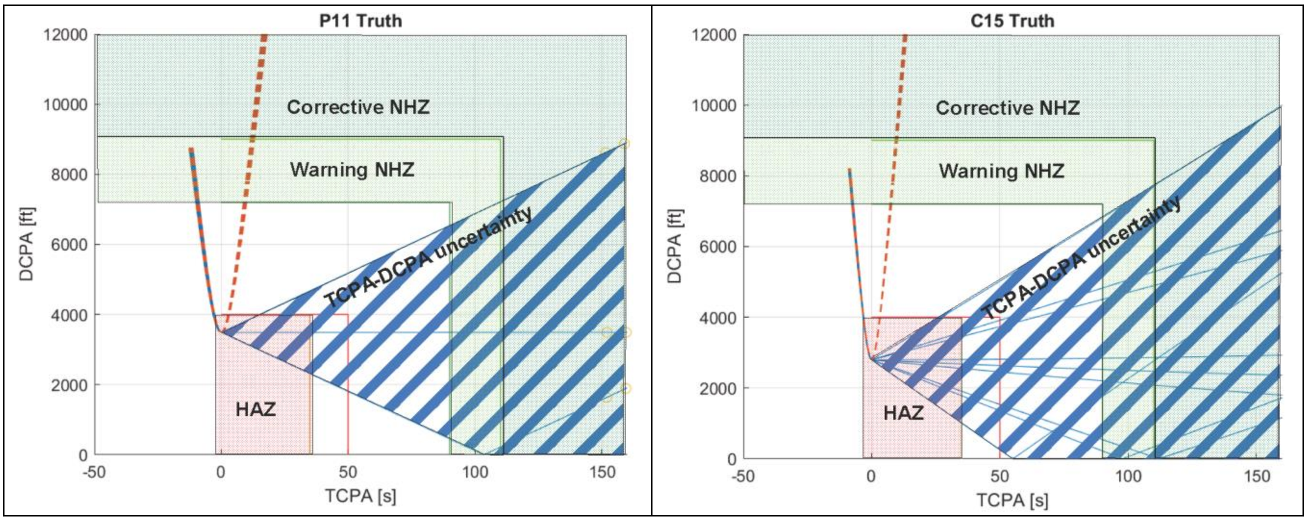

Figure 7 – TCPA-DCPA uncertainty space for P11 and C15 using a velocity of 20 kts as 2-sigma value

The 2-sigma contours intersect the vertical NHZ lines before the (extended) horizontal thresholds would be crossed. This indicates that for the P11 and the C15 must-alert test-vectors, even with the limited velocity accuracy (20 kts), fewer than 5% of the TCPA-DCPA points will be in the NHZ determined by the HMD[5]. Since a designer can choose a spatial alerting threshold that is larger than the HAZ threshold, the insight obtained from this TCPA-DCPA depiction also helps to understand why increasing the threshold with an increase in TCPA will help increase alerting performance. However, this will come at a cost of an increase is false alerts.

Unfortunately, the position and velocity accuracy of mode C, mode S and ATAR depend on range and conflict geometry. Hence, the simplification used in Figure 7 is only helpful for development/configuration of SUM for ADSB traffic. Figure 8 shows the TCPA-DCPA for C14 truth, ADSB tracker and mode C tracker.

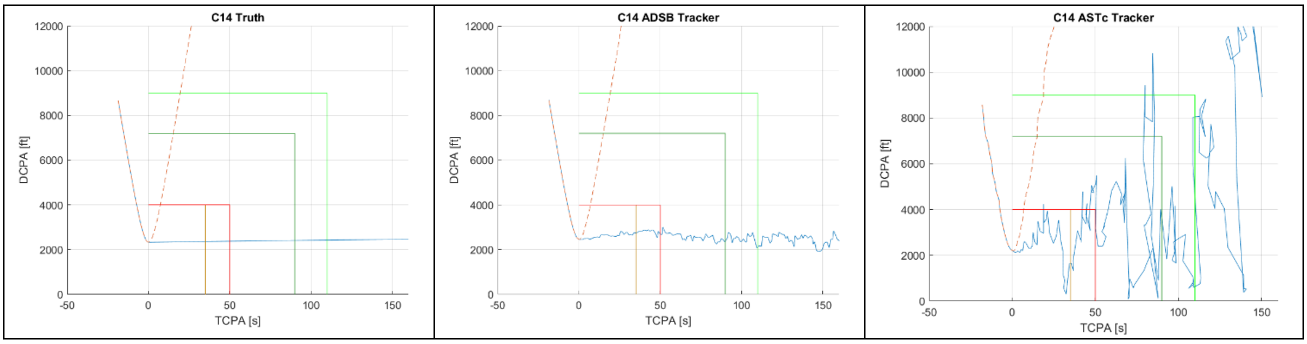

Figure 8 – TCPA-DCPA plots for C14 Truth, ADSB tracker and Mode C tracker

3.1 Visualizing the 2-sigma velocity impact using Monte-Carlo

Whereas a TCPA-DCPA plot of a single degraded test-vector such as depicted in Figure 8 provides an indication about the severity of the sensor noise, a much larger number would be needed to create a graphic from which the contours of the TCPA-DCPA uncertainty space can be observed. The underlying data can also be used to estimate the likelihood that points with a true HMD in the HAZ end up in the NHZ and vice-versa.

In Figure 3 it was illustrated how the original degraded test-vectors were created by SC-228. To create additional degraded test-vectors, both a sensor model/simulation and a tracker model/simulation are required. Since these simulations are not available at the Netherlands Defence Academy (NLDA), a simplified method to generate additional degraded test-vectors has been applied. To generate additional samples of a particular test-vector with representative sensor-noise, the specification of the position- and velocity uncertainty (available in the output from the tracker files for each individual position and velocity sample) has been used as the standard-deviation scaling factor of a random-number generator. For each entry in the truth vector, the associated standard deviations of position- and velocity errors are used to generate new errors which are subsequently added to the truth-vectors. For the examples used in this paper, this process is repeated 20 times for each truth-vector, yielding 20 new degraded test-vectors with the same statistical properties in terms of the distribution of the error[6]. Figure 9 provides an overview of the approach that has been used.

[5] The uncertainty in velocity also causes an uncertainty in TCPA. This can cause early alerting when the actual TCPA is still in the NHZ and the computed TCPA has left the NHZ.

[6] It is recognized that aspects such as correlation between subsequent samples and a bias are not captured by this approach. However, these would only be relevant if the resulting data were used to drive an alerting algorithm. For the application in this paper, the data was only intended to be used to depict the resulting uncertainty in TCPA-DCPA space and generate percentages of TCPA-DCPA points in the HAZ and NHZ.

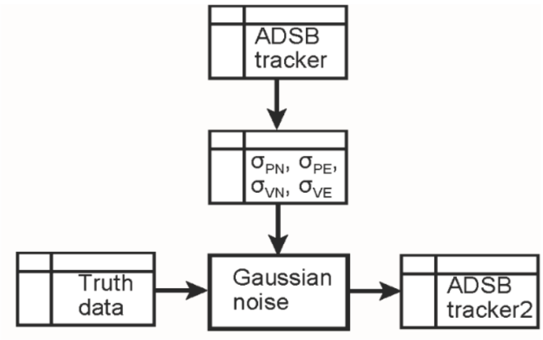

Figure 9 – Approach to create more simulated noisy samples of a single test-vector

Figure 10 presents the results for a conflict geometry with a high closure-rate. The red U-shape indicates range, the blue lines are the DCPA traces as a function of TCPA. The impact of velocity noise will increase with a decrease in relative velocity. This can be seen in Figure 11.

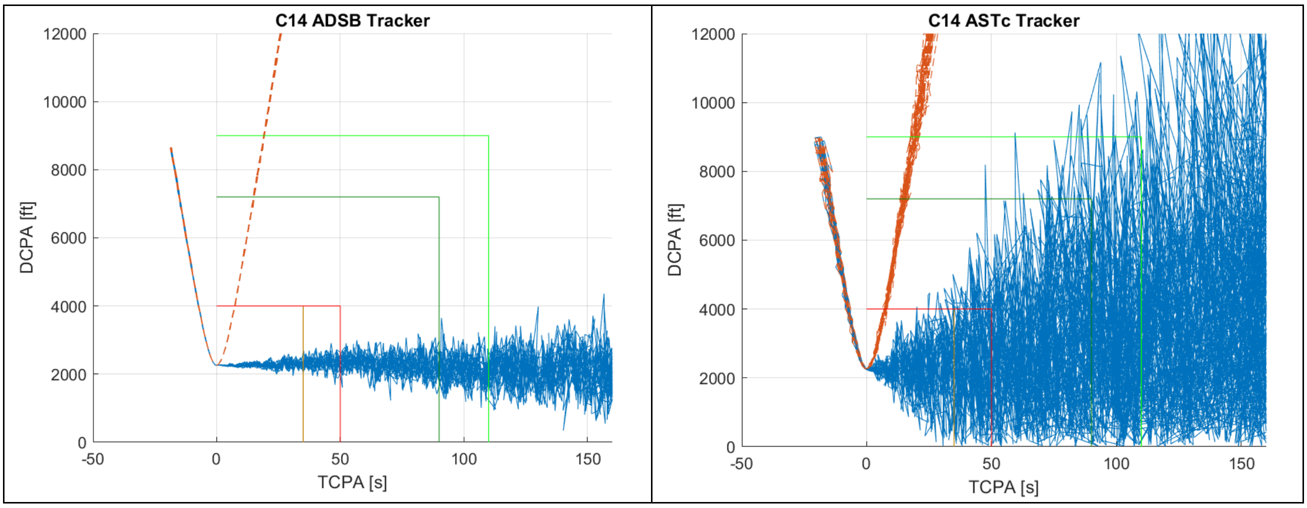

Figure 10 – TCPA-DCPA plot for 20 simulated instances of C14 from ADSB and mode C tracker

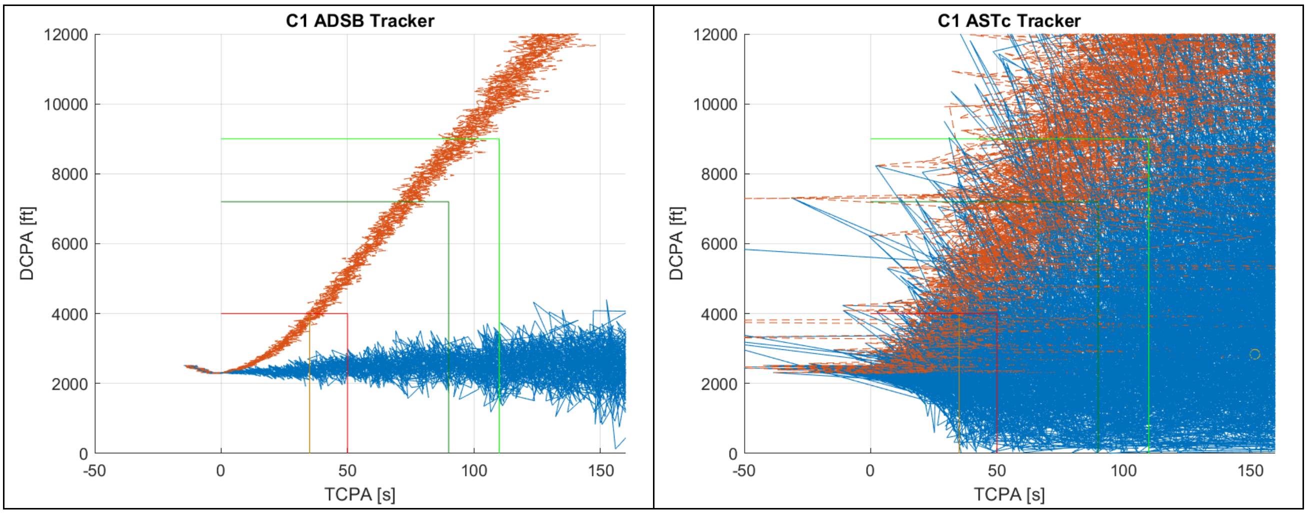

Figure 11 – TCPA-DCPA plot for 20 simulated instances of C1 from ADSB and mode C tracker

Although the right TCPA-DCPA plot in Figure 10 may still look good from a must alert perspective, it also shows that there is a considerable likelihood that DCPA may remain in the NHZ, and this above any reasonable alert threshold. The use of a scaled standard deviation of position and velocity to vary the input to the alerting is a technique used to increase the likelihood of a timely alert in such situations [2,5]. However, this comes at a cost of an increase in false alerts in non-alerting geometries. This becomes clear when looking at an example of a TCPA-DCPA plot for a ‘must not alert’ test-vector (Figure 12).

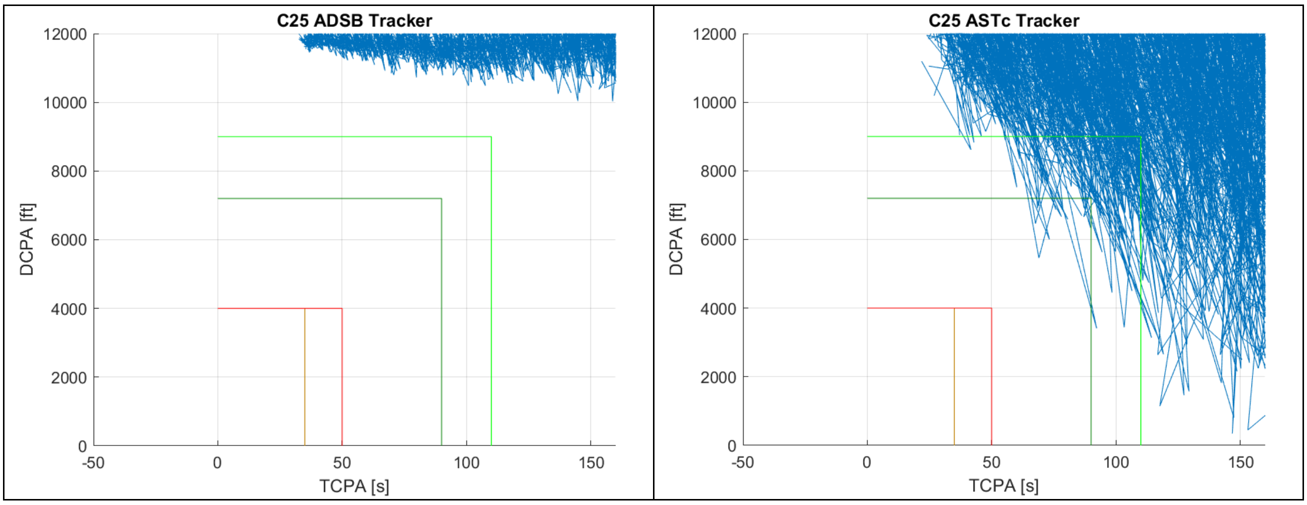

Figure 12 – TCPA-DCPA plot for 20 simulated instances of C25 from ADSB and mode C tracker

The right TCPA-DCPA plot shows that for mode C even without the addition of a scaled standard-deviation a considerable part of the TCPA-DCPA points below TCPA=110 (the early alert threshold) extends from the NHZ to a DCPA smaller than 4000 ft, the minimum spatial threshold that should be used. One take-away from plots such as Figure 12 is that for mode C some occasional false alerts seem unavoidable.

3.2 Sensitivity to noise when close to alerting threshold

In DO-365B Appendix K, agreed exemptions for ACAS-Xu failure to particular test-vectors are provided. As one of the reasons for agreement on an exemption it is indicated that the test-vectors include situations ‘*in which the test vector geometry places the aircraft on or near the various alerting boundaries, or the boundaries are marginally violated. Once noise is applied, the proximity to the alerting boundaries makes passing or failing that encounter highly susceptible to noise*’. The examples in Figures 13 and 14 illustrate this issue for a ‘must-alert’ and ‘must-not-alert’ geometry.

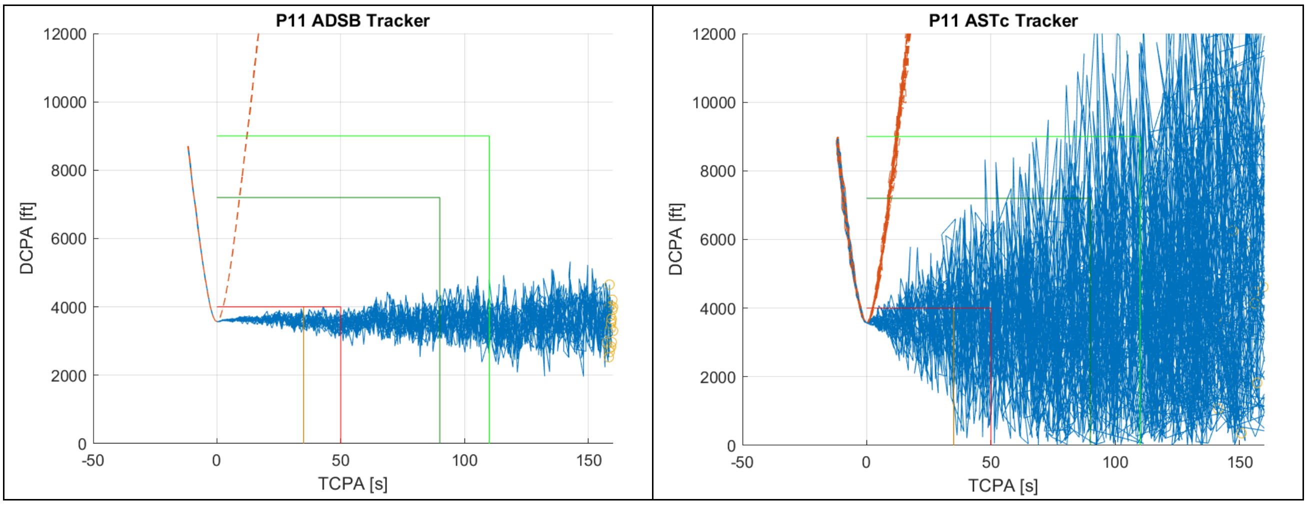

Figure 13 – TCPA-DCPA plot for 20 simulated instances of P11 from ADSB and mode C tracker

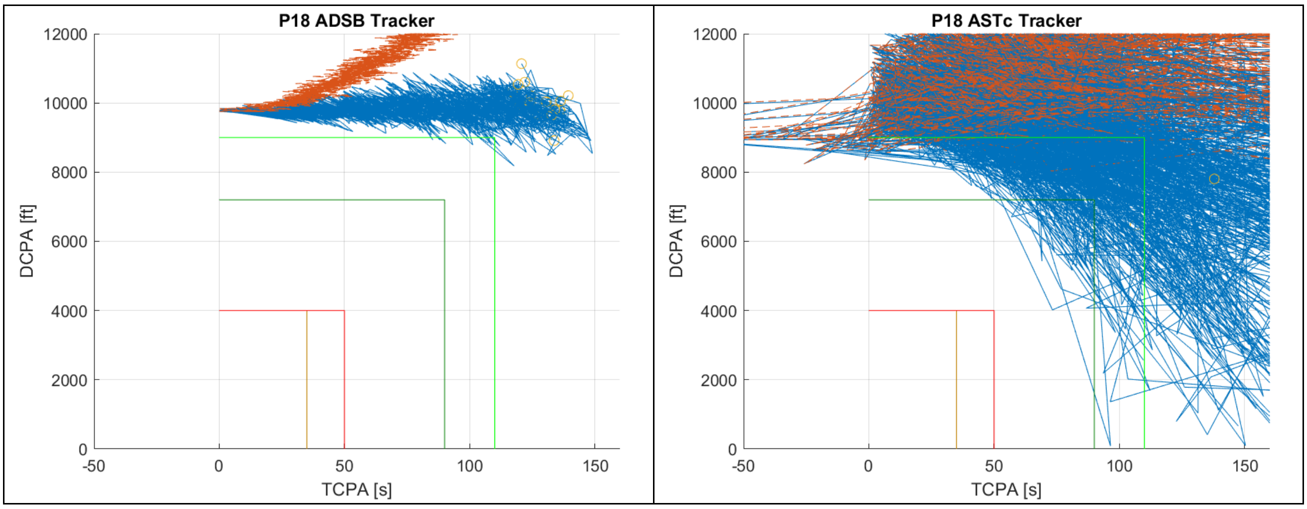

Figure 14 – TCPA-DCPA plot for 20 simulated instances of P18 from ADSB and mode C tracker

4. SUM

4.1 Introduction

In the operational environment, the specification of position- and velocity uncertainty can be available in different forms. ADSB messages may include so-called NACp[7] and NACv[8] data. Such data is not directly available for mode C and mode S traffic. A tracker that estimates traffic position and velocity from subsequent ranges and bearings (from mode C, mode S or ATAR) should be able to provide a measure of uncertainty for these estimates. However, situations may also occur in which no uncertainty data is available.

4.2 SUM strategies

Using the reported uncertainty, variations to the reported position and velocity can be input to the alerting algorithm to find the associated TCPA and DCPA and test whether an alert condition is met for any of these combinations. However, if these variations yield both options in the HAZ and in the NHZ, the uncertainty is too large to declare a reliable alert. To reduce the likelihood of having outcomes both in the HAZ and NHZ, [5] uses range-based scaling of the standard deviations added to the reported data. Another (additional) possibility to deal with too much uncertainty is to use the available margins between the early and late-alert time to reduce the alert-time threshold in case the variations yield both outcomes in the NHZ and HAZ. In case no uncertainty data is available, sensor-specific assumptions can be used to increase the threshold against which HMD_P is tested, and also make it a function of range and/or TCPA.

In the reference implementation from [2], scaled standard deviations are used to create states located on the contour of a hockey puck around the reported state. These new states are used as input to the alerting algorithm. The scaling factors provide the designer with a certain trade-space. Clearly, reducing late- and missed alerts by increasing the scaling factors will also increase the potential for incorrect alerts. The likelihood of such incorrect alerts can be reduced with the application of m-out-of-n filters that will require m alert conditions within n tests to occur before an alert is declared. A larger value of m will increase the capability to reject incorrect alerts but will also increase the time between the occurrence of an alert condition and the actual alert. For non-maneuvering geometries this can partly be compensated by using a larger value for the alert-time. However, the maximum value of m is limited by the alert-time requirements since for maneuvering encounters in which an intruder quickly transitions from the NHZ to the HAZ, an alert is required no later than 5 seconds after entering the HAZ. Using dynamic m and n, e.g. based on TCPA, a faster

[7] NACp is an integer value that references a maximum standard deviation of the reported horizontal position. A lookup table is needed to convert NACp to the standard deviation.
[8] NACv is an integer value that references a maximum standard deviation of the reported horizontal velocity. A lookup table is needed to convert NACv to the standard deviation.

response at lower TCPA can be combined with better reduction of incorrect alerts at higher TCPA (where the impact of velocity uncertainty is also more likely to cause incorrect alert conditions).

It is evident that besides the design of SUM algorithms, the many degrees of freedom provide the designer with a challenge concerning the quantification of the relevant design parameters. These include but are not limited to:

- The scaling factors for the standard deviations of horizontal- and vertical position
- The scaling factors for the standard deviations of horizontal- and vertical velocity
- Any dynamic scaling factors that make the above scaling dependent on range and/or TCPA
- The m and n for each sensor-specific filter
- Any TCPA thresholds to dynamically configure the m-out-of-n, with associated m and n
- The actual HMD used to test HMDp against (this may be larger than the HMD that defines the HAZ), possibly sensor-specific
- The time-thresholds use to declare an alert (these may be made dynamic to reduce the impact of too much uncertainty)

Also, given the required consistency with guidance, it is essential to realize that these quantifications will have impact on the guidance. As illustrated in Figure 1, an increase in uncertainty will increase the directions to be avoided.

Using the methods to analyze the impact of sensor-uncertainty in the TCPA-DCPA domain together with test-vectors from [2] that are both available in truth- and degraded form, the impact of configuration changes can be related to the domain in which the thresholds for alerting are defined. This aids in the development of the rationale underlying the quantification of the design parameters. The next section compares the performance of a prototype SUM implementation against two conceptual performance baselines without SUM.

4.3 SUM performance improvement

4.3.1 Implementation

To realize acceptable alerting performance, a SUM function has been designed and implemented. The configuration has been performed using the results from fast-time simulation with the test-vectors from [2] as input. Given the identified impact of position- and velocity errors in the TCPA-DCPA domain, the following SUM strategies have been implemented to reduce late- and missed alerts, reduce intermittent alerts and keep false alerts to a minimum:

- Dynamic (i.e. sensor type and TCPA dependent) spatial threshold against which DCPA is tested
- Dynamic (sensor type dependent) alert-time
- Position- and velocity variations as input to the alerting function using scaled standard deviations
- Sensor and TCPA-dependent m-out-of-n filter for the transition from non-alert to alert to reduce early and intermittent alerting
- Alert-coasting in the transition from alert to non-alert to reduce intermittent alerting

4.3.2 Establishing the baseline for performance

To assess the performance of the SUM approach, two baselines were established. The first baseline is obtained by using the HAZ thresholds for alerting. Any DCPA larger than the HMD defined by the HAZ is regarded as a non-alert. In other words, none of the available ‘may alert’ trade-space is used. In the current paper, only the results for ADSB and Mode C are presented. Table 1 provides an overview of the ‘Must Alert’ performance using the ADSB tracker test-vectors, Table 2 lists the performance for the ‘Must not Alert’ ADSB tracker test-vectors. For the Mode C tracker test-vectors, Table 3 provides an overview of the ‘Must Alert’ performance and Table 4 for the ‘Must not Alert’ performance.

Table 1 - Must-alert performance using HAZ threshold and ADSB test-vectors

	Early alerts			Late- and missed alerts		
	Early Caution	Early Warning	Early Caution & Early Warning	Late Caution	Late Warning or no alert	Late Caution & Late Warning
Alerts	0	0	0	10	3	23
Percentage	0	0	0	3.3	1.0	7.6
Total %	0			11.9		

Table 2 - Must-not-alert performance using HAZ threshold and ADSB test-vectors

	False alerts		
	Caution	Warning	Caution & Warning
Alerts	1	2	1
Percentage	1.2	2.4	1.2
Total %	4.8		

Table 3 - Must-alert performance using HAZ threshold and Mode C test-vectors

	Early alerts			Late- and missed alerts		
	Early Caution	Early Warning	Early Caution & Early Warning	Late Caution	Late Warning or no alert	Late Caution & Late Warning
Alerts	10	6	8	6	6	20
Percentage	3.3	2.0	2.6	2.0	2.0	6.6
Total %	7.9			10.6		

Table 4 - Must-not-alert performance using HAZ threshold and Mode C test-vectors

	False alerts		
	Caution	Warning	Caution & Warning
Alerts	8	7	8
Percentage	9.5	8.3	9.5
Total %	27.3		

What stands out is the much larger percentage of false alerts for Mode C (alerts in the ‘Must not alert’ set of test-vectors). Given the earlier depiction of variation in the TCPA-DCPA domain (i.e. Figures 12 and 14), this is not surprising.

The second baseline for performance is established by moving the decision threshold from the HAZ to the Warning alert NHZ HMD, i.e. 1.2 NM. Any DCPA below this value will be considered as a trigger for an alert (within the alert time). This should reduce the number of late-and missed alerts, but at the expense of an increased amount of false alerts.

Table 5 provides an overview of the ‘Must Alert’ performance using the ADSB tracker test-vectors, Table 6 lists the performance for the ‘Must not Alert’ ADSB tracker test-vectors. For the Mode C tracker test-vectors, Table 7 provides an overview of the ‘Must Alert’ performance and Table 8 for the ‘Must not Alert’ performance.

Table 5 - Must-alert performance using NHZ threshold and ADSB test-vectors

	Early alerts			Late- and missed alerts		
	Early Caution	Early Warning	Early Caution & Early Warning	Late Caution	Late Warning or no alert	Late Caution & Late Warning
Alerts	1	34	23	4	1	6
Percentage	0.3	11.3	7.6	1.3	0.3	2.0
Total %	19.2			3.6		

Table 6 - Must-not-alert performance using NHZ threshold and ADSB test-vectors

	False alerts		
	Caution	Warning	Caution & Warning
Alerts	1	2	1
percentage	1.2	2.4	1.2
Total %	4.8		

Table 7 - Must-alert performance using NHZ threshold and Mode C test-vectors

	Early alerts			Late- and missed alerts		
	Early Caution	Early Warning	Early Caution & Early Warning	Late Caution	Late Warning or no alert	Late Caution & Late Warning
Alerts	12	26	33	4	2	10
Percentage	4.0	8.6	10.9	1.3	0.7	3.3
Total %	23.5			5.3		

Table 8 - Must-not-alert performance using NHZ threshold and Mode C test-vectors

	False alerts		
	Caution	Warning	Caution & Warning
Alerts	6	10	14
percentage	7.1	11.9	16.7
Total %	35.7		

Table 9 provides a summary of the data in Tables 1 to 8, allowing a comparison of the two baselines for the performance per sensor.

Table 9 – Summary of the baseline performance for thresholds at HAZ and NHZ

		ADSB tracker		Mode C tracker	
	HMD	At HAZ	At NHZ	At HAZ	At NHZ
Must Alert	Early [%]	0	19.2	7.9	23.5
	Late and missed [%]	11.9	3.6	10.6	5.3
Must Not alert	False alert [%]	4.8	4.8	27.3	35.7

The results in Table 9 confirm that indeed the amount of missed- and late alerts reduces significantly when increasing the HMD threshold from the HAZ to the NHZ. For ADSB, the reduction is from 11.9% to 3.6% and for mode C from 10.6% to 5.3%. Furthermore, for Mode C the results in Table 9 confirm what can be expected as the cost for this improvement, i.e. moving the alerting threshold towards the NHZ increases the amount of false alerts. For ADSB however, the amount of false alerts does not increase. This is not too surprising since the margin between the HAZ and the NHZ is significantly larger than the impact of the position- and velocity uncertainty with ADSB. But there is another cost: Both for ADSB and mode C, the amount of Early alerts increases.

4.3.3 Performance using SUM

This section provides an overview of the alerting performance with the application of the SUM function used for the current research. Table 10 provides the performance for the must-alert test-vectors using the ADSB tracker output files from [2], and Table 11 shows the performance on the 'must-not-alert' ADSB tracker output files. Similarly, Table 12 provides the performance for the must-alert test-vectors using the mode C tracker output files from [2], and Table 13 shows the performance on the 'must-not-alert' mode C tracker output files.

Table 10 - Must-alert performance using SUM and ADSB test-vectors

	Early alerts			Late- and missed alerts		
	Early Caution	Early Warning	Early Caution & Early Warning	Late Caution	Late Warning or no alert	Late Caution & Late Warning
Alerts	2	0	0	0	2	0
Percentage	0.7	0	0	0	0.7	0
Total %	0.7			0.7		

Table 11 - Must-not-alert performance using SUM and ADSB test-vectors

	False alerts		
	Caution	Warning	Caution & Warning
Alerts	2	1	1
Percentage	2.4	1.2	1.2
Total %	4.8		

Table 12 - Must-alert performance using SUM and Mode C test-vectors

	Early alerts			Late- and missed alerts		
	Early Caution	Early Warning	Early Caution & Early Warning	Late Caution	Late Warning or no alert	Late Caution & Late Warning
Alerts	1	7	3	5	3	4
Percentage	0.3	2.3	1.0	1.7	1.0	1.3
Total %	3.6			4.0		

Table 13 - Must-not-alert performance using SUM and Mode C test-vectors

	False alerts		
	Caution	Warning	Caution & Warning
Alerts	3	10	11
Percentage	3.6	11.9	13.1
Total %	28.6		

4.4 Discussion

For ADSB targets, just moving the threshold from the HAZ to the NHZ reduces missed / late alerts from 11.9 to 3.6% at the cost of an increase from 0 to 19.2% early alerts. The percentage of incorrect alerts stays at 4.8%.

Applying SUM reduces missed / late alerts from 11.9 to 0.7%, while increasing early alerts from 0 to 0.7% and maintaining incorrect alerts at 4.8%.

For Mode C targets, just moving the threshold from the HAZ to the NHZ reduces missed / late alerts from 10.6 to 5.3% at the cost of an increase from 7.9 to 23.5% early alerts and from 27.3 to 35.7% incorrect alerts.

Applying SUM reduces missed / late alerts from 10.6 to 4.0%, while also reducing early alerts from 7.9 to 3.6% at the cost of an increase in incorrect alerts from 27.3 to 28.6%.

Thus, for must-alert situations, SUM achieves better performance than the best baseline with the HMD threshold at the NHZ (all TCPA-DCPA points in the may-alert zone considered as alerting). And for must-not alert situations, SUM performance is not worse than the baseline with the HMD threshold at the HAZ (all TCPA-DCPA points in the may-alert zone considered as non-alerting).

Unfortunately, a comparison with the performance of the reference implementation from [2] is not possible, since only the performance on the 'must-alert' test vectors is reported. The 'cost' of using that approach to SUM is not quantified in terms of incorrect alerts for the must-not alert test-vectors.

5. Conclusion

In [2], the Minimum Operations Performance Standards for UAS DAA systems are provided and it is indicated that "*Implementing an alerting and guidance scheme using simple predictions and the HAZ values directly from the requirements will result in unacceptable performance*".

This unacceptable performance is characterized by a too high percentage of late- and/or missed alerts, a too high percentage of nuisance alerts and too many intermittent alerts. SUM is intended to achieve the required performance by addressing all three issues.

SUM uses information about the measurement uncertainty to determine whether TCPA-DCPA estimates that lie in the area between the HAZ and the NHZ can be considered as a potential alert state. SUM also comprises filtering (e.g. m-out-of-n) to limit the number of incorrect alerts.

Besides the design of SUM algorithms, the many degrees of freedom resulting from the SUM design parameters provide the designer with a challenge concerning the quantification of these parameters.

By analyzing the impact of sensor-uncertainty in the TCPA-DCPA domain using the test-vectors from [2] that are both available in truth- and degraded form, designers can tune the quantification based on an understanding of the expected impact on the alerting performance and an understanding of fundamental limitations caused by the magnitude of the sensor-noise. If a tracker different than the reference tracker is used, the test-vectors with the raw sensor data can be used as an input to such a tracker. This may result in differences in the reported standard deviations and possibly require a different SUM configuration.

6. Copyright Statement

References

[1] FAA. *TSO-C211 Detect and Avoid (DAA) Systems*, U.S. Department of Transportation, 2017.

[2] RTCA. *DO-365B - Minimum Operational Performance Standards (MOPS) for Detect and Avoid (DAA) Systems Revision B*. 2021.

[3] Tadema, J., E. Theunissen, R.M. Rademaker and M. Uijt de Haag. Evaluating the Impact of Sensor Data Uncertainty and Maneuver Uncertainty in a Conflict Probe. *Proceedings of the 29th Digital Avionics Systems Conference*, pp. 4C3.1-4C3.16, Salt Lake City, UT, 2010.

[4] Theunissen, E., B. Suarez and F. Kunzi. Designing a Robust Detect and Avoid Alerting Function. *Proceedings of the 17th Integrated Communications Navigation and Surveillance Conference*, Herndon, VA, 2017.

[5] Anthony Narkawicz, César Muñoz, and Aaron Dutle. Sensor Uncertainty Mitigation and Dynamic Well Clear Volumes in DAIDALUS, *Proceedings of the 37th Digital Avionics Systems Conference*, London, England, UK, 2018.