

Robust Conflict Detection and Resolution around Top of Descent

Andrew C. Cone,^{*} Aisha R. Bowe[†] and Todd A. Lauderdale[‡]

NASA Ames Research Center, Moffett Field, CA, 94035

One of the complicating factors when detecting and resolving aircraft-to-aircraft conflicts is trajectory prediction uncertainty, which can cause conflicts to be detected without sufficient time to resolve them. Previous work by the authors showed this situation occurs most often in Center Airspace when an aircraft is between about 10 nmi before its top-of-descent point and its arrival fix at the edge of the terminal area. This paper explores a pair of enhanced vertical conflict detection ranges (buffers) for aircraft descending to their meter fix that could provide enough coverage to catch potential losses for multiple types of uncertainty while limiting the increase in false alerts. The specific uncertainties examined include the predicted wind speed, cruise speed, descent speed, top-of-descent location, and the combination of all of those uncertainties plus uncertainty in aircraft fuel weight prediction. Performance metrics include false alerts, missed alerts, losses of separation, number of resolutions issued, and the total system delay caused by aircraft flying conflict avoidance maneuvers. Results show that these vertical buffers reduce the number of losses of separation for arriving aircraft from 207 to 12. However, using the vertical buffers increases the number of resolutions issued by 50% and doubles the delay accrued by aircraft flying conflict resolution maneuvers. Results also suggest that a smaller buffer (80% size) could be used to gain most of the same benefit as the full buffer with less additional delay, while alternative methods might be best suited to remove the last few loss of separation cases during descent. Further, improving the trajectory prediction accuracy combined with a vertical buffer significantly reduced the number of losses of separation, resolutions issued, and delay accrued.

I. Introduction

CONFLICT prediction is an integral part of maintaining safe separation between aircraft in the National Airspace System. This task is currently handled completely by human controllers, but as air traffic demand continues to rise, automation will play a more prominent role in conflict detection and possibly resolution. To enable this transition, it is important to develop automation tools that can operate reliably and with a high degree of accuracy in real-world settings.

One of the primary challenges in predicting conflicts using automation in a realistic setting is dealing with the effects of uncertainty. Automated conflict probes often utilize some type of trajectory generator to build a predicted trajectory, and this predicted trajectory never exactly matches what the aircraft will actually fly. Depending on the type and magnitude of the uncertainty, as well as the capabilities of the trajectory predictor, these errors can range from relatively minor errors, such as differences in turn modeling, to major issues like intent errors that can greatly affect the accuracy of trajectory predictions.¹ These errors, in turn, affect the ability of the conflict probe to detect conflicts and to suggest resolutions that are free from conflict.

There are two primary ways to address the issue of errors in the trajectory prediction: improve the accuracy of the trajectory generator, or compensate for the error by adapting the conflict probe. For this work, the focus is on compensating for errors, rather than trying to improve trajectory predictions. Specifically, the authors are detecting and resolving conflicts at greater than the legal standard for separation of five nmi horizontally and 1,000 ft vertically using a geometric conflict detection scheme. While improving

^{*}Aerospace Engineer, Aviation Systems Division, Mail Stop 210-10, and AIAA member.

[†]Aerospace Engineer, Aviation Systems Division, Mail Stop 210-10, and AIAA member.

[‡]Aerospace Engineer, Aviation Systems Division, Mail Stop 210-10.

trajectory prediction accuracy is beneficial for detecting conflicts when using automated conflict probes,^{2,3,4,5} and validating the accuracy of a trajectory prediction is seen as a necessary part of the future National Airspace System⁶ (NAS), improving trajectory prediction performance often requires some form of equipment on the aircraft and/or data sharing, which can make the solution more difficult and expensive to implement in actual operations. There are alternative approaches that try to use adaptive algorithms to improve the accuracy of trajectory predictions, particularly during climb, by adapting either the modeled aircraft thrust^{7,8} or weight.⁹ These approaches show promise, but are not yet mature. On the mitigation side, there are many studies that look at using increased separation criteria, or “buffers,” to deal with uncertainty or error. Some examples include studies examining enlarged horizontal detection ranges in the presence of cruise speed errors,¹⁰ wind prediction errors,¹¹ and maneuver-initiation time errors.¹² One common theme through much of this work is a primary focus on horizontal errors and the effectiveness of using horizontal buffers or probabilistic conflict detection schemes to account for those errors. Some simulation test beds, such as the Center-TRACON Automation System (CTAS)¹³ use a vertical buffer for aircraft that are transitioning altitudes, but that buffer is often on the order of hundreds of feet, which is not enough to account for uncertainties during the descent phase.

II. Background

The authors have examined the effectiveness of a buffer to aid automated conflict detection and resolution in the presence of trajectory prediction errors using a geometric conflict detection algorithm, but that work was focused on using buffers in the horizontal plane of up to two nmi.¹⁴ In that work, the authors looked at a range of uncertainties and tried to determine what effects these uncertainties had on missed-alert rates, false-alert rates, and losses of separation (LOS) using automated conflict detection and resolution tools in a non-real-time simulation. Missed alerts are conflicts that should have been detected but are not, and false alerts are conflicts that are predicted to occur but do not. That study found, among other things, that there was a very high rate of missed detections and LOS cases near an aircraft’s top-of-descent (TOD) point. Additionally, increasing the horizontal separation for these cases did not have a noticeable effect on the number of LOS observed. Follow-on, unpublished work revealed the difficulty the conflict probe had in the vertical plane due to the very small legal vertical separation requirement. With many aircraft descending fast enough to completely pass through the legal separation in under 20 seconds, it does not take much uncertainty to create a situation that results in an LOS using standard vertical separation. This danger can be reduced by asking the aircraft to provide intent information about its planned descent, such as the anticipated TOD point or its desired descent profile, but an aircraft might not exactly fly the stated profile anyway. As an example, there can be an error on the order of a few nautical miles between the Flight Management System’s (FMS) predicted and actual TOD point.¹⁵ Additionally, depending on the airspace class, aircraft in level flight can pass over another level aircraft 1,000 feet below and be perfectly legal and safe. It is only when aircraft are transitioning altitudes that detections in the vertical plane become a concern.

One solution is to simply clear all of the airspace beneath an aircraft that might be descending soon. This would ensure safety, but could decrease airspace capacity and increase the total delay experienced by aircraft near either their own or some other aircraft’s TOD point. Therefore, it was decided to design vertical conflict detection buffers that would provide just enough warning to ensure safe separation, while minimizing the amount of extra airspace that would need to be cleared in addition to the amount required by the legal separation standard. The question of “How much warning is enough?” is one that is not totally answered. For this study, the vertical separation buffers have an altitude range that covers roughly four minutes of descent at the nominal predicted descent rate for each aircraft. The exact sizes of the buffers, therefore, vary from flight to flight. Four minutes was chosen as the look-ahead time for this study as a reasonable minimum that should allow a resolution tool to resolve a potential conflict before it becomes an LOS.

III. Simulation Environment

A. ACES and AAC Autoresolver

The simulation test bed used for the study is the Airspace Concepts Evaluation System (ACES).¹⁶ This is a non-real-time simulation that uses a four-degree-of-freedom model to create aircraft trajectories based off

performance data and stored flight plan information. The aircraft performance data are derived from the Base of Aircraft Data (BADA)¹⁷ and the flight plans are created from the filed flight plans for days in the National Airspace System (NAS).

For the current study, the traffic scenario for all the data runs consisted of the flight plans for 9,272 flights across the US, representing about four hours worth of takeoffs during the busiest part of a day in 2005. The wind data is RUC data recorded from a day in May of 2002.

The simulation includes a conflict probe that uses knowledge of those flight plans to check for conflicts along an aircraft's predicted trajectory. Conflict resolution is handled by the Autoresolver,¹⁸ which is a component of the Advanced Airspace Concept¹⁹ (AAC). The simulation only examines aircraft during their flight through Center airspace, from departure fix to arrival fix, and does not examine flights inside terminal airspace. The current version of AAC attempts to resolve conflicts at roughly eight minutes until predicted LOS, referred to as its action time, though if both aircraft are headed to the same arrival fix and within 20 minutes of that fix, AAC can attempt to resolve the conflict at 20 minutes to predicted LOS. Typically, AAC tries to issue a resolution that is free of conflicts for up to four minutes beyond its action time.

B. Trajectory Prediction

In this study, two trajectories are created for each aircraft. The first is the “real” trajectory, which is the one that the aircraft will actually fly. From that trajectory a “perturbed” trajectory is created, which includes the prediction errors being tested (figure 1(a) has an example of this). Every minute, the perturbed trajectory is sent to the conflict detection algorithm and then on to the AAC Autoresolver if a resolution is required.

IV. Experiment Setup

The study consists of two sections. The first section examines the performance of the vertical buffer in terms of conflict detection only. This allows the vertical buffers to be examined in repeatable data sets using multiple error types, with every case having the exact same number of actual conflicts because each aircraft will fly the same “true” trajectory in each run. The second study uses the vertical buffers for both detecting and resolving conflicts. Both studies use the same error types.

The legal separation requirement used for all runs is defined as five nmi of horizontal separation and 1,000 ft of vertical separation. For conflict detection, a six nmi horizontal range is used for all cases. When issuing resolutions, the Autoresolver attempts to obtain seven nmi of horizontal separation. The simulations that were run without a vertical buffer use 1,000 ft vertically for the entire flight, while the cases with vertical buffers use a specialized buffer near TOD and during descent, and 1,000 ft elsewhere. The vertical buffers used in this study are described in detail in the next subsection.

A. Method

Table 1 shows the five types of trajectory uncertainty used in this study. The error ranges were chosen to be roughly in line with values used in other studies, though they were chosen to be slightly larger overall. Cruise speed, descent speed, and TOD location are modeled as uniform distributions around zero. In this simulation, “descent speed” includes an error in the predicted descent Mach number and descent CAS. Wind speed errors are modeled as a prediction that is 25% stronger than the actual wind, as read from a RUC wind file, with the direction unchanged. The aircraft fuel weight is used to adjust the aircraft's weight, and is modeled as a uniform distribution applied to predicted aircraft fuel weights around a nominal value. The references for the error ranges are included in the table, while a more detailed description of how the errors are implemented in ACES can be found in previous work by the authors.¹⁴

Table 1. Trajectory Prediction Error Source and Range.

Error Source	Error Range
Top of Descent Location	roughly +/- 10 nmi; ^{20,21}
Descent Speed	+/-10% ²²
Cruise Speed	+/- 5% ²²
Wind Speed	+ 25% ²³
Aircraft Fuel Weight	+/- 20% ²⁰

The conflict detection portion of this study examined six error configurations and had two vertical buffer settings, for a total of 12 data runs. The test matrix for this portion of the study is shown in table 2. Four of the errors were examined individually (TOD location, descent speed, cruise speed, and wind speed), with a fifth case that had all of the errors, including weight, enabled together. Aircraft weight was not examined on its own because previous work has shown it to have a mild, though non-zero, impact on the aircraft trajectory near TOD.¹⁴ A case with no error was also examined to establish the baseline behavior without trajectory prediction errors. All the error configurations had one simulation run with the enhanced vertical buffers disabled and one with them enabled.

Table 2. Test Matrix: Conflict Detection.

Error Type	Enhanced Vertical Buffer Setting
No Error	Disabled; Full Buffer
Top of Descent Location	Disabled; Full Buffer
Descent Speed (CAS and Mach)	Disabled; Full Buffer
Cruise Speed	Disabled; Full Buffer
Wind Speed	Disabled; Full Buffer
All Errors (including Weight)	Disabled; Full Buffer

The resolution portion of the study consisted of five simulations with all the errors enabled. The test matrix for this part of the study is shown in table 3. The first three runs used the full error range with vertical buffers disabled, set to the full size they were in the detection runs, and set to 80% of full size. The last two runs used 50% of the values shown in table 1, and included a case with no vertical buffer and one with a buffer set to 50% of the full value used in the detection runs. This last pair of runs was used to roughly simulate how effective a vertical buffer combined with improved trajectory predictions would be.

Table 3. Test Matrix: Conflict Detection and Resolution.

Error Type	Buffer Size	Enhanced Vertical Buffer Setting
All Errors	Full Error Range	Disabled; Full Buffer; 80% Buffer
All Errors	50% Error Range	Disabled; 50% Buffer

The vertical conflict detection buffer consists of two parts (see figure 1(a) and figure 1(b)). The first part was a buffer around the predicted top of descent point. This buffer was constructed by assuming the aircraft might descend as much as four minutes early or late. Using the predicted average descent for each aircraft, this buffer is extended along the descent for four minutes and should provide a minimum of three minutes warning.

The second part of the vertical detection buffer is implemented after the aircraft has started to descend and is shown in figure 1(b). This buffer is created at the aircrafts current position during each conflict detection cycle and extends forward along the predicted trajectory. The buffer takes the predicted descent rate at the temporal midpoint of its remaining descent, creates a “fast-descent” and “slow-descent” profile, and extends those four minutes into the future from the aircraft’s current position. The fast-descent profile assumes the aircraft is descending 400 fpm faster than predicted, while the slow descent profile assumes a descent rate 200 fpm slower than predicted. These values were based on results from preliminary data

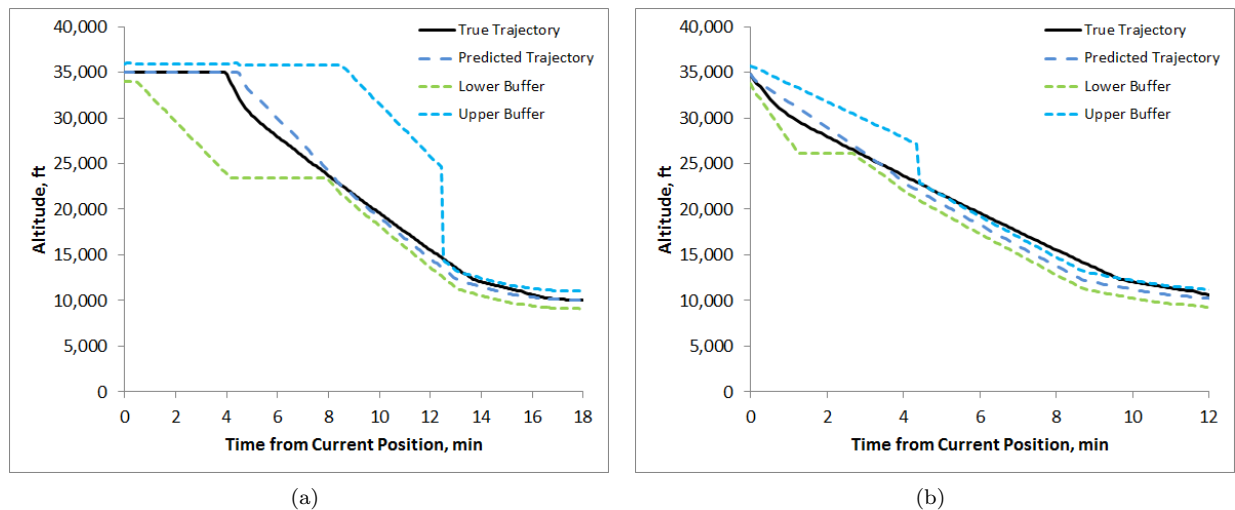


Figure 1. Vertical conflict detection buffer (a) before and (b) after the top-of-descent.

collected for this study. These two buffers comprise the “full” buffer case. The simulations with reduced buffer size simply scaled the early/late descent time and fast/slow climb rate by a percentage value. The look-ahead time was four minutes along the aircrafts predicted descent. Alternative look-ahead times and buffer shapes will be examined in a future study.

B. Conflict Detection Metrics

The main metrics being used in the detection part of the study are the number of missed and false alerts for a specific predicted time until LOS. A missed alert is defined as a case where there is a loss of separation along an aircraft’s true trajectory that is not detected by looking at the perturbed, predicted trajectory. This is recorded by the time until the aircraft would actually have a loss of separation. Depending on the error type, there are generally more missed detections when the aircraft are still 20 minutes apart than when they are closer. In this study, we are most concerned with the missed detections that occur with 3 minutes or less until the time of first loss of separation. These late detections can be very difficult to solve, as neither aircraft has much time to move out of the way. In addition, even though these conflicts do not always result in a loss of separation, they are cases that could more easily become losses, depending on the situation in the surrounding airspace and the capabilities of the person or automation attempting to resolve the conflict.

False alerts are defined similarly to missed, except that false alerts are cases where the perturbed predicted trajectory detects a potential conflict that the true trajectory reveals will not actually occur. Furthermore, for all cases in this study, the legal separation requirement (five nmi horizontally and 1,000 ft vertically) is used to determine whether or not a conflict actually occurred. This means that using any enlarged conflict detection criteria will produce false alerts, even with zero trajectory prediction error. These false alert cases do not impact safety directly, but they can have a large impact on efficiency, as a high rate of these alerts means that many aircraft are being moved to resolve conflicts that would not have actually occurred. This, in turn, adds to the delay for aircraft flying through the area. Therefore, even though some non-zero value should be expected any time an enlarged detection criteria is used, it is desirable to keep this rate as low as possible without degrading safety.

C. Conflict Resolution Metrics

For the portion of the study looking at resolutions, the primary metrics are the number of losses of separation, the number of resolutions issued, and the total delay for aircraft in flight due to conflict resolution maneuvers. The average delay per resolution is also reported.

The LOS metric is the driving one for this study, as it represents failures of the system that could affect safety. It is defined as any case where two aircraft in enroute airspace pass within the legal separation requirement of each other. It should be noted that LOS cases are expected in this simulation, because the

vertical buffers presented here are only a partial solution aimed at significantly reducing or eliminating the number of losses seen in the descent phase of flight. Also, the Autoresolver is only the first level of the multi-layered AAC, so LOS cases in this study could be more accurately described as conflicts that would not be resolved by Autoresolver, and would fall through to the next layer of an overall system. Analysis of those other systems is beyond the scope of the current work, so for simplicity, conflicts that the Autoresolver fails to resolve will be called LOS cases.

The number of resolutions and total delay are ways to quantify the effect of the vertical buffers on system efficiency, as compared to a system with no buffers. Any changes to the system that improve robustness will likely have efficiency penalties, but keeping track of the number of extra resolutions and the amount of extra delay allow for comparisons between options, both now and in future work.

V. Results

A. Conflict Detection

The goal of the first part of the study is to check how many conflicts are detected at least three minutes before predicted loss of separation. As a reminder, a “missed alert” is an alert where the flown “truth” trajectory predicts a loss of separation while the perturbed predicted trajectory does not. A “false alert” is the case where the perturbed predicted trajectory identifies a conflict that would not have occurred, based on the “true” flown trajectory. Also as a reminder, conflicts were only detected, not resolved in this portion of the study, so it is possible for a single conflict to produce multiple false alerts and/or missed alerts.

Figure 2 details the percentage of potential conflicts that were missed conflict alerts, plotted as a function of time until loss for each of the five error cases. The chart on the left uses the standard legal separation vertically for conflict detection, while the chart on the right includes the vertical detection buffers. Both charts use the same horizontal detection range of six nmi, and both charts show all error types, with the case that had no trajectory prediction errors included as a reference. The case without the vertical buffer illustrates the problem caused by these uncertainties. While there are fewer missed detections as the time to loss decreases, many of those missed detections persist until there is very little time to resolve them. This is especially true for the descent speed and top-of-descent position errors which do not really show a decrease in the number of missed alerts until there is less than five minutes until LOS. These two errors help drive the curve for the “all error” case up, so that at four minutes to actual LOS for the all-error case, 8.7% of conflicts are missed by the detection algorithm without the vertical buffer.

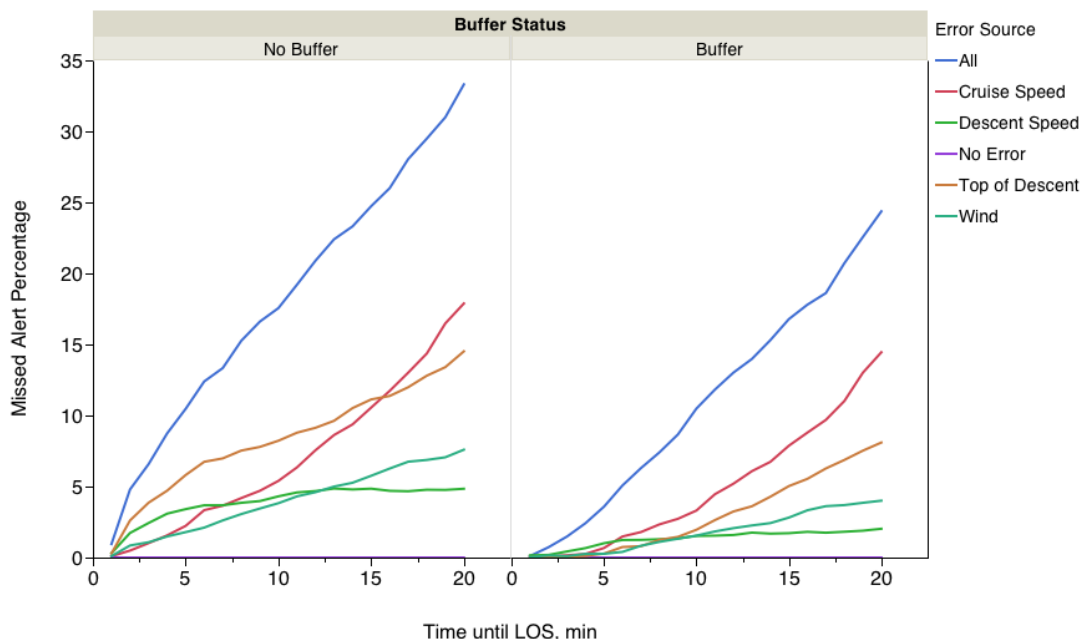


Figure 2. Missed alerts for multiple uncertainties using conflict detection only.

The right chart shows the effectiveness of the current iteration of the vertical detection buffer. It should be noted that these are missed alerts for all phases of flight, so there are some in climb or cruise that vertical detection buffers for descent simply will not address, especially in the case with all the errors combined. There is a noticeable drop in missed alerts for all error types due to the addition of the buffers, especially with less than five minutes until the loss of separation. To continue the example from the previous paragraph, enabling the buffer drops the missed detection value at four minutes to LOS to 2.4% for conflicts in all phases of flight. Considering that the vertical buffer will not directly impact conflicts that do not involve at least one arriving aircraft, that reduction is significant. Looking at an error that directly affects arriving aircraft, the number of missed detections at eight minutes to LOS for TOD error drop from 7.5% without the vertical buffer, to 1.2% with the buffer enabled. At four minutes, there 4.7% of conflicts are missed with TOD error and no vertical buffer and 0.1% missed for the same error with vertical buffers enabled. These results further strengthen the position that conflicts involving aircraft descending into their arrival fix make up a large portion of the conflicts that are difficult to detect, and that a vertical buffer can largely mitigate this.

Figure 3 shows the percentage of alerts that were false alerts for the conflict detection study. This figure shows that adding vertical buffers has a significant effect on false alerts, especially below 10 minutes until predicted LOS. As these were the results with no conflict resolutions implemented, every aircraft flew the same true trajectory in both the left and right charts. Therefore, the overall increase should be entirely attributable to the added vertical detection buffer. To give an example of the scope of the increase, in the case with all errors, at 8 minutes until predicted LOS, the percentage of detections that were false alerts is 30% without the vertical buffers and 53% with the buffers enabled, a jump of 23%. The case without error saw the biggest jump, with the percentage of false alerts moving from 17% without the vertical buffers to 49% with buffers, or a jump of 32%.

The magnitude of these increases implies that there was a lot of traffic around aircraft that were near or past their TOD point. This, in turn, implies that there are many aircraft that could be at risk if an aircraft deviates much from its predicted trajectory near TOD. The net result of a large increase in the number of resolutions issued due to these false alerts is supported by results presented later in the paper. The fact that there are so many more detections also means that it will likely be difficult to find a strategy for mitigating uncertainties in trajectory prediction for this level of trajectory prediction error that does sharply increase the number of false alerts or add significant delay, as there are simply many aircraft in relative proximity during the descent phase for aircraft that are arriving.

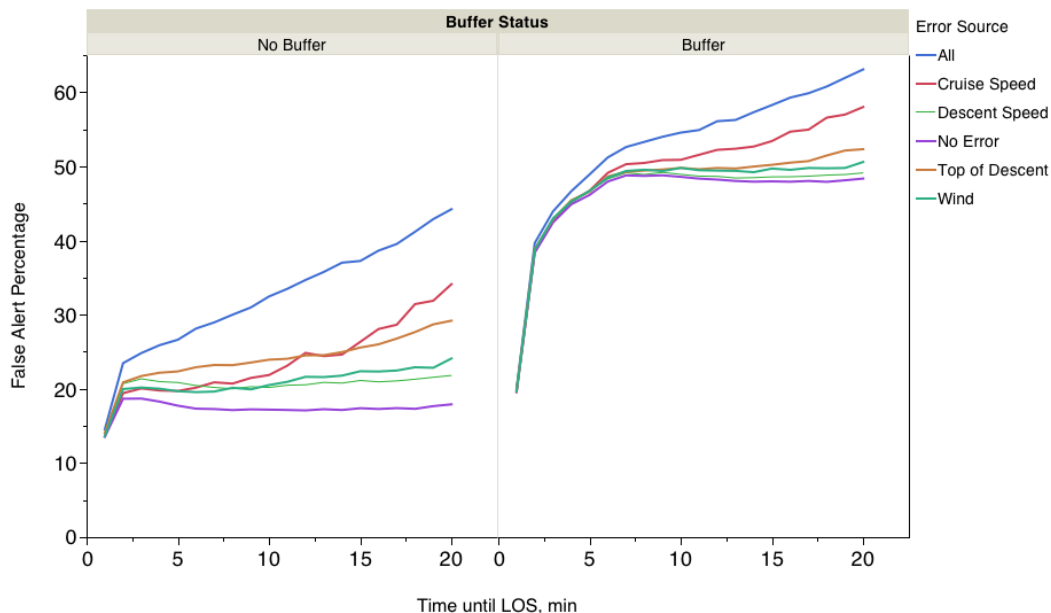


Figure 3. False alerts for multiple uncertainties using conflict detection only.

B. Conflict Resolution

The results with conflict resolutions enabled are discussed in this section. The primary metrics analyzed are the number of losses of separation, the number of resolutions issued, and the total delay and delay per resolution. This section will explore the effectiveness of the vertical buffers and the penalty for using the buffer in terms of the amount of extra delay created and number of resolutions issued.

Figure 4 shows the losses of separation, categorized by flight phases of the two aircraft involved, for the simulations with and without the vertical buffers. All cases detect conflicts at six nmi horizontally, and attempt to obtain at least seven nmi horizontally when issuing resolution maneuvers. The left group of columns are cases where at least one aircraft was climbing, the second group is the case where both aircraft were roughly in their cruise segment, the third group is for cases where one aircraft was an arriving aircraft descending into the Terminal area, and the fourth group is the special case where one aircraft was a descending arrival and the other was a departing aircraft still climbing to cruise altitude. The main point of this plot is to show that simply by using increased vertical separation criteria near and after an aircraft's top of descent point, one can dramatically decrease the number of times uncertainty induces an LOS. However, these buffers alone are not enough to completely remove the problem.

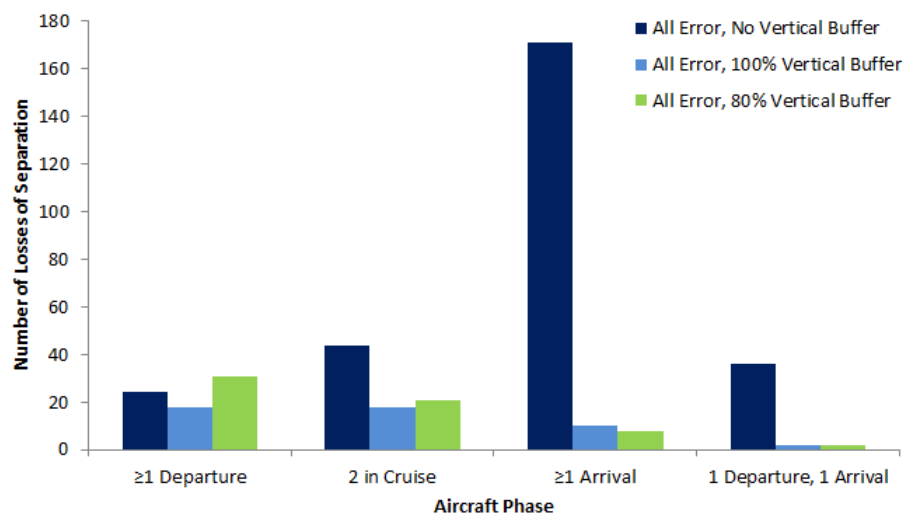


Figure 4. Losses of separation with all uncertainty enabled.

The case without the vertical buffer once again emphasizes the difficulties the conflict detection and resolution algorithms have with descending aircraft, with 209 of the 276 total LOS cases involving at least one aircraft that was descending near its arrival airport, as shown by the right two blue columns. When the buffer is enabled, the number of total LOS cases drops to 49, while the number of losses involving arriving aircraft drops to 12 (right two light blue columns combined). Ideally, the number of LOS for arriving aircraft would be zero with the vertical buffer enabled, but there were a few cases that slipped through. While it should be possible to keep increasing the size of the vertical buffer to cover all of those cases, it is worth investigating a different approach for dealing with those last few losses of separation.

The idea that it might not be efficient to continually increase the vertical buffer to remove all LOS cases came from the results of the 80% buffer run, shown as the green columns in the figure. This scenario actually produced slightly better results than the full buffered case as far as dealing with LOS cases occurring for arriving aircraft, for reasons that are not clear. Unfortunately, due to timing constraints, exploring the causes of these losses of separation involving arriving aircraft that are not resolvable even with the vertical buffers is beyond the scope of this paper. However, this analysis will be done in future work, as accounting for these cases is necessary for making a system that is truly robust to trajectory prediction errors.

One interesting result was the decrease in the number of LOS cases between two aircraft in cruise for both buffer cases. This could simply be the result of these buffers being used for temporary, in cruise descents to avoid conflicts, but further analysis would be required before it could be claimed as a benefit. As the focus of this study is on LOS cases involving an arriving aircraft, the examination of cruise LOS cases will be also

deferred to a future study.

The number of resolutions as well as the delay per resolution are summarized in table 4. It is immediately obvious that the addition of these vertical buffers produces a steep increase in both the number of resolutions as well as the total delay experienced by aircraft in the system, with the number of resolutions issued increasing by 53% and the total delay by 106%. However, one has to consider that the unbuffered case also had a very large number of losses of separation that need to be addressed, so some penalty is likely unavoidable.

The 80% buffer results are more promising, showing an appreciable decrease in the number of resolutions and amount of delay added to the system, increasing the number of resolutions by 42% and the total delay by 69%. The fact that this 80% buffer was able to perform just as well as the full buffer in regards to accounting for losses of separation involving arriving aircraft was a major point. It strengthened the idea that alternative options might be best suited to dealing with the few losses that the vertical detection buffer does not catch, as there seems to be diminishing returns in regards to catching LOS cases when increasing the vertical buffer beyond a certain size. Planned future work includes developing alternative methods for dealing with these last few LOS cases, as well as exploring the use of smaller buffers. This could lead to a system with no LOS involving arriving aircraft and with less system-wide delay caused by resolution maneuvers than would be possible using the enhanced vertical buffers alone.

Table 4. Resolution efficiency metrics by vertical buffer, full uncertainty.

Configuration	Resolutions Issued	Total Delay, min	Average Delay per Resolution, s
All Error; No Buffer	12,574 (Base)	3,844 (Base)	18.3
All Error; Full Buffer	19,180 (Base+53%)	7,912 (Base+106%)	24.8
All Error; 80% Buffer	17,794 (Base+42%)	6,492 (Base+69%)	21.9

This section describes a pair of simulations run with all of the trajectory prediction error ranges set to 50% of their full values, to roughly simulate the effects of improving trajectory prediction accuracy. The first run used no vertical buffer, while the second used a vertical buffer that was also set to 50% of the full value, to take advantage of the reduced error range. As trajectory accuracy improves, one would expect to see fewer losses, a smaller buffer requirement, and more efficiency in terms of the number of resolutions and the delay per resolution, though it is difficult to predict the amount of savings without simply collecting the data. The results of these runs are shown in figure 5, along with the original, full error case with no vertical buffer.




Figure 5. Losses of separation with 50% uncertainty enabled.

It is immediately apparent that, even without the addition of the vertical buffer, cutting the trajectory prediction error in half cuts more than half of the losses off separation in all flight regimes. Furthermore, the addition of a vertical buffer (also half the size of the first one tested in this study), cuts the overall number of LOS cases from 117 in the unbuffered case to 23, and the cases involving arriving aircraft from 87 to 3.

Table 5 shows the results of the half-error cases, with the full-error case (no vertical buffer) provided for reference. Compared to the half-error, no vertical buffer case, implementing the half-sized vertical buffer increased the number of resolutions issued by 42% and increased the total delay due to resolution maneuvers by 82%. These increases are in line with the percentages seen for implementing the vertical buffers in table 4, though the overall numbers are lower because of the reduced number of resolutions and delay in the half-error case without vertical buffers. This illustrates the effectiveness of combining the approaches of improving trajectory prediction accuracy and implementing vertical buffers, as the overall number of LOS cases, the number of added resolutions, and the increase in delay are all significantly reduced when the buffer is implemented in the half-error case.

Table 5. Resolution efficiency metrics by vertical buffer, 50% uncertainty.

Configuration	Resolutions Issued	Total Delay, min	Average Delay per Resolution, s
All Error; No Buffer	12,574 (Full)	3,844 (Full)	18.3
50% Error; No Buffer	9,860 (Base)	2,866 (Base)	17.4
50% Error; 50% Buffer	13,967 (Base+42%)	5,202 (Base+82%)	22.3

VI. Future Work

As stated previously, these vertical buffers are only a first step to making a system that can reliably predict and resolve all conflicts in the presence of uncertainty. The major problem area considered prior to this work was when one or more aircraft were descending towards their arrival fix. However, uncertainties and their resulting trajectory prediction errors can lead to losses of separation in climb and cruise, as well. Additionally, even our relatively large vertical buffers were not sufficient to completely deal with trajectory prediction errors during descent. The next step is to explore ways to remove those last few LOS cases for arriving aircraft. Additionally, more work needs to be done examining different vertical buffer sizes and look-ahead times. Following that, the focus will shift to using buffers and perhaps an adaptive climb algorithm for removing LOS cases during climbs, and then figuring out a way to remove LOS cases in cruise as efficiently as possible. Increasing the range of types of uncertainty is also part of the planned work, with the eventual goal of producing a system that can be made robust to varying levels of trajectory prediction uncertainty in all phases of flight. A secondary goal is to do that while limiting the decrease in system efficiency in terms of delay caused by resolution maneuvers.

VII. Conclusions

The results of this study point to vertical buffers as an effective first step for detecting potential losses of separation involving aircraft descending into terminal airspace in the presence of trajectory prediction errors. Results showed that errors involving predictions during the descent phase are difficult to detect with much more than a few minutes until loss with just a conventional horizontal buffer. Further, implementing vertical separation buffers can be an effective technique for reducing the number of cases that are missed by the detection algorithm, especially for errors in TOD location and descent speed. Results also showed that enabling the vertical buffers decreased the number of missed alerts with four minutes to LOS from 8.7% for all errors without the vertical buffers to 2.4% with buffers. For TOD errors, the missed alerts at 4 minutes until LOS dropped from 4.7% without the vertical buffer to 0.1% with the buffers. Enabling the buffers also increased the percentage of detections that were false alerts by around 20% to 30% for predicted LOS times in the eight minute range for all error types.

When AAC was allowed to resolve detected conflicts, implementing the full vertical buffer produced a significant reduction in the number of losses of separation in the simulation. The number of LOS cases involving arriving aircraft was reduced from 207 in the case with full error and no vertical buffer to 12 when

the full-sized vertical buffer was used. Setting the vertical buffer to 80% of the full size actually reduced the number of LOS for arriving aircraft compared to the full buffer, dropping the number of cases to 10. In terms of the other metrics, the full buffer increased the number of resolutions issued by 53% and the amount of delay accumulated by aircraft executing resolution maneuvers by 106%, while using the 80% buffer increased the number of resolutions by 41% and the delay by 69%. These results imply that vertical buffers can be effective in reducing the number of LOS cases involving arriving aircraft, but there is a point beyond which increases to the buffer size results in larger delay and more resolutions with no real reduction in the number of LOS cases, and alternative methods for catching the remaining LOS cases for arriving aircraft need to be developed.

The study with 50% uncertainty showed that improvements in trajectory prediction significantly improve the ability of the vertical buffer to account for all LOS cases with arriving aircraft. Additionally, the system as a whole runs more efficiently with less resolutions and total delay when there is reduced uncertainty. While this is expected, it suggests that a joint approach of reducing trajectory prediction error and building robust detection schemes is likely the most viable way to achieve a system that has no LOS cases in the presence of multiple trajectory prediction errors.

References

- ¹Mondoloni, S. and Bayraktutar, I., "Impact of Factors, Conditions and Metrics on Trajectory Prediction Accuracy," *6th USA/Europe ATM R&D Seminar*, Baltimore, Maryland, 2005.
- ²Green, S. M., Vivona, R. A., and Grace, M. P., "Field Evaluation of Descent Advisor Trajectory Prediction Accuracy for En-route Clearance Advisories," *AIAA Guidance, Navigation, and Control Conference*, 1998.
- ³Schleicher, D. R., Jones, E., and Dow, D., "Improved Lateral Trajectory Prediction through En Route Air-Ground Data Exchange," *AIAA Aviation Technology, Integration and Operations (ATIO) Forum*, Los Angeles, California, 2002.
- ⁴Bilimoria, K. D., Paglione, M. M., and Lee, H. Q., "Performance Analysis of a Conflict Probe Utilizing Only State Vector Information," *24th International Congress of the Aeronautical Sciences*, 2004.
- ⁵Paglione, M., McDonald, G., Bayraktutar, I., and Bronsvort, J., "Lateral Intent Error's Impact on Aircraft Prediction," *8th USA/Europe Air Traffic Management R&D Seminar*, Napa, California, 2009.
- ⁶Paglione, M. M. and Oaks, R. D., "Implementation and Metrics for a Trajectory Prediction Validation Methodology," *AIAA Guidance, Navigation and Control Conference and Exhibit*, Hilton Head, South Carolina, 2007.
- ⁷Slater, G. L., "Adaptive Improvement of Aircraft Climb Performance for Air Traffic Control Applications," *Proceedings of the IEEE International Symposium on Intelligent Control*, Vancouver, British Columbia, Canada, 2002.
- ⁸Godbole, A. A., *Adaptive Improvement of Climb Performance*, Master's thesis, University of Cincinnati.
- ⁹Schultz, C., Thippavong, D., and Erzberger, H., "Adaptive Trajectory Prediction Algorithm for Climbing Flights," *AIAA Guidance, Navigation, and Control Conference*, No. AIAA 2012-4931, 2012.
- ¹⁰Lauderdale, T. A., "The Effects of Speed Uncertainty on a Separation Assurance Algorithm," *AIAA Aviation Technology, Integration, and Operations Conference*, Fort Worth, Texas, 2010.
- ¹¹Consiglio, M., Hoadley, S., and Allen, B. D., "Estimation of Separation Buffers for Wind-Prediction Error in an Airborne Separation Assistance System," *8th USA/Europe Air Traffic Management R&D Seminar*, Napa, California, 2009.
- ¹²Cone, A., "Effect of Conflict Resolution Maneuver Execution Delay On Losses of Separation," *29th Digital Avionics Systems Conference*, Salt Lake City, Utah, 2010.
- ¹³McNally, D., Mueller, E., Thippavong, D., Paielli, R., Cheng, J., Lee, C., Sahlman, S., and Walton, J., "A Near-Term Concept for Trajectory-Based Operations with Air/Ground Data Link Communication," *27th International Congress of the Aeronautical Sciences*, Nice, France, 2010.
- ¹⁴Lauderdale, T. A., Cone, A., and Bowe, A., "Relative Significance of Trajectory Prediction Errors on an Automated Separation Assurance Algorithm," *9th USA/Europe Air Traffic Management Research and Development Seminar*, Berlin, Germany, 2011.
- ¹⁵Stell, L. L., "Prediction of Top of Descent Location for Idle-thrust Descents," *9th USA/Europe Air Traffic Management Research and Development Seminar*, Berlin, Germany, 2011.
- ¹⁶Meyn, L., Windhorst, R., Roth, K., Drei, D. V., Kubat, G., Manikonda, V., Roney, S., Hunter, G., and Couluris, G., "Build 4 of the Airspace Concepts Evaluation System," *AIAA Modeling and Simulation Technologies Conference and Exhibit*, No. AIAA 2006-6110, 2006.
- ¹⁷Nuic, A., "User Manual for the Base of Aircraft Data (BADA) Revision 3.8," Tech. Rep. 2010-003, EUROCONTROL Experimental Centre, April 2010.
- ¹⁸Erzberger, H., Lauderdale, T. A., and Cheng, Y., "Automated Conflict Resolution, Arrival Management and Weather Avoidance for ATM," *27th International Congress of the Aeronautical Sciences*, Nice, France, 2010.
- ¹⁹Erzberger, H., "Transforming the NAS: The Next Generation Air Traffic Control System," *NASA/TP-2004-212828*, 2004.
- ²⁰Mondoloni, S., Paglione, M., and Green, S., "Trajectory Modeling Accuracy for Air Traffic Management Decision Support Tools," *The 23th Congress of the International Council of the Aeronautical Sciences (ICAS)*, 2002.
- ²¹Stell, L. L., "Predictability of Top of Descent Location for Operational Idle-Thrust Descent," *AIAA Aviation Technology, Integration, and Operations Conference*, Fort Worth, Texas, 2010.

²²Coppenbarger, R. A., Kanning, G., and Salcido, R., "Real-Time Data Link of Aircraft Parameters to the Center-TRACON Automation System (CTAS)," *4th USA/Europe ATM R&D Seminar*, Santa Fe, New Mexico, 2001.

²³Schwartz, B. E., Benjamin, S. G., Green, S. M., and Jardin, M. R., "Accuracy of RUC-1 and RUC-2 Wind and Aircraft Trajectory Forecasts by Comparison with ACARS Observations," *Weather and Forecasting*, Vol. 15, No. 3, 2000, pp. 313–326.