Improvements to Conflict Alert to Account for Discretionary Descents

Bill Holtzman
Ken Leiden
Mosaic ATM

January 2020

Background

One of the goals of the Separation Automation System Engineering (SASE) project is to address performance issues with the ERAM Conflict Alert (CA) tool. CA initiates flashing data blocks on the ERAM display to let the controller know when two aircraft are predicted to lose separation within the CA lookahead time (typically 2.5 minutes). CA uses position data along with derived ground speed and vertical rate to predict the trajectory of the aircraft. The ground speed and vertical rate are assumed to be constant over the CA lookahead time (i.e., dead reckoning). One exception to the dead reckoning assumption exists in the vertical dimension when CA models a level-off associated with assigned or interim[1] altitude clearances. Figures 1 and 2 depict situations without and with the level-off modeling, respectively. In Figure 1, UAL2 is climbing into the path of DAL1 at FL220 and an alert is given. In Figure 2, CA models a level off for UAL2 corresponding to its assigned altitude of FL210 and thus no alert is given. In both of these, the trajectory modeling and the results are appropriate.

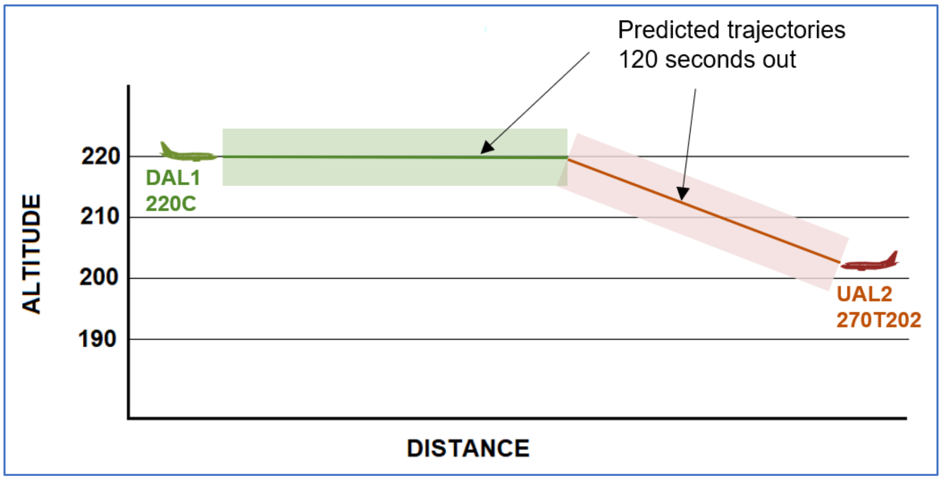

Figure 1: Alert - CA provides alert for aircraft in conflict

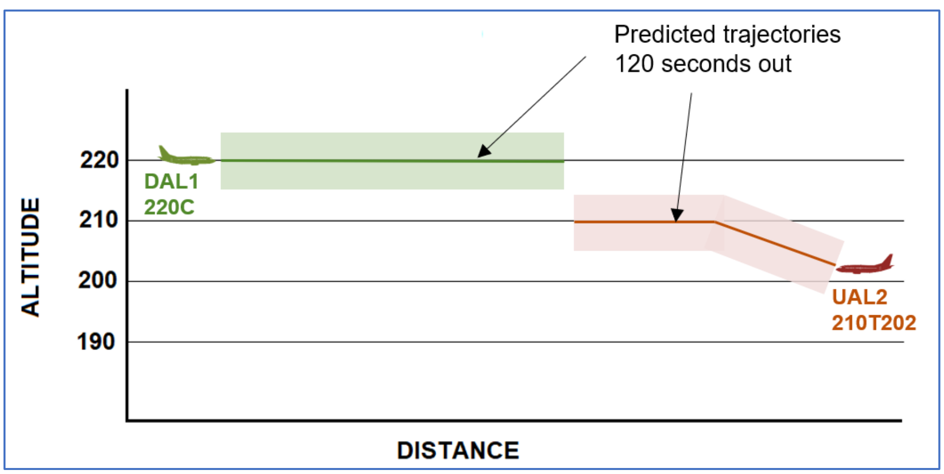

Figure 2: No alert - CA models a level off due to assigned altitude and no conflict is detected

---

[1] Today's Procedure Altitude commands (QQ P) may include an interim altitude as well, but CA only models that altitude and none of the constraints within the associated procedure.

Discussion

Current CA for VNAV Procedures

CA accuracy has been a major concern since the introduction of VNAV procedures into the NAS. These procedures include one or more vertical restrictions along the route such that the altitude entered by the controller into ERAM does not often reflect the aircraft's target altitude. Consider the PHLBO3 arrival into Newark in Figure 3.

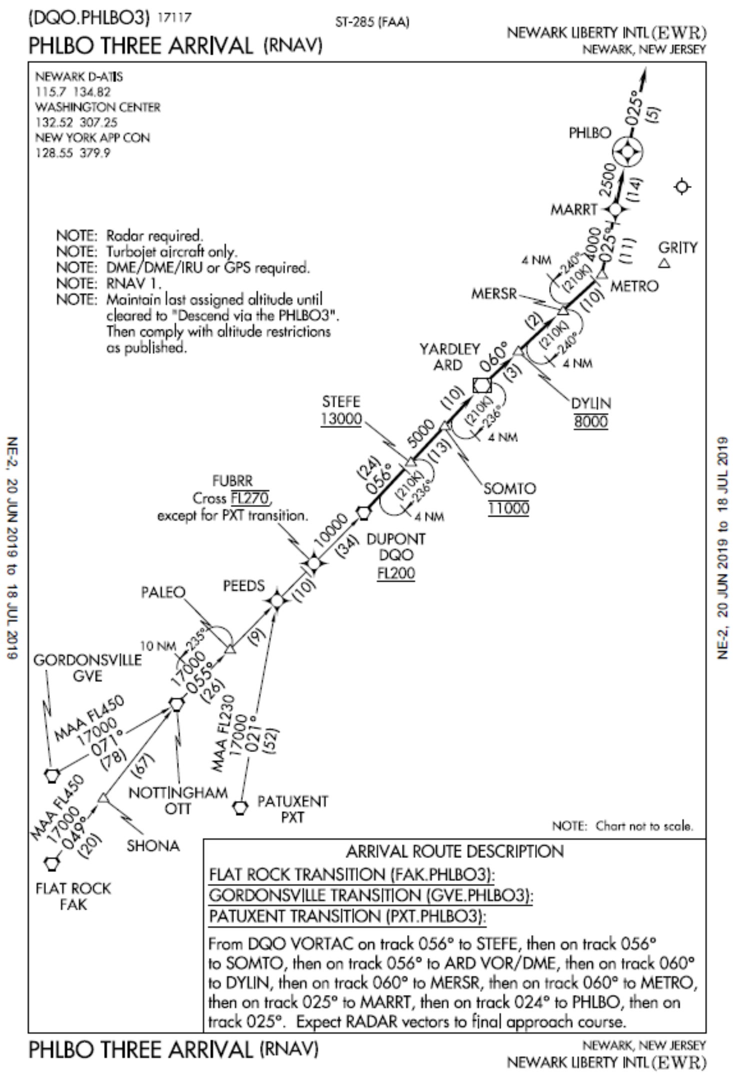

Figure 3: PHLBO3 arrival

The controller may clear an arrival to Descend Via the PHLBO3 prior to FUBRR. The pilot then must make all downstream vertical restrictions, including:

- FUBRR at FL 270
- DQO at or above FL 200
- STEFE at or above 13,000
- SOMTO at 11,000
- DYLIN at 8000

But the controller is not inclined nor capable of repeatedly updating the ERAM interim altitude every time the aircraft passes one of these fixes. The controller typically makes a single entry when issuing the Descend Via clearance which looks like:

QQ P080 DAL1

This tells ERAM that the Descend Via has been issued for DAL1, and that DAL1 can be expected to fly the vertical profile of the PHLBO3 to maintain 8000. However, CA does not currently process this information and does not model the aircraft to remain within the vertical profile of the procedure.

Note: Part of the SASE ERAM Enhancements 3 (EE3) include modeling of trajectory compliance with VNAV restrictions for the purpose of improving Conflict Probe performance. CA is a separate algorithm and is therefore not affected by these enhancements. However, changes to CA proposed below will be able to leverage the processing of restriction information by ERAM.

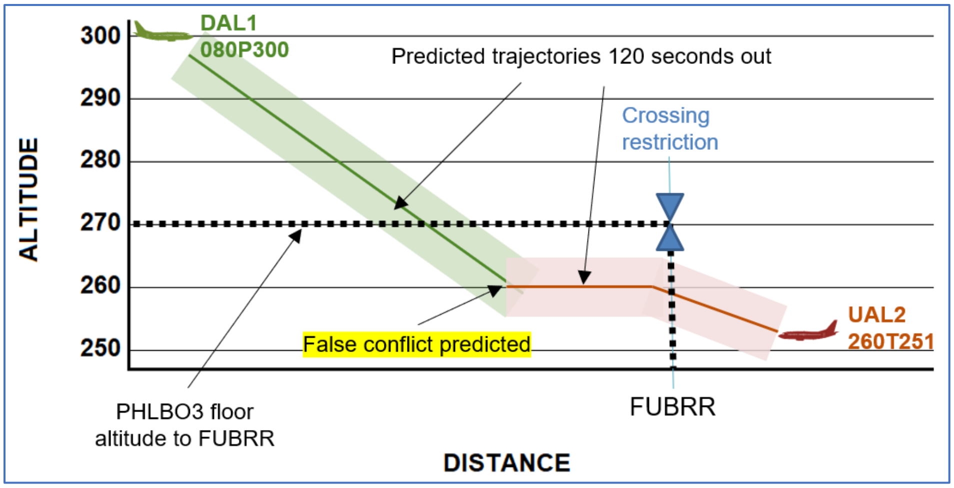

Figure 4: VNAV procedure false alert

An example of CA with the PHLBO3 restriction at FUBRR is illustrated in Figure 4. Unaware of the restriction, CA models a *clearance deviation* for the DAL1 trajectory by continuing its current vertical rate through the FL 270 floor prior to FUBRR. This results in a false alert with traffic at FL 260. Such alerts often cause confusion; they are a significant distraction that can compromise safety. Over time controllers become accustomed to them, and they dilute the value of CA as controllers ignore the alerts.

Previously Proposed CA Enhancements for VNAV Procedures under EE3

One of the enhancements[2] proposed is to apply a level off when the flight is predicted to descend below an "at" or "at-or-above" altitude constraint before crossing the associated fix. This modeling change is similar to the level off illustrated in Figure 2, which is triggered by the presence of an assigned or interim altitude. However, in this case, the trigger is a VNAV procedure altitude constraint. This is illustrated in Figure 5 where a conflict with climbing traffic is identified. Note that this situation would have been a late alert under current CA. Furthermore, this enhancement eliminates the false alert in Figure 4.

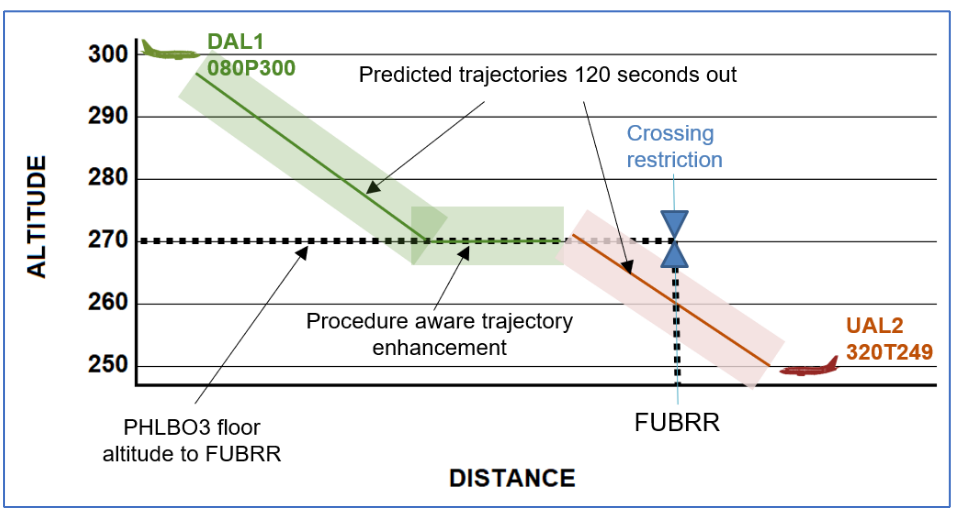

Figure 5: Proposed suppression of VNAV procedure false alerts

[2] Referred to as Candidate Enhancement 2a in the Leidos report [1].

A different dynamic exists should DAL1 opt to remain level at FL 300 for as long as possible, which is typical of VNAV procedure descents. In these cases, CA will model level flight until a descent is detected by position reports. This is depicted in Figure 6. Note that DAL1 is still level, and that the shaded green area represents airspace DAL1 could use in the descent to make FUBRR at FL 270.

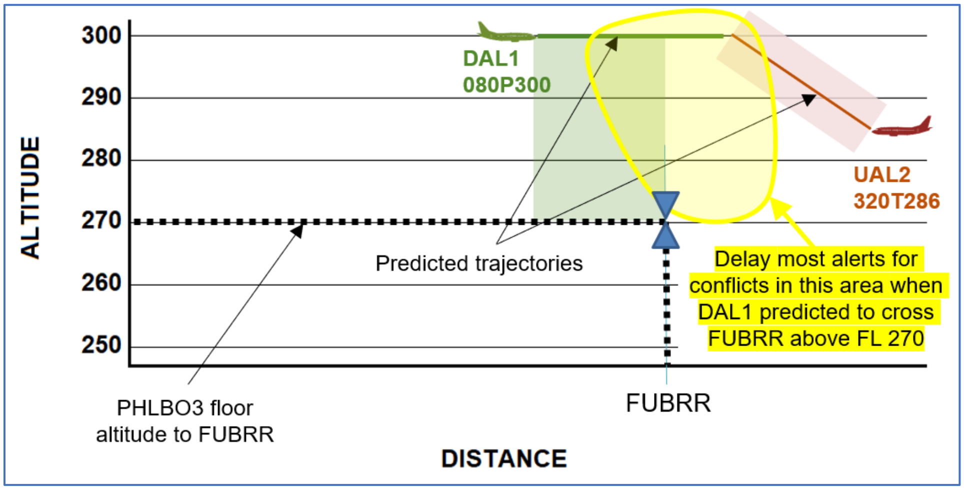

Figure 6: CA modeling approaching FL 270 crossing restriction @ FUBRR

The CA proposed enhancement[3] for this situation is to briefly delay conflict notification when a flight is predicted to cross above a constraining altitude under specific circumstances. The reasoning for this enhancement is that there is an expectation that the aircraft would begin descending (or increase its descent rate if already descending) in the very near-term timeframe. Thus, delay in notification will allow time for the aircraft to come into compliance with its restriction. For procedurally-separated flights, this will mitigate false alerts.

The enhancement includes an exception for cases when the required descent rate to make the crossing restriction becomes unreasonable – for instance should DAL1 enter the yellow area in Figure 6. The exception provides an alert to the controller without delay. This is similar to logic CA employs to identify scenarios in which the vertical rate of a climbing or descending aircraft suggests the aircraft will not level off at the assigned altitude.

Discretion-based CA Background

Note that the proposed enhancement shown in Figure 6 above does not involve a change to the modeled trajectory, only a delay to the alert notification. But the modeled trajectory – maintaining level flight over the lookahead time – deviates from the DAL1 clearance and is therefore not realistic. Implicit in this realization is that there may be a missed or late alert as a result. Figure 7 depicts such a situation with level traffic at FL 270. Even though ERAM has adaptation data on the PHLBO3 arrival and the controller has entered a message to indicate that DAL1 has been issued the Descend Via clearance, CA will not model the descent until position reports indicate the descent has begun. Any controller can quickly recognize that this situation is almost certainly going to be a serious conflict, but CA lacks that ability. This is a clear shortfall in both the current CA capability and those proposed for EE3, and it seriously undermines efforts to encourage controllers to rely more heavily on conflict detection tools.

---

[3] Referred to as Candidate Enhancement 1 in the Leidos report [1].

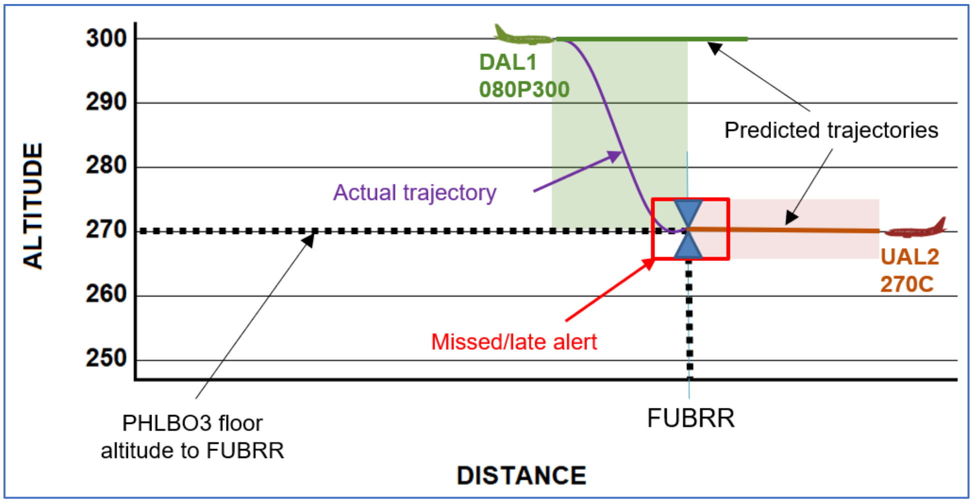

Figure 7: Today's CA failure to identify obvious conflict

A dangerous aspect about the situation in Figure 7 is that CA will eventually provide an alert, but only after DAL1 has fully nosed over and position reports (with their inherent time lag) indicate the descent rate has increased sufficiently to bring DAL1 into conflict with UAL2. At this point, the controller is dealing with a very near-term conflict with aircraft that are on a head-on collision course. The controller has to instantly determine the best resolution of the conflict – go over or under[4] - and to deliver and confirm readback of the appropriate clearances. The unfortunate presence of a stuck mic, radio interference, pilot request or any other occupation of the radio frequency could easily deprive the controller of any opportunity to rescue the situation.

Without prompt response by the controller and pilots or even with it, the aircraft TCAS will likely provide Resolution Advisories (RAs) which could lead to injured crew and/or passengers. The incident will likely disrupt sector operations and cause general malaise in the operation.

By contrast, if CA provided an alert prior to the start of the discretionary descent, the controller could casually amend the Descend Via clearance issued to DAL1 without incident. Where today's response involves crisis vectors, distraction, and a potential safety incident, a *pre-descent* alerting capability enables an even, measured and orderly response that maintains the sector rhythm and is hardly noticeable to any participants but the most attentive.

With a nominal CA lookahead time of 2.5 minutes, the typical aircraft starting an unpredicted discretionary descent can come into conflict with aircraft as much as 9000 ft below it. However, the more critical safety issue is for aircraft within 4000 ft of the discretionary descent aircraft because the

[4] To achieve vertical separation, the controller can either try to stop the descent and climb back up DAL1 while descending UAL2, or "swap out" the aircraft, descending expeditiously DAL1 through FL 270 and climbing UAL2. This is a complex decision requiring a dead reckoning evaluation using the vertical rate of the descending aircraft, speeds, range to the conflict and winds. But the decision also requires consideration of aircraft type, likelihood of expeditious compliance, presence of other traffic, turbulence and frequency congestion. In addition, the vertical rate of the DAL1 as indicated to the controller may lag the actual rate. A bad decision or sluggish execution can lead to a serious incident.

alert notification time can be 75 seconds or less, bordering on the time that TCAS involvement becomes an issue.

Note that the situation illustrated in Figure 7 is a bit contrived to make it easy to view in a two-dimensional graphic. Realistically, the problem would more likely involve an aircraft crossing the arrival at some angle near or downstream of FUBRR.

However, with weather deviations present the head-on situation does become a real possibility. A departure from the airport to which DAL1 is inbound could deviate from the departure sector into the arrival sector, straying directly into the face of DAL1. While a departure is more likely to be climbing, its trajectory could take it over FUBRR at FL 270 and thus be direct traffic for DAL1.

The lack of precision conflict detection tools is one of the reasons arrival sectors tend to shut down during weather rather than continue limited operations.

The scenario above only reveals one aspect of a more fundamental problem with CA. This shortfall involves not just modeling a procedure descent, but modeling any descent involving a pilot's discretion clearance. There are three types of clearances that give the pilot discretion in the vertical dimension:

- Explicit clearance — example: "Descend at pilot's discretion"
- Crossing restriction — example: "Cross FUBRR at FL 270"
- Procedure descent — example: "Descent Via the PHLBO3 arrival"

In each case, the pilot can start a descent whenever they like. They may level off at any altitude but cannot climb at any point. For a crossing restriction and a procedure descent, they must reach specified altitudes before passing the corresponding fix or fixes.

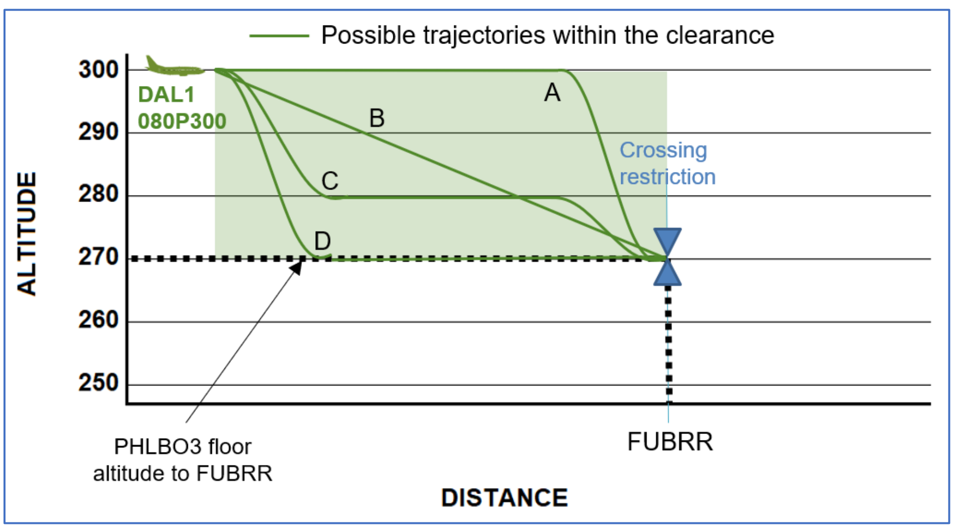

Figure 8: Possible trajectories for a PHLBO3 Descend Via clearance

The complication of a discretionary descent for trajectory and conflict modeling is that CA cannot reliably predict the descent path, only the end points of the descent. This is illustrated in Figure 8. From present position at FL 300, DAL1 has complete discretion to fly any of the trajectories shown or any other trajectory within the shaded green airspace, as long as DAL1 does not climb and as long as DAL1 crosses FUBRR at FL 270.

The reason Optimized Profile Descents exist is to allow aircraft automation to determine (and the pilot to accept) the Top-of-Descent point such that the flight path is optimal based on winds, loading and other performance considerations. But even when the pilot has set the FMS to manage the descent, they may choose to over-ride the planned profile due to turbulence, abruptly swapping the aircraft from the trajectory A to trajectory D in Figure 8.

As previously mentioned, CA trajectory predictions are strictly based on position reports, dead reckoning and assigned or interim altitudes entered into ERAM. These predictions will only include a small portion of the green-shaded airspace shown in Figure 8. Due to the uncertainty implicit in discretionary descents, CA is not able to identify conflicts with some aircraft that might appear within that airspace. Consider Figure 9, in which DAL1 has clearance (either by VNAV procedure or with a crossing restriction) to cross MYFIX in the altitude window shown.

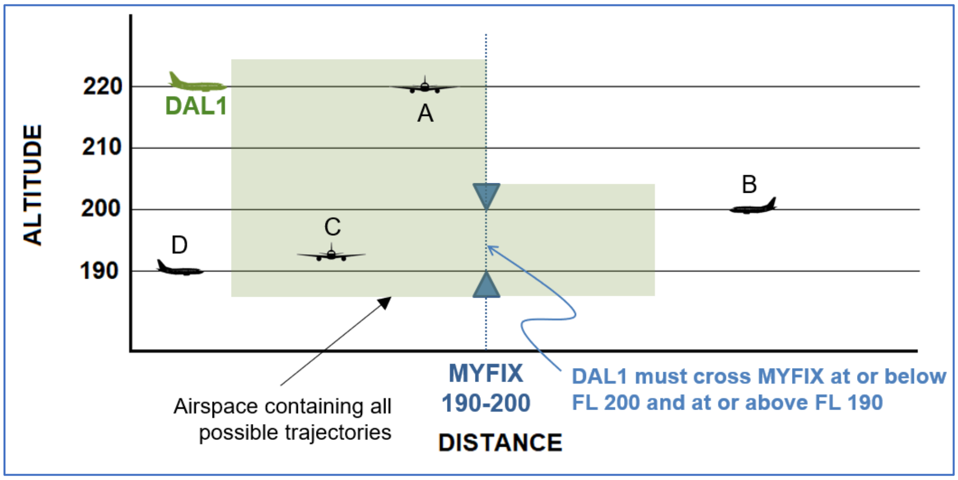

Figure 9: Potential traffic for DAL1, with restriction to cross MYFIX between FL 190 and 200

Suppose that aircraft A and C will arrive in the vertical plane of DAL1 such that they could cause a conflict with DAL1 depending on when DAL1 begins its descent. Also suppose that aircraft D and B will cross MYFIX at the same time as DAL1, and that all four example aircraft will remain level. Until DAL1 begins a descent, today's CA will only alert the conflict between DAL1 and aircraft A. Importantly, all four of these aircraft should be alerted because CA has no reliable way to know the trajectory DAL1 will fly. Such alerts may not rise to quite the same level of concern as one generated by dead reckoning considerations, and for that reason perhaps an alert of this type would have a different presentation.

The identification of potential conflicts with the four aircraft in Figure 9 represents a long-lived, traditional and essential concept in air traffic control, the assignment of airspace to more than one aircraft. This concept is at the heart of non-radar ATC, in which two aircraft that are scheduled to cross the same fix at the same altitude within 10 minutes of each other are considered a conflict.

Figure 10 shows a non-radar conflict as was at one time identified using paper strips and colored pens. A controller who did not identify this conflict would have committed an operational error. It was the responsibility of every controller to never assign the same airspace (fix and altitude) to more than one aircraft at a time. As long as this responsibility was fulfilled, the only way two aircraft could lose separation is if there was a deviation from a clearance.

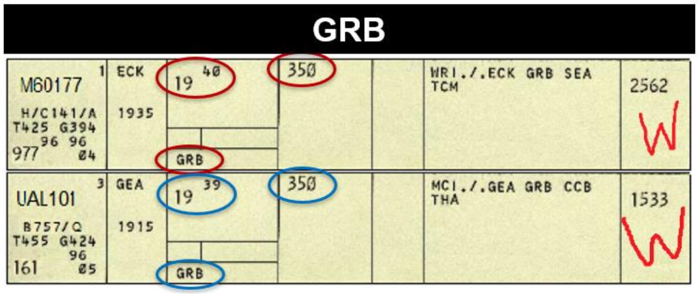

Figure 10: Traditional non-radar conflict

The same concept should be applied to the use of complicated clearances such as VNAV procedures and other discretion-based clearances. The assigned airspace for a VNAV procedure involving multiple windows is not simple to visualize and evaluate in real time. That is an opportunity for automation, which could certainly identify these conflicts between assigned airspace in real time. This is critical to a NAS in which 80,000 flights are handled every day and in which there is a continuous need and desire to push more and more aircraft through VNAV descent procedures.

Discretion-based CA Capability Description
The new capability would:

- Recognize the uncertainty in trajectory predictions of aircraft on discretionary descents.
- Provide alerting with any traffic for whom a loss of separation can be predicted (with realistic speed and vertical rate assumptions) for aircraft complying with their clearances.
- Provide the visual alert in a form distinct from the current CA.
- Require the controller to make an entry into ERAM when crossing restrictions and explicit pilot's discretion descents have been issued, as well as Climb Via clearances.

Concluding Remarks

False alerts:
Since CA was introduced in the 1980's, it has always had certain limitations that have been accepted without question. Since VNAV descent procedures were introduced, these limitations have become much more noticeable and of much greater concern. The commonly viewed response to the issue has been that the use of procedure descents generates false alerts and these simply need to be inhibited.

But a deeper review is revealing. False alerts are a symptom of trajectory modeling that uses incomplete data. ERAM cannot predict the trajectory simply because it doesn't have intent from the flight. Delaying alerts is helpful, but is something of a band-aid until intent can be captured. The issue has not been well fleshed out because controllers aren't really "woke" to potential improvements in CA.

The lesson here is simple: As long as trajectories are treated as single track predictions, conflict probing is a zero-sum game. If there are false alerts being generated, CA is probing in the wrong place and will not identify valid conflicts elsewhere.

Discretionary Descents:
The trajectory prediction in CA is single track; at all times there is either a single straight-line prediction, or a line with one kink for a level off. Along with a lack of intent data, this limited dimensionality is the underlying reason why CA is not equipped to detect conflicts for aircraft with discretionary descents, whether those related to VNAV procedures, crossing restrictions or explicit discretionary clearances. These are also the reasons CA does not provide effective trajectory modeling in the lateral dimension.

With a discretionary descent the controller needs to know not just the conflicts CA sees with a single-trajectory option, but any conflict possible within the bounds of the clearance. This is because the actual trajectory chosen cannot be reliably predicted (due to lack of pilot/aircraft automation intent). The concept is the same in non-radar separation, a core paradigm of air traffic control: No two aircraft are assigned the same airspace at the same time. With today's complex VNAV procedures involving multiple descent windows, controllers are expected to perform the same complex trajectory predictions in their heads that current automation is incapable of, and at the same time endure the distraction of false alerts generated by that automation.

The lack of a capability within CA to model trajectories that accurately bound the range of vertical profiles is a major shortfall that diminishes the value of CA and inhibits controllers from increased reliance on CA. Poor performance of CA is one of the reasons that arrival sectors tend to shut down traffic during weather events. Situations in which aircraft are deviating for weather or otherwise off their route are actually situations when CA could really shine. A single stray aircraft is not always easy to identify when multiple aircraft are deviating, the radio frequency is choked and other distractions exist. When that same aircraft is issued a complicated Descend Via clearance, controllers are loath to guarantee they'll spot every conflict every time and will therefore tend to issue flow restrictions or shut off traffic.

Improvement of CA is critical to support the future Trajectory-Based Operations environment.

References

[1] Separation Assurance System Engineering (SASE) DS#10 - PBN Operations Algorithm – Trajectory Modeling Report, under FAA Contract DTFAWA-16-D-00005, TO00021/TORP 5106, October 31, 2018