# Redistribution of Necessary Delay in the US National Airspace System: Benefits from Trajectory-based Operations

**Almira Ramadani, Martin Durbin, Brian Kravitz, and Dave Knorr**
**Federal Aviation Administration, USA**

## ABSTRACT

"Necessary delay" is the airborne delay imposed on aircraft for maintaining safety and maximizing throughput under unpredictable changes in operating conditions and arrival times. Through experimentations with three years of surveillance data inputs (2016-2018) for 41 key airports in the Continental US, the FAA has developed a standardized methodology for evaluating necessary delay that is currently absorbed in the National Airspace System (NAS). This methodology focuses on aircraft delays during periods when demand exceeds capacity in terminal areas around the NAS, and considers external, non-Air Traffic Management (ATM) factors that also contribute to the necessary delays, including occurrence of convective weather en route, airport meteorological conditions, and equipment outages. The methodology has been used to estimate benefit opportunities from Trajectory-based Operations (TBO) by determining the amount of delay and fuel burn that is currently absorbed in terminal areas, but could be redistributed to an upstream and more cost efficient phase of flight by using TBO tools. In 2018, about 11 per cent of arrivals to the 41 key airports in the US could have derived benefit from delay redistribution, with average fuel savings between 40kg and 245kg per flight. Notably, most delay redistribution is manageable with speed control adjustments. Only three per cent of flights require either en route lateral extensions or ground delay before take-off.

This analysis presents an important example of environmental benefit opportunities from ATM improvements. While smaller in magnitude compared to those from improved technology and alternative fuels, ATM improvements can yield significant environmental improvements as well. The initial investigation focuses on benefits opportunities that are possible under the same efficiency of using the existing airport capacities, and excludes analysis of benefit opportunities during periods with convective weather and better routing options. The FAA will continue to work on both advancing the TBO concept and refining the assessment of the corresponding benefit opportunities.

## Introduction

Stakeholders of the US Federal Aviation Administration (FAA), including airlines, US Congress, and the Office of Management and Budget, continue to demand realistic benefit estimates associated with new programs, technologies, and procedures. Enhancements in underlying historical aircraft trajectory data and analysis techniques now make it possible to generate much more realistic estimates of efficiency improvement opportunities in all phases of flight (see Figure 1).

This paper presents a new methodology for empirically estimating necessary delay. The methodology was used to evaluate such delay for arrivals to 41 key airports in the NAS, determine the amount of the delay that could potentially be redistributed by Trajectory-based Operations (TBO), and estimate the fuel savings associated with this more efficient delay absorption.

Necessary delay is the airborne delay that needs to be absorbed by aircraft to maintain safety and maximize throughput under unpredictable changes in schedules and operational conditions. It may not be possible to reduce the necessary delay in magnitude or in cost through improvements in NAS operational efficiency. TBO's potential impact in variable conditions such as convective weather, winds driving a configuration change at the airport, or unpredictable changes in visibility are not addressed with the analytical method described here. This method focuses on underlying inefficiencies caused by demand for services that exceed capacity of the NAS resources, and improvements enabled by air traffic management (ATM) solutions This paper also discusses some key considerations that need to be addressed when translating efficiency improvement opportunities into benefits, such as the contribution of non-ATM factors, including air traffic demand and performance in convective weather.

As a NextGen ATM method for strategically planning and managing flights in the NAS, TBO builds on advanced time-based management (TBM) tools and performance based navigation (PBN) infrastructure, integrates decision-making across domains and systems, and optimizes delivery of aircraft into terminal areas. With more accurate delivery of aircraft into terminal airspace, TBO pushes delays that are currently absorbed in low altitude airspace during busy periods further back, resulting in delays of the same magnitude but lower cost. With new trajectory management tools both in the terminal and en route control facilities, TBO maintains the same level of runway throughput and moves the necessary delay to en route airspace.

Under TBO, decision support tools help controllers with managing converging and diverging aircraft flows through control points. Ground Interval Management – (GIM-S) aids with improving accuracy of meeting scheduled arrival times to the terminal area boundary by suggesting speed adjustments between the extended freeze horizon and arrival meter fix. When speed advisories alone prove insufficient, the Path Stretch tool provides lateral path extensions or shortenings that may be needed for aircraft to meet metering schedules. In addition, En Route Departure Capability (EDC) aids with reserving a spot in an en route flow to a constrained destination, while Integrated Departure and Arrival Capability (IDAC) helps with integration of departures into the overhead flows above departure airports. All of these tools alleviate the need for vectoring in terminal airspace, resulting in a reduction of overall fuel burn while achieving overall equivalent flight times.

**FIGURE 1:** Phases of Flight

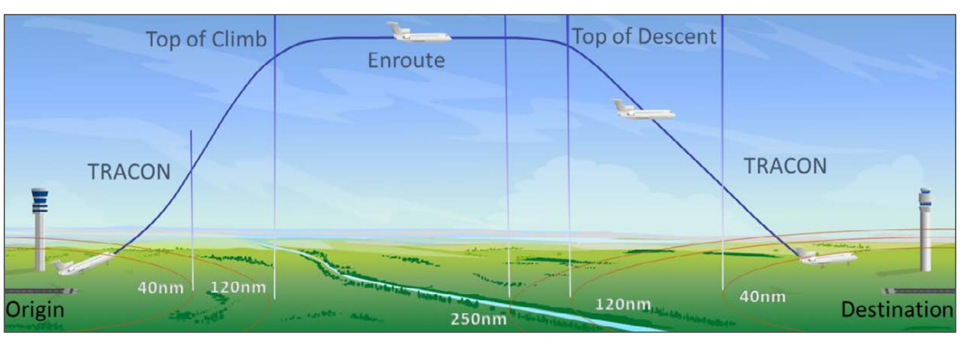

FIGURE 2: Distribution of Arrivals by Great-circle Distance between Origin and Destination

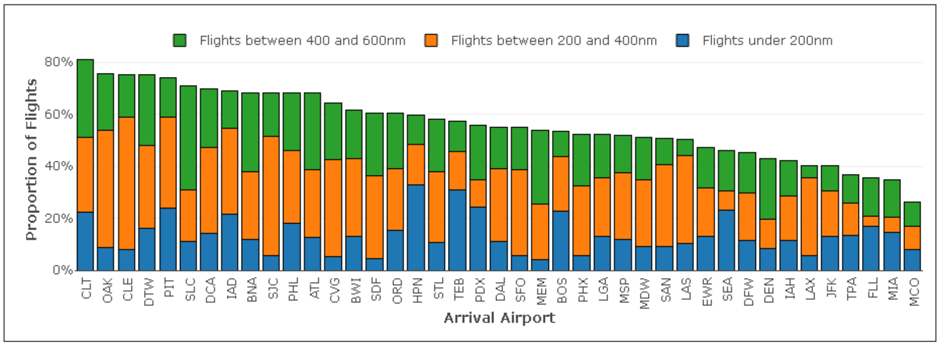


Initial Trajectory-Based Operations (iTBO) is the first phase of TBO that focuses on site-specific deliveries of new capabilities between 2018 and 2023, and on integration of new and legacy capabilities to achieve TBO objectives. Full TBO will complete the development and deployment of TBO capabilities through the end of 2025, and is applicable to all phases of flight. Dynamic TBO includes capabilities that are in the research and development phase that would provide flight-specific TBO capabilities through the end of 2030, which may enable dynamic optimization.

This research touches upon work in four main areas: surveillance data cleaning, flight efficiency, delay redistribution, and fuel efficiency.

## Surveillance Data Cleaning

A significant number of runway assignments in the data were either missing or did not align well with actual flight trajectories when plotted. Runway assignments are an important part of our methodology for assessing flight inefficiency; therefore, it was necessary to develop an algorithm for accurate runway assignments. As in Szurgyi[1], arrival runways were determined by comparing the final flight coordinate to a radius around the airport. Classification was performed by extending the runway centerline to the radius around the airport and determining if the flight was within a certain tolerance. At airports with crossing runways, the radial crossing points were sometimes very close, resulting in inaccurate runway assignments. To resolve this, the radius was gradually increased until the radial points from non-parallel runways were far enough apart to allow accurate determination of the runway assignments.

## Flight Efficiency

In order to redistribute necessary delay to a more cost efficient phase of flight, the amount of necessary delay at each airport first had to be identified. Over the years, many researchers, including CANSO[2], Kettunen, et al.[3], Knorr, et al.[4], and Gouldey[5], have studied this topic. This study is similar to Gouldey[5], but the approach to defining flows differs slightly, including: corner posts are defined by clustering flight tracks, parallel runway groups are used, and aircraft are not grouped by category. The definition of necessary delay that was used in this study is based on distance flown above the 15th percentile for each flow, rather than 105% of the median track distance (as in Gouldey[5]). Knorr's[4] approach was used to extend the definition of necessary delay, and also to assess vertical inefficiency by determining necessary level-flight, or the amount of level distance that exceeds the 15th percentile of distance in level flight for like-flights.

## Delay Redistribution

Using speed control en route to achieve fuel savings has been previously studied by Jones et al.[7], Knorr et al.[4], and Xu[8]. In this study, necessary delay, and necessary delay in level-flight, were both used to determine how much horizontal and vertical inefficiency can be moved from the terminal area altitudes to cruise flight by reducing cruise speed while maintaining a more efficient fuel burn rate at the higher altitudes. When combined with reduced speed, the fuel benefit is realized by both removing the low level excess flight distance/time and saving fuel by flying a little slower en route.

### Fuel Efficiency

Airlines use Cost Index (CI) – a ratio between the unit cost of time and the unit cost of fuel – to optimize the speed of an aircraft. There are two theoretical speed options for the cruise phase of flight, maximum-range cruise (MRC) speed and the long-range cruise (LRC) speed. With the CI of zero, the MRC speed provides the farthest distance and aircraft can reach with a given amount of fuel. The LRC speed is typically 3-5 per cent higher, and requires about 1 per cent higher fuel consumption (see Figure 3). Airlines typically fly at speeds faster than that of the CI zero; while this may be inefficient in terms of fuel burn, business decisions sometimes require prioritization of time over the cost of fuel. Research into fuel efficiency, such as Folse[6], helped determine limits to speed reduction in cruise to avoid unrealistic increases in fuel burn due to slower cruise speeds when absorbing redistributed time.

## STANDARDIZED METHODOLOGY

**FIGURE 3:** The Relationship between MRC and LRC Speeds

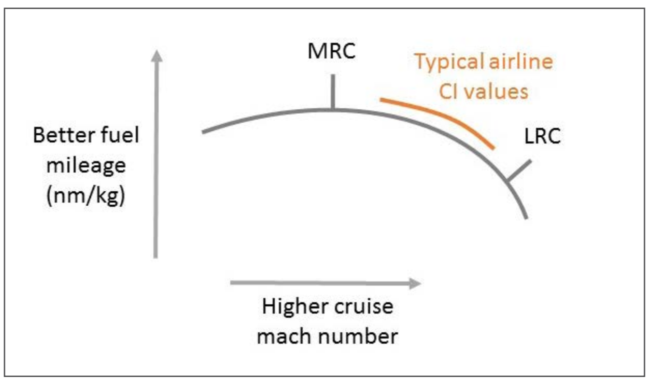


## FOR ESTIMATING EFFICIENCY IMPROVEMENT OPPORTUNITIES FOR ARRIVALS

Lateral and vertical efficiency improvement opportunities exist in cases when a flight's actual distance, time and altitude can be better aligned with those of its optimal ground track and vertical profiles. However, efficiency of a flight between the same origin and destination airports can vary greatly with operating conditions including meteorological conditions, en route weather, airspace closures, route availability, and traffic demand levels. Even when evaluated over just a limited segment such as within a terminal area, flight efficiency of aircraft arriving via the same corner post may still greatly vary with demand level runway configuration at the destination at the time of arrival, as well as with demand levels and runway configuration at the nearby-airports.

As a result, optimal trajectory is not a static construct for flights between the same origin and destination, but a variable one and highly dependent on many ATM and non-ATM factors. Moreover, while optimal distance, time and altitude may be known to the aircraft operator or the service provider at the time of operation, they are not recorded in empirical data archives, and need to be estimated for post-operational assessments. Therefore, the study methodology started by investigating flight parameters, applicable operating conditions, and geometries of flown trajectories, to properly categorize aircraft into groups of like-flights. Optimum distance and distance in level-flight are then estimated for each group of like-flights as an *achievable* distance, and distance in level-flight to the runway, respectively. Since it is based on historical inputs, this achievable optimum is not a theoretical but an empirical estimate that is truly achievable in the applicable airspace, and that incorporates restrictions as applicable to the corresponding group of like flights.

### Empirical Data for Evaluation of Flight Performance

Cleaning and merging of the terminal area and en route surveillance data is a key component of successful evaluation of empirical trajectories, and estimation of achievable unimpeded paths and the benefit potential for improved operations. Additionally, complex analysis of empirical trajectories is necessary to overcome gaps in archived data, such as runways aircraft used to take off from, or land at, their origin and destination airports.

Also, since this analysis focused on improvement opportunities in nominal conditions, flights conducted during periods with convective weather were filtered out, as were flights delayed by "airport turning" – significant changes in runway configuration that happened within 30 minutes before their actual landing. Note that historical data is stored in quarter-hour time bins, hence the actual configuration change could have occurred in a 16-44 minute window prior to landing. Future analysis will focus on convection and potential improved use of available airspace and airport resources with advanced tools.

The methodology used started with dissecting each empirical trajectory into portions flown within 40nm, and between 40nm and 120nm of both the origin and destination airports. This made it possible to roughly capture the segments flown in terminal areas, and segments applicable to the climb and descent phase of flight. To provide for capturing segments applicable to sequencing and merging before descent and approach, it was necessary to analyze portions of empirical trajectories that were flown between 120nm and 250nm of their destinations. These values – 40, 120 and 250 nautical miles – are somewhat arbitrary, but they provide a good set of values that are standardized across the 41 key airports in the NAS. This allowed for consistent evaluation of the intended segments and the corresponding aircraft performance.

## Optimal Horizontal Profiles

For each of the segments flown within 250nm of each arrival airport, flights were categorized into like-flight groups (i.e. those that share the same geometries of actual trajectories), and then study distribution of actual distances flown by each group. For a NAS-wide study that includes flights with significantly different characteristics, using the 15th percentile of actual distance flown assures a reasonable threshold of optimality that has been empirically derived. While the 15th percentile is somewhat arbitrary, it does represent a set of feasible, empirically-confirmed trajectories (see Figure 4). This is illustrated as the red trajectories shown in Figure 5**;** a set of empirical trajectories of optimal length for this group of like-flights.

Aircraft that flew shorter distance than the 15th percentile determined for their corresponding like-flights are considered as efficient as they can possibly be. Efficiency improvement opportunities for the remaining flights are evaluated as the difference between their individual actual distance, and the 15th percentile determined for the like-flight group. Knowing that air navigation service providers strive to provide the most efficient service possible, that may only be second to safety, this improvement opportunity can also be referred to as the necessary delay, or the delay that was necessary to absorb in order to assure safe sequencing and merging of otherwise unrestricted traffic flows. In other words, by evaluating improvement potential relative to an unimpeded distance that was established through historical records, the analysis effectively accounted for the most significant contributors to inefficiency that occur often enough that they cannot be easily eliminated.

**FIGURE 4:** Example Distribution of Actual Distance for Arrivals at Philadelphia Airport, RWY 27L/R via HOGEY (SW corner post)

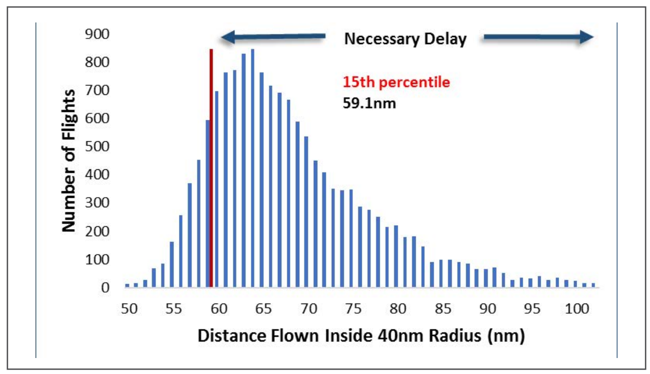


**FIGURE 5:** An Example of Empirical Trajectories for Arrivals to PHL, RWY 27L/R via HOGEY (SW corner post)

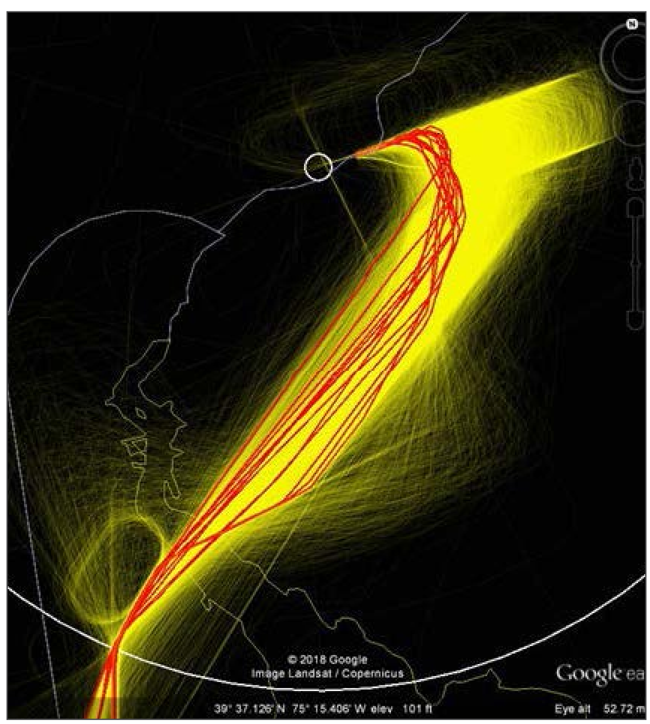

Clearly, a different set of optimal distances and paths would have to be identified for improvements that would enable significant changes in currently flown trajectories. For improvements that rely on the same or similar underlying network of routes, procedures, and fixes, this approach offers realistic and easily customizable thresholds of optimality that are derived from, and validated with empirical records.

## Optimal Vertical Profiles

Since improved metering adherence with TBO tools is expected to increase the use of optimal profile descent procedures, the analysis looked for opportunities to improve vertical profiles as well. For example, arrivals from the southwest corner post to PHL 27L/R fly significantly different horizontal profiles (see Figure 5). Likewise, the vertical profiles of these trajectories are also significantly different (see Figure 6); some of these approaches are continuous descents, but most of the trajectories will experience some degree of level flight.

**FIGURE 6:** Example of Vertical Profiles for Arrivals at Philadelphia Airport, RWY 27L/R via HOGEY (SW corner post)

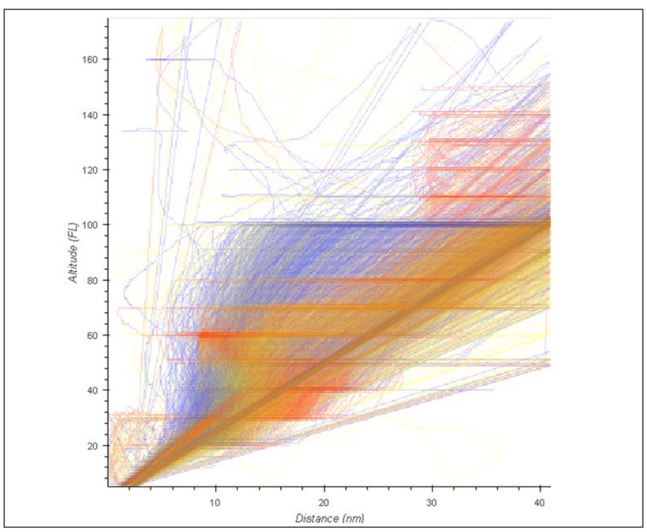


Similar to the optimal lateral profile, an actual optimal vertical profile is driven by many ATM and non-ATM factors, such as de-confliction of air traffic flow and demand levels, respectively; however, neither the underlying factors nor the optimal profiles are recorded in empirical data archives. Therefore, the same empirical approach to estimate vertical flight efficiency was applied as was used

**FIGURE 7:** Variance in Empirical Trajectories within Terminal Airspace of PHL for Flights Arriving via HOGEY and Landing on RWY 27L/R during High and Low Demand Levels in 2016

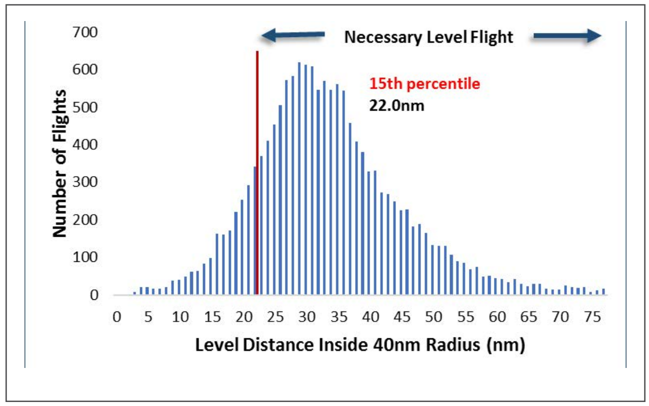


to estimate lateral efficiency—the 15th percentile of actual distance flown in level flight within 40nm of destination (see Figure 7).

It is important to note that, ideally, an aircraft wants to descend without any level-off segments below top of descent. However, interactions between air traffic flow and corresponding procedural restrictions may require aircraft to level-off to remain above departing flows. The empirically derived optimal value prevents overestimating benefit opportunities by allowing some level-offs, as applicable and observed for each like-flight group at each airport.

For the purpose of this assessment, it was assumed that a continuous descent is equivalent to no level-offs longer than 50 seconds inside the arrival terminal area. Additionally, the analysis focused on opportunities to alleviate level-offs below top of descent, which typically happen around 120nm from destination airports.

## Non-ATM Factors Contributing to Flight Inefficiency

Necessary delays currently observed in the NAS are not equivalent to benefit opportunities. Some of these delays are driven by other significant factors outside of ATM influence, such as demand and meteorological conditions at an airport.

Failure to address non-ATM contributions to necessary delays may result in erroneous perceptions of benefit opportunities. For instance, inefficiency in the form of significant vectoring in terminal airspace and long downwinds on approach can be significantly higher during periods with high demand (see Figure 8). On average, throughout the NAS from 2016 to 2018, when compared to periods with high demand, necessary delay was 33% lower during medium demand, and about 50% lower during low demand.

**FIGURE 8:** Example Distribution of Actual Level Distance for Arrivals to PHL, RWY 27L/R via HOGEY (SW corner post)

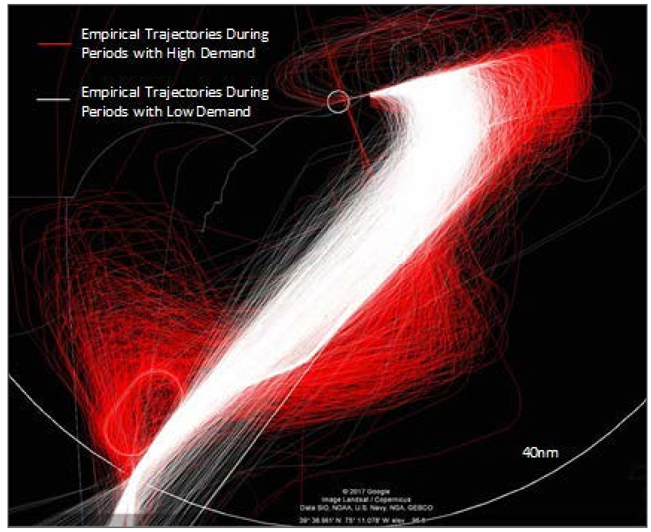


However, unless an ATM improvement can help by increasing the capacity of the system, or by more efficient use of the existing capacity, the improvement may not be able to alleviate this excess vectoring under the same demand levels. At best, with improved awareness of aircraft positions, aircraft speed management and traffic flow metering algorithms, these delays may be shifted upstream where they are more cost efficient, as they would be absorbed at higher altitudes. However, this type of benefit stems from likely reduction in the cost-of-delay rather than in the magnitude of the delay itself.

On the other hand, benefit opportunities can also be easily overstated if one assumes that ATM improvements may improve operations to the point of aircraft behavior observed during periods with low demand also becoming possible during periods of high demand. Direct-to-clearances are often executed in the NAS during periods of low demand (see Figure 9). These are clearly visible as frequent "corner-cutting" that results in a greater variance in empirical trajectories within about 120nm of destination during low demand levels, and therefore shorter overall distances. Also, longer trajectories along predefined standardized arrival procedures are typically flown during periods with higher demand (i.e., clearly outlined by the tight variance in empirical trajectories executed during high and medium demand levels). In addition, similarity in aircraft behavior on descent and approach during periods of medium and high demand may even be greater at other airports.

Hartsfield–Jackson Atlanta International Airport (ATL) is currently one of the most successful users of time-based

**FIGURE 9:** Variance in Empirical Trajectories within 120nm of PHL for Flights Arriving via HOGEY and Landing on RWY 27L/R during Different Demand Levels in 2016

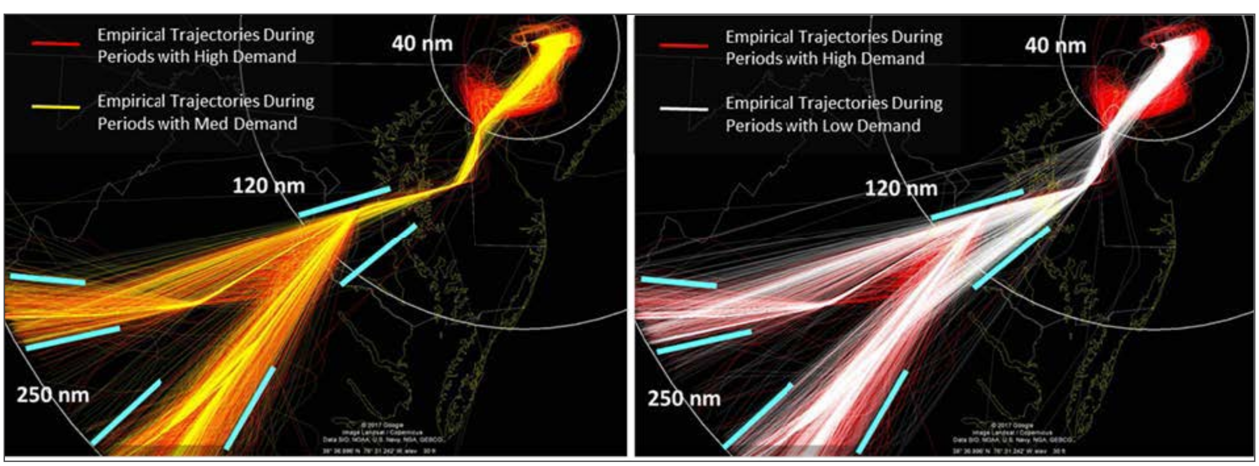

metering in the US. For instance, there is very little difference in actual trajectories during medium (yellow) and high (red) demand, resulting in a sea of orange with only a few red segments, indicating heavier vectoring and longer down-winds in high demand periods (see Figure 10). Such "predictability" of flown tracks is necessary for the existing metering algorithms to produce a feasible metering schedule. During low (blue) demand periods, evidence of shortcuts exists.

**FIGURE 10:** Variance in Empirical Trajectories for Arrivals to ATL via NW Corner Post during West Airport Flow and for Different Demand Levels in 2016

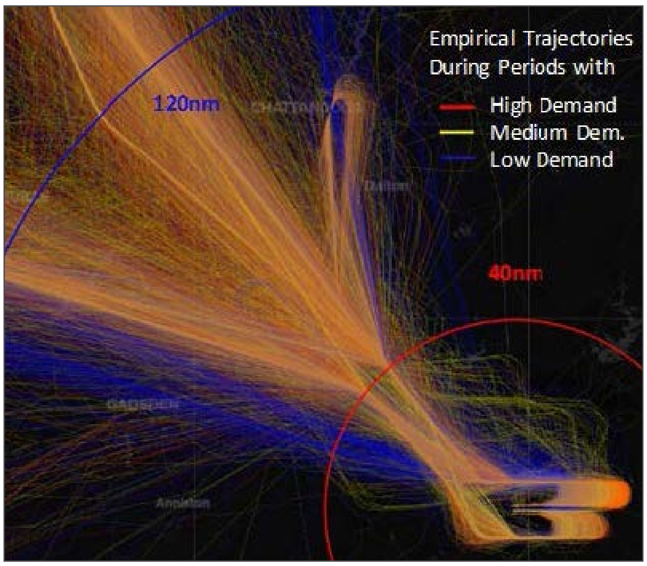


**FIGURE 11:** Variance in Vertical Profiles of Empirical Trajectories for Flights Arriving to ATL via NW Corner Post during West Airport Flow and for Different Demand Levels in 2016

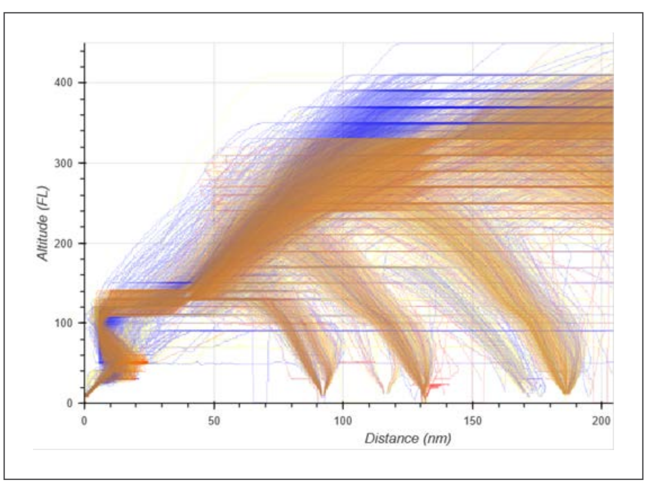


The vertical profiles tell a similar story (see Figure 11). Again, very similar behavior can be seen during both high and medium pressure periods, while during low demand periods, aircraft remain at a higher altitude for a longer time, and initiate their descents closer to their destinations.

There are two very important points highlighted by the above examples. First, conformance to published procedures sometimes results in longer actual distances; however, as demand builds up, that firm structure, and even the "extra distance", are quite necessary to maintain safe merging and spacing, especially in areas with limited airspace or significant flow interactions. Second, as new ATM solutions are considered for implementation, a great effort must be made to preserve efficiencies that are currently possible because of controller flexibility and the ability to react quickly in the moment. This may translate into a requirement for self-adaptive automation that adjusts to demand and other operational conditions, or for a simpler set of business rules for proper timing of automation use. Either way, further investigations are very much needed to better understand the pros and cons of both of these approaches.

## CURRENT NECESSSARY DELAYS IN THE NAS

On average, necessary delay that is currently being absorbed in terminal areas around the 41 key airports in the NAS varies between 2.2 and 9.1nm, and necessary distance in level-flight varies between 3.3 and 12.3nm (see Figure 12). An additional 1.1 to 3.3nm of necessary delay and 7.0 to 17.2nm in distance in level-flight are absorbed between 40 and 120nm of destination.

Airports in the northeast corridor of the US, lead both in terms of average and total necessary delays and distances in level-flight. The highest average delays are observed at George Bush Intercontinental Airport (IAH), Charlotte Douglas International Airport (CLT), Westchester County White Plains Airport (HPN), and Ronald Reagan Washington National Airport (DCA). Total necessary delay and distance in level-flight, however, are the highest at Chicago and Atlanta, driven by the significantly higher operation counts at these two locations.

**FIGURE 12:** Average Necessary Delay and Necessary Level Flight by Destination Airport

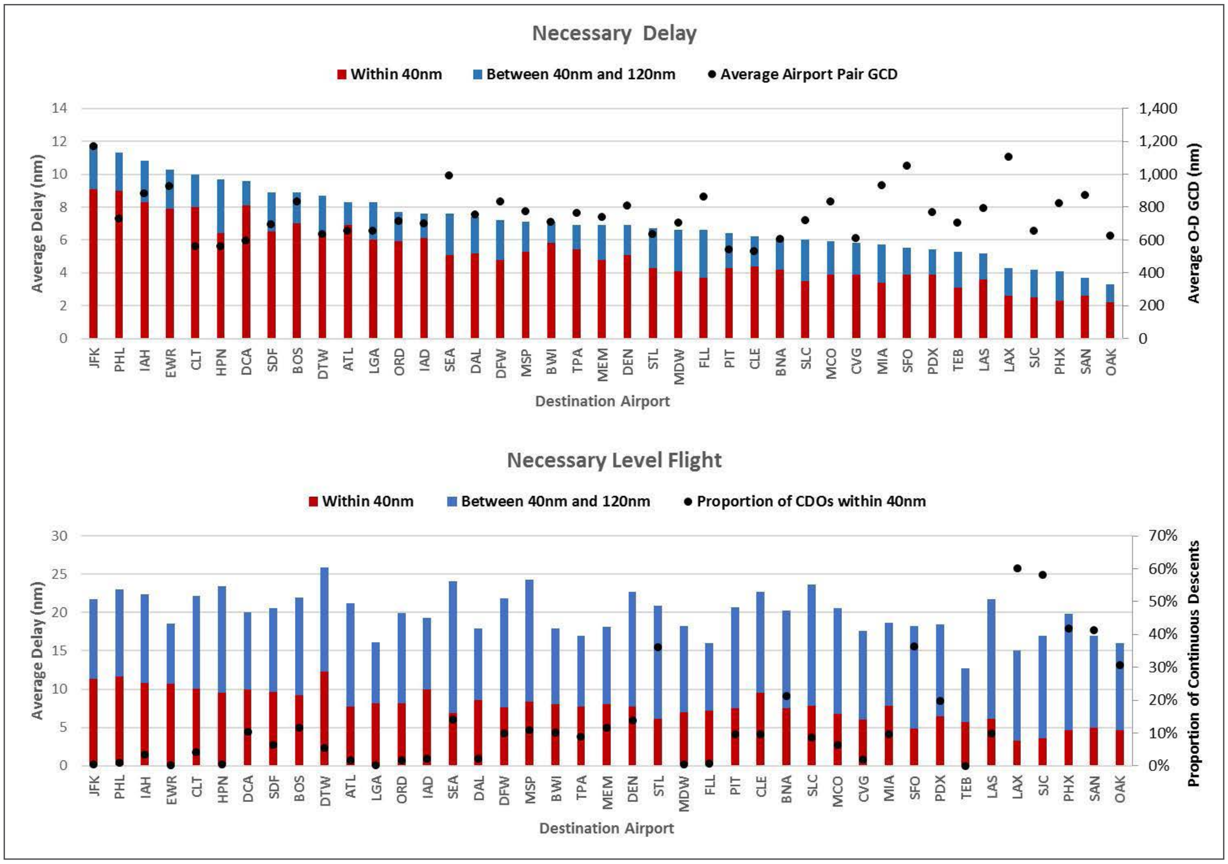


Average distance between origin and destination across all arrivals to an airport is an important indicator of the ability of aircraft to absorb necessary delay en route. That is, the longer the distance between origin and destination for an arrival airport on average, the higher the likelihood that speed reduction will be able to absorb the delays in the en route phase of flight. Of course, this is an issue that needs to be studied carefully by origin-destination pair; however, such theoretical averages at the airport level provide a good "feel" for whether speed reduction could be helpful. The other end of this spectrum is also very important, not only because speed reduction is less likely to be sufficient, but also because of the difficulties with integration of short-haul arrivals into the arrival metering sequence. That is because, long before the short-haul flights are even ready to leave the gate, the arrival metering schedule has already been populated with flights that are already airborne. To assure equity in delay between short-haul and long-haul flights, TBO automation will need accurate information about their planned departure times in order to be able to reserve slots for the flights that haven't taken-off yet but plan to reach the same destination as some of the flights that are already airborne.

The proportion of flights without level-offs within 40nm of destination illustrates vertical efficiencies in the existing system. In fact, at most of the 41 key airports in the NAS, less than 10% of arrivals can descend and land without any level-offs within 40nm of destination. This inefficiency is partially driven by the interaction of air traffic flows to and from the same airport, but also by flows at the nearby airports as well. Clearly, opportunities to improve vertical efficiency through TBO solutions will vary by airport, and depend on their location and air traffic flow complexity.

## BENEFITS FROM REDISTRIBUTION OF NECESSARY DELAY

With TBO tools, delays that are currently absorbed in low altitude airspace will be pushed back upstream, and absorbed at higher altitudes with speed-control and flight paths stretched to the extent possible. Any remaining delay will be pushed to the origin airports. In other words, vectoring that is currently observed in terminal airspace will be replaced by speed control and path stretch en route, and ground delays at the origin, resulting in reduced fuel consumption while achieving the same overall equivalent flight times.

Under TBO, time-based metering is conducted by an integrated automation with three key components: extended metering, coupled scheduling, and pre-departure scheduling. Extended metering places an aircraft into an initial metering schedule when it is about 90 minutes away from its destination (up to about 600nm). Coupled scheduling revises that initial schedule as aircraft gets closer to the terminal entry control-point, and integrates short-haul flights into the arrival metering schedule as well (these flights are known as departures within the same center, and are typically up to 400nm long). Pre-departure scheduling automation reserves a spot in the arrival metering list for the short- haul arrivals before they take off from their origins.

### Methodology for Delay Redistribution

Even under the best operating conditions, it would be unreasonable to expect arriving aircraft to precisely meet TBO schedules. Among other factors, airborne winds and unanticipated delays of other nearby flights can have significant impact on an aircraft's ability to meet its scheduled time of arrival through a control point. To account for these unpredictable variations between scheduled and actual times through a control-point, as well as to provide means for the controllers to line-up aircraft in tight sequences that keep pressure on runways and fully utilize the existing airport capacities, some of the necessary delay will simply remain in the same phase of flight as currently absorbed in the NAS. In fact, one of the requirements for Terminal Spacing and Sequencing (TSAS) automation is to allow for a variability of +/-30 seconds around scheduled terminal entry times. Therefore, the methodology used for estimating TBO delay redistribution opportunities assumes that arrivals will continue to absorb up to a minute of necessary delay within 40nm of destination, and up to a minute of delay in the region between 40nm and 120nm of destination. While these values exceed the requirements, they provide for a more conservative benefits assessment and they ensure that "pressure" on the runway is maintained, thus decreasing the likelihood of missed slots.

For every flight with a necessary delay within 120nm from its destination, the necessary delay is first converted from distance to time by using the reference speed for the affected aircraft type, and Base of Altitude Data (BADA) inputs. To simplify already complex and resource consuming calculations, the reference speed is selected for the distance-weighted altitude (dwAlt) within the same scope. Distance-weighted altitude is a construct similar to the time weighted altitude found in Vempati[9] and Vempati and Ramadani[10], and represents an average altitude for an aircraft, weighted by the proportion of distance in level flight spent at each level-altitude. It is assumed that up to one minute may remain in the terminal area as needed to manage tight aircraft sequences, and temporarily add the remaining delay to that already observed between 40nm and 120nm of the destination. It is then assumed that only up to one minute of that overall delay may remain in the same region, and the remaining amount is moved to the cruise portion of the flight prior to sequencing and merging – roughly outside the 250nm to the destination.

For every flight shorter than 400nm, the amount of delay that could be redistributed from low altitude airspace further upstream is determined. Due to insufficient distance over which speed reduction could be applied, short-haul flights cannot benefit from speed reduction, and would likely benefit most by absorbing their delays at their origin airports.

For every flight longer than 400nm the analysis attempted to absorb the redistributed delay by reducing aircraft speed during the cruise portion of the flight, and before sequencing and merging is typically initiated. Since the extended metering automation determines aircraft sequence up to about 600nm to the destination, and sequencing and merging usually occurs within 250nm from the destination, speed reduction is first attempted outside the 250nm to the destination and within the cruise phase of up to 600nm of destination. The smallest reduction in speed is applied over the available cruise distance (up to a maximum of 350nm). If the necessary delay to be redistributed is too high for even the 10 per cent speed reduction over the available cruise distance, an attempt is made to absorb the remainder between 250nm and 120nm. Since speed reduction during descent would negatively affect its efficiency and could compromise procedural restrictions, it is assumed that it would not be applicable within the 120nm range. Consequently, any delay that could not be absorbed by speed reduction would need to be managed by other means, such as path stretch or ground delay at their origins.

One of the key premises of TBO is to enable aircraft to fly closer to their optimal trajectories, including optimal vertical profiles. However, it will not be possible to eliminate all instances of level flight in descent or approach. To account for this, the amount of the level distance each aircraft flew within 120nm of destination that is greater than the corresponding 15th percentile of its like-group is calculated. This is the necessary level flight that ideally would be moved further back and onto more efficient altitudes before top of descent. However, if the necessary level-flight is shorter than the necessary delay for the same aircraft, it must be left it in the corresponding scope because these lateral and vertical inefficiencies are intrinsically connected and cannot be alleviated separately from each other. Only if the necessary level-flight is longer than the necessary delay for the same aircraft will the calculation attempt to move the difference between the two onto a more efficient cruise-altitude.

For example, suppose two aircraft fly similar trajectories with roughly the same distances within 120nm of destination resulting in 11nm of necessary delay; one of the two aircraft experiences 16nm in necessary level-flight, while the other experiences only three nm. It is reasonable to assume that vertical inefficiency is taken at the same time as the lateral inefficiency to the extent possible. Therefore, the first aircraft will have five nm in necessary level-flight that can be absorbed more efficiently in addition to the 11nm of necessary delay that has already been redistributed to a more cost efficient phase of flight. On the other hand, the second aircraft will absorb the three nm in necessary level-flight along with the 11nm in necessary delay, so no additional efficiency improvement is possible.

## Delay Redistribution Results

In 2018, there were 20 million flights in the NAS, with about 30 perc ent of which arrived to the 41 key airports in the TBO study. Many of these aircraft however, had incomplete records in the data archive used, including:

- Aircraft type was unknown;
- There were surveillance gaps in 4D trajectories;
- The aircraft flew during severe weather, affecting flight times;
- Destination airport changed runway configuration in use, resulting in delays that couldn't be alleviated by TBO; or
- There were insignificant occurrences of other like-flights.

**FIGURE 13:** Proportion of Long-haul Flights by Delay Redistribution Opportunity and by Arrival Airport

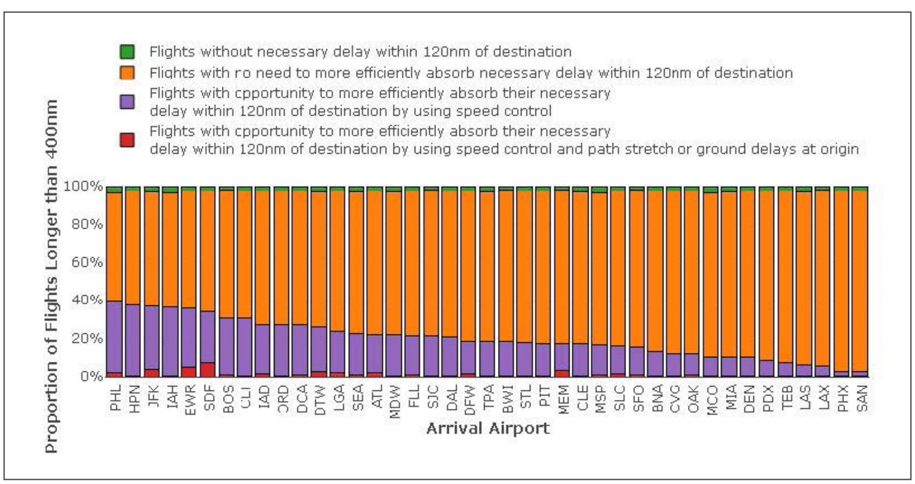

**FIGURE 14:** Average En route Speed Reduction Required for Redistribution of Delay

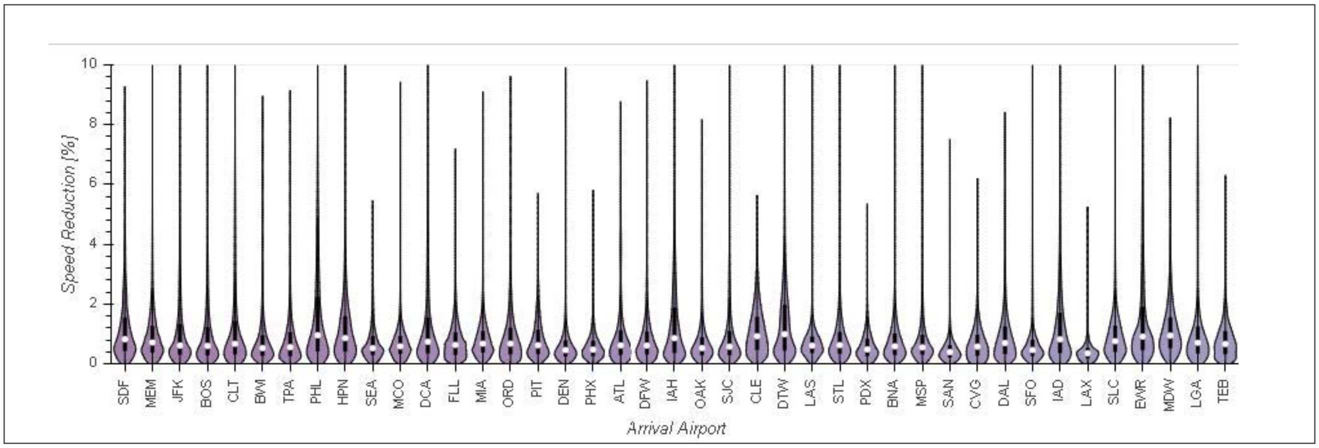


**FIGURE 15:** Average Redistributed Delay by Arrival Airport for Arrivals Shorter than 400nm

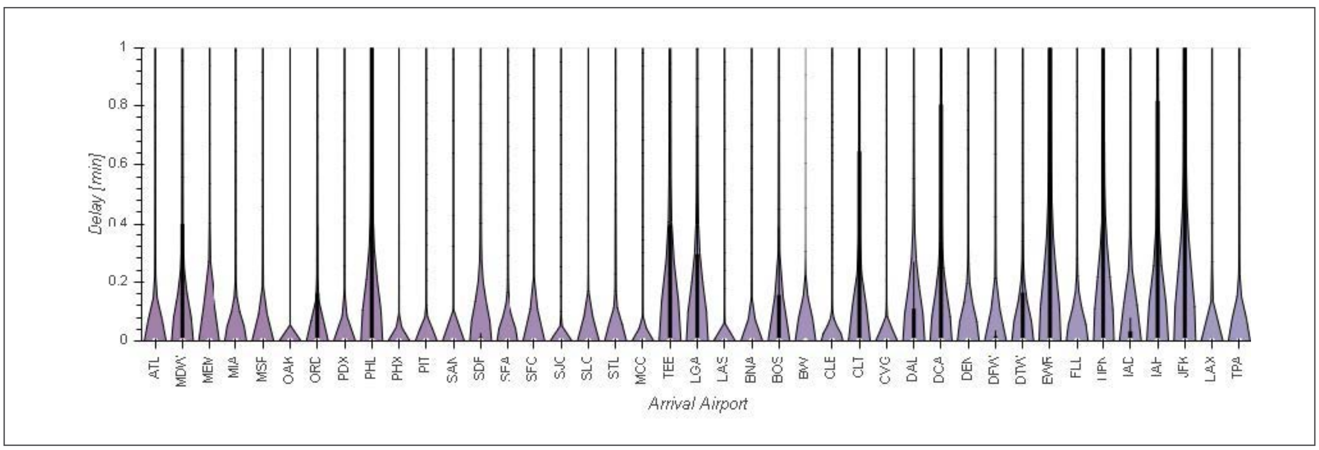


To prevent outliers from potentially skewing the findings, the analysis included only those flights with complete and validated records with a minimum of thirty other like-flights – flights with similar characteristics. Of the almost 3 million flights analyzed, 49 per cent were longer than 400nm.

The study findings are summarized and presented by airport in Figure 13. The average speed reduction by arrival airport for the flights which were required to redistribute delay is shown in Figure 14. Note that the overall average speed reduction was just under one per cent across arrivals to all of the 41 airports. Also, the chart contains the median as a white dot, and the width of each purple bubble represents the distribution of speed control decreases.

On average across the 41 airports, only 17.4 seconds of delay for the short-haul flights could be better managed by redistributing them further upstream. Arrivals to the airports in the northeast corridor suffered the highest inefficiencies with, on average, just under one minute of delay that could potentially be absorbed via more cost efficient means (see Figure 15). The analysis indicated that more than 6,800 hours of delay observed in low altitude airspace in 2018 could have been redistributed through the use of TBO tools such as pre-departure scheduling.

Additionally, about 10 per cent of long-haul flights included in the study required no redistribution of delay simply because their necessary delay within 40nm of destination was lower than one minute, as was the total redistributed delay between 40 and 120nm. For almost 86 per cent of the long-haul flights, the redistributed delay

**FIGURE 16:** Average Fuel Savings for Arrivals Longer than 400nm by Airport

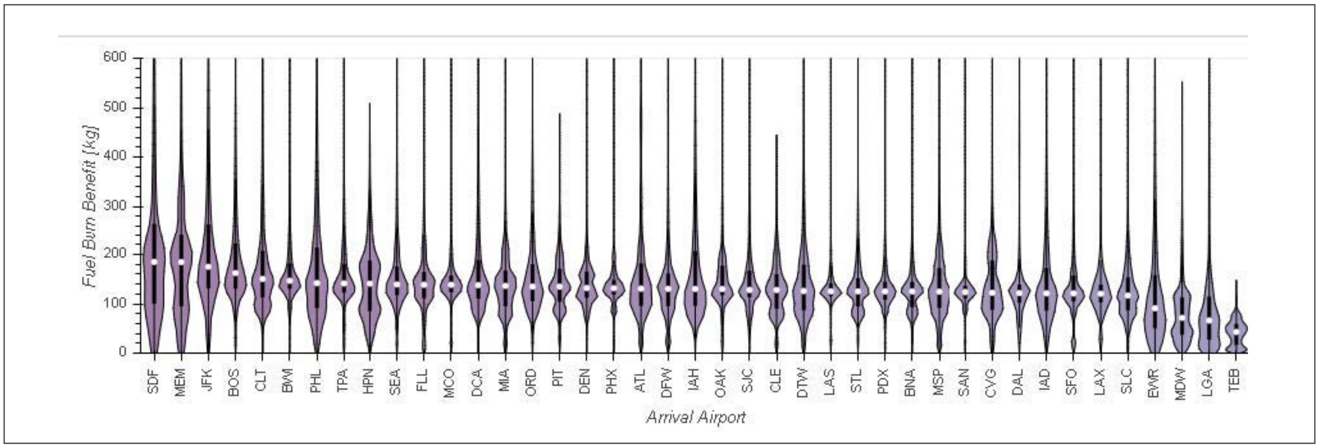


**FIGURE 17:** Average and Total Fuel Savings Benefit by Airport

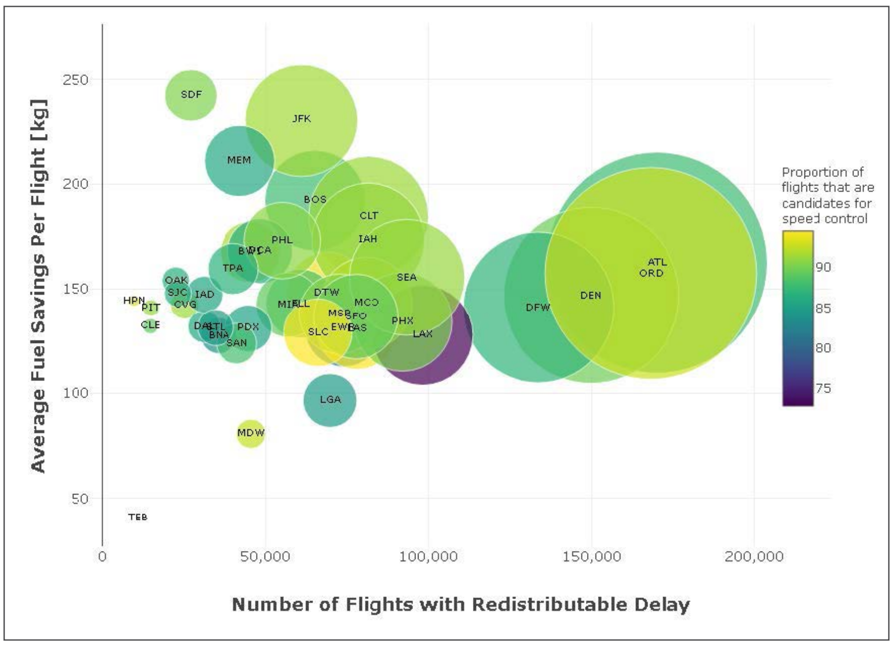


was possible to fully absorb in a more efficient manner by applying en route speed reduction. Finally, about three per cent of long-haul flights required other means to absorb the redistributed delay, such as path-stretch or ground delay at their origins.

Finally, for flights longer than 400nm, the analysis indicated that in 2018, about 1,220 hours of delay could have potentially been redistributed using speed control and additional means of delay management such as path stretch and ground delay before departure.

## Fuel Reduction Resulting from Delay Redistribution

To process more than 3 million flights in a reasonable amount of time, some simplifications are required. One of these simplifications was to use the dwAlt of a particular scope (i.e., inside 40nm, between 40 and 120nm, on route) to represent the altitude of the trajectory during the entire scope. Another simplification was the use of BADA 3.12 inputs to model the impact of speed decreases. While granularity and lack of accuracy of BADA inputs for very low and very high altitude analyses is a known issue, these inputs are considered a standard in the civil aviation industry, and include more aircraft types than any other sources available. A third simplification was to ignore winds and to substitute BADA reference speed for actual aircraft true air speeds to evaluate fuel flow. Speed reduction was applied in increments of one per cent, up to a maximum of ten per cent.

Interestingly, the highest speed reduction is rarely advantageous as a fuel saving option, and is often not necessary to begin with. With BADA data, some aircraft types and altitude combinations resulted in no beneficial speed reduction at all. If an aircraft is flying at optimal fuel flow during cruise, speed reduction will result in higher fuel consumption at that flight level. In fact, this may even result in increased fuel burn in cases when the difference between the fuel flow at reference and at the reduced speed in cruise is higher than the fuel burn for the same amount of necessary delay when flown in low altitude airspace.

It was assumed that under TBO, delay redistribution and absorption through speed reduction would only be attempted if resulting in decreased fuel burn; otherwise, such delays would attempt be absorbed at origin and before take-off. The resulting fuel savings by airport are displayed in Figure 16.

Fuel savings results for flights in this study vary from an average of nearly 245 kg per flight (which had necessary delay and excess level flight) at Louisville International Airport (SDF) to about 40 kg per flight at Teterboro Airport (TEB). Figure 17 displays the average fuel savings per flight for each arrival airport. The color of the circle represents the proportion of flights that are candidates for speed control for each arrival airport. The size of the circle indicates the total fuel savings for each arrival airport.

## CONCLUSIONS

This study was the first ever of its type that involved the magnitude and complexity of processing three years of NAS surveillance data. It involved controlling for numerous variables including: air traffic demand levels, severe on route and terminal weather, airport meteorological conditions, changes in runway configurations, etc. Through simple adjustments in parameter values or aircraft grouping and filtering, these empirical outcomes lend themselves to investigations of additional improvement opportunities, such as the gap in performance during IMC and VMC at an airport, including the extent to which this gap may be closed. In this first application of the new methodology however, the focus was on how much of the current necessary delay in the NAS could be more efficiently absorbed from the fuel consumption and cost of operation perspectives. Since the analysis described above included only those flights for which there was complete, validated and statistically significant records – about half of all the arrivals at the 41 key airports in the NAS in 2018 – the findings presented herein are a conservative estimate of delay redistribution opportunities and the corresponding fuel savings.

Nevertheless, significant work remains to determine if and how such redistribution of delay could be handled by the system. For instance, can the system truly keep the pressure on runways and fully utilize airport capacity given the variability between scheduled and actual times through a control point? Or, can the on route airspace structure and controllers handle additional complexity of delay being absorbed on route? How much of the redistributed delays of short flights can be absorbed on the ground at each origin? Which of the airports needs additional gate or apron capacity to handle such increase in ground delays? Could redistribution of delay lead to a new, unforeseen bottleneck in the system and potentially result in even higher delays?

As new requirements and solutions are perfected through field evaluations and operational use of initial TBO capabilities at select facilities across the NAS, the FAA continues to work on these truly complex and challenging analyses, as it aims to continue to improve both the TBO concept and its assessment of the corresponding benefit opportunities.

## REFERENCES


1. S. F. Szurgyi, "MADAT v2 - Arrival Assignment Algorithm," 2007.
2. CANSO, "Methodologies for Calculating Delays/Improvement Opportunity Pools By Phase of Flight," May 2013.
3. T. Kettunen, J.-C. Hustache, I. Fuller, D. Howell, J. Bonn and D. Knorr, "Flight Efficiency Studies in Europe and the United States," June 2005.
4. D. Knorr, X. Chen, M. Rose, J. Guilding, P. Enaud and H. Hegendorger, "Estimating ATM efficiency pools in the descent phase of flight: Potential saving in both time and fuel," in 9th USA/Europe Air Traffic Management R&D Seminar, Berlin, Germany, June 2011.
5. D. Gouldey, "A Method for Determining Excess Distance in Terminal Airspace," in 15th annual Integrated Communication, Navigation and Surveillance Conference, Herndon, VA, April 2015.
6. J. C. Jones, D. J. Lovell and M. O. Ball, "En Route Speed Control Methods for Transferring Terminal Delay," in USA/FAA Air Traffic Management R&D Seminar, Chicago, IL, 2013.
7. Y. Xu, R. Dalmau and X. Prats, "Effect of speed reduction in climb, cruise, and descent phases to generate linear holding at no extra fuel cost," in 7th ICRAT, Philadelphia, PA, June 2016.
8. S. A. Folse, H. A. Tran, L. L. Jensen and R. J. Jansman, "Cruise Altitude and Speed Optimization Implemented in a Pilot Decision Support Tool," in 16th AIAA Aviation Technology, Integration, and Operations Conference, Washington, DC, June 2016.
9. L. Vempati, "Observed impact of traffic and weather on contnuous descent and continuous climb operations," in 2015 IEEE/AIAA 34th Digital Avionics Systems Conference (DASC), Prague, Czech Republic, September 2015.
10. L. Vempati and A. Ramadani, "O Point-to-Point Operations: How flexible are they?," in 13th AIAA Aviation Technology, Integration, and Operations Conference, Los Angeles, California, August 2013.