



Operating cost based cruise speed reduction for ground delay programs: Effect of scope length



Luis Delgado ^{a,*}, Xavier Prats ^b

^a University of Westminster, 35 Marylebone Road, London NW1 5LS, United Kingdom

^b Castelldefels Telecommunications and Aeronautical Engineering School (EETAC), Office C3-104, UPC Av. Esteve Terradas, 5. Castelldefels 08860, Catalonia, Spain

ARTICLE INFO

Article history:

Received 7 May 2014

Received in revised form 25 September

2014

Accepted 25 September 2014

Keywords:

Ground delay program

Speed reduction

Ground delay program scope

Delay savings

Airborne delay

ABSTRACT

Ground delay programs typically involve the delaying of aircraft that are departing from origin airports within some set distance of a capacity constrained destination airport. Long haul flights are not delayed in this way. A trade-off exists when fixing the distance parameter: increasing the 'scope' distributes delay among more aircraft and may reduce airborne holding delay but could also result in unnecessary delay in the (frequently observed) case of early program cancellation. In order to overcome part of this drawback, a fuel based cruise speed reduction strategy aimed at realizing airborne delay, was suggested by the authors in previous publications. By flying slower, at a specific speed, aircraft that are airborne can recover part of their initially assigned delay without incurring extra fuel consumption if the ground delay program is canceled before planned. In this paper, the effect of the scope of the program is assessed when applying this strategy. A case study is presented by analyzing all the ground delay programs that took place at San Francisco, Newark Liberty and Chicago O'Hare International airports during one year. Results show that by the introduction of this technique it is possible to define larger scopes, partially reducing the amount of unrecovered delay.

© 2014 Elsevier Ltd. All rights reserved.

1. Introduction

In the presence of capacity shortfalls and/or demand peaks, in North America, regulations, i.e. ground delay programs (GDP) are defined at airports. By assigning slots at the congested arrival airport, the demand is not surpassed and costly holding delays around the congested infrastructure are minimized. In order to meet their assigned controlled time of arrival (CTA), flights are requested to wait on ground prior to their departure, and controlled time of departure (CTD) are assigned accordingly. GDPs are implemented within a distance scope; this means that aircraft taking off from an airport outside the scope radius have a slot reserved for them but they are not required to be delayed on ground. The scope has an impact on the amount of delay that is served on ground and on the holding delay that will be required at the destination. The scope also affects the amount of delay that can be recovered if the regulation is canceled beforehand: a large scope distributes the required delay among more aircraft and reduces the holding delay at destination; however, the delay that will be recovered, in the case of early cancellation, will be limited, as delay has already been realized at origin (Ball and Lulli, 2004).

A cost based cruise speed reduction strategy aiming at realizing airborne delay and the effect of this strategy on the amount of delay that can be recovered in early cancellation, as a function of the scope of the program is analyzed in this

* Corresponding author. Tel.: +44 (0) 20 7911 5000x67033.

E-mail address: l.delgado@westminster.ac.uk (L. Delgado).

paper. In this analysis, only the cost of fuel is considered; other costs, such as crew, are excluded. However, the amount of delay that is transferred to the air is relatively small, and therefore its impact on crew cost should be minimal.

Speed control for air traffic management purposes has been the subject of several research studies and projects. Some works propose speed control as a mechanism to enable traffic synchronization strategies (Dravecka, 2006; Lowther et al., 2008). In this context, Günther and Fricke (2006) proposed en-route speed reductions to prevent aircraft from performing airborne holding patterns when arriving in the congested airspace. A similar rationale is behind the ATM long-range optimal flow tool developed by Airservices Australia (Airservices Australia, 2008), where aircraft within a 1000 NM radius of Sydney Airport are required to reduce their flight speed in order to prevent them from arriving before the airport is open, and thereby reduce unnecessary holdings. More recently, a joint Federal Aviation Administration (FAA) and Eurocontrol study, estimated that half the terminal area inefficiency in the current system could be recovered through speed control in the cruise phase of flight, without reducing throughput efficiency (Knorr et al., 2011). Similar research was undertaken by Jones et al. (2013), when speed adjustment was used to sequence aircraft arriving at congested terminal maneuvering areas. At a pre-tactical level, some research has also been conducted considering speed control as an additional decision variable to solve the ground holding problem (Bertsimas and Patterson, 1998; Bertsimas and Patterson, 2000). The economic impact, however, (or the impact solely on fuel consumption) caused by these speed variations is seldom investigated.

With the cruise speed reduction strategy presented by the authors in Delgado and Prats (2012), part of the assigned delay can be absorbed while airborne maintaining the fuel consumption as initially planned. This strategy is interesting considering that air traffic flow management (ATFM) regulations are usually canceled before initially planned (Ball et al., 2010; Inniss and Ball, 2002). Thus, if a regulation is canceled, the aircraft that are already airborne can change their speed to the initially planned one and recover part of the delay at no extra fuel consumption. Delgado et al. (2013) applied this strategy to the incoming traffic at San Francisco International airport (SFO). In that study the whole national airspace system (NAS) was considered and only international flights and aircraft which were already in the air at the definition of the GDP were exempt. In this paper, the effect of the GDP's scope on this strategy and the benefits in terms of delay recovered are studied.

The next section discusses the required background for the paper, with special focus given to ground delay programs. Section 3 is devoted to explaining the equivalent speed concept and its applicability to ATFM regulations. In Section 4 the speed reduction technique is applied to San Francisco, Newark Liberty (EWR) and Chicago O'Hare (ORD) airports. The results of the simulations and the effect of the scope of the program are described in Section 5 and discussed in Section 6. Finally the paper concludes with Section 7, where the main findings are summarized and further research highlighted.

2. Ground delay programs

As stated in Ball and Lulli (2004), when an imbalance between demand and capacity takes place, delay is generated. The total amount of delay required to redress this imbalance is approximately constant. For a reduced capacity at an arrival airport, this amount of delay depends only on the airport acceptance rates (AAR) and the demand, as shown in Fig. 1.

In the ATFM community, it is widely accepted that ground delay at origin airports is preferable to delay near the congested sector/airport, from a fuel consumption (and environmental) point of view (Carlier et al., 2007). In the United States of America when a ground delay program is implemented, the Federal Aviation Administration, acting in its role as traffic flow manager, activates a program where aircraft are assigned to available slots following a ration-by-schedule principle (Richetta, 1994). After this assignment, airlines are given an opportunity to reassign and cancel flights based on updated flight status information and their internal business objectives. This is achieved in a collaborative decision making (CDM) process motivated by a need to combine information sources (Ball et al., 2000; Vossen and Ball, 2006). CDM also ensures that the view of the aircraft operators is integrated in the decision making process ensuring that equity is guaranteed. By applying speed reduction during the cruise, a shift in where the delay is realized is performed, but the arrival slots are preserved. Therefore, it is worth noticing that the strategy suggested by the authors in this paper does not affect the equity issues related with different slot allocation algorithms.

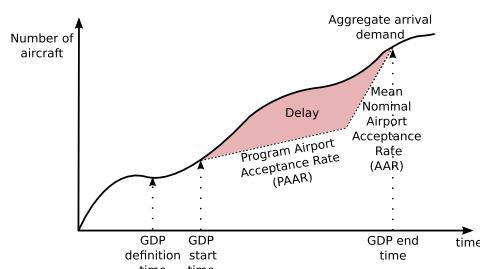


Fig. 1. Ground delay program scheme.

2.1. GDP scope

Not all the aircraft arriving at the airport affected by the GDP are controlled. Some of them are exempted from the FAA assigned delay. A first set of exempted flights are international non-Canadian flights and those airborne at the time the GDP is defined. The second set is GDP dependent and exempts flights taking off outside a certain radius from the affected airport, or outside a number of tiers from the center where the affected airport is located ([Ball and Lulli, 2004](#)). This defines the geographical scope of the program ¹.

One of the main reasons for applying this exemption policy is uncertainty when estimating the arrival capacity at the airport. Predicted capacity reductions are often caused by adverse weather conditions which in turn, are sometimes forecast several hours ahead; overly pessimistic forecasts can lead to excessive ground delays. Since longer flights are delayed well in advance of their expected arrival, if the ground delay is canceled, all the accrued delay is unnecessary. For this reason, most of the delay is usually assigned to shorter-haul flights by exempting flights originating outside the above mentioned scope.

The actual scope is fixed at the GDP implementation and depends mainly on the forecast severity of the capacity reduction and the negotiation during the CDM process. Alternative programs can be defined by changing the scope.

2.2. Unnecessary delays

Unnecessary delays are delays that are ultimately realized unnecessarily, because the regulation is canceled before planned ([Ball and Lulli, 2004](#)). If the scope of the GDP is small, then the majority of the aircraft are exempt from realizing ground delay. However, if the AAR does not increase, i.e., the situation at the airport does not improve, holding delay will be needed. As the radius of scope increases, the pool of flights that serve ground delay increases and therefore, there is a reduction in the total holding delay required. Beyond a certain distance, the holding delay remains almost constant. A program distance shorter than the point where holding delay is minimized is not optimal as unnecessary and expensive holding delay could be transferred to safer and cheaper ground delay. As the scope is increased, the average and maximum delays assigned are reduced, because the total amount of delay is divided between more participants. If the GDP is canceled, the unnecessary delay tends, however, to increase as previously explained.

One of the problems faced when a GDP must be implemented is the estimation of the capacity shortfall and, therefore, the duration of the GDP. Since the predicted capacity at the airport is often subject to uncertainties, airspace managers are typically conservative and the GDP is usually planned to last longer than actually needed. Essentially, it is preferred to have planes waiting on ground, even unnecessarily, and cancel the GDP earlier, rather than having too many flights arriving at the concerned terminal maneuvering area when all of them cannot yet be accommodated.

3. The cruise speed reduction strategy

Flying at the minimum fuel consumption (maximum range) cruise speed for ground delayed aircraft was proposed in [Prats and Hansen \(2011\)](#). In this way, the fuel consumption (and environmental impact) of these flights was reduced while some ATFM delay was absorbed in the air. The impact of this strategy was quantified by analyzing the historical data of all delayed flights to San Francisco International airport over one year.

A different strategy was proposed in [Delgado and Prats \(2012\)](#), where aircraft were allowed to fly at the lowest possible speed in such a way that the specific range remained the same as initially planned. In this case, by flying slower than the maximum range cruise speed (MRC), higher values of delay absorbed in the air were obtained while exactly the same amount of fuel was consumed as initially planned in the nominal flight plan. This strategy was applied to ground delay programs in SFO without a limiting scope and results are presented in [Delgado et al. \(2013\)](#).

In the current operations, all the assigned delay is realized on ground and there is no control over the actual arrival time at the destination. However, controlled times of arrivals should be imposed to ensure that the assigned cruise delay will be realized. This is expected to be possible as monitoring and enforcement of the times of arrival are envisaged for NextGen and SESAR.

3.1. The equivalent speed concept

The specific range (SR) is defined as the distance flown per unit of fuel burnt, and it is usually measured in NM/kg or NM/lb. Typically, there is a speed that maximizes the SR (the MRC speed) due to the fact that the fuel flow increases monotonically with the true airspeed, for typical flight altitudes and aircraft weights.

Since typical operating speeds are higher than the MRC speed, as the aircraft operator also considers the cost of time ([Airbus, 1998](#)), the actual specific range (SR_0) is lower than the maximum. The equivalent speed (V_{eq}) is defined as the minimum speed that produces the same specific range as flying at the nominal speed V_0 , as shown in [Fig. 2](#). The margin between V_0 and V_{eq} depends on the shape of the specific range curve which is aircraft, flight level and weight dependent. It is worth

¹ At present, the NAS is divided into 20 centers, and for each center a first and a second tier are defined. The first tier is the set of all centers immediately adjacent to the center in consideration, and the second tier is the first tier with the centers immediately adjacent to the first tier centers, an so on.

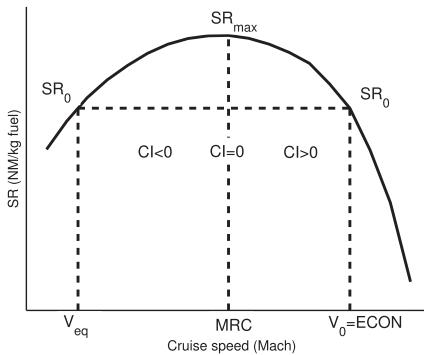


Fig. 2. Specific range as a function of cruise true air speed and V_{eq} definition.

mentioning that V_{eq} might be limited by the minimum operating speed. In this paper, a typical minimum margin against stalling at a load factor² of 1.3 g is considered when computing the minimum operational speed for a given weight and altitude; this ensures good aircraft maneuverability, while preventing the aircraft from stalling (European Aviation Safety Agency, 2011). In the presence of wind, the equivalent speed can be computed considering the specific range with respect to the ground speed. However, the effect of wind is beyond the scope of this paper and its effects and the results of the amount of airborne delay in the presence of wind environments are presented in Delgado and Prats (2013).

3.2. Speed reduction applicability

Notice that in the ground holding problem literature, when the term *airborne delay* is used, it mainly refers to (fuel consuming and undesired) holdings or to path stretching. In this paper, however, the term *airborne delay* is used to define the delay that can be absorbed during the cruise by flying at the equivalent speed without incurring extra fuel consumption, in what is a linear holding, and the term *holding delay* is used to define the delay realized on arrival to the airport, due to lack of capacity.

Flying at V_{eq} , the airborne delay realizable without incurring extra fuel consumption is maximized. Yet, only a few minutes of delay can be performed in the air by flying at this speed during the cruise. For example, in a Chicago O'Hare to Newark Liberty flight with an A320 (626 NM), 7 min of airborne delay can be realized without using extra fuel. See Delgado and Prats (2012) for other examples in the European context. Thus, the airborne delay will be typically lower than the total assigned delay due to an ATFM regulation and therefore, the total assigned delay will be divided between some ground delay, at the origin airport, and airborne delay while flying slower, as depicted in Fig. 3. GDP controlled flights are expected to arrive at a given CTA at the destination airport. With the current GDP implementation, this means delaying the flight at the origin airport by D minutes. After this delay, the nominal flight plan is executed with a total flight time of T_{V_0} minutes. With the enroute speed reduction strategy, the aircraft incurs a ground delay of GD minutes ($GD \leq D$), takes off at a new departure time (CTD') and flies slower than initially planned. In this way, it will take $T_{V_{eq}}$ minutes to reach the destination airport, being $GD + T_{V_{eq}} = D + T_{V_0}$. For a particular flight, if the GDP is not canceled before the aircraft arrives at the destination airport the same amount of delay occurs in the baseline and in the speed reduction scenarios. Moreover, the same amount of fuel is burned in both cases (according to the V_{eq} definition). The benefits of this strategy occur when the assigned delay is reduced due to cancellation of the regulation: in this scenario it is possible to recover part of the delay by speeding up to V_0 , without incurring extra fuel costs over the initially planned flight.

With the current concept of operations, if all the delay is realized on ground, the recovered delay can only be the delay that is not yet accrued. As seen in Fig. 3, for a given flight the delay recovered will be all the initially assigned delay if the cancellation time occurs before the flight's estimated time of departure (ETD) or the difference between the cancellation time and the controlled time of departure if the GDP cancels between the ETD and the CTD. No delay is recovered for a flight that has already taken off, without the flight speeding up over V_0 , leading to more fuel consumption than initially planned (as studied, for instance in Cook et al. (2009)). Yet, this paper focuses on the case where delay recovery is performed at no extra fuel consumption. Thus, in this scenario aircraft flying at V_{eq} can speed up to V_0 to recover some delay if the GDP is canceled when airborne at no extra fuel cost. This strategy ensures that the cost of fuel is constant for the airline operations; other costs such as crew cost are not considered in this study: the absorption of delay during the flight will imply longer flights and therefore, in some cases, higher crew wages. As the amount of delay realized during the cruise is relatively small these costs have not been considered in this research.

Note that with the suggested strategy, all the delay is absorbed during the cruise, therefore, the aircraft arrive at the top of descent at the same time as if the delay is realized purely on ground prior to the departure. This could be combined with tactical sequencing and synchronization strategies at the TMA such as the strategy described in Jones et al. (2013).

² The load factor is defined as the ratio of the lift to the aircraft weight.

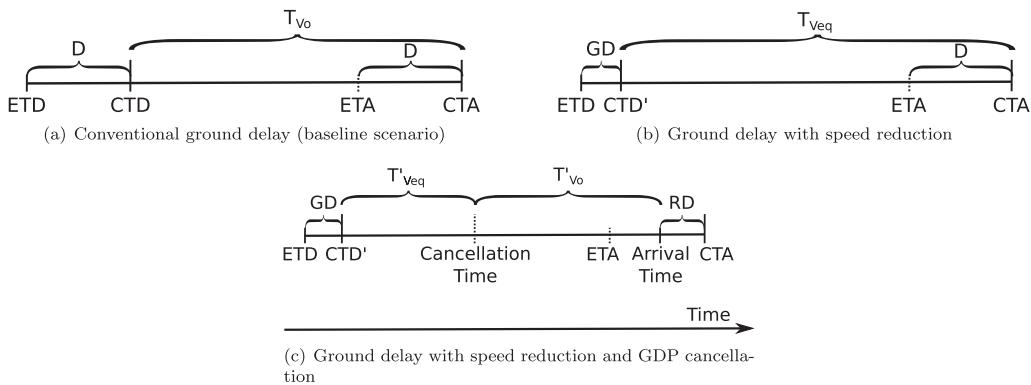


Fig. 3. Application of speed reduction concept on ground delay programs.

The use of the speed reduction implies that there will be a higher variability on cruising speeds. This can represent an increase in the workload the en-route air traffic controllers might face. However, this speed variability is becoming more common as airlines individually optimize their flights costs (Rumler et al., 2010). In Delgado et al. (2013) the authors analyzed the effect of the speed reduction strategy on the number of aircraft airborne to assess if that would represent an overload for the air traffic controllers. Due to the limited delay that can be absorbed during the cruise the variation in number of additional aircraft airborne was not significant.

It is worth noticing that if an extension on the duration of the GDP is needed, the aircraft will still arrive at their destination as planned, realizing the whole assigned delay. Thus, the effect of this strategy on a possible GDP extension is minimal but if the extended GDP is canceled before planned, the benefits of the strategy are again relevant.

The definition of a scope radius, instead of using the whole NAS, has two major impacts on the use of this speed reduction strategy. Firstly, the average assigned delay is increased, as fewer slots are available for the aircraft performing the delay, and secondly, the flying distances available to realize airborne delay are reduced.

4. Case of study: simulation

In the United States, according to the CDM archival database, a total of 1,052 GDPs were defined in 2006 by a total of 49 airports. The first three airports by number of GDPs implemented were Newark Liberty, San Francisco and Chicago O'Hare International with 148 (14.07% of the total), 131 (12.45% of the total) and 120 (11.4% of the total) GDPs set respectively. The GDPs of these three airports represent more than 37% of all the GDPs defined during 2006. EWR and ORD are used as a hub by United Airlines, thus when a GDP is issued in one of those airports a high number of aircraft are affected, leading to high quantities of assigned delay. In 2006, EWR accumulated a total of 2,591,987 min of delay and ORD is the airport with the most number of minutes of GDP served during that year, with a total of 4,533,341 min (25.14% of all the delay generated due to GDPs), and a total of 92,816 aircraft affected. Due to their different locations, center, west and east of the NAS, these three airports present different traffic patterns, as will be discussed in Section 4.1. Due to these characteristics, these three airports have been selected to be studied in this paper.

Simulations are performed using the Future ATM Concept Evaluation Tool (FACET), developed by NASA-Ames (Bilimoria et al., 2000), and the Airbus Performance Engineer's Program (PEP) suite in order to obtain accurate cruise performance data to model the specific range functions of the aircraft and compute, in this way, their equivalent speed (V_{eq}).

4.1. Scenario definition and assumptions for the simulations

The Enhanced Traffic Management System (ETMS) data from August 24th–25th, 2005, is used to generate the traffic information required for the simulations. From the CDM archival database, all GDPs defined in 2006 in SFO, EWR and ORD are analyzed and clustered according to their main characteristics.

4.1.1. Traffic data sample

A total of 5240 flights are simulated to generate the demand for the three airports (1012 for SFO, 1382 for EWR and 2846 for ORD). As only the Airbus family's performances are available, aircraft are grouped into six different families, corresponding to six different Airbus aircraft models: A300, A320, A321, A330 and A340. These families of aircraft are created based on the characteristics of the aircraft, in such a way that all aircraft in the same category have similar performances (see Delgado and Prats (2012) for details).

After this grouping, 4,068 flights are simulated with Airbus performance, representing 77.6% of the total traffic: 723 flights (71.4%) for SFO, 975 flights (70.5%) for EWR and 2370 (83.3%) for ORD. The remaining traffic is not considered for

the speed reduction strategy, either because the aircraft are already flying when the simulation started, or because they are notably different from any of the Airbus models available, i.e., small business jets, turboprops and propeller driven aircraft. All these flights, however, are simulated according to their nominal flight plan to correctly represent the demand at the airport. If any of those flights has delay assigned, it will be performed completely on the ground, as in the current concept of operations.

Fig. 4 presents the histograms of the arrivals to SFO, EWR and ORD as a function of the flight plan distance for the two simulated days. Note that, even if all the traffic is considered in the simulations, only the flights that take off during the simulation from the United States of America or Canada are represented in the figure as these are the flights that can potentially have assigned delay if a ground delay program is implemented. Different traffic patterns are present for the three airports:

- San Francisco International airport is located on the west coast. It receives traffic from the surrounding airports, mainly from Los Angeles International Airport (LAX) and McCarran International Airport in Las Vegas (LAS). Due to its location, the amount of medium haul flights is relatively low. As presented in **Fig. 4(a)**, there is a gap between 1000 NM and 1300 NM distance. However, SFO has a considerable amount of long haul flights from the east coast, for example there are 34 flights from John F. Kennedy International Airport in New York (JFK), and 35 from the Hawaii islands. Thus, for SFO, the demand is divided between short and long flights. 56% of the arrival traffic is generated closer than 1000 NM and 21.7% comes from airports located further than 2000 NM.
- Newark International Airport is located on the east coast, and similar traffic to SFO might be expected. However, as depicted in **Fig. 4(b)**, there is more traffic from closer airports: 75.2% of the traffic comes from airports within a 1000 NM radius. The reason for this is that there is a higher number of major airports located on the east coast than on the west coast, like Washington Dulles International (IAD), Hartsfield–Jackson Atlanta International Airport (ATL) or Chicago O'Hare International. Only 9.1% of the arrival traffic comes from further than 2000 NM, mainly coast to coast flights from Los Angeles International (28 flights), San Francisco International (16 flights) and Seattle-Tacoma International Airport (SEA) (15 flights). For these reasons, EWR mainly has short and medium traffic with a limited number of long flights.

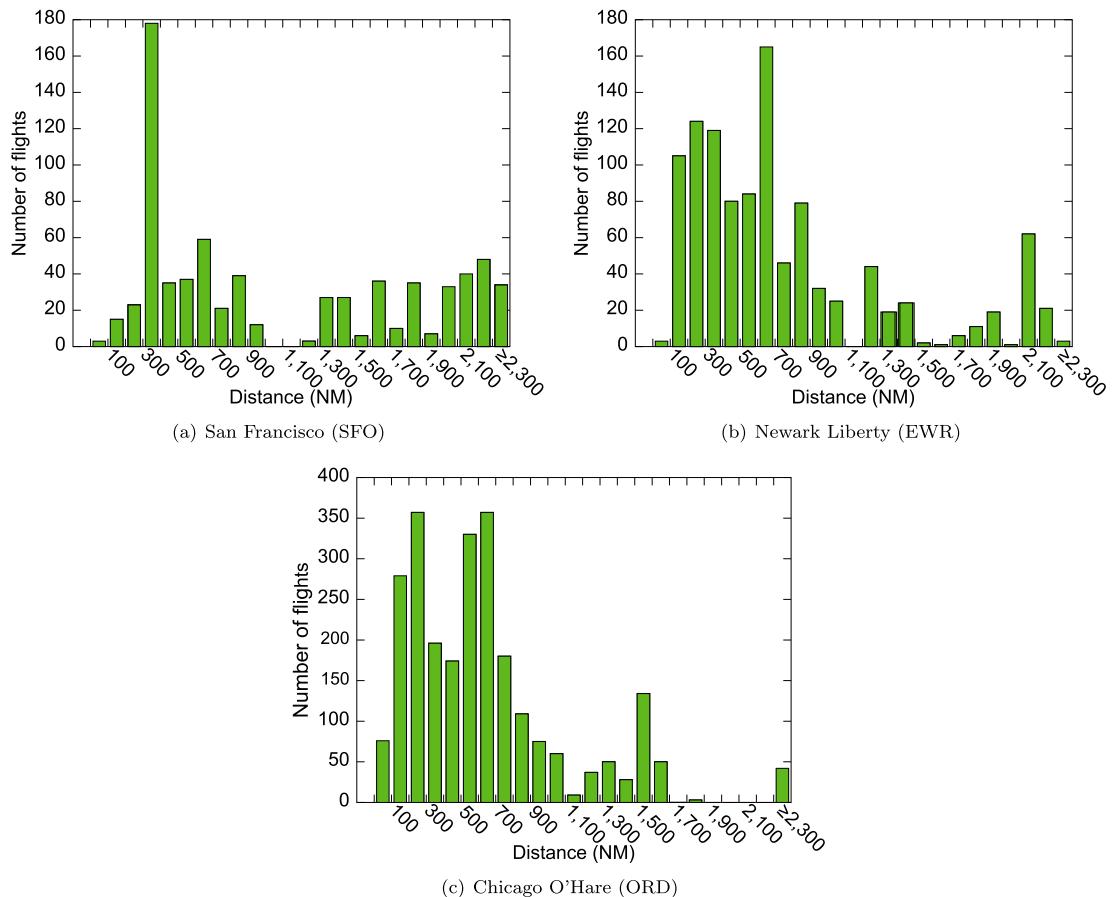


Fig. 4. Histogram of traffic arrivals as a function of flight plan distance to SFO, EWR and ORD for August 24th–25th 2005 and considering only traffic taking off from United States of America and Canada.

- Finally, Chicago O'Hare International generally has short and medium flights; 82.6% of the traffic is from airports closer than 1000 NM, only 1.8% of the traffic comes from further than 2000 NM, all of which are flights from the Hawaii islands. From ORD, SFO is located at 1650 NM, JFK at 660 NM and ATL at 550 NM.

A cost index of 60 kg/min is used for all the flights except the A330 and A340 families, where a cost index of 120 kg/min is selected. These values are representative of common operations (Airbus, 1998). The CI expresses the ratio between the cost of the flight time and the cost of fuel. By choosing the CI the pilot is changing the ratio of cost between fuel and time and, therefore, is determining the speed which minimizes the total cost (Cost = Fuel + CI · Time). It should be noted that the CI value not only affects the cruise speed but determines the whole flight trajectory. This means that the optimal flight level may change and that the climb and descending profiles might also be different for different CI settings.

The same payload has been considered for all the flights realized by the same aircraft model. For A319 and A320 it is assumed that only passengers are carried with a load factor of 80%; for long haul flights (A300, A330 and A340) it is assumed that freight is also transported and their payload has been fixed at 80% of the maximum payload of each aircraft model (ELFAA - European Low Fares Association Members, 2008).

It should be noted that only the change of speed during the cruise is considered, therefore, the flight levels are maintained as defined in the original flight plan. Finally, the simulations are carried out in wind calm conditions, in order to avoid masking the results with the uncertainty associated to actual wind conditions. The reader is referred to Delgado and Prats (2013) for an assessment of the effect of the wind on the proposed speed reduction strategy.

The maximum delay that can be recovered is computed assuming that the aircraft that are delayed on ground, at the cancellation time, can immediately take off and that the airborne aircraft, which are flying at V_{eq} , can speed up immediately to V_0 .

4.1.2. Ground delay programs data sample

In 2006, approximately 75% of all the GDPs were defined due to weather related issues. Different scenarios in airport arrival acceptance rate reflect in most cases well-identified weather patterns in the regions where the airports are located (Liu et al., 2008). For example, in the case of SFO, it is common to have marine stratus which usually burn off around the middle of the day. There are days, however, when the capacity remains at reduced values throughout the day. In addition, in SFO, some reductions in the airport arrival acceptance rate are due to the rainy periods in the winter season.

It is expected that representative GDPs for a given airport can be determined due to the common reasons that trigger the definition of a GDP. The K-means clustering algorithm (Macqueen, 1967) is used in this paper to group all GDPs defined in each of the three airports under study during 2006. The clustering is computed with an iterative process from two to eight clusters and the silhouette coefficient is used to determine the best clustering (Calinski and Harabasz, 1974). The GDPs are characterized by their filed, starting, planned ending and actual cancellation times; being the Euclidean distance between these times considered. It should be noted, however, that in this clustering the AARs are not used as they were not available. Finally, the centroids of each cluster are assumed to be representative of their category and are presented in Table 1. It is worth observing that, in an operational scenario, the GDPs might have a wide variability from one implementation to another. By considering the centroid of the cluster as a representative of the GDPs, the authors are able to generate realistic GDPs which the suggested speed reduction strategy can be applied to.

Only two airport acceptance rates are defined, a reduced one (PAAR), which is considered while the capacity is limited, and a nominal airport acceptance rate. However, there are other possibilities for AARs during GDPs, as each airport has a wide variety of runway configurations.

It is assumed that once the GDP is canceled, the capacity at the airport is unconstrained. Even if this is not always true, since the GDP has shifted the demand, the natural spread of flights times and schedules seems to allow traffic management to use this criterion quite extensively in practice. This assumption is similar to the cancellation policies defined in Ball et al. (2010).

The values considered for the capacity at the airports are in accordance with the runway capacities and operations defined by the FAA. The time where the airport capacity changes from the PAAR to the AAR, see Fig. 1, is computed in such a way that the end of the GDP, i.e., when the traffic demand is equivalent to the available capacity, corresponds to the planned end time defined by the cluster.

Table 1

Clusters centroids for the 2006 GDPs of SFO, EWR and ORD (Hours in local time).

Airport	GDP group	GDPs	Filed time	Start time	Planned end time	Cancel time
SFO	Morning GDPs	91	6h31	8h59	13h55	11h30
	All-day GDPs	24	6h12	8h58	22h32	20h08
	Afternoon GDPs	15	15h42	17h08	22h51	21h06
EWR	All-day-night GDPs	68	10h20	12h36	00h18	22h58
	All-day-evening GDPs	64	11h59	13h30	21h49	19h59
	Afternoon GDPs	16	16h53	16h55	23h14	21h20
ORD	All-day GDPs	65	8h28	9h52	22h19	20h13
	Afternoon GDPs	43	14h58	15h26	22h15	19h58
	Early cancel GDPs	12	7h49	9h02	18h33	9h53

The 131 GDPs that were issued in SFO during 2006 are clustered into three categories with a silhouette coefficient of 0.70. The silhouette coefficient ranges from -1 to 1 which represents the best clustering possible (Rousseeuw, 1987). Therefore, the results show that in SFO the three categories are clearly identified. The first cluster contains the majority of the year's GDPs (91), corresponding to *Morning* GDPs caused by low ceilings. These GDPs are typically declared early morning and canceled when the fog burns off which, on average, is around 2 h and 25 min before initially planned, this types of GDPs are found all year round. The second group are *All-day* GDPs (24) that are also filed in the early morning, but extend throughout the whole day because the meteorological conditions do not improve, having a planned duration of 13 h and 34 min and being canceled, on average, 2 h and 24 min ahead. Finally, the third category of GDPs, correspond to *Afternoon* GDPs (15) with an average duration of 5 h and 43 min and canceled around 1 h and 45 min before planned. These latter two categories correspond to GDPs declared mainly during the winter season. This clustering is in line with the results presented in Liu et al. (2008), where airports were characterized by their AAR during the day; and the centroids GDPs duration and cancellation times are consistent with the values from Cook and Wood (2010). For San Francisco International Airport, the PAAR is considered to be 30 aircraft per hour and the nominal AAR is fixed to 60 aircraft per hour; as the two parallel arrival runways of SFO cannot be independently operated when the visibility is reduced (Janić, 2008; Federal Aviation Administration, 2012a).

In EWR, the three clusters found correspond to *All-day-night* GDPs (68), GDPs that are declared in the morning and planned to last until midnight; *All-day-evening* GDPs (64), which are GDPs filed in the morning, as in the first category, but planned to end around ten in the evening; and finally *Afternoon* GDPs (16) implemented in the afternoon. The layout of EWR airport requires ceilings of at least 2,500 feet and visibility of 5 miles or more for the approaches to runway 11. If these conditions are not met, the capacity is reduced to 29 aircraft per hour, which is the PAAR considered; in nominal conditions, the AAR is set to 46 aircraft per hour (Federal Aviation Administration, 2012b). Runway 11 is usually needed in the afternoon as the traffic is higher (Snell and Tamburro, 2011), thus, the three categories incorporate the afternoon period. Note that on average, for the three cluster categories, the GDPs are canceled between 1 h and 20 min and 1 h and 50 min before planned. The silhouette coefficient for the EWR clustering is 0.35.

Finally, for ORD, three categories emerge with a silhouette coefficient of 0.46. The first cluster, by number of regulations (65), includes the GDPs planned to last throughout the day (*All-day* GDPs). The second group (43) is formed by the GDPs implemented in the afternoon (*Afternoon* GDPs). These regulations are planned to be extended, on average, until 22h15. These two categories are, on average, canceled around 2h00 ahead of plan. Finally, some days in 2006 (12), regulations where declared early in the morning but they were probably not needed; they compose the third category of ground delay programs (*Early cancel* GDPs). These GDPs are, in some cases, even canceled before their start time, their average cancellation time is 7 h and 40 min before planned, only 3 h and 14 min after being filled and 1 h and 51 min after their start. For Chicago O'Hare International Airport, a PAAR of 84 aircraft per hour, in reduced capacity conditions, and a nominal AAR of 112 aircraft per hour are assumed (Federal Aviation Administration, 2012c).

As stated in Ball and Lulli (2004), for each GDP, there are an infinite number of distances that can be selected for the scope. However, the finite set of airports to be included or excluded from the program naturally reduce these possibilities into a discrete set of options. There is no interest in considering an additional distance if it does not encompass a new set of airports. However, for the study undertaken in this paper, three different radii are considered in nautical miles from the airports: 400 NM, 800 NM and 1200 NM. These distances are selected to present a progressive increment in the scope. The distance of the flight plan is considered to decide if an aircraft is affected or not by the radius of the scope of the ground delay program.

4.2. Simulation setup

After analyzing and adapting the ETMS traffic, as explained in Section 4.1, the trip distance of each flight is determined. For this purpose, the flight plans, as defined in the original traffic file, are considered. Therefore, the distance from an origin airport might be different for two flights depending on the actual route flown. Then, by using the Airbus PEP suite and the assumed cost indexes and payloads, the nominal parameters for each flight are computed: initial cruise weight, cruise flight level(s) and speed(s) with the required cruise steps if needed.

The initial traffic is simulated twice. In the first simulation the speed and flight levels of the aircraft are maintained to their nominal values. At each step (one minute) of the simulation the fuel flow is computed according to the Airbus performances of the aircraft. If the flight plan requires a step climb at some point of the cruise, it is adequately simulated. The result of this simulation at V_0 , is the initial arrival demand at the airport. The application of the GDP to this demand, in order to keep it below the airport capacity, results in the amount of delay assigned to each aircraft.

In the second simulation, all the aircraft reduce their cruise speed to V_{eq} . As the equivalent speed varies with the weight, at each simulation step, the equivalent speed is recomputed for all the airborne flights considering their current weight. If a particular aircraft had a change in cruise altitude in the nominal flight, it will also be performed in this second simulation. By comparing the arrival times between the two simulations, the maximum airborne delay that each aircraft can realize without incurring extra fuel consumption is determined.

After this second simulation, the assigned delay is divided into ground and airborne delay. In the case that a particular flight is assigned a delay smaller than the maximum airborne delay realizable by flying at V_{eq} , a new speed (between V_{eq} and V_0) is selected, so that the flight fulfills the CTA. Consequently, those (non frequent) flights will save some fuel.

5. Simulation results

5.1. Application of ground delay program to traffic

Fig. 5 shows the division between holding delay, ground delay and airborne delay for each of the distances and each of the simulated GDPs. The percentage indicates, for each GDP, the relative value of the airborne delay with respect to the total delay assigned. It is worth noticing that, as expected, the total amount of delay needed to accommodate the demand to the airport capacity is approximately constant, as shown in **Fig. 1**. With the increment of the scope, a transfer of where the delay is served is produced.

Once the holding delay is minimized, an extension of the scope does not reduce its value any further. For San Francisco International, **Fig. 5(a)** shows that a radius larger than 800 NM does not significantly reduce the holding delay any more. For example, in the *Morning* GDP, the total holding delay decreases from 508 min for a 400 NM to 88 min for a 800 NM; decreasing to 75 min for a 1200 NM; and only reducing its value by 34 min, to 41 min of holding delay, if all the NAS is considered in the GDP. In ORD similar values of holding delay as in SFO are assigned as depicted in **Fig. 5(c)** (178 min for the 800 NM scope radius in the *All-day* cluster). However, even with the reduced 400 NM radius the holding delay is already small in comparison with the Newark Liberty and the San Francisco scenarios: 182 min of holding delay in the *Afternoon* GDP for a 400 NM radius and 5 min for the 800 NM radius. Conversely, in EWR higher values of holding delay are generally found, see **Fig. 5(b)**. For example, in the *Afternoon* GDPs cluster, with a scope radius of 800 NM, 735 min of holding delay are required to accommodate the demand to capacity.

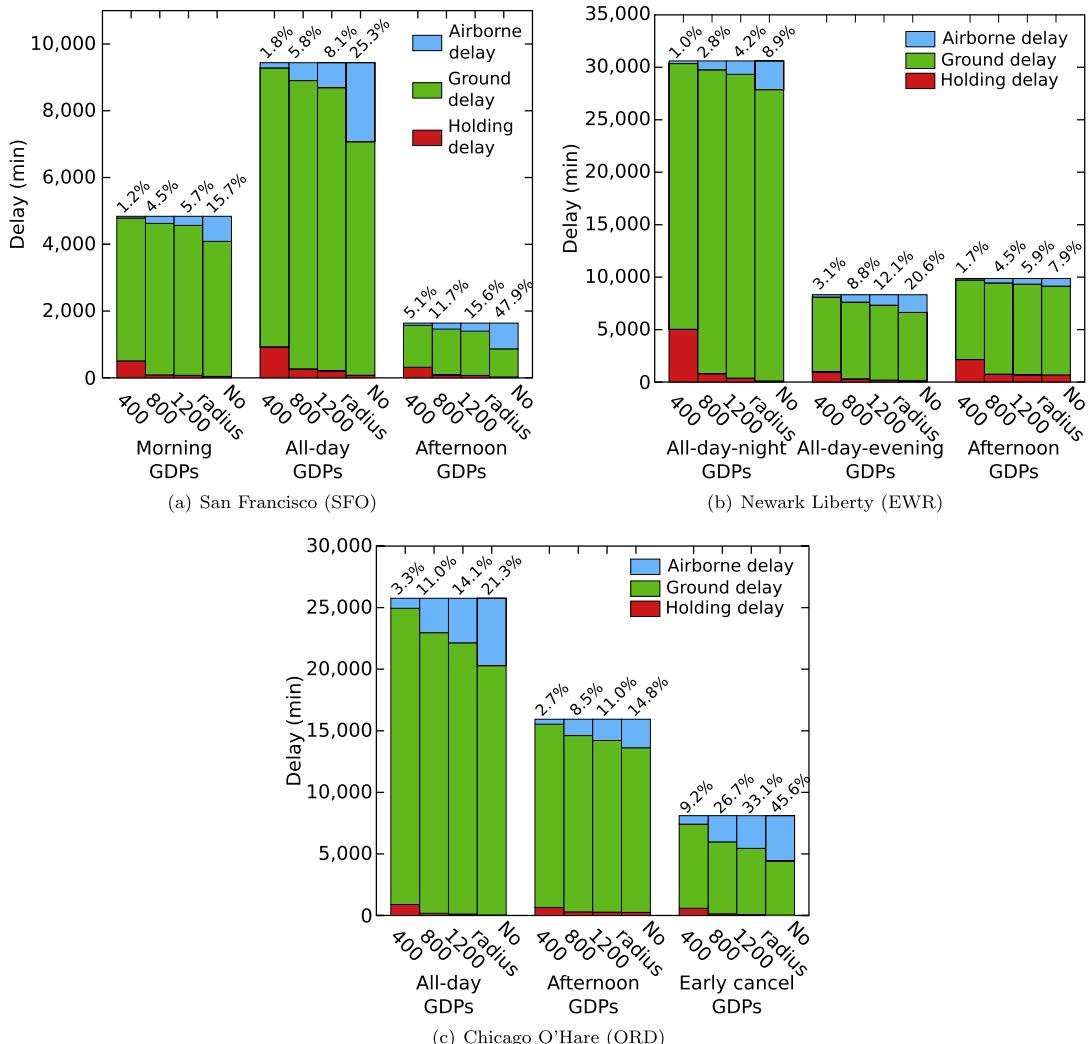


Fig. 5. Division between holding, ground and airborne delay. Percentage of airborne delay over the total assigned delay (ground and airborne delay).

The rest of the delay (ground and airborne delay) would be realized completely on ground if the speed reduction technique were not applied. If the cruising speed is adapted to V_{eq} , part of that delay can be realized while airborne.

The increment of the scope leads to more aircraft in the pool of aircraft which can potentially have delay assigned and that can realize part of that delay airborne, as depicted in Fig. 6. Those aircraft are the ones which can potentially recover part of their assigned delay without incurring extra fuel consumption by speeding up to their nominal speed once the ground delay program is canceled. Those aircraft will also come from further distances and therefore can potentially realize higher amounts of airborne delay during their longer cruise. For average radii lengths (800–1200 NM), the amount of airborne delay varies between 9% and 15% of the total assigned delay, but it is possible to observe that if the whole NAS is used the airborne delay is around 20%, reaching up to almost 50% in some cases. As an example, in Chicago O'Hare International for the *Afternoon* GDPs, the ground delay and the airborne delay assigned for the 400 NM radius is 14,876 min and 408 min respectively (2.67% of the total assigned delay can be realized airborne); for the 800 NM 14,317 min and 1326 min (8.48%); for the 1200 NM 13,946 min and 1726 min (11.01%); and 13,374 min and 2317 min when no scope is defined (14.77%).

Another effect of the increment of the scope is that the maximum and the average delay are reduced as the total amount of delay is divided between more flights. As a consequence, the number of flights that can absorb all their assigned delay airborne, i.e., flying at a speed between V_0 and V_{eq} and therefore saving fuel, increases. Fig. 6 presents these tendencies. For example, in the *All-day* GDPs of ORD, the maximum assigned delay decreases from 166 min to 83 min by setting a scope radius of 800 NM instead of 400 NM. The average delay assigned by aircraft also decreases, from 77.2 min to 36.5 min. This reduction in the maximum and average assigned delay facilitates the absorption of delay while airborne. Thus, when the average delay is only of 10.7 min, as in the whole NAS case of the *Early cancel* GDPs for ORD, 45.6% of the total assigned delay can be realized while airborne by flying at V_{eq} . Notice that the aircraft which realize all their assigned delay on ground are either very short flights which do not have time to realize any airborne delay or aircraft with very different performances from Airbus types, i.e., turboprop, small business jet or propeller driven aircraft.

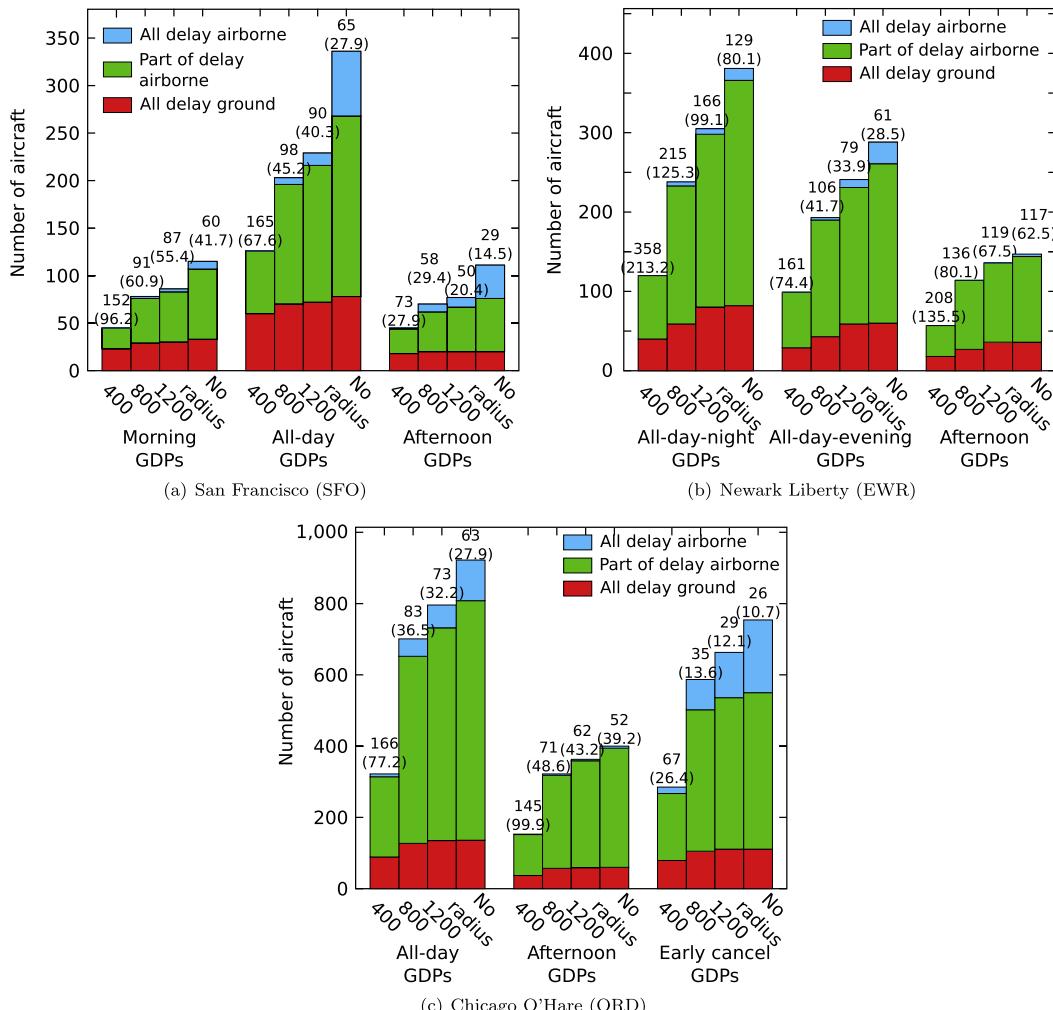


Fig. 6. Aircraft affected by GDP, maximum delay assigned in minutes and, between brackets, average delay per aircraft serving delay in minutes.

The increment on the amount of airborne delay with the length of the radius is a tendency but it is not proportionally related. In some airports, as in Chicago O'Hare, the division between ground and airborne delay increases gradually, while in others, as in San Francisco, there is a large increment when switching from 1200 NM radius to whole NAS scope. The underlying reason is that in order to increase the amount of airborne delay realized, there is a need to increase the pool of aircraft with delay assigned. There is a direct relationship between the type of traffic an airport generally has and the amount of aircraft which are able to realize all their assigned delay airborne: aircraft types and average flight length are as important as the arrival demand.

Fig. 7 presents the distribution of delay (assigned, airborne and ground) as a function of flight plan distance for a representative ground delay program (i.e., Chicago O'Hare All-day GDP).

The delay assigned to the flights on a ground delay program is independent of the distance of the flight plan, as it is only based on the arrival time at the airport. This can be seen in **Fig. 7(a)**. However, as shown in **Fig. 7(b)** there is a relationship between flight plan distance and airborne delay realizable. For very short flight the airborne realizable delay decreases as flight plan distance increases; this is due to the fact that the flights are too short and do not have enough distance to reach an optimal flight level. In general, for distances greater than 200 NM longer flight plans allow more distance to realize airborne delay. For these reasons, the longer the scope the higher the potential of realizing delay in the air (as shown in **Fig. 7(c)**).

For some airports, such as in the Chicago O'Hare case, due to their location, there is a progressive inclusion of airports as the scope distance increases, see **Fig. 4(c)**. Thus, as depicted in **Fig. 7**, the increase of the scope of the GDP implies a smooth inclusion of airports and aircraft which can potentially realize airborne delay.

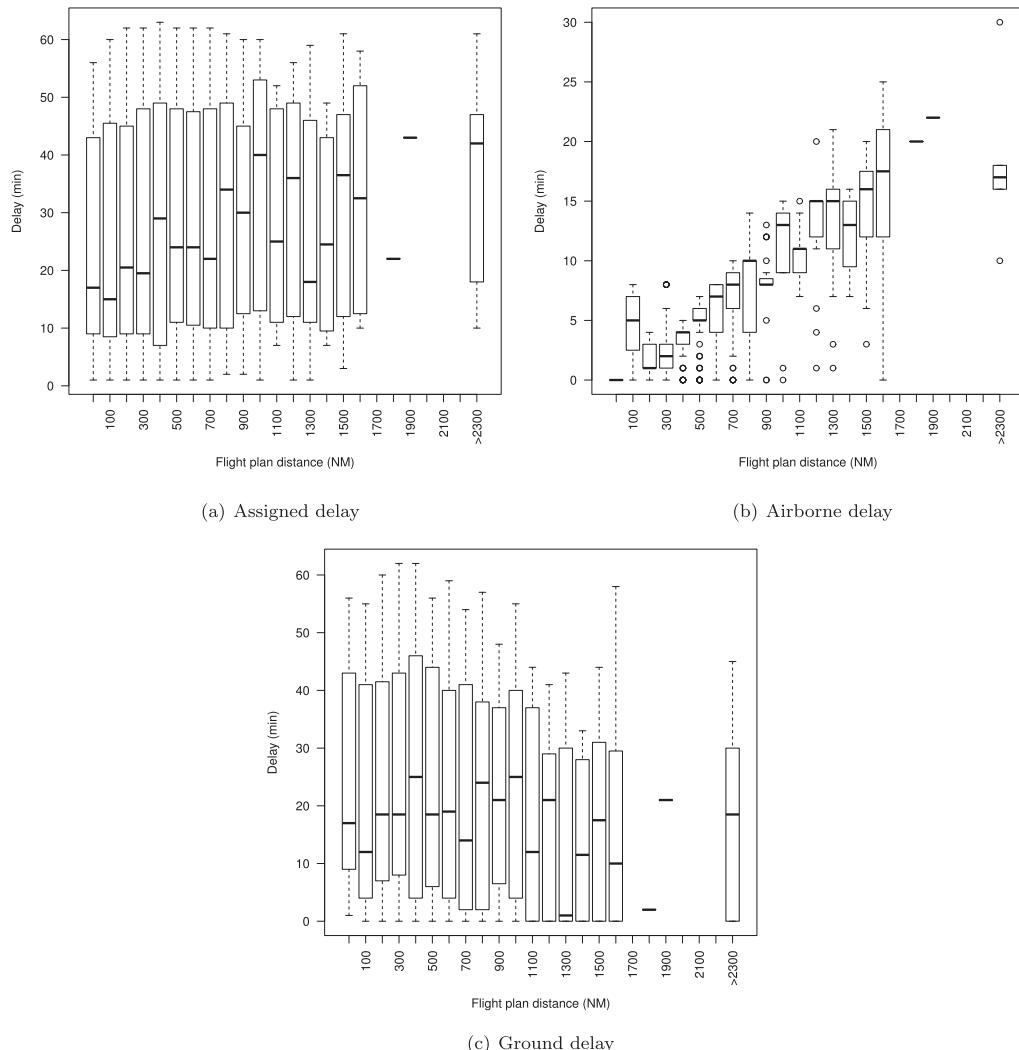


Fig. 7. Delay assigned, ground delay and airborne delay realized as a function of the flight plan distance Chicago O'Hare All-day cluster.

For other airports there might not be a smooth inclusion of traffic such as in ORD. For example in San Francisco International there is a gap between 900 NM and 1200 NM (see Fig. 4(a)), and there is a significant number of long haul flights coming from the east coast, Hawaii, and important hubs, such as Hartsfield-Jackson Atlanta International Airport, which are excluded in the 1200 NM radius. For this reason, a big difference is observed in the amount of airborne delay realizable between having a 1200 NM radius scope or using the whole NAS (from 745 min to 2366 min in the *All-day* cluster).

Finally, the east coast airports, such as EWR, are generally also affected by the long haul demand, but its effect is more limited as the demand on the east coast is mainly composed of shorter flights (see Fig. 4(b)). Thus, with a majority of short and medium haul flights, the amount of airborne delay is relatively small; the effect of the long haul flights can be appreciated in some GDPs such as the *All-day-evening*.

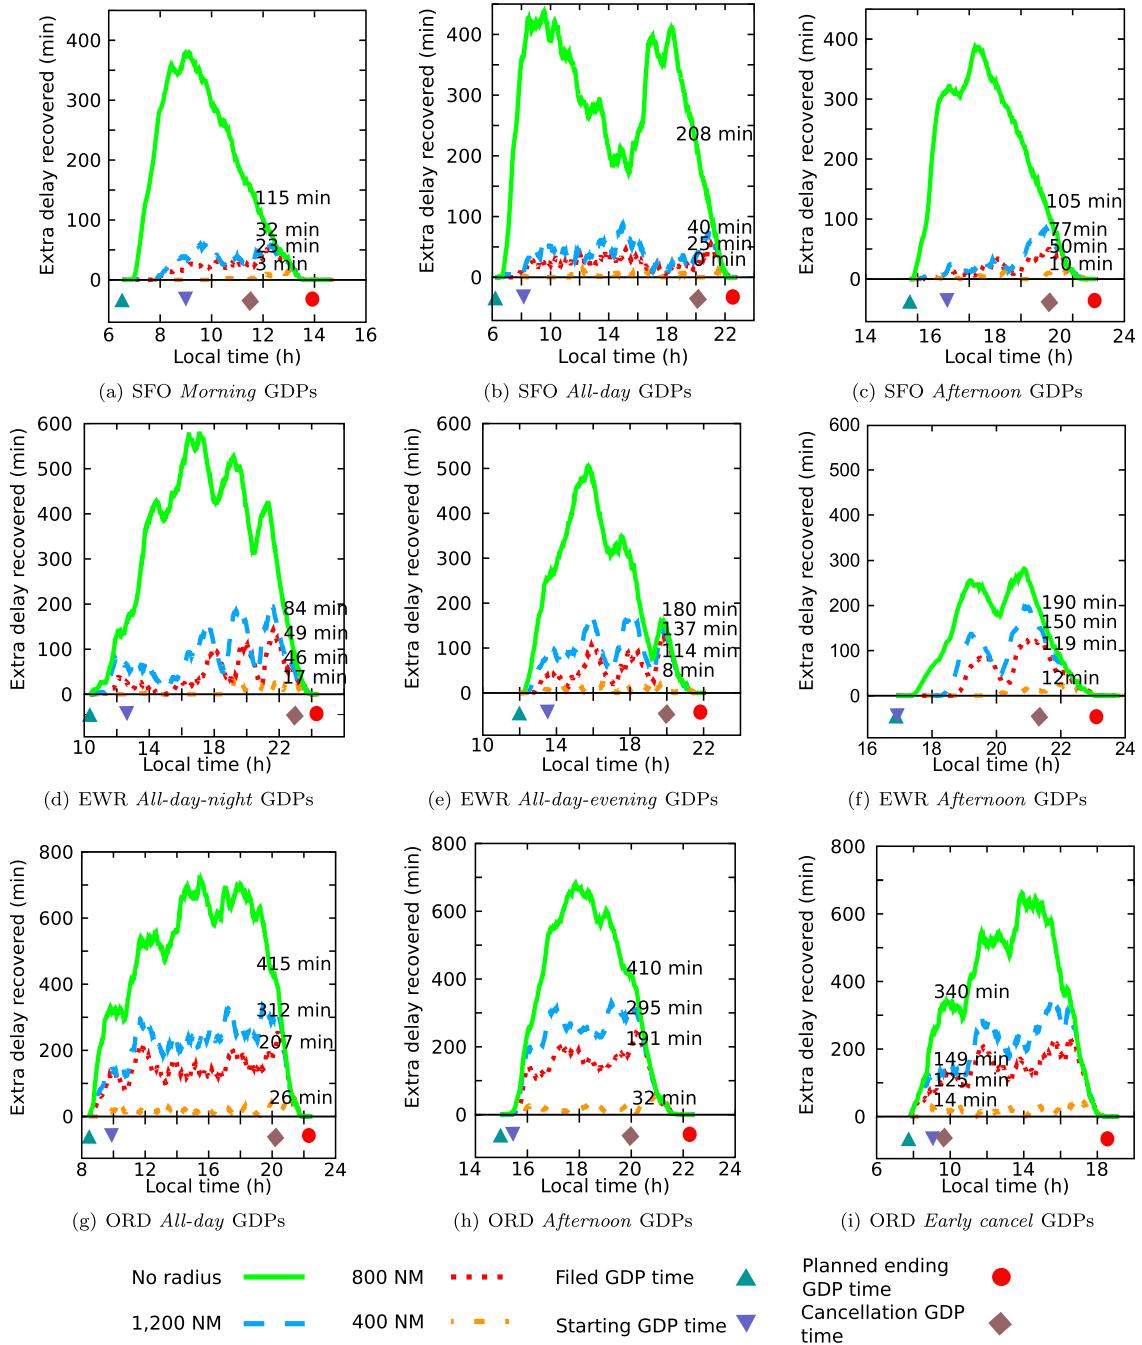


Fig. 8. Extra delay recovered as a function of cancellation time and scope radius. Value the amount of extra delay at the cancellation time according to the centroids of each GDP.

5.2. Delay saved if ground delay program canceled

For each simulated GDP, the amount of delay that is recovered if the regulation is canceled before planned (and the aircraft which are flying at V_{eq} speed up to V_0) is computed. Fig. 8 shows the difference between the delay recovered if the speed reduction technique is implemented and the nominal case where all the delay is realized on ground, if the ground delay program was canceled at different times. Hence, this figure indicates the extra delay that is recovered by the airborne aircraft flying at V_{eq} . Each plot also indicates the amount of extra delay recovered at the cancellation time computed after the GDP clustering process and detailed in Table 1.

The total number of aircraft affected in the ground delay programs applied to ORD is high, but due to the large capacity of the airport, the average delay per aircraft is relatively small for long scope radii. This means that the amount of aircraft absorbing part of the assigned delay airborne is very high, as presented in Fig. 6(c). Therefore, if the regulation is canceled before initially planned, there are a significant number of aircraft in the air that can potentially increase their speed to their nominal one and recover part of the delay. In our simulations, the total airborne delay recovered can be up to 717 min.

As expected, the larger the scope, the more airborne delay is recovered. As mentioned in the previous section, the increment of airborne delay as a function of the radius depends on the traffic and location of the airport. For Chicago O'Hare the benefit of using the speed reduction strategy increases gradually with the length of the scope radius (see Figs. 7 and 8). For example in the Morning GDPs, the extra delay recovered due to this technique increases, respectively, from 26 min to 207 min, 312 min and finally 415 min, as the scope increases gradually from 400 NM to the whole NAS. While for other airports, such as San Francisco, there is, in general, a gap of more than 200 min of extra delay recovered between using a

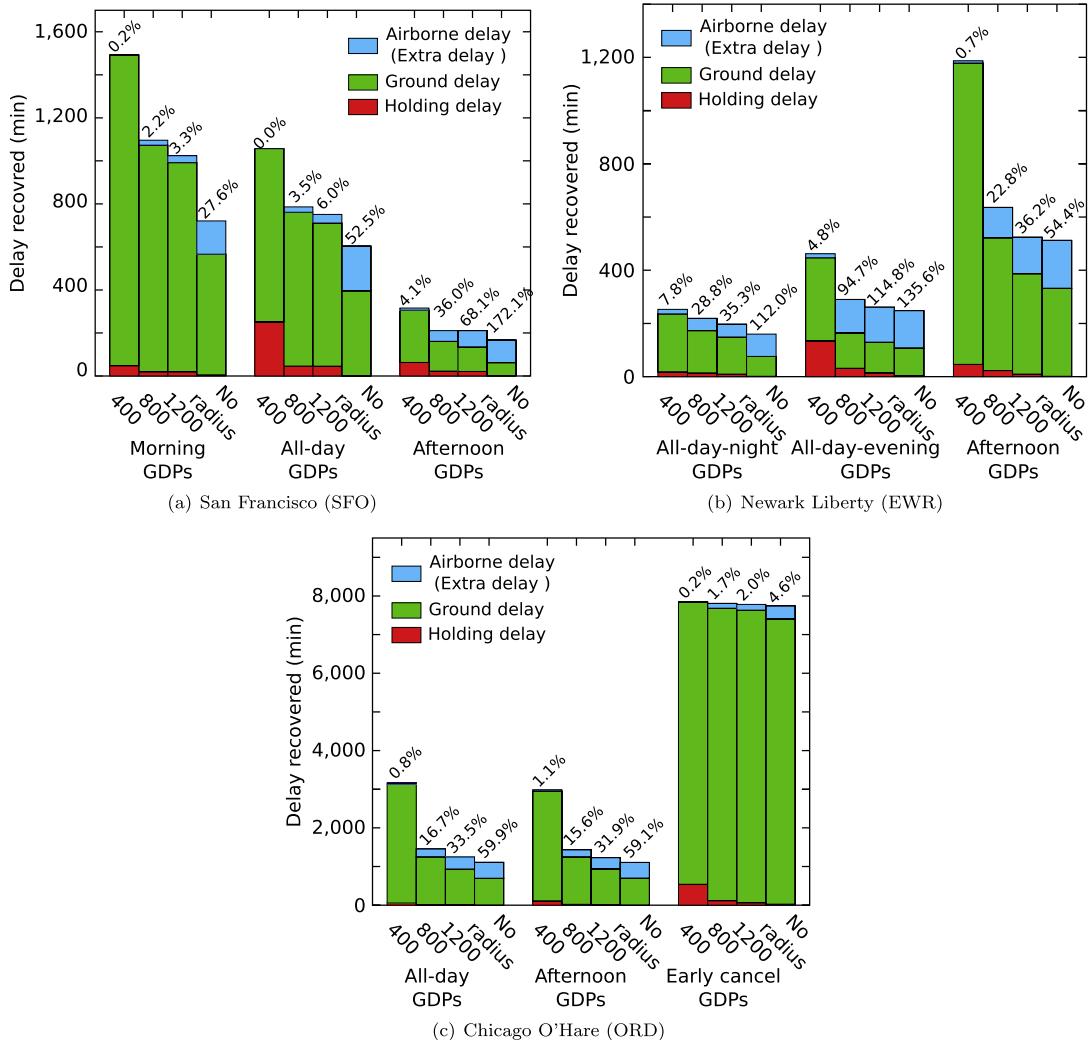


Fig. 9. Results at cancellation time according to the cluster centroids. The percentage indicate the saved airborne delay with respect to the ground saved delay.

1200 NM radius or the whole NAS. Finally for Newark, it is possible to see how the influence of having a radius is of greater and lesser importance as a function of the time of day when the GDP is issued.

From these results, it could be concluded that, if the speed reduction strategy is implemented, the FAA should define the longest possible scope, in order to maximize this extra recovery. In general, a long radius implies less delay recovered, as aircraft have already served delay before the regulation is canceled, if the total delay saved is considered (ground and holding delay) (Ball and Lulli, 2004).

Fig. 9 presents the total delay recovered if the ground delay programs are canceled at their cancellation time according to the clustering performed. As depicted, the extra airborne delay recovered can represent up to more than 150% of the total ground delay recovered. As a direct implication of the use of a scope radius, the longer the radius, the lower the holding delay, and therefore, less holding delay is recovered if the GDP is canceled ahead of planned. Similar behavior is appreciated for the ground delay, as more ground delay is accrued with a larger scope radius due to the longer flight distance the aircraft need to fly to attain their assigned slots, less delay can be recovered. The speed reduction strategy adds an extra amount of delay that is saved.

6. Discussion of the results

When defining the scope of a GDP, the air navigation service provider has to consider the associated trade-offs. On one hand, defining a small scope minimizes the unnecessary delay, thus, if the regulation is canceled before planned, more delay can be recovered. On the other hand, however, a large scope ensures a reduction of the maximum and the average assigned delay as the pool of affected aircraft is increased. This leads to a fairer and less costly solution and the expensive and undesired holding delay is also minimized, as more delay is served on ground. As stated in Ball and Lulli (2004): “*the total ground delay and the total cost may not be related in a simple manner. As the delay assigned to a flight increases, it becomes more likely that passengers will miss connections, that crews will timeout, that the delayed availability of aircraft will cause delays on subsequent flights, etc. Thus, the cost to an airline of 20 flights, each incurring 15 min of delay, as a rule, is less than the cost of 5 flights each incurring 60 min of delay*”. The speed reduction strategy presented in this paper allows a larger scope to be used, minimizing the negative impact on the total amount of delay recovered.

Considering the speed reduction strategy, the larger the scope the more distance is available for the aircraft to realize airborne delay and therefore more delay can be absorbed during the cruise phase without incurring extra fuel consumption. Moreover, more aircraft realizing airborne delay implies that the number of flight that can potentially recover extra delay by speeding up during their cruise if the regulation is canceled, is increased. Finally, as larger scope imply lower average assigned delays, the number of aircraft that can realize all their assigned delay in the air, and therefore save some fuel with respect to their initially planned flight plan, is also increased.

The aggregate extra delay saved using this speed reduction strategy is computed for each radius and presented in **Table 2**. At this aggregate level values are significantly high even if for a single GDP the extra delay recovered is relatively small.

Even if the total extra delay recovered using the speed reduction strategy is relatively small, the use of this technique has an interesting benefit: the difference in recovered delay when the regulation is canceled beforehand is reduced between two consecutive studied radii of scope. For example, as depicted in **Fig. 9(b)**, in Newark Liberty airport, for the *All-day-night* GDPs, when the radius increases from 400 NM to 800 NM, if no speed reduction is realized there is a difference of 63 min of delay recovered. However, if the aircraft are realizing part of their delay airborne, the total delay recovered if the GDP is canceled and the radius is set to 400 NM is only 34 min higher than for the 800 NM radius. Moreover, having a 800 NM radius for that GDP, and using the speed reduction strategy, leads to only 17 min less delay recovered than setting the GDP with a 400 NM radius and realizing all the assigned delay on ground. For ORD, as depicted in **Fig. 9(c)**, if no speed reduction technique is implemented the difference between the delay recovered with a 400 NM radius in the *All-day* GDPs and a 800 NM radius is 1885 min. In the speed reduction scenario, the difference between them is 1704 min (181 min less). In the *Afternoon* case, the difference between the 800 NM and the 1200 NM radius is 309 min if all the aircraft fly at their nominal speed and 205 min if the speed reduction strategy is implemented (a reduction of 104 min). This behavior is found in all the cases studied for all the airports and GDPs.

If the 400 NM radius GDPs are dismissed due to their high amount of undesired holding delay, in general, realizing a longer radius and airborne delay lead to the same amount of delay recovered (or even more) than setting a smaller radius without the speed reduction strategy. As an example, in the *All-day* GDP cluster of Chicago O'Hare airport, only 2 min of extra delay is saved with an 800 NM radius and no speed reduction implemented with respect to a 1200 NM scope radius where the cruise speed reduction technique is realized. In the *All-day* GDP cluster, 172 min of extra delay are recovered if the whole

Table 2

Aggregated extra delay saved for all GDPs during one year per scope radius.

Airport	400 NM rad (min)	800 NM rad (min)	1200 NM rad (min)	No radius (min)
San Francisco International	423	3443	5027	20,672
Chicago O'Hare International	3246	23,168	34,753	48,685
Newark Liberty International	2244	13,016	13,972	17,616

NAS is used with the speed reduction technique with respect to the total delay recovered with a 1200 NM radius and no speed reduction strategy.

7. Conclusion

Defining a large scope when setting a GDP implies that the number of aircraft serving delay is large, therefore, the average delay assigned per aircraft is relatively low and the amount of holding delay required is minimized. On the other hand, the delay is realized considerably ahead of the arrival slot, and if the regulation is canceled, it is not possible to recover some of the delay as it has already been accrued at the origin airport.

The speed reduction strategy analyzed in this paper benefits from larger scopes because aircraft realizing airborne delay will have longer cruises and a lower average assigned delay. More delay can therefore be absorbed in the air. All the aircraft flying at the reduced speed can potentially recover part of their assigned delay by speeding up to their nominal cruising speed.

Airborne delay management compensates partially for the negative effect on delay recovery in early cancellations when large GDP scopes are used: the delay is not realized entirely ahead and it can be partially recovered without fuel penalty. Results show that the total amount of delay recovered for large radii of scope is similar than the one obtained for shorter radii without the speed reduction technique implemented. A less costly and a fairer ground delay program can be implemented, as the total needed delay can be assigned between more aircraft, while the benefit of uncertainty is maintained, i.e. early cancellation. The effects of the use of a scope are airport and demand dependent.

Finally, in this paper only the cruise speed has been modified, as it is easier from an operational point of view. The results presented are considering that the aircraft realize the airborne delay during their entire cruise, therefore an upper bound of the amount of delay that can be absorbed and recovered without incurring extra fuel consumption is given. In [Delgado and Prats \(2012\)](#), however, it was shown that if the flight level is optimized in order to maximize the airborne delay realizable the amount of airborne delay increase significantly. Therefore, instead of solely adjusting the cruise speed, a whole trajectory optimization could be considered to increase the airborne delay while maintaining the fuel consumption. The use of more realistic scenarios including wind should also be considered.

Acknowledgments

The authors thank Dr. Banavar Sridhar, Mr. Alex Morando, Dr. Avijit Mukherjee and Dr. Hok Kwan Ng from NASA Ames Research Center for their comments and help with the use of FACET, Airbus for the use of PEP program, which allowed us to undertake realistic aircraft performance simulations and P. Slocombe for helping with English corrections.

References

- Airbus, 1998. Getting to grips with Cost Index, Flight Operations Support & Line Assistance. Customer services directorate., Blagnac, France.
- Airservices Australia, 2008. Annual Report: 2007–2008, Technical Report, Airservices, Australia.
- [Ball, M.O., Lulli, G., 2004. Ground delay programs: optimizing over the included flight set based on distance. Air Traffic Control Q. 12 \(1\), 1–25.](#)
- [Ball, M.O., Hoffman, R., Chen, C.-Y., Vossen, T., 2000. Collaborative Decision Making in Air Traffic Management: Current and Future Research Directions, Technical Report, NEXTOR –National Center of Excellence in Aviation Operations Research.](#)
- [Ball, M.O., Hoffman, R., Mukherjee, A., 2010. Ground delay program planning under uncertainty based on the ration-by-distance principle. Transp. Sci. 44 \(1\), 1–14. <http://dx.doi.org/10.1287/trsc.1090.0289>, ISSN 0041-1655.](#)
- Bertsimas, D., Patterson, S.S., 1998. The air traffic flow management problem with enroute capacities. Oper. Res. 46 (3), 406–422. <http://dx.doi.org/10.1287/opre.46.3.406>, ISSN 0030-364X.
- Bertsimas, D., Patterson, S.S., 2000. The traffic flow management rerouting problem in air traffic control: a dynamic network flow approach. Transp. Sci. 34 (3), 239–255. <http://dx.doi.org/10.1287/trsc.34.3.239.12300>, ISSN 0041-1655.
- [Bilimoria, K.D., Sridhar, B., Chatterji, G.B., Sheth, K.S., Grabbe, S., 2000. FACET: future ATM concepts evaluation tool. Air Traffic Control Q. 9 \(1\), 1–20.](#)
- Calinski, T., Harabasz, J., 1974. A dendrite method for cluster analysis. Commun. Stat. – Theory Methods 3 (1), 1–27. <http://dx.doi.org/10.1080/03610927408827101>.
- Carlier, S., de Lépinay, I., Husatche, J.-C., Jelinek, F., 2007. Environmental impact of air traffic flow management delays. In: 7th USA/Europe Air Traffic Management Research and Development Seminar (ATM2007), Barcelona, Spain.
- [Cook, L.S., Wood, B., 2010. A model for determining ground delay program parameters using a probabilistic forecast of stratus clearing. Air Traffic Control Q. 18 \(1\).](#)
- Cook, A., Tanner, G., Williams, V., Meise, G., 2009. Dynamic cost indexing – managing airline delay costs. J. Air Transp. Manage. 15 (1), 26–35. <http://dx.doi.org/10.1016/j.jairtraman.2008.07.001>, ISSN 09696997.
- Delgado, L., Prats, X., 2012. En-route speed reduction concept for absorbing air traffic flow management delays. J. Aircraft 49 (1), 214–224. <http://dx.doi.org/10.2514/1.C031484>, ISSN 0021-8669.
- Delgado, L., Prats, X., 2013. Effect of wind on operating cost based cruise speed reduction for delay absorption. IEEE Transact. Intell. Transp. Syst. 14, 918–927. <http://dx.doi.org/10.1109/TITS.2013.2246864>, ISSN 1524-9050.
- Delgado, L., Prats, X., Sridhar, B., 2013. Cruise speed reduction for ground delay programs: a case of study for San Francisco International Airport arrivals. Transp. Res. Part C: Emerg. Technol. 36, 83–96. <http://dx.doi.org/10.1016/j.trc.2013.07.011>.
- [Dravecka, L., 2006. Operational feasibility of traffic synchronisation in central Europe – results. In: 25th Congress of the International Council of the Aeronautical Sciences \(ICAS2006\). ICAS, Hamburg, Germany.](#)
- ELFAA - European Low Fares Association Members, 2008. European Low Fares Airline Association, Members' Statistics December 2008, Technical Report, ELFAA.
- [European Aviation Safety Agency, 2011. Certification Specifications and Acceptable Means of Compliance for Large Aeroplanes CS-25, 11. European Aviation Safety Agency.](#)

- Federal Aviation Administration, 2012a. San Francisco International Airport (SFO) Acceptance Rate, URL <http://www.fly.faa.gov/Information/west/zoa/sfo_sfo_aar.htm>.
- Federal Aviation Administration, 2012b. Newark Liberty International Airport (EWR) Acceptance Rate, URL <http://www.fly.faa.gov/Information/east/zny/ewr/ewr_aar.htm>.
- Federal Aviation Administration, 2012c. Chicago O'Hare International Airport (ORD) Acceptance Rate.
- Günther, T., Fricke, H., 2006. Potential of speed control on flight efficiency. In: 2nd International Conference on Research in Air Transportation (ICRAT), vol. 1, Belgrade, pp. 197–201.
- Inniss, T.R., Ball, M.O., 2002. Estimating One-Parameter Airport Arrival Capacity Distributions for Air Traffic Flow Management, Technical Report, The National Center of Excellence for Aviation Operations Research (NEXTOR).
- Janić, M., 2008. Modelling the capacity of closely-spaced parallel runways using innovative approach procedures. *Transp. Res. Part C: Emerg. Technol.* 16 (6), 704–730. <http://dx.doi.org/10.1016/j.trc.2008.01.003>, ISSN 0968090X. URL <<http://linkinghub.elsevier.com/retrieve/pii/S0968090X08000223>>.
- Jones, J., Lovell, D., Ball, M., 2013. En route speed control methods for transferring terminal delay. In: Proceedings of the 10th USA/Europe Air Traffic Management Research and Development Seminar (ATM2013). Eurocontrol/FAA, Chicago, IL, USA.
- Knorr, D., Chen, X., Rose, M., Guldin, J., Enaud, P., Hegendoerfer, H., 2011. Estimating ATM efficiency pools in the descent phase of flight. In: 9th USA/Europe Air Traffic Management Research and Development Seminar (ATMf 2011), Berlin.
- Liu, P.-C.B., Hansen, M., Mukherjee, A., 2008. Scenario-based air traffic flow management: from theory to practice. *Transp. Res. – Part B: Methodol.* 42 (7–8), 685–702.
- Lowther, M.B., Clarke, J.-P.B., Ren, L., 2008. En route speed optimization for continuous descent arrival. In: AIAA Guidance, Navigation and Control Conference. Georgia Institute of Technology.
- Macqueen, J., 1967. Some methods for classification and analysis. In: University of California Press (Ed.), 5th Berkeley Symposium on Mathematical Statistics and Probability, Berkeley, pp. 281–297.
- Prats, X., Hansen, M., 2011. Green delay programmes, absorbing ATFM delay by flying at minimum fuel speed. In: Proceedings of the 9th USA/Europe Air Traffic Management Research and Development Seminar (ATM2011). Eurocontrol/FAA, Berlin, Germany.
- Richetta, O., 1994. Dynamic solution to the ground-holding problem in air traffic control. *Transp. Res. Part A: Policy Pract.* 28 (3), 167–185. [http://dx.doi.org/10.1016/0965-8564\(94\)90015-9](http://dx.doi.org/10.1016/0965-8564(94)90015-9), ISSN 09658564.
- Rousseeuw, P.J., 1987. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.* 20, 53–65. [http://dx.doi.org/10.1016/0377-0427\(87\)90125-7](http://dx.doi.org/10.1016/0377-0427(87)90125-7), ISSN 03770427.
- Rumler, W., Günther, T., Fricke, H., Weißhar, U., 2010. Flight profile variations due to the spreading practice of cost index based flight planning. In: 4th International Conference on Research in Air Transportation (ICRAT), Budapest, Hungary, pp. 209–216, doi: 978-1-4507-1468-6.
- Snell, D., Tamburro, R., 2011. New York city airspace overview. In: 64th Annual Meeting & Convention. National Business Aviation Association (NBAA), Las Vegas.
- Vossen, T.W.M., Ball, M.O., 2006. Slot trading opportunities in collaborative ground delay programs. *Transp. Sci.* 40 (1), 29–43. <http://dx.doi.org/10.1287/trsc.1050.0121>, ISSN 0041-1655.