

# Cruise speed reduction for ground delay programs: A case study for San Francisco International Airport arrivals

Luis Delgado<sup>a,\*</sup>, Xavier Prats<sup>b</sup>, Banavar Sridhar<sup>c</sup>

<sup>a</sup>Castelldefels Telecommunication and Aeronautical Engineering School (EETAC), Office C3-120, UPC. Av. Esteve Terradas, 5. Castelldefels 08860, Catalonia (Spain)

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

<sup>c</sup>NASA Ames Research Center, Moffett Field, California 94035

---

## Abstract

Ground Delay Programs (GDP) are sometimes cancelled before their initial planned duration and for this reason aircraft are delayed when it is no longer needed. Recovering this delay usually leads to extra fuel consumption, since the aircraft will typically depart after having absorbed on ground their assigned delay and, therefore, they will need to cruise at more fuel consuming speeds. Past research has proposed speed reduction strategy aiming at splitting the GDP-assigned delay between ground and airborne delay, while using the same fuel as in nominal conditions. Being airborne earlier, an aircraft can speed up to nominal cruise speed and recover part of the GDP delay without incurring extra fuel consumption if the GDP is cancelled earlier than planned. In this paper, all GDP initiatives that occurred in San Francisco International Airport during 2006 are studied and characterised by a  $K$ -means algorithm into three different clusters. The centroids for these three clusters have been used to simulate three different GDPs at the airport by using a realistic set of inbound traffic and the Future Air Traffic Management Concepts Evaluation Tool (FACET). The amount of delay that can be recovered using this cruise speed reduction technique, as a function of the GDP cancellation time, has been computed and compared with the delay recovered with the current concept of operations. Simulations have been conducted in calm wind situation and without considering a radius of exemption. Results indicate that when aircraft depart early and fly at the slower speed they can recover additional delays, compared to current operations where all delays are absorbed prior to take-off, in the event the GDP cancels early. There is a variability of extra delay recovered, being more significant, in relative terms, for those GDPs with a relatively low amount of demand exceeding the airport capacity.

*Keywords:* ground delay program, speed reduction, airborne delay, delay recovery, fuel consumption,  $K$ -means clustering

---

## 1. Introduction

Air traffic flow management (ATFM) aims to compensate capacity shortfalls and/or demand peaks, either at an airport or in an air traffic control (ATC) sector; and impose traffic management initiatives that delay aircraft in such a manner that airborne traffic flows do not exceed what can be handled with available resources. In the ATFM community, it is widely accepted that ground delay at origin airports is preferable than airborne delay near the congested sector/airport, from a fuel consumption (and environmental) point of view (Carlier et al., 2007). This statement assumes that airborne delays are in the form of re-routings, air holding stacks or path stretching in terminal manoeuvring areas (TMA). Airborne delay, however, can

---

\*Corresponding author. Tel: +34.93.413.41.26  
Email address: luis.delgado@upc.edu (Luis Delgado)

also be absorbed by slowing an aircraft from its nominal cruise speed, thus intentionally increasing the trip time.

This speed reduction strategy, aimed at partially absorbing ATFM delays while airborne, was presented by the authors in previous publications. Prats and Hansen (2011) proposed that ground delayed aircraft could fly at the minimum fuel speed (the maximum range cruise speed). In this way, the fuel consumption (and gaseous emissions) of these flights were reduced at the same time as some ATFM delay was absorbed in the air. The impact of this strategy was quantified by analysing the historical data of all delayed flights to San Francisco International Airport (SFO) over one year. Results showed values ranging from 5% to 15% of the initially assigned delay that could have been absorbed in the air, leading to fuel savings in the order of 4-7% for each individual flight, if compared with the nominal situation.

A different strategy was proposed by Delgado and Prats (2012), where aircraft were allowed to fly at the lowest possible speed in such a way that the specific range (SR) (NM/kg fuel) remained the same as initially planned. In this case, the aircraft speed being slower than the maximum range cruise speed, higher values of delay absorbed in the air were obtained while exactly the same fuel as initially planned in the nominal situation was consumed. This strategy is interesting if we consider the fact that ATFM initiatives (as ground delay programs) can be cancelled before their initially planned ending time, as is often the case (Ball et al., 2009; Mukherjee et al., 2012; Inniss and Ball, 2002). Thus, if a GDP is cancelled, 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, as shown by Delgado and Prats (2011). If weather clears (and the ground delay program is cancelled), the aircraft that are already airborne, and flying slower, can change their speed to the initially planned one and recover part of the delay at no extra fuel consumption. Thus, the suggested strategy, which can be used in conjunction with the current practise of distance-based exemptions, gives a new alternative that can be used by airlines when dealing with imposed delay due to ATFM.

In this paper, all GDP initiatives that occurred in San Francisco International Airport during 2006 are studied and characterised by a  $K$ -means algorithm into three different clusters. The centroids for these three clusters have been used to simulate three different GDPs at the airport by using a realistic set of inbound traffic and the Future Air Traffic Management Concepts Evaluation Tool (FACET). The amount of delay that can be recovered using this cruise speed reduction technique, as a function of the GDP cancellation time, has been computed and compared with the delay recovered with the current concept of operations. The paper is structured as follows: section 2 discusses the required background for the paper with special focus given to explaining the ground delay program initiatives and the cruise speed reduction concept. Section 3 is devoted to explaining the different simulations that have been conducted. The different assumptions, data and architecture of the simulations is explained. Section 4 shows the application of the ground delay programs to the traffic and the results of the simulation of the GDPs with the speed reduction strategy. The introduction of abnormally slow traffic might have an impact on the air traffic control and air traffic management. Therefore, a brief assessment of the impact on the air traffic management is also performed. Finally the paper concludes with section 5, where the conclusions are summarised and further research highlighted.

## 2. Background

Speed control for air traffic management (ATM) purposes has been the subject of several research studies and projects. The majority of the applications focus on a tactical level, where speed adjustments are used to resolve (or mitigate) aircraft conflicts (Chaloulos et al., 2010). Some other works also propose speed control as a mechanism to enable traffic synchronisation strategies (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 1,000 NM radius of Sydney Airport are proposed to reduce their flight speed in order to prevent them from arriving before the airport is open, and thereby reducing unnecessary holdings. More recently, a joint Federal Aviation Administration (FAA) and Eurocontrol study, estimated that half of the terminal area inefficiency

in the system today could be recovered through speed control in the cruise phase of flight, without reducing throughput efficiency (Knorr et al., 2011).

At a pre-tactical level, some research has also been conducted considering speed control as an additional decision variable (in addition to the amount of time of ground holding) to solve the Ground Holding Problem (Bertsimas and Patterson, 1998, 2000). However, the economic impact (or solely the impact on fuel consumption) caused by these speed variations is seldom investigated. In previous publications (Prats and Hansen, 2011; Delgado and Prats, 2011, 2012, 2013), the authors have studied the impact that changes in speed may have on the fuel consumption and how different speed reduction strategies may be used to enhance ATFM initiatives. This section gives a brief overview of ground delay programs and the principal concepts that are behind the speed reduction strategy described in this paper.

### *2.1. Ground delay programs*

In the US, a ground delay program (GDP) is implemented when an airport is expected to have insufficient arrival capacity to accommodate forecast arrival demand. The FAA, acting in its role as traffic flow manager, proposes a program in which flights are assigned to slots. Some flights are exempted from the FAA assigned delay. A first set of exempted flights are those airborne at the time the GDP is implemented and international non-Canadian flights. The second set is GDP dependent and exempts flights originating outside a certain radius from the affected airport (Ball and Lulli, 2004). One of the main reasons for applying this exemption policy is the uncertainty when estimating the arrival capacity of the airport. Predicted capacity reductions are often caused by adverse weather conditions which in turn, are sometimes forecast several hours before. Thus, too pessimistic forecasts can lead to excessive ground delays. Since flights originating further from the airport must execute their ground delay well in advance of their arrival, most of the delay is usually assigned to shorter-haul flights by exempting flights originating outside the above mentioned radius. The actual value of this radius is fixed at the GDP implementation and depends mainly on the forecast severity of the capacity reduction. In this paper, the use of a radius of exemption is not considered. By doing so, the distance of the aircraft serving delay is maximised, thus the maximum benefits of this strategy are analysed. However, notice that the suggested strategy could be used in conjunction with the current practice of distance-based exemptions. A detailed description of the different parameters defining a GDP can be found in (Vossen et al., 2011).

For each non-exempt (or controlled) flight in a GDP, a controlled time of arrival (CTA) or arrival *slot* is assigned at the destination airport. Based on filed flight plans and weather forecasts, trip times can be estimated with a reasonable accuracy and consequently, the CTA is translated to a controlled time of departure (CTD) at the origin airport. Thus, the CTD is the CTA minus the trip time, and the ground delay is the CTD minus the estimated (scheduled) time of departure (ETD). The assignment of the slots are done in a first-schedule first-served rationing. This is known as Ration by Schedule (RBS) (Richetta and Odoni, 1994; Kotnyek and Richetta, 2006). After this assignment, individual airlines are given an opportunity to reassign and cancel flights based on updated flight status information and their internal business objectives.

Besides ground delay, other strategies can also be initiated in order to solve capacity-demand imbalance problems, such as rerouting or air holding, all of them less desirable because of higher operating costs if compared with ground delays (mainly due to fuel consumption). For this reason, an equivalent slot allocation initiative is implemented when a capacity demand imbalance is detected in an airspace sector: the airspace flow program (AFP).

In Europe, a similar process is implemented in order to deal with airspace or airport capacity constraints. Eurocontrol, through the central flow management unit (CFMU), manages the slot allocation system based also on a RBS basis (Eurocontrol Central Flow Management Unit, 2011). The main difference with the north-american GDP or AFP is that all CFMU controlled flights are affected regardless their origin airport (no delay exception if the flight originates outside a certain radius). Moreover, in Europe, a finite number of administrative slots are also given to airlines to schedule flights in the majority of European airports, aiming to keep demand below maximum capacities while unforeseen situations (such as severe weather affecting a particular airport) are absorbed via the CFMU. However, with very few exceptions, administrative slots are not imposed in the United States, where GDP initiatives can be a frequent issue in some airports where

capacity can be highly reduced if weather degrades. For instance, in San Francisco Airport, in California, when low ceiling clouds are present, capacity drops from sixty planes per hour to only thirty. Due to restrictions on independent parallel runway configurations under instrumental meteorological conditions (IMC) (Janic, 2008).

One of the problems that is faced when a GDP must be implemented is the estimation of the capacity shortfall and therefore, the duration of the GDP initiative. For example, for GDPs caused by degraded weather, if weather clears before forecast it will lead to under use of capacity at the airport and result in unnecessary delays. Conversely, if reduced capacity conditions last longer than expected the GDP will have to be extended and/or inefficient air holdings will be necessary near the destination airport. Since the predicted capacity at the airport is often subject to uncertainties, airspace managers are typically conservative with these scenarios and the GDP is usually planned to last longer than actually needed. Essentially, it is preferred to have planes waiting on ground, even if not necessary, and cancel the GDP earlier rather than having too many flights arriving at the concerned TMA when the available capacity cannot yet accommodate all of them. As can be seen in Table 1, on average at SFO the GDPs are cancelled almost two hours before the initial planned duration. This leads to a under use of capacity at the airport and to unnecessary ground delays (Cook and Wood, 2010).

Although this paper focuses on some GDPs that were implemented in SFO, the speed reduction strategy would be also valid for AFP (or CFMU) initiatives. All these ATFM initiatives assign ground delays to a subset of flights and therefore, by flying slower, part of the assigned delay could be performed airborne at no extra fuel consumption.

Table 1: Average statistics for SFO (2005-2007). (Cook and Wood, 2009)

|  |       |
|--|-------|
| Initial average affected flights   | 79    |
| Initial total delay (min)  | 3,642 |
| Initial average maximum delay (min)  | 98    |
| Initial average delay (min)  | 44    |
| Planned average overall duration   | 4h51  |
| Actual average duration (Cancellation time minus start time)                 | 2h52  |
| Early average cancellation time (Planned end time minus actual GDP duration) | 1h59  |

## 2.2. The cruise speed reduction concept

For a given flight, three types of costs are present: fuel costs, time-dependent costs and fixed costs, which are independent of the time or fuel consumption (such as landing fees or aircraft ground handling). As shown in 1(a), fuel and time-dependent costs depend on the flight cruise speed.

The optimal speed that gives the minimum fuel consumption for a given flight distance is the Maximum Range Cruise (MRC) speed. On the other hand, time-related costs decrease as speed increases, since trip times become shorter. Depending on the importance given by the operator to time related costs, the optimal speed for a given flight will change. To help the operator in assessing this trade-off, the Flight Management System (FMS) of the aircraft allows the pilot to enter a cost index (CI) parameter (Airbus, 1998).

The CI expresses the ratio between the cost of the flight time and the cost of fuel. Thus, a CI set to zero means that the cost of fuel is infinitely more important than the cost of the time, and the aircraft will fly at the MRC speed. On the other hand, the maximum value of the CI gives all the importance to flight time<sup>1</sup>. In this case, the aircraft will fly at the maximum operating speed with, in general, some safety margins. By choosing the CI the pilot is changing the ratio of cost between fuel and time and therefore, is determining

---

<sup>1</sup>Strictly speaking, CI is defined as the cost of time divided by the cost of fuel and multiplied by a scalar. Depending on the FMS vendor, this scalar might be different and, therefore, the actual value of the maximum CI too. Typical CI maximum values are 99 kg/min or 999 kg/min

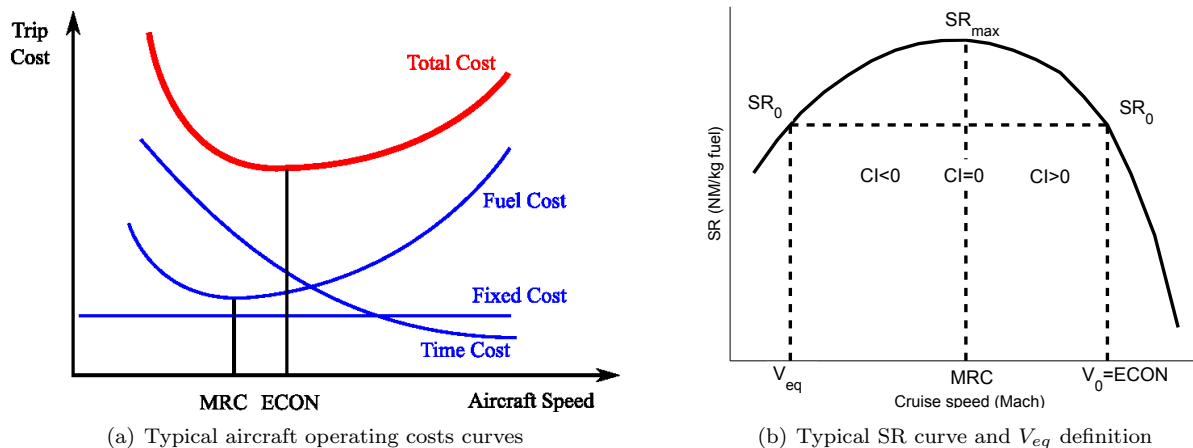


Figure 1: Aircraft operating costs and specific range (SR) as a function of the cruise speed

the speed which minimises the total cost. This speed is usually called the ECONomic speed and will be denoted as  $V_0$  in this paper (see Figure 1(a)). It should be noted that the CI value not only affects the cruise speed but also 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.

### 2.2.1. The equivalent speed

Given a flight distance, a payload weight and a cost index, the optimal flight level, the optimal cruise speed ( $V_0$ ) and consequently, the fuel needed for that particular flight (block fuel), are fixed. Figure 1(b) shows the relationship of the specific range (SR) with the cruise speed. The specific range is defined as the distance that can be flown per unit of fuel burnt, and it is usually measured in NM/kg or NM/lb. The maximum SR is achieved when flying at the MRC speed which is the same as minimising the fuel consumption per unit of distance flown. Since typical operating speeds (ECON speeds) are higher than the MRC speed, the actual specific range will be lower than the maximum one.

The *equivalent speed* ( $V_{eq}$ ) is defined as the minimum speed that produces the same specific range ( $SR_0$ ) as flying at the nominal speed  $V_0$  (see figure 1(b)). 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. As the aircraft flies, fuel is burned and therefore its weight changes leading to changes in the  $V_{eq}$  speed. It is worth mentioning that  $V_{eq}$  might be limited by the minimum speed of the aircraft at that given flight level and weight with some safety margins. In this paper, a typical minimum margin against stalling at a load factor<sup>2</sup> of 1.3g has been considered when computing the minimum operational ( $V_{min}$ ) speed for a given weight and altitude<sup>3</sup>.

The goal of the speed reduction strategy proposed in this paper is to maximise the airborne delay but without incurring extra fuel consumption with respect the initially planned flight plan. Yet, only a few minutes of delay can be generated in the air by flying at  $V_{eq}$ . For example, in a typical Frankfurt International Airport - Madrid Barajas flight (769 NM), 7 min of airborne delay can be realised without using extra fuel consumption (see in (Delgado and Prats, 2012) other example flights). Therefore, the airborne delay will be typically lower than the total assigned delay due to an ATFM regulation (such as GDP). Thus, the total assigned delay will be divided between some ground delay, at the origin airport, plus airborne delay while flying slower.

<sup>2</sup>The load factor is defined as the ratio of the lift of an aircraft to its weight.

<sup>3</sup>In order to ensure good aircraft manoeuvrability, while preventing the aircraft from stalling, the minimum operational speed is set to the stall speed at a given load factor. This load factor is typically chosen at 1.3g. (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 defining the SR as ground NM/kg or ground NM/lb. However, the effect of wind is out of scope of this paper. Results of the amount of airborne delay in the presence of wind and its effects are presented in (Delgado and Prats, 2013).

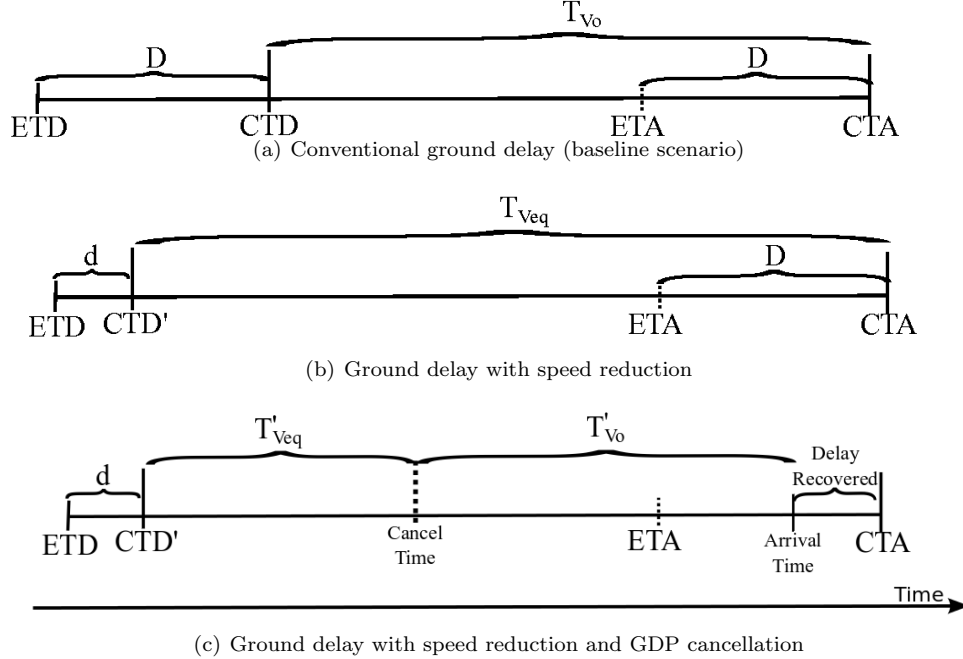


Figure 2: Schematic representation of the delays in the baseline and speed reduction scenarios and the delay recovery in case of GDP cancellation

### 2.2.2. Speed reduction applicability

Figure 2 compares the speed reduction strategy with the current ground delay strategy (baseline scenario). GDP controlled flights are expected to arrive at a given CTA at the destination airport, with a given time window or slot. With the current GDP implementation, this requires 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, as depicted in Figure 2(a).

With the en-route speed reduction strategy, the aircraft incurs a ground delay of  $d$  minutes (with  $d \leq D$ ), takes off at a new departure time (CTD') and flies slower than initially planned, as shown in Figure 2(b). In this way, it will take  $T_{V_{eq}}$  minutes to reach the destination airport, in such a way that  $d + T_{V_{eq}} = D + T_{V_0}$  (i.e. the aircraft is arriving at the same CTA as in the baseline scenario).

It should be noted that the aircraft will still experience the imposed GDP delay at the arrival airport, since this delay has been distributed by waiting on ground at the origin airport and by flying slower during the route. Therefore, the fairness aspects due to the assignment of the delay, regarding different aircraft of different companies, are not affected by this speed reduction strategy.

For a particular flight, if the GDP is not cancelled 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). If the GDP is cancelled while the aircraft is still flying, however, some benefits arise.

Since flying at  $V_0$  has the same fuel consumption per distance flown than flying at  $V_{eq}$ , at the moment the GDP is cancelled, airborne controlled aircraft can accelerate to  $V_0$  and recover part of the delay without incurring extra fuel costs over the initially planned fuel cost (see Fig. 2(c)). In the research presented in this

paper, it is considered that the application of this airborne delay strategy is decided at a pre-tactical phase, and it is assumed that the airline maintains its initial computation of block fuel as its target fuel for the flight. The fact that a CTA has been issued and that some delay can be potentially recovered, might vary the amount of fuel the airline considers adequate for the flight. However, by maintaining the fuel consumption as initially planned this paper presents a conservative approach on the amount of delay that can be recovered. Other strategies, out of the scope of this paper, include the realisation of airborne delay at other speeds rather than  $V_{eq}$  to save fuel or perform even more airborne delay and the recovery of delay by increasing the cruise speed above this nominal speed at the expense of more fuel consumption than initially planned (as studied, for instance in (Cook et al., 2009) ).

With the current concept of operations (baseline scenario), where an aircraft absorbs the total amount of assigned delay on ground, delay recovery can only be done by speeding up over  $V_0$  leading to more fuel consumption for that trip than initially planned. It is worth mentioning that speed reduction strategies are difficult to implement with the current concept of operations since CTA are still not enforced. Thus, even with the current baseline scenario, some companies may decide to accelerate their delayed aircraft, trying to recover part of the delay previously performed on ground (incurring higher fuel costs) and not meeting the assigned arrival slot (Knorr et al., 2011). Nevertheless, in the near future, more accurate control of the trajectory will be available to aircraft operators in the context of SESAR<sup>4</sup> and NextGen<sup>5</sup> projects. Thus, it is expected that CTAs will be effectively enforced on aircraft and the speed reduction strategy at a pre-tactical level as proposed in this paper will be useful.

### 3. Simulations

We have simulated the 24th-25th August 2005 inbound flights to SFO, subject to different GDP scenarios, by 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. This section gives more details on the data used and assumptions made for the different simulations.

#### 3.1. Simulated traffic

The August 24th-25th, 2005 Enhanced Traffic Management System (ETMS) data was used to generate traffic information required to perform the simulations. A total of 1,011 flights were simulated to generate the demand. Accurate cruise performances have been obtained by using the Airbus aircraft databases from the PEP suite. As only the Airbus family performances were available, aircraft were grouped into six different families, corresponding to six different Airbus aircraft models: A300, A320, A321, A330 and A340. The families of aircraft types were created based on the performances of the aircraft, in such a way that all aircraft in the same family had similar performances.

Table 2 shows this grouping: 725 flights are simulated with Airbus performance, representing 71.7% of the total traffic. The 28.3% remaining aircraft are not considered for the speed reduction strategy either because they were 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 aircraft, however, are simulated to correctly represent the demand at the airport, but are excluded from the speed reduction strategy. If any of those flights has some assigned GDP delay it will be done completely on the ground, as in the current concept of operations. During the simulation of the ground delay programs, only international (non-Canadian) and aircraft already flying are exempt from serving delay. No radius of exemption has been considered to study the effect of the  $V_{eq}$  strategy. Aircraft that are already flying when the simulation starts are kept on their original aircraft type when simulated with FACET as it is not necessary to know their accurate cruise performance, since they are exempt from the GDP.

As stated in section 2.2, the equivalent speed depends on the chosen Cost Index and the payload mass of the aircraft. A Cost Index of 60 kg/min has been used for all the flights except for the A330 and A340

---

<sup>4</sup><http://www.sesar.eu>

<sup>5</sup> <http://www.faa.gov/nextgen>

Table 2: Number of aircraft simulated and grouping according to equivalent Airbus types

| Aircraft Family                                | Aircraft Types  | Absolute number of flights | Relative number of flights |
|--|---|----------------------------|----------------------------|
| A300   | A300, A310  | 2                          | 0.2 %                      |
| A319   | A319, B727, B737-200, B737-300, B737-500, DC-9, MD-90, E-145, CRJ-200, CRJ-700, CRJ-900 | 289                        | 28.6 %                     |
| A320   | A320, B737-400, B737-800, B737-900, MD-80   | 193                        | 19.1 %                     |
| A321   | A321, B757  | 147                        | 14.5 %                     |
| A330   | A330, B767, B777, DC-10   | 77                         | 7.6 %                      |
| A340   | A340, B747  | 17                         | 1.7 %                      |
| Total Airbus-like aircraft simulated           |   | 725                        | 71.7 %                     |
| Aircraft without equivalence or already flying |   | 286                        | 28.3 %                     |
| Total of simulated aircraft                    |   | 1,011                      | 100 %                      |

families where a Cost Index of 120 kg/min has been selected. These values are representative for common operations (Airbus, 1998). Finally, to estimate the payload, an 80% of passenger load factor has been assumed for A320 and A319 flights, while for long haul flights (A300, A330 and A340) 80% of the total payload has been assumed (including also freight) (ELFAA - European Low Fares Association Members, 2008).

### 3.2. Simulated ground delay programs

Different scenarios in airport arrival capacity, or arrival acceptance rate (AAR), reflect in most cases well-identified weather patterns in the regions where the airports are located (Liu et al., 2008). 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, where the capacity remains at reduced values throughout the day. Finally, some reductions on the airport arrival acceptance rate are produced due to the rainy periods in the winter season.

In this paper, we have analysed all 130 GDPs that occurred in SFO during 2006 (CDM archival database). By using a  $K$ -means clustering algorithm (Macqueen, 1967), the GDPs have been grouped into three different categories. To cluster the GDPs, the Euclidean distance has been used and they have been characterised by their filed time, starting time, planned end time and actual cancellation time. The number of clusters has been determined by the silhouette coefficient with an iterative process from two to eight clusters. The best clustering has been obtained with three clusters. The centroids of the three resulting clusters are shown in Table 3. 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 cancelled when the weather clears, which on average, is around 3h25 before initially planned. The second group are *All-day* GDPs that also defined early morning, but expand during the whole day because the meteorological conditions do not improve. Finally, the third category of GDPs, correspond to *Afternoon* GDPs. GDPs of the first category are found during the whole year, whilst the GDPs of the second and third category are mainly declared only during the winter season. The duration and cancellation time that have been obtained for the centroids of the clusters of the GDPs are consistent with the values from (Cook and Wood, 2010). Moreover, this clustering is in line with the results presented in (Liu et al., 2008), where airports were characterised by their AAR during the day. It should be noted, however, that in our clustering we have not used AAR data but times related to the GDP definition and cancellation.

The centroids of the three GDP clusters are considered representative of their category and were used to create the three different GDPs simulated in this paper. In these simulations, it has been considered that in nominal conditions (and when the GDP is cancelled), the AAR is 60 aircraft per hour, while during the



GDP the capacity is reduced to 30 aircraft per hour; as the two parallel arrival runways of SFO cannot be independently operated when the visibility is reduced (Janic, 2008). It is assumed that once the GDP is cancelled, the capacity at the airport is unconstrained. Actually, this is not always true since the GDP has shifted the demand and, at the cancellation time, the forecast arrival demand at the airport might occasionally exceed the airport new arrival capacity. Moreover, the maximum delay that could be recovered has been computed assuming that the aircraft that are, at the cancellation time, delayed on ground can immediately take-off and that the airborne aircraft, which are flying at  $V_{eq}$ , can speed up immediately to  $V_0$ .

Table 3: Cluster centroids for the 2006 SFO ground delay programs (hours in local time)

| ID | GDP group             | Number of GDPs | Filed time | Starting time | Planned ending time | Cancellation time |
|----|-----------------------|----------------|------------|---------------|---------------------|-------------------|
| 1  | <i>Morning</i> GDPs   | 91             | 6h31       | 8h59          | 13h55               | 11h30             |
| 2  | <i>All-day</i> GDPs   | 24             | 6h12       | 8h58          | 22h32               | 20h08             |
| 3  | <i>Afternoon</i> GDPs | 15             | 15h42      | 17h08         | 22h51               | 21h06             |

### 3.3. Architecture

As explained in section 3.1, ETMS traffic has been analysed and modified for our simulations. Figure 3 shows the process followed to compute the initial traffic and the nominal parameters of the flights. In order to compute the initial traffic, the aircraft types have been replaced by Airbus aircraft when applicable. For these flights, the trip distances from their origin airport to SFO have been determined. For this purpose, the flight plan of each flight, as defined in the original traffic file, has been considered. Therefore, the distance between two airports might be different for two different 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 have been computed: initial cruise weight, cruise flight level(s) and speed(s) with the required cruise steps if needed<sup>6</sup>.



Figure 3: Nominal traffic generation

The initial traffic has been simulated twice, as depicted in the diagram of figure 4. In the first simulation the speed and flight levels of the aircraft have been kept to their nominal values. The result of this simulation is the initial arrival demand at the airport. In the second simulation, the aircraft reduce the cruise speed to  $V_{eq}$ . The second simulation represents the demand at the airport if all the aircraft fly at their equivalent speed. By doing a comparison of the arrival times, it is possible to compute the maximum airborne delay that each aircraft contributes without incurring in extra fuel consumption.

FACET uses the Base of Aircraft Data (BADA) database (Eurocontrol Experimental Centre, 2011) to compute the performances of the different aircraft. However, it was necessary to accurately control the speed of the aircraft during the cruise. As the diagram of figure 5(a) shows, the climb and descent are simulated directly by FACET. But in the nominal simulation, at the beginning of the cruise the flight

<sup>6</sup>As long as the aircraft burns fuel and loses weight, the optimal flight altitude increases. Therefore, as function of the aircraft type and payload, there exist a certain distance where it becomes optimal to perform a step climb and change the cruise altitude.




Figure 4: Generation of the demand at SFO

level ( $FL_0$ ), speed ( $V_0$ ) and weight ( $W_0$ ) are initialised with the parameters from the PEP computations. These values will be kept constant during the cruise and updated only when a change of cruise altitude is needed according to the nominal flight plan. At each iteration of the simulation the fuel flow is computed according to the Airbus performances of the aircraft. Recalling to section 2.2.1, the equivalent speed varies with the weight, therefore in the simulation of the reduced speed at each simulation step (one minute) the equivalent speed is recomputed for all the airborne flights considering their current weight (see figure 5(b)). If a particular aircraft had a change in cruise altitude in the nominal flight, it will also be performed in the second simulation.

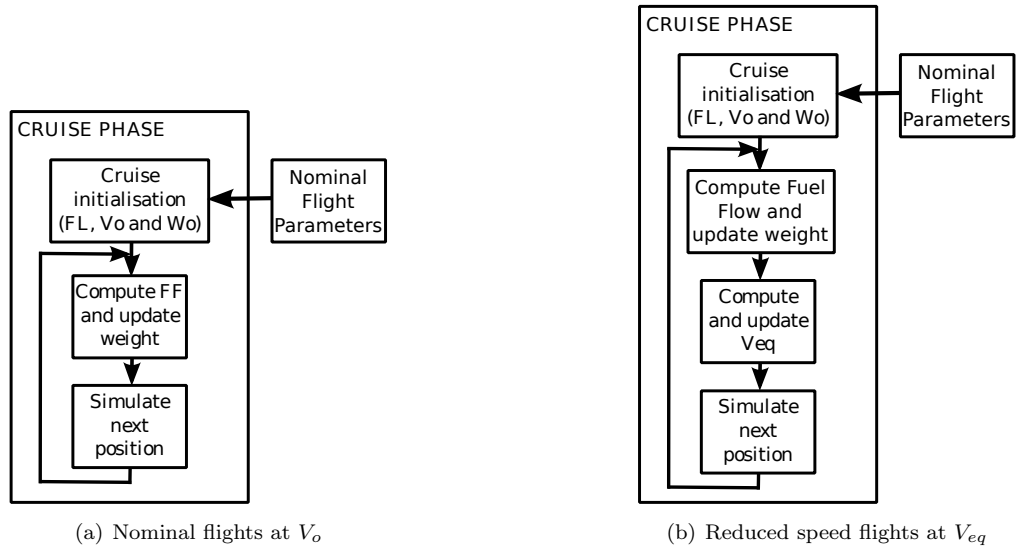


Figure 5: Diagram of the simulations

The application of the GDP to the initial arrival demand, in order to keep the demand below the airport capacity, will result in the amount of delay assigned to each aircraft. This delay assignment is done following a ration by schedule policy. The time when the capacity of the airport is changed from 30 to 60 aircraft per hour is computed in order to finish the regulation according to the centroids. Having previously computed the maximum airborne delay that each flight can perform by flying at  $V_{eq}$ , the assigned delay is divided into ground delay and airborne delay. In the case that a particular flight has been assigned a delay smaller than the maximum airborne delay it can do by flying at  $V_{eq}$ , a new speed (between  $V_{eq}$  and  $V_0$ ) is selected. With this new speed, the CTA is fulfilled and consequently, all the assigned delay is done in the air while saving some fuel with respect to the nominal situation. This process is represented in figure 6.

Finally, if a GDP is cancelled before its defined ending time some delay might be recovered. In this paper

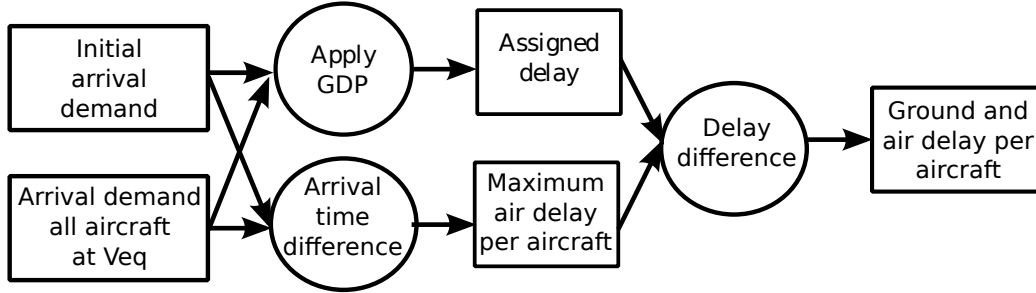


Figure 6: Delay generation

the authors have computed the amount of delay that could be saved by assuming that when the GDP is cancelled the aircraft which are doing ground delay can take off at that moment and the aircraft which are flying at  $V_{eq}$  can speed up to their nominal speed ( $V_0$ ).

#### 4. Results

The results of the application of the three simulated GDPs are presented in Table 4, showing the total, the maximum and average delays assigned. As commented previously, in the speed reduction scenario, part of this total assigned delay can be realised airborne. Table 5 presents these division of the delay, between airborne and ground, in absolute and relative values. The amount of delay that can be airborne varies between 15.7% (*Morning* GDP) and 47.9% (*Afternoon* GDP) with respect the whole assigned delay. In the *Afternoon* GDP the average assigned delay is usually smaller than in the other GDPs, and since the amount of airborne delay that an aircraft can realise is usually small, the percentage of air delay assigned is larger than in the other scenarios. More than 71% of all the aircraft with assigned delay can realise part of it airborne. However, only 9.8% of the total traffic in the *Morning* GDP can perform all the delay assigned airborne, while 38.5% of the traffic in the *Afternoon* is in this situation. The main reason for the difference, is that *Afternoon* GDPs have smaller average delays for each flight (as seen in Table 4) and maximum airborne delays can reach up to 20 min in the best case as was reported in previous publications (Delgado and Prats, 2011, 2012).

Table 4: Results of the application of the simulated GDPs

| ID | GDP group      | Number of affected aircraft | Total assigned delay (min) | Maximum delay assigned (min) | Average delay assigned (min) |
|----|----------------|-----------------------------|----------------------------|------------------------------|------------------------------|
| 1  | Morning GDPs   | 117                         | 4,798                      | 60                           | 41.0                         |
| 2  | All-day GDPs   | 347                         | 9,363                      | 65                           | 27.0                         |
| 3  | Afternoon GDPs | 116                         | 1,615                      | 29                           | 13.9                         |

According to Table 3, the *Morning* GDPs (group 1) and the *Afternoon* GDPs (group 3) have a similar duration. As a consequence, the amount of airborne delay that can be realised is also similar (see Table 5). However, the total amount of delay is higher in the *Morning* GDP due to the fact that the arrival demand is greater. Thus, in the *Afternoon* GDPs almost half of the delay can be realised airborne, while in the *Morning* GDPs only 15.7% of the delay can be absorbed during the cruise. Notice that all the air delay is realised with the same amount of fuel consumption as in the baseline scenario since aircraft are flying at  $V_{eq}$  (see section 2).

Table 5: Division between ground and airborne delay for the simulated GDPs

| ID | GDP group      | Total ground delay (min) | Total airborne delay (min) | Airborne delay | Aircraft realizing airborne delay <sup>†</sup> | Aircraft realizing only airborne delay <sup>*</sup> |
|----|----------------|--------------------------|----------------------------|----------------|--|---|
| 1  | Morning GDPs   | 4,046                    | 752                        | 15.7%          | 71.3%  | 9.8%  |
| 2  | All-day GDPs   | 6,997                    | 2,366                      | 25.3%          | 76.8%  | 26.4%   |
| 3  | Afternoon GDPs | 842                      | 773                        | 47.9%          | 82.0%  | 38.5%   |

<sup>†</sup> Percentage over the the total number of aircraft.

<sup>\*</sup> Percentage over the total number of aircraft doing airborne delay.

#### 4.1. GDP cancellation

The delay that could be saved if the GDP is cancelled has been computed for the three different scenarios. Figure 7 shows the results of these computations. Notice the two different scales in the y-axis. As a function of the time, it represents the accrued delay realised by all the aircraft. The recovered delay achieved if the GDP is cancelled at each time is also presented for both scenarios. In the baseline scenario, where all the delay is performed on ground, the recovered delay can only be the delay that has not been accrued yet. For a given flight, see figure 2(a), the delay recovered will be all the initially assigned delay if the cancellation time is before the flight’s ETD. It will be the difference between the cancellation time and the CTD if the GDP cancels between the ETD and the CTD. If the flight has already taken off, however, no delay will be recovered for that flight since we are assuming that the flight will cruise at  $V_0$ . With the speed reduction strategy, the recovered delay is increased by the time that can be gained by speeding up to  $V_0$  (i.e. not using extra fuel) for the aircraft that are already flying at  $V_{eq}$  when the GDP is cancelled. The extra delay being recovered due to the speed reduction strategy is the difference between these recovered delays and is presented in figure 7.

At the beginning of the GDP, none of the delay has been accrued and therefore, if the GDP is cancelled all the delay can be recovered. As time advances, more delay has already been realised and therefore less delay can be saved if the GDP is cancelled. The benefit of the speed reduction strategy applied to the GDP programs depends on when the GDP is actually cancelled. If the GDP is finally not cancelled before initially planned, this strategy will lead to a change on where all the assigned delay is realised. In addition, the plots in figure 7 show the filed, start, ending and actual cancellation times of each GDP according to Table 3.

The amount of extra delay that will be recovered with respect the baseline scenario depends on the number of aircraft that are at that time in the air realizing airborne delay. Figure 8 shows this dependency. There is a correlation between the number of aircraft in the air flying at  $V_{eq}$  and the extra savings of delay if the GDP is cancelled. The curve showing the number of aircraft is indeed shifted to the right: when the aircraft are flying at  $V_{eq}$ , the later the GDP is cancelled the more delay is already realised and the smaller is the distance available to recover delay. For this reason, close to the GDPs ending time, there are aircraft in the air that are doing airborne delay but there is no extra delay recovered, in this case those aircraft are already on their descending phase and have already realised the whole assigned delay. Table 6 shows the maximum number of aircraft that are at the same time in the air doing the speed reduction strategy and the maximum extra delay recovered with respect the baseline scenario for the three GDPs.

Even if the aggregated total amount of airborne delay that can be realised in the three GDPs is very different, the maximum extra delay recovered with the speed reduction strategy is very similar for the cases, around 430 min. The extra time that can be recovered depends on the number of aircraft that are flying at  $V_{eq}$  at the cancellation time, which it is very similar among the three GDPs.

For the *Morning* and the *Afternoon* GDPs there is a maximum value of the extra delay that can be recovered (around 9h00 and 18h30 respectively) and then, it decreases until the end of the GDP. Conversely, for the *All-day* GDP, we observe two peaks of extra delay recovered: one in the morning and the other in the afternoon (see figures 7 and 8). This result shows the dependency between the amount of extra delay that is recovered and the actual demand at the airport. The maximum extra delay recovered is achieved

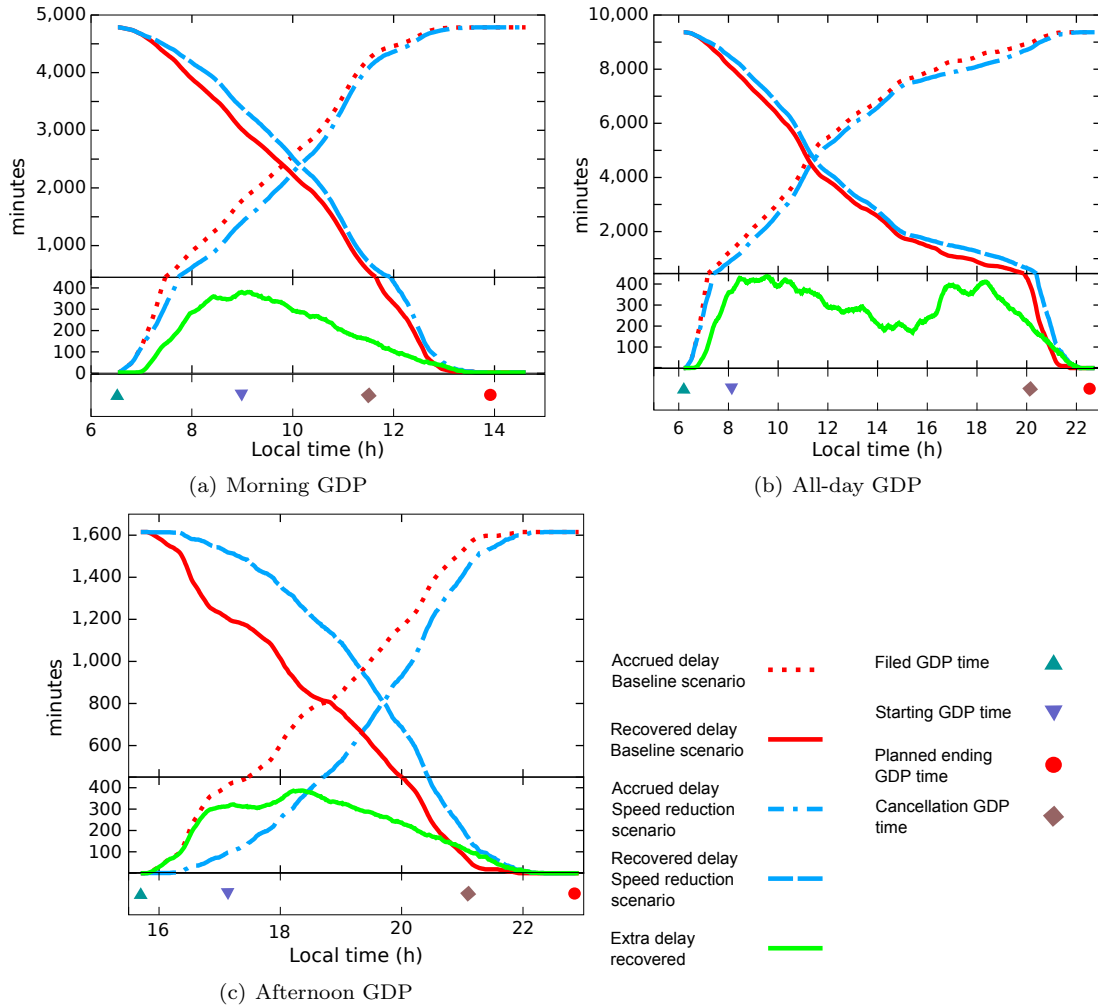


Figure 7: Delay incurred and recovered for the baseline and speed reduction scenarios

before the demand at the airport attains its maxima which is around 10h00 in the morning and 21h00 in the afternoon.

As stated before, the simulations are computed without considering a radius of exemption, by doing so, the benefits of the strategy suggested in this paper are maximised. Figure 9 presents the average delay assigned for each flight as a function of its flight plan distance for all the GDPs studied grouped by 500 NM. It also presents the average ground and airborne delay per aircraft. It is possible to observe that the delay assigned per aircraft is independent on the flight plan distance as it is based only on the schedule of arrival. The maximum airborne delay realisable increases with the distance available for the flight, as more airborne delay can be realised for longer flights, and consequently the ground delay realised decreases as the flight plan distance increases. With a radius of exemption the aircraft that are further from the airport would be excluded. However, this strategy might be still interesting as the airborne delay recovered depends on the number of aircraft in the air flying at  $V_{eq}$  at the cancellation time.

#### 4.1.1. Results at the actual cancellation time

As presented in (Mukherjee et al., 2012; Cook and Wood, 2010), a probability distribution function of fog clearance time can be computed for SFO. This function could be used to compute the average extra delay recovered from the results presented in the previous section. However, in this paper, as the GDPs have

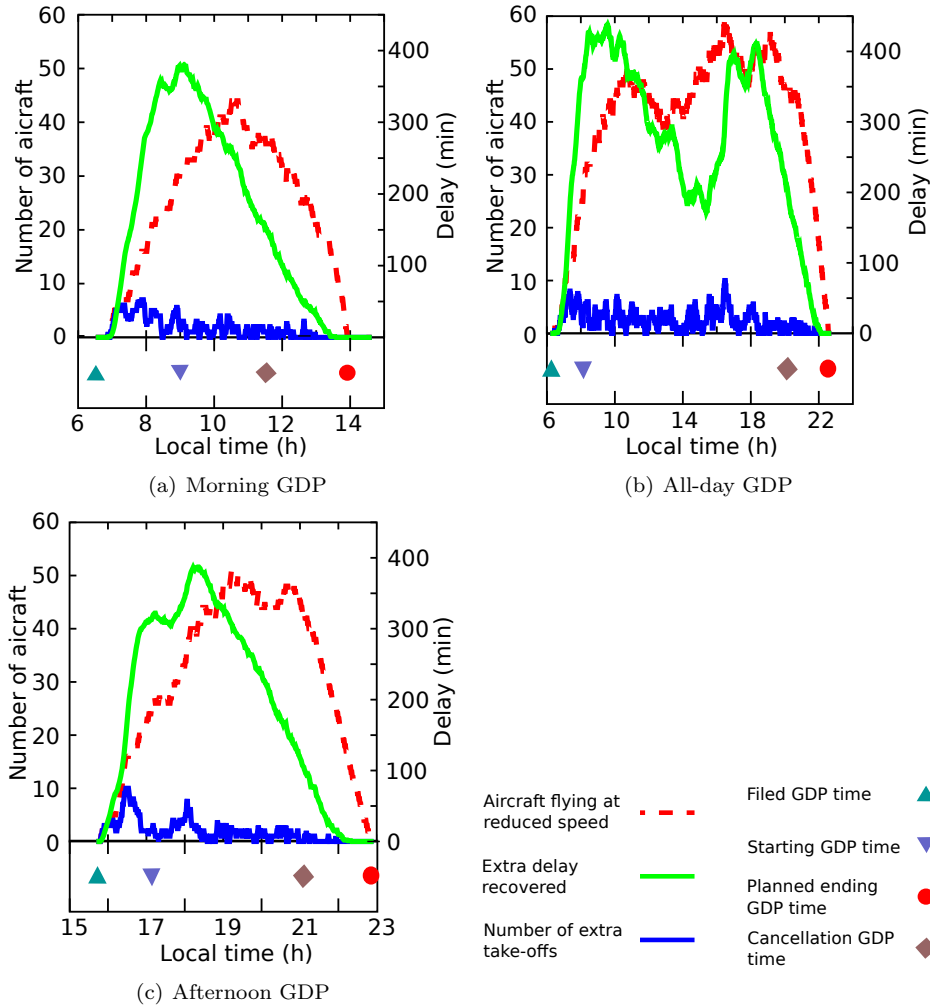


Figure 8: Extra delay recovered, aircraft flying at  $V_{eq}$  and extra take offs

been clustered, the cancellation time according to those centroids (presented in Table 3) is used in order to compute the delay that could be extra saved. This is interesting as not all the GDPs defined in SFO in 2006 where due to low ceiling clouds, the clustering represents this variety in GDP causes.

Table 7 shows the extra delay that is recovered with respect the baseline scenario at the actual cancellation time. In the *Morning* GDP, 155 min of extra delay are recovered, representing 27.6% of the total delay that can be saved in the baseline case. This percentage is increased for the GDPs that are forecast for all the day to 52.5% (208 min) and becomes very significant for the *Afternoon* GDP, with 172.1% of extra delay recovery (105 min).

This extra delay has been saved by the aircraft that were at that moment flying at  $V_{eq}$ . Therefore by dividing the extra amount of delay saved by the number of aircraft that were at that time doing airborne delay it is possible to obtain an average recovered delay per aircraft. In the *Morning* and *All-day* GDP the average delay recovered by aircraft flying at the cancellation time at  $V_{eq}$  is around 4.5 min while in the *Afternoon* GDP this value is reduced to 2.5 min. This recovered delay would be done at no extra fuel consumption as the flights only use  $V_{eq}$  and  $V_0$  speeds (see section 2).

At aggregate level, if we multiply the extra delay recovered with the speed reduction strategy for each GDP centroid by the number of GDPs present in each cluster, we obtain a total of 20,672 min of extra delay

Table 6: Maximum number of aircraft flying and doing airborne delay at the same time, number of extra take-offs and extra delay recovered

| ID | GDP group      | Maximum number of aircraft realizing airborne delay at same time | Maximum number of extra take offs | Maximum extra delay recovered <sup>†</sup> (min) |
|----|----------------|--|-----------------------------------|--|
| 1  | Morning GDPs   | 45   | 7                                 | 379  |
| 2  | All-day GDPs   | 60   | 10                                | 437  |
| 3  | Afternoon GDPs | 51   | 10                                | 387  |

<sup>†</sup> With respect the baseline scenario.

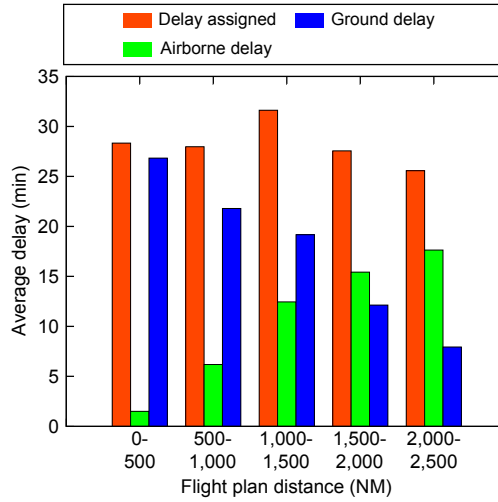


Figure 9: Average delay assigned, ground delay and airborne delay realised as a function of the flight plan distance for all the GDPs studied

recovery in a year, with contributions of: 14,105 min for the *Morning* GDPs, 4,992 min for the *All-day* long GDPs and 1,575 min for the *Afternoon* GDPs.

#### 4.2. ATM effect

In order to do a first assessment of the impact of this strategy on the air traffic system, the number of extra aircraft in the air with the speed reduction technique has been computed.

The use of the speed reduction strategy will lead to more aircraft in the air because the aircraft will be flying slower to realise part of the delay in the air. Figure 8 shows the number of aircraft that are flying at  $V_{eq}$  at every moment during the simulations and the number of extra take-offs, if compared with the baseline scenario. As shown in these figures (and also in Table 6), the number of aircraft that are in the air flying at  $V_{eq}$  varies with the GDP and can be at its maximum value between 45 and 60 aircraft. However, this does not mean that there are 60 extra aircraft in the air that otherwise (in the baseline scenario) would be on the ground.

A flight facing airborne delay and the same flight incurring all the assigned delay as ground delay have the same CTA at the destination airport. The distance between the position of the flight doing airborne delay and the position of the flight at the same time if flying at nominal speed gets reduced as the flight progresses. At the beginning of the descent, the aircraft doing airborne delay will be at the same position as where it would be if flying in the baseline scenario, therefore the impact on this strategy on the TMAs is small. There will be an overlap between the time when both aircraft are in the air. And almost all the time both flights will be airborne at the same time. This leads to the fact that the number of extra take-offs, i.e.

Table 7: Results of the simulated GDPs at the actual cancellation time

| ID | GDP group | Aircraft flying at $V_{eq}$ | Delay saved in the baseline scenario (min) | Delay saved in the speed reduction scenario (min) | Extra delay recovered (min) | % Extra delay recovered <sup>†</sup> | Average extra delay recovered* (min) |
|----|-----------|-----------------------------|--|---|-----------------------------|--------------------------------------|--------------------------------------|
| 1  | Morning   | 37                          | 562  | 717   | 155                         | 27.6%                                | 4.2                                  |
| 2  | All-day   | 45                          | 396  | 604   | 208                         | 52.5%                                | 4.6                                  |
| 3  | Afternoon | 42                          | 61   | 166   | 105                         | 172.1%                               | 2.5                                  |

<sup>†</sup> With respect the delay saved in the baseline scenario.

\* By aircraft recovering part of its delay by speeding up to  $V_0$ , considering that all the aircraft are still in their cruise phase.

the number of aircraft that are in the air with the speed reduction technique that would be on ground in the baseline scenario, is very small (as shown in Figure 8 and Table 6). During the entire simulations there are less than 10 aircraft in the air that otherwise would be on the ground. The amount of airborne delay is high enough to make this strategy interesting but small enough to not increase significantly the number of aircraft in the air and therefore not resulting in a significant increase in the number of aircraft controlled by the ATC.

If more than one GDP is defined at the same time, the total number of flight applying this technique might be higher, but due to the natural spread of the airports and the limited impact on the ATM of an aircraft realising this strategy the effect is limited.

If instead of cancelling earlier a GDP, it is extended, the number of aircraft excluded because they are airborne, when in the baseline scenario would be still on ground, is also very small. Therefore, this strategy does not have a negative impact on the replanning of the GDP initiatives.

Finally, it is worth mentioning that the slow aircraft might interfere with the rest of the traffic due to their abnormally slow cruise speed. This might require some training from the ATC or the need of use offset tracks to avoid the saturation of the nominal tracks.

## 5. Conclusion

The en-route speed reduction strategy presented in this paper has proven useful to recover delay if an air traffic flow management (ATFM) initiative is cancelled before planned (as is usually the case). Simulations were performed using the Future Air Traffic Management Concepts Evaluation Tool (FACET), along with performance data from the Airbus Performance Engineer’s Program suite which allowed us to obtain accurate data of specific range and fuel consumption. In order to simulate representative and realistic ATFM scenarios, all the ground delay programs (GDPs) from 2006 at San Francisco International Airport (SFO) have been characterised by using a  $K$ -means algorithm, obtaining three different types of GDPs: *Morning* GDPs, *All-day* long GDPs and *Afternoon* GDPs. The *Morning* GDPs have an average planned duration of 4h56 and are cancelled in average 2h25 before planned, the *All-day* GDPs are planned for 13h34 and cancelled 2h24 before scheduled and finally, the *Afternoon* GDPs have an average duration of 5h43 and are cancelled around 1h45 before planned.

As reported earlier by the authors, the amount of airborne delay that can be performed for an individual flight using the suggested strategy is not very high. Supposing a scenario where operators have enhanced control of their 4-D trajectories, this strategy is feasible and can complement and make more efficient current ATFM initiatives and do not increase the complexity of the airspace management. Only few aircraft are in the air that would be otherwise still on ground (less than 10 in the presented study). Simulation results show values of extra delay recovery (specially at aggregate level) that encourage further research on the proposed strategy. Around 20,000 min of extra delay could have been saved in 2006 only for SFO GDP initiatives. An average of 39 aircraft get benefit of this strategy by recovering around 4 min of extra delay on each GDP, and without using more fuel than the initially planned for the flight. It is worth noticing that this values



have been calculated assuming no radius of exemption is applied. Thus, the effects of unrecoverable delay should be studied in further research. However, regardless of the radius of exemption, the suggested strategy will add some extra delay recovering to all the aircraft flying at  $V_{eq}$  at the cancellation time of the GDP.

Logically, if a GDP is cancelled as predicted when filed, the baseline scenario and the speed reduction scenario will lead to the same amount of delay and fuel burned. The only difference is that in the baseline scenario all the delay will be on the ground, at the origin airports, while in the other scenario some of it will be realised airborne, although costs other than fuel should also be considered. Moreover, even if the GDP cancels before planned, it is very unlikely that it will be cancelled a short time after its implementation. Therefore, it is very likely that the initial demand will have to perform all the assigned delay, and no benefits may arise from the cruise speed reduction strategy for these flights. A possible idea, for further research, is to start the GDP applying the classical ground delay strategy and, from a certain moment, change to the speed reduction strategy as long as the GDP cancellation time becomes more probable.

It also seems interesting to formulate a ration by schedule policy at the beginning of the GDP (minimising the amount of delay assigned when the GDP is very unlikely to be cancelled) and transition to another policy, which maximise the potentially recovered delay if the GDP finally cancels. In this case, the effect on the total delay assigned should be analysed, as a trade-off will exist between the delay that can be potentially extra recovered and the total delay assigned.

With the simulations, accurate results have been obtained and presented in this paper. However, it might be possible to analytically relate the flight plan distance with the maximum airborne delay realisable. As further research, this could be used to analyse the GDPs, minimising the number of simulations needed.

## Acknowledgments

The authors thank Airbus for the use of PEP program, which allowed us to undertake realistic aircraft performances simulations. We would also like to thank the Aviation System Division from NASA Ames, specially Dr. Hok Kwan Ng and Mr. Alex Morando and Dr. Avijit Mukherjee. Finally, we would like to acknowledge Dr. Avijit Mukherjee, Dr. Tasos Nikoleris and Dr. Mark Hansen for their comments on the paper and P. Slocombe for helping with English corrections.

## References

- Airbus, 1998. Getting to grips with the cost index. Tech. Rep. II, Blagnac, France.
- Airservices Australia, 2008. Annual Report: 2007-2008. Tech. rep.
- Ball, M., Lulli, G., 2004. Ground delay programs: Optimizing over the included flight set based on distance. *Air Traffic Control Quarterly* 12 (1), 1–25.  
URL <http://www.nextor.org/pubs/Ball-Lulli.pdf>
- Ball, M. O., Hoffman, R., Mukherjee, A., 2009. Ground Delay Program Planning Under Uncertainty Based on the Ration-by-Distance Principle. *Transportation Science* 44 (1), 1–14.  
URL <http://transci.journal.informs.org/cgi/doi/10.1287/trsc.1090.0289>
- Bertsimas, D. J., Patterson, S. S., 1998. The air traffic flow management problem with enroute capacities. *Operations Research* 46 (3), 406–422.
- Bertsimas, D. J., Patterson, S. S., 2000. The traffic flow management rerouting problem in air traffic control: a dynamic network flow approach. *Transportation Science* 34 (3), 239–255.
- Bilimoria, K. D., Sridhar, B., Chatterji, G. B., Sheth, K. S., Grabbe, S., 2000. FACET: Future ATM Concepts Evaluation Tool. *Air Traffic Control Quarterly* 9 (1), 1–20.  
URL [http://d.wanfangdata.com.cn/NSTLQK\\_NSTL\\_QK5018044.aspx](http://d.wanfangdata.com.cn/NSTLQK_NSTL_QK5018044.aspx)

- Carrier, S., de Lépinay, I., Husatche, J.-C., Jelinek, F., Jul. 2007. Environmental impact of air traffic flow management delays. In: 7th USA/Europe Air Traffic Management Research and Development Seminar (ATM2007). Barcelona, pp. 1–13.
- Chaloulos, G., Crück, E., Lygeros, J., Dec. 2010. A simulation based study of subliminal control for air traffic management. *Transportation Research Part C: Emerging Technologies* 18 (6), 963–974.  
URL <http://linkinghub.elsevier.com/retrieve/pii/S0968090X10000288>
- Cook, A. J., Tanner, G., Williams, V., Meise, G., 2009. Dynamic cost indexing: managing airline delay costs. *Journal of Air Transport Management* 15 (1), 26–35.
- Cook, L., Wood, B., 2009. A Model for Determining Ground Delay Program Parameters Using a Probabilistic Forecast of Stratus Clearing. In: Eighth USA/Europe Air Traffic Management R&D Seminar. Napa.  
URL [http://www.atmseminarus.org/seminarContent/seminar8/papers/p\\_125\\_W.pdf](http://www.atmseminarus.org/seminarContent/seminar8/papers/p_125_W.pdf)
- Cook, L. S., Wood, B., 2010. A Model for Determining Ground Delay Program Parameters Using a Probabilistic Forecast of Stratus Clearing. *Air Traffic Control Quarterly* 18 (1).
- Delgado, L., Prats, X., 2011. Simulation of airborne ATFM delay and delay recovery by cruise speed reduction. In: Eurocontrol (Ed.), 1st SESAR Innovation days. Toulouse, France, pp. 1–7.  
URL <http://sesarinnovationdays.eu/files/SIDs/SID2011-20.pdf>
- Delgado, L., Prats, X., 2012. An en-route speed reduction concept for absorbing air traffic flow management delays. *Journal of Aircraft* 49 (1), 214–224.
- Delgado, L., Prats, X., 2013. Effect of wind on operating-cost-based cruise speed reduction for delay absorption. *IEEE Transactions on Intelligent Transportation Systems* 14 (2), 918–927.
- ELFAA - European Low Fares Association Members, 2008. Members’ statistics. Tech. rep.
- Eurocontrol Central Flow Management Unit, Mar. 2011. Basic CFMU Handbook - General & CFMU Systems, 15th Edition.
- Eurocontrol Experimental Centre, 2011. User manual for the Base of Aircraft Data (BADA) revision 3.9. Tech. rep.
- European Aviation Safety Agency, 2011. Certification Specifications and Acceptable Means of Compliance for Large Aeroplanes CS-25, 11th Edition.
- Günther, T., Fricke, H., Jun. 2006. Potential of Speed Control on Flight Efficiency. In: Second International Conference on Research in Air Transportation (ICRAT2006). 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. asdfasdf, The National Center of Excellence for Aviation Operations Research (NEXTOR).
- Janic, M., Dec. 2008. Modelling the capacity of closely-spaced parallel runways using innovative approach procedures. *Transportation Research Part C: Emerging Technologies* 16 (6), 704–730.  
URL <http://linkinghub.elsevier.com/retrieve/pii/S0968090X08000223>
- Knorr, D., Chen, X., Rose, M., Gulding, J., Enaud, P., Hegendoerfer, H., Jun. 2011. Estimating ATM Efficiency Pools in the Descent Phase of Flight. In: Ninth USA/Europe Air Traffic Management Research and Development Seminar (ATM2011). Berlin.  
URL <http://www.atmseminar.org/seminarContent/seminar9/papers/116-Chen-Final-Paper-4-18-11.pdf>
- Kotnyek, B., Richetta, O., May 2006. Equitable Models for the Stochastic Ground-Holding Problem Under Collaborative Decision Making. *Transportation Science* 40 (2), 133–146.  
URL <http://transci.journal.informs.org/cgi/doi/10.1287/trsc.1050.0129>

- Liu, P., Hansen, M., Mukherjee, A., Aug. 2008. Scenario-based air traffic flow management: From theory to practice. *Transportation Research Part B: Methodological* 42 (7-8), 685–702.  
URL <http://linkinghub.elsevier.com/retrieve/pii/S0191261508000106>
- Lowther, M. B., Clarke, J.-P. B., Ren, L., 2008. En Route Speed Optimization for Continuous Descent Arrival. In: *AIAA Guidance Navigation and Control*. Georgia Institute of Technology, pp. 1–19.  
URL [https://smartech.gatech.edu/bitstream/1853/28271/1/lowther\\_marcus\\_b\\_200805\\_mast.pdf](https://smartech.gatech.edu/bitstream/1853/28271/1/lowther_marcus_b_200805_mast.pdf)
- Macqueen, J., 1967. Some methods for classification and analysis. In: University of California Press (Ed.), *Proceedings of 5th Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley, pp. 281–297.
- Mukherjee, A., Hansen, M., Grabbe, S., Aug. 2012. Ground delay program planning under uncertainty in airport capacity. *Transportation Planning and Technology* 35 (6), 611–628.  
URL <http://www.tandfonline.com/doi/abs/10.1080/03081060.2012.710031>
- Prats, X., Hansen, M., Jun. 2011. Green Delay Programs. In: *Ninth USA/Europe Air Traffic Management Research and Development Seminar (ATM2011)*. Berlin.  
URL <http://www.atmseminar.org/seminarContent/seminar9/papers/83-Prats-Final-Paper-4-18-11.pdf>
- Richetta, O., Odoni, A. R., May 1994. Dynamic solution to the ground-holding problem in air traffic control. *Transportation Research Part A: Policy and Practice* 28 (3), 167–185.  
URL <http://linkinghub.elsevier.com/retrieve/pii/0965856494900159>
- Vossen, T., Hoffman, R., Mukherjee, A., 2011. Air Traffic Management. In: Barnhart, C., Smith, B. (Eds.), *Quantitative Problem Solving Methods in the Airline Industry*. Springer, pp. 385–454.