Proceedings of the 2023 Winter Simulation Conference
C. G. Corlu, S. R. Hunter, H. Lam, B. S. Onggo, J. Shortle, and B. Biller, eds.

TACTICAL MINIMIZATION OF THE ENVIRONMENTAL IMPACT OF HOLDING IN THE TERMINAL AIRSPACE AND AN ASSOCIATED ECONOMIC MODEL

Aditya A. Paranjape
Anwesha Basu

Tata Consultancy Services Ltd.
Tata Research, Development and Design Centre
Sahyadri Park 2, Hinjewadi
Pune, MH 411057, INDIA

ABSTRACT

Minimization of the carbon footprint of aviation is an active area of interest to the industry and policy makers alike. Optimization of the individual flight phases is an important step in that direction. This paper considers the holding phase, wherein aircraft hold in the terminal airspace of airports prior to approach and landing during times of busy operation or when the arrival capacity is reduced due to factors such as bad weather. We propose a tactical method to allocate landing slots while minimizing the environmental impact of holds. An environmentally-driven policy can be perceived as unfair, particularly by airlines whose environmentally friendly aircraft which might need to hold longer than they would under a fair first-come-first-served policy. To alleviate this challenge, we propose a number of economic reward schemes, including one based on a linear programming problem obtained by applying complementary slackness to the dual of the assignment problem.

1 INTRODUCTION

It is well known that aviation contributes significantly to the world's carbon footprint, accounting for $2.5-3\%$ of the total CO_2 emissions by some estimates (Soomer and Franx 2008). A major proportion of this footprint is derived from routine operation of aircraft and can be reduced using a combination of scheduling and flight path optimization. A typical flight can be broken down into representative segments such as climb, cruise, hold, etc., each with its characteristic thrust requirement and consequent emissions profile (Stettler et al. 2011). Of these, the environmental impact of the climb and the cruise is the greatest due to the high thrust requirement. The optimization of these phases for a given aircraft and routing constraints is addressed through speed and altitude management (Folse et al. 2016). The pre-arrival hold phase requires a relatively high thrust setting, comparable to cruise, because of the low descent rates. The time spent in hold by an aircraft depends significantly on the traffic and the number of runways available for handling arrivals. As a result, when the traffic is heavy and the average holding times are high, it would be environmentally beneficial to manage the holding time of aircraft as a function of their environmental footprint rather than following a fair first-come-first-served policy or heuristics that directly minimize the average holding time. At the same time, a policy that appears unfair to airlines that operate environmentally friendly aircraft needs to be supported by a fair economic reward. This paper presents a class of methods for determining an environmentally-optimal holding policy and an economic reward concurrently.

1.1 Overview of the Literature

Problems related to tactical scheduling the arrival of aircraft have been usually solved using a combination of heuristics (Beasley et al. 2001; Bäuerle et al. 2007; Zhang et al. 2020) often based around queue-theoretic

979-8-3503-6966-3/23/$31.00 ©2023 IEEE

analysis or optimization techniques based on integer programming, for instance (Ikli et al. 2020; Higasa and Itoh 2022; Murca and Muller 2015). The decision to allocate landing slots can be taken by directly computing target landing times that satisfy the necessary separation constraints (Ding and Valasek 2007), by modeling and optimizing the mean arrival rates in the airspace (Itoh and Mitici 2019; Higasa and Itoh 2022) using an underlying stochastic model for the arrival process, or by managing the holding times for individual aircraft whose arrival into the terminal airspace is modeled as a single queue (Bayen et al. 2004). The objective of the allocation could be to either minimize the deviation from the desired times of arrival (Ding and Valasek 2007) or minimizing metrics such as the arrival time of the last aircraft (Bayen et al. 2004). A related problem is that of the allotting landing slots keeping airlines' preferences in mind: this can be either a strategic or a tactical problem. Towards that end, auction-based methods have been proposed for allocation of landing slots (Le et al. 2003; Soomer and Franx 2008) which allow airlines to specify their cost functions and even bid competitively for slots.

One particular problem of interest in the context of growing relevance of "green aviation" is that of minimizing the net carbon footprint of the arrival process through a combination of trajectory optimization and scheduling (Mesgarpour et al. 2010; Zhang and Fillippone 2022). It is evident how a scheduling method that prioritizes the carbon footprint might appear unfair to airlines that operate environmentally friendlier aircraft (Zhang and Fillippone 2022). As a result, it is essential to support the scheduling scheme with an economic model, such as a carbon or emissions trading schemes proposed for aviation at large (UK Civil Aviation Authority 2013; Scheelhaase and Grimme 2007). While the complexities of setting up a tactical trading scheme (even within the scope of an aviation-wide scheme) are a challenge in their own right, it is rather straight-forward to define and compute rewards or penalties based on engineering metrics such as the fuel consumption or the actual net emissions.

1.2 Contribution

This paper presents a tactical landing slot allocation method for minimizing the environmental impact of aircraft when they are required to hold in the terminal airspace. The net environmental impact of holds can be determined by summing over the impact of individual aircraft which depends on the aircraft type, its engines, as well as its weight. As a result, if a certain number of aircraft arrive at a holding location, it is desirable to allow older, heavier aircraft with a larger environmental footprint to exit the queue before newer aircraft. This requirement conflicts with operational requirements and fairness, in the sense that airlines would prefer that their aircraft be allowed to land in their order of arrival into the terminal airspace. In order to restore fairness, the proposed method includes a calculation of the economic cost of holding, and particularly that of the additional amount of time an environmentally-friendly aircraft might be required to hold. By presenting an economic reward, the apparent unfairness of an environmentally optimal scheme can be mitigated, bringing its fairness on par with typical baseline schemes used by air traffic control.

The rest of the paper is organized as follows. After presenting the necessary preliminaries in Section 2, we present the problem formulation in Section 3. The slot allocation algorithm, including candidate accompanying economic reward mechanisms, are presented in Section 4. Simulation results for a few representative cases are presented in Section 5, followed by the concluding discussion in Section 6.

2 PRELIMINARIES

2.1 Notation: Terminal Airspace

The arrival of an aircraft consists of four phases, as illustrated in Figure 1: descent from cruising altitude, pre-approach hold (if applicable), final approach, and touchdown. When the traffic is sparse, it is possible to bypass the pre-approach hold altogether and proceed to the final approach which starts at the final approach fix (FAF) located approximately 10 miles from the arrival runway.

The number of aircraft that can enter the final approach corridor during a given time window is constrained by the necessity to ensure appropriate inter-aircraft separation (IAS, not to be confused with the indicated airspeed). The minimum permissible inter-aircraft separation depends on the relative size of the two aircraft, the weather, the visibility on the ground, and ground movement constraints at the airport. The pre-approach hold is used to regulate the flow of aircraft into the final approach corridor when the number of aircraft arriving in the terminal airspace exceeds the rate at which aircraft can enter the final approach

The holding pattern involves flying one or more circuits around well-defined individual way points while holding or reducing altitude in prescribed steps. Thus, the holding area is also referred to as the holding stack. As shown in Figure 1, each stack consists of a series of closed paths separated by constant altitude. Each aircraft in the stack is allocated a fixed altitude. An arrival runway is served by a multitude of such stacks. As an example, two of the four stacks used at London Heathrow (LHR) are shown in Figure 2 (UK Civil Aviation Authority 2023).

A final point to be noted is that there are two types of path from any given stack to the FAF. A *direct path* is the shortest path from the stack to the FAF, and involves a steep descent with the minimal thrust setting. The alternative is a longer *waypoint path* (WP path for short), wherein the aircraft flies over a series of prescribed waypoints or through a a prescribed combination of distances, turns, and bearings as shown in Figure 2. Since the aircraft's descent rate is relatively smaller in magnitude along the waypoint path, the thrust required to fly the waypoint path is higher than the direct path.

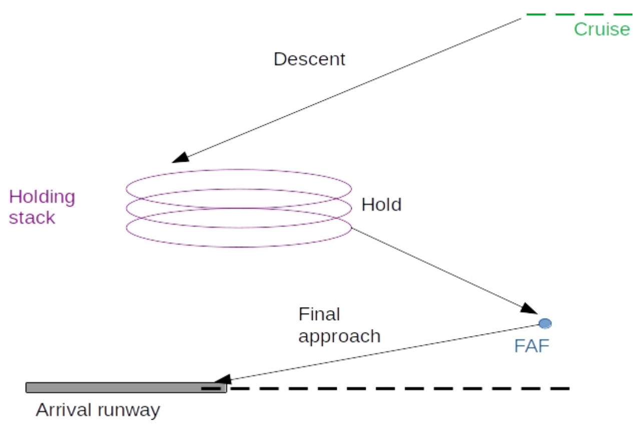

Figure 1: Typical arrival phases.

2.2 Model for Fuel Consumption

We restrict our attention to jet aircraft in this paper, although the extension to propeller-powered aircraft (and even electric aircraft) is straight-forward. The rate of fuel consumption of a jet aircraft is given by

$$\dot{m}_f = \eta_T T \tag{1}$$

where η_t is called the thrust-specific fuel consumption (TSFC) and T is the net thrust produced by the engines. The TSFC is engine-specific while T depends on the weight W and the flight profile as follows:

$$T = W\left(\sin\gamma + \frac{C_D}{C_L}\cos\gamma\right) \tag{2}$$

where C_L/C_D is the lift-to-drag ratio, a measure of the aerodynamic efficiency of the airframe, and γ is the flight path angle defined by $\gamma = \tan^{-1}(dz/dl)$, where z is the altitude of the aircraft and l denotes the path length in the horizontal plane. Combining the equations (1) and (2), we get

$$\dot{m}_f = \eta_T W\left(\sin\gamma + \frac{C_D}{C_L}\cos\gamma\right) \tag{3}$$

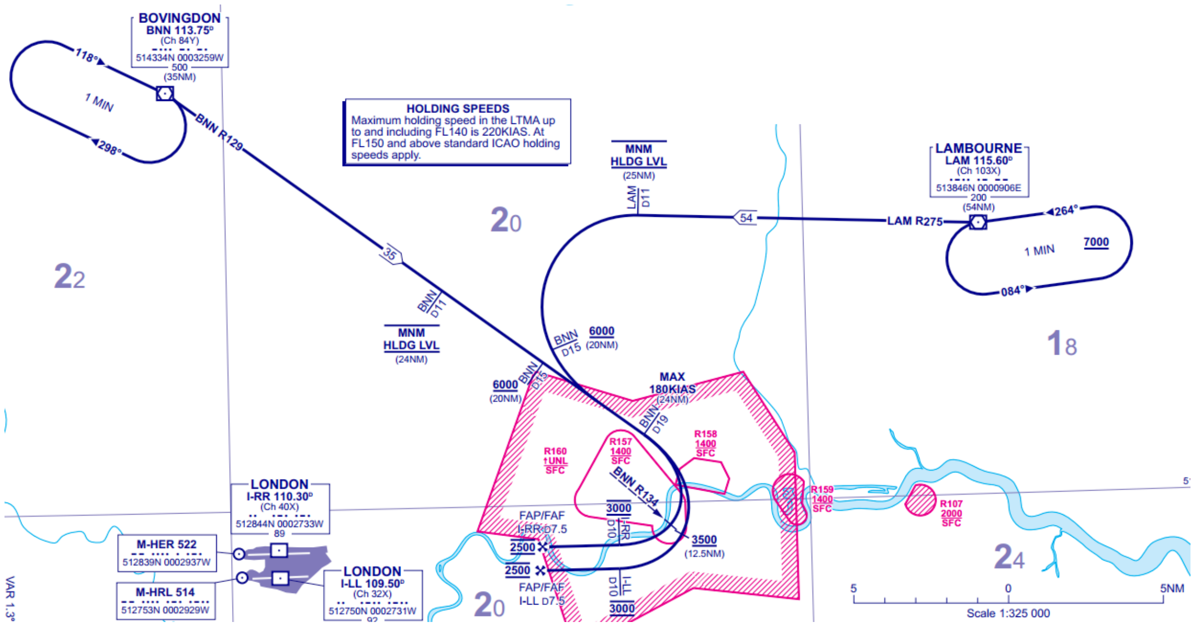

Figure 2: An aeronautical chart showing Bovingdon (left) and Lambourne (right), two of the four stacks that service London Heathrow, together with paths leading to FAFs for runways 27L and 27R.

Assuming that the carbon footprint of an aircraft is proportional to $\dot{m}_f$, it is clear that the carbon footprint is decided by the characteristics of the airframe and the engine via C_D/C_L and η_T, respectively. Alternately, one may estimate it using a combination of data from aircraft and engine manufacturers and model-based prediction (Sun et al. 2018). We provide numerical examples to illustrate the typical fuel burn rate for various classes of commercial jet aircraft in Table 1. Aircraft typically hold and approach at speeds between 70% and 90% of the minimum-drag flight speed, so that the corresponding value of C_L/C_D is approximately $0.9(C_L/C_D)_{\max}$ (Anderson 2016). The examples in Table 1 are calculated using (3) with $\gamma = -0.02$ rad (approximately -1.2 deg). The values of C_L/C_D and TSFC are usually not disclosed by the manufacturers, and we have used estimates from (Martinez-Val et al. 2005).

In a real-world setting, the rate of fuel burn can be obtained from an aircraft in real time (UK Civil Aviation Authority 2013). The fuel burn rate in a holding pattern with gradual descent rates is a predictable fraction of the fuel burn rate at the end of the cruise. This value of the fuel consumption rate and the corresponding altitude can be communicated to the air traffic management system, either automatically or by the pilot, and used in the algorithm presented in the paper.

Table 1: Typical fuel consumption for commercial aircraft. MLW stands for the maximum landing weight. The flight path angle is assumed to be -1.2 deg.

Aircraft	Maiden flight [year]	Typical # passengers	$(L/D)_{\max}$	η_T	$\dot{m}_f/W$ [kg/s/(10^6 N)]	MLW [10^6 N]	$\dot{m}_f$ at MLW [kg/s]
Airbus A320-200 (A320)	1987	160	16.3	9.4	0.407	0.66	0.269
Boeing 747-400 (B747)	1988	380	15.5	9.4	0.437	2.7	1.18
Airbus A330-300 (A333)	1992	280	18.1	9.0	0.335	1.93	0.627
Boeing 777-200ER (B772)	1996	300	19.3	8.16	0.276	2.13	0.588

3 PROBLEM FORMULATION

3.1 Queuing Dynamics in the Holding Stack

Consider an airport which is served by N holding areas (or stacks) $h_1, \ldots, h_N$. Let $q_i[k]$ denote the number of aircraft that would enter the stack i at time instant k. With each aircraft j, we associate the states

$$s_j[k] = (\text{id}, \tau_j^h[k], z_j[k], u_j[k], v_j[k]) \tag{4}$$

where *id* is a unique identifier assigned to the aircraft, $z_j[k]$ denotes the aircraft's altitude, and $\tau_j^h[k]$ denotes the total time spent in the holding stack until time k. The governing equations for τ_j^h and z_j are given by

$$\begin{aligned}
\tau_j^h[k+1] &= \begin{cases} \tau_j^h[k]+1 & \text{aircraft in hold} \\ \tau_j^h[k] & \text{aircraft released from hold} \end{cases} \\
z_j[k+1] &= \begin{cases} z_j[k] & \text{holding slot at } z_j[k]-\Delta z \text{ unavailable and } u_j[k]=v_j[k]=0 \\ z_j[k]-\Delta z & \text{otherwise} \end{cases} \\
u_j[k] &= \begin{cases} 1 & \text{aircraft released from hold via WP path at} k \\ 0 & \text{otherwise} \end{cases} \\
v_j[k] &= \begin{cases} 1 & \text{aircraft released from hold via direct path at} k \\ 0 & \text{otherwise} \end{cases}
\end{aligned} \tag{5}$$

where Δz denotes the difference between successive altitude levels in the holding stack. Furthermore, there typically exist altitude thresholds z^u and z^v such that $u_j[k] = 0$ necessarily if $z_j[k] > z^u$ and likewise for $v_j[k]$. Let $x_i[k]$ denote the number of aircraft in stack h_i at time k. Then, we have that

$$0 \leq x_i[k+1] = x_i[k] + q_i[k] - \sum_{j \in h_i} (u_j[k] + v_j[k]) \tag{6}$$

Let τ_i^c and τ_i^w denote the time to reach the FAF from the holding area i along direct and WP paths, respectively. To ensure that at most one aircraft can arrive at the FAF at an instant k, we get

$$\sum_{h_i} \sum_{j \in h_i} (u_j[k - \tau_i^c] + v_j[k - \tau_i^w]) \leq 1 \tag{7}$$

Equations (5), (6) and (7) completely describe the dynamics of the queue in the holding stack. We do not prescribe any specific dynamics for $q_i[k]$, the number of aircraft entering stack i at time k.

3.2 Optimal Control Problem

The objective of the air traffic control policy for managing the queues is as follows:

1. Minimize the the total amount of fuel burned by the aircraft during the time that it holds in the stack and then when it flies to the FAF. We note that an aircraft may have two routing options between the stack and the FAF.
2. Minimize the net time spent by the the aircraft between the point of entry into the stack and arrival at the FAF. The net holding time represents routine operational metrics (e.g., on-time arrival statistics or the time available for passengers transferring to onward flights) rather than fuel burn.

It is evident from Table 1 how these two objectives can be contrary to each other. As an example, consider the case where a fully loaded B747 arrives at a stack just after a fully loaded Boeing 772. While a fair holding policy would prioritize the B772, an emissions-minimizing policy would prioritize the B747 since its fuel burn rate is just over twice as that of the B772 (see Table 1).

Recall, from (4), that $\tau_j^h[k]$ denotes the cumulative time for which an aircraft j has been in hold in a stack at a time instant k. Let $f_j^h[k]$ denote the cost of holding an aircraft j for another instant (i.e., up to time $k+1$), calculated using (3) in terms of the fuel consumption. Let $f_j^f[k]$ denote the cost of flying the aircraft to the FAF if it were released at time instant k, also calculated using (3). Note that the cost would depend on whether the aircraft opts for a CDA or a WPF path. For completeness, we note that an operational cost of holding, independent of the fuel consumption and dependent on the cumulative delay, can be defined as $f_j^o[k] \triangleq f^o(\tau_j^h[k])$. For brevity, let $u[0:T]$ denote the vector $[u_1[0],\ldots,u_h[0],u_1[1],\ldots,u_h[1],\ldots,u_1[T],\ldots,u_h[T]]$ (see (6), and likewise for $v[0:T]$. The optimal control problem (OCP) of interest is as follows:

$$\min_{u[0:T],v[0:T]} \sum_{k=0}^{T} \sum_{j} (f_j^h[k] + f_j^f[k] + f_j^o[k]) \tag{8}$$

subject to the dynamics (5), (6) and (7). Here T is the terminal time. It is worth noting that the set of aircraft is not fixed at time $t=0$; rather, it time-varying. Thus, $\sum_j$ needs to be carried out over *all* aircraft that arrive at the holding areas during the time window $[0,T]$. Finally, we note that the calculation of f^o is ad hoc (compared to the physically unambiguous definition of f^h and f^f) and we ignore f^o in this paper.

3.3 Economic Model

The traditional control policy used by the ATC involves a combination of first-come-first-served (FCFS) together with a heuristic grouping of similar-sized aircraft to reduce the average holding time. This baseline solution is commonly accepted as *fair* by the airlines. The solution to the OCP (8) need not be fair in this sense, as illustrated by the example in Section 2.2. Thus, in order to ensure fairness, we propose that an economic incentive be provided to the participating airlines as explained in Section 4.3.

4 ASSIGNMENT ALGORITHM

4.1 General Assignment Algorithm

Consider an *ordered* set $\mathscr{O}$ of $m>0$ objects, and suppose that these are to be allocated to n individuals in the set $\mathscr{I}$ such that every individual gets at most one object. Since $\mathscr{O}$ is ordered, we label the elements of $\mathscr{O}$ as $1,\ldots m$, where the magnitude of the label is consistent with the underlying order. Let $a=[a_1,\ldots,a_m]\in\mathbb{N}_\emptyset^m \triangleq (\mathbb{N}\cup\{\emptyset\})^m$ denote the vector of assignments; i.e., $a_j=i$ implies that the $j\in\mathscr{O}$ has been assigned to $i\in\mathscr{I}$, and $a_j=\emptyset$ implies that the object is unassigned. For each individual i, let the feasible assignments (i,j) be captured by $F_i\in\mathbb{R}^m$, where $F_{ij}=1$ if (i,j) is feasible and $F_{ij}=0$ otherwise.

With each individual $i\in\mathscr{I}$, we associate a variable (property) $t_i^{\lim}>0$ (possibly infinite). Likewise, for each object $j\in\mathscr{O}$, we associate a variable $t_j^{\text{stamp}}>0$ which satisfies, additionally, $t_j^{\text{stamp}}>t_k^{\text{stamp}}$ if $j>k$. In the context of this paper, these variables stand for, respectively, the upper limit on an aircraft's holding time and the time stamp of an arrival slot (measured at the FAF).

Assumption 1 (Well-posedness) For all $j\le m$, we assume that $\text{card}\{i \mid t_i^{\lim}\le t_j^{\text{stamp}}\}\le j$. This is necessary for the assignment problem, defined presently, to be well-posed.

Definition 2 (Valuation) An individual $i\in\mathscr{I}$ associates a value of $V_{ij}\in\mathbb{R}$ with every object $j\in\mathscr{O}$. Furthermore, $V_{ij}=-\infty$ if $F_{ij}=0$.

Problem 1 (Assignment problem) We state the assignment problem as follows: determine the assignment $a^*\in\mathbb{N}_\emptyset^m$ such that the following objectives are met

1. Feasibility: $a_j^*=i$ only if $F_{ij}=1$

2. Mandatory assignment requirements (MARs) met: for each $i' \in \mathscr{I}$ satisfying $t_{i'}^{\text{lim}} \leq t_m^{\text{stamp}}$, there exists a unique j' satisfying $F_{i'j'} = 1$, $t_{j'}^{\text{stamp}} \leq t_{i'}^{\text{lim}}$, and $a_{j'} = i'$.
3. Reward maximization: $a^* = \arg\max_a \sum_j V_{a_j,j}$

Note that the mandatory assignment requirement (MAR) can be combined with feasibility by setting $F_{ik} = 0$ if $t_i^{\text{lim}} < t_k^{\text{stamp}}$ (with a strict inequality). However, MARs are technically soft requirements unlike hard feasibility constraints such as the time required to reach the FAF from a holding area. The assignment problem in Problem 1 can be solved efficiently using auction (Bertsekas 1979; Bertsekas 1981) which has the added benefit of calculating an optimal price for each object. In the context of this paper, the optimal price translates into the economic reward given to airlines. The auction algorithm, modified to account for the fact that $m \neq n$ (the number of objects does not equal the number of individuals), is listed in Algorithm 1.

Algorithm 1 Assignment algorithm, called using `Assignment`

Require: set of individuals $\mathscr{I}$ and the set of objects $\mathscr{O}$
Require: value matrix V, feasibility vectors $\{F_i\}_{\mathscr{I}}$, and initial price vector $p = 0$
  Initialize: for all $i \in \mathscr{I}$, $j \in \mathscr{O}$, $F_{ij} = 0$ if $t_i^{\text{lim}} < t_j^{\text{stamp}}$
  Initialize: $\mathscr{U} \leftarrow \mathscr{I}$, $\mathscr{V} \leftarrow \mathscr{O}$
  Initialize: for all j, $a_j = \emptyset$
  {Step 1: assign objects with tight feasibility constraints}
  while $\min_{i \in \mathscr{U}} \sum_j F_{ij} = 1$ **do**
    $k = \arg\min_{i \in \mathscr{U}} (\sum_j F_{ij} = 1)$, $j_k = \arg_j (F_{kj} = 1)$
    Assignment: $a_{j_k} = k$, $p_k = 0$
    Update: $\mathscr{U} \leftarrow \mathscr{U} \backslash \{k\}$, $\mathscr{V} \leftarrow \mathscr{V} \backslash \{j_k\}$
    Remove k^{th} row and j_k^{th} column from V: $V \leftarrow V((-k),(-j_k))$
  end while
  {Step 2: call the auction algorithm if more than one individual is unassigned}
  if card$(\mathscr{U}) > 1$ **then**
    (a,p) = Auction$(\mathscr{U},\mathscr{V},V,p)$ {the auction algorithm can be found in (Bertsekas 1979)}
  else
    $i \in \mathscr{U}$, $j_i = \arg\max_j V_{ij}$, $a_{j_i} = i$
  end if

We extend the assignment algorithm in Algorithm 1 to the case where the feasibility is not just intrinsic (i.e., dependent on the individual-object pair) but also a function of the assignment of the other objects. In particular, recall the mandatory assignment requirement captured by the variable $t_{i \in \mathscr{I}}^{\text{lim}}$. Given the assignment a, we define a simulator function which computes for each $i \in a$ a variable $t_i^{\text{act}} > 0$ (in the context of landing slot allocation, the actual time of arrival of an aircraft at the FAF, given all slot assignments and the inter-aircraft separation requirements). If $t_i^{\text{act}} > t_i^{\text{lim}}$ (with some tolerance), then the assignment $a_j = i$ is deemed infeasible.

Definition 3 (Simulator function) Consider the assignment defined by the vector $a \in \mathbb{N}_{\emptyset}^m$. For each $a_j \neq \emptyset$, the function `Simulator` calculates a variable $t_{a_j}^{\text{act}}$ which depends on a (i.e., in general, all of the assignments). The assignment is feasible if and only if $t_{a_j}^{\text{act}} \leq t_{a_j}^{\text{lim}} + \delta_a$ for all $a_j \neq \emptyset$, where $\delta_a \geq 0$ is a user-defined tolerance.

Algorithm 2 lists the pseudocode for an iterative auction algorithm based on Algorithm 1 and the function `Simulator`.

Algorithm 2 Iterative assignment algorithm, called using `IterativeAssignment`

Require: The set of individuals $\mathscr{I}$ and the set of objects $\mathscr{O}$
Require: Mandatory limits $\{t_i^{\text{lim}}\}_{i\in\mathscr{I}}$, nominal values of $t_{j\in\mathscr{O}}^{\text{stamp}}$
Require: value matrix V, feasibility vectors $F_{i\in\mathscr{I}}$, and initial price vector $p = 0$
  Initialize: converged = 0
  while converged == 0 **do**
    Run Algorithm 1: (a, p) = `Assignment`$(V, p, F, t_{i\in\mathscr{I}}^{\text{lim}}, t_{j\in\mathscr{O}}^{\text{stamp}})$
    Calculate $t_{j\in\mathscr{O}}^{\text{act}}$ = `Simulator`(a), converged = 1
    for $j \in \mathscr{O}$ **do**
      if $a_j \neq \emptyset$ `and` $t_{a_j}^{\text{act}} > t_{a_j}^{\text{lim}} + \delta_a$ **then**
        converged = 0 {don't exit the for loop without checking all assignments}
        $V_{a_j,j} = -\infty$, $F_{a_j,j} = 0$
      end if
    end for
  end while
  Output: assignment and price

4.2 Application to the Aircraft Holding Problem

Consider the problem of releasing n aircraft holding in h stacks, with n_i aircraft in stack i. A fair solution to release the aircraft is a first-come-first-served (FCFS) scheme. We note that the time required to fly from a holding area to the FAF is not identical for all holding areas and each holding area may have a different number of aircraft in hold at any given moment. As a result, a precise definition of FCFS is needed: we define "served" in FCFS in the sense of being offered a choice. The choice made by the aircraft that enters earlier is strictly honored, except in the event of an emergency, and regardless of its implications for aircraft that enter the holding stacks thereafter.

Definition 4 (Baseline FCFS) A scheme to allocating FAF slots is said to be FCFS if and only if it ensures the following for *any* two aircraft. Suppose that the two aircraft arrive at times t_1 and $t_2 > t_1$ in either the same or different holding stacks. Let $T_b^h[1]$ denote the time for which aircraft 1 holds in its stack, and $T_{b,i}^h[1]$ the time for which it would have had to hold had aircraft 2 not arrived at all. Then $T_b^h[1] = T_{b,i}^h[1]$.

We note that FCFS guarantees neither that any given aircraft spends a shorter time in hold nor that it arrives sooner at the FAF than aircraft arriving at a holding stack after it. One may define, alternately, a scheme for allocating FAF slots that minimizes the net holding time of aircraft. Either of these cases can be considered as the *baseline* against which the proposed environmentally-optimal scheme will be evaluated.

Due to the non-stationary nature of the problem, we adopt a receding-horizon approach for solving the optimal problem. Recall that a receding horizon scheme takes the following form:

1. Starting with the current time, solve the OCP over a time interval of duration $T_p > 0$, where T_p is called the planning horizon. We refer to this as the short-term OCP (ST-OCP).
2. Implement the solution for time $T_h < T_p$ or until an external signal for replanning is triggered (e.g., an emergency involving an aircraft), and return to step 1 after updating the current time.

Problem 2 (ST-OCP) At time $t = 0$ (without loss of generality), decide which aircraft will be released during time $[0, T_p]$, such that:

1. The arrival time of individual aircraft at the FAF ensures the desired inter-aircraft separation.
2. The net environmental impact of the hold is minimized, in the sense of (8).
3. Constraints on the arrival time of individual aircraft at the FAF, to be specified presently, are satisfied.

Algorithm 3 Receding horizon solver

Require: : at least one aircraft in the holding pattern
  Initialize: time $t = t_0 =$ clock time; planning and replanning horizons T_p, T_h {in minutes}
  Update_Flag $= 1$, Interrupt $= 0$ {Interrupt is an external signal}
  while Update_Flag $== 1$ **do**
    Determine $\mathscr{I}$ {the set of aircraft}
    Determine $\mathscr{O}$ and $t^{\text{stamp}}_{j\in\mathscr{O}}$ {the set of all available FAF slots}
    Determine $t^{\text{lim}}_{i\in\mathscr{I}} = \texttt{Baseline}(\mathscr{I}, \mathscr{O}) + \delta_B$ {Definition 4, with δ_B user defined}
    For each aircraft $i \in \mathscr{I}$, compute F_i, the set of all feasible assignments
    Compute V, the value matrix for each (aircraft, feasible slot) pair
    Assignment: (a, p) = $\texttt{IterativeAssignment}()$
    Calculate $n' = \text{card}(\{j : a_j \neq \emptyset\})$ {the number of allocated slots}
    Compute replanning period: $T_h = \min\{T_p/2, n'/2\}$
    while $t < t_0 + T_h$ and Interrupt == 0 **do**
      t += 1
    end while
    $t_0 = t$, Interrupt = 0
    if `isempty`(holding stacks) == TRUE **then**
      Update_Flag = 0
    end if
  end while

The slot allocation logic based on receding horizon is listed in Algorithm 3. The problem ST-OCP is solved inside the outer `while` loop. All aircraft in the holding areas at the planning instant are eligible to participate in the auction. The number of slots released for allocation are user-specified (in this case, by the air traffic control), and each slot is identified by its nominal time stamp $t^{\text{stamp}}_{j\in\mathscr{O}}$. Thus, no additional calculation is needed to determine $t^{\text{stamp}}_{j\in\mathscr{O}}$. On the other hand, the calculation of $t^{\text{lim}}_{i\in\mathscr{I}}$ relies on the calculation of a baseline allocation, which is accomplished using the function `Baseline`. The FCFS scheme of Definition 4 is one possible baseline scheme, albeit not the only one, and we use it for the simulation study described in the next section. The value matrix, V, is calculated using an equation similar to (8), where

$$V_{ij} = -A_{ij}, \;\; A_{ij} = f^h_{ij}(T^h[i,j]) + f^f_{ij} \tag{9}$$

and $T^h[i,j]$ denotes the cumulative hold time of the i^{th} aircraft when it accepts the slot j, measured from the time it enters the stack (rather than the start of the planning horizon), f^h_{ij} denotes the amount of fuel consumed during the hold, and f^f_{ij} denotes the amount of fuel consumed during the flight to the FAF when the slot j is accepted. Thus, any scheme that maximizes $\sum_j V_{a_j j}$ minimizes the total fuel consumption, which is assumed to be proportional to the carbon footprint of the aircraft.

4.3 Calculation of the Economic Reward

In order to compute the economic reward, we start by recalling the complementary slackness condition (Bertsekas 1981). Let $\mathscr{O}_{\text{assg}}$ denote the set of objects which have been assigned; $a_j \neq \emptyset$ for any $j \in \mathscr{O}_{\text{assg}}$. For an optimal assignment (a_j, j), the optimal price vector p^* satisfies

$$V_{a_j j} - p^*_j = \max_k (V_{a_j k} - p^*_k | F_{a_j k} = 1) \;\Leftrightarrow\; A_{a_j j} + p^*_j = \min_k (A_{a_j k} + p^*_k | F_{a_j k} = 1) \;\forall j \in \mathscr{O}_{\text{assg}}$$

which is based on the notation in (9) and reflects the fact that the fuel consumption is actually a cost. The price vector p is still interpreted as the price (in terms of abstract emissions). There are two distinct ways to map it to an economic incentive:

1. Use the price vector to determine the *cost paid by the airlines* for slot allocation. Practically, this translates into an environmental surcharge on *all* landing aircraft.
2. Adjust the price to determine *a fair, equivalent reward to be paid to the airlines* that operate environmentally friendly aircraft. In practical terms, this represents credit given to possibly a subset of the airlines - either in terms of a reduction in the airport landing fees or equivalent carbon credits.

We posit that the equivalent reward can be calculated by the solving the following linear programming problem (LP), where we have imposed the requirement that $p_{(\cdot)}^{*,\text{adj}} \leq 0$:

Problem 3 (Equivalent reward) Minimize $\sum_{j\in\mathscr{O}_{\text{assg}}} p_j^{*,\text{adj}}$ subject to

$$\begin{aligned} A_{a_j j} + p_j^{*,\text{adj}} &\leq A_{a_j k} + p_k^{*,\text{adj}} \quad \forall j \in \mathscr{O}_{\text{assg}} \text{ and } \forall k \neq j \text{ s.t. } F_{a_j k} = 1 \\ p_j^{*,\text{adj}} &\leq 0 \quad \forall j \in \mathscr{O}_{\text{assg}} \end{aligned}$$

The solution $p_{j\in\mathscr{O}}^{*,\text{adj}} \leq 0$ can be mapped to an equivalent reward so that the aircraft a_j receives a reward $R[a_j] = p_j^{*,\text{adj}}$. A simpler scheme for calculating the economic reward is based on the assumption that an aircraft spends a longer time in hold because of its smaller environmental footprint compared to other aircraft. Therefore, the reward can be calculated by comparing the holding time with that for a fair baseline (e.g., FCFS). Let $T^h[j]$ and let $T^f[j]$ denote, respectively, the total holding time of the j^{th} aircraft and the time taken by that aircraft to fly to the FAF. Let $T_b^h[j]$ and $T_b^f[j]$ denote the corresponding values if a baseline policy were employed. Then, the economic reward given to the j^{th} aircraft is

$$R[j] = \max(0, \alpha(T^h[j] + T^f[j] - T_b^h[j] - T_b^f[j])), \quad \alpha > 0$$

The constant α can be chosen to equal the fuel burn rate or the rate of emission of the specific aircraft.

5 SIMULATION RESULTS

We present a few illustrative simulation results in this section. Each simulation runs in discrete time, with time steps of 1 minute. At each time step, each stack receives an aircraft with a probability of 50%. The type of the aircraft is decided by drawing from a uniform distribution, with the types restricted to those in Table 1. Based on the slot allocation calculated by Algorithm 3, aircraft are released from the holding areas only if their altitude is less than or equal to 9000 ft. For simplicity, we use a hierarchical approach for ensuring that the inter-aircraft separation (IAS) requirement is met wherein the ordering prescribed by Algorithm 2 is retained and aircraft are simply moved to later slots if necessary. We assume that only two stacks are used for holding. Moreover, the time of flight from each stack to the FAF is 5 min. We compare the statistics of policies that minimize fuel consumption and holding time, respectively. In both cases, assignment is carried out using the auction algorithm. Results representing 10 simulations of each policy, each of which is carried out over 40 minutes, are summarized in Table 2. The replanning horizon is set to 2 minutes. It is evident that the minimum-fuel policy helps reduce the environmental cost by nearly

Table 2: Statistics over 10 simulations.

Policy ($\rightarrow$)	Holding time		Environmental cost	
Statistic /Variable	Aircraft released	Env cost	Aircraft released	Env cost
Average	23.8	331	28.7	299
Std deviation	4.92	83	3.3	54

10% per aircraft released from hold. Interestingly, the environmentally-driven policy also leads to a larger throughput: we believe that this is the outcome of our hierarchical approach to rearranging aircraft for meeting the IAS requirement. This result, however, demonstrates how the environmentally-driven approach groups heavier aircraft into successive slots indirectly, by prioritizing low emissions.

Next, we investigate how the rewards are distributed between the participating aircraft. The rewards are computed by solving the LP in Prob. 3. In Table 3, we list the fraction of aircraft of a given type that receive rewards through assignment. Since an aircraft may participate in multiple rounds of assignment, we record the rewards assigned only during the last round when that aircraft participates. Since the rate of fuel consumption of the A333 is similar to that of the B772 in Table 1, we group them together for the purpose of statistics. We observe, as expected, that the fraction of aircraft from a given family receiving a reward reduces as the fuel consumption rate of that family increases. The B747, which has the highest rate of fuel consumption per Table 1, does occasionally receive rewards when the arrivals traffic is heavy. We note such reward would not be accrued if the holding time were compared with a baseline policy which minimizes the holding time. Although it is not evident from the table, the raw data suggests that aircraft do not receive any reward when the stacks are sparsely populated, while the calculated economic rewards are non-zero only when each stack is populated with more than 2 aircraft at the time of slot assignment.

Table 3: Fraction of aircraft that receive a reward during a given simulation, over 10 simulations.

Statistic (↓) Aircraft (→)	A320	B747	A333 + B772
Average	0.8	0.18	0.49
Std deviation	0.13	0.12	0.07

Finally, in Table 4, we illustrate how the statistics in Table 2 change when the distribution of the arriving aircraft is skewed towards the lighter A320, as is the case at several major airports. It is clear that the throughput increases due to the reduced pressure on inter-aircraft separation. The environmentally-driven policy delivers a 13% improvement in the environmental cost, comparable to that seen in Table 2.

Table 4: Statistics over 10 simulations with a greater number of A320 aircraft.

Policy (→)	Holding time		Environmental cost	
Statistic /Variable	Aircraft released	Env cost	Aircraft released	Env cost
Average	36	182	38.2	196.7
Std deviation	5.35	27.57	2.3	48.09

6 CONCLUSION

In this paper, we presented an algorithm which allocates landing slots at airports while minimizing the environmental impact of the holds and, simultaneously, determines a reward which can be given to airlines whose aircraft hold longer in order to minimize the environmental impact. We noted that a reward scheme, such as the one presented here, is essential in order to make environmentally-friendly holding policies viable in a practical setting. We used an auction algorithm to determine the optimal assignment. The price computed by the auction algorithm can proposed as a substitute for the emissions-related component of the landing fees. Alternately, it was shown how the optimal assignment can be used to obtain a linear programming problem for calculating a fair reward for airlines with environmentally friendly aircraft. The problem of optimal allocation of landing slots has been widely addressed in the literature, and the present paper adds to that repertoire by contributing an economic angle. Future work would need to focus on simulation with richer, realistic data and field testing. The optimization techniques employed in this paper, and elsewhere in the literature, are mature and proven enough for this purpose. The practical viability of these methods rests on whether the economic rewards calculated on the basis of these methods are deemed acceptable by airlines and airport operators alike.

REFERENCES

Anderson, J. D. 2016. *Introduction to Flight*. 8 ed. New York, NY: McGraw-Hill Education.

Bäuerle, N., O. Engelhardt-Funke, and M. Kolonko. 2007. "On the Waiting Time of Arriving Aircrafts and the Capacity of Airports With One or Two Runways". *European Journal of Operational Research* 177(2):1180–1196.

Bayen, A. M., C. J. Tomlin, Y. Ye, and J. Zhang. 2004. "An Approximation Algorithm for Scheduling Aircraft with Holding Time". In 43rd *IEEE Conference on Decision and Control (CDC)*, 2760–2767. December 14th - 17th, Atlantis, Paradise Island, Bahamas.

Beasley, J. E., J. Sonander, and P. Havelock. 2001. "Scheduling Aircraft Landings at London Heathrow Using a Population Heuristic". *Journal of the Operational Research Society* 52(5):483–493.

Bertsekas, D. P. 1979. "A Distributed Algorithm for the Assignment Problem". Technical report, Lab. for Information and Decision Systems Working Paper, Massachusetts Institute of Technology.

Bertsekas, D. P. 1981. "A New Algorithm for the Assignment Problem". *Mathematical Programming* 21(1):152–171.

Ding, Y., and J. Valasek. 2007. "Aircraft Landing Scheduling Optimization for Single Runway Noncontrolled Airports: Static Case". *Journal of Guidance, Control, and Dynamics* 30(1):252–255.

Folse, S., H. Tran, L. Jensen, and R. J. Hansman. 2016. "Cruise Altitude and Speed Optimization Implemented in a Pilot Decision Support Tool". In 16th *AIAA Aviation Technology, Integration, and Operations Conference*. June 13th - 17th, Washington D.C., AIAA 2016-4211.

Higasa, K., and E. Itoh. 2022. "Controlling Aircraft Inter-Arrival Time to Reduce Arrival Traffic Delay via a Queue-Based Integer Programming Approach". *Aerospace* 9:663.

Ikli, S., C. Mancel, M. Mongeau, X. Olive, and E. Rachelson. 2020. "Coupling Mathematical Optimization and Machine Learning for the Aircraft Landing Problem". In 9th *International Conference for Research in Air Transportation*. June 23rd - 26th, Tampa, FL.

Itoh, E., and M. Mitici. 2019. "Queue-Based Modeling of the Aircraft Arrival Process at a Single Airport". *Aerospace* 6:103.

Le, L., S. Kholfi, G. Donohue, and C. Chen. 2003. "Proposal for Demand Management Using Auction-Based Arrival Slot Allocation". In 5th *USA/Europe Air Traffic Management R&D Seminar*. Budapest, Hungary.

Martinez-Val, R., E. Perez, and J. F. Palacin. 2005. "Historical Evolution of Air Transport Productivity and Efficiency". In 43rd *AIAA Aerospace Sciences Meeting and Exhibit*. January 10th - 13th, Reno, NV, AIAA 2005-121.

Mesgarpour, M., C. N. Potts, and J. A. Bennell. 2010. "Models for Aircraft Landing Optimization". In 4th *International Conference on Research in Air Transportation*. June 1st - 4th, Budapest, Hungary.

Murca, M. C. R., and C. Muller. 2015. "Control-based Optimization Approach for Aircraft Scheduling in a Terminal Area with Alternative Arrival Routes". *Transportation Research Part E: Logistics and Transportation Review* 73:96–113.

Scheelhaase, J. D., and W. G. Grimme. 2007. "Emissions Trading for International Aviation - An Estimation of the Economic Impact on Selected European Airlines". *Journal of Air Transport Management* 13:253–263.

Soomer, M. J., and G. J. Franx. 2008. "Scheduling Aircraft Landings Using Airlines' Preferences". *European Journal of Operational Research* 190(1):277–291.

Stettler, M. E. J., S. Eastham, and S. R. H. Barrett. 2011. "Air Quality and Public Health Impacts of UK Airports, Part 1: Emissions". *Atmospheric Environment* 45:5415–5424.

Sun, J., J. Ellerbroek, and J. M. Hoekstra. 2018. "Aircraft Initial Mass Estimation Using Bayesian Inference Method". *Transportation Research Part C: Emerging Technologies* 90:59–73.

UK Civil Aviation Authority 2013. "Practicalities of Airspace Based Emissions Trading System". Technical Report CAP 1055.

UK Civil Aviation Authority 2023. "The Integrated Aeronautical Information Package". https://nats-uk.ead-it.com/cms-nats/opencms/en/Publications/AIP/. accessed 13th April 2023.

Zhang, J., P. Zhao, Y. Zhang, X. Dai, and D. Sui. 2020. "Criteria Selection and Multi-Objective Optimization of Aircraft Landing Problem". *Journal of Air Transport Management* 82:101734.

Zhang, M., and A. Fillippone. 2022. "Optimum Problems in Environmental Emissions of Aircraft Arrivals". *Aerospace Science and Technology* 123:107502.

AUTHOR BIOGRAPHIES

ADITYA A. PARANJAPE is a Senior Scientist at TCS Research, Tata Consultancy Services Ltd. He is also Honorary Lecturer in the Department of Aeronautics at Imperial College London and Visiting Associate Professor at the Indian Institute of Technology (IIT) Bombay. His research interests include control theory, nonlinear dynamics and multi-agent systems. His e-mail address is aditya.paranjape@tcs.com.

ANWESHA BASU was a researcher at TCS Research. She is currently a graduate student at Texas A & M University. Her email address is anwesha042@gmail.com.