information

MDPI

Article

Digital Assistant for Arrival Scheduling with Conflict Prevention Capabilities

Francesco Nebula, Roberto Palumbo, Gabriella Gigante * and Angela Vozella

Italian Aerospace Research Center (CIRA), 81043 Capua, Italy; f.nebula@cira.it (F.N.); r.palumbo@cira.it (R.P.); a.vozella@cira.it (A.V.)
* Correspondence: g.gigante@cira.it

Abstract: Nowadays, in view of the growing traffic volume, an appropriate aircraft sequencing in the arrival sector is needed to maintain safety levels and improve the performance of the runway system and flight times. This paper presents a digital assistant supporting the air traffic controller in aircraft sequencing by providing suggestions for next waypoints, speed adjustments and altitude holdings. On the one hand, the suggested paths are such to preserve safety by ensuring the prescribed minimum separation, while also promoting environmental benefits through continuous descent operations (CDO). On the other hand, the suggestions aim to reduce landing times, improving the runway throughput. The proposed tool exploits multipath planning, for which a global optimization technique is used in conjunction with the dynamic time warping distance metric and a reinforcement learning approach to resolve conflicts through speed modulation and/or altitude holding. The performances of the assistant are assessed by means of a multi-agent simulator tailoring its reasoning on the procedures of Olbia airport (Italy). The analysis of a stream of many random aircraft has revealed its effectiveness in terms of arrival time reduction against a standard first-come-first-served strategy, usually adopted by controllers, and strong conflict reduction while considering a CDO-like adherence. Additionally, the man/machine interaction is investigated through an analysis of the overall latency from the suggestions provided by the digital assistant up to the actual aircraft maneuvers.

Keywords: arrival sequencing and scheduling problem; trombone procedure; decision support system; terminal maneuvering area; conflict detection and resolution; reinforcement learning; continuous descent operation; fast time simulations

check for updates

Citation: Nebula, F.; Palumbo, R.; Gigante, G.; Vozella, A. Digital Assistant for Arrival Scheduling with Conflict Prevention Capabilities. *Information* **2023**, *14*, 216. https://doi.org/10.3390/info14040216

Academic Editors: Emil Pricop, Grigore Stamatescu and Jaouhar Fattahi

Received: 27 February 2023
Revised: 23 March 2023
Accepted: 28 March 2023
Published: 1 April 2023

1. Introduction

The goal of air traffic control (ATC) is not only to manage air traffic safely and efficiently but also to plan and coordinate it to optimize the use of available airspace and infrastructure. In particular, terminal management areas (TMAs) or, more generally, terminal control areas (CTAs), are critical areas where flows of departing and arriving aircraft with varying altitudes, speeds and directions converge, increasing traffic volumes and overall complexity [1].

Since air traffic congestion is one of the causes of flight delays, improving traffic management in the terminal area can reduce arrival delays and allow better utilization of runways. Specifically, the approach controller (APP) plays a critical role in ensuring safe, orderly and timely traffic control of aircraft during landing and takeoff operations, based on the current regulations. One of the concerns of the APP controllers is to determine the order and the times in which aircraft are to land by properly arranging the minimum longitudinal spacings on the final path. This task, known as *sequencing* and *scheduling*, is becoming increasingly difficult as traffic volumes increase and the need to optimize air traffic flow increases. However, it is evident that the proper organization of queue sequencing at an earlier time horizon can result in considerable advantages in relation to both runway capacity and throughput. Consequently, these improvements can lead to a decrease in fuel burn and consequent cost savings for the airlines.

The constraints of the so-called *arrival sequencing and scheduling problem* (ASSP) vary depending on the airport level and on the potential impact of their violation at the ATC segment level. They are usually grouped in hard and soft constraints [2].

The *hard constraints* are typically related to safety and must always be respected. One of the most limiting factors for take-off and landing frequencies is represented by wake turbulence. The wake vortex magnitude is proportional to the size of the leading aircraft and its effects are more relevant if the trailing aircraft is small [3]. Consequently, the mean separation times (MSTs) depend on the types of both aircraft, and the landing order plays an important role for the maximization of the runway capacity [4]. The MST values are derived from the ICAO minima in-trail separations and reported in Table 1, where "1" represents aircraft with a maximum weight of 7 t, "2" if their weight is more than 7 t but less than 136 t, "3" if their weight is greater than 136 t.

Table 1. MSTs (s) according to ICAO wake vortex safety rules [5].

		Trailing Aircraft Category		
		1	**2**	**3**
Leading Aircraft Category	**1**	82	69	60
	2	131	69	60
	3	196	157	96

It must be pointed out that, recently, a finer categorization of the wake vortex separation minima has been introduced by EASA as a result of an extensive research collaboration between EUROCONTROL and FAA. With this new information, aircraft have been assigned to one of six new categories (A through F) under the RECAT-EU six-category scheme [6]. In this study, however, we will focus exclusively on the categories listed in Table 1. In fact, the inclusion of additional categories would not alter the constraints of the proposed approach but would only consider other values of separation minima.

Figure 1 reports an example of the reduction of the queue length (*makespan*) when a proper aircraft sequence is imposed: the optimal ordering leads to the minimum makespan achieving a reduction of about 33%. While optimal solutions are desirable, very often even sub-optimal solutions are used to improve runway throughput.

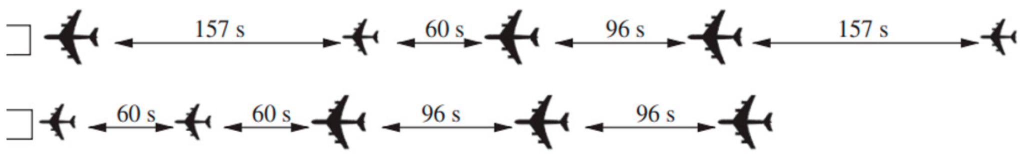

Figure 1. Effects on runway throughput of the optimal aircraft sequence [7].

In contrast, *soft constraints* are constraints that can be violated leading, in this case, to a suboptimal solution to the ASSP. A typical example is represented by the planned time slots (available or assigned) for landing in airports with limited capacity. In these airports, it is acceptable to adjust the planned landing times of some aircraft in order to achieve an overall optimized solution for the ASSP [8].

To accommodate higher traffic volumes and serve the changing needs of airlines and ATC, specific arrival route structures—such as "trombone procedures" or "point merge procedures"—have been introduced based on the performance-based navigation (PBN) concept [9]. These RNAV-based arrival (or departure) procedures have been intentionally designed to support the controllers in guaranteeing a common approach to the above-mentioned tasks ensuring a certain degree of predictability. These procedures were introduced to improve and ensure greater efficiency and safety of air traffic in the terminal regions because aircraft follow specific trajectories, and their possible collisions are minimized in the early design phase of the procedures themselves. In more detail, these procedures have been demonstrated to improve lateral dispersion and distance flown, allowing, at the same time, the arriving flights to remain at higher altitudes. Two immediate benefits of this are controller workload reduction and improvement of airport capacity [10]. Such procedures are currently adopted in several large European airports.

Nevertheless, under some conditions, performing the procedure along the designed nominal routes may be highly inefficient. For this reason, the ATC is authorized (and even encouraged), when possible, to use shortcuts towards the existing waypoints to expedite traffic, while ensuring safety and maximizing the throughput. Actually, a proper management of aircraft scheduling may reduce delays, increase the number of operations per unit time and minimize fuel consumption [11]. Scheduling a specific aircraft is related to ensuring adequate airport capacity and predictability of operations. Inadequate scheduling of flight operations leads to the lack of traffic smoothness and consequently to the onset of the congestion problem. Unfortunately, the controller's ability to achieve optimal aircraft scheduling and queue sequencing can be a highly challenging task that demands a considerable level of expertise and concentration, especially for large airports or when traffic levels in the terminal area exceed the system's capacity. Furthermore, the risk of conflicts increases in the vicinity of final approach fixes due to the bottlenecks they create.

The above-described topics compete on the horizontal plane, but a further aspect has to be taken into consideration to complete the problem description, the so-called continuous descent operation (CDO) [12]. It is an advanced flight technique to descend continuously from cruising altitude (top of descent, TOD) to the final approach fix (FAF) without levelling and with an idle thrust setting. According to this procedure, an aircraft could stay as high as possible for a longer time than with a conventional "step descent", expanding the vertical distance over ground and thus significantly reducing the noise levels for populated areas near airports. Furthermore, idle engine settings lead to a reduction of fuel consumption and harmful pollutant emissions [13,14]. Nevertheless, CDOs are not widely implemented, especially during high density operations, due to safety constraints. Actually, these procedures may require an increase in separation between aircraft arrivals, which may affect the airport arrival rate and runway throughput [12]. The larger separation spacing for a CDO aircraft is mainly due to two reasons: the difficulty for air traffic controllers to predict the future position of an aircraft with significantly variable speed and the inability of the pilot to quickly decelerate during descent [15]. Although CDO has been proven to be feasible and without increasing the required spacing between aircraft under light traffic conditions (such as nighttime operations), currently, aircraft flying CDO are most likely to be spaced further apart under heavy traffic conditions.

All these constraints to be considered by the APP lead to the need for new tools to assist him in giving the proper instructions. Considering the current perspective of digitalization in ATM, such new tools push the concept of a digital assistant intended as "a specialized intelligent artificial agent that helps users to do their activities" as an "intermediary between humans and other agents in a multi-agent environment" [16]. This paper presents a preliminary development of a digital assistant—named ACOP (Arrival digital assistant with COnflict Prevention)—aimed at helping the controller in the operational management of traffic during the arrival phase from the TOD up to the FAF, following trombone procedures. ACOP could support the ground controller in the specific task of shortcutting the trombone keeping safety levels, maximizing the throughput and minimizing the changes to ensure CDO adherence. ACOP uses artificial intelligence (AI) techniques and is structured in two main functionalities. The first one is based on multi-path planning triggered by any arriving flight, for which a global optimization technique is used, in conjunction with a distance metric based on the *dynamic time warping* (DTW) technique [17]. The second one prevents possible loss of separation (LOS) emerging at the tactical level, following a reinforcement learning approach.

The remainder of this paper is structured as follows. Section 2 reports the main current available solutions for the described problems of arrival scheduling and conflict prevention.

In Section 3, the ACOP tool is described in depth, while in Section 4, the obtained simulation results are shown under the specified conditions, and in Section 5, some operational aspects are discussed. Finally, in Sections 6 and 7, future work is addressed, and the most relevant conclusions are drawn, respectively.

2. State of the Art

The urgent need for increasing the efficiency of the ATM process, as pursued by the SESAR program [18] and FAA NextGen plan [19], is leading to intense efforts in designing automatic and semi-automatic tools to support and alleviate the work of the ACC controllers.

Generally, the two above mentioned functionalities of a ground digital assistant can be grouped in two hierarchical levels: a layer capable of solving the ASSP and one implementing a CDR capability, where the former is designed as a sporadic task, running on specific triggering events, while the latter is continuously running. In the next sections, a short overview on the main existing solutions for these two levels are provided.

2.1. Arrival Sequencing and Scheduling Problem

The ASSP aims to develop a scheduler minimizing the times of arrivals or their delays with respect to the schedule, maximizing in this way the runway throughput in compliance with safety and operational constraints.

This problem becomes a very complex task when all operational constraints are considered. Examples of these constraints are the number of delayed flights or their maximum delays, an essential aspect in terms of safety because it may be associated with a dangerous low fuel level. It is widely recognized that an exact algorithmic solution of the full ASSP is not feasible and, probably, not possible.

It is also worth noting that the scheduler must always generate an updated outcome, but several unexpected events may occur (change in airport configuration, in airport capacity, etc.), requiring an update in acceptable time (few seconds). Consequently, the scheduler should be also characterized by a low computational load [20].

Furthermore, the controllers, being driven by the general principles of *fairness* and *safety*, may be limited in flexibility to order the incoming aircraft. The strategy that often meets both principles is the so-called first-come-first-served (FCFS) one [21], playing the role of the "standard" sequence, against which all other strategies are usually referenced.

Unfortunately, the straightforward FCFS approach generally produces relevant delays; therefore, different strategies have been developed. In practice, the ATCOs use the *constrained positioning shifting* (CPS) technique. In particular, the k-CPS is a methodology requiring the determination of a parameter k representing the maximum number of shifts (generally less than three) of any aircraft with respect to the FCFS order.

It has been demonstrated that the CPS improves the runway throughput with low computational effort. In [22,23], two algorithms have been proposed, but no one takes into account the time slots for landing. In [24], an additional "relative" parameter is introduced to bound the workloads of controller and pilot during a re-sequencing. In [25], the authors describe a new procedure for a real-time dynamic programming by means of the optimal sequence computation that removes most of the issues of the k-CPS technique and minimizes the time makespan and the average delay. This new approach, named *dynamic positioning shifting* (DPS), allows for matching more quickly the operative constraints. However, the algorithm complexity increases linearly with the number of aircraft and runways, and the value of k is computed depending on the traffic level. In conclusion, the DPS behaves similarly to FCFS for a "normal" traffic volume, while it is most similar to the CPS technique for an intense one.

Another interesting line of research deals with models and algorithms based on *job shop scheduling*. From this point of view, the terminal area is considered as a machine [20] or as a job scheduling problem [26]. In the former work, it is shown that the ASSP corresponds to the *cumulative travelling salesman problem* (CTSP), where the objective is to minimize the sum of the arrival times at the customers (which are the runways, in our analogy) rather than the total travel time, as it happens in the classical TSP. The problem becomes more complex when both re-scheduling and re-routing are included in delay minimization, and for this, the job shop is a more useful approach. In this case, some heuristics are required to quickly compute good quality solutions [27,28]. In [29], the authors develop and compare different models for scheduling and re-routing, assuming high traffic volumes.

Other proposed methods include *fuzzy programming* [30], simplex-based [31] and genetic algorithms [32–34], which are particularly competitive in terms of solution quality and robustness. In particular, in [35], a *genetic local search* (GLS) algorithm has been proposed, but no real application was shown.

Finally, it should be noted that some works consider arrivals, departures and ground operations as parts of a more complicated but realistic problem [26,36].

The above works are generally difficult to use in a real application due to their non-negligible computation time. Some works try to fill this gap [28,37,38], but they are not fully satisfactory, and the ASSP is basically considered still as an unsolved problem, at least in real scenarios.

2.2. Conflict Detection and Resolution

Aircraft must always maintain horizontal and vertical safety distances from each other. If a potential LOS is detected in advance, the controller must provide resolution instructions to one or both aircraft to resolve the conflict. Typically, conflict resolution maneuvers involve a change in flight path, speed or flight level.

Consequently, software tools are strongly desired to ensure safety in high traffic density scenarios, often affected by high uncertainties too. Many mathematical models are proposed for conflict resolution [39]. In [40,41], the authors use reachability zones to represent the possible future aircraft positions, and conflict mitigation is achieved by separating these sets using flight dynamics relationships. However, with this approach, the computational time strongly increases with the number of aircraft and the spatial grid density. *Model predictive control* (MPC) is a promising approach for conflict resolution. In [42], the MPC performs trajectory prediction and conflict resolution simultaneously, but the mathematical model is quite complex, and the quality of the solution depends on the quality of the available historical data and models. MPC is also used in [43], where the authors propose many conflict resolution models that prescribe minimizing the costs associated with the maneuver. In [44], a preliminary analysis of the surrounding traffic is performed, introducing the concept of "aircraft ecosystem" and an assessment of traffic complexity. The individual aircraft pairs are then considered one-by-one for conflict resolution.

However, the conflict resolution models proposed in the above works have several common limitations. Complete knowledge of the conflict scenarios in terms of speeds and trajectories is required. For this reason, the models are very complex, and the solutions are characterized by low quality in presence of relevant uncertainties. Moreover, the input scenarios used should be well standardized. *Machine learning* (ML) may overcome these problems since it does not require prior knowledge of how to efficiently solve a conflict, and the algorithm can improve itself when exposed to scenarios not seen before. Unlike model-based approaches, ML can exploit the historical data considering the environmental uncertainties, and it is not required to map the actions to each possible scenario [45]. For decision problems such as conflict resolution, the breadth and continuity of state and action spaces are also important for ML methods. *Reinforcement learning* (RL), developed for solving various board games, can, in fact, help in this regard [46].

3. ACOP Description

The primary objective of ACOP is to assist air traffic controllers in the effective operational management of traffic during the arrival phase from the top of descent (TOD) up to the final approach fix (FAF) while following the trombone procedures. Specifically, ACOP helps in the critical task of shortcutting the trombone, while ensuring safety levels, maximizing the throughput and minimizing the changes to guarantee a significative continuous descent operations (CDO) adherence. In fact, ACOP provides two types of suggestions joining both sequencing and scheduling and conflict management functionalities issuing "Direct To" and "Conflict resolution" commands.

ACOP has been developed as a proof-of-concept for sequencing optimization in trombone procedures, with ENAV (the Italian air navigation service provider) providing the use case. To the best of our knowledge, there are currently no operational ground tools available to assist controllers in such procedures (including point merge or trombone), especially in regard to direct-to directives. As a result, ACOP solutions are compared to an air traffic control technique commonly used as a reference in research works [21], with the assumption being that a virtual controller directs traffic according to a first-come-first-served sequencing.

ACOP provides the following significant scientific contributions:

1. The tool has an intrinsic capability to reduce conflict occurrences already at the sequencing and scheduling level, by employing the DTW distance concept;
2. The tool also includes a preliminary automatic procedure for generating admissible paths starting from the STAR procedure of any airport;
3. Finally, the tool employs reinforcement learning techniques for the CDR process in order to train an agent to detect and resolve conflicts in real-time by learning from past experiences.

3.1. Functional Architecture

The right side of Figure 2 shows a functional architecture of ACOP. The dashed lines indicate the required preliminary input data (RNAV procedures of the airport and their initial off-line elaboration), while the solid lines identify the runtime modules and their interactions.

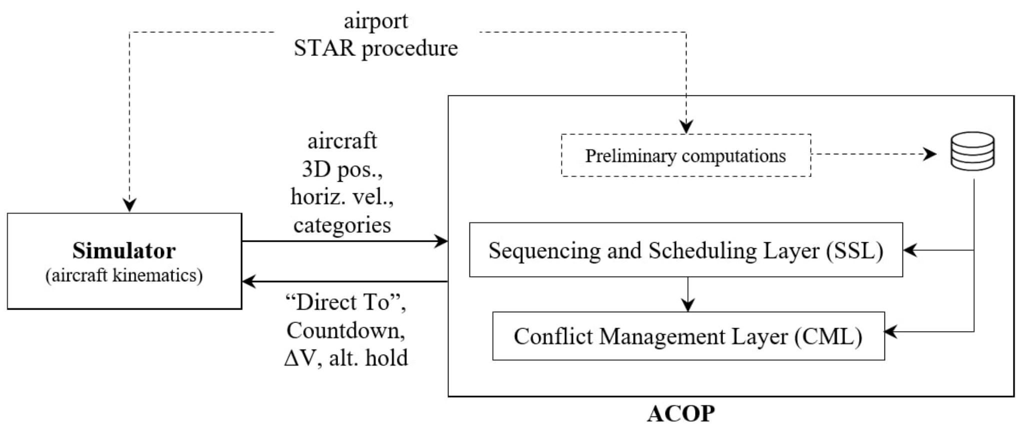

Figure 2. ACOP functional architecture.

The inputs to ACOP are represented by the aircraft three-dimensional positions, horizontal velocities and categories (see Section 1). The outputs of ACOP consist instead of real-time suggestions for path changes via appropriate "Direct To" commands (including a running countdown, i.e., a time limit for the controller to provide them to the pilots) and commands to reduce speed and/or hold altitude to avoid possible conflicts.

Figure 2 also highlights the ACOP *sequencing and scheduling layer* (SSL) and the *conflict management layer* (CML). The first layer provides a solution to the ASSP, while the second layer implements the CDR functionality. This division is useful because the time horizons of the two problems are very different, and the two functionalities are decoupled. Even though the SSL guarantees separation at FAF and spreads aircraft trajectories as much as possible, some LOSs could still occur if a dedicated CDR capability is not included. Consequently, a CML must be designed to comply with the required safety levels.

Finally, the left side of Figure 2 depicts a multi-agent simulator providing the positions of all aircraft at each time step, ideally following the sequence of the given waypoints (see Section 4.3 for further details).

3.2. Design Assumptions and Limitations

In this paragraph the main ACOP design assumptions have been collected.

1. "Direct To" starting points: The starting points of the "Direct To" commands are always the current next waypoints. In other words, an aircraft cannot leave its leg at an intermediate point.
2. Operational constraints: Each aircraft cannot be shifted in the landing sequence if this implies a delay at FAF greater than a user-defined value.
3. CDR approach: The available conflict resolution actions are only deceleration and levelling (no path stretch or vectoring, for example). The CDR approach is pairwise. However, if two aircraft are both already engaged in conflict resolution maneuvers, their conflict is considered not solvable.
4. Speed assumption: The ACOP ASSP solution involves the utilization of average speeds to calculate the ETAs in advance. This design assumption is necessary, and it may potentially result in slightly downgraded outcomes (for details, refer to Section 3.4.1).
5. Other limitations: The following factors are not considered by ACOP:
    - Previous delay cumulated by aircraft before entering the terminal area;
    - Departures;
    - Emergencies (for example, due to ground traffic or maintenance, adverse weather conditions, presence of drones, onboard health issues, commercial agreements with some airlines, etc.) that would give a higher priority to some specific aircraft.

3.3. Preliminary Computations

As shown in Figure 2, some preliminary airport data are required for the ACOP design. In more detail, both SSL and CML need the list of the admissible flown paths from a generic waypoint up to the FAF in order to provide only valid solutions to the human controller. The admissible paths are obtained imposing the following conditions:

1. Flight direction in each leg must be as specified in the STAR procedure;
2. Turns on each WP cannot be made with an angle less than 90°.

3.4. Sequencing and Scheduling Layer

The multipath planning is triggered when a new aircraft arrives on an initial fix, as no information is available on the upstream traffic (see Section 4.3). Therefore, the same aircraft can be rescheduled more than once because it is involved in multiple optimization sessions. In more detail, the SSL—developed within the MATLAB environment [47]—is composed of two in-series optimizations. They are required (1) to minimize the times of arrival and (2) to reduce the scenario complexity. Hereafter, they will be briefly described.

3.4.1. Flight Times Minimization

To maximize the runway throughput and capacity, a constrained combinatorial optimization problem must be solved.

Generally, the trajectory planning does not involve all the inflight aircraft but only those that are not already directed towards a preselected set of waypoints. At a minimum, this set of waypoints is composed only of the FAF. Additional waypoints nearby the FAF may be included in this set, but this will strongly reduce the state space and hence the possibility of finding good global solutions. In any case, all the aircraft directed to these specific waypoints will, of course, continue to be considered in the safety constraints, even if their trajectories are frozen. In our settings, no additional waypoints have been added to the abovementioned group, leaving the FAF as the only waypoint in the set.

It is assumed that the independent variable x is an integer vector composed of (1) the index of the path from the current next waypoint up to FAF for each of the n aircraft involved in the optimization process and (2) the arrival order on FAF, i.e.,:

$$x = [i_1 \ldots i_n k] \tag{1}$$

where $k \in [1; n!]$ represents the row index of the matrix M_{perm} containing all possible permutations of the n aircraft, with reference to chronological arrival order on the initial fixes. For example, for $n = 3$, this matrix looks like:

$$M_{perm} = \begin{bmatrix} 3 & 2 & 1 \\ 3 & 1 & 2 \\ 2 & 3 & 1 \\ 2 & 1 & 3 \\ 1 & 3 & 2 \\ 1 & 2 & 3 \end{bmatrix} \tag{2}$$

and $k = 5$ means that the arrival sequence on FAF is 1-3-2. The ETA on FAF point of the j-th aircraft can be written as:

$$ETA_j = \mathrm{t} + \frac{DTG_{P_{curr\,j}}^{WP_{next\,j}}}{V_j} + \frac{\left(DTG_{WP_{next\,j}}^{FAF}\right)_{i_j}}{\left(V_{avg}\right)_{i_j}} \tag{3}$$

where:

- t is the time in which an aircraft arrives in the terminal area, and the optimizer is triggered;
- $WP_{next\,j}$ is the current next waypoint of the j-th aircraft;
- V_j is the current horizontal velocity of the j-th aircraft;
- $DTG_{P_{curr\,j}}^{WP_{next\,j}}$ is the distance from the current position of the j-th aircraft up to $WP_{next\,j}$;
- $\left(DTG_{WP_{next\,j}}^{FAF}\right)_{i_j}$ is the distance from $WP_{next\,j}$ up to FAF, computed along the i_j-th path;
- $\left(V_{avg}\right)_{i_j}$ is the average velocity along the i_j-th path, with respect to the allowed values in each leg (note: a different actual speed profile due, for example, to a CR maneuver, might lead to results slightly downgraded).

The cost function to be minimized is assumed to be the sum of the planned of the estimated time of arrival (ETA) of each aircraft, i.e., $f_{obj} = \sum_{j=1}^{n} ETA_j$. This implies that each single ETA will be minimized too, all being positive quantities. Different choices of the cost function would not ensure this result because an optimization may override the previous one. Cleary, only the last term of the above equation can actually be optimized, involving the selection of the "future" paths after the current leg.

The problem constraints are represented by the ICAO wake vortex rules for spacings (Table 1) and by the maximum admissible delay D_{max} for each aircraft:

$$diff\left(ETA_{vec}^{k}\right) \geq f\left(Table_{ICAO}, Cat_{vec}^{k}\right) \tag{4}$$

$$ETA_{vec}^{k} - ETA_{vec_{nom}}^{k} \leq D_{max} \tag{5}$$

where:

- ETA_{vec}^{k} and $ETA_{vec_{nom}}^{k}$ are $(n \times 1)$ vectors composed by the computed and nominal ETAs of the k-ordered aircraft, respectively;

- $f\left(Table_{ICAO}, Cat_{vec}^{k}\right)$ represents a function providing the $(n-1)$ minima spacings for the k-ordered aircraft. $Table_{ICAO}$ is reported in Table 1, while Cat_{vec}^{k} is a $(n \times 1)$ vector containing the k-ordered aircraft categories;
- the $diff$ operator calculates differences between adjacent elements of its vectorial argument.

A genetic algorithm [48] is used because it has proven its effectiveness for the described application.

3.4.2. DTW-Based Filtering

If the optimization described above yields multiple solutions with the same cost function value, i.e., with the same total flight time, a further selection is performed using the *dynamic time warping* (DTW) algorithm.

DTW is an algorithm for measuring the distance between two discrete sequences, even if they consist of a different number of samples. It was originally introduced to calculate the distance of time series. Basically, DTW recursively searches all possible points that lie between two trajectories for the point with the smallest distance. Further details can be found in [17].

In our study, the DTW distance is initially calculated for each pair of trajectories of each ex aequo scenario. Then, the DTW distances are averaged for all pairs, and finally the scenario with the largest average DTW value is selected and proposed to the controller. For example, Figure 3 shows the DTW-based selection for two equivalent sets of paths. The selected scenario is the one on the right, which contains paths further apart. This mechanism clearly helps reduce in advance LOS risks, especially near the FAF.

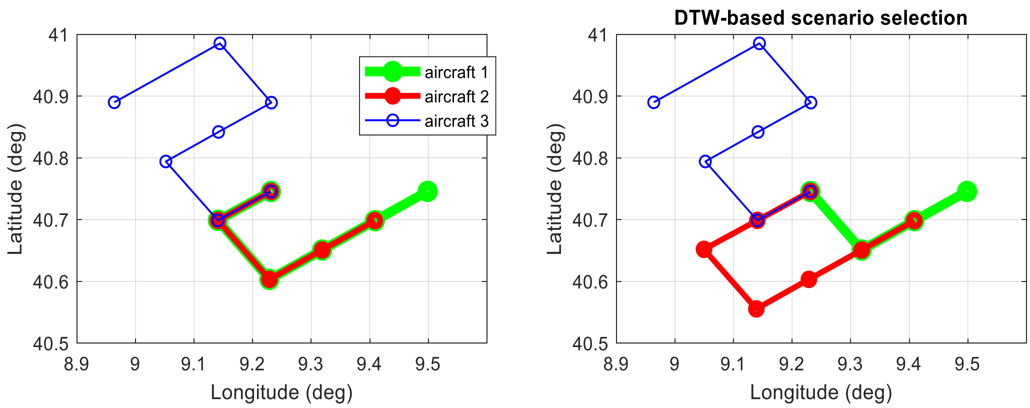

Figure 3. DTW-based selection from equivalent scenarios.

3.5. Conflict Management Layer

The CDR function is composed of two different sub-functions:

- Conflict detection between all possible pairs of aircraft and;
- Conflict resolution to separate aircraft by reducing speed and/or holding altitude.

3.5.1. Conflict Detection

The loss of separation (LOS) is a three-dimensional event defined as the violation of the horizontal and vertical separation minima at the same time.

When considering separation conditions, a key parameter to be evaluated is the expected closest point of approach (CPA) between two aircraft [49]. However, in contrast to the usual situation, the trajectory prediction here is performed along the paths previously found by the SSL.

Prediction is carried out propagating the current aircraft positions along the expected path both horizontally and vertically to check for the following separation condition:

$$\Delta h > \Delta h_{min},\ d > d_{min} \tag{6}$$

Δh being the altitude difference between the aircraft and d their relative distance. Prediction considers the current speed and flight-path angle as well as velocity limitations along the legs. Once the algorithm detects the potential LOS, coordinates of the infringement points along each path are determined. In Figure 4, an example of the "three-dimensional" CPA is reported for our application.

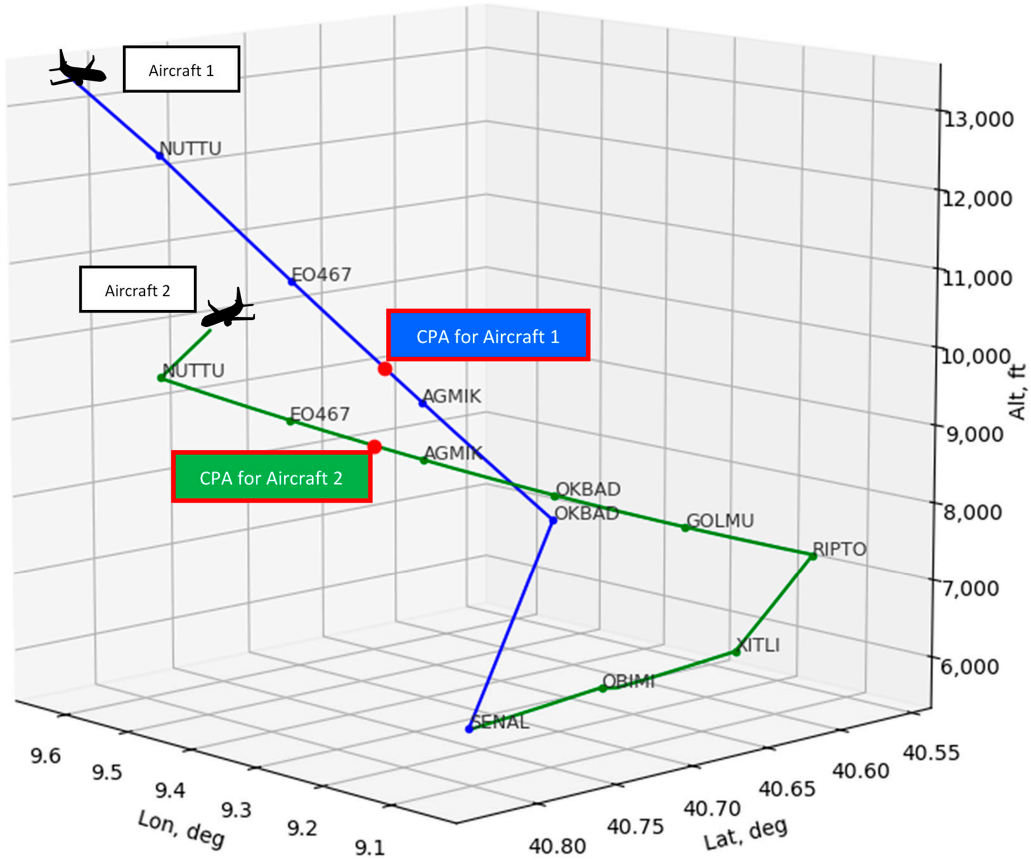

Figure 4. Pictorial view of the three-dimensional CPA.

3.5.2. Conflict Resolution

In the CR, only two actions are supposed available: speed reduction and altitude holding. When the CPA is defined (i.e., a conflict is detected), three different cases can be distinguished:

1. Immediate speed reduction within the admissible speed range of $[v_{min}; v_{max}]$ kts is applied to resolve the conflict. In the presented approach, it is assumed that the controller and pilot tasks are highly automated with minimal latency. However, if an estimation of the average response times and latencies is known, ACOP is capable of considering it in its computations and accounts for possible delays. In Section 5 of this study, we leverage this functionality to analyze the CDR in the context of current operational procedures with respect to different response times settings.
2. If there is no feasible speed reduction value within the admissible range, the altitude of the highest aircraft is held until the CPA is reached. More specifically, the holding altitude is set at 1000 feet above the expected conflict altitude. Of course, this maneuver breaks the CDO-like profile if the altitude holding lasts longer than 20 s [12]; however, it is necessary to avoid conflict and maintain safety.

3. If there is no speed reduction value capable of solving the detected conflict and the aircraft are only vertically separated (the aircraft are "one above the other"), we combine a predefined deceleration with an altitude holding of the highest airplane.

In the first case, the applied speed reduction is the minimum possible to avoid the horizontal infringement and is determined using a *reinforcement learning* (RL) approach [45,46]. When a potential conflict is detected, the trained AI model is called and returns an appropriate speed reduction for one aircraft only. The aircraft "to be controlled" is chosen considering the vehicle that is at a greater time distance to the expected CPA. If the RL solution results in a velocity value outside the allowable range, this means that speed reduction is not a feasible way to prevent the conflict, and another approach must be considered.

The RL agent training process for conflict resolution is briefly explained in the following. First of all, conflict scenarios are generated for aircraft pairs and presented to the agent using a custom learning environment. The agent, guided by the RL algorithm, learns to solve these conflicts by applying a speed reduction also given the environmental uncertainty. For such a maneuver, the agent receives a reward as performance feedback, and the value of the reward depends on the quality of the maneuvers. The learning goal is to maximize reward, and the agent is considered trained when consistently achieving high rewards for solving "never seen" conflict scenarios [46].

More formally, let P be the set of admissible paths. Let A and B be a pair of aircraft, and let v_A and v_B be their initial velocities. Let $d_{CPA\,A}$, $t_{CPA\,A}$ and $d_{CPA\,B}$, $t_{CPA\,B}$ be their distances and estimated times to the CPA, respectively, and let d_{rel} be their relative distance. Let t_{sol} be the time it takes for the controlled aircraft to reach the estimated CPA after the speed reduction, and let $t_{episode}$ be the episode length. Let s be a generic state and s_f the state at the end of the episode. Finally, let R be the reward function. We can now describe the process steps within each episode:

- Reset the environment:
  - Select two random paths $p_1, p_2 \in P$.
  - Generate aircraft A and B at random positions on the first leg of each path.
  - Assign random initial velocities v_A and v_B between 210 kts and 230 kts.
- Calculate CPA and determine $d_{CPA\,A}$, $t_{CPA\,A}$, $d_{CPA\,B}$, $t_{CPA\,B}$.
- Check for potential conflict. If no conflict, skip scenario with probability ε and reset.
- Generate state observation $s = [d_{CPA\,A}, d_{CPA\,B}, d_{rel}, v_A, v_B]$ and normalize values to [0, 1].
- Select aircraft to control: choose A if $t_{CPA\,A} \geq t_{CPA\,B}$ otherwise choose B.
- Select the action: choose a speed reduction δ in the integer range [10, 30] kts.
- Apply speed reduction to the selected aircraft.
- Receive reward:
  - If loss of separation, $R\left(s_f\right) = -10$.
  - If safe arrival, $R\left(s_f\right) = -t_{sol}/t_{episode}$.

The reward for safe arrival is such that it encourages the agent to choose the minimum speed reduction that is sufficient to solve the conflict. Otherwise, the naïve solution would be to reduce speed as much as possible. Each training episode ends either when a separation violation occurs or when the CPA is safely passed.

Since most of our generated scenarios were conflict-free, we decided to filter the episodes to feed the training loop mainly with scenarios that required conflict resolution. The training environment was implemented in Python3 [50]. We trained a DQN agent taking advantage of OpenAI Gym [51] APIs in order to be able to build a custom environment and use the StableBaselines3 RL algorithm library [52]. The DQN agent is trained to approximate the action–value function considering the MSE loss between predicted value and the target value. Satisfactory behavior was achieved after 400 k timesteps.

Figure 5 shows an example of the trained agent within a training scenario (details on the use case scenario are given in Section 4): once the potential conflict is detected (the red dots on the path indicate the CPAs for both aircraft), Aircraft A is instructed to reduce speed by 12 kts since it is temporally the farthest from the CPA (a). The black circles show the d_{min} buffer around the aircraft, while the light blue dot and circle are purely informative and show the position and buffer of Aircraft A without the RL model correction. When the CPA is reached, it can be seen that the conflict is resolved (b), and the episode ends once the two aircraft are safely away from the CPA (c).

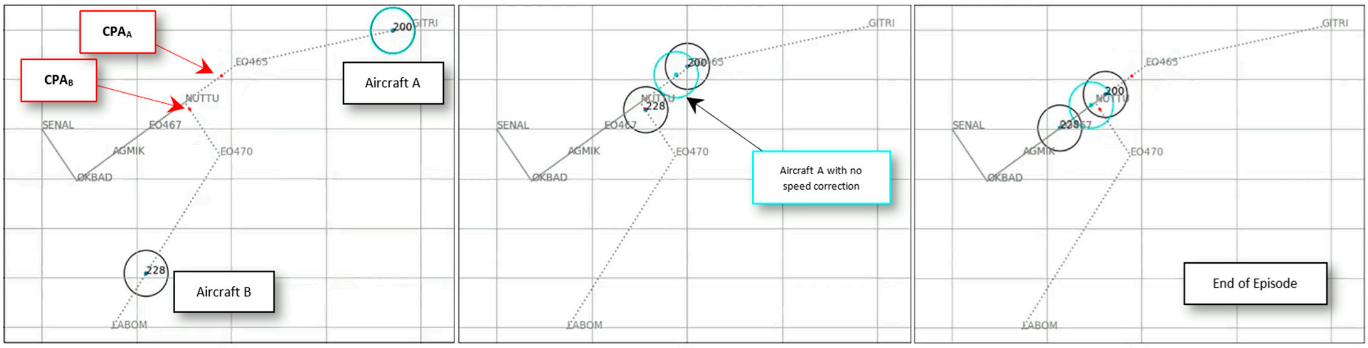

(a) (b) (c)

Figure 5. An example of the trained agent within the training scenario: (**a**) conflict detection and speed reduction; (**b**) conflict is correctly solved; (**c**) end of episode.

4. ACOP Performance Assessment

4.1. ACOP Settings

This study focuses on the STAR RNAV1 RWY05 procedure, in force at the Olbia Costa Smeralda (LIEO) airport [53] reported in Figure 6, where the FAF is represented by the SENAL waypoint.

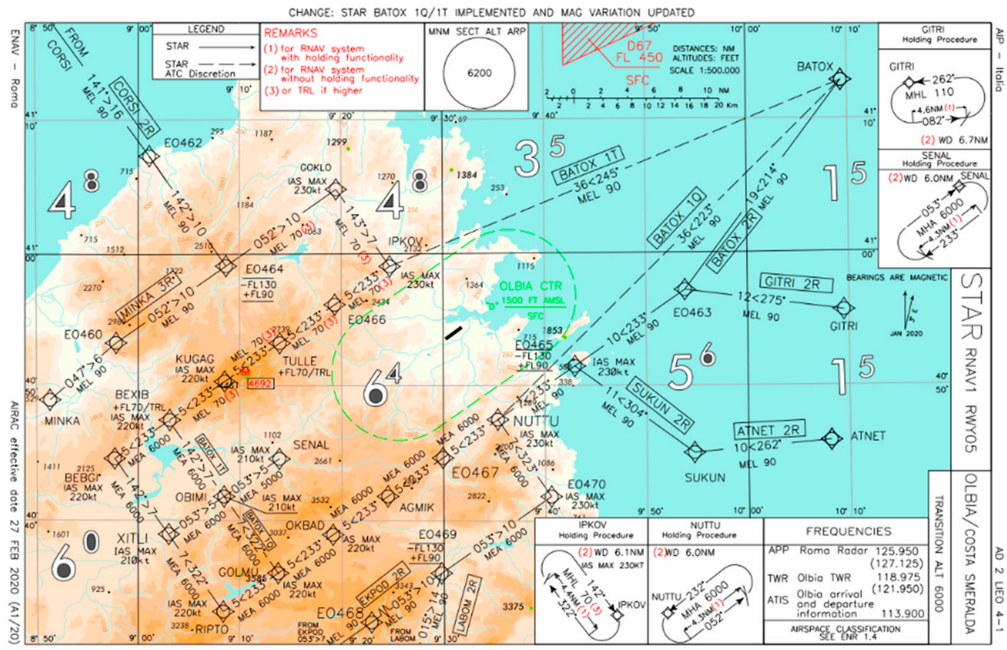

Figure 6. Trombone procedure—Costa Smeralda, OLBIA—STAR RNAV1 RWY 05 [53].

Preliminary offline computations on this procedure (see Section 3.3) allow the determination of valid paths from the final point of each leg to the FAF. The total number of admissible paths are 347. For example, in Figure 7, the 20 admissible paths starting from the CORSI-EO462 leg are shown.

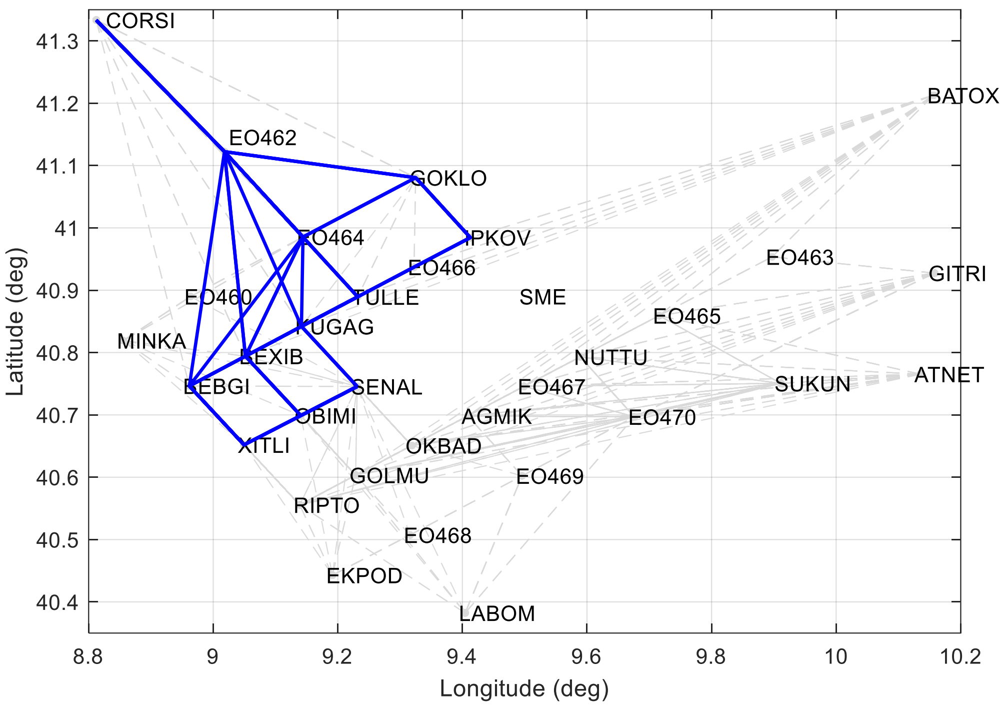

Figure 7. Admissible paths (in blue) from the end of CORSI-EO462 leg to SENAL (FAF).

The mentioned STAR procedure is also characterized by upper speed limits ranging from 210 to 230 kts (imposed by the next waypoints). The horizontal and vertical separation minima are assumed to be 5 NM (equal for the longitudinal and lateral directions) and 1000 ft (based on the ATM vertical separation minima), respectively. Finally, the controller can enter the following information in the ACOP user interface:

1. The list of waypoints to which the aircraft are assumed non-reschedulable (see Section 3.4.1), as a kind of "frozen horizon" (in our settings, only the planes directly heading to the FAF are considered non-reschedulable);
2. The maximum delay on FAF tolerated by each aircraft, with respect to its nominal arrival time (in our settings, it is assumed equal to 15 min for all aircraft);
3. The target altitude on FAF for all aircraft (in our settings, it is assumed equal to 5500 ft for all aircraft).

4.2. Reference Strategy

An ideal evaluation of ACOP's performance should require an analysis of the controller's actual behavior. In fact, we have obtained historical data concerning LIEO airport arrivals in diverse traffic level conditions, encompassing winter and summer seasons.

Nonetheless, we decided against utilizing these data because:

- The historical data do not capture all the significant factors that influence the controller's actions;
- Given the low traffic volume conditions of LIEO (less than two arriving aircraft per hour), a sequencing strategy is not mandatory, rendering a legitimate comparison with the output of ACOP unmeaningful.

For this reason, the reference solution to which ACOP has been compared is a simulated "ideal controller" capable of continuously implementing a pure FCFS strategy [21] while ensuring the minimization of ETAs on FAF. This means to impose the aircraft arrival order (i.e., $k = n!$) in the formulation reported in Section 3.4.1. Furthermore, the reference strategy does not have the additional capability of "spreading" the horizontal trajectories to maximize their "relative distance" (an innovative ACOP feature, described in Section 3.4.2). On the other hand, the low-level functionality of CDR is the same as that of ACOP. Therefore, the FCFS strategy also allows for a clear evaluation of the capabilities of ACOP to reduce in advance the conflict occurrences.

Finally, it is emphasized that the reference sequencing is quite demanding for a human controller, as it requires a continuous search for all routes with minimum travel times, ensuring at the same time FAF spacings and minima separations along the paths.

4.3. Simulation Features and Justifications

Hereafter the main simulation choices are listed.

1. Olbia Costa Smeralda airport: The assumed STAR procedures are related to the LIEO airport [53]. These procedures are defined within the Sardegna CTA [54]. As seen in Sections 3.2 and 3.3, some airport data are used for the ACOP design, but actually they are used for the simulation too. In particular, the STAR procedure also provides the upper bounds for altitude and IAS for each leg.
2. Aircraft simulation model: The behavior of each aircraft is modeled as a point mass with three translational degrees of freedom. The aircraft minimum approach speed is assumed to be equal to 200 kts. A sample time of one second is adopted.
3. No turn dynamics: The aircraft reach waypoints without simulating the turn dynamics.
4. Initial paths: The initial paths of all aircraft are the nominal sequences, specified in [53] for each initial fix.
5. Traffic level: A rule of thumb for the estimation of the maximum number of the aircraft allowed in the trombone pattern is given by $(D + B)/10 = (25 + 7)/10 = 3.2 \cong 3$, being the lengths of the downwind and base legs for LIEO airport equal to $D = 25$ NM and $B = 7$ NM, respectively. Of course, this condition is an indirect requirement for the traffic level simulation. In Figure 8, the trombone occupancies are reported for the assumed 1000 aircraft stream. The north and south trombone average occupancies result, respectively, on 0.56 and 1.42 for FCFS and 0.36 and 1.24 for ACOP. The related initial fix arrival times distribution is shown in Figure 9. In more detail, the CTA time intervals are selected uniformly random between 0 and 8 min but, to avoid unrecoverable collisions at the beginning, never less than 3 min for same or adjacent initial fixes (we assume that this condition is ensured by the upstream controller). This is the reason for the peak in Figure 9 in correspondence of this specific value.
6. Weather conditions:
    - No wind. The zero-wind assumption is adopted by both the simulation and ACOP sides as we aim to isolate and evaluate the functioning of ACOP exclusively. In any case, the trombone geometry is aligned with prevailing wind directions and runway orientation. Therefore, neglecting the wind would result in a constant offset to the velocities of the involved aircraft, with minimal alteration to their relative positions. In addition, speeds received by ACOP are ground speeds. No input information on forecast wind is usually available to the controllers; therefore, the internal ACOP computations are carried out using inertial data only, and the resulting deceleration commands are also with respect to ground. It is worth noting that, even in a real operational environment, a hybrid situation exists where the speed limits on the STAR procedure and the speed considered by pilots are both indicated airspeeds (IAS), while the controller visualizes ground speed. This discrepancy between IAS and inertial speed can result in tracking errors and is a well-known limitation of the current interaction between the controller and aircraft, which is compensated by the controller's expertise.
    - International standard atmosphere [55]. Upon arrival, all aircraft must adopt a common altimeter setting. Specifically, in the case of Olbia airport and many other airports, the QNH setting must be used, which provides all altitudes referenced to the mean sea level (MSL). As a result, the QNH altitudes will not be affected by the current atmospheric pressure. In fact, it is essential to have a common

altitude reference for all aircraft to ensure comparable results and to conduct an accurate assessment of ACOP performances. Therefore, the assumption of an international standard atmosphere (ISA) is useful only to homogenize altitudes and flight levels.

7. Flight phase definition: Each aircraft appears on an initial fix with an altitude ranging from 110 to 190 FL with steps of 1000 ft. Then, it will end its flight in correspondence with FAF (SENAL waypoint, see Figure 6), in this way excluding approach and landing from our problem.
8. CDO-like profiles: Accurate CDO simulation is out of scope in this work. Our scenarios begin when the aircraft reach the arrival sector, while CDOs are usually established at TOD. We therefore only emulate a continuous descent behavior assuming that the optimal FPA is the one that is determined considering the DTG to the FAF and the relative altitude between the aircraft and the FAF. While usually the resulting FPAs are compatible with typical vertical speeds, the FPA is nonetheless lower saturated to $-5°$. The descent profile is suspended only if the minimum altitude of the current leg is reached or if a leveling command is issued by ACOP to solve a conflict. Once the minimum altitude constraint is removed (or the conflict is solved), the aircraft recomputes the FPA towards the FAF and continues its descent.
9. ACOP inputs: The ACOP inputs are supposed not affected by any measurement error or delay.
10. Latency between ACOP outputs and their actuation: The simulation results presented in Section 4.5 are obtained using zero latency. However, in Section 5, non-zero latencies are considered.

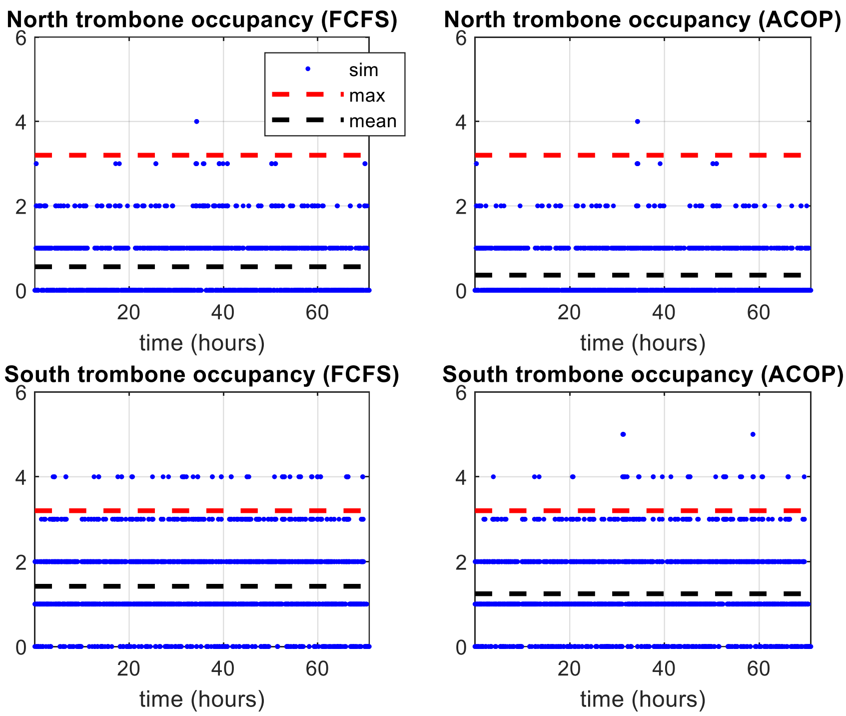

Figure 8. North and south trombones occupancies for FCFS and ACOP scenarios.

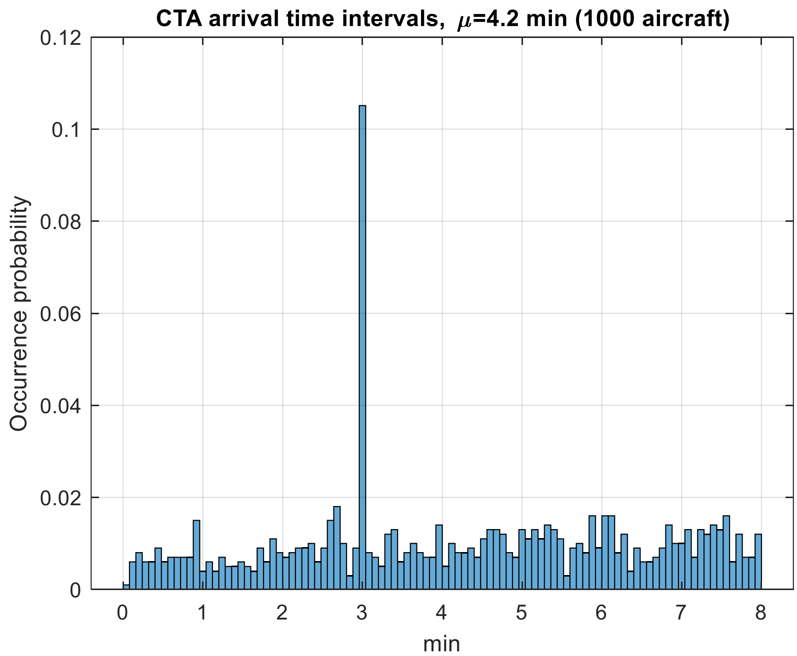

Figure 9. CTA arrival times distribution.

The above assumptions were shared with ENAV (the Italian ANSP) and are considered compatible with the assessment of the ACOP proof-of-concept feasibility, at least at this preliminary design stage. Future developments (see Section 6) will try to remove some of these assumptions and validate the digital assistant in a more relevant environment.

4.4. Metrics Definition

The ACOP general objective is to operationally support the approach (APP) controller in making decisions aimed at maintaining safety and optimizing the aircraft flow in the terminal area. However, the improvement of terminal area traffic management will necessarily have to go through the introduction of some metrics (or KPI), i.e., quantitative indications of the achievement of the declared objectives. In the following, the used metrics are introduced, while in Section 4.5, they will be evaluated for the case study. Note that, in all the mentioned metrics, the evaluations end at FAF (SENAL point).

1. A performance indicator for the efficiency could be trivially represented by the ACOP objective function, i.e., the sum of the arrival times on FAF (see Section 3.4.1). In Section 4.5.1, flight time reduction for each aircraft due to ACOP SSL will be quantified. It is worth noting that these ETA differences are strongly related to the well-known KPIs about runway capacity and throughput [56]. However, in our simulation environment, the improvements of capacity and throughput cannot be immediately appreciated because the CTA arrival frequency is kept fixed, regardless of the downstream management.
2. The predictability is defined as the percentage of the aircraft arriving on FAF with a time difference not greater than ± 15 min with respect to the nominal ETAs [56]. It is clear that this indicator penalizes both too early and too late aircraft. This metric will be verified in Section 4.5.1.
3. Regarding the safety aspects, usually the main indicator is the *number of LOSs*, and it will be evaluated in Section 4.5.2.
4. Finally, regarding the environmental impact, a relevant indicator is related to the adherence to CDO procedures, which results from the amount of time (larger than or equal to 20 s) spent in level flight from the top of descent (TOD) to the FAF [12]. In our case, we consider only the aircraft descent from the initial fix of the STAR procedure to the FAF. We also considered the number of performed horizontal flight segments

not shorter than 20 s. Of course, it should be noted that an accurate CDO simulation is out of the scope of this work, therefore we assumed that the vertical profile is simply computed by imposing a FPA that constantly points to SENAL at an altitude of 5500 ft (a tunable parameter).

4.5. Results

In this Section, a performance comparison between ACOP and the FCFS reference strategy is carried out under the same conditions. In fact, in both cases a flow stream of 1000 random aircraft is imposed using the simulation settings reported in Section 4.3. The arrivals cover a time span of 70.4 h with a flow rate of 14.2 per hour. The resulting mean arrival time on initial fixes is 4.2 min, a quite small value with respect to the considered airport. Although this scenario cannot be considered as a stress test, nonetheless, the assumed flow rate is about ten times greater than the observed one in the historical data of the same airport and for different seasons and mixed conditions (see Section 4.3).

It is noteworthy that the computational time required by ACOP is minimal (~1 s), facilitating its seamless integration into a real-world operational environment.

4.5.1. Flight Time Reduction

The FCFS strategy cannot provide solutions (i.e., paths not compliant with the ICAO constraints) in about 1.8% of cases, while ACOP in 4%. This difference can be explained considering that the FCFS strategy is a special case of the ACOP algorithm, facing a much less difficult combinatorial problem. In a real operational environment, the holding pattern procedures could be of help to solve these scenarios.

In 46.3% of the provided valid suggestions, ACOP returns solutions with the same aircraft sequence of FCFS, mainly in correspondence with low-traffic periods, as expected.

The mean aircraft queue length suggested by ACOP is 11.6 s shorter that the FCFS one, thanks to its reordering capability. In Figure 10, the histograms of the queue length margins over the minimum (optimal) values are reported for the same groups of aircraft.

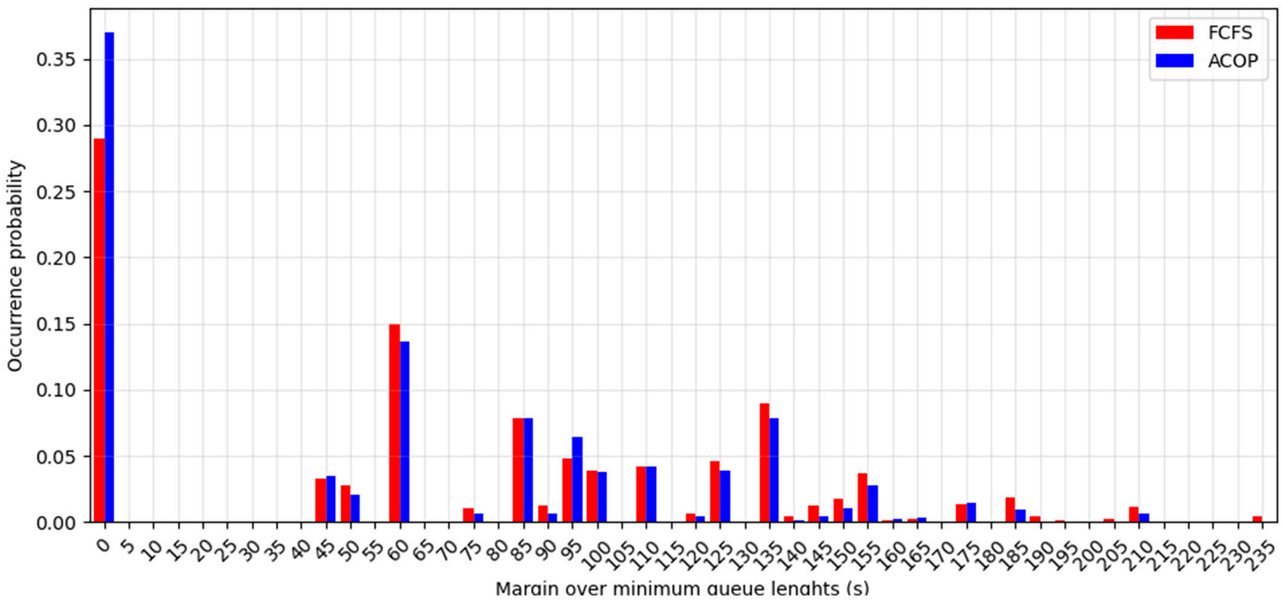

Figure 10. Time margins with respect to the optimal queue lengths.

In Figure 11 shows that ACOP performs a better exploitation of both runway and airspace, reducing the flight time of each aircraft from initial fix to SENAL. The time reduction has a mean of 1.44 min and a standard deviation of 3.73 min, and, after 1000 random aircraft, the "time saving" achieves the cumulated value of more than 24 h (total flight time is reduced by 10.7%).

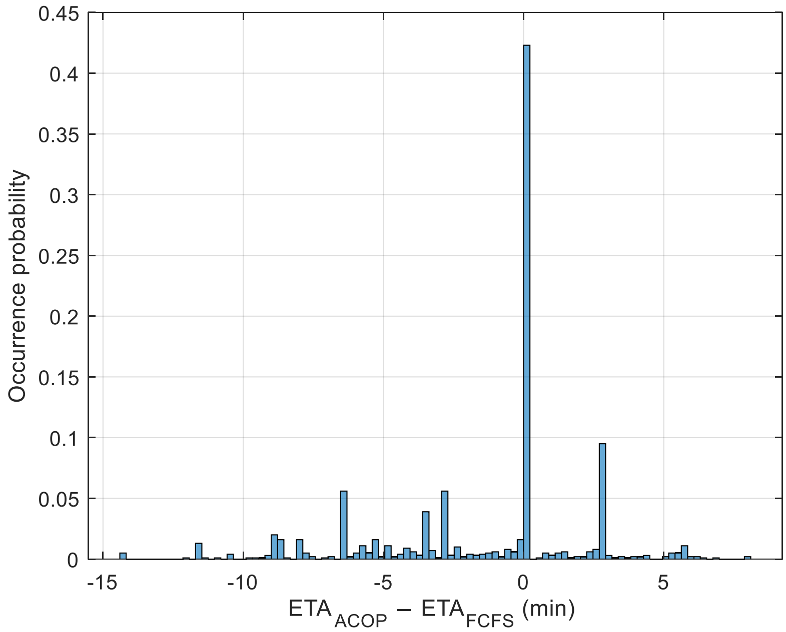

Figure 11. ACOP flight time reductions.

It has also been verified that, when ACOP provides larger flight times (in about 10% of cases), the involved aircraft headed straight to FAF (and thus excluded from optimization) came from quite far waypoints.

For both FCFS and ACOP, the predictability—defined in Section 4.3—results equal to 100%, also thanks to the constraint on the maximum delay tolerated by each aircraft (see Section 3.2 point 2 and Section 4.3). Instead, analyzing historical data from real world, this metric goes down to 98.4%.

The higher ACOP performance is obtained coming closer to the safety limits. In Figure 12, it is shown that ACOP generally provides solutions within the safety margins addressed by ICAO minima spacings defined according to the wake turbulence intensity of the preceding aircraft (Table 1). The negative values reported in the figure correspond to invalid solutions, i.e., not compliant to the safety requirements and for this reason not suggested at all.

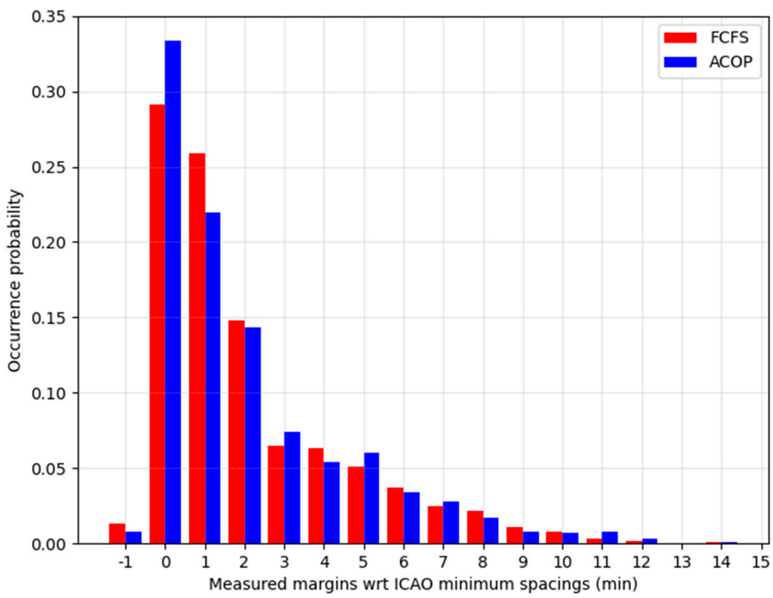

Figure 12. Safety margins over the ICAO minima separations.

4.5.2. Conflict Prevention

As reported before, ACOP has a twofold conflict prevention capability: one related to the innovative multi-path SSL feature (Section 3.4.2) and one related to CDR low-level functionality (Section 3.5).

Table 2 shows the conflict reduction capability of the SSL only. To build this table, 50 random aircraft are simulated arriving to some initial fixes with time intervals from 2 to 14 min. The FCFS and ACOP behaviors are evaluated with the CML disabled.

Table 2. Effect of the ACOP conflict reduction capability at SSL.

Arrival Time Interval (Min)	Nr. of LOSs (FCFS SSL)	Nr. of LOSs (ACOP SSL)	Difference
2	35	32	−3
3	16	15	−1
4	9	7	−2
5	5	3	−2
6	6	4	−2
7	5	4	−1
8	3	2	−1
9	2	2	0
10	1	1	0
11	1	1	0
12	1	1	0
13	1	1	0
14	1	1	0

It is shown that the number of LOS events decreases applying the ACOP strategy, at least until the aircraft come so rarely (time intervals larger than 8 min) that the FCFS order is suggested by ACOP too. The main reason why the infringements decrease using ACOP is the aircraft shorter permanence in the CTA.

On the other hand, both FCFS and ACOP use the same CML for CDR capability, in order to automatically detect the possible residual LOS and, in the case, suggest a maneuver to avoid it. In Table 3, the main CDR results are collected.

Table 3. CDR results.

Description	FCFS	ACOP
Nr. of LOSs	5	1
Nr. of LOS due to failed CDs	1	0
Nr. of LOS due to failed CRs	4	1
Nr. of detected conflicts	175	119
Percentage of successful CDRs	92%	99%
Nr. of DTW-based path selections	N/A	9

As expected, it is shown that ACOP helps decrease—under the same simulation conditions of FCFS—the LOS number from 5 to 1.

The "CD fail" cases are due to the assumed CDR "one-by-one" approach (see Section 3.2). In fact, when traffic levels are high, CDR is often hindered by the fact that the two involved aircraft may be both already involved in other conflict resolution maneuvers. In this case, no action is taken, and this could lead to a LOS. Actually, in all "CDR fail" cases, the south trombone occupancies were four or five, exceeding the maximum admissible value (Figure 8).

On the other hand, the "CR fail" cases are four in the FCFS scenario, while one in the ACOP one. The failed resolutions are due to the fact that as the traffic level is quite high and the FCFS order is too strict, the aircraft cannot decelerate at speeds lower than the lower limit of 200 kts, leading to a loss of separation.

4.5.3. CDO-Like Facilitation

As reported in Section 4.4, an effect of ACOP is to contain the number of levelings during the execution of the CDO-like vertical profiles (see Section 4.3 point 8 for their definition in this work). Note that we mainly focus on the ACOP improvements with respect to FCFS scenario; therefore, the exact definition of the CDO profiles fades into the background.

Table 4 reports the main CDO results for each scenario. It also contains some historical data provided by the Italian ANSP (ENAV) for the same airport. We do not know if these historical data are influenced by resolution maneuvers only or also by other factors (departures, ground operations, weather conditions, etc.). In any case, we can appreciate a significant improvement in the CDO metrics (see Section 4.4) using ACOP solution with respect to the available historical data.

Table 4. CDO-like results.

Description	FCFS	ACOP	Historical Data
Perc. of aircraft involved in levelings $\geq$ 20 s	8%	7.5%	54%
Nr. of levelings $\geq$ 20 s	105	85	87
Nr. of levelings < 20 s	3	0	36
Nr. of levelings imposed by STAR procedure	288	174	60
Total leveling time ($\geq$20 s)	264.8 min	335.1 min	124.5 min
Total flight time	13,546 min	12,101.7 min	726.5 min
Perc. total lev. time ($\geq$20 s) over tot. flight time	1.9%	2.8%	17.14%

The flight time reduction and the maximization of the CDO-like procedure determine tangible benefits with respect to air pollution (related to fuel consumption) and noise emissions [12].

Finally, in Figure 13, a bar chart of the altitudes reached on SENAL is reported. It is noted that the historical flights often arrive on SENAL around the minimum allowed altitude (5000 ft). Instead, in FCFS and ACOP scenarios, this altitude parameter is set to 5500 ft, an intermediate value between the minimum allowed altitude and the imposed one by previous waypoint (6000 ft) to avoid the risk of breaking the minimum altitude constraint. A small percentage of aircraft (~3.5% for both FCFS and ACOP) reach SENAL at high altitudes, i.e., >6500 ft. This happens because in certain conditions some aircraft must maintain vertical separation to avoid conflicts with lower aircraft almost up to the FAF.

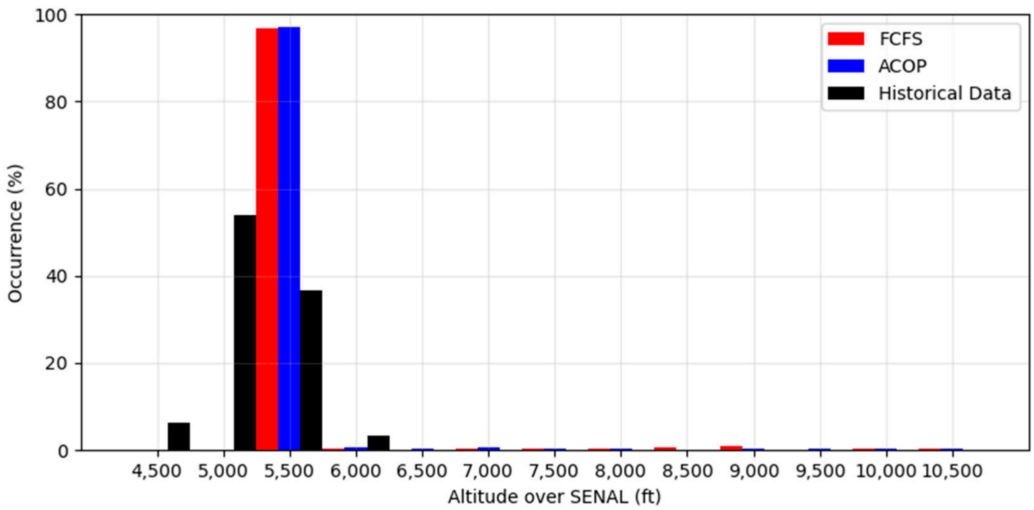

Figure 13. Occurrence probability of the altitudes reached over SENAL.

4.6. ACOP Strengths and Weaknesses

The performance comparison between ACOP and the reference strategy FCFS has shown the following pros and cons.

Strengths:

1. Reduced flight times and increased runway throughput, thanks to the global optimal sequencing;
2. Reduced infringements, thanks to the DTW-based conflict reduction capability;
3. Optimized resolution maneuvers, thanks to reinforcement learning approach;
4. The aircraft paths are known in advance because the possible resolution maneuvers do not involve vectoring or path stretching, resulting in a better situational awareness;
5. Adherence to CDO-like procedures, i.e., reduced number of levelings;
6. Adaptability to any airport, regardless of its size;
7. Very low computational load.

Weaknesses:

1. Information about the airport STAR procedure must be collected before use;
2. Pairwise conflict acceptance and workload related to the use of ACOP are not yet evaluated;
3. The controller acceptance and workload related to the use of ACOP are not yet evaluated.

5. Operational Aspects

ACOP provides two types of suggestions: "Direct To" commands and "Conflict resolution" commands.

"Direct To" commands are required to be issued by the controller within the designated countdown period, as indicated by the supplementary output from ACOP. The statistical distribution of this "remaining time" is reported in Figure 14.

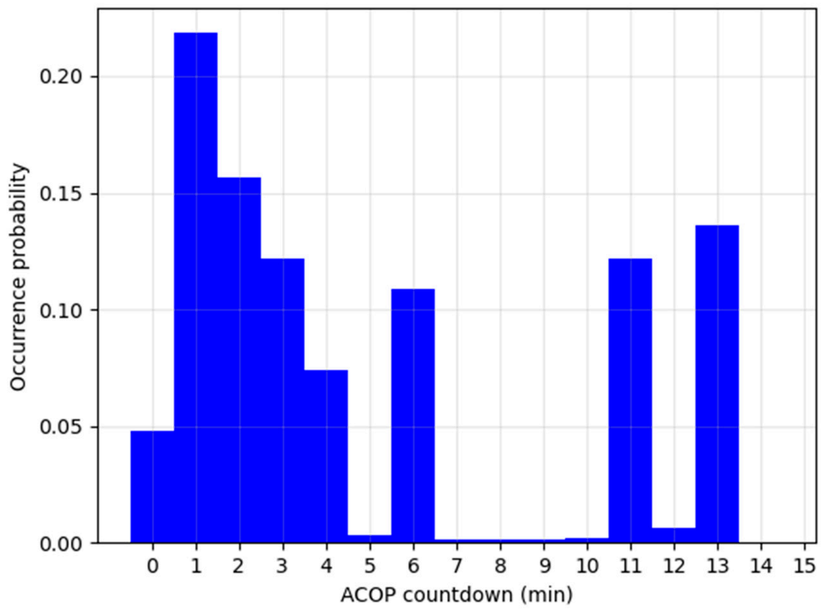

Figure 14. Histogram of the expiration times for the ACOP "Direct To" suggestions.

It is worth noting that the time for execution is considerable, showing an average value of 5.7 min. It is probable that these values exceed the real delays between the suggestion of the commands and their implementation. In any case, in the event of a missed implementation, the aircraft sequencing would remain unchanged, resulting solely in a degradation of system-level metrics, but would not pose any safety concerns.

"Conflict resolution" commands are supposed immediately actuated (Section 4.3, point 10). However, in order to consider current operational aspects of air traffic management and account for limited levels of automation, a dedicated sensitivity analysis was carried out to evaluate and quantify the robustness of the ACOP approach to CDR with regards to delays resulting from human–machine interaction. Specific evaluations of human factors were not conducted in this study. However, conservative assumptions concerning the response time of air traffic controllers and pilots can be derived from prior research.

Table 5 summarizes the average response times derived from a review of several literature papers [57–60].

Table 5. Characteristic response times.

Actor	Response Time (s)
Pilot	between 10 and 25
Controller	between 5 and 10
Voice communications	less than 10
Transmission delay	less than 1

By utilizing the information reported in the table above, it was assumed that the actuation of ACOP suggestions may occur with a latency ranging from 15 to 55 s, modeled as a uniform distribution. This delay encompasses various factors, including the behaviors of controllers and pilots, communication latency, and flight dynamics transients.

The sensitivity analysis was therefore carried out considering five settings as detailed in Table 6. The first setting represents the nominal configuration, in which there is no latency, and both the simulation engine and ACOP respond immediately to the computed suggestions. The next two settings account for a simulated latency modeled as a uniform distribution with values ranging from 15 to 35 s and from 35 to 55 s. These latencies affect the behavior only of the simulation engine. Therefore, in these cases, while the controller and pilot require some time to implement the suggestions, the digital assistant operates as if the commands are issued instantly. The remaining two settings allow ACOP to account for the expected latency (a mean latency of 20 s and 40 s, respectively), potentially enhancing its robustness to response times.

Table 6. Latency sensitivity analysis (ACOP only).

Simulated Latency (s)	ACOP Average Latency Parameter (s)	Nr. of LOS
0	0	1
[15, 35] unif. distrib.	0 and 25	2 and 1
[35, 55] unif. distrib.	0 and 45	7 and 6

Table 6 reports the number of LOS incidents associated with the delay in executing the resolution maneuvers, using the same simulation configurations. The first row reports the conditions under which the results of Section 4.5 are obtained.

The results of the sensitivity analysis demonstrate that incorporating the knowledge of the average latency, if available, enhances the performance of the system (a decrease from 2 to 1 in the number of LOS incidents). However, this improvement is contingent on the latency not being too large, as beyond a certain threshold (approximately 35 s), the LOS becomes often unrecoverable.

6. Future Work

Possible future improvements of this work are the application on another airport (to validate the adaptability of the proposed digital assistant) and the removal of certain assumptions under which ACOP has been built.

Furthermore, in a future operational scenario, such a tool will accept real traffic data and output one or more actions that represent suggestions to the controller while optimizing, always in compliance with safety requirements:

1. Financial aspects (throughput, fuel consumption, delays);
2. Environmental aspects (air pollution and noise pollution).

Finally, some real-time validation sessions with air traffic controllers could be planned to demonstrate the operational effectiveness of the proposed approach.

7. Conclusions

In this paper, a digital assistant—named ACOP (Arrival digital assistant with COnflict Prevention)—has been proposed to assist air traffic controllers in the complex task of arrival traffic management. The operational requirements of ACOP have been provided by the Italian ANSP (ENAV), acting as an end-use.

The proposed solution provides near real-time suggestions for dealing with the following problems:

1. Sequencing and scheduling, which consists of determining the sequence of aircraft and their arrival times at intermediate (or final) approach fix, in order to make the best use of the terminal area space and to increase overall runway throughput and capacity. The provided solution must always meet the ICAO requirements for minima longitudinal separations, depending on the type of aircraft pairs (Table 1).
2. Conflict detection and resolution (CDR), which consists of preventing loss of separations (LOS), i.e., simultaneous violations of ICAO minima horizontal and vertical separation distances, assumed equal to 5 NM and 1000 ft, respectively. To achieve this objective, each aircraft was assumed to be able to decelerate and/or level its flight (the latter action is used the minimum necessary, to maximize compliance with CDO-like procedures).

For the first problem, a genetic algorithm has been used, while for the second problem, a reinforcement learning approach has been chosen.

The performance evaluation of ACOP is carried out using a simulation environment that assumes a continuous traffic flow of 1000 aircraft generated at the initial fixes at random instants, with random altitudes and categories. The assumed traffic volume has been found to be approximately ten times the historical volume at the same airport. The ACOP performances are then evaluated with respect to the FCFS reference strategy applied under the same simulation conditions.

The obtained numerical results are the following:

1. The average flight time to FAF of each aircraft is reduced by 1.44 min, with a cumulative time saving of more than 24 h. In addition, the average time length of the arrival queue is reduced by about 11.6 s. This better use of airspace opens the possibility of increasing runway capacity and throughput by properly tuning the upstream traffic.
2. Regarding the safety aspects, it should be noted that:
   - Thanks to the ACOP conflict reduction capability of the SSL—performed by the dynamic time warping (DTW) technique—the number of conflict detections decreases from 175 to 119, while the percentage of the successful CDRs increases from 92% to 99%;
   - The applicable ICAO safety rules are always respected, except for one case, where both aircraft were already involved in a conflict resolution maneuver. This situation can be tolerated because it is strongly related to the traffic level, which is clearly too high for the adopted airport.
3. Regarding the environmental aspects focusing on CDO, the leveled segments in the CDO-like procedures are minimized to about 2%.
4. The short ACOP response time (~1 s) and the high usability of the user interface is also an important feature, which is particularly promising with respect to the next real applications.

5. Finally, a dedicated sensitivity analysis was conducted to evaluate ACOP's robustness to CDR with respect to potential delays resulting from human–machine interaction.

Author Contributions: Conceptualization: F.N., R.P., G.G.; writing and software: F.N., R.P.; review: G.G., A.V. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Italian Ministry of Research within the Italian National Aerospace Research Program (PRORA) managed by CIRA, project NUVASC.

Data Availability Statement: Restrictions apply to the availability of used data. Data was obtained from ENAV and are available on request to the authors with the permission of ENAV.

Acknowledgments: We thank L. Babusci, G. Berti and D. Teotino (ENAV, the Italian ANSP) for their support in defining the ACOP operational concept and for the historical data used for its performance assessment.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

ACOP	Arrival digital assistant with COnflict Prevention
AI	Artificial intelligence
ANSP	Air navigation services provider
APP	Approach control unit
ASSP	Arrival sequencing and scheduling problem
ATC	Air traffic control
ATM	Air traffic management
ATCO	ATC Operator
CDO	Continuous descent operation
CDR	Conflict detection and resolution
CML	Conflict management layer
CPA	Closest point of approach
CR	Conflict resolution
CPS	Constrained positioning shifting
CTA	Control area
CTSP	Cumulative traveling salesman problem
DPS	Dynamic positioning shifting
DQN	Deep Q-network
DTG	Distance to go
DTW	Dynamic time warping
ENAV	Ente Nazionale per l'Assistenza al Volo
ETA	Estimated time of arrival
FAA	Federal Aviation Administration
FAF	Final approach fix
FCFS	First come, first served
FL	Flight level
FPA	Flight path angle
GLS	Genetic local search
IAS	Indicated air speed
ICAO	International Civil Aviation Organization
KPI	Key performance indicator
LOS	Loss of separation
ML	Machine learning
MPC	Model predictive control
MST	Mean separation time
NextGen	Next generation air transportation system
RL	Reinforcement learning
RNAV	aRea NAVigation
SESAR	Single European Sky ATM Research
SSL	Sequencing and scheduling layer

STAR	Standard terminal arrival route
TMA	Terminal maneuvering area
TOD	Top of descent
TSP	Travelling salesman problem

References

1. Kistan, T.; Gardi, A.; Sabatini, R.; Ramasamy, S.; Batuwangala, E. An Evolutionary Outlook of Air Traffic Flow Management Techniques. *Prog. Aerosp. Sci.* **2017**, *88*, 15–42. [CrossRef]
2. Anagnostakis, I.; Clarke, J.-P.; Bohme, D.; Volckers, U. Runway Operations Planning and Control Sequencing and Scheduling. In Proceedings of the 34th Annual Hawaii International Conference on System Sciences, Maui, HI, USA, 6 January 2001. [CrossRef]
3. Artiouchine, K.; Baptiste, P.; Dürr, C. Runway Sequencing with Holding Patterns. *Eur. J. Oper. Res.* **2008**, *189*, 1254–1266. [CrossRef]
4. Bäuerle, N.; Engelhardt-Funke, O.; Kolonko, M. On the Waiting Time of Arriving Aircrafts and the Capacity of Airports with One or Two Runways. *Eur. J. Oper. Res.* **2007**, *177*, 1180–1196. [CrossRef]
5. Xu, B.; Ma, W.; Huang, H.; Yue, L. Weighted Constrained Position Shift Model for Aircraft Arrival Sequencing and Scheduling Problem. *Asia-Pac. J. Oper. Res.* **2016**, *33*, 1650028. [CrossRef]
6. EASA. Easy Access Rules for ATM-ANS (Regulation (EU) 2017/373). 2022. Available online: https://www.easa.europa.eu/en/document-library/easy-access-rules/online-publications/easy-access-rules-atm-ans-regulation-eu?page=2#_DxCrossRefBm2104645389 (accessed on 27 March 2023).
7. Ma, J.; Delahaye, D.; Sbihi, M.; Scala, P.; Mujica Mota, M.A. Integrated Optimization of Terminal Maneuvering Area and Airport at the Macroscopic Level. *Transp. Res. Part C Emerg. Technol.* **2019**, *98*, 338–357. [CrossRef]
8. Cao, J.-M.; Kanafani, A. The Value of Runway Time Slots for Airlines. *Eur. J. Oper. Res.* **2000**, *126*, 491–500. [CrossRef]
9. Hardell, H.; Lemetti, A.; Polishchuk, T.; Smetanová, L. Evaluation of the Sequencing and Merging Procedures at Three European Airports Using Opensky Data. *Eng. Proc.* **2022**, *13*, 13. [CrossRef]
10. Sáez, R.; Prats, X.; Polishchuk, T.; Polishchuk, V. Traffic Synchronization in Terminal Airspace to Enable Continuous Descent Operations in Trombone Sequencing and Merging Procedures: An Implementation Study for Frankfurt Airport. *Transp. Res. Part C Emerg. Technol.* **2020**, *121*, 102875. [CrossRef]
11. Post, W.; de Jonge, H.W.G. *Free Flight in a Ground-Controlled ATM Environment*; National Aerospace Laboratory NLR: Amsterdam, The Netherlands, 1997.
12. EUROCONTROL. *European CCO/CDO Action Plan*; Eurocontrol: Brussels, Belgium, 2020.
13. Clarke, J.-P.B.; Ho, N.T.; Ren, L.; Brown, J.A.; Elmer, K.R.; Tong, K.-O.; Wat, J.K. Continuous Descent Approach: Design and Flight Test for Louisville International Airport. *J. Aircr.* **2004**, *41*, 1054–1066. [CrossRef]
14. Coppenbarger, R.A.; Mead, R.W.; Sweet, D.N. Field Evaluation of the Tailored Arrivals Concept for Datalink-Enabled Continuous Descent Approach. *J. Aircr.* **2009**, *46*, 1200–1209. [CrossRef]
15. Weitz, L.A.; Hurtado, J.E.; Barmore, B.E.; Krishnamurthy, K. An Analysis of Merging and Spacing Operations with Continuous Descent Approaches. In Proceedings of the 24th Digital Avionics Systems Conference, Washington, DC, USA, 30 October 2005–3 November 2005; Volume 1, pp. 2–21. [CrossRef]
16. Knote, R.; Söllner, M.; Leimeister, J.M. Towards a Pattern Language for Smart Personal Assistants. In Proceedings of the 25th Conference on Pattern Languages of Programs, PLoP '18, Portland, OR, USA, 25 October 2018; 2020; pp. 1–16.
17. Su, H.; Liu, S.; Zheng, B.; Zhou, X.; Zheng, K. A Survey of Trajectory Distance Measures and Performance Evaluation. *VLDB J.* **2020**, *29*, 3–32. [CrossRef]
18. EUROCONTROL. *European ATM Master Plan-SESAR Joint Undertaking*; Eurocontrol: Brussels, Belgium, 2015.
19. Federal Aviation Administration. *NextGEN AVS Work Plan 2016*; Federal Aviation Administration: Washington, DC, USA, 2016.
20. Bianco, L.; Dell'Olmo, P.; Giordani, S. Minimizing Total Completion Time Subject to Release Dates and Sequence-dependentprocessing Times. *Ann. Oper. Res.* **1999**, *86*, 393–415. [CrossRef]
21. Erzberger, H. Design Principles and Algorithms for Automated Air Traffic Management. *Change* **1995**, *7*, 2.
22. Psaraftis, H.N. A Dynamic Programming Approach for Sequencing Groups of Identical Jobs. *Oper. Res.* **1980**, *28*, 1347–1359. [CrossRef]
23. Trivizas, D.A. Optimal Scheduling with Maximum Position Shift (MPS) Constraints: A Runway Scheduling Application. *J. Navig.* **1998**, *51*, 250–266. [CrossRef]
24. Bianco, L.; Dell'Olmo, P.; Giordani, S. Scheduling Models for Air Traffic Control in Terminal Areas. *J. Sched.* **2006**, *9*, 223–253. [CrossRef]
25. Malaek, S.M.B.; Naderi, E. A New Scheduling Strategy for Aircraft Landings under Dynamic Position Shifting. In Proceedings of the 2008 IEEE Aerospace Conference, Big Sky, MT, USA, 1–8 March 2008; pp. 1–8. [CrossRef]
26. Bennell, J.A.; Mesgarpour, M.; Potts, C.N. Airport Runway Scheduling. *4OR* **2011**, *9*, 115–138. [CrossRef]
27. Samà, M.; D'Ariano, A.; Corman, F.; Pacciarelli, D. Metaheuristics for Efficient Aircraft Scheduling and Re-Routing at Busy Terminal Control Areas. *Transp. Res. Part C Emerg. Technol.* **2017**, *80*, 485–511. [CrossRef]
28. D'Ariano, A.; Pacciarelli, D.; Pistelli, M.; Pranzo, M. Real-Time Scheduling of Aircraft Arrivals and Departures in a Terminal Maneuvering Area. *Networks* **2015**, *65*, 212–227. [CrossRef]

29. Samà, M.; D'Ariano, A.; D'Ariano, P.; Pacciarelli, D. Optimal Aircraft Scheduling and Routing at a Terminal Control Area during Disturbances. *Transp. Res. Part C Emerg. Technol.* **2014**, *47*, 61–85. [CrossRef]
30. Tavakkoli-Moghaddam, R.; Yaghoubi-Panah, M.; Radmehr, F. Scheduling the Sequence of Aircraft Landings for a Single Runway Using a Fuzzy Programming Approach. *J. Air Transp. Manag.* **2012**, *25*, 15–18. [CrossRef]
31. Ernst, A.T.; Krishnamoorthy, M.; Storer, R.H. Heuristic and Exact Algorithms for Scheduling Aircraft Landings. *Networks* **1999**, *34*, 229–241. [CrossRef]
32. Bencheikh, G.; Khoukhi, F.; Baccouche, M.; Boudebous, D.; Belkadi, A.; Ouahman, A. Hybrid Algorithms for the Multiple Runway Aircraft Landing Problem. *Int. J. Comput. Sci. Appl.* **2013**, *10*, 53–71.
33. Hansen, J.V. Genetic Search Methods in Air Traffic Control. *Comput. Oper. Res.* **2004**, *31*, 445–459. [CrossRef]
34. Hu, X.-B.; Di Paolo, E. An Efficient Genetic Algorithm with Uniform Crossover for Air Traffic Control. *Comput. Oper. Res.* **2009**, *36*, 245–259. [CrossRef]
35. Liu, Y.-H. A Genetic Local Search Algorithm with a Threshold Accepting Mechanism for Solving the Runway Dependent Aircraft Landing Problem. *Optim. Lett.* **2011**, *5*, 229–245. [CrossRef]
36. Deau, R.; Gotteland, J.-B.; Durand, N. *Airport Surface Management and Runways Scheduling*; ATM Seminar: Savannah, GA, USA, 2009.
37. Samà, M.; D'Ariano, A.; D'Ariano, P.; Pacciarelli, D. Scheduling Models for Optimal Aircraft Traffic Control at Busy Airports: Tardiness, Priorities, Equity and Violations Considerations. *Omega* **2017**, *67*, 81–98. [CrossRef]
38. Samà, M.; D'Ariano, A.; Pacciarelli, D. Rolling Horizon Approach for Aircraft Scheduling in the Terminal Control Area of Busy Airports. *Transp. Res. Part E Logist. Transp. Rev.* **2013**, *60*, 140–155. [CrossRef]
39. Kuchar, J.K.; Yang, L.C. A Review of Conflict Detection and Resolution Modeling Methods. *IEEE Trans. Intell. Transp. Syst.* **2000**, *1*, 179–189. [CrossRef]
40. Yang, Y.; Zhang, J.; Cai, K.-Q.; Prandini, M. Multi-Aircraft Conflict Detection and Resolution Based on Probabilistic Reach Sets. *IEEE Trans. Control Syst. Technol.* **2017**, *25*, 309–316. [CrossRef]
41. Hao, S.; Cheng, S.; Zhang, Y. A Multi-Aircraft Conflict Detection and Resolution Method for 4-Dimensional Trajectory-Based Operation. *Chin. J. Aeronaut.* **2018**, *31*, 1579–1593. [CrossRef]
42. Yokoyama, N. Decentralized Conflict Detection and Resolution Using Intent-Based Probabilistic Trajectory Prediction. In *2018 AIAA Guidance, Navigation, and Control Conference*; AIAA SciTech Forum; American Institute of Aeronautics and Astronautics: Reston, VA, USA, 2018. [CrossRef]
43. Jilkov, V.P.; Ledet, J.H.; Li, X.R. Multiple Model Method for Aircraft Conflict Detection and Resolution in Intent and Weather Uncertainty. *IEEE Trans. Aerosp. Electron. Syst.* **2019**, *55*, 1004–1020. [CrossRef]
44. Radanovic, M.; Piera Eroles, M.A.; Koca, T.; Ramos Gonzalez, J.J. Surrounding Traffic Complexity Analysis for Efficient and Stable Conflict Resolution. *Transp. Res. Part C Emerg. Technol.* **2018**, *95*, 105–124. [CrossRef]
45. Sutton, R.S.; Barto, A.G. *Reinforcement Learning: An Introduction*, 2nd ed.; A Bradford Book; The MIT Press: Cambridge, MA, USA, 2015.
46. Wang, Z.; Pan, W.; Li, H.; Wang, X.; Zuo, Q. Review of Deep Reinforcement Learning Approaches for Conflict Resolution in Air Traffic Control. *Aerospace* **2022**, *9*, 294. [CrossRef]
47. The Mathworks, Inc. *MATLAB Version 9.10.0.1613233 (R2021a)*; The Mathworks, Inc.: Natick, MA, USA, 2021.
48. Goldberg, D.E. *Genetic Algorithms in Search, Optimization and Machine Learning*, 1st ed.; Addison-Wesley Longman Publishing Co., Inc.: Boston, MA, USA, 1989.
49. Errico, A.; Di Vito, V. Methodology for Estimation of Closest Point of Approach between Aircraft in ATM. In *AIAA Aviation 2019 Forum*; AIAA AVIATION Forum; American Institute of Aeronautics and Astronautics: Reston, VA, USA, 2019. [CrossRef]
50. Van Rossum, G.; Drake, F.L. *Python 3 Reference Manual*; CreateSpace: Scotts Valley, CA, USA, 2009.
51. Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; Zaremba, W. OpenAI Gym. *arXiv* **2016**. Available online: http://arxiv.org/abs/1606.01540 (accessed on 11 February 2023).
52. Raffin, A.; Hill, A.; Gleave, A.; Kanervisto, A.; Ernestus, M.; Dormann, N. Stable-Baselines3: Reliable Reinforcement Learning Implementations. *J. Mach. Learn. Res.* **2021**, *22*, 12348–12355.
53. ENAV. AIP-Italia AD 2 LIEO 4-1. 2020. Available online: https://www.enav.it/services/list/aip (accessed on 27 March 2023).
54. ENAV. AIP-Italia ENR2.1.1.4.5-2. 2018. Available online: https://www.enav.it/services/list/aip (accessed on 27 March 2023).
55. ICAO. *Manual of the ICAO Standard Atmosphere: Extended to 80 Kilometres (262 500 Feet)*, 3rd ed.; International Civil Aviation Organization: Montreal, QC, Canada, 1993.
56. APACHE Project Consortium. *Review of Current KPIs and Proposal for New Ones*; SESAR Joint Undertaking: Brussels, Belgium, 2017.
57. Feigh, K.M. An Airspace Simulator for Air Traffic Management Research. Master's Thesis, Department of Power, Propulsion and Aerospace Engineering, Cranfield University, Bedford, UK, 2003.
58. Blom, H.A.; Bakker, B.G. Can Ground-Based Separation Accommodate Very High En Route Traffic Demand as Well as Advanced Self-Separation? In Proceedings of the 15th AIAA Aviation Technology, Integration, and Operations Conference, Dallas, TX, USA, 22–26 June 2015; p. 3180.

59. Ziccardi, J.; Roberts, Z.; O'Connor, R.; Rorie, C.; Morales, G.; Battiste, V.; Strybel, T.; Chiappe, D.; Vu, K.-P.L.; Shively, J. Measuring UAS Pilot Responses to Common Air Traffic Clearances. In *Human Interface and the Management of Information. Information and Interaction for Health, Safety, Mobility and Complex Environments*; Yamamoto, S., Ed.; Springer Berlin Heidelberg: Berlin/Heidelberg, Germany, 2013; pp. 606–612.
60. Palumbo, R.; Errico, A.; Pascarella, D.; Gargiulo, F.; Filippone, E. Modeling Approach for Resilience Engineering of the Future ATM System. In Proceedings of the 15th AIAA Aviation Technology, Integration, and Operations Conference, Dallas, TX, USA, 22–26 June 2015.