Optimising Flight Trajectories under Thunderstorm Conditions

Huijuan Yang
OPTIM
École Nationale de l'Aviation Civile
Toulouse, France
School of Airport Engineering
Civil Aviation Flight University of China
Guanghan, China
huijuan.yang@enac.fr

Daniel Delahaye
OPTIM
École Nationale de l'Aviation Civile
Toulouse, France
delahaye@recherche.enac.fr

Ying Huo
Department of General Aviation
Civil Aviation Management Institute of China
Beijing, China
huoying@camic.cn

Li Lu
School of Air Traffic Management
Civil Aviation Flight University of China
Guanghan, China
luli@cafuc.edu.cn

***Abstract*—Given the operational disruptions caused by thunderstorms, this paper introduces an innovative approach to flight trajectory optimisation within the context of air traffic management, with a particular focus on minimising disruptions due to adverse weather conditions. The proposed framework integrates wake turbulence separation requirements, node conflict evaluations, and real-time operational constraints to address the complexities inherent in trajectory planning. Unlike previous studies that addressed stochastic optimal control problems for conflict resolution offline, this study presents an enhanced Selective Simulated Annealing (SSA) algorithm for online conflict mitigation. A case study conducted in the approach area of Chengdu Shuangliu International Airport (CTU), involving over 900 flights, demonstrates the applicability and effectiveness of the proposed algorithm. This result underscores the algorithm's potential as a powerful tool for trajectory optimisation, offering significant improvements in operational performance in a timely manner. Additionally, sensitivity analysis of the slot shift ranges confirms the robustness of the selected parameters, highlighting their role in minimising delays and maintaining system efficiency. By deactivating decisions individually, this study also discusses the sensitivity of the algorithm to the proposed decision combinations.**

***Keywords*—Sequencing and Merging; Aircraft Trajectory Optimisation; Thunderstorm Conditions; Selective Simulated Annealing**

I. Introduction

The inherently uncertain nature of convective weather, commonly referred to as storms or thunderstorms, poses significant safety risks to aviation. Intense, conflicting updrafts and downdrafts within thunderstorms can result in severe turbulence. Even when a storm is not fully mature, flying over a developing storm can lead to substantial turbulence due to strong updrafts above the visible cloud top. Additionally, hazards such as hail, icing, and lightning can inflict considerable damage to aircraft equipment and windshields when flying through a thunderstorm [1].

When encountering thunderstorms, aircraft can safely fly over them only if their altitude is well above the turbulent cloud tops. However, the most intense and turbulent storms are often the tallest storms [2]. To avoid strong convective activities, pilots are instructed to deviate from their initial flight plans and navigate around them. These multiple deviations can cause significant air traffic disruptions, including blockage and overcrowding of neighbouring airspace, aircraft holding patterns, delays, and cancellations. Such disruptions pose substantial challenges to operational efficiency and lead to reduced time and cost efficiencies [3].

Moreover, weather has become the largest cause of air traffic delay in the United States, accounting for 75.48 % of system-impacting delays of more than 15 minutes from 2017 to 2022 [2]. In Europe, adverse weather accounted for 29.7% of the en-route Air Traffic Flow Management (ATFM) delay [4]. Given that ATFM delay costs 100 euros per minute [5], effective aircraft trajectory planning and management at the tactical level is crucial in daily airline and air traffic control operations.

Trajectory optimisation is a critical component of aviation and aerospace operations, focusing on determining the most efficient path for an aircraft to follow. For instance, Dalmau and Prats [6] optimised aircraft trajectory using the Cost Index (CI) to reduce fuel consumption, trip time and gaseous emissions during the eventual continuous cruise climb operations. A multiple-phase Optimal Control (OC) approach is also

employed to model different flight phases, each with its own constraints or flight modes. This approach relies on various methods for phase separation, such as the knotting method [7]. Other formulations have been proposed, including dynamic programming [8], where the trajectory is characterised as a set of static parameters optimised by assuming a typical trajectory pattern.

More recently, trajectory optimisation was extended to stochastic approaches to better handle uncertainty. Air Traffic Management (ATM) applications include optimisation with conflict resolution in the presence of positional uncertainty of the aircraft, convective weather and weather forecast uncertainties. For instance, Matsuno et al. [9] proposed a stochastic optimal control method for three-dimensional conflict-free aircraft trajectory optimisation under wind uncertainty. Their focus was on the wind correlation and its effects on aircraft separation. Considering the computational complexity, they suggested solving the stochastic optimal control problem for conflict resolution offline, allowing the computation of aircraft trajectory, separation, and current conflict probability in real time to apply to daily operations. They also highlighted that the accuracy and convergence rate of the proposed framework largely depend on the wind model, noting that assuming a time-invariant wind error neglects the temporal variations. Hentzen et al. [1] accounted for meteorological uncertainty with a tailored stochastic storm model and presented an optimal control algorithm to maximise the probability of reaching a waypoint while avoiding hazardous storm regions. García-Heras et al. [10] proposed a robust open-loop optimal control methodology to reduce reroutings during flight execution, incorporating meteorological uncertainties at the flight planning level. Andrés et al. [11] proposed a novel heuristic approach to address the aircraft trajectory optimisation problem, incorporating uncertainties in the evolution of convective cells. They adopted the Augmented Random Search (ARS) algorithm to search for the optimal path, testing their method over 4,000 iterations to reveal its convergence. Sáez et al. [12] present a framework for flight trajectory optimisation in the en-route phase with several weather constraints to reduce the environmental impact, while increasing the resilience of air operations to weather phenomena. González-Arribas et al. [13] introduced a novel simulation-based flight planning methodology for weather-optimal 4D flight trajectory options under uncertain meteorological and operational context. Baneshi et al. [14] accounted for anthropogenic climate change, and proposed an environmental-oriented aircraft trajectory optimisation framework. Their experiment with 1,005 flights indicates that climate-optimal trajectories inevitably increased the operational cost and the number of conflicts. Resolving these conflicts required compromises between climate impact and cost. Based on robust tracking optimal control theory, Simorgh et al. [15] considered climate impact and introduced an aircraft trajectory optimisation framework that accounts for meteorological uncertainty.

While most incidents and accidents occur during the approach and landing phases at the Terminal Manoeuvring Area (TMA), previous literature has seldom investigated flight trajectory optimisation during those critical stages, particularly under extreme weather conditions.

The TMA is a highly congested airspace where trajectory deviations can easily lead to conflicts with other aircraft. Additionally, TMAs often cover densely populated areas, making safety and noise abatement critical considerations when navigating through them. Thunderstorm cells passing through a TMA can disrupt Standard Instrument Departures (SID) and Standard Terminal Arrival Routes (STAR) for extended periods (e.g., 20 to 30 minutes). For departures, this issue is typically managed by holding aircraft at the gate. For arrivals, Air Traffic Control Officers (ATCOs) often vector aircraft affected by weather along alternative routes, guiding them to the runways through successive rerouting manoeuvres. Another option for aircraft further upstream in the sequence is to enter holding patterns until the weather clears and the STARs become available again.

This paper addresses this gap by designing and implementing a generalised framework for aircraft trajectory optimisation, identifying safe detours for approach and landing while minimising a weighted combination of flight time and fuel consumption, and accounting for the uncertainty of thunderstorm development regions. The objective of this research is to develop a decision support tool to assist ATCOs in managing aircraft during such situations. Leveraging an optimisation algorithm, this tool will assign aircraft to Mach regulation, rerouting within the STAR network, and speed regulation within the TMA to minimise aircraft holding times.

The subsequent parts of this paper are organised as follows: Section II introduces the mathematical formulation, and Section III explains the design and implementation of the optimisation algorithm. Section IV presents the computational experiments to validate the applicability and effectiveness of the proposed model. Finally, Section V discusses the findings, contributions, limitations and future works of the study, and concludes the research.

II. The integrated Mixed Integer Programming model

This section presents an integrated Mixed Integer Programming (MIP) model designed to adjust flight trajectories under thunderstorm conditions. The model employs a graph-based approach, where the aircraft trajectories are represented by straight segments, referred to as links (l), which are connected by intersections, denoted by nodes (n). For each flight f, a route (r_f) corresponds to a sequence of connected flight links.

Thunderstorms are incorporated as constraints that influence flight path availability. Specifically, thunderstorms are modelled as dynamic no-fly zones or restricted airspace regions, which are translated into spatial and temporal constraints, such as link blockages during specific time periods, preventing flights from entering the affected areas.

Lastly, weather-induced delays are captured by adjusting aircraft speeds and route entry times, accounting for deviations or holding patterns required to avoid thunderstorms.

A. Decision variables

For each flight $f \in \mathcal{F}$, several key decisions are made based on its location and speed during thunderstorms. These decisions include the route (r_f), route entry time (t_f), and speed (v_f):

- r_f: The route assigned to flight f, selected from a predetermined set of routes ($\mathcal{R}$).
  The route set $\mathcal{R}$ is defined by the available links and nodes at a given time, considering weather conditions, various link combinations, as well as route lengths. The availability of a route is determined by the location of the aircraft and the accessibility of links at specific times. The shortest route from $\mathcal{R}$ is then selected and assigned to ensure operational efficiency during deviations.

$$r_f = \begin{cases} 1 & \text{if flight } f \text{ chooses this route option,} \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

- t_f: The assigned route entry time of flight f, adjusted by the number of time slots Δt.
  The route entry time can be shifted either forward or backward. The maximum allowable earliness and tardiness relative to the initial route entry time are denoted as $t_f^{\min}$ and $t_f^{\max}$, respectively. This study assumes that $|t_f^{\max}| \geq |t_f^{\min}|$, reflecting the fact that there is generally more margin for slowing down an aircraft during the en-route phase (Mach regulation) than for speeding it up.
  Mathematically, the assigned route entry time can be formulated as:

$$t_f \in T_f := \{t_f^o + i \times \Delta t \mid i \in \mathbb{Z},\, t_f^{\min} \leq i \times \Delta t \leq t_f^{\max}\} \tag{2}$$

  where T_f represents the set of all possible route entry times derived from the initial time (t_f^o) for a flight f by adding or subtracting multiples of Δt, constrained within the bounds $t_f^{\min}$ and $t_f^{\max}$.

- v_f: The assigned route entry speed of flight f, represented by a user-defined speed increment Δv with respect to the initial route entry speed (v_f^o).
  The minimum and maximum allowable speeds for flight f are denoted by $v_f^{\min}$ and $v_f^{\max}$, respectively. The route entry speed can be formulated as follows:

$$\begin{aligned} v_f \in V_f := \{&v_f^o + j \times \Delta v \mid \\ &j \in \mathbb{Z},\, v_f^{\min} \leq v_f^o + j \times \Delta v \leq v_f^{\max}\} \end{aligned} \tag{3}$$

  where V_f denotes the set of all possible speeds for flight f.
  This formulation defines a range of possible speeds for a flight, based on small adjustments around a nominal speed, while ensuring that these adjustments stay within a safe and allowable range.

While the variables are selected from sets of possible values, they are treated as decision vectors in the optimisation process, respectively. For instance, t_f is selected from the set T_f, which contains multiple possible entry times, each adjusted by integer multiples of Δt. This flexibility reflects the vector nature of the decision-making process. Similarly, v_f is chosen from the set V_f, which includes a range of speed options around the nominal speed. This range-based selection further reinforces the vector formulation of v_f. Subsequently, the decision vector is denoted as $\mathbf{x} = (r_f, t_f, v_f)$, where r_f represents the route vector, t_f represents the time vector, and v_f represents the speed vector.

B. Constraints

Two types of conflicts are detected to maintain safety operations during rerouting, including link conflicts and node conflicts.

1) Conflicts on links: Wake turbulence separation requirements are employed to detect link conflict and ensure operational safety in each segment. A violation of these rules leads to a link conflict.

For two consecutive aircraft, f and g ($f, g \in \mathcal{F}$), using the same link l, the wake turbulence separation requirements (s_{fg}) are based on the aircraft categories (c_f and c_g) (see Table I).

Although separations are typically denoted in Nautical Miles (NM), link conflicts are evaluated based on the time (t_f^l and t_g^l) and speed (v_f^l and v_g^l) when aircraft enter the link, as expressed in Equation (4).

$$C_{fg}^l(\mathbf{x}) = \begin{cases} 1 & \text{if } \left[(t_g^l \times v_g^l - t_f^l \times v_f^l < s_{fg})\right. \\ & \text{or } \left.(t_f^l \times v_f^l - t_g^l \times v_g^l < s_{fg})\right] \\ & \text{and } l_f = l_g, \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

where $C_{fg}^l(\mathbf{x})$ equals 1, if flights f and g use the same link and the separation requirement is not met. Consequently, the minimum wake turbulence separation is respected when $\sum_{\substack{(f,g)\in\mathcal{F}\times\mathcal{F} \\ f\neq g}} C_{fg}^l(\mathbf{x}) = 0$.

TABLE I. WAKE TURBULENCE SEPARATION MINIMA FOR DIFFERENT CATEGORIES OF AIRCRAFT (IN NM).

Operation-category		Trailing aircraft		
		Light	Medium	Heavy
Leading aircraft	Light	3	3	3
	Medium	4	3	3
	Heavy	6	5	4

2) Conflicts on nodes: A node conflict occurs when two consecutive aircraft, f and g ($f, g \in \mathcal{F}$), pass through the same node at the same time. To prevent node conflicts, a protection area is established around each node. This area is represented by a disc with a radius, denoted as R, equal to the required separation distance - typically 3 NM in the TMA (see Fig. 1).

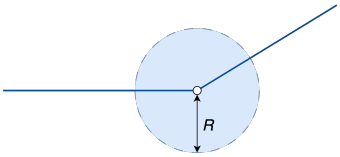

Figure 1. Illustration of the disc area.

For the leading flight f, the node time, represented by t_f^n, is calculated by adding the link entry time t_f^l to the link travel time, determined by dividing the link length (D^l) by the assigned speed (v_f^l) (see Equation 5).

$$t_f^n = t_f^l + \frac{D^l}{v_f^l} \tag{5}$$

where n denotes the node at the end of the link. In this sense, $t_f^n \leq t_g^n$ denotes the sequence of two consecutive aircraft f and g, where flight f is followed by flight g.

Subsequently, the entering and exiting times of the disc area associated with a given node (t_f^{in} and t_f^{out}) for a flight f can be estimated using the node time (t_f^n), adjusted by adding or subtracting the travel time within the disc area (see Equation 6).

$$t_f^{\text{in}} = t_f^n - \frac{R}{v_f^l} \tag{6a}$$

$$t_f^{\text{out}} = t_f^n + \frac{R}{v_f^l} \tag{6b}$$

Equation (7) denotes the total number of node conflicts.

$$C_{fg}^n(\mathbf{x}) = \begin{cases} 1 & \text{if } t_g^{\text{in}} \leq t_f^{\text{out}}, t_f^n \leq t_g^n \text{ and } n_f = n_g, \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

where $C_{fg}^n(\mathbf{x})$ equals 1 when a node conflict occurs between flights f and g. t_g^{in} and t_g^n represent the node entry time and the node time for the trailing flight g, respectively.

Thus, the minimum node separation is respected when $\sum_{\substack{(f,g)\in\mathcal{F}\times\mathcal{F} \\ f\neq g}} C_{fg}^n(\mathbf{x}) = 0$.

C. Objectives

As discussed in Section I, flights are often deviated from their initial flight plans to navigate around thunderstorms and avoid strong convective activities. During those deviations, conflict mitigation becomes a priority. A conflict-free trajectory can be ensured with Equation (8), based on the selected route and decision variables.

$$\sum_{\substack{(f,g)\in\mathcal{F}\times\mathcal{F} \\ f\neq g}} C_{fg}^l(\mathbf{x}) + \sum_{\substack{(f,g)\in\mathcal{F}\times\mathcal{F} \\ f\neq g}} C_{fg}^n(\mathbf{x}) = 0 \tag{8}$$

To control the optimisation process more precisely, conflicts are prioritised based on their associated costs, denoted as link and node evaluations, respectively. Link evaluation is established with Equation (9), while Equation (10) introduces a time-based evaluation for the disc area as node evaluation.

$$\varphi_l(\mathbf{x}) = \sum_{\substack{(f,g)\in\mathcal{F}\times\mathcal{F} \\ f\neq g}} \left[\frac{s_{fg} - \left| t_g^l \times v_g^l - t_f^l \times v_f^l \right|}{s_{fg}} + C_{fg}^l(\mathbf{x}) \right] \tag{9}$$

$$\varphi_n(\mathbf{x}) = \sum_{\substack{(f,g)\in\mathcal{F}\times\mathcal{F} \\ f\neq g}} \left[-\frac{t_g^{\text{in}} - t_f^{\text{out}}}{\max\left(\frac{t_f^{\text{out}} - t_f^{\text{in}}}{2}, \frac{t_g^{\text{out}} - t_g^{\text{in}}}{2} \right)} + C_{fg}^n(\mathbf{x}) \right] \tag{10}$$

where $\varphi_l(\mathbf{x})$ and $\varphi_n(\mathbf{x})$ denote aircraft performance on links and nodes, respectively.

Furthermore, conflicts with thunderstorms must be incorporated into the link evaluation. If an aircraft crosses a link impacted by a thunderstorm, a significant penalty is applied. The link must remain clear of adverse weather from the time the aircraft enters until it exits; otherwise, a large penalty is incurred. For blocked links, the evaluation process can be expressed with $\varphi_l(\mathbf{x}) + p$, where p represents the penalty for conflicts with thunderstorms.

Beyond conflict mitigation, this study also aims to minimise deviations in time slots and speeds, while reducing flight distances to minimise flight time and fuel consumption during detours by selecting the shortest route. Accordingly, three sub-objectives are introduced:

$$\varphi_t(\mathbf{x}) = \sum_{f\in\mathcal{F}} \left| t_f - t_f^o \right| \tag{11a}$$

$$\varphi_v(\mathbf{x}) = \sum_{f\in\mathcal{F}} \left| v_f - v_f^o \right| \tag{11b}$$

$$\varphi_d(\mathbf{x}) = \sum_{f\in\mathcal{F}} D^l \quad \forall\, l \in \mathcal{L} \tag{11c}$$

where $\varphi_t(\mathbf{x})$ denotes the changes in absolute value of the arrival time slots, $\varphi_v(\mathbf{x})$ represents the changes in absolute value of speeds, and $\varphi_d(\mathbf{x})$ indicates the total distance flown by the aircraft.

The objective function to be minimised is given by:

$$\min \left\{ \alpha \times \varphi_t(\mathbf{x}) + \beta \times \varphi_v(\mathbf{x}) + \gamma \times \varphi_d(\mathbf{x}) \right\} \tag{12}$$

where α, β and γ are weighting coefficients for each sub-objective, ensuring that all sub-objectives are on the same order of magnitude.

To ensure a conflict-free solution, the above-mentioned conflicts (Equation 8) are incorporated into the objective function as penalties through link and node evolutions (see Equations 9 and 10). Consequently, the original model is solved as a multi-objective minimisation problem (see Equation 13):

$$\min \Big\{ \alpha \times \varphi_t(\mathbf{x}) + \beta \times \varphi_v(\mathbf{x}) + \gamma \times \varphi_d(\mathbf{x}) + \lambda \times \left[\varphi_l(\mathbf{x}) + \varphi_n(\mathbf{x}) \right] \Big\} \tag{13}$$

where λ is a big weight assigned to prioritise the minimisation of conflicts during optimisation.

III. Optimisation algorithm design and implementation

Due to the computational complexity involved, the tractability of large-scale real-world flight trajectory optimisation problems remains challenging. As solution spaces rapidly expand and interactions between operating resources become extensive, traditional exact or near-exact algorithms often struggle to achieve optimal solutions within a reasonable timeframe. Therefore, scholars recommended adopting heuristic-based approaches to obtain near-optimal solutions with manageable computational resources [16].

Given the substantial memory requirements typically associated with real-world simulation environments, this paper builds upon the standard Simulated Annealing (SA) method and introduces an advanced problem-specific Selective Simulated Annealing (SSA) algorithm, which demonstrates enhanced computational efficiency while upholding the integrity of the resultant solution [17].

A. The mechanism of the standard SA

SA is a renowned meta-heuristic method widely applied to the air transport industry, including air traffic management [18], flight scheduling [19] and airport capacity management [20].

The standard SA algorithm begins with an initial feasible solution and iteratively explores the neighbourhood to find better solutions until no further improvements can be made [21]. The algorithm simulates the annealing process of solid substances, starting from an initial temperature with an initial solution $\overrightarrow{X_0}$. During each iteration, a transition is performed by generating a neighbour solution $\overrightarrow{X_j}$ from the current solution $\overrightarrow{X_i}$. If $\overrightarrow{X_j}$ improves the objective, it is accepted as the current best solution. Otherwise, the Metropolis criterion is established to decide whether to accept $\overrightarrow{X_j}$ based on a probability characterised by the Boltzmann statistical distribution, which provides the probability for each state i of energy E_i at the temperature T [22]. The Metropolis acceptance criterion helps overcome the limitations of the Monte-Carlo approach, making SA particularly suitable for complex simulations and large-dimensional state spaces. [23].

At the beginning of the process, the high temperature allows the algorithm to accept transitions with high objective degradation. This temporary setback helps the algorithm explore the state space more thoroughly and escape the local optimum. As the temperature decreases during the cooling phase, the probability of accepting transitions with high objective degradation decreases, allowing only transitions with relatively low deterioration. The algorithm continues this process until it reaches a predefined temperature threshold ε.

B. The implementation of the SSA and model transformation

To improve the algorithm performance, SSA introduces a *come back* operator to modify a vector $\overrightarrow{X_i}$ without duplication. Precisely, if the new solution is not accepted, the modified component reverts to its original value. The only information that needs to be stored is the index i of the changed component and its corresponding value $\overrightarrow{X_i}$. Therefore, this method avoids creating full copies of the state space during solution generation.

Moreover, the SSA enhances the performances of the standard SA by initially focusing on aircraft that are most involved in conflicts, using a performance-based neighbourhood operator. Particularly, the algorithm evaluates a vector of decisions along with their associated costs (see Equation 13). It then identifies the decision with the highest cost and computes the ratio between the cost of each decision and the highest cost. As a result, solutions with higher costs are more likely to be selected by the neighbouring operator.

The SSA algorithm is initialised with a historical flight dataset. Each flight is assigned a decision comprising three sub-decisions: route selection, arrival slot shifting, and speed adjustment. Consequently, a decision vector $\overrightarrow{X_f} = [\overrightarrow{r_f}, \overrightarrow{t_f}, \overrightarrow{v_f}]^T$ is created to denote all the sub-decisions for the selected flight f. To minimise the overall cost, the cost of each aircraft is proportional to its likelihood of being selected. When a new decision is evaluated through the neighbourhood operator, the previous cost associated with that decision is removed and updated. This process is repeated until the algorithm reaches the final temperature step (ε) as described earlier in the standard SA process.

While the optimisation algorithm aims to minimise the overall cost, penalty terms are introduced to discourage violations of both the objective and constraints. In this sense, Equation (13) provides a trade-off between constraint satisfaction and optimisation performance, allowing the algorithm to explore a broader solution space.

Therefore, the simulation-based evaluation can be described as follows: the SSA algorithm controls the vector of decision variables ($\overrightarrow{X_i}$) used by the simulation process to compute the performance (y) of those decisions (see Fig. 2).

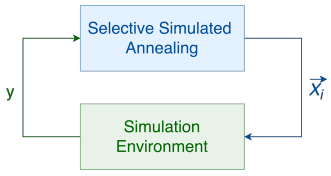

Figure 2. Objective function evaluation based on a simulation process.

IV. Case study

A. Background and dataset

To address the uncertainty associated with thunderstorm development and the high incidence of incidents and accidents during the approach and landing phases, this paper chooses the approach area of Chengdu Shuangliu International Airport (IATA code: CTU) as a case study to validate the proposed framework.

Fig. 3 presents a two-dimensional overview of the aircraft trajectories in this area, based on actual operational practices. In this representation, straight segments are denoted

as links (l), while intersections or fixes are represented as nodes (n). Dark green arrows indicate the Standard Terminal Arrival Routes (STARs), which are essential for ensuring safe operations, particularly under adverse weather conditions or low visibility. Light green arrows represent alternative routes with green dots marking intersections and grey dots indicating runways. In total, the terminal area under study comprises 5 entry nodes, 43 links, and 78 routes that sequentially connect flight segments.

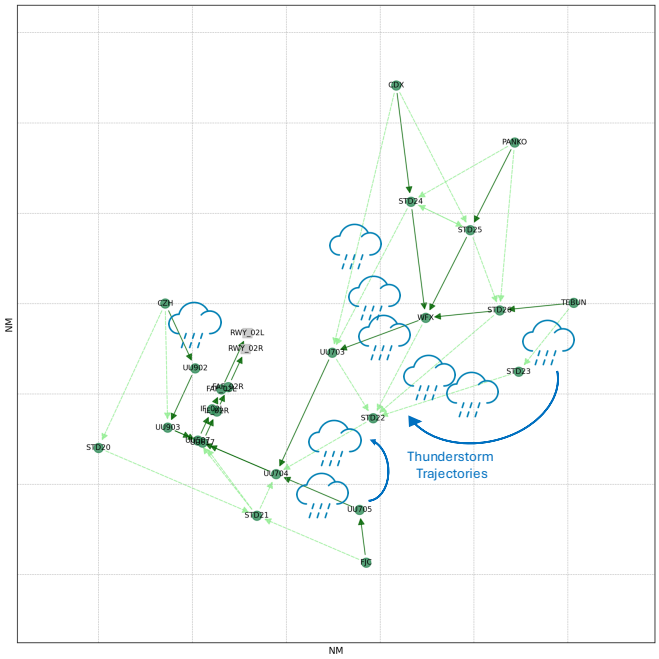

Figure 3. The STARs, alternative routes and thunderstorm trajectories in CTU approach area.

This graph-based representation is designed to be dynamic, enabling real-time updates to reflect airspace changes, such as closed sectors, reroutes, or temporary path restrictions. This model can be applied in situations where storm cells emerge and dissipate rapidly. For instance, when a storm develops along certain trajectories, the affected paths can be marked as temporarily unavailable, while the model re-optimises flight paths accordingly within a short time window.

Table II illustrates these variations in segment availability based on the Notice to Air Missions (NOTAM).

TABLE II. THUNDERSTORM DURATION AND LINK AVAILABILITY

Thunderstorm duration (by seconds)	Blocked links
[8000, 10000]	CZH-UU902, FJC-UU705, CZH-UU903, CZH-STD20, FJC-STD21, TEBUN-STD23, STD23-STD22
[10000, 11800]	UU705-UU704, STD26-STD22, WFX-STD22, STD26-WFX
[11800, 12500]	WFX-UU703, STD24-WFX, STD22-UU704, STD25-WFX
[13000, 14000]	CDX-UU703, STD24-UU703

Instead of replicating a real-time weather event, this study aims to evaluate the algorithm's performance in terms of speed and accuracy under high-load conditions. Therefore, a historical flight dataset consisting of 451 arrival flights over an 8-hour period from a single operational day is selected to stress-test the proposed algorithm in terms of computational efficiency and scalability. To simulate a more challenging operational scenario, the dataset is duplicated, yielding 902 flights within the same 8-hour window to ensure a higher traffic density for exploratory analysis.

Additionally, the model can be applied in shorter timeframes by running iteratively in rolling windows (e.g., every 10-15 minutes) to capture the dynamic nature of weather events and arrival management in real-world operations.

B. Parameter settings

The proposed algorithm is implemented in Java, and all experiments are conducted on a MacBook Air with an 8-core Apple M2 Chip and 16 GB of memory.

Typically, the time slot Δt is set to be 5 seconds. Given that delays are more frequent than early arrivals, the initial settings are $t_f^{\min} = -10$ minutes and $t_f^{\max} = 15$ minutes, respectively. However, preliminary tests revealed that a 15-minute delay was insufficient to resolve all conflicts, leading to an adjustment of the upper limit to 30 minutes for the time shift range.

Meanwhile, the following parameters are used for speed adjustments: $\Delta v = 0.01 \times v_f^o$, $v_f^{\min} = 0.9 \times v_f^o$, and $v_f^{\max} = 1.1 \times v_f^o$.

To assess node overlapping times between aircraft, this study applies a common radar separation standard of 3 NM.

The penalty for using a blocked link (p) is set to 500, while the weight for conflict mitigation (λ) is set to 50. The parameters α, β and γ are each set to be 1.

The algorithm initiates the decision with historical data, before generating a random decimal fraction between 0 and 1 to guide the neighbourhood operator. Specifically, if the decimal is less than 0.5, a route decision is made; if the decimal is between 0.3 and 0.7, a slot decision is made; and if the decimal is greater than 0.6, a speed decision is made (see Fig. 4). As a result, each flight has a 50% chance of a route change, a 40% chance of a slot change, and a 40% chance of a speed change.

Due to the overlapping ranges, the final probabilities for each scenario are adjusted accordingly. For each flight, there is a 30% chance of a route-only decision, a 10% chance of a slot-only decision, a 30% chance of a speed-only decision, a 20% chance of a route and slot decision overlap, and a 10% chance of a slot and speed decision overlap. While the total probability sums up to 100%, all possible outcomes have been considered.

In adapting the SSA algorithm to the proposed aircraft trajectory optimisation model, the parameters were fine-tuned using historical flight data. The initial temperature T_0 is strategically determined by a heat-up phase, where the temperature is progressively increased until 80% of the solutions are

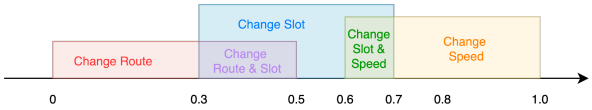

Figure 4. Decision combinations.

accepted. For each temperature level, the algorithm generates 2,000 neighbours to thoroughly explore the state space and capture the diverse possibilities inherent in aircraft trajectory optimisation. During the cooling process, the temperature decreases exponentially by a factor of 0.995 at each iteration ($T_{n+1} = T_n \times 0.995$), until it reaches a threshold ($\varepsilon = 0.0001 \times T_0$). This threshold symbolises the convergence of the algorithm, ensuring precision in the aircraft trajectory optimisation process.

C. Model performances

This study highlights the efficiency of the SSA algorithm in obtaining near-optimal solutions for the aircraft trajectory optimisation problem. Particularly, the proposed algorithm successfully computes trajectories for 902 flights within 70 seconds, delivering timely results for stakeholders to make informed decisions.

The convergence curves shown in Fig. 5 illustrate the evolution of the cost functions, providing a comprehensive view of the optimisation process. Precisely, the cost functions are represented by the multi-objective minimisation function and all the sub-objectives outlined in Section II, with their associated values on the y-axis, and the temperature on the x-axis. Among them, *Conflicts* represents the total number of conflicts detected on the links and nodes, expressed as $\sum_{(f,g)\in\mathcal{F}\times\mathcal{F}, f\neq g} C^l_{fg}(\mathbf{x}) + \sum_{(f,g)\in\mathcal{F}\times\mathcal{F}, f\neq g} C^n_{fg}(\mathbf{x})$. Meanwhile, *EvalLinks* and *EvalNodes*, *EvalDelay*, *EvalSpeed* and *EvalRoute* represent the evaluation results for $\varphi_l(\mathbf{x})$, $\varphi_n(\mathbf{x})$, $\varphi_t(\mathbf{x})$, $\varphi_v(\mathbf{x})$, and $\varphi_a(\mathbf{x})$, respectively. Notably, the significant gap between the objective function and other sub-objectives is primarily due to the weight ($\lambda = 50$) applied to prioritise conflict mitigation.

A consistent decrease in the cost functions is obvious over time across all curves, with a noticeable plateau observed in the early stages of computation. This rapid initial decline emphasises the efficiency of the search process, while the smooth and steady decrease indicates a robust and effective optimisation procedure. Notably, the predominant decrease in the cost function is primarily attributed to a significant reduction in the total number of conflicts among blocked links affected by thunderstorm cells, as the algorithm prioritises conflict mitigation before proceeding to aircraft performance evaluation. Further, the substantial decrease in sub-objectives related to aircraft trajectory evaluation further confirms the algorithm's capability to deliver high-quality solutions (see the zoomed-in part of Fig. 5).

Within the slot shifting range of [-10, 30] minutes, the algorithm optimised flight trajectories by primarily identifying the shortest routes, shifting slots and adjusting speeds. Fig. 6

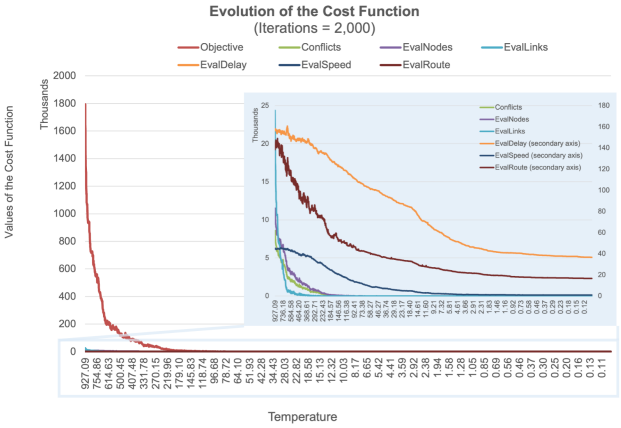

Figure 5. Evolution of the cost function.

visualises the distributions of decisions made by the proposed algorithm. As shown in Fig. 6a, 36.14% of flights experienced slot shifts within the range of [-1, 1] minutes, and 86.47% were shifted within [-5, 5] minutes, indicating minimal changes to slot allocations. These results also indicate that holding tracks in the TMA is not required to manage such a weather event. Meanwhile, Fig. 6b reveals that only 4.66% of flights required speed adjustments. This relatively low percentage suggests that switching routes and shifting slots are generally sufficient for effective trajectory optimisation under thunderstorm conditions.

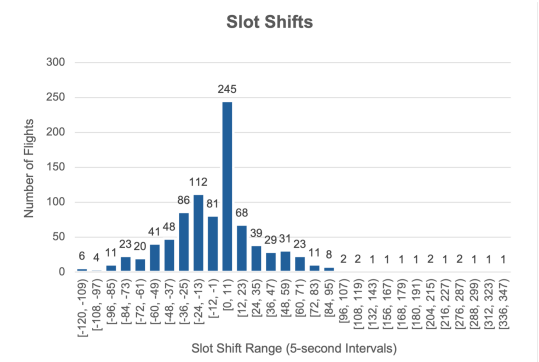

(a) Slot shift distribution.

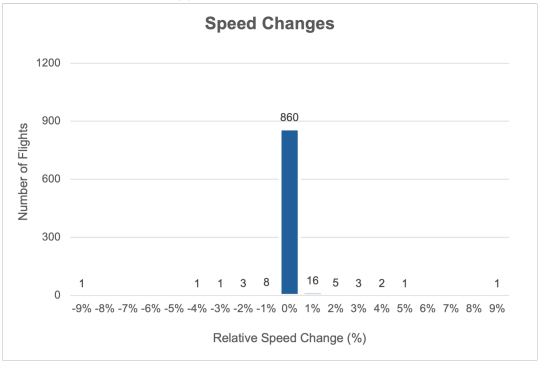

(b) Relative speed change distribution.

Figure 6. Decision distributions.

D. Discussions

1) Sensitivity analysis for slot shift ranges: One key decision in trajectory optimisation involves determining the appropriate slot shift range for the arrival flights. This study conducts a sensitivity analysis to assess the robustness of the selected slot range and evaluate how the objective and sub-objective functions respond to those variations.

The time windows are chosen to explore the theoretical upper bound for slot shifting, aiming to investigate the model's flexibility and limits in terms of conflict resolution. Fig. 7 compares the cost functions under different slot shift range combinations, categorised by their respective lower and upper limits. Generally, as the slot shift range expands, delays increase, leading to a higher total cost function, which the objective aims to minimise. More specifically, Fig. 7a divides the shift range into 5 groups based on their lower limits. An upward trend in the objective function is observed as the upper limits expand, which can be explained by the diminishing effectiveness caused by an over-extended acceptable slot shift range. For instance, a shift range of [-50, 60] increases the likelihood of delays, making it less efficient than a smaller range of [-10, 30]. Similarly, when the upper limits are fixed, a comparable trend is observed (see Fig. 7b). These findings confirm that the selected slot shift range effectively balances conflict mitigation with delay minimisation, demonstrating the algorithm's robustness across different slot shifting scenarios.

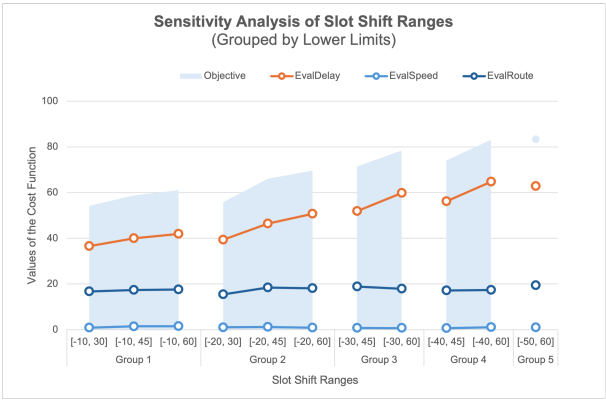

(a) Grouped by lower limits.

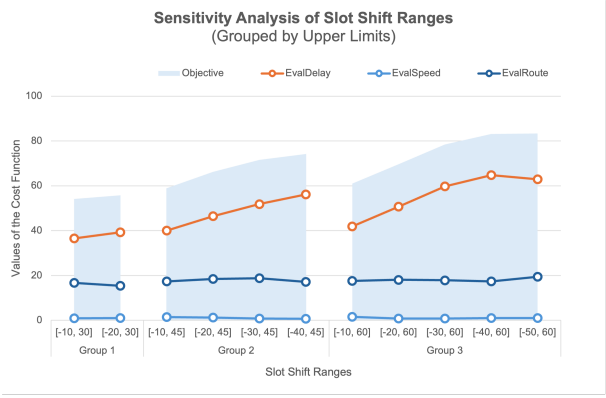

(b) Grouped by upper limits.

Figure 7. Comparative analysis of slot shift ranges.

2) Sensitivity analysis for decision combinations: To verify the sensitivity of the algorithm to the proposed decision combinations, the parameters used in Section IV.B 1) are considered as the baseline. Subsequently, route, slot and speed decisions are deactivated individually to observe the algorithm's response.

As demonstrated in Table III, combining any two of the sub-decisions is insufficient for resolving the aircraft sequencing and merging problem under thunderstorms, as none fully mitigate conflicts without fine-tuning additional parameters. Furthermore, the increase in delays indicates that incorporating aircraft holdings could be beneficial for finding solutions under those conditions. Compared to route and slot decisions, speed decisions have the least impact on the results. However, this also suggests that applying a minimal level of speed adjustments can enhance the overall performance and effectively regulate flights.

V. Conclusion

The presence of thunderstorms can severely disrupt standard flight operations, particularly during critical phases such as approach and landing, where precision and stability are paramount. To address those challenges, this study presents a generalised approach to optimising flight trajectories under thunderstorm conditions, delivering timely results for stakeholders to act upon.

The proposed framework prioritises conflict mitigation during flight deviations and assesses node and link conflicts to account for the dynamic nature of thunderstorms and other adverse weather conditions.

Flight trajectories are optimised primarily through route selection, minor slot shifts and speed adjustment, ensuring operational flexibility and timely landings. Moreover, the algorithm can be customised to prioritise specific decisions. For instance, if minimising holding is a priority, the algorithm will primarily adjust aircraft routes to reduce arrival delays, which would otherwise be managed through holdings when aircraft are at high altitudes.

Unlike previous studies that addressed stochastic optimal control problems for conflict resolution offline, this study presents an enhanced SSA algorithm for online conflict mitigation. Incorporating the SSA algorithm into the framework enhances its ability to adjust flight paths and speeds within predefined slot shift ranges. This underscores SSA's robustness and adaptability in real-world scenarios, making it a powerful tool for air traffic management with significant improvements in trajectory planning and conflict resolution.

Sensitivity analysis further illustrates the impact of slot shift ranges on the total cost function, validating that the selected shift range effectively minimises delays while maintaining overall system efficiency. By deactivating decisions individually, this study also discusses the sensitivity of the algorithm to the proposed decision combinations. The results show that applying a minimal level of speed adjustments can enhance the overall performance and effectively regulate flights.

TABLE III. PERFORMANCE METRICS FOR DIFFERENT DECISION DEACTIVATIONS

	Objective	Conflicts	EvalNodes	EvalLinks	EvalDelay	EvalSpeed	EvalRoute
Baseline	54.13	0.00	0.00	0.00	36.59	0.87	16.66
Deactivating Route Decision	173384.61	321.00	418.22	3045.64	121.52	20.97	49.10
Deactivating Slot Decision	547331.39	1516.00	3351.37	7591.77	0.00	64.15	109.99
Deactivating Speed Decision	463.25	4.00	7.82	0.00	49.16	0.00	23.17

While the uncertainties in aircraft entry times and storm forecasts present significant challenges when optimising over long time horizons, the decision to optimise over an 8-hour period was made to capture broader operational trends and minimise computational burden, ensuring a global view of network-wide disruptions. In this sense, a sequential sliding window approach for addressing real-time uncertainties more effectively will be explored in future work to improve the model's adaptability to dynamic conditions.

Future research could also explore additional key variables, such as altitude considerations and vertical conflict detection, using more detailed datasets. The algorithm could also be extended to measure the benefit of runways assignment. Additionally, the accuracy of weather forecasts may influence the effectiveness of the proposed framework, a factor to be integrated into future analyses for a more comprehensive evaluation.

ACKNOWLEDGMENT

The authors would like to thank the three reviewers for providing insightful comments to elevate this paper to the next level.

REFERENCES

[1] D. Hentzen, M. Kamgarpour, M. Soler, and D. González-Arribas, "On maximizing safety in stochastic aircraft trajectory planning with uncertain thunderstorm development," *Aerosp. Sci. and Technol.*, vol. 79, pp. 543–553, 2018.

[2] Federal Aviation Administration. (2024, 03) Weather delay. [Online]. Available: https://www.faa.gov/nextgen/programs/weather/faq

[3] D. González-Arribas, M. Soler, M. Sanjurjo-Rivo, M. Kamgarpour, and J. Simarro, "Robust aircraft trajectory planning under uncertain convective environments with optimal control and rapidly developing thunderstorms," *Aerosp. Sci. and Technol.*, vol. 89, pp. 445–459, 2019.

[4] EUROCONTROL, "Performance review report an assessment of air traffic management in europe," Performance Review Commission, EUROCONTROL, 96, rue de la Fusée, B-1130 Brussels, Belgium, Tech. Rep., 2024.

[5] A. Cook and G. Tanner, "European airline delay cost reference values," University of Westminster, 309 Regent St., London W1B 2HW, UK, Tech. Rep., 2015.

[6] R. Dalmau and X. Prats, "Fuel and time savings by flying continuous cruise climbs: Estimating the benefit pools for maximum range operations," *Transp. Res. Part D: Transp. and Environ.*, vol. 35, pp. 62–71, 2015.

[7] S. Kamo, J. Rosenow, H. Fricke, and M. Soler, "Fundamental framework to plan 4d robust descent trajectories for uncertainties in weather prediction," *Aerosp.*, vol. 9, no. 2, 2022.

[8] Y. Miyazawa, N. K. Wickramasinghe, A. Harada, and Y. Miyamoto, *Dynamic Programming Application to Airliner Four Dimensional Optimal Flight Trajectory*, ser. Guidance, Navigation, and Control and Co-located Conferences. American Institute of Aeronautics and Astronautics, 2013.

[9] Y. Matsuno, T. Tsuchiya, J. Wei, I. Hwang, and N. Matayoshi, "Stochastic optimal control for aircraft conflict resolution under wind uncertainty," *Aerosp. Sci. and Technol.*, vol. 43, pp. 77–88, 2015.

[10] J. García-Heras, M. Soler, D. González-Arribas, K. Eschbacher, C.-H. Rokitansky, D. Sacher, U. Gelhardt, J. Lang, T. Hauf, J. Simarro, A. Valenzuela, A. Franco, and D. Rivas, "Robust flight planning impact assessment considering convective phenomena," *Transp. Res. Part C: Emerg. Technol.*, vol. 123, p. 102968, 2021.

[11] E. Andrés, D. González-Arribas, M. Soler, M. Kamgarpour, M. Sanjurjo-Rivo, and J. Simarro, "Iterative graph deformation for aircraft trajectory planning considering ensemble forecasting of thunderstorms," *Transp. Res. Part C: Emerg. Technol.*, vol. 145, p. 103919, 2022.

[12] R. Sáez, B. G. Cabrera, M. Melgosa, A. Riccio, R. Verbeek, and N. van den Dungen, "Alternative 4d trajectories for the avoidance of weather- and contrail-sensitive volumes." Budapest: 12th SESAR Innovation Days, 2022.

[13] D. González-Arribas, F. Baneshi, E. Andrés, M. Soler, A. Jardines, and J. García-Heras, "Fast 4d flight planning under uncertainty through parallel stochastic path simulation," *Transp. Res. Part C: Emerg. Technol.*, vol. 148, p. 104018, 2023.

[14] F. Baneshi, M. Soler, and A. Simorgh, "Conflict assessment and resolution of climate-optimal aircraft trajectories at network scale," *Transp. Res. Part D: Transp. and Environ.*, vol. 115, p. 103592, 2023.

[15] A. Simorgh, M. Soler, S. Dietmüller, S. Matthes, H. Yamashita, F. Castino, and F. Yin, "Robust 4d climate-optimal aircraft trajectory planning under weather-induced uncertainties: Free-routing airspace," *Transp. Res. Part D: Transp. and Environ.*, vol. 131, p. 104196, 2024.

[16] E. Hernańdez-Romero, A. Valenzuela, and D. Rivas, "Strategic and probabilistic aircraft conflict detection and resolution for three-dimensional trajectories." 10th SESAR Innovation Days, 2020.

[17] J. Lavandier, A. Islami, D. Delahaye, S. Chaimatanan, and A. Abecassis, "Selective Simulated Annealing for large scale airspace congestion mitigation," *Aerosp.*, vol. 8, no. 10, 2021.

[18] Y. Huo, D. Delahaye, and M. Sbihi, "A probabilistic model based optimization for aircraft scheduling in terminal area under uncertainty," *Transp. Res. Part C: Emerg. Technol.*, vol. 132, p. 103374, 2021.

[19] H. Yang, C. Buire, D. Delahaye, and M. Le, "A heuristic-based multi-objective flight schedule generation framework for airline connectivity optimisation in bank structure: An empirical study on Air China in Chengdu," *J. of Air Transp. Manag.*, vol. 116, p. 102571, 2024.

[20] P. Scala, M. M. Mota, C.-L. Wu, and D. Delahaye, "An optimization-simulation closed-loop feedback framework for modeling the airport capacity management problem under uncertainty," *Transp. Res. Part C: Emerg. Technol.*, vol. 124, p. 102937, 2021.

[21] N. A. Ribeiro, A. Jacquillat, A. P. Antunes, and A. Odoni, "Improving slot allocation at Level 3 airports," *Transp. Res. Part A: Policy and Practice*, vol. 127, pp. 32–54, 2019.

[22] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equation of state calculations by fast computing machines," *J. Chem. Phys.*, vol. 21, no. 1087, 1953.

[23] D. Delahaye, S. Chaimatanan, and M. Mongeau, *Simulated Annealing: From Basics to Applications*. 11 W 42nd St Fl 15 New York, NY 10036, USA: Springer, 2019, pp. 1–35.