

Prescriptive Analytics System for Long-Range Aircraft Conflict Detection and Resolution

Samet Ayhan*
University of Maryland
College Park, Maryland
sayhan@cs.umd.edu

Pablo Costas
Boeing Research & Technology
Europe
Madrid, Spain
pablo.costas@boeing.com

Hanan Samet
University of Maryland
College Park, Maryland
hjs@cs.umd.edu

ABSTRACT

At the present time, there is no mechanism for Air Navigation Service Providers (ANSPs) to probe new flight plans filed by the Airlines Operation Centers (AOCs) against the existing approved flight plans to see if they are likely to cause conflicts or bring sector traffic densities beyond control. In the current Air Traffic Control (ATC) operations, aircraft conflicts and sector traffic densities are resolved tactically, increasing workload and leading to potential safety risks and loss of capacity and efficiency.

We propose a novel Prescriptive Analytics System to address a long-range aircraft conflict detection and resolution (CDR) problem. Given a set of predicted trajectories, the system declares a conflict when a protected zone of an aircraft on its trajectory is infringed upon by another aircraft. The system resolves the conflict by prescribing an alternative solution that is optimized by perturbing at least one of the trajectories involved in the conflict. To achieve this, the system learns from descriptive patterns of historical trajectories and pertinent weather observations and builds a Hidden Markov Model (HMM). Using a variant of the Viterbi algorithm, the system avoids the airspace volume in which the conflict is detected and generates a new optimal trajectory that is conflict-free. The key concept upon which the system is built is the assumption that airspace is nothing more than horizontally and vertically concatenated set of spatio-temporal data cubes where each cube is considered as an atomic unit. We evaluate our system using real trajectory datasets with pertinent weather observations from two continents and demonstrate its effectiveness for strategic CDR.

CCS CONCEPTS

• **Information systems** → **Data analytics**; • **Computing methodologies** → **Dynamic programming for Markov decision processes**; • **Applied computing** → **Aerospace**;

KEYWORDS

Prescriptive Analytics; Hidden Markov Model; Time Series

*Senior Engineer at Boeing Research & Technology.

1 INTRODUCTION

In the present Air Traffic Management (ATM) system, Airline Operations Center (AOC) personnel file a flight plan for a particular flight. The filed flight plan usually contains 2D coordinates of the fixed way points and the planned initial cruise speed and cruise altitude in addition to their speed and level changes along the route. It does not include the time information for the way points and the existing information is not routinely updated by the ATM systems. Moreover, the filed flight plan is not checked against other flight plans to probe potential interference with other aircraft or identify sector traffic complexity before the aircraft departs. Airspace sector complexity and aircraft separation assurance are dealt with tactically just a few minutes before the aircraft enters the sector or is likely to have a conflict with other aircraft.

It is estimated that by 2040 the USA alone can expect an increase of more than 68% in the commercial air traffic [10]. Hence, a new concept of operations is needed to accommodate this increase in volume. To meet this challenge ahead of us, new technologies and procedures for next generation ATM are being developed in the context of the Next Generation Air Transportation System (NextGen) [8] in the USA and the Single European Sky ATM Research (SESAR) [32] in Europe. The SESAR and NextGen concept of operations require a paradigm shift from a highly structured and fragmented system that is heavily reliant on tactical decision making and with few strategic planning functions based on uncertain information, to an integrated one based on collaborative strategic management of trajectories and information sharing [35]. In the future ATM systems to be built under NextGen and SESAR, the trajectory becomes the fundamental element of new sets of operating procedures collectively referred to as Trajectory-Based Operations (TBO). The underlying idea behind TBO is the concept of business trajectory which is the trajectory that will best meet airline business interests. The TBO concept of operations and the notion of business trajectory will result in more efficient 4D trajectories. Thus, the future ATM system should offer flexibility to accommodate business trajectories, while bearing the primary goal of safety in mind. Overall, the next generation ATM paradigm shift requires more safe and efficient CDR due to fact that maintaining the separation minima among the vast volume of aircraft that fly their versions of most optimal business 4D trajectories becomes more challenging.

Hence, we introduce a Prescriptive Analytics System to address the long-range aircraft CDR problem. The system performs CDR in two steps: In the first step, given a set of predicted trajectories in the form of a 4D joint spatio-temporal data cubes on a 3D grid network, the system declares a conflict if one of the aircraft's predicted

trajectory segment overlaps with the other's protected zone at the same time interval in the future. Obviously, the more accurate the predicted trajectories, the more accurate the detection of the conflict as the separation minima will be on the containment volume of each predicted trajectory. In the second step, the system builds a stochastic model, HMM that learns from the historical trajectories and their correlation with the pertinent weather parameters. Using a variant of the Viterbi algorithm, the CDR System prescribes an optimal solution by perturbing at least one of the 4D trajectories involved in the conflict.

In summary, the contributions of this paper are as follows:

- We propose a cube-shaped protected zone surrounding each aircraft that can be expanded by joining the neighboring cubes horizontally and vertically, yielding a protected zone with a desired size. This idea resonates with the assumption that an airspace is nothing more than a set of concatenated spatio-temporal data cubes around a 3D grid network, where each cube is considered as an atomic unit.
- We propose a scalable Prescriptive Analytics System to strategically address the aircraft CDR problem. Given a set of predicted trajectories, the system declares a conflict when a protected zone of an aircraft on its trajectory is infringed upon by another aircraft at the same time interval in the future. Upon a conflict, the system executes our conflict resolution algorithm and prescribes a solution that resolves the conflict by perturbing at least one of the 4D trajectories involved in the conflict.
- We conduct extensive experiments based on real trajectory and weather data from two continents and demonstrate that our CDR System can detect and resolve conflicts with lateral and vertical accuracies that are within the boundaries of conventionally accepted minimum separation values, set by the ANSPs. This translates to the fact that, ANSPs can now detect and resolve potential conflicts before the aircraft depart, resulting in safer and greener airspace with more efficiency and capacity, and thereby reducing the air traffic controller workload.

Our Prescriptive Analytics System can be used as a ground-based strategic CDR system by air traffic flow managers to resolve potential interference among large volume of aircraft and identify high density and complex sector traffic before the aircraft depart. The set of optimized resolutions should improve ATM automation and reduce the workload of air traffic controllers. The rest of the paper is organized as follows. Section 2 reviews related work. Section 3 introduces preliminary concepts followed by Section 4 where our Prescriptive Analytics System is described. Section 5 presents the results of our experiments, while Section 6 discusses the results. Concluding remarks are drawn in Section 7.

2 RELATED WORK

Trajectory data plays a large role in the problem of aircraft CDR. Trajectories have been the subject of much work in the spatial domain with a focus on cars along roads [30]. The focus has been on their generation (e.g., [31]), queries (e.g., [18, 20, 26, 28, 29]), and matching (e.g., [12, 17, 27]). This data is collected continuously and is quite voluminous. Instead, we address the aircraft flight

domain and on conflict (i.e., collision) detection and resolution. An excellent survey of various CDR methods is presented in [13]. In this survey, Kuchar et al. propose a taxonomy to categorize the basic functions of CDR modeling methods. The proposed taxonomy includes: dimensions of state information (lateral, vertical, or three-dimensional); method of state propagation (nominal, worst-case, or probabilistic); conflict detection threshold; conflict resolution method (prescribed, optimized, force field, or manual); maneuvering options (speed change, lateral, vertical, or combined maneuvers); and management of multiple aircraft conflicts (pairwise or global).

Conflict detection methods can be classified as nominal, worst-case, and probabilistic techniques. The nominal technique projects the current states into the future along a single trajectory without taking uncertainties into account [9, 37]. The worst-case technique assumes that an aircraft will perform any of a set of maneuvers and a conflict is predicted if any of the maneuvers could cause a conflict [34, 38]. The disadvantage of the worst-case technique is that it can declare a conflict as soon as there is a minimum likelihood of a conflict within the definition of the worst-case trajectory model thereby leading to false positives. The probabilistic approach offers a balance between relying on either a single trajectory model as in the nominal technique or a set of worst-case maneuvers. Instead it models uncertainties to describe potential changes in the future trajectory [14, 21, 40].

Once a conflict has been detected, the next step in the CDR process is to initiate the conflict resolution phase by determining the course of action. The conflict resolution method and the maneuvering options are two major factors in defining the course of action. Conflict resolution methods can be categorized as *a)* prescribed, *b)* optimized, *c)* force-field, and *d)* manual. The prescribed resolution method provides a fixed maneuver based on a set of predefined procedures [5]. Hence, depending on the nature of the conflict, the predefined resolution maneuver is automatically performed, minimizing the response time. However, it does not compute the optimal resolution path for the aircraft, resulting in a less efficient trajectory. The optimized method provides a conflict resolution strategy with the lowest cost based on a certain cost function (separation, fuel, time, workload, etc.) [9]. In the force-field resolution method, each aircraft is treated as a charged particle and the resolution maneuvers are defined using repulsive forces between the aircraft [41]. Although it is practical when properly applied, it may require a high level of guidance on the flight deck especially when the aircraft vary their speed over a wide range. The manual resolution method allows users to generate potential resolution options and provide feedback if the option is viable [40]. During the resolution phase, some CDR approaches only offer a single maneuver [9], while others offer a combination of maneuvers [15]. Obviously, the more maneuvering options the CDR approach offers, the more likely an efficient solution can be provided to a conflict. Our CDR approach has some similarities with [6] due to fact that both approaches consider the airspace as a set of cubic cells, called the grid model, declare a conflict if one of the aircraft's predicted trajectory segments overlaps with the other's protected zone, and propose an optimized 4D trajectory for the conflict resolution. However, unlike our study, they attempt to address tactical CDR (short-term and medium-term) by using Particle Swarm Optimization [6].




Figure 1: A sample crossing conflict and a PZ. (left) regular representation, (right) our unique representation.

Summarizing the above literature and comparing to the proposed prescriptive analytics approach, we make the following remarks. First, our method considers dimensions of state information as 4D. Second, our system is scalable, i.e. it can address CDR problem among multiple aircraft. Given a set of aircraft trajectories in 4D, our CDR System declares a conflict if any of the aircraft’s predicted trajectory segment overlaps with the other’s protected zone at the same time interval in the future. In our approach, one or more cubic cells forming the airspace is considered to be a conflict detection threshold. Next, we compute and prescribe an optimized solution which is a conflict-free 4D trajectory. Our approach addresses the CDR problem strategically over a time horizon of several hours to compute conflict-free 4D trajectories for optimal flight plans and less complex sector traffic densities.

3 PRELIMINARIES

The primary concern of the ANSPs, for example; the FAA in the USA and EUROCONTROL in Europe is to assure safety, which is quantified by the number of resolved conflicts.

Definition 3.1. A **conflict** is an event in which two or more aircraft come closer than a certain distance to one another.

Definition 3.2. **Separation minima** are encoded by lateral and vertical separation, forming a bounding volume around each aircraft, a *protected zone(PZ)*. Currently, the minimum lateral separation for en-route airspace is 5nm. It is 3nm inside the terminal radar approach control (TRACON) area. The minimum vertical separation is 2000 ft above the altitude of 29000 ft (FL290) and 1000 ft below FL290. Due to fact that lateral and vertical separation are specified by single distance values, the resulting PZ becomes a cylinder. Each aircraft is assumed to be surrounded by a PZ that moves along the aircraft.

Definition 3.3. **Conflict detection** is a process that evaluates the separation between any pair of aircraft, by comparing the distance between them with the separation minima. Formally, given a pair of predicted aircraft trajectories formed by a set of aircraft positions $T_i = [p_{i1}, p_{i2}, \dots, p_{im}]$, $T_j = [p_{j1}, p_{j2}, \dots, p_{jn}]$ where each point p is defined by its 4D spatio-temporal parameters (*latitude, longitude, altitude, and timestamp*), distance values between them $d_{i,j}$ are computed and compared with the separation minima d_s and a conflict is declared if any of distance values is less than the separation minima $d_{i,j} < d_s$.

Definition 3.4. **Conflict resolution** is a process that generates a feasible safe alternative trajectory by fulfilling the separation




Figure 2: Simplified depiction of a pairwise CDR.

minima criteria. Formally, given a pair of predicted aircraft trajectories formed by a set of aircraft positions $T_i = [p_{i1}, p_{i2}, \dots, p_{im}]$, $T_j = [p_{j1}, p_{j2}, \dots, p_{jn}]$, upon conflict resolution, all the distance values $d_{i,j}$ between the pairs of aircraft positions are greater than the separation minima $d_{i,j} > d_s$.

Definition 3.5. **Long-range CDR** is a process in which conflict detection and resolution is carried out several hours before the potential conflict occurs. Hence, long-range CDR is strategically performed before the departure for better planning, whereas mid-range and short-range CDR are tactically performed while the aircraft is airborne.

Unlike online CDR approaches in which distance between predicted future aircraft positions are constantly computed and compared with separation minima upon receipt of each new aircraft position, our approach is offline and uses a 3D grid network as a reference system. In our approach, raw trajectories are transformed into *aligned trajectories* causing aircraft to move along grid points. This results in conflict queries being computed at grid points only. In addition, unlike most other CDR approaches, our system considers a cube shaped PZ surrounding each aircraft. This idea resonates with the fact that our approach creates virtual data cubes around grid points, forming an overall airspace. Each cube is defined by its centroid, the original grid point, and associated weather parameters that remain homogeneous within the cube during a period of time. With this vision, we define trajectories as a set of 4D joint cubes.

This uncommon representation of 4D trajectories enables us to view conflicts and PZ from a unique perspective. Hence, in our view, PZ is a cube that can be expanded by joining the neighboring cubes horizontally and vertically. The process yields a PZ with a desired size. In our study, we use a PZ of variable size. It can be made up of a single cube or expanded by a number of cubes in each direction on each axis reaching a larger volume. Figure 1 illustrates a PZ in two different forms. The PZ on the left is formed by a cylinder. The PZ on the right is formed by 27 cubes. Figure 2 illustrates a sample pairwise CDR. Aircraft #1 departs before aircraft #2. Both aircraft move one cube at a time. In Figure 2a, aircraft #1 and #2 are located at cube J7 and K6, respectively. This causes a conflict as aircraft #2 intrudes aircraft #1’s PZ outlined in red. In Figure 2b, the conflict is resolved by a lateral shift to cube L6 by aircraft #2. No conflict occurs from here on as aircraft #1 follows cubes in gray (K7, L7, M7, N7, O7, P7, R7, S7) and aircraft #2 follows cubes in blue (M7, N8, N9, N10, O11, P11, R11, S11, T11) until they land at their pertinent airports.

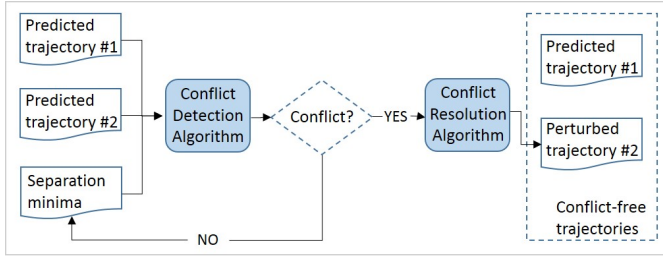


Figure 3: Overview of the CDR System.

4 PRESCRIPTIVE ANALYTICS SYSTEM FOR CDR

Figure 3 shows overview of the proposed Prescriptive Analytics System for a simplified pairwise CDR. Predicted trajectory #1 and #2 are generated by our Aircraft Trajectory Prediction System [2]. Our conflict detection algorithm takes predicted trajectories along with separation minima as input. The output of the process is a data cube defined by its 4D position, where conflict, if any, occurs. The next process in the pipeline is the conflict resolution, where a variant of the Viterbi algorithm is performed to avoid the conflicting trajectory segment. The process perturbs at least one of the trajectories involved in the conflict, generating a new optimized path and thereby resulting in conflict-free trajectories.

4.1 Pairwise Conflict Detection

The predicted trajectories along with the size of the PZ are fed into our pairwise conflict detection algorithm, presented in Algorithm 1. The algorithm declares a conflict if a PZ of an aircraft is infringed upon by another aircraft at the same time interval in the future.

Formally, given a pair of predicted trajectories $T_i = [p_{i1}, p_{i2}, \dots, p_{im}]$ and $T_j = [p_{j1}, p_{j2}, \dots, p_{jn}]$ where each trajectory is formed by a set of segments defined by their 4D spatio-temporal centroid parameters *latitude*, *longitude*, *altitude*, and *timestamp*, along with PZ for trajectory T_i in data cubes, we want to return the very first trajectory segment p_{js} , if a conflict occurs, *null* otherwise. Note that each aircraft position along predicted trajectories are recorded once every minute. To compute this, we start with the departure point of the aircraft that departs beforehand, and keep moving forward, one grid point at a time. With the departure of the second aircraft, we compare the trajectory segment p_{js} with PZ_{is} at each time instance t_s to see if p_{js} overlaps with PZ_{is} . Note that Algorithm 1 assumes that all the PZs have the same definition.

4.2 Pairwise Conflict Resolution

Our conflict resolution approach shares a common ground with our previous Trajectory Prediction System [2] as they both attempt to address an optimization problem; *given a set of weather observations, what is the most likely sequence of aircraft positions?* However, unlike the previous system [2], the current conflict resolution algorithm avoids the spatio-temporal data cubes where the conflict is detected. To recapitulate our approach to the optimization problem, we briefly review our previous Trajectory Prediction System [2] here.

Given a set of historical trajectories along with pertinent weather observations, the system works based upon an assumption that the

Algorithm 1: Pairwise Conflict Detection

Result: Detected conflict or no conflict

Input : Trajectory pairs T_i, T_j , Protected Zone PZ_i

Output : Conflicting trajectory segment p_{js} or *null*

```

1  $T_i \leftarrow [p_{i1}, p_{i2}, \dots, p_{im}]$ 
2  $T_j \leftarrow [p_{j1}, p_{j2}, \dots, p_{jn}]$ 
3 foreach  $t_s \in (p_{is} \cap p_{js})$  do
4   if  $p_{js} \subset PZ_{is}$  then
5     return  $p_{js}$ 
6   end
7 end
8 return null
    
```

weather observations are realizations of hidden aircraft positions i.e. trajectory segments and the transitions between the underlying hidden segments following a Hidden Markov model [22]. This assumption considers a finite set of states, each of which is associated with a probability distribution over all possible trajectory segments. Transitions among the states are managed by a set of probabilities. The states are not visible, but the pertinent observations are. Given a sequence of observations, the system trains an HMM, and derives hidden states, aircraft positions that correspond to the weather observations. The system computes the most likely sequence of aircraft positions in three steps:

- In the training data processing step, the system transforms raw trajectories into aligned trajectories and fuses weather parameters for each grid point along aligned trajectories. To achieve this, the system uses a 3D grid network with a spatial resolution of 6km x 6km as a reference system.
- In the test data processing step, the system resamples the weather parameters to generate buckets with distinct ranges and feeds them into the time series clustering algorithm [3] to produce input observations.
- In the final step, the HMM parameters generated in the first two steps and the flight time computed by our Estimated Time of Arrival (ETA) Prediction System [1] are used as input to the Viterbi algorithm. The output is the optimal state sequence, joint 4D cubes defining aircraft trajectories.

During the conflict resolution stage, our current CDR System makes use of the first two steps, outlined above. However, unlike the 3rd step of the process, our current CDR System avoids the cubes where the conflict is detected. Hence, the CDR System prescribes an optimized solution by perturbing at least one of the 4D trajectories, involved in the conflict. The process executes as follows: In addition to the regular HMM parameters of transition, emission, and initial probabilities, the system uses conflict-free probabilities where each state is assigned a probability value indicating how conflict-free it is. The parameters are fed into a variant of the Viterbi Algorithm, in which the system computes the optimal state sequence by considering the maximum HMM probabilities. Due to fact that the trajectory segments that are part of the first aircraft's trajectory are assigned low conflict-free probabilities, the system avoids selecting

them during the Viterbi process, yielding a conflict-free trajectory for the second aircraft.

Now, we present a variant of the Viterbi algorithm. Note that our previous Trajectory Prediction System [2] characterized an HMM by the following elements:

- N , the number of states in the model. States $S = \{S_1, S_2, \dots, S_N\}$ are represented by reference points' coordinates (*latitude, longitude, altitude*) that form aligned trajectories. We denote state at time t as q_t .
- M , the number of distinct observation symbols per state. Observations $V = \{v_1, v_2, \dots, v_M\}$ are represented by weather parameters (*temperature, wind speed, wind direction, humidity*) recorded at grid points.
- The state transition probability distribution $A = \{a_{ij}\}$ is the probability of an aircraft discretely transitioning from one state i to another j along its aligned trajectory, where

$$a_{ij} = P[q_{t+1} = S_j | q_t = S_i], \quad 1 \leq i, j \leq N$$

- The observation symbol probability distribution in state j , $B = b_j(k)$ is the probability of discrete weather parameters having been observed at that specific state, where

$$b_j(k) = P[v_k \text{ at } t | q_t = S_j], \quad \begin{array}{l} 1 \leq j \leq N \\ 1 \leq k \leq M \end{array}$$

- The initial state distribution $\pi = \{\pi_i\}$ is the probability of an aligned trajectory beginning at a state i , where

$$\pi_i = P[q_1 = S_i], \quad 1 \leq i \leq N$$

These parameters form an HMM, compactly denoted by $\lambda = \{A, B, \text{ and } \pi\}$. Now, we propose an additional parameter;

- The conflict-free probability distribution in state j , $C = c_j(k)$ is the probability of a conflict not occurring at that specific state, where

$$c_j(k) = P[v_k \text{ at } t | q_t = S_j], \quad \begin{array}{l} 1 \leq j \leq N \\ 1 \leq k \leq M \end{array}$$

Hence, the lower the conflict-free probability for a particular state, the lower likelihood of that state to be included in the most probable path. With this new parameter, an HMM can be expanded and denoted by $\lambda = \{A, B, \pi \text{ and } C\}$

The next step in the process is to choose a corresponding state sequence $Q = q_1, q_2, \dots, q_T$ that best explains the observation sequence $O = O_1, O_2, \dots, O_T$ given the model λ . A variant of the Viterbi algorithm [39] that is based on dynamic programming addresses this problem. The key component in the algorithm is the *optimal probability*, $\delta_t(j)$ is computed as follows:

$$\delta_t(j) = \max_{q_1, \dots, q_{t-1}} \pi_{q_1} b_{q_1}(o_1) c_{q_1}(o_1) \prod_{j=2}^t (a_{q_{j-1}, q_j} b_{q_j}(o_j) c_{q_j}(o_j))$$

Due to their low conflict-free probabilities, a variant of the Viterbi algorithm avoids trajectory segments where the conflict is detected, and generates a new optimized path.

Algorithm 2: Conflict Detection for Multiple Aircraft

Result: Detected conflict or approved optimal trajectory

Input : A new predicted trajectory T_i , constraints list CL_{PZ}

Output: Conflicting trajectory segment p_i or updated constraints list CL_{PZ}

```

1  $T_i \leftarrow [p_{i1}, p_{i2}, \dots, p_{in}]$ 
2  $CL_{PZ} \leftarrow [PZ_1, PZ_2, \dots, PZ_m]$ 
3 foreach  $t_s \in (p_i \& CL_{PZ})$  do
4   PruneAndSearch
5   if  $p_i \subset CL_{PZ}$  then
6     | return  $p_i$ 
7   end
8   else
9     | Insert  $p_i$  into  $CL_{PZ}$ 
10  end
11 end

```

4.3 CDR for Multiple Aircraft

Our pairwise CDR solution can be scaled to address CDR for multiple aircraft which can be used towards better planning of airspace sector densities. Similar to the current flight planning procedures, we propose predicted trajectories to be processed on a first come first served basis. For the sake of simplicity, consider an empty airspace. The airspace will be fully available to the first predicted trajectory. Hence, the first trajectory's optimal set of data cubes will be reserved in the airspace. This process will introduce a set of constraints in the form of PZ s by the first trajectory to the second and following trajectories during its flight time. Any conflict between the first and second trajectory will be detected and resolved using our pairwise CDR solution. Once resolved, a new set of constraints will be introduced by the second trajectory. The next trajectory in the queue will need to satisfy the constraints introduced by the previous trajectories and so on. This also means that the next trajectory will need to avoid the PZ s in the constraints list while generating its optimal set of data cubes during the conflict resolution phase. This process will be repeated until one or more flights land, which will result in the removal of all pertinent constraints from the list which will free up the particular sections of airspace.

Formally, given a new predicted trajectory $T_j = [p_{j1}, p_{j2}, \dots, p_{jl}]$ and a time series of existing constraints list $CL_{PZ} = [PZ_1, PZ_2, \dots, PZ_i, \dots, PZ_m]$, where $PZ_i = [pz_{i1}, pz_{i2}, \dots, pz_{ik}]$ along their existing trajectories $T = [T_1, T_2, \dots, T_i, \dots, T_n]$, where $T_i = [p_{i1}, p_{i2}, \dots, p_{ik}]$, we want to detect and resolve conflicts among these trajectories T_j and T (i.e. *one vs. all*). Note that PZ_i is formed by a set of data cubes, each defined by 4D spatio-temporal parameters around its centroid p_{ij} . With this approach, the implementation and management of constraints list is of central importance. In our implementation, we map time instances to data cubes forming PZ s, which are stored in an Octree data structure, where there is one Octree for each time instance. Hence, when a new trajectory comes in, the pertinent Octree is located based on the trajectory's time instance. We declare a conflict if the trajectory segment is found in the Octree. We process all pairs of possibly conflicting data cubes using spatial indexing

techniques to prune the search (e.g., [23, 25]). Our scalable conflict detection algorithm for multiple aircraft is presented in Algorithm 2. To keep the constraints list up-to-date, we periodically check and delete the pertinent Octree if any of the time instances has expired. Once a conflict has been detected, a variant of the Viterbi algorithm is performed where all the data cubes included in the constraints list for the pertinent time instance are avoided to find a conflict-free, optimized path for the new trajectory. To achieve this, we assign a minimum number to the conflict-free probabilities of those data cubes included in the constraints list.

Algorithm 3: Conflict Resolution for Multiple Aircraft

Result: A conflict-free optimized trajectory

Input : # of states N , # of distinct observations M ,
 transition probabilities A , emission probabilities B ,
 initial probabilities π , constraints list CL_{PZ}

Output: A conflict-free optimized trajectory T_o and
 updated constraints list CL_{PZ}

```

1  $S \leftarrow [p_1, p_2, \dots, p_i, \dots, p_n]$ 
2  $CL_{PZ} \leftarrow [PZ_1, PZ_2, \dots, PZ_m]$ 
3 foreach  $t_s \in ((p_i \in S) \& CL_{PZ})$  do
4   PruneAndSearch
5   if  $p_i \in CL_{PZ}$  then
6      $c_{p_i} \leftarrow 1 \times 10^{-100}$ 
7   end
8   else
9      $c_{p_i} \leftarrow 1$ 
10  end
11 end
12 VariantOfViterbi
13  $T_o \leftarrow \max_{s_1, \dots, s_{t-1}} \pi_{s_1} b_{s_1}(o_1) c_{s_1}(o_1) \prod_{j=2}^t (a_{s_{j-1}, s_j} b_{s_j}(o_j) c_{s_j}(o_j))$ 
14 foreach  $t_s \in ((p_o \in T_o) \& CL_{PZ})$  do
15   Insert  $p_o$  into  $CL_{PZ}$ 
16 end

```

Formally, given the number of states N , number of distinct observations per state M , state transition probability distribution A , observation probability distribution B , and initial state probability distribution π , along with a conflict-free probability distribution for all data cubes included in the constraints list at time instances $t = [t_1, t_2, \dots, t_n]$, we want to form an HMM λ , and train it to find the optimized path that is conflict-free. Our conflict resolution algorithm for multiple aircraft is presented in Algorithm 3.

5 EVALUATION

To evaluate our system, we generated a number of test cases using real trajectory and weather data from Europe and USA. Table 1 shows the European and USA airports forming the routes we used in our evaluation. Although the ideal evaluation would use real trajectories in actual conflicts, this is infeasible due to the nature of ATC operations, where the controller would interfere and separate the aircraft as soon as they are likely to infringe upon one another,

Table 1: A set of European and U.S.A. airports.

<i>AirportCode</i>	<i>AirportName</i>
LEAL	Alicante–Elche Airport
LEBL	Barcelona–El Prat Airport
LECO	A Coruña Airport
LEIB	Ibiza Airport
LEMD	Adolfo Suárez Madrid-Barajas Airport
LEMG	Málaga Airport
LEMH	Menorca Airport
LEPA	Palma de Mallorca Airport
LEVC	Valencia Airport
LEVX	Vigo–Peinador Airport
LEZL	Seville Airport
KATL	Hartsfield-Jackson Atlanta International Airport
KBOS	Boston Logan International Airport
KDFW	Dallas/Fort Worth International Airport
KJAX	Jacksonville International Airport
KLGA	NewYork LaGuardia Airport
KMIA	Miami International Airport
KORD	Chicago O’Hare International Airport
KPIT	Pittsburgh International Airport

so there would be no conflicts to find. Hence, we used a total of 16 European and USA pairs of flights that were in close proximity to cause potential conflicts when a minimal perturbation was applied.

5.1 Setup

The raw trajectory data from Europe was provided by Spanish ANSP, ENAIRE using a radar surveillance feed with a 5 seconds update rate. The raw data was wrangled as part of the Data-driven Aircraft Trajectory prediction research (DART) project under the SESAR Joint Undertaking Work Programme [33]. The European trajectory data contains all commercial domestic flights for Spain, a total of 119,563 raw trajectories and 80,784,192 raw trajectory points for the period of January through November 2016. The fields of the raw trajectory data are as follows: *flight no*, *departure airport*, *arrival airport*, *date*, *time*, *aircraft speed* in X , Y , Z directions, and position information (*latitude*, *longitude*, *altitude*). Note that, as a preprocessing step, we downsampled raw trajectory data from the original resolution of 5 seconds to 60 seconds and aligned them to our 3D reference grid [2]. The raw trajectory data from the USA was extracted from an Aircraft Situation Display to Industry (ASDI) data feed [11] which is recorded once in every 60 seconds and provided in near real-time by the FAA. The USA trajectory data contains flights between 8 major airports, a total of 4,628 raw trajectories and 450,919 raw trajectory points for the period of May 2010 through December 2015. The fields of the raw trajectory data are as follows: *source center*, *date*, *time*, *aircraft Id*, *speed*, *latitude*, *longitude* and *altitude*. Both European and USA weather data were extracted from the Global Forecast System (GFS), provided by the NOAA [16]. The original data has 28-km spatial and 6-hour temporal resolution and it contains over 40 weather parameters including *atmospheric*, *cloud* and *ground* attributes for each grid point as part of its 3D weather model. Hence, for this study’s geographic volume and time period of interest, over 160TB of weather data was collected.

Table 2: Sizes of training and test datasets.

TestCase#	Route#1	TrainingSetSize		TestSetSize		Route#2	TrainingSetSize		TestSetSize	
		#trjs	#pts	#trjs	#pts		#trjs	#pts	#trjs	#pts
1	LEAL-LEBL	1118	55116	200	9860	LEMD-LEIB	2572	125623	200	9769
2	LEAL-LEBL	1118	55116	19	937	LEMD-LEMH	1056	68141	19	1226
3	LEAL-LEBL	1118	55116	152	7493	LEPA-LEMD	5116	306128	152	9095
4	LEBL-LEMG	1704	127451	43	3216	LEMD-LEMH	1056	68141	43	2775
5	LEBL-LEMG	1704	127451	180	13463	LEPA-LEMD	5116	306128	180	10771
6	LEBL-LEZL	2404	183343	41	3127	LEMD-LEAM	1434	70128	41	2005
7	LEBL-LEZL	2404	183343	46	3508	LEMD-LEMH	1056	68141	46	2968
8	LEBL-LEZL	2404	183343	164	12508	LEMG-LEMD	1403	75408	164	8815
9	LEBL-LEZL	2404	183343	210	16016	LEPA-LEMD	5116	306128	210	12566
10	LEIB-LEBL	1360	53443	259	10178	LEPA-LEMD	5116	306128	259	15498
11	LEIB-LEBL	1360	53443	158	6209	LEPA-LEVC	1426	50467	158	5592
12	LEMG-LEBL	1563	114767	38	2790	LEMD-LEAM	1434	70128	38	1858
13	LEMG-LEBL	1563	114767	46	3378	LEMD-LEIB	2572	125623	46	2247
14	LEZL-LEBL	2380	186299	40	3131	LEMD-LEIB	2572	125623	40	1954
15	KDFW-KJAX	1355	153970	19	2159	KATL-KMLA	1296	108005	19	1583
16	KLGA-KORD	1155	131454	62	7056	KPIT-KBOS	822	57490	62	4339

Due to fact that our current system aims at addressing the CDR problem before departure, at least a pair of predicted trajectories are needed as input. Hence, we used our previous system [2] to generate a pair set of predicted trajectories. Next, we searched and found a number of trajectory points between the two flights’ trajectories in the first and second pair, where the *date*, *latitude*, *longitude*, and *altitude* values matched, and *time* mismatched. By perturbing one of the flights’ departure time we *virtually* created conflicts, where both aircraft traversed the same trajectory point at the same time. Table 2 shows the final size of training and test data in number of trajectories (#trj) and points (#pts). Note that trajectory data alone contains over 4 million trajectory points. In Table 2, the test data represents the conflicting trajectories. Aside from these 1,677 trajectory pairs, we also bootstrapped by drawing 100 additional trajectory pairs with replacement from the trajectory set to test for the false positive cases. Hence, we evaluated our CDR System’s effectiveness with a total of 3,554 (3,354 + 200) test trajectories on all 16 test cases.

5.2 Results

We evaluated our CDR System by comparing the output of conflict detection and conflict resolution with the ground truth on 16 test cases. Figure 4 illustrates the pairs with pertinent training and test data in white and yellow lines for these test cases. To evaluate our conflict detection capability we used trajectory prediction accuracy metrics as outlined in [19] and computed horizontal and vertical errors e_{horiz} , e_{vert} . To compute the errors, we fed a pair set of predicted trajectories generated by our previous system [2] into our conflict detection algorithm and compared the locations of *virtual* conflicts versus locations of conflicts detected by our CDR system. Next, we created 4 bin sizes, where each bin size is an integer multiple of 5nmi of lateral and 2000ft of vertical distances, conventionally accepted as minimum separation values for enroute airspace by ANSPs. Table 3 presents the pertinent condition for each bin. Using horizontal and vertical error values, we counted the

number of conflicts in each case and found which bin they belong to. The outcome is presented as a set of histograms in Figure 5. Our algorithm detected **87.2%** of the conflicts within the first bin size and **99%** of the conflicts within the first two bin sizes on all 16 test cases. Note that the conflict detection can only be as accurate as the predicted trajectories. Hence, these errors can be attributed to the accuracy of our Trajectory Prediction System [2], defined by the horizontal, and vertical error of **14.981nmi** and **1589.452ft** respectively along the entire test trajectories.

To resolve the conflicts, we ran our conflict resolution algorithm as highlighted in Section 4.2. Figure 6 is a closer look at one of the detected and resolved conflicts by our system. In all figures, the yellow cubes and white cubes represent the first and second flight’s, respectively, predicted trajectories. The red line parallel to the white cubes represents the first flight’s actual trajectory and the red line parallel to the yellow cubes represents the second flight’s actual trajectory. As both aircrafts move one cube at a time in the flight direction, the *PZ* illustrated in the cyan cube around the current position of the first flight also moves forward in the form of a sliding window. The first figure on the far left shows where each aircraft is at time t_{s-1} , represented respectively by the solid white cubes for the first and yellow cubes for the second flight. The second figure from the left captures the conflict detected at time interval t_s by our system illustrated with a solid red cube. The actual conflict occurs at where the red lines intersect. Note that both the predicted conflict position and the actual conflict position are within the first flight’s *PZ*. The third figure from the left is the 3D view of the second figure from the left. The actual and predicted conflict positions are only 2 cube sizes away from each other, considering the center of the cube as the predicted position of the aircraft. The figure in the far right illustrates the optimized solution to the conflict by our system. The first flight’s trajectory has been perturbed vertically and the conflict has been resolved by our system. The altitude of the second flight has been elevated resulting in conflict being resolved. Note that due to assignment

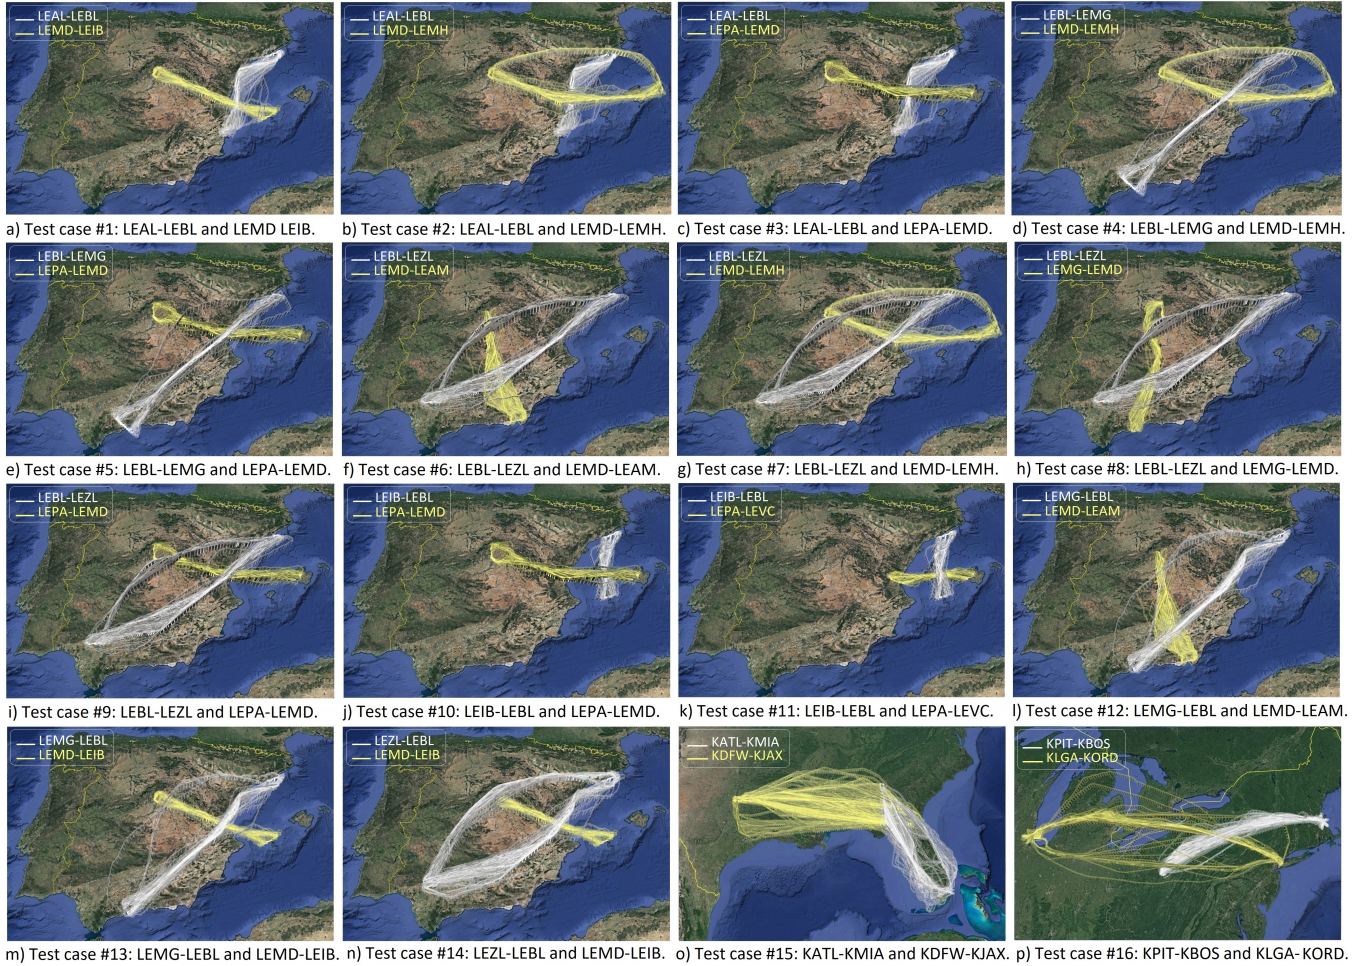


Figure 4: Visual representation of training and test data for conflicting pairs of flights.

of low conflict-free probabilities to the 3D grid points inside of the PZ , the system avoids selecting them during the conflict resolution stage, generating conflict free trajectories.

As the final evaluation step, we executed our conflict resolution algorithm on all conflicts formed by 1677 trajectory pairs and computed accuracy values based on the number of successful resolutions. Due to the fact that our conflict resolution algorithm resolved the vast majority of conflicts on all test cases, we provide aggregated results overall, rather than provide results for each test case. Table 4 presents accuracy values for each bin. Note that, to resolve the conflicts, we treated each bin differently based on their varying sizes so that only relevant 3D grid points were avoided. The process perturbed the trajectory for the second flight, generated an optimal alternative and yielded conflict-free trajectories.

6 DISCUSSION

Our conflict detection algorithm reached **6.012nmi** of mean horizontal error. Note that this value is considerably less than 14.981nmi, the mean horizontal error by our Trajectory Prediction System [2] along the entire trajectory points including climb, cruise and descent phases of a flight. This is due to two major facts: 1) Trajectory

Table 3: Horizontal and vertical error ranges for the bins.

Condition	Bin
$0\text{nmi} \leq e_{\text{horiz}} \leq 5\text{nmi} \wedge 0\text{ft} \leq e_{\text{vert}} \leq 2000\text{ft}$	5nmi,2000ft
$5\text{nmi} < e_{\text{horiz}} \leq 10\text{nmi} \vee 2000\text{ft} < e_{\text{vert}} \leq 4000\text{ft}$	10nmi,4000ft
$10\text{nmi} < e_{\text{horiz}} \leq 15\text{nmi} \vee 4000\text{ft} < e_{\text{vert}} \leq 6000\text{ft}$	15nmi,6000ft
$15\text{nmi} < e_{\text{horiz}} \leq 20\text{nmi} \vee 6000\text{ft} < e_{\text{vert}} \leq 8000\text{ft}$	20nmi,8000ft

prediction accuracy during the cruise phase is often considerably higher than the climb and descent phases of the flight. 2) All conflicts we used in our experiments took place during the cruise phase of the flights. With 6.012nmi of mean horizontal error on the conflict positions, our conflict detection algorithm found **99%** of conflicts within the first two bins of separation minima (10nmi, 4000ft). We also verified that between none of the additional 100 trajectory pairs where the *virtual* conflict has never occurred was falsely detected as a conflict by our system. Our conflict resolution algorithm’s mean accuracy was over **97%** on all test cases. Though, it is interesting to see that it was unable to reach 100% accuracy, given the fact that all it had to do was avoid the selected 3D grid points when generating the new conflict-free optimal trajectory.

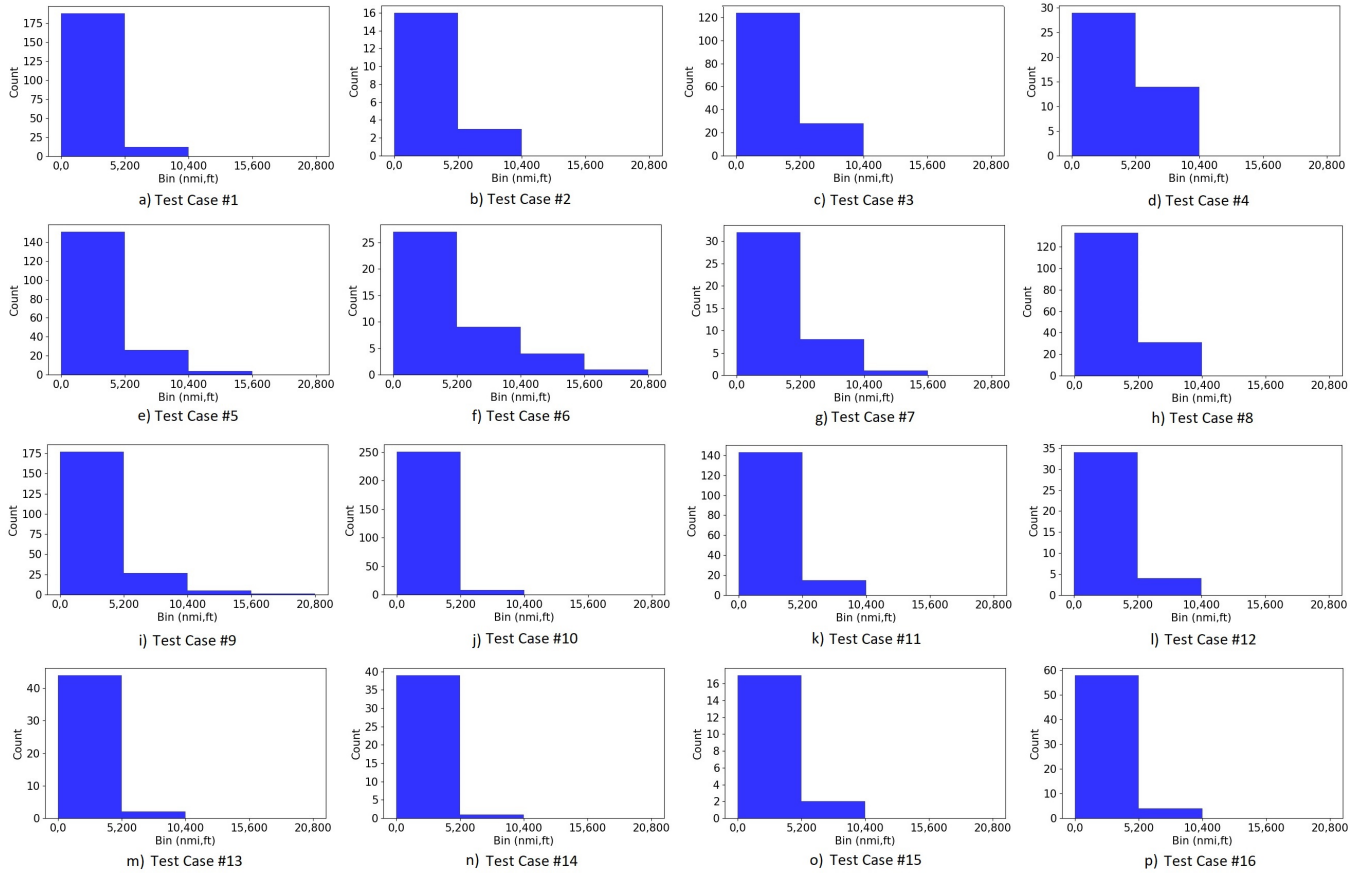


Figure 5: Conflict detection histograms showing how many conflicts were captured by our CDR System in each bin.

The reason for that was the sparse distribution of data cubes, i.e. lack of data over the 3D grid network. The algorithm was unable to connect the new trajectory segments from start to end due to disconnection. Hence, our CDR System craves for more data to reach higher accuracy values.

These results validate the effectiveness of our system on the long-range CDR problem. However, this is not to say that strategic CDR will detect and resolve conflicts once for all and no more conflicts will occur during the flight. There will likely be some convective weather patterns shaping after departure. These sudden changes causing potential conflicts should be addressed tactically by short and or medium-range CDR systems while the aircraft is airborne. Although we were not able to find, there may also likely be some exceptional cases where false positive conflicts may be found. These cases should also be addressed tactically by short and or medium-range CDR systems. Note that the larger the selected PZ value, the higher probability of finding and resolving the conflicts between trajectories. However, that also means less denser sectors resulting in inefficient use of airspace. Hence, the tradeoff should be handled carefully by the ANSPs.

Overall, we propose our system to be used by AOCs to file more realistic flight plans. Our system can also be used by ANSPs to validate that the filed flight plans do not cause conflicts with previously approved flight plans or increase any sector traffic complexities.

These goals can be achieved in the planning phase before aircraft depart, improving ATM automation and reducing the air traffic controller workload.

Table 4: Accuracy of our conflict resolution algorithm.

	Accuracy
Bin1 (5nmi, 2000ft)	98.4%
Bin2 (10nmi, 4000ft)	97.4%
Bin3 (15nmi, 6000ft)	93.7%
Bin4 (20nmi, 8000ft)	100.0%

7 CONCLUSION

We have presented a novel Prescriptive Analytics System addressing a long-range CDR problem. Our experiments on real trajectory and weather datasets from two continents verify that our system achieves lateral and vertical accuracies that are within the boundaries of conventionally accepted minimum separation values, set by the ANSPs. With our system, ANSPs can detect and resolve potential conflicts long before the aircraft depart, resulting in safer and greener skies with higher efficiency and capacity, and thereby reducing the air traffic controller workload.

Some future work could involve adding a spatial browsing capability [4, 7, 24] for the trajectories as well as incorporating our methods in a distributed spatial environment [36].

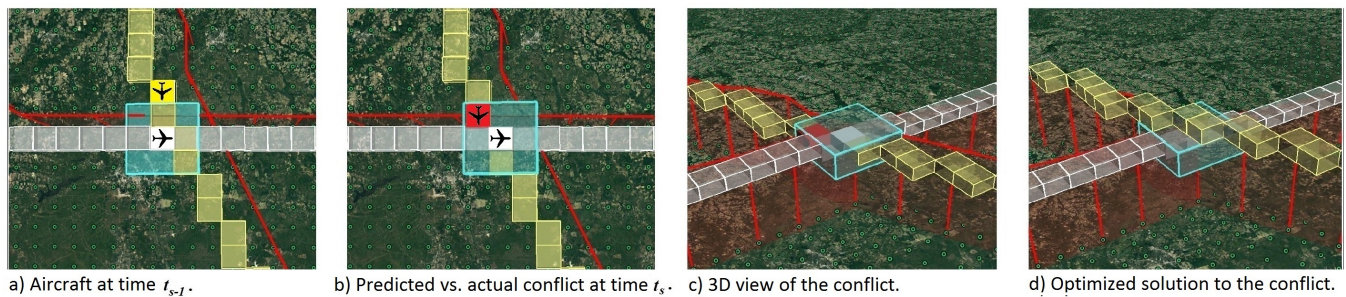


Figure 6: Illustration of a detected and resolved pairwise conflict by our system.

8 ACKNOWLEDGEMENTS

This work was supported in part by the NSF under Grants IIS-13-20791 and IIS-18-16889. We are grateful to the Boeing Research & Technology Europe, a member of the project DART [33] team for providing the European trajectory data used in this study. The project DART has received funding from the SESAR Joint Undertaking under Grant Agreement No. 699299 under the European Union’s Horizon 2020 Research and Innovation Programme.

REFERENCES

- [1] S. Ayhan, P. Costas, and H. Samet. 2018. Predicting Estimated Time of Arrival for Commercial Flights. In *Proceedings of the 24th ACM SIGKDD Int’l Conference on Knowledge Discovery and Data Mining*. London, UK, 33–42.
- [2] S. Ayhan and H. Samet. 2016. Aircraft Trajectory Prediction Made Easy with Predictive Analytics. In *Proceedings of the 22nd ACM SIGKDD Int’l Conference on Knowledge Discovery and Data Mining*. San Francisco, CA, 21–30.
- [3] S. Ayhan and H. Samet. 2016. Time Series Clustering of Weather Observations in Predicting Climb Phase of Aircraft Trajectories. In *Proceedings of the 9th ACM SIGSPATIAL IWCTS*. San Francisco, CA, 25–30.
- [4] F. Brabec and H. Samet. 2007. Client-based spatial browsing on the world wide web. *IEEE Internet Computing* 11, 1 (January/February 2007), 52–59.
- [5] B. Carpenter and J. Kuchar. 1997. Probability-based Collision Alerting Logic for Closely-spaced Parallel Approach. In *Proceedings of the 35th Aerospace Sciences Meeting and Exhibit*. Reno, NV, 1–8.
- [6] J. Cobano, D. Alejo, G. Heredia, and A. Ollero. 2013. 4D Trajectory Planning in ATM with an Anytime Stochastic Approach. In *Proceedings of the 3rd Int’l Conference on ATACCS*. Naples, Italy, 1–8.
- [7] C. Esperança and H. Samet. 2002. Experience with SAND/Tcl: a scripting tool for spatial databases. *Journal of Visual Languages and Computing* 13, 2 (April 2002), 229–255.
- [8] FAA. 2009. *The NextGen*. <https://www.faa.gov/nextgen/>
- [9] FAA. 2012. *Traffic Alert and Collision Avoidance System (TCAS II)*. https://www.faa.gov/other_visit/aviation_industry/airline_operators/airline_safety/info/all_infos/media/2012/InFO12010.pdf
- [10] FAA. 2015. *TAF*. <http://taf.faa.gov/Downloads/TAFSummaryFY2015-2040.pdf>
- [11] FAA. 2016. *Aircraft Situation Display to Industry*. <http://www.fly.faa.gov/ASDI/>
- [12] E. Jacox and H. Samet. 2008. Metric space similarity joins. *ACM Transactions on Database Systems* 33, 2 (June 2008), 7.
- [13] J. Kuchar and L. Yang. 2000. A Review of Conflict Detection and Resolution Modeling Methods. *IEEE Transactions on Intelligent Transportation Systems* 1, 4 (December 2000), 179–189.
- [14] W. Liu and I. Hwang. 2011. Probabilistic Trajectory Prediction and Conflict Detection for Air Traffic Control. *Journal of Guidance, Control, and Dynamics* 34 (November 2011), 1779–1789.
- [15] P. Menon, G. Sweriduk, and B. Sridhar. 1999. Optimal Strategies for Free-Flight Air Traffic Conflict Resolution. *Journal of Guidance, Control, and Dynamics* 22 (March 1999), 202–211.
- [16] NOAA. 2016. *NCEP Global Forecast System*. <https://www.ncdc.noaa.gov/data-access/model-data/model-datasets/global-forecast-system-gfs>
- [17] S. Nutanong, E. H. Jacox, and H. Samet. 2011. An incremental Hausdorff distance calculation algorithm. *PVLDB* 4, 8 (August 2011), 506–517.
- [18] S. Nutanong and H. Samet. 2013. Memory-efficient algorithms for spatial network queries. In *Proceedings of the 29th IEEE Int’l Conference on Data Engineering*. Brisbane, Australia, 649–660.
- [19] M. Paglione and R. Oaks. 2007. Implementation and Metrics for a Trajectory Prediction Validation Methodology. In *AIAA GNC Conference and Exhibit*. Hilton Head, SC, 1–18.
- [20] S. Peng, J. Sankaranarayanan, and H. Samet. 2016. SPDO: High-Throughput Road Distance Computations on Spark Using Distance Oracles. In *Proceedings of the 32nd IEEE Int’l Conference on Data Engineering*. Helsinki, Finland, 1239–1250.
- [21] M. Prandini, J. Hu, J. Lygeros, and S. Sastry. 2000. A Probabilistic Approach to Aircraft Conflict Detection. *IEEE Transactions on Intelligent Transportation Systems* 1, 4 (December 2000), 199–220.
- [22] L. Rabiner. 1990. *Readings in Speech Recognition*. Morgan Kaufmann, San Francisco, CA, Chapter A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition, 267–296.
- [23] H. Samet. 2006. *Foundations of Multidimensional and Metric Data Structures*. Morgan-Kaufmann, San Francisco, CA.
- [24] H. Samet, H. Alborzi, F. Brabec, C. Esperança, G. R. Hjaltason, F. Morgan, and E. Tanin. 2003. Use of the SAND spatial browser for digital government applications. *Commun. ACM* 46, 1 (January 2003), 63–66.
- [25] H. Samet and M. Tamminen. 1986. An improved approach to connected component labeling of images. In *Proceedings of Computer Vision and Pattern Recognition’86*. Miami Beach, FL, 312–318.
- [26] J. Sankaranarayanan, H. Alborzi, and H. Samet. 2005. Efficient query processing on spatial networks. In *Proceedings of the 13th ACM Int’l Symposium on Advances in Geographic Information Systems*. Bremen, Germany, 200–209.
- [27] J. Sankaranarayanan, H. Alborzi, and H. Samet. 2006. Distance Join Queries on Spatial Networks. In *Proceedings of the 14th ACM Int’l Symposium on Advances in Geographic Information Systems*. Arlington, VA, 211–218.
- [28] J. Sankaranarayanan and H. Samet. 2009. Distance oracles for spatial networks. In *Proceedings of the 25th IEEE Int’l Conference on Data Engineering*. Shanghai, China, 652–663.
- [29] J. Sankaranarayanan and H. Samet. 2010. Query processing using distance oracles for spatial networks. *IEEE Transactions on Knowledge and Data Engineering* 22, 8 (August 2010), 1158–1175.
- [30] J. Sankaranarayanan and H. Samet. 2010. Roads belong in databases. *IEEE Data Engineering Bulletin* 33, 2 (June 2010), 4–11.
- [31] J. Sankaranarayanan, H. Samet, and H. Alborzi. 2009. Path oracles for spatial networks. *PVLDB* 2, 1 (August 2009), 1210–1221.
- [32] SESAR. 1999. *Single European Sky ATM Research*. <http://www.eurocontrol.int/dossiers/single-european-sky>
- [33] SESAR. 2017. *DART*. <http://dart-research.eu/the-project/>
- [34] M. Shewchun, J. Oh, and E. Feron. 1997. Linear Matrix Inequalities for Free Flight Conflict Problems. In *Proceedings of the 36th IEEE Conference on Decision and Control*. San Diego, CA, 2417–2422.
- [35] M. Soler, A. Olivares, E. Staffetti, and J. Cegarra. 2011. Multi-phase Optimal Control Applied to 4D Business Trajectory Strategic Planning in Air Traffic Management. In *Proceedings of the 1st Int’l Conference on ATACCS*. Barcelona, Spain, 68–78.
- [36] E. Tanin, A. Harwood, and H. Samet. 2005. A distributed quadtree index for peer-to-peer settings. In *Proc. of the 21st IEEE Int’l Conf. on Data Engineering*. Tokyo, Japan, 254–255.
- [37] A. Valenzuela and D. Rivas. 2009. Conflict Detection and Resolution in Converging Air Traffic. In *9th AIAA ATIO Conference*. Hilton Head, SC, 1–13.
- [38] A. Vink, S. Kauppinen, J. Beers, and K. Jong. 1997. Medium-term Conflict Detection in EATCHIP Phase III. In *Proceedings of the 16th DASC. AIAA/IEEE Digital Avionics Systems Conference*. 45–52.
- [39] A. Viterbi. 1967. Error Bounds for Convolutional Codes and an Asymptotically Optimum Decoding Algorithm. *IEEE Transactions on Information Theory* 13, 2 (April 1967), 260–269.
- [40] L. Yang and J. Kuchar. 1998. Using Intent Information in Probabilistic Conflict Analysis. In *AIAA GNC Conference and Exhibit*. Boston, MA, 797–806.
- [41] K. Zeghal and E. Hoffman. 1999. Design of Cockpit Displays for Limited Delegation of Separation Assurance. In *Proceedings of the 18th Digital Avionics Systems Conference*. St. Louis, MO, 1–8.