

An Agent Based Model of the Air Traffic Management

C. Bongiorno¹,

S. Miccichè¹, R. N. Mantegna^{1,2}

¹Dipartimento di Fisica e Chimica,
Università di Palermo
Viale delle Scienze
I-90128 Palermo, Italy

²Center for Network Science and
Department of Economics,
Central European University,
Nador utca 9, 1051 Budapest, Hungary

G. Gurtner³, F. Lillo^{1,3,4}, L. Valori³

³Scuola Normale Superiore di Pisa
Piazza dei Cavalieri 7
I-56126 Pisa, Italy,

⁴Santa Fe Institute
1399 Hyde Park Rd.
Santa Fe NM 87501, USA.

M. Ducci⁵, B. Monechi⁵, S. Pozzi⁵

⁵Deep Blue s.r.l.
Piazza Buenos Aires 20
00100 Roma, Italy.

Foreword - *This paper describes a project that is part of SESAR Workpackage E, which is addressing long-term and innovative research.*

Abstract—**The path to a deep understanding of how the forthcoming introduction of the SESAR trajectory based scenario will impact on the future air traffic management procedures goes through a better understanding of the actual air traffic system. To this end the WP-E ELSA project aims at developing an empirically grounded agent based model that describes some of the stylized facts observed in the Air Traffic Management of the European airspace. The model itself has two main parts, a Strategic layer and a Tactical layer, which aims at emulating the two main steps relevant in the ATM: (i) The strategic layer, focused on the interaction between the Network Manager and the Air Operators and (ii) the tactical layer, focused on aircraft and controllers behaviour in Air Traffic Control (ATC) sectors.**

The preliminary results for the strategic layer show that when we have a mixing of re-routing and shifting companies, the overall satisfaction can even increase together with the number of flights, which is an effect not observed when one have in the system one type of companies only. The preliminary results for the tactical layer indicate that when shocks in the system are confined in small areas, the interplay between the re-routing and change of flight level strategies may even lead to trajectory modifications that give smaller average delays as long as the number of shocks increases.

I. INTRODUCTION

In the future of Air Traffic Management (ATM) it is expected to observe an increase of traffic demand and new business challenges that will bring the current ATM system to its capacity limits within the 2013-2015. As a consequence, an overall productivity improvement is urgently needed [1], [2], [3], [4]. The structure of ATM system, as it is known today, will therefore change in many aspects. One of the key enabler to the productivity and efficiency shift foreseen by SESAR will be the business-trajectory concept [1]. In the future SESAR scenario airspace users will not fly along structured routes. On the contrary, they will be able to fly a 4D trajectory selected on the basis of their own business and efficiency needs. Within this major change not only the ATM productivity should be

drastically enhanced, but consequently also the ATM system safety and resilience standards will have to be improved.

The path to a deep understanding of how these aspects will impact the future air traffic management procedures goes through a better understanding of the actual air traffic system and its management procedures. To this end the WP-E ELSA project aims at developing an empirically grounded agent based model that describes some of the stylized facts observed in the Air Traffic Management of the European airspace. The model itself has two main parts, a Strategic layer and a Tactical layer, which aims at emulating the two main steps relevant in the ATM: (i) The strategic layer, focused on the interaction between the Network Manager and the Air Operators and (ii) the tactical layer, focused on aircraft and controllers behaviour in Air Traffic Control (ATC) sectors. This model will then be used as a scenario simulator in order to understand which benefits the trajectory based scenario will bring to the ATM world. Here we will discuss the features of the model only within the current ATM scenario.

The paper is organized as follows: in section II we will review the main features of the strategical and tactical layers of the model. In section III we will summarize the main inputs we have gathered from operational experts about some of the ATM features to be implemented in the model. In section IV we will show the main preliminary results of the model and in section V we will finally draw our conclusions.

II. THE MODEL

As mentioned above, the model is organized in two main layers: a Strategic layer and a Tactical layer. In the current implementation of the model these two layers are still independent from each other. However, they are logically connected in the sense that the results of the strategic layer should feed the tactical layer. Such integration stage is currently ongoing.

A. The Strategic layer

The strategic layer of the ABM aims at modelling the events taking place from the planning of the flight plan by the Air Operators to the final acceptance by the network manager. Its real temporal scale goes from months down to a few hours before departure. On the other hand, its spatial extension concerns the whole ECAC space down to the sectors. The main aim of the strategic layer is to model how and why the airspace is filled such as it is. The “final state” is a result of different actors, but also of the structure of the airspace itself: sectors, national airspaces, etc. In the following, we focus on the role of the former, the airspace itself being kept fixed, either by using an artificially generated network of sectors or the real one.

There are two main types of agents in the system. On one hand, Air Operators try to get the best trajectories for their flights. This best trajectory has two components: the geometrical length (in 3 dimension, and taking into account dominant winds) and the times of departure/arrival. Different Air Operators, which may have different goals, are competing one with each other for the “best” slots and trajectories, based on business considerations and constrained by the structure of the airspace. On the other hand, the network manager tries to fill the airspace as best as possible, its main concern being to avoid to overload the airspace, in order to guarantee safety. In the current scenario, the network manager is very passive, takes only propositions from the Air Operators and tries to fill the airspace. In the future SESAR scenario, the network manager might have more proactive behavior, for instance submitting counter-propositions to the Air Operators. This is a feature we plan to add to the modeling, which currently only describes the current ATM scenario.

The main object of the model is the flight plan. It is defined as a pair $fp = (t, \mathbf{p})$, where t is the time of departure and \mathbf{p} is a vector containing the list of sectors followed by the aircraft. t is a real number, hence our model is not a continuous time model.

The aircraft are bound to travel on a network, whose nodes represent sectors and links exists if two sectors share a common boundary. Moreover, a weight is associated with the links: it represents the time of travel between two sectors. These different times are clearly related one to each other in reality – because of the physical extension of the sectors. However in this paper, we draw weights as independent normal random variables with mean τ and standard deviation σ , in order to be able to perform semi-analytic calculations. A second metric is associated to the nodes themselves: an integer representing the capacity of the sector. This capacity is the maximum number of aircraft that can be simultaneously present in a sector. In reality, it is related to many features of the sector, like its area, its volume, etc. Here we chose a constant capacity (equal to 5) for each sector.

We used mainly a simulated network of sectors for the simulations, but we tested our models also on real airspace networks, obtaining similar results (not shown here). The

airspace is built by using a Voronoi tessellation of a set of 90 points randomly thrown on the plane. After the tessellation, we connect each node with its geometrical neighbors – thus building the Delaunay triangulation. Then we choose the distribution of weights on the network (normal), the capacities (constant) and the airports (randomly chosen, here we pick only two airports). The structure of this network, even though it is planar, has some strong similarities with the real network of sectors, as we show in section IV-A.

In the simulations we performed, we have several agents of type AO (Air Operator). Each of them has a unique form for the cost function for its flights, which is

$$c(t, \mathbf{p}) = \alpha|\mathbf{p}| + \beta(t - t_0),$$

where t_0 is the desired time of departure, $|\mathbf{p}|$ the weighted length of the path on the network and α and β two parameters defining the main characteristics of the company. Please notice that flight are only shifted ahead in time, therefore $t \geq t_0$. A high value of β simulates a company eager to have its flights on time. On the other hand, a high α simulates a company more preoccupied by the length of the trajectory.

Each AO begins with the generation of k flight plans for each “flight”, defined by a pair of airports and a desired departure time t_0 . They are generated by finding the k (k is set to 10 in the simulations) best flights plans, ranked by increasing cost, among all the possible paths \mathbf{p} connecting the two airports and times of departure t . Among them, the best one is the one with the shortest path \mathbf{p}_{sp} on the network and the desired time ($t = t_0$), with associated cost $c_{best} = \beta|\mathbf{p}_{sp}|$. More specifically, the AO can shift the flight plan in time by a constant increment, that we fix always equal to the average time of travel τ between sectors, hence giving the natural time scale of the system.

Once all the flight plans have been generated, a company is randomly drawn and it submits the k flight plans of one flight to the second type of agent, the network manager (NM). Following this queue, the NM tries to fill the airspace. For each flight, it takes the best flight plan and tries to fill it on the network. If one sector or more overreaches its capacity, the flight plan is rejected. Then the next flight plans, with higher costs, are tried, until one is accepted or all are rejected.

A key parameter of the model is the desired departing times t_0 . We tested several patterns for the density of desired departing times by all the companies, ranging from a totally uniform distribution to very peaked distributions. The results in terms of occupation of the space and satisfaction of the air operators are very different, as we show in the results part (see IV-A).

Finally, we define a metric to measure the satisfaction of the system. First, we choose to define the satisfaction of a single flight by computing $s_f = c_{best}/c(fp_{accepted})$, where c_{best} is the cost of the best flight plan (hence the one with the smallest cost) and $c(fp_{accepted})$ is the cost of the accepted flight plan. We choose also to take $s_f = 0$ if all its flight plans have been rejected. We define also the satisfaction of the overall system

as the average satisfaction over all flights:

$$S = \frac{1}{N_f} \sum_f s_f,$$

where N_f is the total number of flights.

We implemented the model in Python, using the library `networkx`.

B. The Tactical layer

The agents of the tactical layer of the ELSA agent-based model are aircraft/pilots and controllers who are active at the level of ATC sectors. In this layer of the agent based model we model and simulate the events that make a planned flight plan, recorded in the so-called M1 files, transform into an actual one, recorded in the so-called M3 files. The aim is that of investigating the issues that affect the predictability of the last filled flight-plan within the ATM system. The specific scientific questions we are investigating are:

- What are the issues that affect the predictability of the last filled flight-plan within the ATM system? How is the predictability affected by these issues?
- Can sectors capacity be improved by a more efficient management of conflicts?

Within the ELSA projects these issues are seen in the perspective of the new SESAR scenario of a business-trajectory regulation scheme. Here we will address the problem of predictability and capacity improvements only within the current ATM scenario.

1) *General features of the model:* The interaction between the agents is needed in order to manage the tactical changes occurring in the system due to unforeseen events, i.e. weather events, congestions, limitation of sectors capacity, etc. Moreover, the ATC sectors are the places where flight trajectories are made conflict free. Consequently, we will try to model and simulate the events that make M1 transform into M3 by considering agents that operate at the level of ATC sectors.

In the current version of the model the aircraft/pilots have a limited intelligence. We are also giving these agents the opportunity of modifying the aircraft velocity in order to solve safety events, even though this strategy is actually used by controllers only in specific cases. However, we decided to explore this possibility because this will nevertheless be a possible strategy adopted by companies and controllers in the future SESAR scenario. Indeed, we have constructed the code in a modular way that allows to swap the priority of the strategies adopted by the controllers. In fact, we can easily modify the code in such a way that controllers first check for the possibility of doing re-routings and then change the flight altitude. Another major feature of the model is to take into account the possibility that the de-conflicting module changes the aircraft velocity in order to solve the conflicts.

The model take into account that M1 trajectories are not conflict free. Thus one main task to be performed within the model is to deconflict trajectories. Moreover, we simulate shocks in the system and see how the system reacts to it. Specifically we simulate a shock in an area around a navigation

point. We assume that the shock lasts for a certain time window. Operatively, this means that for a certain time window a certain area of the ATC sector can not be crossed by flights. This might correspond to a situation where an extreme weather event occurs as well as to a situation when a certain area is highly congested and therefore the air traffic must be deviated. As a result, another task of the model is to change one or more flight trajectories in order to avoid the shocked areas. The way we model this step is to deviate the flight trajectories along new navigation points that are external to the restricted area and with the constraint that (i) we want to minimize the length of the deviated trajectory and (ii) the deviated trajectory must be conflict free. We will perform different simulation experiments changing the statistical features of the shocks.

These tasks are accomplished as follows

- We take into account three different situations:

- 1) *there is the onset of a shocking area.* In this case we select a new navigation point for each flight trajectory such that the new trajectory has the minimal distance from the planned one. However, we allow for changing the aircraft velocity in the modified trajectory segments, if necessary. This enables us to select trajectories characterized by better fitness measures. This algorithm is therefore essentially based on re-routing, possibly augmented with the possibility of changing velocities.
- 2) *there is a possible conflict of trajectories that nevertheless do not intersect one with each other* - We are implementing a two-level algorithm. This algorithm mixes together the search of a new navigation point along the lines discussed below with an approach that, based on a genetic algorithm, searches the optimal velocity in the mutated trajectory segment, as to ensure that the new trajectory exists. The advantage of this two-level algorithm is that it is fast. Indeed, it must be taken into account that aircraft can change their velocity within a range that depends on the flight level, the type of aircraft and other variables. As a preliminary choice we will consider velocity changes of the order of 5 % of the previous en-route velocity.
- 3) *there is a possible conflict of trajectories that intersect with each other* - When solving head-on conflicts we use the algorithm illustrated above, which proves to be efficient also in this case. From a logical point of view there is no difference from the previous case. The algorithm treats these two situations in the same way. We keep this case distinct from the previous one to emphasize that the present case usually occurs in the planned trajectories, while the previous one usually occurs when one of the two conflicting trajectories has already been deviated.

- Based on the inputs from the italian ENAV operational experts, we have introduced the feature such that when a trajectory is deviated, then it is not sent back to

its planned trajectory. Rather, it is sent to the planned exit navigation point of the sector. This helps us in implementing the fact that airplanes are given directs within a sector.

2) *Implementation of the Model:* The code that implements the model is written in Python [7]. However, some modules have been written in C [8] in order to improve the computational efficiency of the ABM. Below we describe the modules that compose the tactical layer of the Agent based model.

In the current version, the tactical layer of the model works at the level of a single ATC sector. Having that, in the current version of the model we consider shocked areas which are totally included in the sector and that do not involve navigation points on the boundaries. A future release of the model will consider multiple ATC sectors and this will also give us the opportunity to treat the case when shocked areas are overlapping different sectors or do involve navigation points on the boundaries. As we mentioned above, the model simulate the events that make the planned flight-trajectories transform into the actual one by considering agents that operate at the level of ATC sectors. In the current implementation the planned trajectories are assumed to be existing. Specifically we will consider the flight trajectories are recorded in the DDR (Demand Data Repository) M1 files we have access to [5].

Navigation Points module - Given the sector, we populate it with navigation points. On one hand, part of the navigation points selected are real ones (e.g., 28), i.e., those crossed by the flights according to their M1 last-filled flight-plan. On the other hand, other navpoints are generated randomly from an uniform distribution (e.g., 1000) and we have considered only those falling inside the polygon of the given sector (i.e., we have excluded those on the boundaries). These new navigation points could be seen as temporary points (!-points) in the M3 flight plan. An illustration of all the navigation points we use is provided in figure 1. Those colored in blue are the real ones, in red are the randomly generated ones. We have not generated navigation points on the boundaries because these are sensible points and need to be treated separately. Indeed, since these points connect two or more sectors, it is required coordination between the reference controllers. Usually, controllers gives directs inside their area of reference changing the navigation points a flight has to pass through, while they rarely substitute those navpoints that are located on the boundaries and that can affect the air traffic management of the neighbor controller. In other words, different aircraft can cross sector borders in different places. However, a specific aircraft, even though it was re-routed, will cross the border in the planned navpoint.

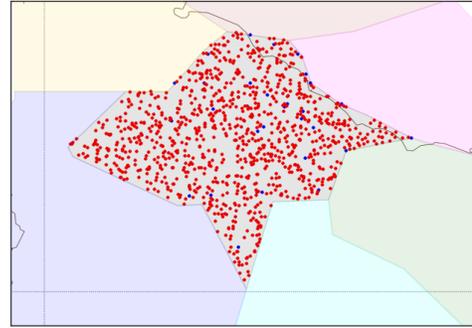


Fig. 1. Visualization of the navigation points inside the selected sector. The points in blue are the real one derived from the M1 last-filled plan, those in red are randomly generated.

Flight List module - Once the sector has been populated with navigation points, we create a list FL_k of flights active in that time-step in the considered sector. Such list will be reshuffled in the next time-step. Within this list we check whether the flights are crossing a shocked area and whether or not they conflict with other trajectories. Specifically, the i -th aircraft in the list will be checked against all other $j < i$ flights. When a trajectory modification is needed, it will affect the i -th flight. The reason for shuffling the list at each time step is in order to avoid that the trajectory modifications are always applied to the same aircraft.

Collision module - In order to check for collision between two flights, we use a data structure that considers the aircraft localized inside the trajectory segments travelled within the given time-step Δt . For this purpose, we introduce a finer subdivision of the time-step into N elementary time increments δt and compute the real space-time position of the aircraft at each elementary time increment, by assuming a constant velocity. The collision algorithm will have to simply calculate the positions of the aircraft for each of the elementary time increments and then compute the distances between the two aircraft at these positions. Suppose we are now checking if the i -th flight trajectory is conflicting with all other f_j trajectories, with $j < i$. We are therefore considering a number i of flight trajectories. For each of them we have an array \mathcal{P}_j , $j = 1, \dots, i$ of positions computed according to the algorithm illustrated above. For each of the elementary time-increments, we compute an array of distances d choosing as a value for each array element the minimum distance between the i -th aircraft and all the other aircraft in the list FL_i with $j < i$. From the distances array it is possible to estimate a *fitness value* to maximize. By assuming that the safety distance threshold is d_{thr} , and $X = \{d_i\}$ is the subset of distances below d_{thr} , the fitness value is defined as $\mathcal{F}_1 = \sum_{i \in X} (d_{thr} - d_i)$. If this value is different from zero then there is a conflict and the algorithm proceeds to the next module that performs the de-conflicting of trajectories. As a result the computational

time increases linearly with the number of aircraft. Moreover, this method allows us to introduce a simple way to manage possible changes of aircraft velocity. Of course the collision module is not in place when the necessity for re-routing is due to the existence of a shocked area.

De-conflicting module - After the check for collision has been done, this module searches for a new conflict-free trajectory. It is conceived as a three-step algorithm that acts on the velocities of the aircraft, the search of a new trajectory (re-routing) and the change of flight level, in case conflict exists. The order by which the three steps are applied might be changed. Here we describe them.

The first step of the module we present here is the one that performs the re-routing. The procedure is illustrated in Fig. 2. We first identify the two navigation points B and A which are before and after the collision (crossed circle in the figure), respectively. The idea is to (i) keep B , (ii) substitute A and (iii) eliminate all the other subsequent navigation points but the last one L . To do that, we take the previously generated temporary navigation points T_k (squares in the figure) and we order them with respect to the angle that the segment connecting B and T_k forms with the original trajectory. We select the temporary navigation point that have the smallest angle. We admit a maximum angle of 45° . Having this new navigation point we compute again all the distances with the $j < i$ trajectories and compute again \mathcal{F}_1 . If $\mathcal{F}_1 = 0$, then we select this navigation point, otherwise we go to the navigation point with the second smallest angle. This procedure is iterated until we find $\mathcal{F}_1 = 0$. If we find only navigation points with an angle larger than 45° , the algorithm exits this module and go to the next one.

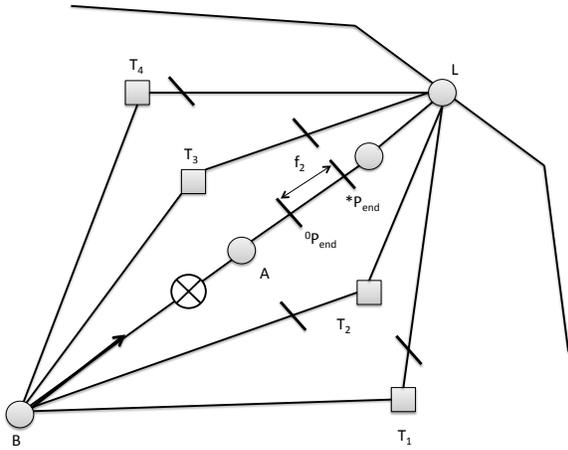


Fig. 2. The figure illustrates the techniques of rerouting after the rescheduling of velocity. Every possible rerouting is tested with the genetic algorithm for the velocity. This would define a point P_{end} . The route selected will be the one such that minimizes the fitness \mathcal{F}_2 , i.e. it minimizes the distance between the point P_{end}^* of the original route and the P_{end} of the new route.

The second step of this module involves changes of flight level. The model implements three possible flight levels. All

flights are initially considered to be active in the central flight level. Therefore they can move upwards or downwards whenever the re-routing is not feasible. The choice of the new level is done by considering the one where there is less probability of having conflicts. This is assessed by computing the sum of the \mathcal{F}_1 functions for all the flights in the two external flight levels and choose the flight level with the smallest one. If the two sums are equal then we choose the upper flight level. If no level is available then the algorithm exits this module and go to the next one.

The third step for changing the flight velocity is mainly based on a genetic algorithm [6] that minimizes two fitness functions by using a mix of cross-over and mutation operators. Suppose we are now checking if the i -th flight trajectory is conflicting with all other f_j trajectories, with $j < i$. We are therefore considering a number $i - 1$ of flight trajectories. For each of them we have an array \mathcal{P}_j , $j = 1, \dots, i - 1$ of positions computed as illustrated above. When the collision module identifies a subset of points within the j -th array \mathcal{P}_j such that these points are below the safety distance threshold, the ABM tries to solve the collision using a new module that implements a genetic evolutive algorithm. This module is written in C language in order to make it computationally efficient. In our data structure we use a characteristic speed for each elementary segment between two navigation points. Assuming that before going out from the sector the aircraft will run into N_{nvp} navigation points, the aircraft will therefore change $N_{nvp} - 1$ velocities during his trajectory. In this step we first generate a population N_{pop} of velocities arrays each of length $N_{nvp} - 1$. We generate velocity arrays \mathcal{V}_s , $s = 1, \dots, N_{pop}$ filled with v_s velocities randomly selected under the condition that $v_{min} < v_s < v_{max}$, where v_{min} and v_{max} are the minimal and maximal velocity acceptable for the considered aircraft flying in the considered sector. For each element of this population we compute the array of positions \mathcal{P}_s , $s = 1, \dots, N_{pop}$ with dimension N . For each \mathcal{P}_s , $s = 1, \dots, N_{pop}$ the algorithm computes the \mathcal{F}_1 fitness value, according to the procedure illustrated in the re-routing module, see Fig. 2. The population of possible solutions (velocities arrays) will be sorted according to this fitness value. Another target to pursue is to minimize the delay with respect to the time scheduled in the file $M1$. To this end, we sort the elements with degeneration in the fitness value associating an estimator of this delay. Assuming that P_{end} and P_{end}^* are the end points of the aircraft according to the considered population and the scheduled flight plan, respectively, the $\mathcal{F}_2 = d(P_{end}, P_{end}^*)$ will be the second fitness to be minimized. The genetic algorithm searches for the optimal solution by applying to the population of velocities three different operators. Each one has a characteristic probability to be applied. Moreover, in order to avoid the selection of elements of the population that are too similar with each other, these operators will not be applied to the best element, but according to a probability function (a Gaussian distribution defined on the positive real axis, peaked in zero and normalized to unity in the interval from 0 to N_{pop}). Each

new element generated with these operators will take the place of the worst element in the population. These operators are applied on a random basis with different probabilities as to ensure that local minima are avoided.

3) *Expected results:* For each simulated flight we will monitor two variables. First we will consider the number $A_{c,f}^{(0)}$ of actions that controller $c = 1, \dots, C$ has performed for each flight $f = 1, \dots, F$. This will simply be the number of changes (velocity, position, ...) operated on the planned trajectory. We will also consider positive actions $A_{c,f}^{(+0)}$ the ones when the controller gives a direct. We will consider negative actions $A_{c,f}^{(-0)}$ all the others. As a second variable we will consider the arrival time $T_{3,f}^{(0)}$ of each flight $f = 1, \dots, F$.

We believe that the model is realistic enough and efficient from a computational point of view. It will allow us to perform different types of simulation experiments aiming at understanding to which extent the sector capacity can be improved. In fact, having an automated algorithm for de-conflicting the trajectories might lead to an improvement of the capacity performances within sectors. These capacity stress-test will be done by increasing the number N of flights entering the sector in a certain time window and investigating up to which value of N the algorithm is able to find a de-conflicted solution. We believe that investigating capacity issue in connection with the optimal de-conflicting of trajectories is a relevant issue, especially in view of the new SESAR scenario.

Moreover, the modular structure by which the model has been implemented will allow us to easily switch from the current to the future SESAR scenario by eliminating some module, such as the Navigation points module and/or changing the order by which some tasks are executed, i.e. some strategies are adopted.

III. VALIDATION ACTIVITIES

A. Operational Input for the Strategic Layer

The main operational inputs for the tactical layer have been collected during interviews with Alitalia Flight Dispatchers that work at the Alitalia Operation Center (OCC). They are the professional figures in charge of defining the flight plans and monitoring the flight execution phase.

The Alitalia Operation Center is responsible of coordinating and managing almost 700 flight per day, of which around 70 are long-haul flights. For each of these flights a flight plan has to be produced by the OCC and then submitted and approved by the CFMU. Long-haul flights' planning is handled manually and starts 6 hours before the scheduled departure time while short and medium-haul flights are handled using an automatic procedure. Dispatchers have to intervene only if the system flags an exception. This process starts 2 hours before the scheduled departure time. In both cases flight dispatcher make use of a dedicated software tool called LIDO Flight.

The planning phase starts by collecting information about the flight such as weather at destination and on the route or the aircraft performances and possible limitations and failures on board. On the basis of the information collected a flight

plan is prepared by optimizing the overall cost of the flight and by ensuring at the same time the safe execution of the flight. For example, the occurrence of a weather perturbation is considered to be an unsafe event and it will be always be avoided even at the cost of travelling a longer route. The costs taken into account always include fuel and ATC fees. Costs related to delays are not taken into account by the software tool but can be evaluated on a case-by-case basis by the flight dispatcher. At this stage no information about other flight trajectories is taken into account. As a result, flight trajectories might be not conflict free.

After the flight plans are prepared, manually or automatically, they are submitted according to the ICAO format to the CFMU through a dedicated system (SITA). The ICAO format contains the take-off and landing times and a list of navigation points with the related flight level. The CFMU recalculates the flight plan using their own models. These models differ from those used by Air Operators. In fact they do not consider the differences in performance that aircraft of the same type may have and they manage the vertical profile of the trajectory in a different way. If the flight plan is rejected the dispatcher is noticed; CFMU gives the reason of the rejection but they do not suggest an alternate solution. Moreover the flight dispatcher is unaware of what other companies are doing. When a flight plan is rejected there is no bargaining between CFMU and the dispatcher. The dispatcher simply submits an updated flight plan and waits for its approval. This process is iterated until the dispatcher has a flight plan approved. Communication between CFMU and the dispatcher takes place almost exclusively through the SITA system. In some case the final flight plan can be discussed with CFMU through a phone call. A schematic representation of the process is shown in figure 3. The information collected that are more relevant for

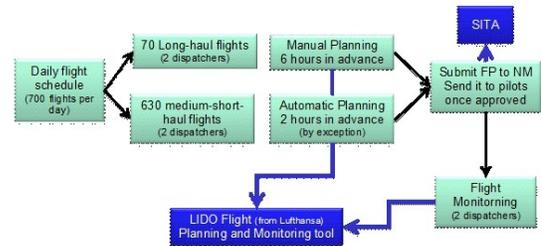


Fig. 3. The process of flight planning and monitoring in the Alitalia OCC.

the development of the strategic ABM are mainly related to:

- The timeframe of the flight plan definition process (6 hours in advance for long-haul flights and 2-hours in advance for medium and short-haul flights)
- The costs taken into account for the flight plan optimization.
- The interactions between the Flight Dispatcher and the Network Manager and the fact that the flight dispatcher is unaware of other companies strategies.
- The flight plan submission process, how flight plans are rejected and submitted again for the final approval

- The criticalities related to the planning phase such as the exceeding of capacity of one or more sector, bad weather avoidance, partial or total closure of destination airport or unpredictable events like strikes, big events, wars.

B. Operational Input for the Tactical Layer

The sector chosen for the calibration of the tactical layer was the sector “LIRROV” in the Rome Area Control Center (ACC), presented in figure 4. An operational expert working in the Rome ACC has been interviewed regarding the features of such sector and most of the operative knowledge acquired has been integrated into the tactical layer.

The sector is crossed by North-South and East-West over-flight traffic as presented in figure 4. The flights inside the sector are mainly commercial, with few exceptions of military flights that however behave like commercial flights. As a result of these traffic pattern, several critical areas emerge from the crossing of these traffic flows, highlighted by red circles in figure 4.

The sector can operate in two main configurations presented in table I. During summer the traffic load is usually high so the sector operates with the configuration “B” in which it is vertically split in order to increase its capacity. On the other hand, during winter when the traffic is lower, it usually operates in configuration “A” where it is composed by just one volume and lesser capacity.

The strategies used to avoid conflicts are both horizontal and vertical. In the first case one of the aircraft involved in a possible conflict is deviated from its original route to achieve horizontal separation, while in the other case a small variation in flight level of just 10FL is used to achieve vertical separation. Combinations of these two strategies are also possible. Despite the fact that horizontal deviations are more convenient in terms of fuel consumption, a small vertical deviation is usually preferred. Moreover in order to reduce the amount of traffic to be managed and the delay generated by their action, controllers usually send aircraft directly to the exit point of the sector after any deviation and whenever possible. Another possibility for the controllers to reduce the traffic load of the sector is to apply a “direct”, i.e. to send an aircraft directly to a point in the next sector of its flight plan. However, since directs require the coordination of the controllers in the involved sectors and thus an increase of their workload, they are considered unlikely events. The tactical layer implements all these findings regarding the traffic patterns inside the sector. Main routes and critical areas are reproduced and also the seasonality of traffic has been considered in the calibration.

Adverse weather conditions occur on a daily basis and are not a negligible effect inside the sector and the system in general. This kind of events does not represent a challenge for the controllers that are always supposed to be capable of handling them. We have been able to identify two major classes of perturbations depending on their dimension:

- Small shocks ($\approx 5 NM$ of radius), with a fast dynamics and a short lifetime ($\approx 1 h$), usually occurring during summertime.

- Large shocks (around $60 NM \times 20 NM$), which can be considered static and with a lifetime that goes from 8 h to 10 h. This kind of perturbations represents big storms occurring during winter.

While the only possible way to manage a small shock is to avoid it, it is possible that the biggest one could be crossed by an aircraft instead of being avoided and thus generating a small delay instead of a large one.

These kinds of shocks have been implemented into the tactical layer, following the discussion with the operational expert. The correlations between size, lifetime and dynamics of the events have been introduced, so that both small dynamical shocks and large static ones are present. In both cases shocks can be avoided using horizontal deviations as well as vertical ones if the considered shock does not affect all the flight levels. Since aircraft might fly through a large perturbation instead of completely avoiding it, a probability of being crossed is assigned to each large perturbation. Moreover in order to simulate correctly the seasonality of the disturbances, small shocks are more likely to occur when simulating summertime while large shocks occur more frequently when simulating winter. After any redirection, due to separation or adverse weather conditions, controllers try to send the aircraft directly towards its exit point in the sector.

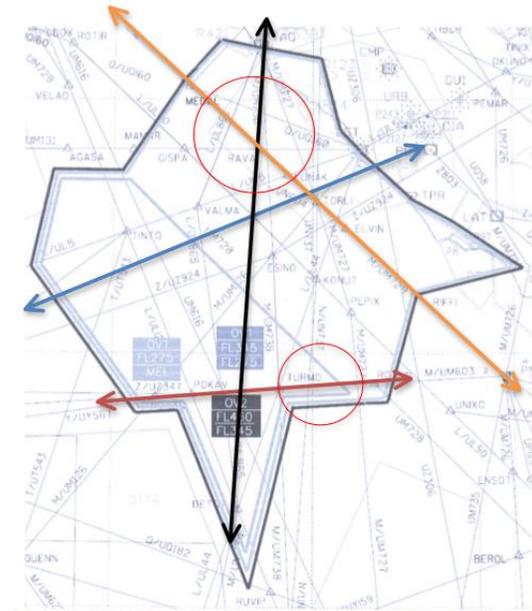


Fig. 4. Projection of the LIRROV sector on a map including the navigation points. Major routes and their directions are indicated by the blue, red, orange and black arrows while the critical areas emerging from their intersection are marked with red circles.

IV. RESULTS

A. The Strategic layer

In this section we present the results obtained so far with the strategic layer. We explore the output of the model by modifying three parameters independently.

TABLE I
OPERATION CONFIGURATIONS OF SECTOR LIRROV. MEL STANDS FOR
MINIMUM EN-ROUTE LEVEL.

Configuration	Sector/s	Heights	Capacity
A	OV1 + OV2	MEL - 460 FL	52
B	OV2	350 + FL	44
B	OV1	MEL - 345 FL	44

The first is the ratio β/α , defining the behavior of the AO. In particular, we use the two extreme configurations: $\beta/\alpha \gg 1$, for which the company cares only about punctuality and takes any path that guarantees the desired departing time, and $\beta/\alpha \ll 1$, for which the company cares only about the length of the path, thus taking the shortest one, possibly shifted in time. In the following, we will call the first type of company “R” (for rerouting) and the second one “S” (for shifting).

Secondly, we tested different patterns of the desired times of departure. Specifically, we use a “wave” structure, as it is the case currently in the air traffic. The wave itself has a duration equal to the average time τ of crossing between sectors, which is also the increment of shifting of the flight plan for the AO. We denote with Δt the time between the end of a wave and the beginning of the next one, using τ as unit. Moreover, we choose 24τ as the interval of possible departure times, meaning that the waves need to occur between $t = 0$ and $t = 23\tau$.

Finally, the last parameter is simply the total number of flights submitted to the NM.

1) *Pure population*: We begin by showing the results of simulations with “pure” populations. This means that each AO has the same cost function, i.e. the same ratio β/α within each simulation.

The top panel of Figure 5 shows how the total satisfaction of the system varies with the number of flights, for different values of the ratio β/α , with $\Delta t = 23$ (only one peak at $t = 0$). As expected, we see that the satisfaction is monotonically decreasing with the total number of flights. Moreover, we see that there are some differences between different values of β/α . Indeed, the initial plateau is more extended for small values of the ratio. Overall, company of type S, with small ratio, tend to be more satisfied that company of type R.

This is actually due to the pattern of desired departing times. On the bottom panel of figure 5, we present the variations of the satisfaction against the ratio β/α , for different values of Δt . As one can see, for big values of Δt , i.e. few and well separated waves, the satisfaction decreases monotonically with the ratio. Company S is always doing better than company R. On the contrary, when there are many waves (small Δt), the satisfaction is increasing with the ratio, i.e. companies R are better off. This is the case because in this situation the companies which shift in time find other companies ahead: it is thus better to change the route instead. It is also interesting to see that there is an intermediate situation ($\Delta t = 1$) for which none of the extremes is better: the companies doing a compromise between length of path and time of departure have a higher satisfaction.

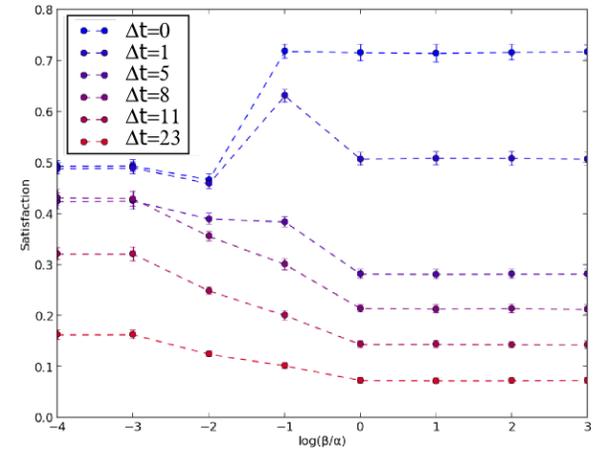
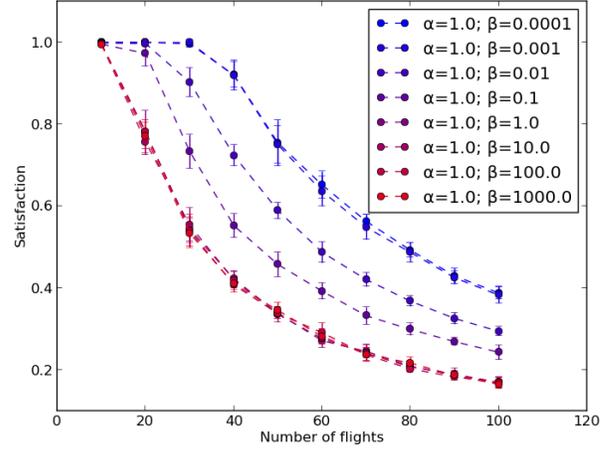


Fig. 5. Top: Satisfaction against number of flights, for different ratios β/α with $\Delta t = 23$. Bottom: Satisfaction against ratio β/α , for different value of Δt and 240 flights. Each point is the results of the average over 100 simulations, with standard errors taken as error bars.

2) *Mixed population*: Now we consider a system where there are two types of companies competing with different cost functions. A fraction of the flights, called “S”, are operated by AOs with $\beta/\alpha = 10^{-3}$. On the other hand, flights called “R” are operated by AOs with $\beta/\alpha = 10^3$.

Figure 6 shows the dependance of the satisfaction of companies from the fraction of “S” flights. This figure, which should be compared to figure 5 (bottom) is strikingly different. In fact, it is clearly not trivial to find a pure population with a behavior similar to the mixing of two extreme companies. In particular, the satisfaction is hardly monotonic now, except for very small values of Δt . We note that the overall tendency is that a uniform distribution of departing time increases the satisfaction, except for very pure populations, as we saw before. Moreover, even a high value of Δt does not strongly favor S companies. In fact, the plot is almost entirely flat for $\Delta t = 23$ and thus the result is insensitive to the fraction of S or R companies.

In figure 7 we show more in details how the satisfaction of each type of company depends on the mixing of the population. In the top panel, the satisfaction of company R is quite simple. First, it is maximum when f_S is close to 1, thus displaying a “the loner, the better” effect: a company gets a higher satisfaction when it is surrounded by companies of the other type, rather than its own. We can understand this result by considering that if everybody wants to shift the flight plan, it is better for a company to reroute, because the secondary routes will be left free. In the bottom panel, we see that the satisfaction of company S displays the same type of behavior: the satisfaction is maximum when the company is alone. However, the variations are more complex here. As one can see, the different curves corresponding to different values of Δt are crossing each other. This means that for different values of f_S , it is better sometimes to have a small Δt , and sometimes a big Δt . For instance, for $f_S = 0.4$, it is better to have $\Delta t = 20$ rather than to have $\Delta t = 1$, which is perfectly normal for S companies. On the other hand, for $f_S = 0.8$, the contrary happens, and suddenly it is better for S companies to have a more uniform distribution of departing times. The whole picture is even more complicated by the fact that this behavior is not monotonic with Δt . With the same example, with $f_S = 0.4$, it is better for S to have $\Delta t = 20$ rather than 1, but it is much worse to have $\Delta t = 23$. This non trivial effect of mixing different companies gives a rich behavior in terms of optimization of the total satisfaction, as noted above.

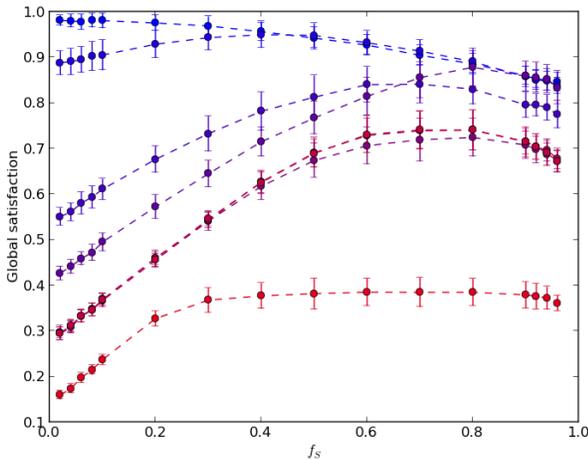


Fig. 6. Total satisfaction against the fraction of flight “S” for different values of Δt and 120 flights. The color legend is on the bottom panel of figure 7.

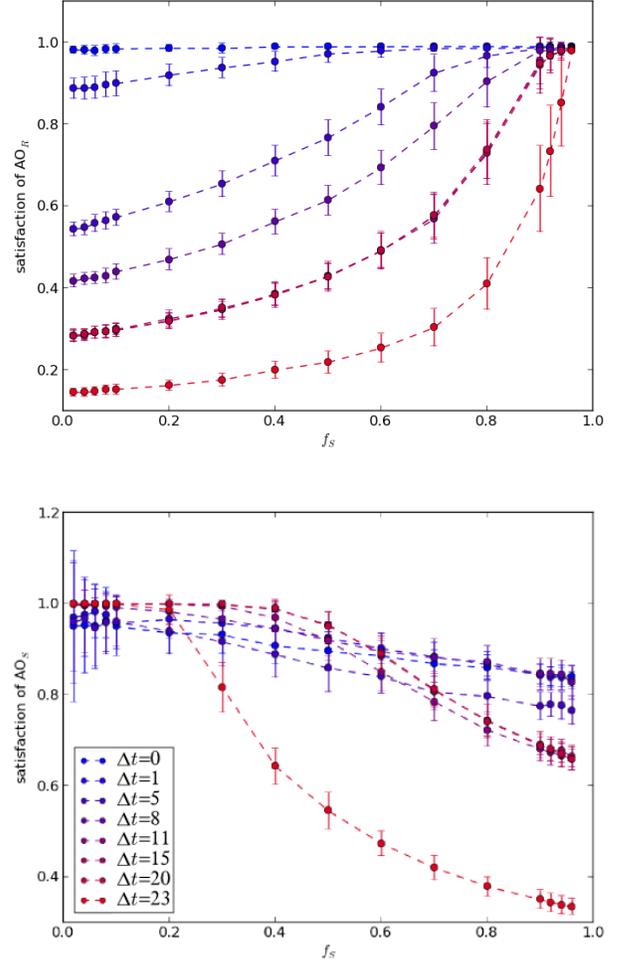


Fig. 7. Top: Average satisfaction of companies R against the fraction of S-flights. Bottom: Average satisfaction of companies S against the fraction of S-flights. The total number of flights is equal to 120. The same color legend (on the bottom) applies to both plots.

3) *Shocks*: Finally, we present some simulations in order to study the resilience of the system when some sectors are closed. With the same network of the above results, we simulate a series of “shocks”. These are obtained by deleting some nodes from the network, simulating closed sectors (for every time). The exact procedure is the following: after a normal simulations, when we have already all the flight plans, we randomly delete a node. We detect all the flights concerned, which recompute alternative flight plans on the new network. Then we delete the second node, and so on until reaching the desired number of shocks.

Figure 8 shows the evolution of the satisfaction during this procedure of shocks. In the top panel we show the results of simulations with pure populations. Here we see that the satisfaction decreases monotonically with the number of shocks, with a small kink for the S company. This kink happens at a value slightly smaller than the percolation threshold (which is 47 ± 8) of this network. In the bottom panel, we show the

results of simulations involving mixed populations. As one can see, in this configuration, the company S begins with an advantage and increases this advantage when more sectors are shocked. Hence, S companies seem to be more resilient to this type of shocks than R companies.

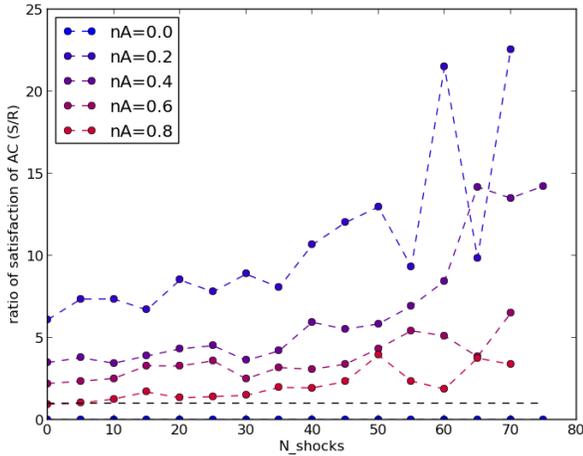
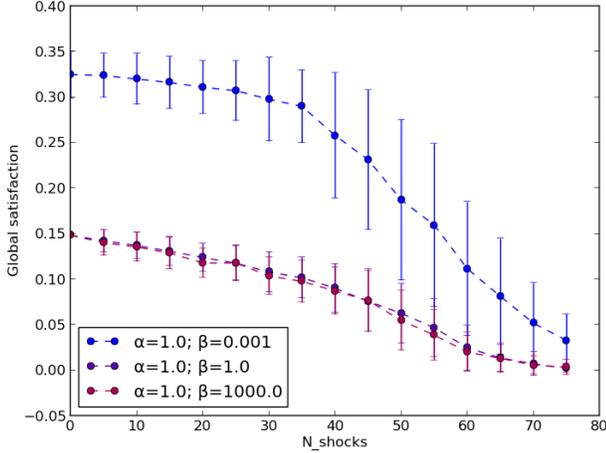


Fig. 8. Top: satisfaction for three types of companies against the number of shocks. The total number of flights is equal to 100 and $\Delta t = 23$. Bottom: evolution of the ratio of satisfaction between companies of type S and type R with the number of shocks, for different values of mixing.

B. The Tactical layer

We model the trajectories of the flights within a single specific sector, i.e. LIRROV2, which is an ATC sector located close to the Fiumicino airport in central Italy, as shown in Fig. 9. We have chosen this specific sector because, from the point of view of its vertical structure, none of its neighbor sectors have a portion in its vertical area, as it can happen in other cases. Moreover, we also consider real M1 trajectories taken from the DDR data already available within the project. We are considering all flights in this sector active on day 06 May 2010. We have a total number $F = 172$ of flights during this day, above the flight level threshold of 200. We will simulate

three main flight levels in the sectors. Initially we will populate the central one. We will then allow aircraft to change upwards or downwards their flight levels in case of conflict.

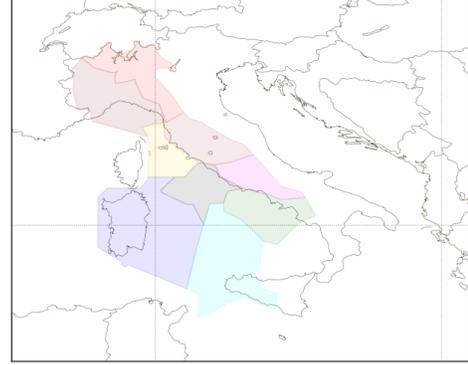


Fig. 9. Sample of Italian sectors. The one we choose is the gray one, fully surrounded by other sectors.

As mentioned above, flights can change their flight level and undergo re-routings. In some special case, when the previous two strategies fail, they can also change velocities. Due to the feedback from operational experts, in the current version of the model the way by which controllers try to solve possible conflicts or avoid shocked areas by firstly do a re-routing, then by changing the flight level and finally, if necessary, by changing the aircraft velocity.

In Fig. (10) we give a graphical representation of a simulation at a certain time instant. The green lines represent the M1 flight trajectories. The blue circles represent the aircraft active in the sector at that time instant in the central flight level. The magenta circles represent the aircraft active in the sector at that time instant in the lowest or highest flight level. The red spot represents a shocked area. The circles representing the aircraft have a radius equal to 2.5 Nautical miles. This is half of the minimum distance at which aircraft can pass close to each other without generating a safety alarm. In the figure one can see that the rightmost aircraft in blue is slightly moving away from its M1 trajectory. In fact, it is starting to modify its trajectory in order to overcome the incoming shocked area in red.

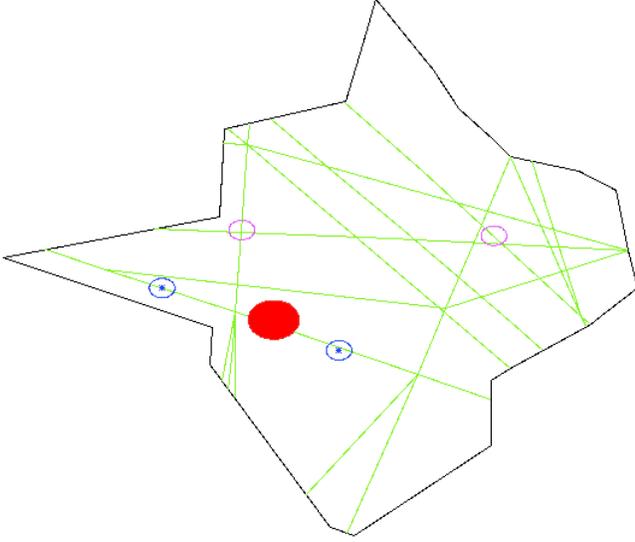


Fig. 10. In the figure we give a graphical representation of a simulation at a certain time instant. The green lines represent the M1 flight trajectories. The blue circles represent the aircraft active in the sector at that time instant in the central flight level. The magenta circles represent the aircraft active in the sector at that time instant in the lowest or highest flight level. The red spot represents a shocked area.

1) *Preliminary Results I:* In the current version of the model shocks are still totally randomly occurring within the sector. Although un-realistic, this is an idealized situation that allows us to investigate how the delays in the flights depend on the number and size of shocks.

In a first set of simulations we considered a situation where we have an average number N_S of shocks per time step. Each time-step is $\Delta t = 300s$ and each shock has a fixed radius of 5 NM. For each simulation experiment we perform $N_E = 5000$ runs. In all performed simulations the change of velocity module did not operate.

As expected, as long as N_S increases, both the percentage of aircraft that re-routed or change their flight level increases, as shown in Table II. The first point $N_S = 0$ corresponds to the case when we do not have shock areas in the model.

TABLE II

PERCENTAGE OF AIRCRAFT THAT ARE RE-ROUTED OR CHANGE THEIR FLIGHT LEVEL. THESE SIMULATIONS ARE PERFORMED ASSUMING THAT THE SHOCKED AREAS HAVE A FIXED RADIUS OF 5 NM.

N_S	re-routed	flight level changes
0	0.159343	0.0977733
1	0.168806	0.100785
2	0.174769	0.101217
3	0.178548	0.103119
4	0.187014	0.106441
5	0.212934	0.114358
6	0.258917	0.134408
7	0.30309	0.151264
8	0.326998	0.161436
9	0.413615	0.227133

A counter-intuitive effect is observed when we consider the average delay of the aircraft. In the top panels of Fig. (11) we show the average (top-left) and the standard deviation (top-right) of the delays experienced by the F flights in the N_E simulations. Average and standard deviations are computed on the re-routed (and therefore delayed) aircraft only. One can observe a trend by which as long as N_S increases the delay gets smaller. This is the opposite of what one would have expected. A possible explanation for this is that as long as the number of shocked areas increases, the percentage of aircraft that change their flight level increases as well, leaving more room in the central flight level and therefore enlarging the possibility that aircraft follow their planned trajectory. In other words, when N_S is large, more aircraft need to be re-routed. However, more aircraft change flight level as well. As a result, the re-routed ones suffer, on average, a small delay because the change in their trajectories is smaller, due to the fact the more room is freed after many aircraft changed flight level. However, one should also notice that such decay is not monotonic.

In the bottom panels of Fig. (11) we show the same variables for the case when the shocked area radius is doubled to 10 NM. Also in this case a non-monotonic behavior is observed. However, in this case we observe a general trend towards an increase in the average delay, despite the fact that many aircraft change flight level. This might be due to the fact that now the area to be overcome in the re-routing is larger.

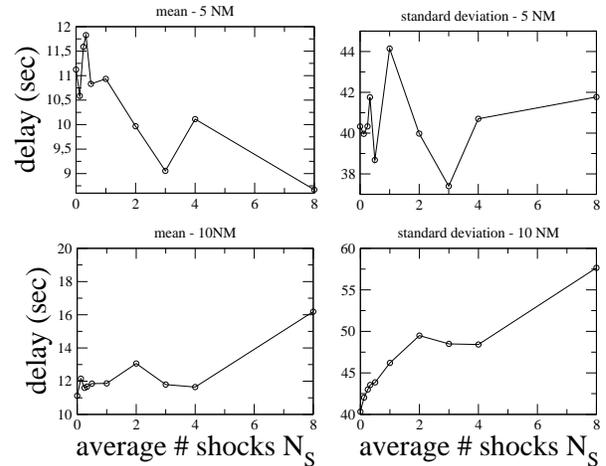


Fig. 11. In the figure we show the average (top-left) and the standard deviation (top-right) of the delays experienced by the F flights in the N_E simulations. These simulations are performed assuming that the shocked areas have a fixed radius of 5 NM in the top panels and 10 NM in the bottom panels.

2) *Preliminary Results II:* In a second set of simulations we considered a situation when we have an average number N_S of shocks per time step and the shocked area is inversely proportional to N_S . We also assume that when $N_S = 1$ then the radius is 5 NM. Each time-step is $\Delta t = 300s$ and for each simulation experiment we perform $N_E = 5000$ runs. In this

way the part of sector not available to aircraft is essentially the same as long as N_S increases. This allows us to test whether delays are depending from the available sector space or from the fact that there are many shocked areas. Again, in all performed simulations the change of velocity module did not operate.

As long as N_S increases, both the percentage of aircraft that re-routed or change their flight level increases, as shown in Table III. It is worth noticing that the percentage of re-routed aircraft is smaller while the percentage of aircraft that change flight level increases with respect to the previous case. The first point $N_S = 0$ corresponds to the case when we do not have shocked areas in the model. As expected, the point $N_S = 1$ has essentially the same re-routings and flight changes as before.

TABLE III

PERCENTAGE OF AIRCRAFT THAT ARE RE-ROUTED OR CHANGE THEIR FLIGHT LEVEL. THESE SIMULATIONS ARE PERFORMED ASSUMING THAT THE SHOCKED AREAS IS INVERSELY PROPORTIONAL TO N_S AND SUCH THAT WHEN $N_S = 1$ THE RADIUS IS 5 NM.

N_S	re-routed	flight level changes
0	0.159343	0.0977733
$\frac{1}{8}$	0.218005	0.166307
$\frac{1}{4}$	0.166115	0.187858
$\frac{1}{3}$	0.171248	0.191772
$\frac{1}{2}$	0.177301	0.199638
1	0.189807	0.216972
2	0.210333	0.241837
3	0.222334	0.258345
4	0.227505	0.267824
8	0.274455	0.326434

In the top panels of Fig. (12) we show the average (top-left) and the standard deviation (top-right) of the delays experienced by the F flights in the N_E simulations. Average and standard deviations are computed on the re-routed (and therefore delayed) aircraft only. The non-monotonic behavior previously observed is again visible. As a general trend it seems that larger values of N_S here give rise to a smaller delay. A possible explanation for the observed effect is that, when the same amount of not available space is fractioned into smaller pieces, the model will change the flight level more frequently and therefore the ones that experience a re-routing will suffer an overall smaller delay. A similar situation is shown in the bottom panels of the figure, where we show the same variables for the case when the radius of the shocked areas is doubled.

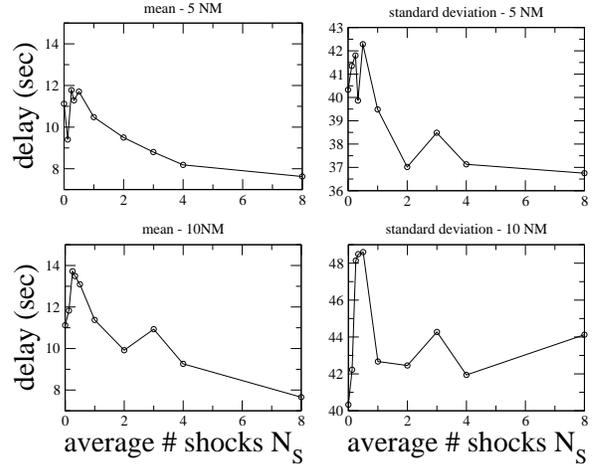


Fig. 12. In the figure we we show the average (top-left) and the standard deviation (top-right) of the delays experienced by the F flights in the N_E simulations. These simulations are performed assuming that the shocked areas have a radius inversely proportional to N_S and such that when $N_S = 1$ the radius is 5 NM in the top panels and 10 NM in the bottom panels.

Further investigations are needed in order to better characterize the non-monotonic behavior observed above. In particular, it is expected that when there are numerous shocking areas they might overlap, thus making the sector space available to aircraft larger than one would have expected if they were not overlapping. However, this effect should not be a major one. In fact, the non-monotonic behavior is observed either in the simulations shown in Fig. 11 and in those of Fig. 12 where the shocked areas have a smaller radius when N_S is large. Another feature worth of further investigation is the role of the directs that again improve the capacity of the considered flight level.

V. CONCLUSIONS AND FUTURE WORK

In this paper we have described the main features and the preliminary results of an agent based model aiming at modeling the management procedure of the air traffic system in the current ATM scenario. The model mainly consists of two layers, the strategic and the tactical one. The strategic layer aims at modeling the interactions between the Air Operators and the network manager in the process that leads to the generation of the flight plans. The tactical layers aims at describing the interactions between air traffic controllers and aircraft/pilots in the process that lead to the actual flights trajectories.

The preliminary results for the strategic layer show that the proto “low-cost” and “traditional” air operators have different advantages depending on the departure pattern and the level of mixture. Even if it is always better for companies to be surrounded by companies of the other type – hence displaying a “the loner, the better” effect, the dominant strategy might also depend on the distribution of departure times, in a non trivial way. For instance, on the contrary of the general case, it is sometimes better for “S” companies to have a more

uniform distribution. Finally, we find that the advantage of “S” companies over “R” companies increases in a constrained environment, leading to a higher resilience for the type of shocks we considered.

The preliminary results for the tactical layer indicate that when shocks in the system are confined in small areas, the interplay between the re-routing and change of flight level strategies may even lead to trajectory modifications that give smaller average delays as long as the number of shocks increases.

Future work will be needed to integrate the two layers into a truly single model: the output of the strategic layer would be a set of flight plans detailed at the level of navigation points to be feed into the tactical layer. On the other side, the tactical layer should be generalized as to take into account the possible interactions between sectors. Moreover, in the tactical layer correlations between shocks should be introduced at the level of their temporal as well as spatial occurrence.

ACKNOWLEDGMENT

This work is co-financed by EUROCONTROL on behalf of the SESAR Joint Undertaking in the context of SESAR Work Package E - ELSA research project.

REFERENCES

- [1] SESAR, “Definition of the future atm target concept - d3,” 2007.
- [2] “Commission regulation (EU) no 691/2010,” 2010.
- [3] EUROCONTROL, “Final report on european commissions mandate to support the establishment of functional airspace blocks (fabs),” 2005.
- [4] Complex World - SESAR WP-E Research Network for the theme “Mastering Complex Systems Safely” , *The ComplexWorld Position Paper* , Version: August 2012, (Restricted audience)
- [5] F. Lillo, S. Miccichè, R. Mantegna, V. Beato, and S. Pozzi, “Elsa project: Toward a complex network approach to atm delays analysis,” *Proceedings of the SESAR Innovation Days*, 2011.
- [6] David E. Goldberg *Genetic Algorithms in Search, Optimization, and Machine Learning* (Addison-Wesley Professional, 1 edition, 1989)
- [7] G Van Rossum; Fred L. Drake Jr *Python reference manual* (Centrum voor Wiskunde en Informatica, Netherlands, 1995)
- [8] Brian W. Kernighan; Dennis M. Ritchie *The C Programming Language* (Prentice Hall, England, 1988)