

Optimized flight level allocation at the continental scale

Cyril Allignol, Nicolas Barnier, Alexandre Gondran

► **To cite this version:**

Cyril Allignol, Nicolas Barnier, Alexandre Gondran. Optimized flight level allocation at the continental scale. ICRAT 2012, 5th International Conference on Research in Air Transportation, May 2012, Berkeley, United States. pp xxxx, 2012. <hal-00938896>

HAL Id: hal-00938896

<https://hal-enac.archives-ouvertes.fr/hal-00938896>

Submitted on 1 Apr 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimized Flight Level Allocation at the Continental Scale

Cyril ALLIGNOL, Nicolas BARNIER and Alexandre GONDRAN
Laboratoire de Mathématiques Appliquées, Informatique, Automatique pour l'Aérien
École Nationale de l'Aviation Civile
Toulouse, France
alignol,barnier,gondran@recherche.enac.fr

Abstract—As acknowledged by the SESAR program, current ATC systems must be drastically improved to accommodate the predicted traffic growth in Europe. This study aims at assessing the performance of 4D-trajectory planning and strategic deconfliction, two of the key concepts identified to meet SESAR capacity objectives.

Among the possible degrees of freedom, we focus here on the flight level (FL) optimization to avoid en-route conflicts between intersecting flights. The resulting problem can be reduced to Graph Coloring with a specific cost function minimizing the discrepancies to requested FLs.

This FL allocation leads to very large combinatorial optimization problems when applied at the continental scale, especially when considering temporal uncertainties. The instances were solved with a Tabu Search algorithm in a few seconds to a few minutes, depending on size and conflict density. Our results shows that the global conflict resolution workload is alleviated by at least 20 %, while bounding the individual FL discrepancies to three levels for a small proportion of the traffic.

Keywords: Deconfliction, Flight Level Allocation, Graph Coloring, Tabu Search

I. INTRODUCTION

In an already saturated European sky, the predicted growth of air traffic volume (+50 % in 2025 compared to 2011 according to [1]) urges to improve Air Traffic Management (ATM) efficiency, as attested by the European Single Sky program SESAR. Current ATM optimization strategies, like reducing the size of control sectors or the distance of separation (RVSM, P-RNAV), seem to have reached the structural limits of the system, while the automation of Air Traffic Control (ATC) has known few significant improvements over the last decades [2].

Among the key concepts identified to meet SESAR performance objectives, the planning of 4D-trajectories would allow to increase en-route capacity, while preserving the current level of safety. Several degrees of freedom on the trajectories can be exploited to regulate

the traffic in order to make the number of potential conflicts decrease and be able to improve the airspace capacity. Previous studies have considered re-routing [3], [4], ground holding to satisfy sector capacities [5], [6] or to deconflict the traffic [7], but few are dedicated to the optimization of the cruise flight level (FL) to vertically separate intersecting flows [8], [9], avoiding any potential conflict between two aircraft of distinct flows. Moreover, recent studies [10] have assessed that more than 30 % of potential conflicts occur between pair of flights that are both in their cruise phase (for traffic data in the French airspace), so that an optimized FL allocation could significantly lower the conflict resolution workload of controllers.

In the current tactical regulation system managed by the Central Flow Management Unit (CFMU) of Eurocontrol, hardly any optimization is attempted¹ on the cruise Flight Levels. So, with each flight plan filed by airlines, a FL depending on the type of aircraft and the length of the trajectory considered is requested and then granted by Eurocontrol. As the performances of aircraft for the same range are likely to be similar, the same FLs are requested over and over, which generates high traffic loads on specific layers of airspace. However, the Requested Flight Level (RFL) of a flight plan usually corresponds to an optimal cruise altitude w.r.t. fuel consumption or CO₂ emission. Thus the distance between the allocated flight level and the RFL should be kept as low as possible, should such a scheme be accepted by airlines.

Of course, flow-based FL allocation schemes, which approximate trajectories in the horizontal plane without taking the climb or descent phases into account and gathers flights sharing the same origin, destination (and possibly RFL) only prevent conflicts occurring between

¹Only the static semi-circular rule that avoids face-to-face conflicts impacts flight level allocation.

distinct leveled flows in their cruise phase. So conflicts involving at least one vertically evolving trajectory, or catch up flights belonging to the same flow, will still have to be resolved by other means (ground holding, ATC actions...). Furthermore, [10] has shown that such a flow-based FL allocation corresponds to very dense conflict graphs and tends to generate very costly solution requiring large discrepancies from the RFLs.

To partially overcome these issues, we propose a FL allocation scheme considering each flight individually and taking into account the temporal dimension of conflicts, whereas previous works considered any given route corresponding to a flow as occupied all day long or at least during the whole flight time of each aircraft of the flow. Such a scheme allows to take catch up conflicts into account and generates much more sparse conflict graphs, depending on the uncertainties assumed on the temporal following of the trajectory.

Several optimization paradigms have been evaluated to solve the corresponding allocation problem: evolutionary algorithm, tabu search and constraint programming. We will focus here on the tabu approach as it appeared to consistently outperform the two others in our experiments, producing solutions with maximal discrepancies of only two or three FL for a limited proportion of the traffic (typically 30 %) in a few minutes of computation time.

In the following sections, we first briefly introduce the current ATC and Air Traffic Flow Management (ATFM) system in Europe and review the related research projects on FL allocation. Then we describe our conflict model and the resulting FL allocation combinatorial optimization problem taking each aircraft individually into account. Next, our results on instances of expanded sizes over one day of traffic in the ECAC² airspace, with varying amounts of temporal uncertainties, are presented. We end with planned further works to enhance the approach before concluding.

II. CONTEXT AND RELATED WORKS

The next section briefly describes current ATFM policies applied by Eurocontrol in the ECAC airspace, through its operational unit, the CFMU, which centralizes the filing of flight plans and issues ATFM regulations to prevent the overload of control sectors. However, no optimization is performed on the RFLs during this process, leading to saturated airspace layers.

In the following section, we review the few research projects that have attempted to address this issue with flow-based models in order to improve airspace capacity, and compare them with our approach.

A. ATC and ATFM

Air Traffic Control (ATC) is a ground-based service provided to ensure the safety and efficiency of the flow of aircraft. The first goal of ATC is to maintain aircraft *separated*: outside Terminal Areas (TMA) around airports (where specific procedures are applied), two aircraft should remain distant from each other at least by 5 NM horizontally and 1000 ft vertically, as illustrated by the safety volume of figure 1.




Fig. 1. Vertical and horizontal separation. Another aircraft cannot be inside the cylinder at the same time.

The overall system currently implemented in Europe to achieve this goal can be conceptually divided in several layers or filters with decreasing time horizon w.r.t. the flight date of the traffic concerned:

- 1) *Strategic* (several months), ASM (Air Space Management): design of routes, sectors and procedures (e.g. reduced separation (RVSM) since 2002, Area Navigation (RNAV) with fictive beacons...).
- 2) *(Pre-)Tactical* (a few days to a few hours), ATFM: definition by each ATC Centre of open sectors with their capacities (opening schedules). To respect these constraints, the CFMU computes and updates flow regulations and reroutings according to the posted flight plans and resulting workload excess.
- 3) *Real time* (5/15 min), ATC: surveillance, coordination with adjacent centres, conflict resolution by various simple manoeuvres (heading, flight level, speed) transmitted to the pilots.
- 4) *Emergency* (less than 5 min), safety nets: ground-based (Short Term Conflict Alert, Minimum Safety Altitude Warning) and airborne (Traffic Alert and Collision Avoidance System, Ground Proximity Warning System).

Our FL optimization approach would be integrated during the pre-tactical phase, with possible updates a

²European Civil Aviation Conference

few hours before takeoff, as the computation times of our resolution algorithm would be compatible with the current practice of the CFMU.

B. Research Projects on Flight Level Allocation

Even if CFMU experts can balance traffic load by separating flights of the same flows over alternative routes, no real optimization of the FL allocation is currently performed at the strategic or pre-tactical level in Europe, except for a few static rules attributing the parity of FL according to the heading of flights to prevent face-to-face conflicts. At the real time level, Air Traffic controllers can issue temporary FL changes to separate aircraft, and this kind of vertical manoeuvre has been integrated in the CATS traffic simulator and conflict solver [11] developed at the French Civil Aviation R&D centre (DSNA/DTI).

Our work is more similar to the flow-based approaches presented in [8], [9] where graph coloring techniques are applied to optimize the FL allocation of direct routes for the French or European traffic, or in [4] where a Genetic Algorithm and an A* algorithm are used to optimize traffic flows with possible changes of FL along the route (as well as lateral deviations for direct routes). However, as our optimization scheme allocate a FL for each flight individually and take the temporal dimension accurately into account, our instances are much larger and the conflict graph much sparser.

Furthermore, contrary to these studies, our model does not fit well within standard graph coloring problems as the choice of FL is very much restricted for operational reasons (detailed in section III-B), so each flight will only have a couple of possible FL available above or below its RFL. Moreover, the FL allocation phase should handle over-constrained instances as well to obtain the best possible solution, even if some conflicts remain (which will have to be solved by other ATC or ATFM actions). Thus, classic coloring techniques providing lower bounds on the cost of solutions like the use of cliques or maximal flow algorithms (cf. the “all-different” constraints used in [9]) cannot be used to speed up the search. Furthermore, the cost of solutions is measured w.r.t. the number of remaining conflicts and the sum of discrepancies from RFLs, the latter not being a conventional criteria for graph coloring problems and algorithms. Consequently, our tabu search algorithm has been tailored (as described in section IV) to handle such variations on the standard graph coloring problem, whose main optimization criterion is the number of colors (i.e. FLs).

III. DECONFLICTION WITH FLIGHT LEVEL ALLOCATION

The flight level allocation is aimed at vertically separating intersecting aircraft, in order to reduce the complexity of the traffic. This FL allocation is performed according to the following steps:

- computation of intersections between aircraft in the horizontal plane;
- allocation of a FL to each flight in such a way that two intersecting flights are given a different FL (if possible).

This allocation does not take the climbing or descending parts of trajectories into account, so that only conflicts occurring in the cruise phase will be solved, which accounts for approximately 25 % of the total number of conflicts in our experiments. The remaining conflicts can be avoided with other strategic methods such as the ground holding scheme presented in [7], or with standard ATC procedures used by controllers.

A. Conflict Detection

Our data are provided by the CATS³ simulator [12], which takes as input all filed flight plans related to the ECAC airspace for a given day of traffic and the relevant airspace configuration. It outputs 4D-trajectories sampled every 15 s, which is the largest time interval that guarantees the detection of the shortest possible conflict (i.e. facing aircraft at maximum speed).

To compute the constraints of the model, trajectories are pairwise probed for conflicting points. Given a flight i , we note:

- $\{p_i^k\}$ the chronologically ordered sequence of the 3D-points of its trajectory;
- t_i^k the time at which flight i will be at point p_i^k .

For any geometrically conflicting points p_i^k and p_j^l such that the separation norm is violated (d_h being the distance in the horizontal plane and d_v in the vertical plane):

$$d_h(p_i^k, p_j^l) < 5 \text{ NM} \quad \text{and} \quad d_v(p_i^k, p_j^l) < 1000 \text{ ft}$$

we must temporally ensure that $t_i^k \neq t_j^l$, which can be rewritten: $0 \neq t_i^k - t_j^l$.

If we consider the whole conflicting zone started at points p_i^k and p_j^l , we obtain the constraint:

$$0 \notin [lb_{ij}, ub_{ij}]$$

where lb_{ij} and ub_{ij} are respectively the lower and upper bounds of the $t_i^k - t_j^l$ for the considered conflict.

³The Complete Air Traffic Simulator developed at DSNA/DTI.

A given pair of flights i and j may conflict several disjoint times over their trajectories, thus leading to several such intervals. The conflict constraint for flights i and j , conflicting σ times, is therefore:

$$0 \notin [lb_{ij}^1, ub_{ij}^1] \cup \dots \cup [lb_{ij}^\sigma, ub_{ij}^\sigma] \quad (1)$$

We note $\mathcal{C}_{ij} = [lb_{ij}^1, ub_{ij}^1] \cup \dots \cup [lb_{ij}^\sigma, ub_{ij}^\sigma]$.

The resulting conflict model is *exact*: a violated constraint necessarily entails a conflict during the simulation and, conversely, any conflict detected between two trajectories in the pre-processing phase will yield a \mathcal{C}_{ij} constraint. More details about this model can be found in [13].

However, aircraft might not follow their 4D-trajectory precisely, so that the above conflict detection shall be carried out with the *real* transit time at point p_i^k :

$$\theta_i^k = t_i^k + \varepsilon_i^k$$

where ε_i^k is a temporal uncertainty.

The constraint for flights i and j becomes:

$$[-\varepsilon_{\max}, \varepsilon_{\max}] \cap \mathcal{C}_{ij} = \emptyset \quad (2)$$

where ε_{\max} is the maximum amplitude of temporal uncertainties: $\forall i, k, |\varepsilon_i^k| \leq \frac{\varepsilon_{\max}}{2}$. For example, $\varepsilon_{\max} = 6$ states that all aircraft are able to follow their 4D-trajectory within a ± 3 min uncertainty on each point. Such an accuracy is out of reach of the current flight management systems, but is within the performance objectives of the SESAR program.

B. Model

Standard models for flight level allocation, such as in [8] or [14], aggregate flights with the same route and possibly the same RFL into flows. FL allocation is then performed considering intersections among flows. However, we propose in this study to allocate FL to flights individually. This new model generates instances of higher dimensions but with a lower density conflict graph, leading to better solutions, i.e. with lower deviation from RFL, as depicted on figure 2 which presents the distribution of discrepancies from RFLs for several days of traffic in the French airspace (optimized by a Constraint Programming algorithm, see [10]).

The variables of the model are the flight levels FL_i for each flight i . In order to keep the operational cost of the allocation as low as possible, flight levels are bounded by a tight interval around RFLs – typically 2 or 3 FL above or below – which induces acceptable over-consumption

of a few percent for airlines (see [15]). If we note RFL_i the RFL for flight i , we ensure that:

$$FL_i \in [RFL_i - \Delta_{\max}, RFL_i + \Delta_{\max}] \quad (3)$$

where Δ_{\max} is an instance parameter defining the maximum deviation from RFL.

Additionally, due to aircraft performances, each FL_i is bound by a value FL_i^{\max} corresponding to the aircraft's ceiling, so that:

$$FL_i \in [RFL_i - \Delta_{\max}, \min(RFL_i + \Delta_{\max}, FL_i^{\max})]$$

Constraints for this model are directly derived from the conflicts presented in section III-A, i.e. whenever two flights i and j are in conflict (w.r.t. equation 1), they must be assigned to different flight levels:

$$0 \notin \mathcal{C}_{ij} \Rightarrow FL_i \neq FL_j \quad (4)$$

As aforementioned, the antecedent of the implication 4 can be replaced by equation 2 to take a given amount of temporal uncertainty into account, thus tightening the density of the conflict graph.

To solve this problem, the primary aim of our search algorithm is to satisfy the greatest possible number of constraints. Then, it tries to minimize a secondary cost that takes airlines preferences into account as described in the next section.

C. Cost of the FL Allocation

When registering a flight plan, airlines must indicate at which FL they wish to fly. This requested flight level usually corresponds to the optimal cruise altitude w.r.t. to the operational cost of the airlines. Flying at a much higher or lower FL might lead to a significant increase in fuel consumption or even be above the aircraft ceiling, thus allocated flight levels must be chosen as close as possible to the RFL.

In order to evaluate the performance of the FL allocation, we define a cost function that measures the deviations from RFL:

$$\text{cost}_{FL} = \sum_i |FL_i - RFL_i|$$

It would also be possible, given a realistic performance model for aircraft, to define a cost function that takes fuel consumption (or CO₂ emission) into account:

$$\text{cost}_{\text{consumption}} = \sum_i c_i(FL_i) - c_i(RFL_i)$$

where c_i gives the operational cost for flight i (depending on aircraft type, travel distance, load factor...) as a function of the FL.




Fig. 2. Discrepancy to RFL for allocation of FL to flows (left) and individual flights FL allocation (right). Flights that obtained their RFL are depicted in light green, flights with a 10, 20, 30 discrepancy are depicted respectively in green, orange and red.

IV. RESOLUTION WITH TABU SEARCH

The flight allocation problem described in the previous section can be considered as a Graph Coloring Problem (GCP): each flight corresponds to a vertex, each conflict detected between two flights (cf. section III-A) to an edge and two connected vertices must be allocated different FLs, which correspond to colors. Usually, the only added constraint is a maximal number of colors k and the problem is then named k -GCP. In our case, the number of FLs is limited for each vertex by operational constraints (cf. section III-B): each flight must be colored with a very tight FL range (compared to the global domain of FLs) around the RFL (relation 3). This problem is known as *list-coloring* problem.

We define in the following sections our slightly modified version of *list-coloring*, then we detail our adaptation of the Tabucol algorithm to solve the FL allocation problem.

A. List-Coloring

Given a graph $G(V, E)$ with vertex set V and edge set E , and, for each vertex u , a color domain D_u , a *list-coloring* of G is a function that allocates to each vertex $u \in V$ a color $c \in D_u$. If two adjacent vertices have the same color, the two vertices and the corresponding edge are said *conflicting*. A solution to our problem is a list-coloring without conflicting edges, named a *legal solution*, that minimizes the discrepancies from RFLs:

$$\begin{aligned} & \underset{S}{\text{minimize}} && \sum_{FL_i \in S} |FL_i - RFL_i| \\ & \text{subject to} && FL_i \in D_{FL_i}, \quad \forall FL_i \in S \\ & && FL_i \neq FL_j, \quad \forall (FL_i, FL_j) \in E \end{aligned}$$

where $S = \{FL_i, \forall i\}$, E is the set of conflicting pairs of flights according to equation 2 and $D_{FL_i} = [RFL_i -$

$\Delta_{\max}, RFL_i + \Delta_{\max}]$ as explained in relation 3.

B. Adaptation to the FL Allocation Problem

To solve our problem, we have adapted the well-known *Tabucol* algorithm [16], an instance of tabu search [17] tailored for k -GCP. [18] improves this algorithm and presents a very good survey of local search methods for graph coloring. As the search space in Tabucol contains all k -colorings, legal or not, the domain constraint of our optimization problem must be relaxed in the objective function.

To adapt Tabucol to our non-conventional cost function, we modify the objective as follows:

$$\begin{aligned} & \underset{S}{\text{minimize}} && \left(\sum_{(FL_i, FL_j) \in E} \delta_{FL_i, FL_j}, \sum_{FL_i \in S} |FL_i - RFL_i| \right) \\ & \text{subject to} && FL_i \in D_{FL_i} \quad \forall FL_i \in S \end{aligned}$$

δ_{FL_i, FL_j} being the Kronecker function which returns 1 if $FL_i = FL_j$ and 0 otherwise. This tailored strategy features two objectives which are considered in lexicographic order. The first one counts the number of conflicting edges in a list-coloring and the second is the sum of discrepancies from RFL.

This modification of the objective allows to find solutions that minimize the number of conflicts – which is the most important goal in practice – even if there is no legal list-coloring (i.e. the first objective equals to zero) or if Tabucol cannot reach one before a reasonable time limit.

C. Principle of the Tabucol Algorithm

As an instance of tabu search, Tabucol is a local search algorithm that attempts to iteratively improve the current solution by selecting (generally) the best candidate in its

neighborhood, i.e. the set obtained by slightly modifying the current solution.

For Tabucol, the neighborhood structure is based on *critical 1-moves*: two list-colorings are neighbors if they have exactly the same FL allocation except for one flight. A critical 1-move then consists in changing the FL of only one flight involved in a conflicting vertex.

The algorithm is initialized by taking the coloring corresponding to the RFL, i.e. $FL_i = RFL_i$, as the current solution. Then, at each iteration of Tabucol, the best critical 1-move which is not tabu and minimizes the objective function is selected. Next, the reverse move of the chosen critical 1-move is added to the tabu list for a given number of iterations, called the tabu *tenure*. The tabu tenure is managed dynamically to prevent all moves from being forbidden as the number of critical neighbors may vary at each iteration.

We take the same parameters as [18] for the tabu tenure and use also the same *aspiration* criteria, i.e. a forbidden move in the tabu list is nevertheless selected if it improves on the currently-known best solution. The algorithm stops when the first objective is equal to zero (no remaining conflict) or after a given number of iterations without improvement of the best solution encountered so far.

V. RESULTS

This model has been implemented and experimented on three data sets of ECAC-wide traffic. The initial data set consists of one day of 2002 traffic, containing around 22 500 flights. The other data sets are build upon this day of traffic by making it uniformly 25 % and 50 % denser respectively, leading to instances with up to 32 000 flights and 700 000 constraints. Those three instances will be referred as ECAC-0, ECAC-25 and ECAC-50 in the following. Table I sums up a few dimensions of the instances, as the number of flights and conflicts. The number of constraints corresponds to the number of edges in the conflict graph (“potential” conflicts) and the number of conflicts is obtained by allocating each flight to its RFL. The resulting conflict graphs have a density on the order of 1 %. As Tabucol is a stochastic algorithm, it is run 10 times for each instance, with a fixed number of one million iterations without improvement, in order to show the stability of solutions. The results displayed in the following are, for each instance, the best solution encountered (in terms of solved conflicts) and the average over the 10 different runs.

TABLE I
INSTANCE DIMENSIONS.

Instance (# flights)	ε_{\max}	Conflicts	Constraints
ECAC-0 (22 453)	0	8 213	153 594
	1	12 620	222 473
	2	17 013	290 562
	3	21 290	358 296
ECAC-25 (27 310)	0	12 047	223 217
	1	18 882	326 674
	2	25 609	428 602
	3	32 327	529 847
ECAC-50 (32 156)	0	16 565	305 191
	1	26 241	449 381
	2	35 736	591 101
	3	45 270	731 969

TABLE II
IMPACT OF Δ_{\max} ON REMAINING CONFLICTS AND COST_{FL}
($\varepsilon_{\max} = 0$).

Instance	Δ_{\max}	Remaining conflicts			cost _{FL}	
		best	(aver.)		best	(aver.)
ECAC-0	30	0	(0)	0 %	6637	(6656)
	20	10	(13)	0.1 %	6609	(6612)
	10	289	(310)	3.5 %	7622	(7482)
ECAC-25	30	0	(0)	0 %	9674	(9724)
	20	24	(30)	0.2 %	9962	(9864)
	10	688	(727)	5.7 %	10201	(10071)
ECAC-50	30	0	(0)	0 %	13553	(13634)
	20	39	(47)	0.2 %	14437	(14365)
	10	1445	(1513)	8.7 %	12773	(11199)

A. FL Allocation

All instances where no uncertainty is considered ($\varepsilon_{\max} = 0$) are solved entirely for $\Delta_{\max} = 30$, with 73 % to 80 % flights obtaining their RFL and only 7 % to 12 % flights being deviated of more than one FL.

Table II gives trade-offs between the number of remaining conflicts and the value of the Δ_{\max} parameter. As expected, the number of remaining conflicts increases when flexibility is reduced (i.e. instances cannot be solved with very low discrepancy from RFL). However, reducing Δ_{\max} generally leads to lower values of cost_{FL} as unsolved conflicts do not contribute to this criterion. The distributions of discrepancies to RFL for different values of the maximum deviation are shown on figure 3. We observe that the variation of Δ_{\max} has little influence on the number of flights obtaining their RFL.

B. Robustness to Temporal Uncertainty

The results presented in the previous section consider no uncertainty on the trajectories. We tested our solutions by introducing random uncertainties (uniformly chosen in the interval $[-\varepsilon_c, \varepsilon_c]$) on takeoff times in the CATS

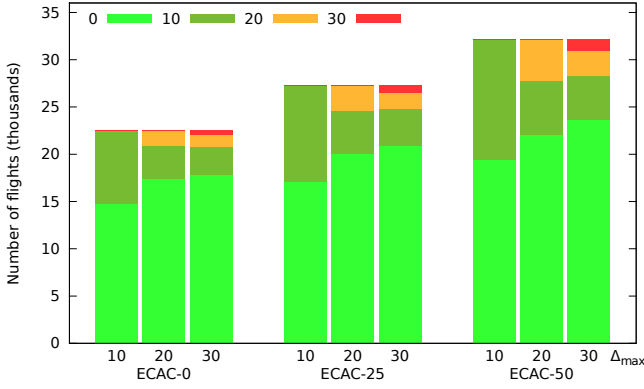


Fig. 3. Deviation from RFL vs Δ_{\max} .




Fig. 4. Remaining en-route conflicts during simulations vs ϵ_c for various values of ϵ_{\max} .

simulator and measuring the number of remaining conflicts.

Figure 4 shows the percentage of remaining en-route conflicts observed during the CATS simulations for different values of ϵ_{\max} and ϵ_c .

It could be expected to observe no remaining conflicts for instances where $\epsilon_c \leq \epsilon_{\max}$ (at least for fully solved instances, see table III). However, changing flight levels leads to a slight modification of trajectories (cruise phase is shortened when FL is increased and the ground speed modified accordingly), which impacts the initial conflict detection. For higher values of the ϵ_c ($\epsilon_c \geq 2$), handling uncertainties as presented in section III-B gives good results as the number of remaining conflicts is reduced when ϵ_{\max} increases.

Nevertheless, increasing ϵ_{\max} leads to harder instances, which has a direct influence on the ability to solve all conflicts and on the cost of the solutions. Table III sums up the number of non-satisfied constraints and the operational cost for all instances processed. The

TABLE III
IMPACT OF HANDLING A ϵ_{\max} UNCERTAINTY ON REMAINING CONFLICTS AND COST_{FL} ($\Delta_{\max} = 30$).

Instance	ϵ_{\max}	Remaining conflicts			cost _{FL}	
		best	(aver.)		best	(aver.)
ECAC-0	0	0	(0)	0%	6637	(6656)
	1	0	(0)	0%	10976	(11089)
	2	0	(1.6)	0%	17344	(17479)
	3	96	(110)	0.5%	22915	(22620)
ECAC-25	0	0	(0)	0%	9674	(9724)
	1	4	(10)	0.02%	17720	(17610)
	2	80	(96)	0.3%	26907	(26758)
	3	790	(822)	2.4%	31034	(30931)
ECAC-50	0	0	(0)	0%	13553	(13634)
	1	16	(28)	0.06%	26481	(26359)
	2	527	(573)	1.5%	35783	(35071)
	3	2184	(2313)	4.8%	40321	(39189)




Fig. 5. Deviation from RFL vs ϵ_{\max} .

constraint graph for handling a ± 3 min uncertainty is two to three times denser than with no uncertainty, which explains the observed results: cost_{FL} is twice higher and up to 9% of the constraints cannot be satisfied. The distributions of discrepancies to RFL associated to these solutions are depicted on figure 5. For the biggest instances, there is still more than 50% of the flights obtaining their RFL, but almost 30% are deviated of two or three flight levels.

Computing times are highly correlated to the instance size and density, and vary from a few seconds to 50 min on a standard workstation (Opteron 6128 at 800 MHz).

VI. CONCLUSION AND FURTHER WORK

We have presented a novel flight level allocation scheme to avoid the conflicts occurring during the cruise phase of intersecting flights, which can amount to 20-30% of all conflicts in a day of traffic in the ECAC airspace, and thus enhance the airspace capacity to accommodate the foreseen traffic growth in Europe. Our

model is based on an exact temporal characterization of each potential conflict computed during a pre-processing phase and, unlike all other approaches known to the authors, allocates a level to each flight individually. This scheme yields much better results than flow-based aggregated ones as reported in [8], [9], and is therefore able to take into account, at a reasonable cost, a given amount of time uncertainty on the following of the trajectories. Furthermore, the discrepancy between the allocated flight level and the one requested in the flight plan is tightly bounded in our model to keep the operational cost of such regulation as low as possible for airlines.

The application of our scheme at the continental scale in Europe, with a traffic comprising more than 30 000 flights daily, leads to huge combinatorial optimization problems, with more than 700 000 conflict constraints for the largest instance. Nevertheless, we were able to solve all instances with at most three levels of discrepancy, from one day of traffic of 2002 up to an expanded instance with 50 % more flights, roughly corresponding to current forecast for 2025. Among the various optimization techniques tried, the best results in terms of number of shifted aircraft, overall cost and computation times were consistently obtained by a tailored version of the Tabucol graph coloring algorithm.

However, our approach does not yet take real-time updates into account, like the cancellation or postponement of a flight, but the pre-processing of conflicts and the low computation time of our resolution algorithm would allow our allocation scheme to be integrated within a dynamic framework like the sliding window algorithm presented in [10]. Another shortcoming of our model is the linear contribution of FL discrepancy in the objective function, which does not exactly correspond to the airlines operational cost in terms of fuel consumption or CO₂ emissions. Our algorithm could though very easily take into account more sensible cost functions, provided we are given enough configuration and performance data by airlines and aircraft manufacturers. Eventually, our approach could be combined with other regulation schemes, as proposed in [13], or real-time conflict solvers, to further reduce the conflict resolution workload of ATC.

REFERENCES

- [1] STATFOR. Eurocontrol long-term forecast: Ifr flight movements 2010-2030. Technical report, Eurocontrol – Air Traffic Statistics and Forecasts, Brussels, décembre 2010.
- [2] Jean-Marc Garot and Nicolas Durand. Failures in the automation of air traffic control. In *Colloque de l'AAA*, 2005.
- [3] Carolyn Sorensen, Diana Liang, and Ian Crook. Limited dynamic re-routing (LDRR). In *USA/Europe Air Traffic Management R&D Seminar ATM2003*, Budapest, Hungary, June 2003.
- [4] David Gianazza and Nicolas Durand. Assessment of the 3d-separation of air traffic flows. In *USA/Europe Air Traffic Management R&D Seminar ATM2005*, Baltimore (MD), USA, June 2005.
- [5] M. Dalichamp, E. Petit, U. Junker, and J. Lebreton. Innovative slot allocation (ISA). Technical report, Eurocontrol, 1997.
- [6] Nicolas Barnier, Pascal Brisset, and Thomas Rivière. Slot allocation with constraint programming: Models and results. In *International Air Traffic Management R&D Seminar ATM2001*, Santa Fe (NM), USA, December 2001.
- [7] Nicolas Barnier and Cyril Allignol. 4D-Trajectory deconfliction through departure time adjustment. In *International Air Traffic Management R&D Seminar ATM2009*, Napa (CA), USA, July 2009.
- [8] Vincent Letrouit. *Optimisation du réseau de routes aériennes en Europe*. PhD thesis, Institut National Polytechnique de Grenoble, 1998.
- [9] Nicolas Barnier and Pascal Brisset. Graph coloring for air traffic flow management. In *CPAIOR'02*, pages 133–147, Le Croisic, France, March 2002.
- [10] Cyril Allignol. *Planification de trajectoires pour l'optimisation du trafic aérien*. PhD thesis, Mathématiques Informatique Télécommunications, Toulouse, December 2011.
- [11] Géraud Granger, Nicolas Durand, and Jean-Marc Alliot. Optimal resolution of en route conflicts. In *International Air Traffic Management R&D Seminar ATM2001*, Santa Fe (NM), USA, December 2001.
- [12] Jean-Marc Alliot, Jean-François Bosc, Nicolas Durand, and Lionel Maugis. CATS: A complete air traffic simulator. In *16th DASC*, 1997.
- [13] Nicolas Barnier and Cyril Allignol. Combining flight level allocation with ground holding to optimize 4D-deconfliction. In *USA/Europe Air Traffic Management R&D Seminar ATM2011*, Berlin, Germany, June 2011.
- [14] Nicolas Barnier and Pascal Brisset. Graph coloring for air traffic flow management. In *CPAIOR'02: Fourth International Workshop on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimisation Problems*, pages 133–147, Le Croisic, France, March 2002.
- [15] Airbus. Getting to grips with fuel economy. Technical Report 3, Flight Operations Support and Line Assistance, Toulouse, juillet 2004.
- [16] Alain Hertz and Dominique de Werra. Using tabu search techniques for graph coloring. *Computing*, 39:345–351, 1987.
- [17] Fred Glover. Tabu search - part i. *ORSA Journal on Computing*, 1(3):190–206, 1989.
- [18] Philippe Galinier and Alain Hertz. A survey of local search methods for graph coloring. *Computers and Operations Research*, 33:2547–2562, September 2006.