



Effects of increased flexibility for airspace users on network performance

D5.4

COCTA

Grant:

699326

Call:

H2020-SESAR-2015-1

Topic:

SESAR-05-2015

Consortium coordinator:

Univerzitet u Beogradu – Saobraćajni fakultet

Edition date:

14 December 2018

Edition:

01.00.00

Founding Members



EUROPEAN UNION



EUROCONTROL



Authoring & Approval

Authors of the document

Name/Beneficiary	Position/Title	Date
Arne Strauss (AS)/UW	Principal Researcher (OR)	13/12/18
Xin Fei (XF)/UW	Researcher (OR)	13/11/18
Stefano Starita (SS)	Post-doc Researcher (OR)	13/11/18
Nikola Ivanov (NI)/UB-FTTE	Researcher (ATM)	13/12/18
Goran Pavlović (GP)/UB-FTTE	Researcher (Data & ATM)	13/12/18
Frank Fichert (FF)/HW	Principal Researcher (Economics)	13/12/18
Thi Thuy An Vo (AN)/HW	Task manager	13/11/18

Reviewers internal to the project

Name/Beneficiary	Position/Title	Date
Obrad Babić (OB)/UB-FTTE	Senior expert (ATM)	13/11/18
Radosav Jovanović (RJ)/UB-FTTE	Principal Researcher (ATM)	14/12/18

Approved for submission to the SJU By — Representatives of beneficiaries involved in the project

Name/Beneficiary	Position/Title	Date
Radosav Jovanović/UB-FTTE	Principal Researcher (ATM)	14/12/18
Arne Strauss/UW	Principal Researcher (OR)	14/12/18
Frank Fichert/HW	Principal Researcher (Economics)	14/12/18

Rejected By - Representatives of beneficiaries involved in the project

Name/Beneficiary	Position/Title	Date
------------------	----------------	------

Document History

Edition	Date	Status	Author	Justification
00.00.01	06/08/18	Initial draft	AS, XF, SS	Draft
00.00.02	20/08/18	Second draft	AS, SS, XF, RJ, NI, GP, OB, FF, AN	Draft (inc. choice model)
00.00.03	28/08/18	Final draft	AS, SS, RJ, NI, GP, FF, AN	Final draft
00.01.00	07/09/18	Final	AS, SS, RJ, NI, GP, FF, AN	Final
00.02.00	13/11/18	Final updated	AS, SS, NI, GP, FF, AN	Final updated to include results from

2 The opinions expressed herein reflect the author's view only. Under no circumstances shall the SESAR Joint Undertaking be responsible for any use that may be made of the information contained herein.



				testing the new model
01.00.00	14/12/18	Final	AS, SS, NI, GP, FF, AN	Final version with answers to the SJU comments

COCTA

COORDINATED CAPACITY ORDERING AND TRAJECTORY PRICING FOR BETTER-PERFORMING ATM

This deliverable is part of a project that has received funding from the SESAR Joint Undertaking under grant agreement No 699326 under European Union's Horizon 2020 research and innovation programme.



Abstract

In this deliverable, we incorporate recommendations received from stakeholders, airspace users in particular. Namely, we update the COCTA model, in line with refined COCTA concept of capacity and demand management. In the refined concept airlines will be able to define their preferred trajectories and the flexibility required for each flight, as well as to decide on the final trajectory. We analyse and evaluate the effects on network performance.



Table of Contents

Abstract	4
1 Introduction	6
2 A new Premium Trajectory product	8
3 A Surrogate-assisted Solution Algorithm for the COCTA strategic-ordering model	10
3.1 Motivation	12
3.2 Surrogate Function Approximation	13
3.3 Product choice behavior	16
4 Numerical Experiments Design	20
4.1 Strategic phase	20
4.2 Pre-tactical phase	22
5 Results	23
6 Discussion and conclusions	29
References	30
Appendix A A new solution approach – working paper	31

1 Introduction

In this deliverable, we incorporate recommendations received from stakeholders, Aircraft Operators (AOs) in particular. To increase the acceptance by AOs, it was suggested that a new trajectory product - Premium Trajectory (PT) – shall be introduced, in addition to the two trajectory products already defined in previous COCTA deliverables: Purchase Specific Trajectory (PST) and Flexibly Assigned Trajectory (FAT). For the sake of simplicity, we rename PST as Standard Trajectory (ST) and FAT as Discounted Trajectory (DT).

Both ST and DT are structurally the same: an AO that purchases either of them will acquire the right to fly a specific origin-destination combination for a specified charge, but the network manager (NM) retains the right to decide shortly prior to the departure day which trajectory exactly will be flown (within agreed margins). The only difference is that the margins (spatial or temporal) for DT are wider than for ST, and hence DT will be offered at a discount. PT, however, has a quite different structure. Since it is now the AO who has the right to decide on a trajectory shortly prior to departure, it introduces a higher level of uncertainty on the capacity side. The NM has to account for this additional source of uncertainty in the capacity management (ordering) process at the strategic level.

The added levels of uncertainty are to be addressed by creating a larger number of traffic scenarios, for each of which we would need to solve the capacity ordering problem. Similar to our approach in the deliverable D5.3 (COCTA consortium, 2018), this would then give us a set of sector opening schemes (one for each scenario) from which one then can infer the capacity ordering decision. Essentially, the more uncertainty in the problem, the more scenarios should be considered. However, the existing approach used in the D5.3 relies on a heuristic that is too slow to be used to solve the capacity ordering problem for such a large number of different traffic scenarios. Therefore, we propose a new approximation to this problem that relies on machine learning ideas. It allows to separate the traffic assignment decisions from the sector opening scheme decisions, and thereby makes it possible to (approximately) solve the problem for many scenarios in very little time. The algorithm requires training (learning) which takes a long time, but only needs to be done once. We elaborate further on this in Section 3.

From the strategic to pre-tactical level, i.e. in the trajectory booking horizon, the NM defines prices for all three products. The price of ST forms the base airport-pair charge, chosen in a way so as to recover overall expected capacity provision cost. The product DT is discounted relative to ST to reflect the greater uncertainty and cost to the AO of potential imposed delay or re-routing. Conversely, PT is a superior product and as such will need to be charged at a premium relative to ST. Overall, the aim is to balance discounts and surcharges so that overall capacity provision costs are (just) met.

In Section 2, we summarize the proposed changes in the ATM value chain and introduce a new trajectory product. We define a new mathematical method in Section 3 to solve the strategic

6 The opinions expressed herein reflect the author's view only. Under no circumstances shall the SESAR Joint Undertaking be responsible for any use that may be made of the information contained herein.



capacity ordering problem, followed by the description of research-experiment design. A preliminary discussion is presented in Section 6.

2 A new Premium Trajectory product

The COCTA mechanism combines capacity and demand management actions to optimise network performance. Within the COCTA framework, the mechanism is primarily designed for strategic (6 months in advance) and pre-tactical stages (7 days in advance), while the tactical stage is considered to a certain extent only. In addition, we also discuss long-term (5 years) capacity planning and ordering and in this section, we provide a brief overview of the process as a whole.

The NM carries out capacity management at the network level. Due to relatively long lead times related to the capacity planning and provision process (Tobaruela et al., 2013), the COCTA network capacity planning and management process spans over a 5-year horizon. Similar to the current practice, we assume that the NM and ANSPs agree on capacity profile which needs to be delivered on an annual level over the long term (5 years), with the difference that this agreement is based on contracts in the COCTA concept. This capacity profile is based on long term traffic forecasts and serves as a foundation for ANSP's decisions affecting capacity (e.g. staff training and technical equipment).

When AOs publish schedules, around six months in advance of a schedule season, the NM has more precise information on O&D pairs and respective times of operations. Based on information of scheduled traffic and accounting for a portion of non-scheduled demand - which is associated with a higher level of uncertainty in terms of O&D pairs, times of operations and overall traffic levels - the NM defines capacity orders within the capacity profile sketched above. Therefore, about six months in advance, the NM refines and makes a more precise capacity order from the ANSPs, aligned with the long-term order. The NM asks for capacity from ANSPs, which is measured using sector-hours. Depending on assumed flexibility of capacity provision in terms of ANSPs' staffing practices, i.e. how much in advance ATCOs rostering is defined, the NM can define its initial order as a sector-opening scheme (less flexible) or as a total sector-hours to be delivered on that day, including maximum number of sectors to be open and duration at maximum configuration (more flexible). The capacity management process continues after this decision, with options to slightly adjust the initial capacity order, in line with flight intentions information received/updated subsequently, again, depending on the assumed flexibility of capacity provision.

In the redesigned ATM value-chain, we also foresee a novel approach to demand management, which becomes trajectory (product) management. The trajectory management process (lifecycle) starts at strategic level and spans until a flight has been executed.

At strategic level, the NM demand management is used primarily to establish a cost-efficient balance between demand and capacity. Namely, the NM evaluates if it is more cost-efficient to delay or re-route flights in certain parts of the network, instead of asking ANSPs to provide more capacity. Moreover, in some parts of the network and during certain periods (peak hours), demand profile might be such that even maximum (structural) capacity might not be sufficient to accommodate

8 The opinions expressed herein reflect the author's view only. Under no circumstances shall the SESAR Joint Undertaking be responsible for any use that may be made of the information contained herein.



anticipated demand without delays (or re-routings). Therefore, using available information on flight intentions (scheduled carriers) and anticipated/forecasted level and spatio-temporal distribution of non-scheduled flights (e.g. charters), the NM evaluates what is the scope of demand management actions, combined with capacity ordering (management), which minimises total cost to AOs. As a result from this analysis, the NM has information on capacity needed per ANSP and the scope of delays and re-routings of flights/flows in the network, which establishes a cost-efficient balance between anticipated demand and capacity ordered.

After the initial capacity order, the NM starts defining trajectory products to incentivise AOs' route/trajectory choice to maintain, to the extent possible, the strategically established balance between demand and capacity, which minimises total cost to AOs. By defining and offering different trajectory products to AOs at differentiated prices, the NM aims to (re)distribute flights in a system optimum manner in the network.

For instance, ST is associated to the shortest route between two airports, including relatively narrow and pre-agreed spatio-temporal trajectory margins, necessary for trajectory fine tuning at a later stage (e.g. shortly before take-off). This product comes at a base charge and is tailored for flights/flows which are not using airspace which is likely to be congested, i.e. these flights will most likely not be subject to demand management actions. On the other hand, by choosing DT, an AO gets a lower charge compared to ST, but delegates the decision to the NM to delay or re-route its flight within pre-agreed margins (usually wider than those for ST). With PT, AOs have an option for last minute trajectory changes, either in space or time, within agreed margins; this option comes at a higher charge compared to the ST.

To sum up, the NM offers a range of trajectory products, at differentiated charges, to incentivise AOs' trajectory/route choices to the extent possible, to achieve required network performance.

For more details on COCTA capacity and demand management actions and re-designed ATM value chain, reader is referred to (COCTA consortium, 2017a, 2017b).

3 A Surrogate-assisted Solution Algorithm for the COCTA strategic-ordering model

Sets:

O	Set of origin-destination pairs
F ,	The set of all flights
R_f	The set of routes available to f
U	Time horizon
A	Set of airspaces
C^a, S^a	Set of configurations and elementary sectors for airspace a
P^c	Partition of elementary sectors corresponding to a configuration

Indices:

f	Flights
od	Origin and destination airports
u	Time index
r	Route
a	Airspace
c, c'	Airspace's configuration
p	Airspace sector (collapsed or elementary)

Parameters:

γ_a	Variable cost of providing one sector-time unit for airspace a
K_p	Maximum capacity of airspace portion p
\bar{h}_{ac}	Number of sector-time units consumed by airspace a working in configuration c
d_r^f	Displacement cost of route r for flight f
b_{frpu}	Is equal to 1 if route r uses collapsed sector p at time u , 0 otherwise



Variables:

$$z_{acu} = \begin{cases} 1, & \text{if airspace } a \text{ configuration is } c \text{ at time } u \\ 0, & \text{otherwise} \end{cases}$$

$$y_r^f = \begin{cases} 1, & \text{if flight } f \text{ is assigned to route } r \\ 0, & \text{otherwise} \end{cases}$$

As highlighted in the introduction, the additional uncertainty introduced by the PT product calls for more scenarios to be evaluated in order to appropriately take uncertainty into account. Since the heuristic used so far is too slow to evaluate large numbers of sampled traffic scenarios, we developed another solution method for the COCTA strategic-ordering model (1-6) as defined in D5.3. For the sake of convenience, let us state this problem here again:

The problem is formulated below as a linear binary programme:

$$\min_{z,y} \sum_{a \in A} \gamma_a \sum_{u \in U} \sum_{c \in C^a} \bar{h}_{ac} z_{acu} + \sum_{f \in F} \sum_{r \in R_{od_f}} d_r^f y_r^f \quad (1)$$

$$\text{s.t. } \sum_{r \in R_{od_f}} y_r^f = 1 \quad \forall f \in F \quad (2)$$

$$\sum_{c \in C^a} z_{acu} = 1 \quad \forall a \in A, \quad u \in U \quad (3)$$

$$\sum_{f \in F} \sum_{r \in R_f} b_{frpu} y_r^f \leq K_p z_{acu} + |F| \sum_{c' \neq c} z_{ac'u} \quad \forall a \in A, \quad c \in C^a, \quad p \in P^c, \quad u \in U \quad (4)$$

$$z_{acu} \in \{0, 1\} \quad \forall a \in A, \quad c \in C^a, \quad u \in U \quad (5)$$

$$y_r^f \in \{0, 1\} \quad \forall f \in F, \quad r \in R_{od_f} \quad (6)$$

Note that we have two type of decisions: assignments of trajectories r to flights f via y_r^f , and sector opening scheme decisions z_{acu} over time. The problem becomes particularly challenging because these two decisions are linked via the constraints (4) that express the capacity limits. Breaking up this dependency led us to the new approach as motivated below.

3.1 Motivation

A surrogate-assisted algorithm is a solution approach that uses the actor-critic paradigm in the machine learning community to resolve large-scale optimisation problems. The solution algorithm is based on a generalised policy iteration framework that learns from the past trial results and uses that experience to improve the current solution. In the COCTA strategic-ordering model, we can divide the overall decision problems into two parts: master problem and sub-problem. The master problem is concerned with optimising traffic distribution in a network to minimise the displacement cost, i.e. cost of delays and re-routings, and cost of capacity provision.

Let $Q: \{0,1\}^{|F| \times |R_{odf}|} \rightarrow R$ denote a projection from a trajectory-flight assignment vector \vec{y} to a real value. The projection can be any surrogate function that returns the (approximate) cost of capacity provision (sector-opening scheme) for a given traffic distribution in the network. The master problem can be written as:

$$\mathbf{Master:} \min_{\vec{y}} \sum_{f \in F} \sum_{r \in R_{odf}} d_r^f y_r^f + Q(\vec{y}) \quad (\text{A.1})$$

$$\text{s. t.} \sum_{r \in R_{odf}} y_r^f = 1 \quad \forall f \in F \quad (\text{A.2})$$

$$y_r^f \in \{0,1\} \quad \forall f \in F, r \in R_{odf} \quad (\text{A.3})$$

As presented above, the master problem does not explicitly model the capacity limitations across the network; however, the assignment of a trajectory for each flight would influence the value of the surrogate function and thus capacity cost implications are indirectly captured. Note that we assume that all flights in F have got a set of trajectory options R_{odf} associated with them, and that we can decide on which one will be assigned to the flight. With the introduction of the new premium trajectory product PT, AOs decide on their trajectories. To incorporate this into our framework, we assume that certain flights (for instance, transatlantic and business aviation) have a preferred (single) trajectory associated with them which is beyond our control (i.e. the NM does not decide on this trajectory). In the model, this boils down to having only a single option in the set R_{odf} which would be randomly sampled using historic distributions to represent an educated guess as to what trajectory these flights might take. This reflects again that we require a large set of traffic scenarios (including which route these flights would take) to ensure that we have a good approximation of the overall traffic distribution.

Since the decision variable y_r^f is binary, the number of all possible feasible solutions is finite. Ideally, if we can obtain the surrogate function value of each potential solution (traffic distribution in the network), then we can find the best one associated with the minimum cost for the system as a whole. However, the number of potential solutions is usually very large in the network. To illustrate the complexity, let us consider the following example. If a network has 10,000 flights, each of which has 10 potential trajectories, then there are $10^{10,000}$ different possibilities to distribute this traffic in the network. One of them should be the best traffic distribution that will yield the minimum cost

12 The opinions expressed herein reflect the author's view only. Under no circumstances shall the SESAR Joint Undertaking be responsible for any use that may be made of the information contained herein.



over the whole system. Obviously, the full enumeration search method that evaluates the surrogate function values of all potential solutions is computationally infeasible. Instead, we aim for an algorithm that can return a valid sector opening scheme to the COCTA strategic-ordering model, even if it is interrupted due to the limited computational time. We expect the solution algorithm to find better and better sector opening schemes the longer it keeps running.

We propose a method to approximate the surrogate function values by using parametric regression. The approximate surrogate function represents an approximation of the capacity cost function in dependence of traffic assignments. This allows us to pre-compute this approximation so as to quickly solve the strategic phase model for various simulated traffic patterns. In other words, we now can cover many more traffic samples than beforehand. This is the basic motivation of surrogate-assisted solution algorithm. In the following, we will discuss the approximation of the surrogate function.

3.2 Surrogate Function Approximation

In the simplest setting, we can take the assignment of flight trajectories as the explanatory term of our parametric regression model. Let M_{acu} be the penalty when the maximum network capacity cannot accommodate given traffic distribution $y_r^f, \forall f \in F, r \in R_{od_f}$ and let x_{acu} represent additional necessary capacity in the partition p of configuration c of airspace a at the time period u . This capacity increase can be thought of a means to assist the model to overcome infeasible solutions: solutions with capacity violations can be accepted, but at a high penalty cost expressed by M_{acu} . The value of M_{acu} is assumed to be much larger than the value \bar{h}_{ac} because we aim to force the model using the existing network capacity. Suppose that we are given a traffic distribution $\vec{y} = \{y_r^f | \forall f \in F, r \in R_{od_f}\}$, its surrogate function value $Q(\vec{y})$ is the optimal objective function value of the following decision problem:

$$\text{Sub: } \min_{z_{acu}, x_{acu}} \sum_{a \in A} \gamma_a \sum_{u \in U} \sum_{c \in C^a} (\bar{h}_{ac} z_{acu} + M_{acu} x_{acu}) \quad (\text{A.4})$$

$$\text{s. t. } \sum_{c \in C^a} z_{acu} = 1 \quad \forall a \in A, \quad u \in U \quad (\text{A.5})$$

$$\sum_{f \in F} \sum_{r \in R_f} b_{frpu} y_r^f \leq (K_p + x_{acu}) z_{acu} + |F| \sum_{c' \neq c} z_{ac'u} \quad \forall a \in A, \quad c \in C^a, \quad p \in P^c, \quad u \in U \quad (\text{A.6})$$

$$z_{acu} \in \{0, 1\}, x_{acu} \in \{0, 1, 2, 3, \dots\} \quad \forall a \in A, \quad c \in C^a, \quad u \in U \quad (\text{A.7})$$

Since the traffic distribution y_r^f has been fixed, the subproblem is only related to finding a sector opening scheme z_{acu} , which can be solved quickly. Let κ denote the optimal objective function value of the above subproblem. The training sample can be written as a two-tuple $(\vec{y}, \kappa(\vec{y}))$.

If we can collect enough training samples, then we will obtain a good approximation of surrogate function to solve the COCTA strategic capacity ordering model. This allows a problem to be solved relatively quickly for a fixed set of flights, and also to quickly evaluate the consequences of different trajectory distribution in the network. However, one major shortcoming of using this approach is that we have to re-train the surrogate function if the underlying set of flights is changed since we have to re-train the parametric regression model (note that the input parameter b_{frpu} of the sub-problem depends on the flights in the traffic scenario).

To overcome this shortcoming, we propose a group of so-called basis functions to construct an approximation of the surrogate function. The basis functions are intended to avoid the explicit dependency on the set of flights in the considered traffic scenario by capturing some basic problem features. Let $\varphi_{a,c,p,u}$ denote a basis function that reflects the traffic flow in the partition p of configuration c of airspace a at the time period u . We can transform the traffic distribution y_r^f to the basis function $\varphi_{a,c,p,u}$ as follows,

$$\varphi_{a,c,p,u} = \sum_{f \in F} \sum_{r \in R_f} b_{frpu} y_r^f, \forall a \in A, c \in C^a, p \in P^c, u \in U \quad (\text{A.8})$$

The basis function implicitly includes the flight distribution information and the input parameter b_{frpu} of the sub-problem so that we achieve some aggregation when constructing parametric regression model. Accordingly, the data structure of training sample can be changed to

$$(\{\varphi_{a,c,p,u} | \forall a \in A, c \in C^a, p \in P^c, u \in U\}, \kappa(\varphi)).$$

In other words, we no longer try to establish a direct relationship between the traffic assignment \vec{y} and the resulting capacity cost $\kappa(\vec{y})$, but instead seek to establish a relationship between traffic flows in specific partitions per unit of time and resulting capacity cost $\kappa(\varphi)$. With this, we do not have to specify individual flights; instead, only the aggregated traffic entry counts in space and time are of interest.

The overall procedure is as follows: we first need to generate the training traffic sample and generate an approximation of the surrogate function $Q(\vec{\varphi})$ (Algorithm 1 below). Then, as described in Algorithm 2, we use this approximation (denoted $\hat{Q}(\vec{\varphi})$) to solve the **Master** problem for a given set of flights to obtain the trajectory assignment decisions for all flights, and finally we solve **Sub** for this assignment to obtain the sector configuration \vec{z} .



Algorithm 1 Approximation of Surrogate Function

0. initialise iteration counter $i \leftarrow 0$
1. **WHILE** {maximum runtime has not yet been exceeded}
2. set $i \leftarrow i + 1$
2. randomly sample trajectory assignment for each flight using a given distribution over all available trajectories (given a discrete set of options for each flight) to obtain \vec{y}
3. solve **Sub** with assignment \vec{y} as input to obtain the optimal objective function value κ^{itr} representing capacity (and potentially penalty) costs of this traffic configuration.
4. calculate the basis function $\varphi_{a,c,p,u}^i$ as in (A.8) for all $a \in A, c \in C^a, p \in P^c, u \in U$
5. save the training sample $(\{\varphi_{a,c,p,u}^i | a \in A, c \in C^a, p \in P^c, u \in U\}, \kappa^i)$
6. **END WHILE**
7. Obtain the parameters $\vec{\beta}$ of a parametric linear regression model using the training samples so as to approximate the surrogate function $Q(\vec{\varphi})$:

$$Q(\vec{\varphi}) \approx \tilde{Q}(\vec{\varphi}) = \beta_0 + \sum_{a,c,p,u} \beta_{acpu} \varphi_{acpu}$$

Output: approximated surrogate function $\tilde{Q}(\vec{\varphi})$

Algorithm 2 Solution of Master and Sub problem using Approximation of Surrogate Function

- 1) Solve the decision problem **Master** by using the approximate surrogate function $\tilde{Q}(\vec{\varphi})$ and return the traffic configuration \vec{y}^*
- 2) Solve the decision problem **Sub** by using the traffic configuration \vec{y}^* and return the sector operation scheme \vec{z}^*

Outputs: traffic distribution \vec{y}^* and sector operation scheme \vec{z}^* for the COCTA strategic-ordering model

This approach has been further refined in our paper Starita et al. (2018), attached in Appendix A. We do not reproduce the formulae here; instead, let us briefly outline the main changes between the above version and the one put forward in our working paper. The underpinning idea stays the same; however, we discovered that we can write both master and sub-problem in an equivalent but more efficient form. Both problems can be stated in a way such that their constraint matrices have a certain structure (called “totally unimodular”) that guarantees that we can relax the integer constraints on the decision variables whilst still obtaining an optimal integer solution. This has huge advantages on runtime since it is much more computationally expensive to solve discrete optimization problems than continuous ones.

3.3 Product choice behavior

In previous deliverables, we used a binary logit model that allows only two choice alternatives. With three options, we need to use something else, such as the multinomial logit. Assuming all three options (ST, PT, DT) are available, the purchase probability for DT is given by:

$$Prob_{DT}(\{ST, PT, DT\}) = \frac{\exp(\beta_0^{DT} + \beta_p \frac{p_{DT}}{p_{ST}})}{1 + \exp(\beta_0^{DT} + \beta_p \frac{p_{DT}}{p_{ST}}) + \exp(\beta_0^{PT} + \beta_p \frac{p_{PT}}{p_{ST}})}$$

Here, β_0^{DT} and β_0^{PT} can be regarded as the base utility of the DT and PT product, respectively. The price sensitivity is measured by β_p . Note that, as in previous deliverables, we measure utility not against the absolute prices, but against prices relative to the price of ST. This is because the AO needs to choose one product, since we assume that no AO would ever choose to cancel the flight because of charges, and this choice is made against the reference point of ST charges. The attractiveness of ST has been normalized to 1 as a reference point.

The purchase probability for PT looks similar and assuming all three options are available we calculate it as:

$$Prob_{PT}(\{ST, PT, DT\}) = \frac{\exp(\beta_0^{PT} + \beta_p \frac{p_{PT}}{p_{ST}})}{1 + \exp(\beta_0^{DT} + \beta_p \frac{p_{DT}}{p_{ST}}) + \exp(\beta_0^{PT} + \beta_p \frac{p_{PT}}{p_{ST}})}$$

The purchase probability for ST is defined by

$$Prob_{ST} = 1 - Prob_{DT} - Prob_{PT}$$

Note that this probability distribution is by definition well-defined, i.e. all probabilities are greater than or equal to zero, and $Prob_{DT} + Prob_{PT} + Prob_{ST} = 1$.

Purchase probabilities can similarly be defined using the same parameters when other combinations of products are offered, namely {ST,PT} and {ST,DT}. If just one option is offered, it will be chosen with probability 1 regardless of the price (assuming that the price is reasonable).

The parameters β are pre-defined for certain AO segments. As before, we do not have any real data on such choices, hence we arbitrarily set the parameters in this model to reflect what we consider to be reasonable behaviour.

We define the segments in Table 1.

To better understand the assumed behaviour of these segments, assuming all three options are available, we visualized the choice probabilities in dependence of the price ratios in the following figures.

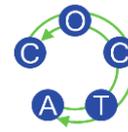
Table 1. AO segments: trajectory product choices

AO segment name	Description of segments	Consideration trajectories set	Parameters
-----------------	-------------------------	--------------------------------	------------

16 The opinions expressed herein reflect the author's view only. Under no circumstances shall the SESAR Joint Undertaking be responsible for any use that may be made of the information contained herein.

Founding Members





			$(\beta_o^{PT}, \beta_o^{DT}, \beta_p)$
<i>Premium</i>	Transatlantic flights and business aviation jet engine a/c. We also include Military flights to this group.	{PT}	Not needed – these AO will always choose PT
<i>Autonomy seekers</i>	AOs with preference for PT and aversion to assigned routes, yet more price sensitive than Premium. Flights include first and last rotation of Full Service Carrier flights, Business aviation with turbo-prop and piston engines.	{PT, DT, ST}	(90,35,-55)
<i>Discount seekers</i>	Very price sensitive AOs with resulting preference for DT provided enough discount is given: Low Cost Carriers, Charters and Cargo.	{ST, DT, PT}	(50,45,-50)
<i>Other</i>	All flights that do not fit into the other segments; preference for ST.	{ST, DT, PT}	(45,30,-40)

DT purchase probability
 $\beta_o^{PT} = 90, \beta_o^{DT} = 35, \beta_p = -55$

PT purchase probability
 $\beta_o^{PT} = 90, \beta_o^{DT} = 35, \beta_p = -55$

ST purchase probability
 $\beta_o^{PT} = 90, \beta_o^{DT} = 35, \beta_p = -55$

Founding Members



© – 2018 – COCTA consortium.

17

All rights reserved. Licensed to the SESAR Joint Undertaking under conditions.

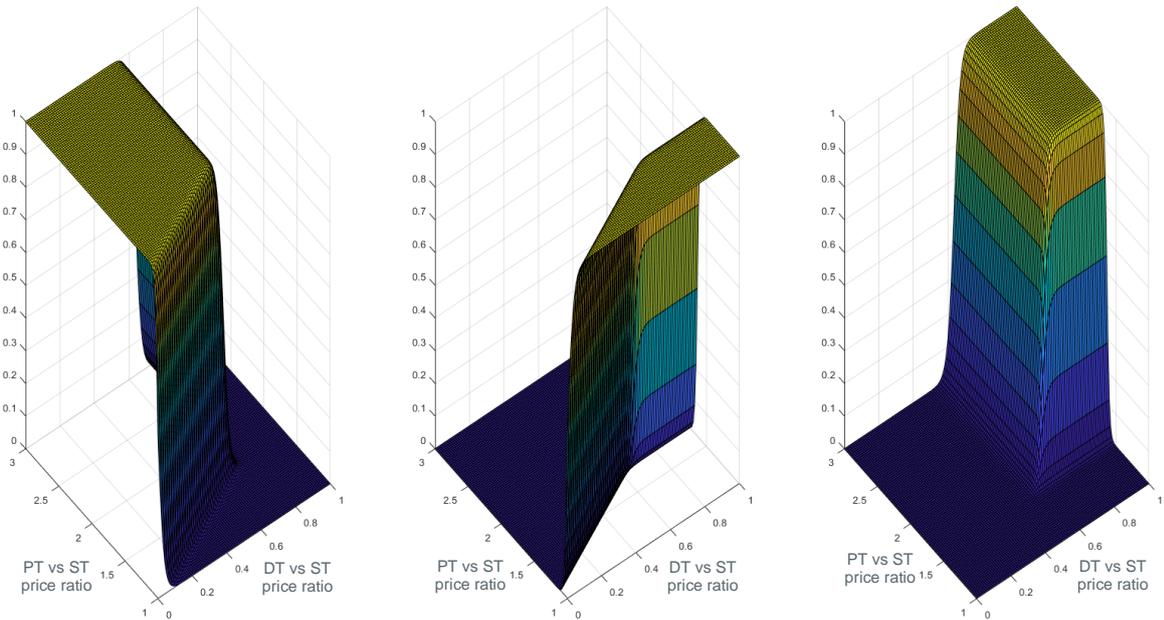


Figure 1. Choice behaviour of Autonomy Seekers

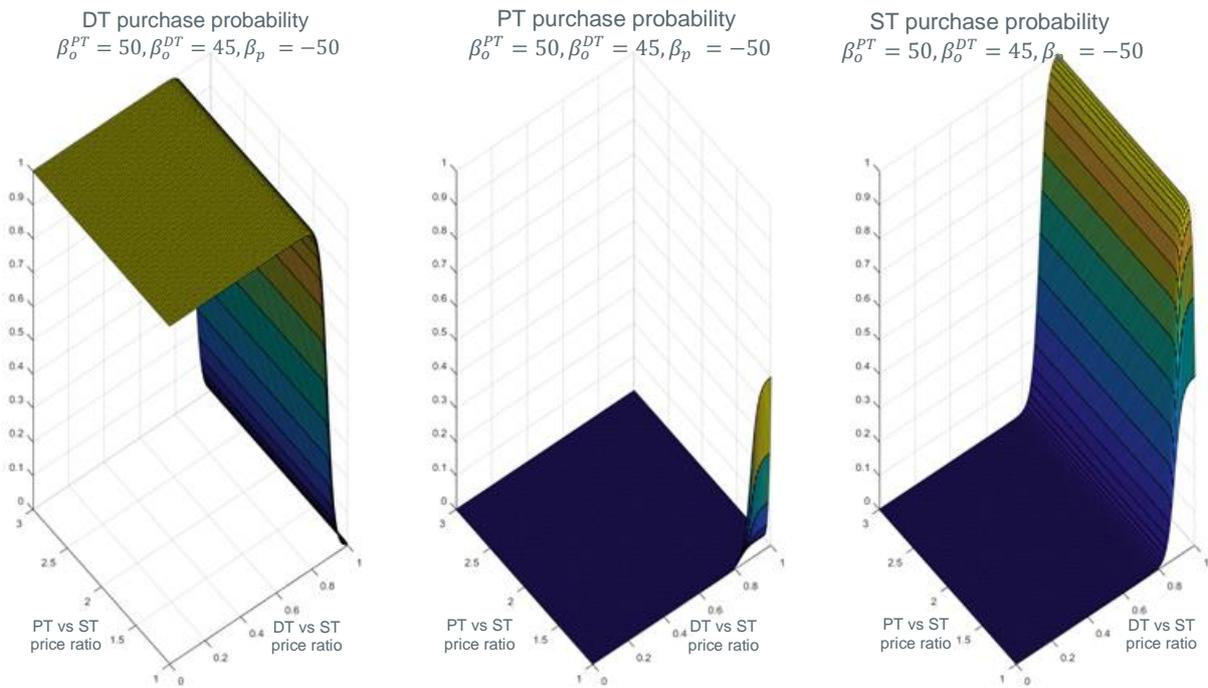


Figure 2. Choice behaviour of Discount seekers

DT purchase probability
 $\beta_o^{PT} = 45, \beta_o^{DT} = 30, \beta_p = -40$

PT purchase probability
 $\beta_o^{PT} = 45, \beta_o^{DT} = 30, \beta_p = -40$

ST purchase probability
 $\beta_o^{PT} = 45, \beta_o^{DT} = 30, \beta_p = -40$

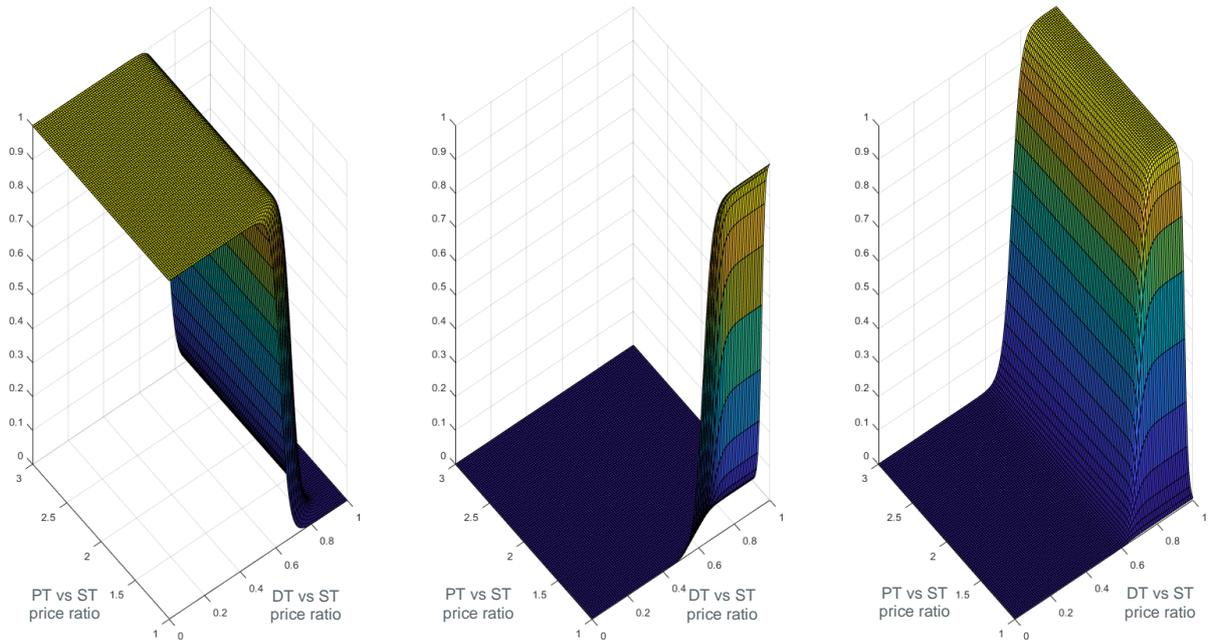


Figure 3. Choice behaviour of Other

4 Numerical Experiments Design

4.1 Strategic phase

At this stage, we want to obtain:

- A decision on how much capacity (in terms of sector hours) to order from each ACC. This decision is derived from the sector opening scheme (SOS) expressed by the z variable in the mathematical model. The sector hours needed (capacity) are captured with variable h_a .
- Base charge, i.e. price for ST product.
- Information on which flights will likely need to be displaced beyond ST boundaries and thus should be offered DT.
- The price of the PT product.
- The price of the DT product.

At this stage, we essentially attempt to build a forecast consisting of large number of traffic scenario instances, and then we find the trajectory and capacity-side decisions that would minimize overall cost (using model (1-6)) for each scenario. A scenario includes all scheduled flights plus a random sample from the assumed distribution of non-scheduled flights, resulting in different traffic scenarios using the same sampling mechanism as in D5.3.

Each traffic scenario then consists of all scheduled flights plus the random sample of non-scheduled ones. We feed them into model (1-6) without any restrictions on the range of available trajectories that we can choose from; at this stage, we neither explicitly model any of the products, nor choice behaviour. Also, there is no overall capacity limitation except those given by the pre-defined potential airspace configurations. In other words, the K_p is fixed for every p , but we are free to choose any configuration out of \mathcal{C}^a .

In summary, the strategic phase experiment is similar to those in D5.3, except that now we have some segments of AOs which may choose the PT product. We assume that the NM offers AOs trajectory products as presented in Table 1. AOs choose from these options using the choice models defined above. We fix all of the non-scheduled flights in a traffic scenario to a sampled trajectory and exclude them from the NM's decisions on delay and re-routing. Instead, the NM assesses the impact of this additional source of uncertainty to evaluate how much (more) capacity is needed to (cost-efficiently) manage demand. Output of this phase is a vector h of capacity budgets for all air spaces.

The perfect foresight approach gives us the ordering cost $\gamma^T h$. This cost we aim to recover using ATM charges. Assuming that only ST are sold, we thus can set the ST price as $p^{ST} = \sum_{F \in \mathcal{F}} [(\gamma^T h) / |F|] / |\mathcal{F}|$,

20 The opinions expressed herein reflect the author's view only. Under no circumstances shall the SESAR Joint Undertaking be responsible for any use that may be made of the information contained herein.

Founding Members



EUROPEAN UNION



EUROCONTROL



where \mathcal{F} denotes the finite collection of flight scenarios in our forecast. This will serve as a reference price to define the prices of DT and PT. Since ST prices are set so as to cover the expected capacity cost C , total DT discounts should roughly equal total PT surcharges (all relative to ST). For the sake of simplicity, we ignore take-off weight in this analysis. Furthermore, we assume that trajectory charges do not change over time – therefore, we can solve this problem in a static fashion as opposed to simulate the booking horizon (and thus having to deal with re-optimizations based on realized demand).

Since AOs who choose the PT product are assumed to be less price sensitive, we assuming a certain willingness to pay (WTP) for them, and then *setting the PT price equal to this max WTP* (e.g. a pre-determined fixed amount that is independent of route and of OD pair, and which is relatively high relative to what a comparable scheduled flight of similar weight would pay). We express the PT price as a-priori fixed markup over the ST price: $p_f^{PT} = (1 + m^+)p_f^{ST}$. In our experiments, we set the price of PT to be 20% higher than the corresponding ST price.

This determines the expected budget of markups from non-scheduled traffic that we can spend on discounting selected scheduled traffic: let N_F^{ns} be the number of non-scheduled flights in scenario F . Then the budget is $B := \sum_F m^+ p^{ST} N_F^{ns} / |\mathcal{F}|$; this is what we can use for DT discounts.

We only want to offer DT products to flights that we think we likely will need to move. To determine which flights these are, we estimate the probability δ_f of a scheduled(!) flight f in F being displaced outside ST margins by defining $\delta_f^F = 1$ if flight f was assigned to route outside ST margins in solution for scenario F , then calculating $\delta_f = \sum_{F \in \mathcal{F}} \delta_f^F / |\mathcal{F}|$. Only scheduled flights with this probability exceeding a certain threshold (in our experiments we set the threshold to 60%) will be offered DT. Denote this set of scheduled flights Δ .

DT discounts are no longer coupled to their associated cost savings because the cost benchmark from the strategic phase already incorporates flights to be displaced where and when necessary. Instead, we want to maximize the probability of AOs accepting DT offers for flight/flow that are likely to be displaced. We are only limited by the overall surcharge budget B that we re-distribute towards DT.

On the one hand, we want to discount the price of DT as much as possible to maximize the chance that an AO actually chooses DT over alternative products. On the other, we cannot offer too much discount over the ST base charge because this would otherwise leave us with insufficient funds to cover capacity costs. To address this trade-off, we solve the problem below for the best markdown value m^- . We assume that the markdown can be at most 80% relative to the price of ST. The markdown needs to set so as to reduce the price as much as possible for all flights that we likely need to move, i.e. all f in Δ , with a higher weight on those flights that are more likely in need to be moved. We have budget B to spend on awarding discounts, hence the markdown m^- is set such that:

$$\begin{aligned} & \max \sum_{f \in \Delta} \delta_f \text{Prob}_f^{DT}(m^-) \\ & s. t. \quad \sum_{f \in \Delta} \text{Prob}_f^{DT}(m^-) m^- p^{ST} \leq B \\ & \quad \quad 0 \leq m^- \leq 0.8 \end{aligned}$$

This concludes the strategic phase (phase 1 in the simulation). We have now the capacity budget and all prices.

4.2 Pre-tactical phase

For a given capacity ordering decision, the output of the strategic phase is now considered fixed.

In the pre-tactical phase (stage 2 in the simulation), we sample the traffic materialization \hat{F} from the true traffic distribution. For each scheduled flight in that materialization, we sample the AO's choice under the choice model as defined above given the prices. This defines the set of routes R_f for all flights. All non-scheduled flights are assumed to choose PT only.

The resulting materialization of flights and product choices is fed into our optimization model (1-6) to determine the demand management measures, namely the sector opening scheme, re-routing and delay decisions (within the given margins prescribed by the set of trajectory options R_f). In this optimization, we add a dummy route for every flight with zero capacity requirements in all sectors but high penalty costs in the objective to ensure feasibility of a solution.

In the next section, we discuss the numerical results obtained by applying this approach to a medium-sized case study involving over a thousand flights.

5 Results and discussion

5.1 Case study

We test this approach on a case study involving Central European airspaces and over a thousand flights during a one-hour time period. Our main objective is to quantify the impact of allowing AOs to choose their own trajectories (i.e., the PT product). It is intuitive that this will have detrimental effects on overall cost performance, but it is unclear how severe these effects will be, and to what extent we can remedy the situation by targeted selling of DT products.

In this deliverable, we use the same case study as in our working paper Starita et al. (2018) – currently under review at *Transportation Science* – which is attached in the appendix. The main difference between the experiments in that paper and this deliverable is that the paper does neither feature the different trajectory product types, nor does it involve pricing decisions. Its focus is solely on the optimization procedure to make capacity budget decisions in the face of uncertainty over the materialization of non-scheduled flights.



Figure 4: Snapshot of scheduled flights and airspaces in the simulation at a fixed point in time

The selected region includes a large part of en-route airspace in Central and Western Europe, including 8 ANSPs and 15 ACCs/sector groups¹ as depicted in Figure 4. The figure shows a snapshot of scheduled traffic at a particular point in time and illustrates the presence of different traffic densities. Based on historical usage of configurations in 2016, for each ACC or sector group, we select configurations with different numbers of sectors which were most frequently used: in total, we have 173 different configurations for the 15 ACCs/sector groups.

Delay and re-routing costs are obtained by drawing on (Cook and Tanner, 2015) and (EUROCONTROL, 2018). Delay costs are calculated per aircraft type and per duration of delay (non-linear with time). Re-routing costs include fuel costs, crew costs, passenger soft and passenger hard costs, as well as maintenance costs. Cost parameters for ANS provision are derived in the following way: based on information provided in the ATM Cost-Effectiveness Report 2017 (EUROCONTROL, 2017), we calculated average ATCO costs per sector-hour in 2015 for the different ANSPs (and in one case for an ACC). In the model, we treat those costs as variable costs. Although ATCO costs might be considered fixed costs in the short term, reducing the daily number of ATCO hours will reduce the total number of ATCOs needed for ANS provision, consequently reducing staff costs.

We use flight data of 9th September 2016, which was the busiest day in 2016 in European airspace. We select flights based on their last filed flight plans, which cross the selected airspace between 10 and 11 AM. In total, we have 910 scheduled flights that are considered fixed in our network in all scenarios. Out of all non-scheduled flights on that day, we select those that cross the selected airspace at any time (1,569 in total). Since we test the model for 10-11 AM period, we change their airspace entry times from the original flight plan to a time uniformly sampled over the selected period. Each traffic scenario is created by uniformly sampling a subset of 160 from this set of 1,569 flights and adding them to the set of scheduled flights. We create 100 traffic scenarios in this manner. Flights can be either delayed or re-routed (only one demand management measure per flight) to improve total cost-efficiency subject to hard capacity budget constraints. Delay options are discrete and the same for each flight, namely 5, 15, 30 or 45 minutes. Each flight has a number of alternative spatial routes as well, all generated using the NEST tool. Overall, the problem is modeled over a two-hours time horizon to account for flights being delayed beyond 11:00 AM. We chose this data to test our approach since the selected airspace covers a large part of the so-called 'core area' (e.g. MUAC and DFS) of the European airspace, as well as a part of airspace which is not as congested (Hungaro Control).

5.2 Results

We test three different capacity decision policies that are defined in full detail in the attached working paper Starita et al. (2018):

- *AV*: in the averaging policy, the capacity decision is obtained by averaging the capacity order decision h_a^F for airspace a and scenario F (that results from the foresight approach) across all scenarios.

¹ For instance, the Maastricht Upper Area Control Centre (MUAC) is divided into three sector groups: Deco, Hannover and Brussels, each with its own sectorization and sector configurations.

- *Eps-5 / Eps-20*: In the risk-based policy, the capacity decision is obtained by setting h such that the sample probability of encountering a flight scenario in which we had better planned for more capacity in at least one airspace is less than a given epsilon (set to either 5% or 20%), where the sample probability distribution of h has been computed by the perfect foresight approach.

Each simulation consists of two stages, corresponding to the strategic and pre-tactical stages described in Section 4. In each simulation run, we start in stage 1 with obtaining the capacity budget h using a given decision policy as well as prices for DT, ST and PT. In stage 2, we sample the actual traffic materialization and trajectory product choices of the AOs. This serves as an input to the optimization of demand management decisions, as well as the sector opening scheme subject to the available capacity budgets. We repeat the simulation 200 times and report average results.

To assess the effect of granting AOs who purchase the PT product the permission to decide themselves on their trajectory, we consider two scenarios. In the first, we assume that all flights that chose PT have a random trajectory assigned to them that is not under the influence of the NM. In the second scenario, we assume that the NM can still assign PT flights to all routes incorporated in the range of route options for ST. The latter scenario could represent the case of asking late-arriving trajectory requests to pay a surcharge whilst not granting them additional benefits over ST. In our simulation, all non-scheduled flights (that often would not be able to book a trajectory in advance) and only on average about 2% of scheduled flights select PT.

In Table 2, we see that the total cost (capacity cost plus displacement cost) is minimized for the most risk-averse policy *Eps-20*. Moreover, this policy also always procured sufficient capacity to accommodate all 1,070 flights. The displacement cost is three to four times the capacity cost depending on our decision policy, which suggests that substantial demand management measures are being applied to serve all flights within the given capacity limits (further discussion is presented in the next section).

Table 2: Average costs when PT flights may choose their own trajectory

Policy	Capacity cost	Displacement cost	Total cost	Avg (#non-assigned flights)	#flights
AV	51,854	205,447	257,301	0.07	1,070
Eps-20	57,808	196,783	254,591	0.00	1,070
Eps-5	61,877	196,760	258,637	0	1,070

Now we compare this to the situation where the NM retains decision power over flight trajectory assignments (even for PT) in Table 3.

Table 3: Average costs when PT flights are assigned a trajectory option like for ST

Policy	Capacity cost	Displacement cost	Total cost	Avg (#non-assigned flights)	#flights
AV	51,854	6,544	58,398	0.02	1,070
Eps-20	57,808	661	58,469	0.00	1,070
Eps-5	61,877	627	62,504	0.00	1,070

A huge difference can be observed: displacement costs are nearly non-existent for the risk-based policies! The capacity costs are exactly the same in both scenarios since they are not affected by our differing assumptions regarding PT. Therefore, the ability to retain the power to assign flights to trajectories from the ST range of options even for PT has a huge effect on displacement cost reduction. This is in line with our earlier findings in D5.3 that the COCTA mechanism has the potential to greatly reduce demand management-related costs.

The risk-based policy Eps-20 again performs best in that it produces the lowest cost (on par with AV) and accommodates all traffic within its ordered capacity. The discount under the given choice models was calculated in all scenarios to be 25% relative to the ST charge.

Cost recovery is roughly achieved: as reported in Table 4, the AV policy generates income slightly in excess of the capacity costs incurred, whereas the risk-based policies Eps-5 and Eps-20 both generate slightly less. Overall, costs and income are within a few percentage points of each other.

Table 4: Income generated when PT flights are being assigned a trajectory from ST set

Policy	Total income	#DT	#ST	#PT	Income from DT	Income from ST	Income from PT
AV	53,049	444	444	182	17,763	23,653	11,634
Eps-20	55,856	444	444	182	18,702	24,904	12,249
Eps-5	58,881	444	444	182	19,716	26,253	12,913

The situation is nearly the same when PT flights are allowed to choose their own trajectory as shown in Table 5.

Table 5: Income generated when PT flights can choose their own trajectory

Policy	Total income	#DT	#ST	#PT	Income from DT	Income from ST	Income from PT
AV	53,044	445	443	182	17,783	23,619	11,642
Eps-20	55,850	445	443	182	18,724	24,868	12,258
Eps-5	58,876	445	443	182	19,738	26,215	12,922

5.3 Discussion

The main motivation of investigating the introduction of a premium trajectory product was that AO representatives at one of our stakeholder workshops suggested that this would be an attractive option for them which would help to increase the attractiveness of the COCTA concept as a whole to AOs. Intuitively, granting AOs more decision power is expected to lead to overall higher costs. In this study, we quantify this cost increase on the example of a case study using traffic data from Central and Western Europe.

We have worked on a relatively small case study because of time constraints; whilst the new solution approach that we propose in the attached paper Starita et al. (2018) is scalable (as demonstrated in that paper), running it over hundreds of simulations is still time-consuming. Nevertheless, even on



the scale of about a thousand flights we can clearly observe the effect of allowing greater flexibility to airspace users; in our context, allowing flights to choose their own trajectory with a premium trajectory product. To repeat for the purpose of clarification: in the COCTA context, all AOs choose their trajectory products in line with their business and operational needs, with associated margins for demand management and trajectory charge. The difference between the premium and other two trajectory products is in the final trajectory decision; AOs decide for the former and the NM for the latter two trajectory products (subject to a negotiation process with AOs).

The most notable effect is on displacement costs, which includes both delays and re-routings. Namely, allowing AOs to have additional flexibility to choose trajectories, which are “non-optimal” from the system’s perspective, but are “optimal” from the users’ perspective, leaves limited options for the NM to distribute (in space in time) other flights in the network. In other words, with capacity decisions fixed and without an option to (re)adjust it, offering premium trajectory (PT) products can be expected to lead to substantially higher displacement costs for other AOs. One should also note that the case study used to test the concept of premium trajectory has an extremely challenging demand profile for the capacity available. In a less challenging demand vs capacity circumstance, the displacement cost might be much lower (yet to be tested). Moreover, the additional displacement cost decreases with a declining share of AOs choosing the PT product. This share, however, can be influenced by the NM through the trajectory pricing mechanism. In other words, if the PT product becomes more expensive, fewer airlines will purchase it and therefore displacement costs for the entire system will be reduced.

To avoid very high displacement cost being imposed on airlines, there are several measures which could be taken:

- 1) *Ordering more capacity at the strategic stage.* In case where the COCTA concept/system has been operational for some time, the NM could account for potential (likely) PT purchases at the strategic level and order more capacity based on historical trajectory product choices. By doing so, the NM would ask for more capacity to be provided by ANSPs to ensure a (more) cost efficient operation. In our case study, we could use our choice model and assumptions to simulate potential future PT choices at strategic level and make capacity decisions based on many simulations.

Note also that the capacity planning is carried out assuming that nominal conditions apply, e.g. that capacity is delivered as planned, which is not always the case. Additional capacity, further increasing capacity cost, would likely be needed to account for potential disruptions which will occur on the day of operations.

- 2) *Re-adjusting capacity at pre-tactical or even tactical level.* In the COCTA experiments so far, we have experimented with different degrees of flexibility of capacity provision. In one case, we assumed that the NM asks for “capacity budget” at strategic level, i.e. total number of sector hours which each ANSP should deliver for a day of operations (and in a longer run). At pre-tactical level, with the updated information on demand, the NM decides on sector-opening scheme. In another case which we tested, the NM has to ask for a sector-opening scheme much more in advance (at strategic level), so that ANSPs could decide on their rostering timely, without the possibility to adjust this capacity level at pre-tactical stage, which is more realistic given the current ANSP practices. Although the capacity provision in this deliverable is already assumed to be *flexible*, we hypothesize that an ATM system with even more flexible capacity provision would further improve the overall cost efficiency.

Mobility of airspace and/or ATCOs, potentially cross-border, should unlock more potential from the COCTA concept, as the “spare/unused” capacity can be shifted in the network where and when needed in a coordinated (network-centric) manner.

Also, one should note that the NM estimates cost of delays and re-routings – “true” cost can only be estimated by AOs. AOs who opted for PT, either to prevent high potential cost of delay (business) or to save fuel and flying time (transatlantic flights), will have benefits that we cannot credibly estimate. This solution where airlines can put “more weight” to certain flights is also supported by the User Driven Prioritization Process (UDPP) initiative. Also, a recent trial by Eurocontrol and SWISS airlines, in which SWISS could prioritise two (delayed) flights over others from the same airline in the first morning rotations (without affecting other airlines) showed some promising results².

In the end, these results should be observed in the light of the study design and information/data available. Namely, since the trajectory products still don’t exist, and there is no historical purchase/AO decisions data available to estimate an accurate choice model, we apply a fairly simple approach of AO decision making process (choosing trajectory products). A different choice model might have yielded somewhat different results. Therefore, in a future study one should carry out further analysis regarding trajectory choices and trajectory products definitions (margins and prices).

In the following, we list a few future research directions:

- It would be very useful to conduct a study of AOs’ preferences regarding our trajectory products (e.g. using conjoint analysis). These results could be used to inform a choice model that will help to gain a more realistic view of AOs’ choice behaviour. This should be fed into a more comprehensive simulation study to quantify the effects on the ATM system at a larger scale. Pricing could similarly be more realistic if one had information regarding maximum willingness to pay for the different trajectory products. This type of insight could be elicited in the same study with a sufficiently representative stakeholder group.
- The robustness of the COCTA approach should be tested against disruptions caused by weather or other factors. In our study, we assumed no disruptions; however, our modelling approach as described in the attached paper could be used to include such effects as well (we just would add corresponding disrupted scenarios to our set of scenarios).
- Multi-stage capacity ordering may further decrease costs; we have only considered a single ordering decision in our experiments.
- The costs effects should be simulated over an extended period of time.
- Flexible capacity provision such as cross-border provision should be linked to the COCTA concept.

² More details: <https://www.eurocontrol.int/news/making-connection-first-rotation-hours-optimisation-trial>



6 Conclusions

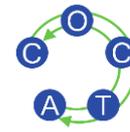
We have presented the concept underpinning a new premium trajectory product and a new modelling approach that should be able to better handle uncertainty due to its much increased ability to work with large numbers of scenarios. A simple approach to modelling choice and setting prices for the trajectory products has been proposed, as well as a procedure for testing the COCTA concept under these new condition in a simulation study.

A major advantage of our proposed optimization method is that it works with any sample of traffic scenarios; in particular, we could use a version of it to test robustness of capacity ordering policies to disruptions in the availability of sectors (e.g. due to weather or military). We do not require a closed-form expression of the distribution of non-scheduled flights or other sources of uncertainty. The efficient formulation allows us to evaluate a large number of scenarios, which improves our ability to make good capacity (and pricing) decisions in advance of demand management decisions.

The numerical experiments indicate that even a moderate number of non-scheduled flights (around 15% of overall traffic) may lead to displacement costs of three to four times the capacity costs if these flights are allowed to choose their own trajectory. On the other hand, if the NM were to retain the decision power regarding trajectory assignment, displacement costs are negligible under our risk-based capacity ordering policy. This type of trajectory product was proposed by practitioners during the first of our stakeholder workshops. Based on our limited evaluation of the new trajectory product, we can conclude that within the context of the experiment, allowing AOs to prioritise some of their flights in the network, assuming given-unchanged initial capacity order, reduces cost-efficiency. However, unlike with the other two trajectory products, the COCTA concept with additional trajectory product has not yet been sufficiently evaluated to come up with any firm conclusions. This primarily includes the currently missing impact (feedback loop in the model testing process) on the capacity ordering decision of offering a new (“premium”) trajectory product. The incorporation of PT product earlier in the COCTA process and an analysis of effects on network performance in such a setting is one of the immediate areas for further investigation.

References

- [1] COCTA consortium, 2018. Choice-based joint capacity ordering and pricing model.
- [2] COCTA consortium, 2017a. Final mechanism design.
- [3] COCTA consortium, 2017b. ATM value-chain redesign.
- [4] Cook, A., Tanner, G., 2015. European airline delay cost reference values. Final Report. University of Westminster for EUROCONTROL Performance Review Unit.
- [5] EUROCONTROL, 2018. Standard Inputs for EUROCONTROL Cost-Benefit Analyses, Ed. 8.0.
- [6] EUROCONTROL, 2017. ATM Cost-Effectiveness (ACE) 2015 Benchmarking Report with 2016-2020 outlook.
- [7] Tobaruela, G., Majumdar, A., Ochieng, W.Y., Schuster, W., Hendrickx, P., 2013. Enhancing cost-efficiency and reducing capacity shortages: strategic planning and dynamic shift management., in: Tenth USA/Europe ATM R&D Seminar.



Appendix A A new solution approach – working paper

Capacity Planning in Air Traffic Management under Traffic Uncertainty

Stefano Starita¹, Arne Strauss², Xin Fei², Radosav Jovanović³, Nikola Ivanov³, Goran Pavlović³, and Frank Fichert⁴

¹Sirindhorn International Institute of Technology, Thammasat University,
Pathum Thani 12120, Thailand

²Warwick Business School, University of Warwick, Coventry CV4 7AL,
United Kingdom

³Faculty of Transport and Traffic Engineering, University of Belgrade,
11000 Belgrade, Serbia

⁴Faculty of Tourism and Transport, Worms University of Applied
Sciences, 67549 Worms, Germany

31 October 2018

Abstract

In air traffic management, a fundamental decision with large cost implications is the planning of future capacity provision. Here, capacity refers to the available man-hours of air traffic controllers to monitor traffic. Airspace can be partitioned in various ways into a collection of sectors, and each sector has a fixed maximum number of flights that may enter a sector within a given time period. Each sector also requires a fixed number of man-hours to be operated; we refer to them as sector-hours.

Capacity planning usually takes place a long time ahead of the day of operation to ensure that sufficiently many air traffic controllers are available to manage the flow of aircraft. However, at the time of planning there is considerable uncertainty regarding non-scheduled air traffic. Once the capacity decision has been made (in terms of committing to a budget of sector-hours per airspace to represent long-term staff scheduling), on the day of operation we can influence traffic by enforcing re-routing and tactical delays. Furthermore, we can modify which sectors to open at what time (the so-called sector opening scheme) subject to the fixed capacity budgets in each airspace. The fundamental trade off is between reducing the capacity provision cost at the expense of potentially increasing displacement cost arising from re-routing or delays.

1

To tackle this, we propose a scalable decomposition approach that exploits the structure of the problem and can take traffic uncertainty into account by working with a large number of traffic scenarios. We propose several decision policies based on the resulting pool of solutions and test them numerically using real-world data.

1 Introduction

Air navigation service providers (ANSPs) are among the key stakeholders in providing safe and efficient air traffic operations. In Europe, there are nearly 40 agencies providing air navigation services to flights operated by airspace users within the boundaries of their respective (national) airspace. To a certain extent, the air navigation service provision is coordinated on a network level by the network manager [European Commission, 2011]. The European en-route airspace is further fragmented into about 70 Area Control Centres (ACC), which are operational units of ANSPs. Airspace of an ACC is divided into sectors, where each sector represents a volume of airspace managed by one or more air traffic controllers (ATCOs). In total, the European airspace is divided into more than 600 of so-called 'elementary' sectors, which can be assembled into larger units called 'collapsed' sectors. Any specific combination of elementary and/or collapsed sectors which cover the entire volume of airspace for an ACC is called sector configuration [Baumgartner, 2007].

ATCOs working in a single sector can safely provide air navigation services to a limited number of aircraft in a given period of time, and this number represents sector capacity [Baumgartner, 2007]. Usually, sector capacity is defined either as a maximum number of aircraft that can enter a sector within a period (so-called 'entry counts') or as a maximum number of aircraft that can be present in a sector within a period ('occupancy counts'). Based on anticipated traffic in their airspace during a day of operations, ANSPs plan how many ATCOs need to be on their positions to safely manage flights, that is, how many sectors should be opened during the course of day. As a rule of thumb, the more sectors are open, for a longer period, the more capacity an ACC can provide [Eurocontrol, 2018a]. One measure of how much aggregated capacity an ACC has provided during the day is the number of sector-hours, i.e. how long sectors were open/active during the day.

Inadequate staffing (i.e. capacity planning decisions) may have significant cost implications. Providing insufficient capacity in certain airspace parts at certain times results in delays being imposed on airspace users and/or re-routing from shortest routes. The costs of en-route imposed delays in the European system were estimated to exceed EUR 900 million in 2017, with more than EUR 550 million associated with lack of air traffic control capacity and staffing [Eurocontrol, 2018a]. Re-routings likewise are undesirable since they result in additional fuel burn and other operating costs for airspace users, as well as in increased CO₂ emissions [Eurocontrol, 2018a].

On the other hand, an overprovision of capacity also comes at a cost imposed on airspace users, owing to cost-recovery elements built into the European route charging scheme. User charges thus reflect the costs of capacity provision, which in Europe amounted to more than EUR 8 billion in 2016 [Eurocontrol, 2018a], averaging about EUR 520 per flight-hour



controlled. Navigation charges (en-route plus terminal) constitute on average about 4.2% of total costs of legacy carriers [IATA, 2015], but this percentage can be as high as 12% for some low-fare carriers [RyanAir, 2018].

Capacity planning for a day of operations starts months in advance [Tobaruela, 2013]. One of the major challenges in this planning process is that there are several sources of uncertainty that often are only revealed shortly prior to or even on the day of operation. Tobaruela [2013] state that these include uncertainty related to the overall number of flights and their spatio-temporal distribution, the impact of adverse weather conditions or the unavailability of airspace due to military use restrictions (among others). They also stress that weather conditions are important, however, they are difficult to incorporate when planning capacity months in advance. The exact number of non-scheduled flights as well as where and when they fly is often only revealed on the day of departure or shortly before that. Therefore, it is important to account for the uncertainty stemming from non-scheduled traffic in capacity planning. In this work, we focus on this type of uncertainty in isolation (i.e. we assume no other sources of uncertainty such as weather, although our approach can also be extended to address these as we outline in the conclusions).

The aim of our work is to provide a methodology for planning capacity provision in the presence of uncertain non-scheduled flights. We assume that capacity budgets (in terms of sector-hours) for each airspace must be acquired *before* the uncertainty regarding non-scheduled flights is revealed. Once the latter information is available, we can only resort to demand management measures, namely re-routing or delaying flights, and to re-arranging the sector opening scheme within the given capacity budgets of each airspace, to ensure that all sectors operate within their capacity limits. This is fairly close to practice in that the full information on non-scheduled flights only becomes available on the day of operation, and the number of available air traffic controllers (and thus the capacity budgets) is essentially fixed at this point.

Our main contribution is to provide an efficient solution approach to this problem, which we model as a somewhat stylized two-stage newsvendor problem with the added difficulty of a cost function that is hard to evaluate. We propose various policies on making capacity decisions in stage 1 that attempt to anticipate the uncertain materializations of non-scheduled traffic in stage 2. The policies are based on solving a version of the problem that assumes full information on all flights is deterministic and known in stage 1; this version of the problem is then solved for many random scenarios using a scalable decomposition approach that separates decisions on re-routing and delay from decisions on airspace configurations. A particular advantage of this approach is that we can apply it to any distribution assumed to underpin the materialization of non-scheduled flights; in fact, even to distribution-free approaches, as long as we can somehow obtain a set of traffic scenarios. We demonstrate the effectiveness of this approach in a numerical study based on real flight data.

The paper is organized as follows: we review related literature in §2 and state the problem in §3. We discuss the solution approach in §4. The approach is tested numerically on data modeled from Central European airspaces and traffic patterns in §5. We draw conclusions in §6.

2 Literature Review

One of the primary research areas in airspace capacity (planning) research is determining sector capacity based on ATCO's workload, as this is a highly relevant practical issue [Majumdar et al., 2005]. Sector capacity (or ATCOs' workload) is a critical decision involved in the airspace capacity planning process and has a large impact on demand management measures (delays, re-routings, cruise speed control, etc) implemented when traffic exceeds capacity. As a consequence, mathematical methods are often introduced to support this decision.

For instance, Sherali et al. [2003] define a mixed-integer programming model to decide on flights' trajectories from a set of alternatives, subject to air traffic control workload, flight safety and airline equity constraints. The framework includes a probabilistic trajectory conflict analysis, the development of air traffic control workload metrics and the consideration of equity among airline carriers in absorbing costs related to demand management measures. In a following paper, Sherali et al. [2006] investigate how their model can be used in both tactical and strategic phases. Regarding strategic applications, authors propose air-traffic control policy evaluations (e.g., revising aircraft separation standards or analyzing the free-flight paradigm where aircraft are permitted to take wind-optimized routes), fixed alternative or dynamic airspace re-sectorization strategies and the construction of a priori plans to respond to various disruption scenarios, among others.

Another well-researched aspect of airspace capacity provision is dynamic airspace management (sectorization), i.e. frequently changing sector boundaries or re-organizing airways to adapt to traffic flows and optimize a criterion function; see Tien and Hoffman [2009], Flener and Pearson [2013], Venugopalan et al. [2018]. However, various limitations to dynamic (sector) management exist as highlighted by Delgado et al. [2015]; for instance, the fact that ATCOs should have a two hours' break per shift. Some ANSPs implement a strategic planning process approach with progressive refinement as more information becomes available, as opposed to a conservative (rigid) staff planning as described by Tobaruela [2013].

Strategic capacity planning on a network level in the European context has not received much attention in the literature, in particular on the cost-efficiency thereof. Tobaruela [2013] investigate how dynamic sector-opening times and a layered planning process can improve the center's cost-efficiency. While the results indicate that the dynamic sector opening improves the ability of matching demand with the available resources, the authors do not account for traffic variability in their research. In a similar research problem, Josefsson et al. [2017] account for a portion of non-scheduled flights in planning which airports will be controlled from remote positions.

In contrast to previous research efforts, we propose strategic capacity planning at the network level. The aim is to make capacity ordering decisions, while minimizing the cost of capacity provision and cost of delaying and re-routing flights. Our approach is similar to the current practice [Eurocontrol, 2013], with two fundamental differences:

- capacity decisions are made by a central planner for all ANSPs in a fully coordinated



manner, taking into account the network as a whole (as opposed to the limited coordination in the current system), and

- the central planner anticipates the impact of re-routing and tactical delay decisions as a demand management measure at the time of making capacity decisions, rather than using re-routing merely to alleviate delays on the day of operations.

Our work is also related to papers dealing with the problem of assigning delays and of re-routing flights under capacity constraints of airports and/or sectors. Tošić et al. [1995] propose a linear model to minimize delay's costs while deciding the assignments of both ground holds and routes. Same problem is tackled by Bertsimas and Patterson [1998]. They also focus on the complexity of the problem and provide practical extensions by considering dependence between arrival and departure capacities, hub and spoke systems, banks of flights and rerouting. Lulli and Odoni [2007] apply the problem to the European airspace, considering issues such as efficiency and equity. Corolli et al. [2010] introduce to the problem the concept of time windows as slots in which flights execute their actions. They propose a two stage model which firstly identifies the set of optimal (i.e., minimum cost) windows' sizes and subsequently selects the optimal solution which grants the largest degree of flexibility to service providers and airports. Castelli et al. [2011] study the single constrained en-route sector problem incorporating the possibility for an airline to pay to reduce delays or obtain delay compensation. Lau et al. [2015] introduce a column generation based algorithm to tackle large-scale network slot allocation. In our work, we also incorporate decisions on assigning delay and re-routing, as well as decisions on the sector opening scheme. However, we consider this in the context of strategic capacity decisions, rather than making tactical demand management decisions. We do not consider the impact of airport congestion.

3 Problem Statement

We consider a central network manager with the mandate to make capacity and demand management (namely re-routing and delaying) decisions across various airspaces. The problem is posed as a somewhat stylized process over two stages: In the first stage, the network manager needs to plan how many sector-hours will be required for each airspace for a specific day in the future (in practice, this corresponds to the strategic planning phase). While we have information on all scheduled flights, non-scheduled flights are unknown at this point. In the second stage, uncertainty regarding non-scheduled flights is revealed (this corresponds to the day of operation). In the light of this information, the network manager needs to decide on re-routing or delaying flights, and on the sector opening scheme subject to the fixed capacity budget available to accommodate all flights. The objective is to minimize the sum of capacity provision cost and expected displacement cost stemming from demand management measures; note that both re-routing and delay incurs displacement costs. Structurally, the problem is related to the well-known newsvendor problem. In the following, we offer a rigorous definition of this problem.

Consider multiple airspaces $a \in A$, each with a finite number of possible sector configurations $c \in C^a$. For a given configuration c , we have a set of sectors $p \in P^c$ that form the elements of the configuration. Each of these sectors p is either an elementary sector or consists of multiple elementary sectors merged together (referred to as a collapsed sector). Each sector p has a fixed capacity of K_p flights that may enter that sector within a given time period $u \in U$. The time periods in U span the day of operation on a uniform grid with spacing chosen such that it is possible to change the configuration of an airspace from one time period to the next (say, 1 hour). Opening configuration c in airspace a for one time period requires \bar{h}_{ac} sector-hours.

Non-scheduled flights may appear in period 2, but they are unknown in period 1. We assume that we do not know the true distribution that governs the materialization of non-scheduled flights, however, we do have a uniform distribution over a finite collection of flight scenarios which we assume to approximate the true distribution. This collection is known in period 1; for instance, this could be defined as the collection of historic non-scheduled flights. We augment this collection by adding a fixed and known set of scheduled flights to every scenario of non-scheduled flights; the resulting collection of flight scenarios is denoted by \mathcal{F} . In other words, every element $F \in \mathcal{F}$ is a set of flights, and the elements only differ by the non-scheduled flights; the scheduled flights are the same in each of them. Note that traffic scenario F relates only to the number of flights, their origins and destinations, planned departure time, and a set of trajectory options R_f for every flight f that represent different demand management measures of re-routing or delaying flight f , including the option for the shortest route without delay. The traffic scenario does not determine the trajectories. In the following, we write the expectation over scenarios F in the understanding that uncertainty only pertains to non-scheduled flights. For a given flight f , route $r \in R_f$, time period $u \in U$ and a sector p , we define $b_{frpu} \in \{0, 1\}$ to be equal to 1 if route r uses sector p at time u , and otherwise 0. Each of these route options $r \in R_f$ comes with an associated displacement cost d_{fr} that reflects the additional fuel cost and delay costs incurred relative to the shortest distance at no delay (d_{fr} for the latter is set to zero). As such, we incorporate not only cost to the air navigation service provider, but also costs to airspace users. The aim is to reduce overall costs.

We face a trade-off between achieving cost savings by decreasing capacity provision cost in stage 1, and increasing costs by potentially increased need for demand management measures depending on the realization of non-scheduled flights. In stage 1, we need to decide on how much capacity budget $\mathbf{h} = (h_a)_{a \in A}$ in terms of sector-hours to acquire for the different airspaces (at unit cost γ_a for each airspace a). In stage 2, we then decide on the sector opening scheme by setting $z_{acu} = 1$ if airspace a gets configuration c at time u , and 0 otherwise. This sector opening scheme is subject to the fixed capacity budget \mathbf{h} . Furthermore, we decide on demand management measures in stage 2: $y_{fr} = 1$ represents assigning flight f to route $r \in R_f$, and 0 otherwise. We summarize the notation in Table 1.

The optimization problem that we tackle can be written as the minimization of expected displacement cost and capacity cost over all possible sets of flight scenarios F by deciding

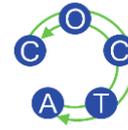


Table 1: Overview of notation.

Sets:	
F	Flight scenario including both scheduled and non-scheduled flights
\mathcal{F}	Finite collection of flight scenarios
R_f	Finite set of re-routing and delay options available to flight f
U	Set of time periods covering the day of operation
A	Set of airspaces
C^a	Set of configurations for airspace a
P^c	Partition of sectors corresponding to a configuration c
Indices:	
f	Flights
u	Time index
r	Route option, fixed in both spatial and temporal terms
a	Airspace
c	Airspace's configuration
p	Airspace sector
Parameters:	
$\gamma = (\gamma_a)_{a \in A}$	Unit cost of one sector hour for airspace a
K_p	Maximum capacity of airspace sector p
$\mathbf{h} = (h_a)_{a \in A}$	Budgets of available sector-hours for all airspaces $a \in A$
\bar{h}_{ac}	Number of sector-hours consumed by airspace a working in configuration c per unit c
d_{fr}	Displacement cost of route r for flight f
$b_{frpu} \in \{0, 1\}$	Indicates whether route r uses sector p at time u
Variables:	
$z_{acu} \in \{0, 1\}$	Indicates whether configuration c is open in airspace a in time period u
$y_{fr} \in \{0, 1\}$	Indicates whether flight f is assigned to route r

on the capacity budget \mathbf{h} :

$$\min_{\mathbf{h} \geq 0} \mathbb{E}_F[G(F|\mathbf{h})] + \gamma^T \mathbf{h}, \quad (3.1)$$

where $G(F|\mathbf{h})$ represents the minimum displacement cost to accommodate flight scenario F under capacity budgets \mathbf{h} . The superscript T denotes the transpose of a vector. To ensure feasibility of $G(F|\mathbf{h})$, we add a dummy configuration c_0 that requires no capacity ($\bar{h}_{ac_0} = 0$ for all a), is used by a dummy route $r_0 \in R_f$ for all flights $f \in F$, and its single sector $p \in P^{c_0}$ has capacity $K_p = |F|$. Using this dummy sector for one time period incurs a high penalty cost of M . With this construct, the displacement cost function $G(F|\mathbf{h})$ for a given flight scenario F is defined by a deterministic integer program:

$$G(F|\mathbf{h}) = \min_{\mathbf{y}, \mathbf{z}} \sum_{f \in F} \sum_{r \in R_f} d_{fr} y_{fr} + \sum_{a \in A, u \in U} M z_{acu}$$

$$\text{s.t.} \quad \sum_{f \in F} \sum_{r \in R_f} b_{frpu} y_{fr} z_{acu} \leq K_p z_{acu} \quad \forall a \in A, c \in C^a, p \in P^c, u \in U \quad (3.2)$$

$$\sum_{u \in U} \sum_{c \in C^a} \bar{h}_{ac} z_{acu} \leq \bar{h}_a \quad \forall a \in A \quad (3.3)$$

$$\sum_{r \in R_f} y_{fr} = 1 \quad \forall f \in F \quad (3.4)$$

$$\sum_{c \in C^a} z_{acu} = 1 \quad \forall a \in A, u \in U \quad (3.5)$$

$$z_{acu} \in \{0, 1\} \quad \forall a \in A, c \in C^a, u \in U \quad (3.6)$$

$$y_{fr} \in \{0, 1\} \quad \forall f \in F, r \in R_f. \quad (3.7)$$

The objective of $G(F|\mathbf{h})$ is to minimize displacement and penalty costs, subject to (3.2): used capacity in a given sector and time period being less than maximum capacity; (3.3): total capacity usage to not exceed the capacity budget \mathbf{h} ; (3.4): every flight being assigned to exactly one trajectory; (3.5): every airspace a having exactly one configuration at every time u ; and binary constraints (3.6–3.7). Note that constraint (3.2) is non-linear; it could be linearized by introducing a significant number of new variables and constraints. Instead, to solve $G(F|\mathbf{h})$ directly, we propose to replace constraints (3.2) with:

$$\sum_{f \in F} \sum_{r \in R_f} b_{frpu} y_{fr} \leq K_p z_{acu} + |F| \sum_{c' \neq c} z_{ac'u} \quad \forall a \in A, c \in C^a, p \in P^c, u \in U. \quad (3.8)$$

Constraints (3.8) have a loose linear programming relaxation but, at least for relatively small instances, they allow solution of $G(F|\mathbf{h})$ with commercial solvers like CPLEX as discussed in the numerical results section. We present $G(F|\mathbf{h})$ with the non-linear constraints (3.2) because we use them in our proposed decomposition approach.

In summary, (3.1) is a type of newsvendor problem that is difficult to solve even for moderately-sized instances due to the challenges in evaluating the expectation. We do not have a closed-form expression of the distribution that underpins the realizations of non-scheduled flights, and we need to solve a large binary program for every such realization to evaluate the costs in stage 2. We discuss our solution approach in the following section.



4 Approximation of the Stochastic Problem

The difficulty in solving this problem rests within the evaluation of the expectation in the problem (3.1). We can approximate the expectation by sampling a large number of non-scheduled flight scenarios, each of which we add to the set of known scheduled flights to form the collection \mathcal{F} of all these sample scenarios. This requires us to solve $G(F|\mathbf{h})$ for a large number of flight scenarios $F \in \mathcal{F}$, which is challenging given the complexity of $G(F|\mathbf{h})$. In the following, we propose two ways of tackling the problem (3.1).

4.1 Naïve Approach

Let us consider the properties of the newsvendor problem (3.1). The function $G(F|\mathbf{h})$ is not generally convex in \mathbf{h} due to the integer variables (see Geoffrion [1974]). In order to obtain a convex lower bound on the optimal objective of (3.1), we can consider the linear program relaxation of $G(F|\mathbf{h})$ with non-linear constraints (3.2) replaced with (3.8); we denote this linear program by $\tilde{G}(F|\mathbf{h})$:

$$\begin{aligned} & \min_{\mathbf{h} \geq 0} \mathbb{E}_F[G(F|\mathbf{h})] + \gamma^T \mathbf{h} \\ & \approx \min_{\mathbf{h} \geq 0} \sum_{F \in \mathcal{F}} \frac{G(F|\mathbf{h})}{|\mathcal{F}|} + \gamma^T \mathbf{h} \\ & \geq \min_{\mathbf{h} \geq 0} \sum_{F \in \mathcal{F}} \frac{\tilde{G}(F|\mathbf{h})}{|\mathcal{F}|} + \gamma^T \mathbf{h}. \end{aligned} \quad (4.1)$$

The problem (4.1) is convex and the gradient with respect to \mathbf{h} corresponds to the dual vector of the capacity budget constraints (see Theorem 4.6 in Sierksma [2002]). Thus, it can be solved e.g. by steepest descent, so starting from some initial \mathbf{h}^1 , in iteration k we update

$$h_a^{k+1} = h_a^k - \alpha \left(\sum_{F \in \mathcal{F}} \frac{\pi_{a,F}^k}{|\mathcal{F}|} + \gamma_a \right) \quad \forall a \in A, \quad (4.2)$$

where $\pi_{a,F}^k \leq 0$ is the dual value of the budget constraint (3.3) for airspace a in $\tilde{G}(F|\mathbf{h}^k)$, and $\alpha > 0$ is the step size. The dual value $\pi_{a,F}^k$ can be interpreted as the marginal reduction in displacement cost when increasing the capacity budget. The algorithm terminates when either $\|\sum_{F \in \mathcal{F}} \frac{\pi_{a,F}^k}{|\mathcal{F}|} + \gamma\| \leq \epsilon$ for a small $\epsilon < 0$, or when a certain maximum number of iterations has been reached. Naturally, this approach suffers from ignoring the integer constraints, and we demonstrate in the numerical results section in §5 that the resulting policy is poor. We only include this approach to emphasize that the integer constraints cannot be dropped.

4.2 Perfect Foresight Approach

We propose another approach under the assumption of perfect foresight, i.e. we assume to have full knowledge of both scheduled and non-scheduled traffic F at the time of deciding

on the budget \mathbf{h} . In that case, we can decide on required capacity and demand actions simultaneously. Given a flight scenario F , the corresponding problem can be formulated as follows:

$$G(F) = \min_{\mathbf{y}, \mathbf{z}} \sum_{a \in A} \gamma_a \sum_{c \in C^a} \sum_{u \in U} \bar{h}_{ac} z_{acu} + \sum_{f \in F} \sum_{r \in R_f} d_{fr} y_{fr} \quad (4.3)$$

$$\begin{aligned} \text{s.t. } & \sum_{f \in F} \sum_{r \in R_f} b_{frpu} y_{fr} z_{acu} \leq K_p z_{acu} \quad \forall a \in A, c \in C^a, p \in P^c, u \in U \quad (4.4) \\ & \sum_{r \in R_f} y_{fr} = 1 \quad \forall f \in F \\ & \sum_{c \in C^a} z_{acu} = 1 \quad \forall a \in A, u \in U \\ & z_{acu} \in \{0, 1\} \quad \forall a \in A, c \in C^a, u \in U \\ & y_{fr} \in \{0, 1\} \quad \forall f \in F, r \in R_f \end{aligned}$$

Note that a solution to $G(F)$ gives the capacity budgets as $h_a := \sum_{c \in C^a} \sum_{u \in U} \bar{h}_{ac} z_{acu}$. If we can solve $G(F)$ for many flight scenarios, then we can use the resulting sample distribution of capacity budgets to derive decision policies.

To that end, consider the structure of $G(F)$: the decisions \mathbf{y} and \mathbf{z} are only linked by the capacity constraints (4.4). We propose to decompose the problem into a master and a sub-problem according to variables \mathbf{y} and \mathbf{z} by assuming that we have a linear function $\tilde{Q}(\mathbf{y})$ that reflects (approximately) the minimal capacity cost under fixed traffic assignment \mathbf{y} . We formulate the master problem as a linear program:

$$\begin{aligned} \text{(Master)} \quad & \min_{\mathbf{y}} \sum_{f \in F} \sum_{r \in R_f} d_{fr} y_{fr} + \tilde{Q}(\mathbf{y}) \\ \text{s.t. } & \sum_{r \in R_f} y_{fr} = 1 \quad \forall f \in F \\ & y_{fr} \in [0, 1] \quad \forall f \in F, r \in R_f. \end{aligned} \quad (4.5)$$

The sub-problem minimizes capacity provision cost for a given flight assignment. To ensure the existence of a feasible solution given traffic assignment \mathbf{y} , we introduce new parameters k_{acu} that represent the capacity shortage in terms of the number of flights that exceed sector capacity limits in airspace a , configuration c and time period u . We define $k_{acu} := \sum_{p \in P^c} (\sum_{f \in F} \sum_{r \in R_f} b_{frpu} y_{fr} - K_p)^+$, where $x^+ := \max\{x, 0\}$. A large penalty cost of M is associated with these parameters to indicate violation of the capacity constraints to the master problem. The sub-problem can be written as the following linear program:

$$\begin{aligned} \text{(Sub)} \quad Q(\mathbf{y}) = & \min_{\mathbf{z}} \sum_{a \in A} \gamma_a \sum_{u \in U} \sum_{c \in C^a} (\bar{h}_{ac} + M k_{acu}) z_{acu} \\ \text{s.t. } & \sum_{c \in C^a} z_{acu} = 1 \quad \forall a \in A, u \in U \\ & z_{acu} \in [0, 1] \quad \forall a \in A, c \in C^a, u \in U \end{aligned}$$



Both master and sub-problem each have a very attractive feature, which is due to their constraint matrices being totally unimodular:

Proposition 4.1. *The optimal solution of the linear program (Master) is integer. Also, the optimal solution of the linear program (Sub) is integer.*

We still need a suitable linear function to approximate the capacity cost $Q(\mathbf{y})$. Moreover, we would like to estimate a function that approximates $Q(\mathbf{y})$ but does not directly depend on \mathbf{y} in as far as we aim to use the same estimator for different traffic scenarios (including different number of flights). This would allow us to evaluate a large number of flight scenarios without having to re-estimate the cost function for each scenario.

To that end, we propose a group of so-called basis functions to construct an approximation of $Q(\mathbf{y})$. The basis functions are intended to avoid the explicit dependency on the set of flights in the considered traffic scenario by capturing some basic problem features. Let ϕ_{acpu} denote a basis function that reflects the traffic flow in the partition p of configuration c in airspace a at time period u . We can transform the traffic assignment \mathbf{y} into the basis function ϕ_{acpu} as follows:

$$\phi_{acpu}(\mathbf{y}) = \sum_{f \in F} \sum_{r \in R_f} b_{frpu} y_{fr}, \quad \forall a \in A, c \in C^a, p \in P^c, u \in U. \quad (4.6)$$

The idea is to approximate the cost function $Q(\mathbf{y})$ with a function $\tilde{Q}(\phi)$ that establishes a direct relationship between traffic flows ϕ in specific partitions per unit of time and resulting capacity cost. With this, we do not have to specify individual flights; instead, only the aggregated traffic counts in space and time are of interest, which allows us to use the same function $\tilde{Q}(\phi)$ for different traffic samples. We define $\tilde{Q}(\phi)$ as a linear function with unknown coefficients β that need to be estimated:

$$\tilde{Q}(\phi, \beta) = \beta_0 + \sum_{a,c,p,u} \beta_{acpu} \phi_{acpu}.$$

The optimal vector β is the one that minimizes the expected squares of deviations from the optimal capacity cost $Q(\mathbf{y}^{*F})$, where the expectation is taken over the distribution of flight scenarios F , and \mathbf{y}^{*F} is the optimal flight assignment under flight scenario F :

$$\min_{\beta} \frac{1}{2} \mathbb{E}_F \left[\tilde{Q}(\phi(\mathbf{y}^{*F}), \beta) - Q(\mathbf{y}^{*F}) \right]^2. \quad (4.7)$$

Since we do not have the optimal assignments \mathbf{y}^{*F} yet, we propose an iterative approach: starting with an initial guess β^k (with iteration counter $k = 0$), we solve the (Master) problem for all samples $F \in \mathcal{F}$ using β^k . This produces a collection of flight-to-route assignment vectors $\mathbf{y}^{k,F}$. Next, we solve $|\mathcal{F}|$ (Sub) problems, one for each solution $\mathbf{y}^{k,F}$. This gives us a collection of cost/traffic flow pairs $(Q(\mathbf{y}^{k,F}), \phi(\mathbf{y}^{k,F}))$. These can be used to update β^k with a stochastic gradient step on objective (4.7) with stepsize α , which are fed into the (Master) problem again, etc. The procedure terminates when β converges or a maximum number of iterations has been reached. We provide the details in Algorithm 1.

Algorithm 1 Perfect foresight approach to approximately solve $G(F)$ for $F \in \mathcal{F}$

Initialize $\beta^0 = \vec{0}$, iteration counter $k = 0$, given traffic samples $F \in \mathcal{F}$.

while do

$\mathbf{y}^{k,F} \leftarrow$ solution of (Master) using β^k for all $F \in \mathcal{F}$

$(Q(\mathbf{y}^{k,F}), \phi(\mathbf{y}^{k,F}), \mathbf{z}^{k,F}) \leftarrow$ solution of (Sub) using $\mathbf{y}^{k,F}$ for all $F \in \mathcal{F}$

$\beta_0^{k+1} \leftarrow \beta_0^k - \alpha \sum_{F \in \mathcal{F}} (\tilde{Q}(\phi(\mathbf{y}^{k,F}), \beta^k) - Q(\mathbf{y}^{k,F})) / |\mathcal{F}|$

$\beta_{acpu}^{k+1} \leftarrow \beta_{acpu}^k - \alpha \sum_{F \in \mathcal{F}} (\tilde{Q}(\phi(\mathbf{y}^{k,F}), \beta^k) - Q(\mathbf{y}^{k,F})) \phi_{acpu}^{k,F} / |\mathcal{F}|$, for all a, c, p, u

break if $\|\beta^{k+1} - \beta^k\| < \epsilon$, or k exceeds maximum iteration limit

$k \leftarrow k + 1$

end while

return collection of solutions $(\mathbf{y}^{k,F}, \mathbf{z}^{k,F})$ for all F

As an output of Algorithm 1, we obtain sector configurations \mathbf{z}^F for all flight scenarios F . For each scenario F , we can define the capacity budget for airspace a as $h_a^F := \sum_{u \in U} \sum_{c \in C^a} \bar{h}_{ac} z_{acu}^F$. This resulting sample distribution of vectors \mathbf{h} over F is used in §5.1 to define decision policies on how to choose the capacity budgets.

5 Numerical Experiments

We apply our proposed methodology to real-world data to get a sense of the speed with which we can reach decisions, and of their quality. The latter is measured in terms of costs incurred as well as in terms of allowing feasible assignments given the capacity budget. In this section, we first define several decision policies in §5.1, describe the simulation study and underpinning real-world data in §5.2, and report our results in §5.3.

5.1 Decision Policies

The ultimate objective of this work is to derive capacity ordering decisions $\mathbf{h}^* = (h_a^*)_{a \in A}$ for the newsvendor problem (3.1). We test the following policies:

- *NA*: In the naïve policy, the capacity decision is obtained directly via solving linear programming relaxation of the problem using the steepest decent method; the iterations are defined in (4.2). The step-size is set to $\alpha = 10^{-4}$.
- *AV*: In the averaging policy, the capacity decision is obtained by defining $h_a^* := \sum_{F \in \mathcal{F}} \frac{h_a^F}{|\mathcal{F}|}$ for all $a \in A$, where \mathbf{h}^F is the perfect foresight solution for scenario F with step-size $\alpha = 10^{-9}$ (small due to scaling issues).
- $\epsilon = x\%$: In the risk-based policy, the capacity decision is obtained by setting \mathbf{h}^* such that the sample probability of encountering a flight scenario in which we had better planned for more capacity in at least one airspace is less than a given ϵ , i.e.



$\text{Prob}(\mathbf{h} \neq \mathbf{h}^*) < \epsilon$, where the sample probability distribution of \mathbf{h} has been computed by the perfect foresight approach over scenarios $F \in \mathcal{F}$ with step-size $\alpha = 10^{-9}$ (small due to scaling issues). The ϵ policy is tested with $\epsilon \in \{0.01, 0.05, 0.10, 0.20\}$. Setting ϵ to large values such as 0.10 or 0.20 reflects an increasing willingness to accept the risk of making capacity decisions which could work poorly under a significant number of traffic scenarios. On the other extreme, setting ϵ to smaller values such as 0.01 or 0.05 represents less risk and thus is more appropriate for risk-averse decision makers.

- *SA*: In the sampling policy, the capacity decision is obtained by sampling a \mathbf{h}^* uniformly from the sample distribution of \mathbf{h} . The distribution is computed by the perfect foresight approach over scenarios $F \in \mathcal{F}$ with step-size $\alpha = 10^{-9}$ (small due to scaling issues).

5.2 Simulation Study

In order to numerically evaluate the performance of each decision policy, we conduct simulation studies of the two-stage planning problem (3.1). In each simulation run, we start in stage 1 with obtaining the capacity budget \mathbf{h} using a given decision policy based on a sample of flight scenarios $\mathcal{F} \subseteq \mathcal{F}'$, where \mathcal{F}' is a finite set of flight scenarios consisting of all flight scenarios that one can encounter in stage 2. We model the true distribution of flight scenarios as being uniform over the finite set \mathcal{F}' . Each flight scenario in \mathcal{F}' is constructed by adding a number of non-scheduled flights sampled from historic data to the same fixed set of scheduled flights.

We keep \mathcal{F} fixed across all simulations and all policies, i.e. the decision policies are all based on the same collection of traffic scenarios. To model the effect of imperfect forecasts of traffic distributions, we vary the collection of possible traffic materializations \mathcal{F}' by adding more traffic scenarios $F \notin \mathcal{F}$ to the collection \mathcal{F}' of traffic scenarios. We emphasize that the number of sampled flights initially stays the same (so we assume that we know in stage 1 how many non-scheduled flights will appear in stage 2), but we do not know when and where they appear; we also experiment with the impact of increasing the number of sampled non-scheduled flights. The assumption $\mathcal{F} = \mathcal{F}'$ corresponds to having full knowledge in stage 1 of the true traffic distribution. When $\mathcal{F} \subsetneq \mathcal{F}'$, we only can exploit partial knowledge of the true traffic distribution, thus reflecting an increasingly poor forecast the larger the cardinality of \mathcal{F}' relative to that of \mathcal{F} . We consider the cases of $\mathcal{F} = \mathcal{F}'_0$, and $\mathcal{F} \subsetneq \mathcal{F}'_1$, where \mathcal{F} features 100 traffic scenarios and \mathcal{F}'_1 has additional 300 scenarios that are not contained in \mathcal{F} .

Given the decision \mathbf{h} from stage 1, we assess in stage 2 the optimal displacement cost $G(\hat{F}|\mathbf{h})$. To do so, we use CPLEX to solve the binary program $G(\hat{F}|\mathbf{h})$ with capacity constraints (3.2) replaced with (3.8). As a result, we obtain the total cost incurred, whether $G(\hat{F}|\mathbf{h})$ was feasible, and if not, how many flights violated capacity constraints at any period of time. This cost evaluation in every simulation is fairly time-consuming, which limits the volume of flights that we can investigate in a simulation. We run the simulation for 200 times to measure the average performance in terms of displacement costs, capacity



Figure 1: Snapshot of scheduled flights and airspaces considered in the simulation at a fixed point in time.

costs and flights exceeding capacity constraints.

The capacity and demand data are obtained via Eurocontrol Demand Data Repository (DDR2) service using the Eurocontrol Network Strategic (NEST) Tool. The selected region includes a large part of en-route airspace in Central and Western Europe, including 8 ANSPs and 15 ACCs/sector groups as depicted in Figure 1. The figure shows a snapshot of scheduled traffic at a particular point in time and illustrates the presence of different traffic densities. The Maastricht Upper Area Control Centre (MUAC) is divided into three sector groups: Deco, Hannover and Brussels, each with its own sectorization and sector configurations. Based on historical usage of configurations in 2016, for each ACC or sector group, we select configurations with different number of sectors which were most frequently used: in total, we have 173 different configurations for the 15 ACCs/sector groups shown in Table 6 in the appendix. The overview of elementary and collapsed sectors used in the simulation is likewise in the appendix, namely Table 7.

Delay and re-routing costs are obtained by drawing on Cook and Tanner [2015] and Eurocontrol [2018b]. Delay costs are calculated per aircraft type and per duration of delay (non-linear with time). Re-routing costs include fuel costs, crew costs, passenger soft and passenger hard costs, as well as maintenance costs. Cost parameters for ANS provision are derived in the following way: based on information provided in the ATM Cost-Effectiveness Report 2017, we calculated average ATCO costs per sector-hour in 2015 for the different ANSPs (and in one case for an ACC). In the model, we treat those costs as variable costs. Although ATCO costs might be considered fixed costs in the short term, reducing the daily number of ATCO hours will reduce the total number of ATCOs needed for ANS provision,



consequently reducing staff costs.

We use flight data of 9th September 2016, which was the busiest day in 2016 in European airspace. We select flights based on their last filed flight plans, which cross the selected airspace between 10 and 11 AM. In total, we have 910 scheduled flights that are considered fixed in our network in all scenarios. Out of all non-scheduled flights on that day, we select those that cross the selected airspace at any time (1,569 in total). Since we test the model for 10-11 AM period, we change their airspace entry times from the original flight plan to a time uniformly sampled over the selected period. Each traffic scenario is created by uniformly sampling a subset of 160 from this set of 1,569 flights and adding them to the set of scheduled flights. We create 100 traffic scenarios, stored in \mathcal{F} , in this manner. Flights can be either delayed or re-routed (only one demand management measure per flight) to improve total cost-efficiency subject to hard capacity budget constraints. Delay options are discrete and the same for each flight, namely 5, 15, 30 or 45 minutes. Each flight has a number of alternative spatial routes as well, all generated using the NEST tool. Overall, the problem is modeled over a two-hours time horizon to account for flights being delayed beyond 11:00 AM.

We chose this data to test our approach since the selected airspace covers a large part of the so-called ‘core area’ (e.g. MUAC and DFS) of the European airspace, as well as a part of airspace which is not as congested (Hungaro Control). To ensure reproducibility of our work and to allow others to conduct benchmark studies against our approach, all data and open source code underpinning the numerical results reported in the paper will be published in the public domain.

5.3 Results and Discussion

The simulation study provides us with some insight on how the policies are working. Algorithms are implemented using C++ and Visual Studio IDE. Libraries provided by the commercial solver CPLEX are used to solve the optimization models. Tests are run on a Windows 8 PC, with 8GB of RAM and AMD Ryzen 7-1700 processor.

Tables 2 and 3 report the average performances of the policies obtained by simulation scenarios under the assumption that we know in stage 1 how many non-scheduled flights will appear, but not when and where. The two tables differ in the level of uncertainty regarding potential flight scenarios: Table 2 shows the results when $\mathcal{F}'_0 = \mathcal{F}$, whereas Table 3 contains the results for $\mathcal{F}'_1 \supsetneq \mathcal{F}$, meaning that the flight materializations for the latter are drawn from a wider set of traffic scenarios containing 300 scenarios in addition to those in \mathcal{F} . The average share of flights that cannot be assigned to any non-dummy route is also shown, providing insights on infeasibility of a given policy.

As expected, policy *NA* results in significantly lower sector-hour budgets than other policies but incurs drastic displacement costs. This is a consequence of the binary variables relaxation underlying this policy since allowing partial flight to routes assignments leads to fewer resources needed to satisfy the capacity constraints. All other policies show consistent behavior across the different simulation scenarios with policy *AV* overall doing best. However, we also notice that uncertainty on the time and location of non-scheduled

Table 2: Simulation results under full knowledge of non-scheduled flight distribution. Exactly the expected number of non-scheduled flights materialize, sampled from $\mathcal{F}'_0 = \mathcal{F}$.

	Capacity Policy	Displacement cost	Total cost	Flights not assigned (%)
<i>NA</i>	29,032	754,855	783,887	0.7
<i>AV</i>	51,854	2,708	54,562	0.0
$\epsilon = 20\%$	57,808	404	58,212	0.0
$\epsilon = 10\%$	61,107	370	61,477	0.0
$\epsilon = 5\%$	61,877	392	62,269	0.0
$\epsilon = 1\%$	64,665	360	65,025	0.0
<i>SA</i>	58,669	384	59,053	0.0

Table 3: Simulation results under low uncertainty. Exactly the expected number of non-scheduled flights materialize, sampled from $\mathcal{F}'_1 \supseteq \mathcal{F}$.

	Capacity Policy	Displacement cost	Total cost	Flights not assigned (%)
<i>NA</i>	29,032	560,948	589,980	2.1
<i>AV</i>	51,854	2737	54,591	0.0
$\epsilon = 20\%$	57,808	417	58,225	0.0
$\epsilon = 10\%$	61,107	404	61,511	0.0
$\epsilon = 5\%$	61,877	393	62,270	0.0
$\epsilon = 1\%$	64,665	370	65,035	0.0
<i>SA</i>	58,669	382	59,051	0.0

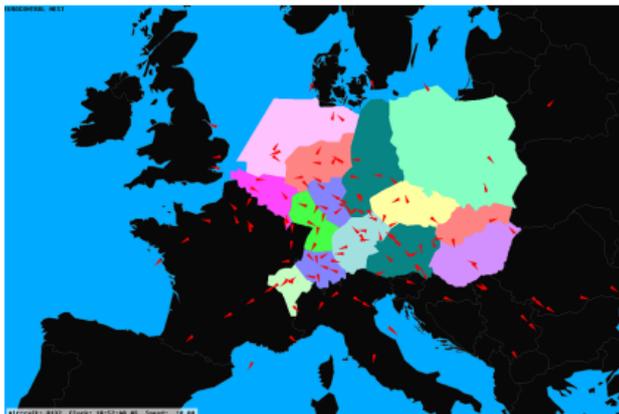


Figure 2: Snapshot of non-scheduled flights at a fixed point in time. This snapshot shows non-scheduled traffic under the low uncertainty scenario, i.e. only 160 flights.

flights does not have a significant impact on the results. This uncertainty is absorbed with the existing capacity or by demand measures. This is not too surprising when we consider a snapshot of the non-scheduled traffic shown in Figure 2: due to the relatively low volume of flights, the non-scheduled traffic is fairly widely distributed so that its impact on individual sectors is minimal.

However, if we additional increase uncertainty over the number of non-scheduled flights, results change significantly as can be seen in Table 4. Specifically, we test the policies assuming that the number of non-scheduled flights materializing is twice the one assumed in the traffic scenarios during the policy building stage. This results in a total of 1,300 flights, 30% of which are non-scheduled.

Table 4: Simulation results under high uncertainty. Twice the expected number of non-scheduled flights materialize, sampled from wider collection $\mathcal{F}'_1 \supseteq \mathcal{F}$.

Capacity Policy	Displacement cost	Total cost	Flights not assigned (%)	
NA	29,032	1,078,160	1,644,242	4.1
AV	51,854	15,894	67,748	0.0
$\epsilon = 20\%$	57,808	3,960	61,768	0.0
$\epsilon = 10\%$	61,107	3,412	64,519	0.0
$\epsilon = 5\%$	61,877	3,408	65,285	0.0
$\epsilon = 1\%$	64,665	2,888	67,553	0.0
SA	58,669	3,432	62,101	0.0

The added uncertainty in traffic volume in the network has a clear impact on displacement costs, which increase by one order of magnitude across all policies. The perhaps

most interesting result is that $\epsilon = 20\%$ yields the lowest total cost; this intuitively makes sense since the ϵ policies are hedging against the risk of underprovision so that with higher levels of uncertainty, we overall perform better. Also, policy *AV* now performs worse than all the other policies (*NA* excluded). This is because ϵ policies result in higher capacity orderings, giving them more room to accommodate the additional traffic. In summary, under high uncertainty on the number of non-scheduled flights to appear we may want to use the risk-based ϵ policies. If the number of non-scheduled flights tends to be fairly stable, we may be best off with the averaging policy since, at least in our case study, uncertainty in the temporal and spatial aspect of these flights can be absorbed with the *AV* capacity budgets and demand management measures. The sampling policy *SA* was dominated by other policies in all scenarios.

In Table 5 we compare the computing time of solving $G(F)$ for a given traffic scenario with CPLEX versus using the perfect foresight (*PF*) algorithm 1. The first column shows the number of flights in F and the second reports the computing time required by CPLEX to solve $G(F)$ to optimality. A single traffic scenario F is sampled from \mathcal{F} and used to test the runtime of CPLEX. Finally, the last column shows the average computing time used by Algorithm 1 per flight scenario (this can vary somewhat depending on how quick the cost function parameters converge in the various scenarios).

Table 5: Average computational times to solve an instance $G(F)$. Note that the times for the perfect foresight approach are averaged over the samples $F \in \mathcal{F}$.

Number of flights	CPLEX time (s)	Average <i>PF</i> time (s)
1,000	532	9.1
1,100	1,678	7.0
1,200	3,938	27.0
1,300	8,506	27.8
1,400	n.a.	25.7

Solving the problem for just one traffic scenario becomes rapidly impossible with CPLEX as the number of flights increases. As for the 1,400 flights case, the solver stopped after failing to find the optimal solution within 250 minutes. This illustrates that a direct solution approach is not scalable, as opposed to proposed decomposition approach.

6 Conclusion

Capacity planning for air traffic navigation service provision is a challenging problem due to uncertain traffic and the volume of decisions to be taken on sector configurations, re-routing and delaying of flights. We propose a decomposition approach that can solve this problem in a manner that is scalable to real-life applications: it is already quite efficient since it only requires the solution of linear programs, but solution times could be further improved by solving the master problems in parallel across all flight scenarios (and likewise for the sub-problem) in every iteration. Once the algorithm has stopped, the resulting output



can be used to construct various capacity planning policies, including one that can reflect the risk attitude of the decision maker. The study has been evaluated on real data from Central European airspaces with over 1,000 flights.

We feel that one of the main advantages of the approach is that we can use it in conjunction with any non-scheduled traffic forecasting approach that can generate traffic scenario samples. This includes traffic predictions that are derived with distribution-free methods such as machine learning approaches. Moreover, it can be used to test for robustness against disruptions: the effects of having to close off a certain sector during a certain period of time in a certain traffic scenario could be embedded by pre-processing the set of trajectory options R_f for each flight f so as to remove all trajectories that would traverse this sector during the closure period. This may be of interest to increase robustness of the decision policy against the effects of random weather events or closures due to exclusive use by military aircraft on capacity planning.

A limitation of this work is the assumption that capacity planning can be reduced to deciding on the number of sector-hours required for each airspace. In practice, it may not always be straight-forward to link such a budget decision to concrete staff rosters. Furthermore, it could be beneficial to also consider airport congestion and propagated delay in the network. A practical complication with including propagated delay is that real data may not permit to easily deduce which two flights are operated by the same aircraft. In future research, it would be interesting to investigate how such constraints could be taken into account.

Acknowledgements

This project has received funding from the SESAR Joint Undertaking under grant agreement No 699326 under European Union's Horizon 2020 research and innovation programme. The opinions expressed herein reflect the author's view only. Under no circumstances shall the SESAR Joint Undertaking be responsible for any use that may be made of the information contained herein. Arne Strauss acknowledges support by the Alan Turing Institute through a Turing Fellowship.

Appendix: ACCs and Sectors in the Simulation

Table 6: ACCs used in the simulation

Country	ANSP	ACC/cluster(s)	ACC/cluster Name
Germany	DFS	EDUUUTAC	Karlsruhe UAC Central
		EDUUUTAE	Karlsruhe UAC East
		EDUUUTAS	Karlsruhe UAC South
		EDUUUTAW	Karlsruhe UAC West
Germany, Belgium, Netherlands	MUAC	EDYYBUTA	MUAC Brussels Sectors
		EDYYDUTA	MUAC Deco Sectors
		EDYYHUTA	MUAC Hannover Sectors
Switzerland	Sky Guide	LSAGUTA	Genève ACC Upper FL 245-9
		LSAZUTA	Zurich (Dübendorf) ACC Upper
Czech Republic	ANS	LKAACTA	Praha CTA 095/125-305
		LKAAUTA	Praha UTA FL 285-999
Slovakia	LPS	LZBBCTA	Bratislava ACC
Austria	Austro Control	LOVVCTA	Wien ACC
Hungary	Hungaro Control	LHCCCTA	Budapest ACC
Poland	PANSA	EPWWCTA	Warszawa ACC



Table 7: Number of elementary and collapsed sectors used in the simulation

Area Control Centre/Cluster	Elementary sectors	Collapsed sectors	Min sectors opened in a configuration	Max sectors opened in a configuration	Number of dif configurations
EDUUUTAC	11	14	1	9	13
EDUUUTAE	10	14	1	9	13
EDUUUTAS	12	29	1	8	13
EDUUUTAW	10	12	1	8	11
EDYYBUTA	8	13	1	6	10
EDYYDUTA	9	12	1	6	7
EDYYHUTA	12	19	1	7	12
EPWWCTA	18	77	1	10	26
LHCCCTA	10	24	1	7	7
LKAACTA	6	9	1	5	6
LKAAUTA	6	9	1	5	8
LOVVCTA	26	58	2	12	21
LSAGUTA	6	9	1	6	11
LSAZUTA	6	7	1	5	7
LZBBCTA	27	69	1	5	8

References

- M. Baumgartner. *European Air Traffic Management : principles, practice and research*, chapter The organisation and operation of European airspace, pages 1–34. Ashgate Publishing Limited, 2007.
- D. Bertsimas and S. S. Patterson. The air traffic flow management problem with enroute capacities. *Operations research*, 46(3):406–422, 1998.
- L. Castelli, R. Pesenti, and A. Ranieri. The design of a market mechanism to allocate air traffic flow management slots. *Transportation research part C: Emerging technologies*, 19(5):931–943, 2011.
- A. Cook and G. Tanner. European airline delay cost reference values. Technical report, University of Westminster, 2015.
- L. Corolli, L. Castelli, and G. Lulli. The air traffic flow management problem with time windows. In *4th International Conference on Research in Air Transportation (ICRAT)*, Budapest, Hungary, 2010.
- L. Delgado, A. J. Cook, and S. Cristóbal. Controller Time and Delay Costs - a trade-off analysis. In *Fifth SESAR Innovation Days*, Bologna, 2015.
- Eurocontrol. Capacity assessment and planning guidance document, Ed. 2.8. Technical report, EUROCONTROL, 2013. URL [http://www.eurocontrol.int/sites/default/files/content/documents/nm/airspace/capacity](http://www.eurocontrol.int/sites/default/files/content/documents/nm/airspace/capacity_assessment_and_planning_guidance_document_2.8.pdf)
- Eurocontrol. ATM cost-effectiveness (ACE) 2016 benchmarking report with 2017-2021 outlook. Technical report, EUROCONTROL, May 2018a.
- Eurocontrol. Standard inputs for eurocontrol cost-benefit analyses. Technical report, EUROCONTROL, 2018b.
- European Commission. Commission regulation (EU) no 677/2011. Official Journal of the European Union, July 2011.
- P. Flener and J. Pearson. Automatic airspace sectorisation : A survey. Technical report, 2013. URL <http://uu.diva-portal.org/smash/record.jsf?pid=diva2%3A675958&dswid=-1083>.
- A. Geoffrion. Lagrangean relaxation for integer programming. *Mathematical Programming Study*, 2:82–114, 1974.
- IATA. Airline cost management group, 2015. URL https://www.iata.org/whatwedo/workgroups/Documents/ACMG/ACMG_Flyer_FY2014_2015.pdf.



- B. Josefsson, T. Polishchuk, V. Polishchuk, and C. Schmidt. A Step Towards Remote Tower Center Deployment: Optimizing Staff Schedules. In *Twelfth USA/Europe Air Traffic Management Research and Development Seminar (ATM2017)*, 2017. URL http://www.atmseminarus.org/seminarContent/seminar12/papers/12th_ATM_RD_Seminar_pape:
- A. Lau, J. Berling, F. Linke, V. Gollnick, and K. Nachtigall. Large-scale network slot allocation with dynamic time horizons. *Eleventh USA/Europe Air Traffic Management Research and Development Seminar*, 2015.
- G. Lulli and A. Odoni. The European air traffic flow management problem. *Transportation Science*, 41(4):431–443, 2007.
- A. Majumdar, W. Y. Ochieng, J. Bentham, and M. Richards. En-route sector capacity estimation methodologies: An international survey. *Journal of Air Transport Management*, 11(6):375–387, nov 2005. ISSN 0969-6997. doi: 10.1016/J.JAIRTRAMAN.2005.05.002. URL <https://www.sciencedirect.com/science/article/pii/S0969699705000396>.
- RyanAir. Annual report. Technical report, RyanAir, 2018. URL <https://investor.ryanair.com/wp-content/uploads/2017/07/Ryanair-FY2017-Annual-Report>
- H. D. Sherali, R. W. Staats, and A. A. Trani. An Airspace Planning and Collaborative Decision-Making Model: Part I—Probabilistic Conflicts, Workload, and Equity Considerations. *Transportation Science*, 37(4):434–456, nov 2003. ISSN 0041-1655. doi: 10.1287/trsc.37.4.434.23272. URL <http://pubsonline.informs.org/doi/10.1287/trsc.37.4.434.23272>.
- H. D. Sherali, R. W. Staats, and A. A. Trani. An Airspace-Planning and Collaborative Decision-Making Model: Part II—Cost Model, Data Considerations, and Computations. *Transportation Science*, 40(2):147–164, may 2006. ISSN 0041-1655. doi: 10.1287/trsc.1050.0141. URL <http://pubsonline.informs.org/doi/abs/10.1287/trsc.1050.0141>.
- G. Sierksma. *Linear and Integer Programming: theory and practice*. Dekker, 2002.
- S.-L. A. Tien and R. Hoffman. Optimizing Airspace Sectors for Varying Demand Patterns using Multi-Controller Staffing. In *Eighth USA/Europe Air Traffic Management Research and Development Seminar (ATM2009)*, 2009.
- G. e. a. Tobaruela. Enhancing cost-efficiency and reducing capacity shortages: Strategic planning and dynamic shift management. In *Tenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2013)*, 2013.
- V. Tošić, O. Babic, M. Cangalovic, and D. Hohlacov. A model to solve en route air traffic flow management problem: A temporal and spatial case. *Unpublished report, University of Belgrade*, 1995.

T. K. Venugopalan, C. S. Y. Wong, S. Suresh, and N. Sundararajan. Simultaneous Optimization of Airway and Sector Design for Air Traffic Management. *Journal of Air Transportation*, 26(1):8–22, jan 2018. ISSN 2380-9450. doi: 10.2514/1.D0090. URL <https://arc.aiaa.org/doi/10.2514/1.D0090>.

