

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/128154>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Air Traffic Control Capacity Planning Under Demand and Capacity Provision Uncertainty

Stefano Starita¹, Arne Strauss^{*2}, Xin Fei², Radosav Jovanović³, Nikola Ivanov³,
Goran Pavlović³, and Frank Fichert⁴

¹School of Manufacturing Systems and Mechanical Engineering, Sirindhorn International Institute of Technology, Thammasat University, Pathum Tani 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

October 16, 2019

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 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 aircrafts. However, at the time of planning there is considerable uncertainty regarding the number and spatio-temporal distribution of non-scheduled flights and capacity provision; the former mainly due to business aviation, the latter usually stemming from the impact of weather, military use of airspaces etc. 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.

To tackle this, we propose a scalable decomposition approach that exploits the structure of the problem and can take traffic and capacity provision 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.

*Corresponding author: arne.strauss@wbs.ac.uk

1 Introduction

Air navigation service providers (ANSPs) are among the key stakeholders in providing safe and efficient air traffic operations. In Europe, there are 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 2019). 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 several hundred so-called 'elementary' sectors, which can be assembled into larger units called 'collapsed' sectors. Any specific combination of elementary and/or collapsed sectors which covers the entire volume of airspace for an ACC is called sector configuration (Baumgartner 2007).

ATCOs working in a single sector can safely provide ATC services to a limited number of aircraft in a given period of time, and this number represents the 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 2018b). 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 preferred 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 2018b). 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 2018b).

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 2018b), 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 et al. 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 et al. (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 they are difficult to incorporate when planning capacity months in advance. Furthermore,

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 both demand and supply in capacity planning.

The aim of our work is to provide a methodology for planning capacity provision in the presence of uncertain non-scheduled flights and uncertain capacity provision. We assume that capacity budgets (in terms of sector-hours) for each airspace must be acquired *before* the uncertainty regarding non-scheduled flights and capacity provision 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. In practice, ATCO rosters are determined several weeks or even months ahead (strategic level) of a day of operations. This leaves no room to call in additional ATCOs at a short notice due to constraints related to ATCO working hours (Conniss et al. 2014). ANSPs send planned sector-opening schemes to the NM at pre-tactical level, i.e. seven days in advance (NM 2019). In our paper, we determine capacity budgets (i.e., in terms of sector hours) by explicitly deciding on sector-opening schemes under several traffic and capacity scenarios. Therefore, the capacity budget has underlying information such as ‘averaged sector-opening scheme’, ‘maximum number of sectors open’ and so on, facilitating roster planning for ACCs.

Note also that ATCOs are usually only licenced to provide ATC services in groups of sectors called ‘sector families’ (Tobaruela 2015). This effectively impacts on the capacity provision, limiting the number of possible sector configurations at any given period and the flexibility of changing from one to another. Modelling this requires to constrain the problem by specifying which ATCOs can control which sectors, hence increasing the problem’s complexity (Conniss et al. 2014). Furthermore, a more generic ATCO training (‘validation’) is being tested in some ACCs and advocated as a way forward for European ANS provision. For instance, the Single European Sky ATM Research (SESAR) project ‘Any Controller, Any Airspace’ (PJ.10-06) aims to reduce the cost of ATC service through more flexible rostering and increase the resilience of the service against local disruptions (SJU 2019a,b). Therefore, in our paper we implicitly assume fairly flexible ATCO rostering.

Our main contribution is to provide an efficient solution approach to this Air Traffic Flow and Capacity Management (ATFCM) 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 and capacity shortfalls in stage 2. The policies are based on solving a version of the problem that assumes full information on all flights and capacity shortfalls, 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 on small synthetic test data as well as on a larger study based on real flight data with over 1,000 flights.

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 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, Staats, and Trani (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, Staats, and Trani (2006) investigate how their model can be used in both tactical and strategic phases. Regarding strategic applications, the 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, Cook, and Cristóbal (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 et al. (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 et al. (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 costs while deciding the assignments of both ground holds and routes. The same problem is tackled by Bertsimas and Patterson (1998). They also focus on the complexity of the same 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, Castelli, and Lulli (2010) introduce to the problem the concept of time windows. 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, Pesenti, and Ranieri (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 ATFCM slot allocation in a large-scale network. 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. Note that we primarily focus on en-route airspace and do not include airports due to complexity of the problem considered.

The topic of our paper is also related more generally to the literature on two-stage robust and stochastic optimization; we refer to the literature review of Gabrel, Murat, and Thiele (2014) that puts these concepts in relation to one another.

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. We further assume uncertainty in the capacity supply which may materialize in the form of reductions of the nominal structural capacities of the airspaces. In the second stage, uncertainty regarding non-scheduled flights and capacity 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 (see Qin et al. (2011)). 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 nominal 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). This is motivated by current practice where a required capacity profile for the following summer season is defined on an hourly basis and ACCs may even change sector configurations more frequently than that (Eurocontrol 2018a). Opening configuration c in airspace a for one time period requires \bar{h}_{ac} sector-hours.

In practice, there is supply-side uncertainty due to factors such as weather conditions, military activity, strikes of air traffic controllers etc. that affect capacity provision. We assume that we do not know the true distribution that governs the materialization of supply-side uncertainty; however, we do have a uniform distribution over a finite collection \mathcal{K} of capacity disruption scenarios $K \in \mathcal{K}$ where some sectors are operating at a fraction of their nominal capacities. We assume that this distribution can be used to approximate the true (unknown) distribution of capacity disruption scenarios. We only observe in period 2 which capacity scenario K we are facing.

Furthermore, there is also demand-side uncertainty over the number and details of non-scheduled flights in period 1 that only gets resolved in period 2. Similarly to modeling the capacity disruption scenarios, we assume that we do not know the true distribution that governs the materialization of non-scheduled flights but that we do have a uniform distribution over a finite collection of flight scenarios which we assume to approximate the true distribution. Both collections are known in period 1; for instance, this could be defined as the collection of historic operating capacities and non-scheduled flights. We augment the collection of non-scheduled flights scenarios \mathcal{F} 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.

Both sources of uncertainty are paired to form scenarios $S = (F^S, K^S)$. Scenarios are collected in set \mathcal{S} . In the following, we write the expectation over scenarios S in the understanding that uncertainty only pertains to non-scheduled flights and capacity scenarios. 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 scenarios S by deciding on the capacity budget \mathbf{h} :

$$\min_{\mathbf{h} \geq 0} \mathbb{E}_S[G(S|\mathbf{h})] + \gamma^T \mathbf{h}, \quad (1)$$

where $G(S|\mathbf{h})$ represents the minimum displacement cost to accommodate flight scenario F^S under capacity budgets \mathbf{h} , given capacities K^S . The superscript T denotes the transpose of a vector. To ensure feasibility of $G(S|\mathbf{h})$, 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 $\kappa_p = |F|$ regardless of the capacity scenario materialization. Using this dummy route incurs a high penalty cost which can be interpreted as the cost of displacing a flight beyond the planning horizon. With this construct, the displacement cost function $G(S|\mathbf{h})$ for a given flight scenario F^S and a given capacity scenario K^S , is defined by a deterministic integer program:

$$G(S|\mathbf{h}) = \min_{\mathbf{y}, \mathbf{z}} \sum_{f \in F^S} \sum_{r \in R_f} d_{fr} y_{fr}$$

$$\text{s.t.} \quad \sum_{f \in F^S} \sum_{r \in R_f} b_{frpu} y_{fr} z_{acu} \leq \kappa_p^S \quad \forall a \in A, c \in C^a, p \in P^c, u \in U \quad (2)$$

$$\sum_{u \in U} \sum_{c \in C^a} \bar{h}_{ac} z_{acu} \leq h_a \quad \forall a \in A \quad (3)$$

$$\sum_{r \in R_f} y_{fr} = 1 \quad \forall f \in F^S \quad (4)$$

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

$$z_{acu} \in \{0, 1\} \quad \forall a \in A, c \in C^a, u \in U \quad (6)$$

$$y_{fr} \in \{0, 1\} \quad \forall f \in F^S, r \in R_f. \quad (7)$$

The objective of $G(S|\mathbf{h})$ is to minimize displacement and penalty costs, subject to (2): used capacity in a given sector and time period being less than maximum capacity; (3): total capacity usage to not exceed the capacity budget \mathbf{h} ; (4): every flight being assigned to exactly one trajectory; (5): every airspace a having exactly one configuration at every time u ; and binary constraints (6–7). Note that constraint (2) is non-linear; a standard linearization of this constraint requires to introduce a new set of binary variables n_{fracu} replacing the products $y_{fr} z_{acu}$, together with a new set of constraints to enforce equivalence. This approach provides a tight relaxation but is not scalable for this problem. As an example, a medium-sized network with 10 airspaces, 10 configurations, 100 flights with 10 routes, studied over 2 time periods,

Table 1: Overview of notation.

Sets:

F	Flight scenario including both scheduled and non-scheduled flights
\mathcal{F}	Finite collection of flight scenarios F
K	Capacity scenario (meaning some sector capacities may be reduced)
\mathcal{K}	Finite collection of capacity scenarios K
S	Scenario defining traffic F^S and capacity K^S scenarios defined by pair (F^S, K^S) , with $F^S \in \mathcal{F}$
\mathcal{S}	Finite collection of uncertain scenarios S
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
κ_p^S	Maximum capacity of airspace sector p under scenario K^S
$\mathbf{h} = (h_a)_{a \in A}$	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 of time
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

requires 200,000 additional variables and 600,000 additional constraints. Consequently, in order to solve $G(S|\mathbf{h})$ directly (for small instances), we replace constraints (2) with:

$$\sum_{f \in F^S} \sum_{r \in R_f} b_{frpu} y_{fr} \leq \kappa_p^S z_{acu} + |F^S| \sum_{c' \neq c} z_{ac'u} \quad \forall a \in A, c \in C^a, p \in P^c, u \in U. \quad (8)$$

Constraints (8) have a loose linear programming relaxation but, at least for small instances, they allow solving of $G(S|\mathbf{h})$ with commercial solvers like CPLEX as discussed in the numerical results section. We represent $G(S|\mathbf{h})$ with the non-linear constraints (2) because we use them in our proposed decomposition approach.

In summary, (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 (1). We can approximate the expectation by sampling a large number of scenarios S . This requires us to solve $G(S|\mathbf{h})$ for a large number of scenarios, which is challenging given the complexity of $G(S|\mathbf{h})$. The function $G(S|\mathbf{h})$ is not generally convex in \mathbf{h} due to the integer variables (see Geoffrion (1974)); dropping the integer restrictions results in this function becoming convex in \mathbf{h} , but our computational experiments showed that the quality of the resulting solutions is very bad and thus we do not consider this approach further.

Our approach of tackling the stochastic problem (1) assumes perfect foresight, i.e. we assume to have full knowledge of both scheduled and non-scheduled traffic F and capacity scenarios K 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 scenario S , the corresponding problem can be formulated as follows:

$$G(S) = \min_{\mathbf{y}, \mathbf{z}} \sum_{a \in A} \gamma_a \sum_{c \in C^a} \sum_{u \in U} \bar{h}_{ac} z_{acu} + \sum_{f \in F^S} \sum_{r \in R_f} d_{fr} y_{fr} \quad (9)$$

$$\text{s.t. } \sum_{f \in F^S} \sum_{r \in R_f} b_{frpu} y_{fr} z_{acu} \leq \kappa_p^S z_{acu} \quad \forall a \in A, c \in C^a, p \in P^c, u \in U \quad (10)$$

$$\sum_{r \in R_f} y_{fr} = 1 \quad \forall f \in F^S$$

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

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

$$y_{fr} \in \{0, 1\} \quad \forall f \in F^S, r \in R_f$$

Note that a solution to $G(S)$ 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 $G(S)$ 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 $G(S)$: the decisions \mathbf{y} and \mathbf{z} are only linked by the capacity constraints (10). 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:

$$\begin{aligned}
(\text{Master}) \quad & \min_{\mathbf{y}} \sum_{f \in F^S} \sum_{r \in R_f} d_{fr} y_{fr} + \tilde{Q}(\mathbf{y}) \\
& \text{s.t.} \quad \sum_{r \in R_f} y_{fr} = 1 \quad \forall f \in F^S \\
& \quad \quad y_{fr} \in [0, 1] \quad \forall f \in F^S, r \in R_f.
\end{aligned} \tag{11}$$

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 $k_{acu} := \sum_{p \in P^c} (\sum_{f \in F^S} \sum_{r \in R_f} b_{frpu} y_{fr} - \kappa_p^S)^+$, 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:

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

Both master and sub-problem have a very attractive feature:

Proposition 1. *The optimal solution of the linear program (Master) and (Sub), respectively, is integer. The optimal solution of (Sub) is given by setting $z_{ac^*u} = 1$ for $c^* = \arg \min_{c \in C^a} [\gamma_a (\bar{h}_{ac} + M k_{acu})]$, and $z_{acu} = 0$ otherwise, for all $a \in A$ and $u \in U$.*

Proof. Proof: Note that the constraint matrices of (Master) and (Sub) are both 0-1 matrices with the consecutive ones property, i.e. the columns can be permuted in a way such that the 1s in every row appear consecutively. It follows that the matrices are totally unimodular, and consequently the optimal solutions to the linear programs (Master) and (Sub) are integer. The optimal solution of (Sub) follows from inspection of the problem. \square

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:

$$\phi_{acpu}(\mathbf{y}) = \sum_{f \in F^S} \sum_{r \in R_f} b_{frpu} y_{fr}, \quad \forall a \in A, c \in C^a, p \in P^c, u \in U. \quad (12)$$

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:

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

With this function, we can exploit the problem structure of (Master) to explicitly state its solution:

Remark 1. *The optimal solution of (Master) is given by setting $y_{fr^*} = 1$ for $r^* = \arg \min_{r \in R_f} [d_{fr} + \sum_{a,c,p,u} \beta_{acpu} b_{frpu}]$, and $y_{fr} = 0$ otherwise, for all flights $f \in F^S$.*

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

$$\min_{\beta} \frac{1}{2} \mathbb{E}_S \left[\tilde{Q}(\phi(\mathbf{y}^{*S}), \beta) - Q(\mathbf{y}^{*S}) \right]^2. \quad (13)$$

Since we do not have the optimal assignments \mathbf{y}^{*S} 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 $S \in \mathcal{S}$ using β^k . This produces a collection of flight-to-route assignment vectors $\mathbf{y}^{k,S}$. Next, we solve $|\mathcal{S}|$ (Sub) problems, one for each solution $\mathbf{y}^{k,S}$. This gives us a collection of cost/traffic flow pairs $(Q(\mathbf{y}^{k,S}), \phi(\mathbf{y}^{k,S}))$. These can be used to update β^k with a stochastic gradient step on objective (13) 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.

As an output of Algorithm 1, we obtain sector configurations \mathbf{z}^S for all scenarios S . For each scenario S , we can define the capacity budget for airspace a as $h_a^S := \sum_{u \in U} \sum_{c \in C^a} \bar{h}_{ac} z_{acu}^S$. This resulting sample distribution of vectors \mathbf{h} over S 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 a synthetic data set as well as a real world one in §5.2, and report our results in §5.3.

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

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

while do

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

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

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

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

break if $\|\beta^{k+1} - \beta^k\| < 0.01$, or k exceeds maximum iteration limit

$k \leftarrow k + 1$

end while

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

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 (1). The foresight approach outlined in Algorithm 1 will produce a set of solutions \mathbf{h}^S for different scenarios S . However, we still require a consensus function that maps from this collection of solutions to a single one to be used. We refer to such a consensus function as a policy and test the following ones:

- *FE*: In the full enumeration policy, the newsvendor problem is solved by computing all the possible combinations of capacity orderings across airspaces. For each combination, the expectation of costs is computed. The capacity ordering decision correspond to the vector \mathbf{h}^* yielding the smallest expected cost. This policy can be used as a benchmark but only for very small problems where full enumeration is computationally feasible.
- *AV*: In the averaging policy, the capacity decision is obtained by defining $h_a^* := \sum_{S \in \mathcal{S}} \frac{h_a^S}{|\mathcal{S}|}$ for all $a \in A$, where \mathbf{h}^S is the perfect foresight solution for scenario S 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. $\text{Prob}(\mathbf{h} \not\leq \mathbf{h}^*) < \epsilon$, where the sample probability distribution of \mathbf{h} has been computed by the perfect foresight approach over scenarios $S \in \mathcal{S}$ 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 $S \in \mathcal{S}$ 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 (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 scenarios $\mathcal{S} \subseteq \mathcal{S}'$ where \mathcal{S}' is a finite set of scenarios consisting of all scenarios that one can encounter in stage 2. We model the true distribution of scenarios as being uniform over the finite set \mathcal{S}' . Each scenario $S \in \mathcal{S}'$ consists of a traffic and a capacity scenario from sets \mathcal{F}' and \mathcal{K}' . 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. The capacity scenario \mathcal{K}' is built by sampling airspace capacities from a given distribution.

We keep \mathcal{S} fixed across all simulations and all policies, i.e. the decision policies are all based on the same collection of scenarios. To model the effect of imperfect forecasts, we consider instances where the expectation is formed over a proper subset $\mathcal{S} \subsetneq \mathcal{S}'$, i.e. we do not anticipate all scenarios that may happen. To reflect the impact of an increasingly poor forecast, we can increase the cardinality of \mathcal{S}' relative to that of \mathcal{S} . To contrast this with having full knowledge in stage 1 about the true distribution of traffic scenarios that we may face, we also look at instances where $\mathcal{S} = \mathcal{S}'$. We emphasize that the number of 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. We consider the cases of $\mathcal{S} = \mathcal{S}'_0$, and $\mathcal{S} \subsetneq \mathcal{S}'_1$, where \mathcal{S}_0 features 100 scenarios and \mathcal{S}'_0 and \mathcal{S}'_1 are two different collections of ground truth scenarios. The first is the same as \mathcal{S} , and \mathcal{S}'_1 has additional 300 scenarios. Similarly, sets \mathcal{F}'_1 and \mathcal{K}'_1 used to build \mathcal{S}'_1 have additional 300 traffic and capacity scenarios.

Given the decision \mathbf{h} from stage 1, we assess in stage 2 the optimal displacement cost $G(\hat{S}, \mathbf{h})$. To do so, we use CPLEX to solve the binary program $G(\hat{S}|\mathbf{h})$ with capacity constraints (2) replaced with (8). As a result, we obtain the total cost incurred, whether $G(\hat{S}|\mathbf{h})$ 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 600 times to measure the average performance in terms of displacement costs, capacity costs and flights exceeding capacity constraints.

A small artificial network and a real-life case study are used for the computational analysis. The small network consists of four airspaces (T, U, R and S) with two elementary sectors each and a fifth airspace (Q) with three elementary sectors (Figure 1). A time horizon of two hours is considered with a total of 120 scheduled flights. A pool of 80 non-scheduled flights is built to generate the traffic scenarios. Flights are scheduled to fly across one of the 8 origin-destination pairs denoted by the arrows. Each flight has a preferred route using the central airspace Q and several more expensive routes avoiding Q or assigning delay up to 30 minutes. As for the capacity uncertainty, we assume each airspace to have the same distribution as the MUAC ANSP described in the appendix (see EDYYBUTA, Table 9). More details about this network can be found in (Starita et al. 2017).

A second larger case study is built using real data. The capacity and demand data are obtained via Eurocontrol Demand Data Repository (DDR2) service using the Eurocontrol Network Strategic Tool (NEST). 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

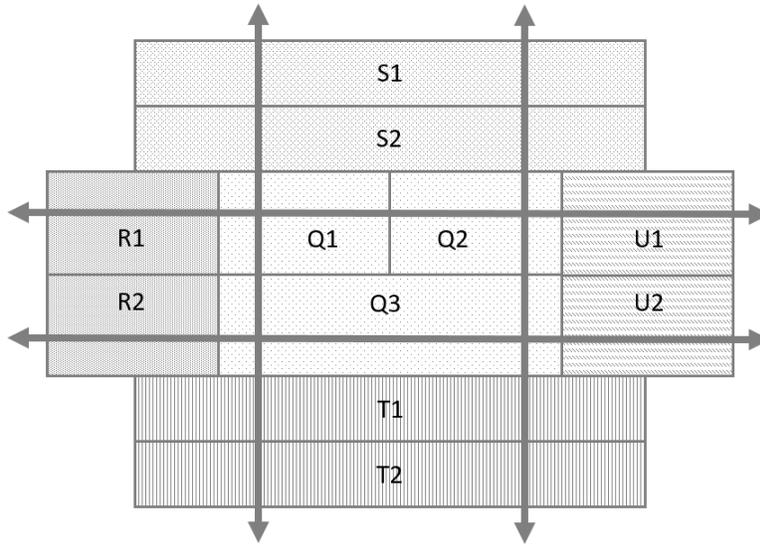


Figure 1: Small artificial network.

in Figure 2. 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 numbers of sectors which were most frequently used: in total, we have 173 different configurations for the 15 ACCs/sector groups shown in Table 7 in the appendix. The overview of elementary and collapsed sectors used in the simulation is likewise in the appendix, namely Table 8.

Delay and re-routing costs are obtained by drawing on Cook and Tanner (2015) and Eurocontrol (2018d). Delay costs are calculated per aircraft type and per duration of delay (non-linear with time). To explain why we use averaged estimated cost of demand management (DM) measures per aircraft type instead of real (estimated) cost for each flight, it is necessary to recall the timeline in capacity and demand management decision-making. Capacity decisions are made weeks and months in advance, where a number of determinants for DM cost calculations (e.g. number of transfer passengers and their connection times) are not known with certainty even to the airlines. One would have to rely on available historical (empirical) distribution for load-factors, number of transfer passengers and a number of connections lost due to different levels of primary delays on a flight basis to calculate passenger-related costs only. The problem is that this data is not available and is considered sensitive business information. Note that the report Cook and Tanner (2015)—that we use as a basis to calculate DM costs—already includes some historical/empirical findings to estimate passenger-related delay costs. Consequently, airlines are the only stakeholders with full information to estimate (their) real cost of delay per flight for a given day of operations. However, using airlines' information for individual flights' cost of delay may lead to 'gaming the system' and raise equity questions: without additional rules and a fair mechanism, airlines might declare higher-than-real costs for delaying/re-routing their flights, since the NM aims to minimize total cost of

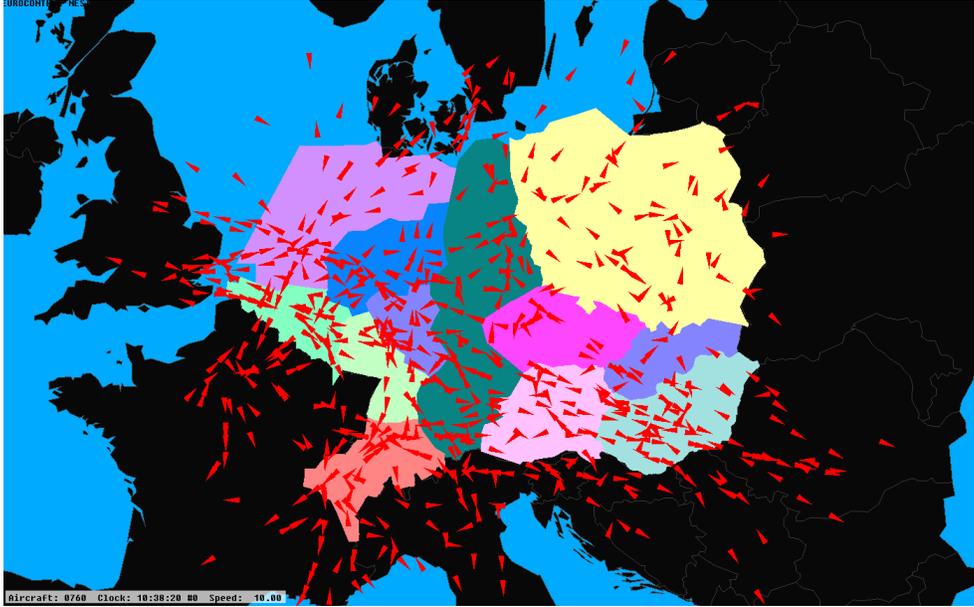


Figure 2: Snapshot of scheduled flights and airspaces considered in the simulation at a fixed point in time.

capacity and demand management actions. This would add another layer of complexity to the model and draw attention from the main focus of the study.

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 the 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. To create non-scheduled flights, we select those out of all historical non-scheduled flights on that day 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.

In practice, capacities planned at strategic level are often not provided on the day of operations due to many reasons: adverse weather conditions, military activity, lower staffing levels than planned, etc. Some of these reasons exhibit some regularities and could be anticipated well in advance (e.g. existing staffing issues in Karlsruhe airspace, see Eurocontrol (2018c)), while some other, like weather, are not easy to predict (especially the impact on capacity). In this paper, we do not differentiate between different sources of capacity provision uncertainty, but rather focus on the outcome such events might have in each ACC. We partially rely on historical ATFCM regulation data to come up with a uniform probability distribution over sectors that a random sector in an ACC will have capacity reduction by 10%, 30%, or 50% or will be closed for flights entirely (zero-rate entry count). We provide the probability distribution in the appendix (Table 9). Capacity scenarios are built by sampling from the given distribution while enforcing that only up to one elementary sector per scenario (and all the collapsed sectors including it) can provide less than nominal capacity. This assumption is to avoid building overly pessimistic scenarios with several uncorrelated capacity reductions.

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 performance of the policies in the two case studies obtained by simulation and under the accurate forecast assumption (i.e., $\mathcal{S}'_0 = \mathcal{S}$). Table 4 shows results on the real network under the imperfect forecast assumption, i.e. the flight materializations and capacity scenarios are drawn from a wider set of traffic and capacity scenarios containing 300 scenarios in addition to those in \mathcal{S} (hence $\mathcal{S}'_1 \supseteq \mathcal{S}$). 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. Results for total costs are shown along with their respective 95% confidence interval.

Results on the small case study show that ϵ policies’ cost performances are close to the cost obtained by the full enumeration policy. In fact, $\epsilon = 20\%$ and $\epsilon = 10\%$ result in the exactly same capacity budget. The relatively high percentage of flights not being assigned is due to the limited menu of routes available in the small case study. For example, all the route options for flights on the north-south flow use sectors S1, S2, T1 and T2. This means that under scenarios where one of these sectors has a capacity shortage, there are no other options than assigning a share of flights to the dummy route.

All policies show consistent behavior across the different simulation scenarios. Policy *AV* results in the minimum capacity ordering. In contrast, the displacement cost returned by this policy is the highest, with an average of 1% flights not being assigned. The best policy in terms of overall cost is $\epsilon = 5\%$. However, we also notice that uncertainty on the time and location of non-scheduled flights and capacity scenarios does not have a significant impact on

Table 2: Simulation results on small network under full knowledge of scenario distribution. Exactly the expected number of non-scheduled flights and capacity materialize (sampled from $\mathcal{S}'_0 = \mathcal{S}$). Small Network. 600 simulation runs.

Policy	Capacity cost	Displacement cost	Total cost	Flights not assigned (%)
<i>FE</i>	29,380	27,153	56,533 \pm 3,718	6.23%
<i>AV</i>	26,240	31,521	57,761 \pm 3,940	6.72%
$\epsilon = 20\%$	29,380	27,502	56,533 \pm 3,718	6.23%
$\epsilon = 10\%$	29,380	27,502	56,533 \pm 3,718	6.23%
$\epsilon = 5\%$	31,800	26,595	58,395 \pm 3,622	6.08%
$\epsilon = 1\%$	32,720	26,595	59,315 \pm 3,622	6.08%
<i>SA</i>	28,460	29,683	58,143 \pm 3,652	6.24%

Table 3: Simulation results on real network under full knowledge of scenario distribution. Exactly the expected number of non-scheduled flights and capacity materialize (sampled from $\mathcal{S}'_0 = \mathcal{S}$). 600 simulation runs.

Policy	Capacity cost	Displacement cost	Total cost	Flights not assigned (%)
<i>AV</i>	43,932	60,959	104,891 \pm 6,689	1.00%
$\epsilon = 20\%$	49,461	44,860	94,321 \pm 5,764	0.83%
$\epsilon = 10\%$	50,350	40,153	90,503 \pm 5,209	0.74%
$\epsilon = 5\%$	53,017	33,804	86,821 \pm 4,427	0.63%
$\epsilon = 1\%$	59,462	30,190	89,652 \pm 4,163	0.57%
<i>SA</i>	49,951	44,796	94,747 \pm 5,676	0.81%

Table 4: Simulation results on real network with somewhat incomplete knowledge of scenario distribution. Exactly the expected number of non-scheduled flights and capacity materialize (sampled from $\mathcal{S}'_1 \supsetneq \mathcal{F}$). 600 simulation runs.

Policy	Capacity cost	Displacement cost	Total cost	Flights not assigned (%)
<i>AV</i>	43,932	66,793	110,725 \pm 7,069	1.14%
$\epsilon = 20\%$	49,461	48,606	98,067 \pm 5,960	0.91%
$\epsilon = 10\%$	50,350	43,272	93,622 \pm 5,381	0.82%
$\epsilon = 5\%$	53,017	38,026	91,043 \pm 4,784	0.72%
$\epsilon = 1\%$	59,462	32,803	92,265 \pm 4,385	0.63%
<i>SA</i>	49,951	47,931	97,882 \pm 5,835	0.88%

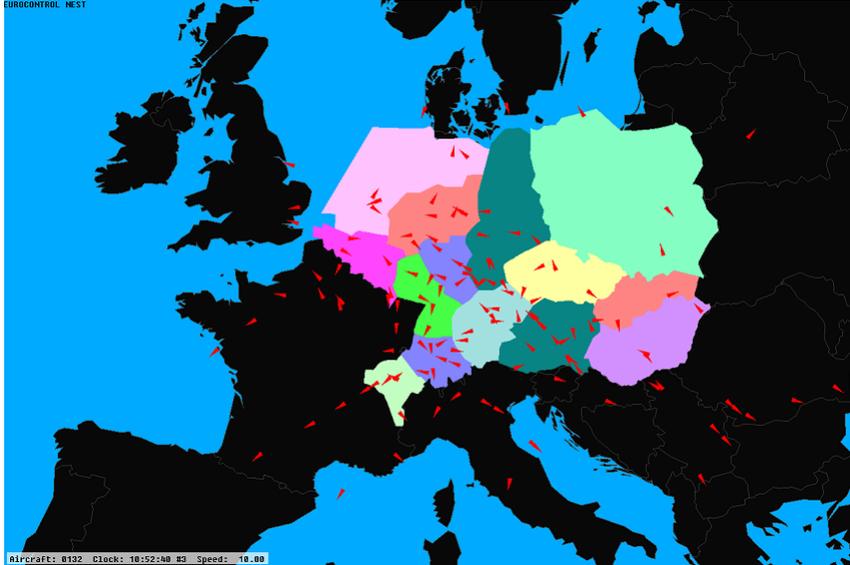


Figure 3: Snapshot of non-scheduled flights at a fixed point in time. This snapshot shows non-scheduled traffic under the low uncertainty scenario, i.e. only 160 flights.

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 3: due to the relatively low volume of flights, the non-scheduled traffic is fairly widely distributed so that its impact on individual sectors is small.

However, if we additionally increase uncertainty over the number of non-scheduled flights and capacity scenarios, results change as can be seen in Table 5. 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 5: Simulation results on real network with highly incomplete knowledge of scenario distribution. Twice the expected number of non-scheduled flights materialize. Scenarios sampled from wider collection $\mathcal{S}'_1 \not\subseteq \mathcal{F}$. 200 simulation runs.

Policy	Capacity cost	Displacement cost	Total cost	Flights not assigned (%)
<i>AV</i>	43,932	123,965	167,897 ± 14,505	1.83%
$\epsilon = 20\%$	49,461	73,767	123,228 ± 12,365	1.05%
$\epsilon = 10\%$	50,350	66,685	117,035 ± 11,067	0.93%
$\epsilon = 5\%$	53,017	56,711	109,728 ± 9,744	0.79%
$\epsilon = 1\%$	59,462	45,931	105,393 ± 9,038	0.68%
<i>SA</i>	49,951	72,028	121,979 ± 11,365	1.00%

The added uncertainty in traffic volume in the network has a clear impact on displacement cost, which increases across all policies. The perhaps most interesting result is that now $\epsilon = 1\%$ 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. As expected, policy *AV* still performs worse than all the other policies. 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 more risk-adverse ϵ policies. If the number of non-scheduled flights tends to be fairly stable, less conservative ϵ policies can be selected. Results also consistently show that *AV* policy fails to provide acceptable displacement costs. The sampling policy *SA* was dominated by the $\epsilon = 10\%$ policy in all scenarios.

In Table 6 we compare the computing time of solving $G(S)$ for a given 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 $G(S)$ to optimality. A single scenario S is sampled from \mathcal{S} 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). The number of decision variables and constraints used by $G(S)$ and iterations required by *PH* are reported as well.

Table 6: Average computational times to solve an instance $G(S)$. Note that the times for the perfect foresight approach are averaged over the samples $F \in \mathcal{F}$.

Number of flights	CPLEX time (s)	Variables	Constraints	Average <i>PF</i> time (s)	<i>PF</i> iterations
1,000	511	11,225	2,693	2.5	15
1,100	1,755	12,313	2,793	3.8	21
1,200	4,655	13,417	2,893	5.9	30
1,300	13,363	14,570	2,993	6.6	32
1,400	n.a.	15,654	3,093	7.2	31

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 the proposed decomposition approach.

6 Conclusion

Inadequate capacity planning in ATM is a real issue with increasingly costly implications upon airspace users. This is probably best condensed in a recent statement of Airlines for Europe (A4E), the association of 15 leading European airlines, accounting together for more than 70% of European passenger journeys: “Nevertheless, when broken down, it is argued that in many cases that much of this [‘unacceptable increase in delay and disruption to the travelling public’ in 2018] was due to sub-optimal staffing both in terms of core numbers and in rostering practices, which resulted in shortfalls in required sector openings to best manage the demand” (NM 2018). With this in mind, knowing that estimated cost of en-route ATFM delay amounted to 1.93 billion EUR (Eurocontrol 2018c), it is clear that any improvement in this respect might have significant positive implications on quality of service provided by the ANSPs by means of decreasing unnecessary delay (and associated costs) imposed on airspace users.

Capacity planning for air traffic navigation service provision is a challenging problem due to uncertainty in traffic and capacity supply 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 very efficient since the decomposition exploits the problem structure in a way such that the resulting sub-problem becomes very easy to solve (specifically by separating decisions on demand management from decisions on sector opening schemes). 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 is evaluated on a small artificial network and on a case study built with 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, we demonstrate that it can be used to test for robustness against disruptions of capacity provision.

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, airport congestion is not explicitly taken into consideration and it is indeed a highly relevant aspect in ATFM research and potentially influential on results. However, accounting for such congestion requires including associated terminal airspace, which connects en-route airspace with airports. This is beyond the scope of this manuscript. Future work in this direction is already under development. Regarding delay propagation, it is partially captured by means of delay costs, i.e. by also taking into account the costs of reactionary instead of solely primary delays. An interesting future research development is to explicitly model delay propagation by including airports and turnaround process.

Acknowledgments

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 7: 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	Skyguide	LSAGUTA	Genève ACC Upper FL 245-999
		LSAZUTA	Zurich (Dübendorf) ACC Upper FL 245-999
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	Hungarocontrol	LHCCCTA	Budapest ACC
Poland	PANSA	EPWWCTA	Warszawa ACC

Table 8: 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 different 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

Table 9: Probability distribution of capacity provision

Area Control	Nominal conditions	10% cap. reduction	30% cap. Reduction	50% cap. Reduction	Zero
EDUUUTAC	0.5	0.2	0.2	0.1	0
EDUUUTAE	0.5	0.2	0.2	0.1	0
EDUUUTAS	0.5	0.2	0.2	0.1	0
EDUUUTAW	0.5	0.2	0.2	0.1	0
EDYYBUTA	0.7	0.05	0.05	0.1	0.1
EDYYDUTA	0.7	0.05	0.05	0.1	0.1
EDYYHUTA	0.7	0.05	0.05	0.1	0.1
EPWWCTA	0.75	0.1	0.05	0.05	0.05
LHCCCTA	0.9	0.05	0.03	0.02	0
LKAACTA	0.95	0.02	0.02	0.01	0
LKAAUTA	0.95	0.02	0.02	0.01	0
LOVVCTA	0.9	0.05	0.03	0.02	0
LSAGUTA	0.9	0.05	0.03	0.02	0
LSAZUTA	0.9	0.05	0.03	0.02	0
LZBBCTA	0.95	0.02	0.02	0.01	0

References

- Baumgartner M, 2007 *European Air Traffic Management : principles, practice and research*. The editor (A. Cook), chapter The organisation and operation of European airspace, 1–34 (Ashgate Publishing Limited).
- Bertsimas D, Patterson SS, 1998 *The air traffic flow management problem with enroute capacities*. *Operations research* 46(3):406–422.
- Castelli L, Pesenti R, Ranieri A, 2011 *The design of a market mechanism to allocate air traffic flow management slots*. *Transportation research part C: Emerging technologies* 19(5):931–943.
- Conniss R, Curtois T, Petrovic S, Burke E, 2014 *Scheduling Air Traffic Controllers*. *10th International Conference of the Practice and Theory of Automated Timetabling* (York, United Kingdom).
- Cook A, Tanner G, 2015 *European airline delay cost reference values*. Technical report, University of Westminster.
- Corolli L, Castelli L, Lulli G, 2010 *The air traffic flow management problem with time windows*. *4th International Conference on Research in Air Transportation (ICRAT), Budapest, Hungary*.
- Delgado L, Cook AJ, Cristóbal S, 2015 *Controller Time and Delay Costs - a trade-off analysis*. *Fifth SESAR Innovation Days* (Bologna).
- Eurocontrol, 2013 *Capacity assessment and planning guidance document, Ed. 2.8*. Technical report, EUROCONTROL, URL <http://www.eurocontrol.int/sites/default/files/content/documents/nm/airspace/capacity-assessment-planning-guidance-doc-v2-8final.pdf>.
- Eurocontrol, 2018a *Annex 5: ACC traffic forecast and capacity plans*. *European Network Operations Plan 2018-2019/22*.
- Eurocontrol, 2018b *ATM cost-effectiveness (ACE) 2016 benchmarking report with 2017-2021 outlook*. Technical report, EUROCONTROL.
- Eurocontrol, 2018c *Performance review report covering the calendar year 2018 (pr 2018)*. Technical report, EUROCONTROL.
- Eurocontrol, 2018d *Standard inputs for eurocontrol cost-benefit analyses*. Technical report, EUROCONTROL.
- European Commission, 2019 *Commission regulation (EU) no 2019/123*. Official Journal of the European Union.
- Flener P, Pearson J, 2013 *Automatic airspace sectorisation : A survey*. Technical report, URL <http://uu.diva-portal.org/smash/record.jsf?pid=diva2%3A675958&dswid=-1083>.
- Gabrel V, Murat C, Thiele A, 2014 *Recent advances in robust optimization: an overview*. *European Journal of Operational Research* 235(3):471–483.
- Geoffrion A, 1974 *Lagrangian relaxation for integer programming*. *Mathematical Programming Study* 2:82–114.
- IATA, 2015 *Airline cost management group*. URL https://www.iata.org/whatwedo/workgroups/Documents/ACMG/ACMG_Flyer_FY2014_2015.pdf.
- Josefsson B, Polishchuk T, Polishchuk V, Schmidt C, 2017 *A Step Towards Remote Tower Center Deployment: Optimizing Staff Schedules*. *Twelfth USA/Europe Air Traffic Management Research and Development Seminar (ATM2017)*, URL http://www.atmseminarus.org/seminarContent/seminar12/papers/12th_ATM_RD_Seminar_paper_70.pdf.
- Lau A, Berling J, Linke F, Gollnick V, Nachtigall K, 2015 *Large-scale network slot allocation with dynamic time horizons*. *Eleventh USA/Europe Air Traffic Management Research and Development Seminar*.
- Lulli G, Odoni A, 2007 *The European air traffic flow management problem*. *Transportation Science* 41(4):431–443.

- Majumdar A, Ochieng WY, Bentham J, Richards M, 2005 *En-route sector capacity estimation methodologies: An international survey*. *Journal of Air Transport Management* 11(6):375–387, URL <http://dx.doi.org/10.1016/J.JAIRTRAMAN.2005.05.002>.
- NM, 2018 *Annual network operations report 2018*. Technical report, Network Manager.
- NM, 2019 *ATFCM Operations Manual, Ed. 23.0*. Technical report, Network Manager.
- Qin Y, Wang R, Vakharia A, Chen Y, Seref M, 2011 *The newsvendor problem: Review and directions for future research*. *European Journal of Operational Research* 213(2):361–374.
- Ryanair, 2018 *Annual report*. Technical report, Ryanair, URL <https://investor.ryanair.com/wp-content/uploads/2017/07/Ryanair-FY2017-Annual-Report.pdf>.
- Sherali HD, Staats RW, Trani AA, 2003 *An Airspace Planning and Collaborative Decision-Making Model: Part I—Probabilistic Conflicts, Workload, and Equity Considerations*. *Transportation Science* 37(4):434–456, URL <http://dx.doi.org/10.1287/trsc.37.4.434.23272>.
- Sherali HD, Staats RW, Trani AA, 2006 *An Airspace-Planning and Collaborative Decision-Making Model: Part II—Cost Model, Data Considerations, and Computations*. *Transportation Science* 40(2):147–164, URL <http://dx.doi.org/10.1287/trsc.1050.0141>.
- SJU, 2019a *A proposal for the future architecture of the European airspace*. Technical report, SESAR Joint Undertaking.
- SJU, 2019b *SESAR Joint Undertaking - any controller, any airspace*. Technical report, SESAR Joint Undertaking.
- Starita S, Strauss A, Jovanović R, Ivanov N, Pavlović G, Babić O, Fichert F, 2017 *Coordinated capacity and demand management in a redesigned ATM value chain: Strategic network capacity planning under demand uncertainty*. *7th SESAR Innovation Days* (Belgrade).
- Tien SLA, Hoffman R, 2009 *Optimizing Airspace Sectors for Varying Demand Patterns using Multi-Controller Staffing*. *Eighth USA/Europe Air Traffic Management Research and Development Seminar (ATM2009)*.
- Tobaruela AG, 2015 *A Framework to assess the ability of automation to deliver capacity targets in European airspace*. *PhD dissertastion*. Ph.D. thesis, Imperial College, London, URL <https://core.ac.uk/download/pdf/77004484.pdf>.
- Tobaruela G, Arnab M, Ochieng YW, Schuster W, Hendrickx P, 2013 *Enhancing cost-efficiency and reducing capacity shortages: Strategic planning and dynamic shift management*. *Tenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2013)*.
- Tošić V, Babić O, Cangalovic M, Hohlacov D, 1995 *A model to solve en route air traffic flow management problem: A temporal and spatial case*. *Unpublished report, University of Belgrade* .
- Venugopalan TK, Wong CSY, Suresh S, Sundararajan N, 2018 *Simultaneous Optimization of Airway and Sector Design for Air Traffic Management*. *Journal of Air Transportation* 26(1):8–22, URL <http://dx.doi.org/10.2514/1.D0090>.