Buffer scheduling for improving on-time performance and connectivity with a multi-objective simulation–optimization model

A proof of concept for the airline industry

van Schilt, Isabelle M.; van Kalker, Jonna; Lefter, Iulia; Kwakkel, Jan H.; Verbraeck, Alexander

DOI
10.1016/j.jairtraman.2024.102547

Publication date
2024

Document Version
Final published version

Published in
Journal of Air Transport Management

Journal of Air Transport Management 115 (2024) 102547

Contents lists available at ScienceDirect

Journal of Air Transport Management

journal homepage: www.elsevier.com/locate/jairtraman

Buffer scheduling for improving on-time performance and connectivity with a multi-objective simulation–optimization model: A proof of concept for the airline industry

Isabelle M. van Schilt [a,*], Jonna van Kalker [b], Iulia Lefter [a], Jan H. Kwakkel [a], Alexander Verbraeck [a]

[a] Faculty of Technology, Policy and Management, Delft University of Technology, Jaffalaan 5, 2628 BX Delft, The Netherlands
[b] KLM Operations Decision Support, Piet Guilonardweg 15, 1182 DB Schiphol Airport, The Netherlands

ARTICLE INFO

Keywords:
Airline scheduling
Buffer allocation
On-time performance
Connectivity
Multi-objective optimization

ABSTRACT

Schedule design in the transportation and logistics sector is a widely studied problem. Transport service providers, such as the train industry and aviation, aim for schedules to be on-time according to the planning (i.e., on-time performance or OTP) in order to increase the service level by ensuring that passengers actually make their connections and to reduce costs. Transportation services also aim for schedules that serve a high variety of destinations and frequency of connections (i.e., connectivity). OTP and connectivity are both highly dependent on buffer time: more lucrative connections can often be offered by reducing the buffer time in the schedule, while more delay can be absorbed by more buffer time. Given strict constraints on the minimum turnaround time of aircraft and minimum (and maximum acceptable) transfer times of passengers, assigning buffer time in an already tightly planned schedule to optimize OTP and connectivity simultaneously is a big challenge. This research presents a novel multi-objective formulation of a daily flight schedule where buffer scheduling is used to ensure the optimal balance between OTP of the schedule and the passenger connections as connectivity, given the tight restrictions. This problem formulation is solved using a simulation–optimization framework. Specifically, we use the Multi-Objective Evolutionary Algorithm (MOEA) BORG. As a proof of concept, a daily European flight schedule of a large international airline is optimized on both OTP and connectivity. The results demonstrate that the presented multi-objective formulation and associated solving through simulation–optimization can result in candidate schedules with both better on-time performance and a higher connectivity.

1. Introduction

Schedule design is a widely studied problem in the transportation and logistics sector. Designing a reliable schedule for transportation services is one of the biggest challenges. A reliable schedule means that the schedule on the day of operation is on-time according to the pre-defined schedule (i.e., on-time performance or OTP) to ensure a high service level for the passengers. At the same time, transportation services, such as rail and aviation, aim for schedules that serve a high variety of destinations and a high frequency of connections (i.e., connectivity) (L'upták et al., 2019). Especially for the aviation sector, scheduling is one of the most challenging and important operations. According to Barnhart and Cohn (2004) and Wu (2006), there are four core problems in airline schedule planning:

1. *Schedule design:* Determine the markets to serve, at what frequency, and how to schedule the flights.
2. *Fleet assignment:* Assign the aircraft to each flight.
3. *Aircraft maintenance routing:* Route the aircraft such that the maintenance requirements are satisfied.
4. *Crew scheduling:* Assign the crew to the flights.

Ideally, the four core problems are solved simultaneously when creating an airline schedule. However, due to the complexity of the scheduling problem, most research focuses on either integrating two core problems or extending one problem (Barnhart and Cohn, 2004; Ageeva, 2000; Şafak et al., 2017; Achenbach and Spinler, 2018). Our paper will also discuss extending one core problem, namely that of schedule design. Creating a schedule in terms of connectivity and OTP takes place in the schedule design phase; the planning of the flights is determined here. Research has shown that the use of buffers improves

* Corresponding author.
E-mail address: i.m.vanschilt@tudelft.nl (I.M. van Schilt).

https://doi.org/10.1016/j.jairtraman.2024.102547
Received 21 June 2022; Received in revised form 24 April 2023; Accepted 28 January 2024
Available online 7 February 2024

the on-time performance of a flight schedule (Wu, 2005; Baumgarten et al., 2014; Ahmadbeygi et al., 2010; Wu, 2006; Fricke and Schultz, 2009). Buffers are allocated to compensate for possible delays that might occur on the day of operation. This research focuses on the use of buffers in the schedule design.

Flight delays are an unavoidable fact and a key concern in aviation. Their impact is both economic and environmental (Sternberg et al., 2017). Delays (for example due to bad weather, technical problems, and crew unavailability) are costly for both the airlines and the passengers (Wu, 2008; Peterson et al., 2013). In particular for hub-and-spoke networks (i.e., locations called spokes connected through an intermediary location called a hub), delays are problematic due to the high intensity of connecting flights, passengers, and crews (Achenbach and Spinler, 2018; Hansen et al., 2001). Therefore, it would be valuable for airlines to incorporate the probability of delays in the design of the flight schedule (Thengvall et al., 2000). In this manner, the flight planning is less likely to be negatively affected by delays on the day of operation (Lee et al., 2007).

However, a more reliable flight schedule often requires more resources or more time, making it more costly for airlines. This leads to a trade-off between OTP and costs (Clausen et al., 2010). For hub-and-spoke airlines, costs are primarily determined by the connecting passengers at the hub airports. An example of this trade-off is when an extra five minutes is added to the ground time of a flight to increase on-time performance in a schedule design. If five minutes is added to the expected arrival time and the flight time does not change, the chances of the flight being on-time is higher, i.e., the on-time performance has increased. However, this shortens the transfer time for connecting flights in the schedule, that could result in not being able to offer certain (profitable) connections anymore. Thus, OTP and connectivity are both highly impacted by buffer time; a higher OTP asks for more buffer time whereas a higher connectivity (i.e., more connecting flights) requires less buffer time, as this shortens transit time and makes connections unfeasible. This research focuses on how to add buffers such that it is most beneficial, both in terms of cost and on-time performance in schedule design. In literature and in industry, optimization tools that analyze one aspect of this complex trade-off have been developed, either maximizing the reliability, or minimizing the cost (Ageeva, 2000; Şafak et al., 2017). This means that first one of the two aspects is optimized, and then the other, making it difficult to actually balance OTP and connectivity (Wong and Tsai, 2012).

The challenge for creating optimal schedules that have the best trade-off between OTP and connectivity becomes even more complex due to tight restrictions. The primary restrictions in airline schedule design is the minimum and maximum acceptable transfer time for passengers on connecting flights (i.e., the time that the passengers has available to transfer from one flight to another flight), and the minimum turnaround time of the aircraft (i.e., the time between landing and taking off). Furthermore, it is also important to consider factors such as slot constraints, crew legalization and resource assignment, which further complicate the allocation of buffers in the schedule. As many airline schedules are already tightly planned to make optimal use of the aircraft, it is a big challenge to assign buffer time to optimize for both OTP and connectivity simultaneously, given these tight constraints.

In this research, OTP and connectivity are optimized simultaneously when designing a schedule since they are both heavily impacted by buffer time. The goal of this research is to investigate how a schedule can be created that balances OTP and connectivity of daily flights by means of buffer scheduling. Revealing the trade-off between OTP and connectivity supports decision makers in creating daily flight schedules. The contribution of this research is to present a novel multi-objective formulation of the schedule design problem where there is an explicit trade-off between OTP and connectivity of the network without the use of aggregation (i.e., without combining the objectives to a single objective using weights). We use a simulation–optimization framework as this has been proven to be a fruitful approach for handling multi-objective optimization for complex problems in other sectors such as inventory management (Tsai and Chen, 2017).

The remainder of this paper is structured as follows. Section 2 presents the related work using the current state-of-the-art literature. Section 3 gives a description of the problem, case study, and scope. Section 4 describes the mathematical formulation of the problem and the model. Section 5 presents the solution approach for this study. Section 6 presents the results of the model performance on a real world case from a large international airline carrier. The paper concludes with a discussion in Section 7, and conclusions and recommendations for further research in Section 8.

2. Related work

The literature on related work is studied on three main topics, namely airline delay management, solving schedule design with a multi-objective optimization, and the use of a simulation–optimization framework.

2.1. Airline delay management

Much research has been performed on airline delay management for the day of operation. Santos et al. (2017) presents a linear programming approach to solve the daily airline delay management problem with capacity constraints and to make decisions on the spot. Jarrah et al. (1993) creates a decision support framework for flight delays and cancellations during the day of operation. Instead of solving delays during the daily operations, delays could also be prevented. Montlaur and Delgado (2017) shows optimization techniques to minimize the flight and passengers delay by including or excluding reactionary delays. Sternberg et al. (2017) shows a review of approaches to predict the flight delays and how machine learning relates to this. In their paper, there is a distinction between root delay (i.e., local delays), or cancellations and delay propagation (i.e., a delay in a flight causes delays in the subsequent flights). Propagated delays are mostly caused by the connected resources within the airline schedule such as the aircraft, crew, passenger, and airport resources (Kafle and Zou, 2016). Propagated delay is a well-known phenomenon in aviation and its impact has been researched extensively. Kondo (2011) compares the impact of propagated delays between hub-and-spoke and point-to-point airports. Churchill et al. (2010) examines the effect of propagated delays on the daily planning. Moreover, Qin et al. (2019) investigates how to optimize the delay propagation in a Chinese aviation network by rescheduling flights. Thus, delay propagation is a crucial element for airline scheduling, especially in hub-and-spoke networks (Achenbach and Spinler, 2018). Next to this, delays can occur either on the ground, i.e., ground delay, or in the air, i.e., en-route delay (Carlier et al., 2007). Ground delay can be defined as the delay during the turnaround of the airplane. En-route delay can be defined as the delay of an airplane between departure (off-blocks) and arrival (on-blocks); this contains taxi and airborne time. En-route delay is also known as block delay (Fricke and Schultz, 2009). This research focuses on allocating buffers to the ground time given the arrival delay of airplanes, as a result of en-route delay.

2.2. Schedule design by multi-objective optimization

Creating a schedule where OTP and connectivity are balanced by means of buffer scheduling is of a multi-objective nature. On the one hand, we want to maximize OTP (i.e., minimize delay) and on the other hand, we want to maximize connectivity. A commonly applied approach is to scalarize the two objectives and formulate a single-objective problem (Deb, 2014). This results in a single optimal solution, thus a single optimal candidate schedule. For this research, there are two main disadvantages to the single-objective approach namely (i) accurately converting OTP and connectivity to the same unit is currently impossible, and (ii) supporting decision making by giving insight in the trade-off between OTP and connectivity, i.e., balancing two conflicting objectives, cannot be fully obtained by only providing a single optimal solution (Kasprzyk et al., 2016). In the same line, Arrow's Impossibility Theorem implies that using aggregated objectives to get one single objective for optimization inadvertently dictates the properties of the optimized candidates solutions in unpredictable ways (Kasprzyk et al., 2016). Thus, optimizing a schedule design that is multi-objective in nature, should be handled as a multi-objective problem where a trade-off exists between the objectives without aggregation.

A literature study revealed that only a few research studies used multi-objective optimization to solve schedule design problems for airlines. Lee et al. (2007) formulates the on-time performance and operation costs as a multi-objective problem in order to improve the robustness of the flight schedule by re-timing the departure times. Burke et al. (2010) uses multi-objective optimization to improve the reliability, i.e., on-time performance, and flexibility of the flight schedule. Katsigiannis et al. (2021) presents a multi-objective optimization that investigates the trade-off between slot allocation and the airport schedule as a whole. These studies show that solving a schedule design with multi-objective optimization is promising. However, there is still limited literature that applies this approach while it has many advantages compared to aggregating objectives. Therefore, this paper contributes to the use of multi-objective optimization with two main conflicting objectives of a schedule design, namely OTP and connectivity.

Research on the integration of OTP and connectivity is limited. Dunbar et al. (2012) presents a new approach on how to integrate aircraft routing and crew scheduling to minimize propagated delay. They focus on the delays caused by missed connection of the crew. Jacquillat and Vaze (2017) designs and assesses a novel approach for scheduling the air traffic congestion of an airport. They want to increase OTP, i.e., mitigate the air traffic congestion at the airport, with network connectivity as a constraint. However, it has not been investigated yet how to handle the trade-off between connecting passengers and OTP in the flight schedule of an airline.

Table 1
Table of notation.

Sets	
AC	Set of aircraft types, $ac \in AC$
K	Set of connection types, $k \in K$
L	Set of fleetlines in the schedule, $l_v \in L$
F	Set of flights in the flight schedule ξ, $F = F^{IN} \cup F^{OUT}$
F^{IN}	Set of inbound flights in the flight schedule ξ
F^{OUT}	Set of outbound flights in the flight schedule ξ
R	Set of all rotations in the flight schedule ξ
S_n	Set of flights over multiple days with the same flight call sign n
Objectives	
C_j^{IN}	Inbound connectivity revenue in euros per rotation
C_i^{OUT}	Outbound connectivity revenue in euros per rotation
$C(\xi)$	Connectivity revenue of the flight schedule ξ in euros
$OTP(\xi)$	Arrival on-time performance of the flight schedule ξ in minutes
Decision variable	
$X_{r_{i,j}}$	Buffer time for preceding rotation $r_{i,j}$ in minutes
Parameters	
ac	Aircraft type subscript
ata_i	Actual time of arrival of inbound flight f_i in minutes
f_i^{IN}	Inbound flight
f_j^{OUT}	Outbound flight
l_v	Fleetline with subscript v as index
n	Flight call sign
$p_k(t)$	Probability that a passenger with connection type k will actually be boarded on the transferring flight as a function of time t
$r_{i,j}$	Rotation with inbound flight f_i and outbound flight f_j
sta_i	Scheduled time of arrival of inbound flight f_i in minutes
std_j	Scheduled time of departure of outbound flight f_j in minutes
$tp_{i,j}$	Expected number of passengers connecting from flight f_i to flight f_j
$tt_{i,j}$	Transfer time between flight f_i and f_j in minutes
$c_{i,j}$	Connectivity revenue of one passenger connecting from flight f_i to flight f_j
$MACT_k$	Maximum acceptable connecting time for connection type $k \in K$ in minutes
MCT_k	Minimum connecting time for connection type $k \in K$ in minutes
MTT_{ac}	Minimum turnaround time for aircraft type $ac \in AC$ in minutes
N_v	Number of rotations in fleetline l_v
T_{l_v}	Total buffer time for a fleetline l_v in minutes
W	Time window for period to calculate OTP
ξ	Flight schedule as a tuple {L, R, MTT, X} for all rotations, for all aircraft types ac, and for all buffer times
τ	Unit of buffer time in minutes

2.3. Simulation–optimization framework

A simulation–optimization framework is used as an approach to solve the complex and multi-objective problem of this research. For optimization of complex systems, two strategies exist. One is to simplify the system's behavior as constraints using the optimization formalism, the other is to use a simulation model of that complex system in the optimization (Andradóttir, 1998). Especially for systems with a stochastic and non-linear behavior over time, simulation models can better represent the actual characteristics over time of the system that is studied. For simulation-based optimization, the goal function of the optimization is expressed as a function of the output variables of the simulation model (single-objective or multi-objective). The optimization algorithm carries out multiple runs of the simulation model, each time setting the input parameters, to figure out which combination of input variables leads to the best outcome of the goal function (Riley, 2013). Thereby, each run of the simulation model generates one data point for the optimization. The main advantages of using a simulation model over coding the system's behavior as constraints are that the simulation model can express complex, non-linear behavior, and that the simulation model can be separately verified and validated for its correct behavior. The optimization of simulation models is also a good option when a simulation model of a complex system already exists, and the user is interested to find those parameter settings for the simulation model that optimize its output variable(s). Main issues are the usually long runtimes as well as the stochastic nature and the non-linear behavior of simulation models. Since most discrete-event and agent-based simulation models are stochastic models, a single run does not express the true outcome of the model. Multiple replications of the model with different random seeds are necessary to estimate the true average of the model's performance indicators with a tight enough confidence interval. Combined with an already long runtime, this can form a blocker for the usage of optimization in combination with simulation. Because of the non-linear behavior of the simulation model, techniques that are specifically taking into account the characteristics of stochastic simulations typically yield better results for this type of optimization than classical optimization techniques (Fu, 2015). In a real-life airline system, OTP is primarily affected by flight delays with a complex, unpredictable, and stochastic character. A simulation model would be suitable to represent the actual characteristics of the airline schedule over time in the system. This makes the combination of simulation and optimization useful for our research. Therefore, our research uses the simulation–optimization framework for analyzing the schedule design.

2.4. Research gaps

The review of the related work shows that there is no research, to our knowledge, that investigates the trade-off between OTP and connectivity for airline schedule design. Most of the previous work focused on optimizing a schedule on OTP with connectivity as a constraint. Here, the impact that buffer time has on OTP and connectivity simultaneously is neglected. Therefore, our research focuses on optimizing both objectives in the schedule design problem. Moreover, the schedule design problem is solved using multi-objective optimization instead of a single-objective optimization as used in most studies. Literature on multi-objective optimization for schedule design is limited though very promising. Thus, our research contributes to the literature on the use of this technique.

Although there are many airline delay types, we explicitly focus on allocating a buffer to the ground time and not to the en-route time. Due to the complex characteristics of the airline system and delays that primarily impact OTP in the schedule design, a simulation–optimization framework is used in this research.

Concluding, our study tackles two research gaps, namely (1) no research, to our knowledge, has combined connectivity and on-time performance in the schedule design problem, and (2) almost no papers formulated the schedule design problem as a multi-objective problem, where a trade-off exists between the objectives without aggregation.

3. Problem description

This section gives a detailed description of the problem addressed in this research. Section 3.1 explains how a typical flight schedule is constructed. Section 3.2 introduces the case study at a large international carrier. Section 3.3 presents the scope of the problem. Table 1 provides an overview of the mathematical notations.

3.1. Flight schedule

The flight schedule for a season is created by the airline's network department. This department determines the destinations, the important connecting flights, the frequency of the flights, and the aircraft types, and therewith the profitability of the schedule. They are also responsible for creating an operationally feasible schedule by taking into account constraints such as slots and crew (Barnhart and Cohn, 2004).

A typical flight schedule consists of fleetlines L and rotations R, as well as the minimum transfer time of all aircraft types and buffer time between each rotation in a fleetline. A fleetline is the sequence of flights scheduled to be performed by one aircraft in a schedule period. A fleetline consists of $l_v = \{r_{i_{v,1},j_{v,1}}, r_{i_{v,2},j_{v,2}}, r_{i_{v,3},j_{v,3}}, \ldots, r_{i_{v,N_v},j_{v,N_v}}\}$ rotations. One rotation generally consists of two flights from the set of flights F, namely a flight from the hub to an outstation as the outbound flight, $f_j^{OUT} \in F^{OUT} \subset F$, and a flight from an outstation to the hub as the inbound flight, $f_i^{IN} \in F^{IN} \subset F$. Each rotation has an unique pair of flights, presented by $r_{i,j} = \{f_j^{OUT}, f_i^{IN}\}$. The first flight of a day is often only an inbound flight to the hub, thus $\{\emptyset, f_i^{IN}\}$. Similar, the last flight of the day is often only an outbound flight from the hub, so $\{f_j^{OUT}, \emptyset\}$. Note that in the remainder of the paper subscript i for a flight is used for the inbound flight f_i^{IN}, and subscript j is used for the outbound flight f_j^{OUT}. Fig. 1 presents a conceptualization of a typical flight schedule.

Each rotation has a scheduled time of arrival and a scheduled time of departure at the hub. The time between the rotations in one fleetline is the scheduled turnaround time of an aircraft; this is the time between the scheduled time of arrival of rotation $r_{i_{v,n},j_{v,n}}$ and the scheduled time of departure of rotation $r_{i_{v,n+1},j_{v,n+1}}$ in fleetline l_v. Each aircraft type has a minimum turnaround time which is the minimum time it takes to prepare an aircraft for the next rotation. Additional turnaround time scheduled on top of this minimum turnaround time is called buffer time. For example, the minimum turnaround time is 45 min and there is 50 min turnaround time scheduled between rotations 1 and 2, which means that the aircraft has a buffer time of 5 min. Buffer times are used to capture possible delay in order to ensure that an aircraft leaves on-time according to schedule. In this research, flights are not swapped between fleetlines and the sequence of the flights within a fleetline does not change, so buffer time is allocated within each fleetline individually. Given that, rotations can be moved between the time frame of the first arriving flight of the day to the last departing flight of the day.

The flight schedule also determines the flight connections that can be offered, because it determines the transfer time. The transfer time needs to be sufficiently long for a connection to be possible, since passengers need to be able to actually make the transfer. The transfer time is the time between an inbound flight f_i^{IN} to an outbound flight f_j^{OUT} for all rotations from the whole set of fleetlines, L. This means that, for example, a connection is possible between rotation 1 of fleetline 1 and rotation 2 of fleetline 2, if there is sufficient time between the flights indicated by the minimum transfer time. The minimum transfer time is the minimum time that a passenger needs to get on the next flight on-time. This time determines whether a connection between two flights is possible. Similar to the minimum turnaround time, buffer time could occur in addition to the minimum transfer time when scheduling. Furthermore, the schedule also has to take into account the maximum acceptable transfer time for passengers. Passengers tend to experience long waiting times as unpleasant, so connections become decreasingly attractive the longer the transfer time is. As a result, there is a maximum to the transfer time passengers are willing to accept.

When creating a flight schedule, it can be designed for optimal connectivity or for optimal OTP. However, it is extremely unlikely that these two optimums occur within the same schedule. In order to maximize OTP, the time between flights should be increased as much as possible. The first flight in a fleetline should take off as early as possible, and the last flight in a fleetline should take off as late as possible. This would maximize the amount of buffer time that can be used to recover from delays. However, in order to maximize connectivity, it is important that the flights that have many connecting passengers have a transfer time that falls within the acceptable range. Flights for which there is large demand from connecting passengers should be scheduled close together, taking into account the transfer time, whereas connecting flights for which there is little demand can be scheduled farther apart. Since both OTP and connectivity are of great importance to an airline, and one usually comes at the cost of the other, two points would be very helpful to airlines:

i. Compromising on one of these objectives should lead to the maximum possible increase in the other.
ii. The trade-off between these two objectives should be made explicit, so that it becomes a clear strategic decision to make when designing a flight schedule.

3.2. Case study

The research presents a Proof-of-Concept applied to a case with real world data of a large international airline carrier. The holding of the airline carrier is leading in terms of international cargo and passenger traffic departing from Europe with more than 2.300 daily flights, and flying to more than 150 destinations. The size of the company makes the network highly complex and challenging to design and operate. The company operates via a hub-and-spoke network. Via the hub, many connections are made which allow passengers to efficiently access major destinations in the world. The hub serves more than 34 million passengers per day.

The large international airline carrier plays a leading role in the European air industry. The short haul flights, carried out within Europe, have a higher intensity than the long haul flights, while there is less buffer time due to the short length of the trips. This makes optimal scheduling and handling delays even more challenging and difficult for European short haul flights than for long haul flights. Due to the tight turnaround windows and high frequency in this part of the network, the need to recover from delays by adding buffers to the schedule, while at the same time maximizing connectivity, is felt most acutely here. Thus, we use a case study of the European flight schedule of this large international airline as the case study for this research.

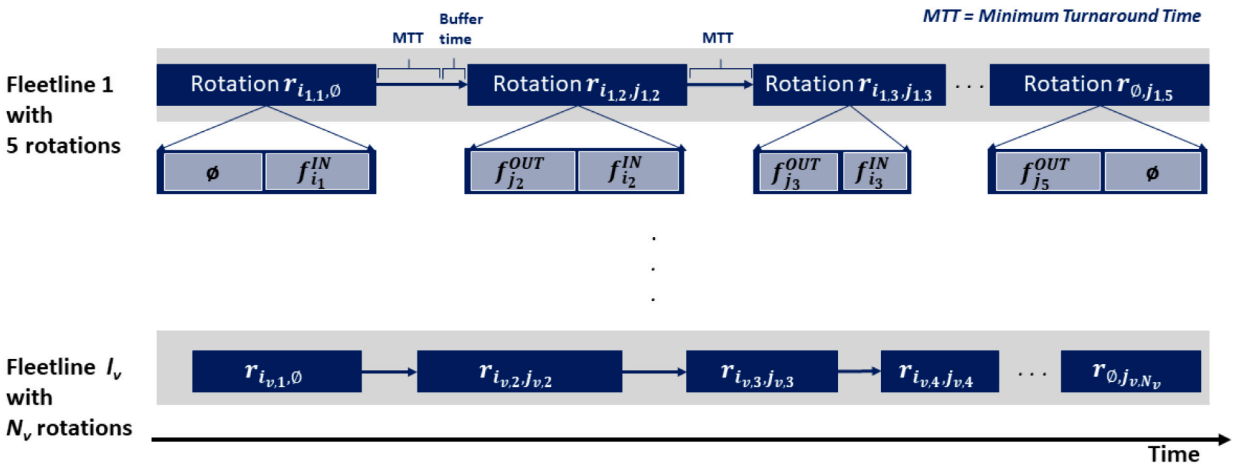

Fig. 1. A typical flight schedule.

3.3. Scope

The focus of this research is on the short haul flight schedules of the month July. Summer is the busiest period for airline carriers, meaning that the frequency of flights in the schedule is very high, leading to a very tight turnaround window. There is an urgent need for minimizing delay with buffers and integrating it with connectivity in this specific season. The base case schedules of this research are the real schedules that were designed and used by the large international airline carrier in the summer of 2019. More specifically, the base case schedules are schedules of two separate summer days of 2019. In the remainder of the paper, we refer to these as Day 1 and Day 2. Historical OTP and connectivity data from summer of 2019 on both of these schedules was provided by the airline. These are European flight schedules, meaning that there is a focus on short haul flights. Long haul flights are taken into account for the connectivity aspect of the short haul flight schedule. However, optimizing the long haul flight schedule itself is not within the scope of this research. As there are typically no overnight flights within the European flight schedule, the schedule can be cut into days. Here we consider the schedule for a single day. Moreover, only the fleetlines of the aircraft type Boeing 737–900 are included in the flight schedule optimized in this research.

Key Performance Indicators (KPIs) are defined to evaluate whether a candidate solution performs better than the base case schedule. OTP and connectivity are two main objectives to optimize in the schedule design for this research, and are thus the most important KPIs. Additional KPIs are determined to give more in-depth insight in the performance of the schedule, which helps decision makers make a choice between candidate schedules. The chosen additional KPIs are total arrival delay minutes of the schedule, percentage of missed connections, and average transfer time of a passenger at the hub in minutes. The total arrival delay minutes of a schedule is a valuable measure of the OTP side since it gives an indication of the overall disruption. Moreover, where the connectivity objective represents the number of connections offered, the percentage of missed connections represents the proportion of those connections that were realized in actuality, taking departure and arrival delay into account. If the initial transfer time is 50 min while the minimum transfer time is 45 min, this connection is possible and will be offered. Should the inbound flight then have an arrival delay of 10 min, the actual transfer time is 40 min and a percentage of connecting passengers would miss their connection. If the connecting outbound flight also has a departure delay of 10 min, the transfer time stays at 50 min and a larger percentage of passengers make their connections. Lastly, the average transfer time of a passenger at the hub is an indicator of customer experience of the flight schedule. A longer average transfer time results in a more negative customer experience since the waiting time for passengers becomes longer.

4. Mathematical model

In order to solve the multi-objective optimization problem, we formulate it mathematically. This section presents the objectives of the problem in Section 4.1 and a novel multi-objective model formulation in Section 4.2. The objectives and model formulation are based on literature (e.g., Wu, 2005; Burke et al., 2010; Danesi, 2006; Lee et al., 2014), and are designed in combination with experts and stakeholders of the large airline carrier.

4.1. Objectives

Let $X_{r_{i,j}}$ be the buffer time for rotation $r_{i,j}$. This is the time by which the arrival and departure of rotation $r_{i,j}$ shifts in fleetline l_v. Buffer time is the primary decision variable of this research. Let τ be the unit of buffer time in minutes. In our research, we use a five-minute buffer time as τ to schedule arriving and departing flights according to the large international airline case.

4.1.1. On-time performance

Combining the insight of previous theories of OTP in the airline industry (Lee et al., 2007; Wu, 2005; Sohoni et al., 2011; Burke et al., 2010) and sectors such as railway (Veiseth et al., 2007; Olsson and Haugland, 2004) and road transport (Chen et al., 2003), the OTP of rotation $r_{i,j}$ can be measured with delay minutes, thus how many minutes a flight of a rotation has "lost" compared to the scheduled time. Airlines mostly focus on the arrival punctuality as on-time performance and therefore, this research only focuses on the on-time performance of the arriving flight of a rotation. An extensive operationalization of on-time performance can be found in Appendix A.1.

The on-time performance objective is defined as the average arrival delay minutes of a rotation $r_{i,j}$ in the flight schedule. Each rotation consists of a unique inbound and outbound flight for that particular day. Each flight has a flight call sign that is not unique over the entire flight schedule over a specific time period, W. Flights with the same flight call sign have the same origin–destination, and are often scheduled on the same day of the week. For example, flight FL1000 to London is scheduled on Monday. Every individual flight FL1000 on Monday during Summer 2019 (W) is part of the collection of flights over which to average.

Let n be the flight call sign, and let S_n be the collection of flights for which we want to average the OTP for time period W. For rotation $r_{i,j} \in l$, the arrival OTP of rotation $r_{i,j}$ is the arrival punctuality of the last flight leg arriving at the hub, the inbound flight f_i^{IN}. Thus, the collection of flights S_n is equal to the outstation-to-hub flights with similar flight call sign n, $f_{i,n}^{IN}$. This gives:

$$OTP(S_n) = \frac{\sum_{f_i \in S_n} max(0, ata_i - [sta_i + X_{r_{i,j}} \tau])}{\# S_n} \tag{1}$$

where

OTP = arrival on-time performance, expressed in delay minutes
ata_i = actual time of arrival of flight $f_i \in S_n$
sta_i = scheduled time of arrival of flight $f_i \in S_n$
$r_{i,j}$ = the rotation belonging to the inbound flight f_i^{IN} and the outbound flight f_j^{OUT} prior to flight f_i^{IN}

The OTP objective for a flight schedule ξ with a set of rotations R is the average OTP for all rotations, i.e., for all inbound flights of the rotations, $f_i^{IN} \in F^{IN}$, and is formulated as:

$$OTP(\xi) = \overline{\sum_{f_i^{IN} \in F^{IN}} OTP(f_i^{IN})} \tag{2}$$

For OTP for various buffer allocations, it is assumed that adding more buffer time to a rotation leads to a better on-time performance and thus, less arrival delay. This means that shifting the rotation "forward", i.e., creating more buffer time, leads to less arrival delay and vice versa.

Another main assumption is that buffer time at outstations is included in the arrival delay. This means that when there is buffer time planned at the outstation, some of the calculated average arrival delay is compensated by this buffer. The arrival delay for a rotation is, thus, the departure delay plus the rotation delay minus the buffer at the outstation. The average arrival delay for a rotation is the historical arrival delay for a rotation averaged over the collection of flights S_n in Summer 2019. Real airline data is used to determine the average arrival delay of a rotation (in minutes). This data includes delays due to uncertain factors such as weather, propagated delay, and late arrival of passengers.

The third main assumption is that early arrivals are considered to be on-time. When an aircraft arrives earlier than the scheduled time of arrival, there is negative delay. Since the turnaround operation often still starts at the scheduled time of arrival and not earlier, the negative delay is set to zero.

4.1.2. Connectivity

Hub connectivity refers to the number and quality of indirect flights available to passengers via an airline hub (Lee et al., 2014; Burghouwt and de Wit, 2005; Danesi, 2006). Hub connectivity is quantified by several studies, mostly by means of indices. Burghouwt and de Wit (2005) defines hub connectivity as the number and quality of the indirect connections generated by the existing flights, and created a weighted indirect connectivity index. Danesi (2006) presents a novel Weighted Connectivity Ratio consisting of the weighted indirect connection number and the approximate number of weighted connections in a purely random situation, during a specific time period. Kim and Park (2012) presents a connectivity index that measures the relationship between arrivals and departures of flights in 24 h. Following the Weighted Connectivity Ratio of Danesi (2006) and Lee et al. (2014) developed the Continuous Connectivity Index for hub-and-spoke operations by adding an extra weighted element to the Weighted Connectivity Ratio and creating a continuous character.

The connectivity index for this research is operationalized building on the existing discrete connectivity index of Danesi (2006) and Continuous Connectivity Index of Lee et al. (2014). For an airline, the actual number of passengers catching their transfer is interesting for determining connectivity. The percentage of passengers that make the transfer is a function of the transfer time, $tt_{i,j} = std_j - sta_i$, and can be extracted from real airline data. The real airline data includes many uncertainties impacting the transfer times and the percentage of passengers making the transfer, such as weather, walking times to the gate of the transferring flight, and waiting times for customs control. This can be translated to the probability that a passenger will actually be boarded on the transferring flight given the scheduled transfer time, $p_k(tt_{i,j})$. This probability differs per transfer type $k \in K$, for example it is faster to transfer from a European flight to another European flight than to transfer from a European flight to an international flight, since the distance between gates for European flights is often smaller. Possible connections types in K are European flights to European flights, International flights to European flights, and European flights to International flights. Hereby, the Schengen versus Non-Schengen connection types also are incorporated.

Not all connecting passengers are equal for the airline, since some bring more revenue than others, and can therefore be considered to be more 'important' from a commercial point of view. Therefore, we chose to include the revenue of a passenger connection to measure hub connectivity for airlines. The expected number of passengers connecting from flight f_i to flight f_j, $tp_{i,j}$, is taken as an approximation for earnings of a connection for an airline. This is multiplied by the revenue of one passenger connecting between flight f_i and f_j, $c_{i,j}$, defined by the commercial branch of the airline, who define revenue of passengers connecting between flights based on historical revenue data for connections.

Airlines have strict constraints regarding the transfer times of passengers. The transfer time of a passenger connection cannot be too short as the passenger will not catch their next flight. On the other hand, the transfer of a passenger cannot be too long as this is seen as undesirable by passengers and thus could lead to a decrease in customer satisfaction. Therefore, the connectivity revenue is only included for flights where the transfer time is between the minimum transfer time of a passenger with transfer type $k \in K$, MCT_k, and the maximum acceptable transfer of a passenger with transfer type $k \in K$, $MACT_k$.

The connectivity objective in this research is total connectivity revenue in euros of all rotations in the flight schedule. The connectivity of rotation $r_{i,j}$ is divided into inbound connectivity and outbound connectivity as shown in Fig. 2.

Inbound connectivity revenue is determined by the connections between inbound flights on hub station $F^{IN} \subset F$ and the first flight of the rotation from hub station to outstation, f_j^{OUT}. Outbound connectivity revenue is determined by the connections between the last flight of the rotation arriving at hub station from outstation f_i^{IN} and the related outbound connections from hub station $F^{OUT} \subset F$. Thus, the outbound flight f_j^{OUT} of rotation $r_{i,j}$ determines the inbound connectivity revenue, C_j^{IN}, and the inbound flight f_i^{IN} of rotation $r_{i,j}$ determines the outbound connectivity revenue, C_i^{OUT}. In combination with the formulated connectivity index, this results in the following definitions of objectives.

Inbound connectivity revenue for flight f_j^{OUT} of rotation $r_{i,j}$ can be defined as

$$C_j^{IN} = \sum_{f_i \in F^{IN}} \begin{cases} p_k(tt_{i,j}) tp_{i,j} \ c_{i,j}, & \text{if } MCT_k \le tt_{i,j} \le MACT_k \\ 0, & \text{otherwise} \end{cases} \tag{3}$$

with $tt_{i,j} = std_j - sta_i + X_{r_{i,j}} \tau$.

Note that the buffer time X can also be negative. Outbound connectivity revenue for flight f_i^{IN} of rotation $r_{i,j}$ can be defined as

$$C_i^{OUT} = \sum_{f_j \in F^{OUT}} \begin{cases} p_k(tt_{i,j}) tp_{i,j} \ c_{i,j}, & \text{if } MCT_k \le tt_{i,j} \le MACT_k \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

with $tt_{i,j} = std_j - sta_i + X_{r_{i,j}} \tau$.

The connectivity revenue objective for a flight schedule ξ with a set of fleetlines L is defined as follows:

$$C(\xi) = \sum_{r_{i,j} \in R} C_j^{IN}(X_{r_{i,j}} \tau) + C_i^{OUT}(X_{r_{i,j}} \tau) \tag{5}$$

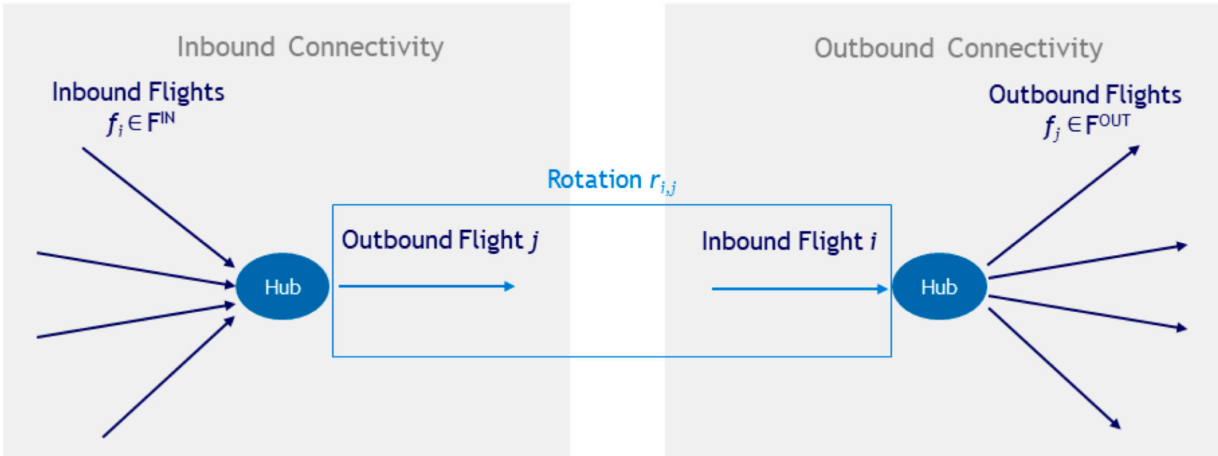

Fig. 2. Inbound and outbound connectivity.

$$\text{minimize} \quad \overline{\sum_{f_i^{IN} \in F^{IN}} OTP(f_i^{IN})}, \quad \text{maximize} \sum_{r_{i,j} \in R} C_j^{IN}(X_{r_{i,j}}\tau) + C_i^{OUT}(X_{r_{i,j}}\tau) \tag{6}$$

$$\text{subject to} \quad \sum_{r_{i,j} \in l_v} (X_{r_{i,j}}\tau) \leq T_{l_v} \qquad \forall l_v \in L \tag{7}$$

$$MCT_k \leq tt_{i,j} \leq MACT_k \qquad \forall i,j \tag{8}$$

$$std_{f_j \in r_{i+1,j+1}} - sta_{f_i \in r_{i,j}} + (X_{r_{i,j}}\tau) \geq MTT_{ac} \qquad \forall r_{i,j} \in l_v, \forall l_v \in L \tag{9}$$

Box I.

An extensive description of previous literature and the quantification of the connectivity objective can be found in Appendix A.2. In the remainder of the paper, we refer to connectivity revenue as connectivity.

4.2. Formulation

To find the optimal combination of $X_{r_{i,j}}$ in a chain of flights, the entire fleetline is optimized simultaneously. The interdependencies between the arrival and departure times are taken into account for this. There are two main objectives for this optimization, namely the average delay minutes of the schedule and the sum of inbound and outbound connectivity revenue of the schedule. This gives the following optimization problem: (see Eqs. (6)–(9) in Box I).

Constraint (7) ensures that no extra time could be added to a day. The total buffer time of the day per fleetline, T_{l_v}, can only be reallocated. This means that the sum of the buffer time of all rotations of one fleetline should be less or equal to the total amount of buffer time that could be reallocated for a fleetline. This applies to all fleetlines in the schedule. Constraint (8) ensures that the transfer time of a passenger between inbound flight f_i and outbound flight f_j is not smaller than MCT_k and not larger than $MACT_k$. Constraint (9) ensures that the turnaround time between two rotations in a fleetline is equal or more than the minimum turnaround time of the aircraft type, MTT_{ac}. The sequence of the rotations in a fleetline l_v is fixed meaning that $r_{i+1,j+1}$ is the subsequent rotation of rotation $r_{i,j}$.

5. Solution approach

In this research, a multi-objective optimization approach is used. Characteristics of this approach are that the objectives are optimized simultaneously without scalarization, and the Pareto optimal front helps to identify the trade-off between objectives (Burke et al., 2010; Emmerich and Deutz, 2018; Kollat and Reed, 2007). The Pareto optimal front is the set of non-dominated Pareto solutions. Non-dominated Pareto solutions are solutions that cannot advance the performance of one objective without deteriorating the other objectives (Emmerich and Deutz, 2018). In the context of this research, a Pareto optimal solution represents a schedule with buffers in which buffers cannot be reallocated in such a way that it improves OTP without decreasing connectivity and vice versa.

A widely used approach for solving multi-objective optimization problems is through using a multi-objective evolutionary algorithm (MOEA), i.e., a population-based search algorithm (Vikhar, 2016). Within the class of MOEAs, genetic algorithms (GA) are known to generate high-quality solutions for optimization problems based on the concept of natural selection in Darwin's theory of evolution (Mitchell, 1996). In this research, we use a generational version of a MOEA called BORG as the optimization algorithm. BORG combines ϵ-dominance archiving, adaptive population sizing, and time continuation, with adaptive operator selection (Hadka and Reed, 2013; Reed et al., 2013; Hernandez-Diaz et al., 2007). The motivation for using the genetic algorithm BORG is twofold. First, BORG is an extension of ϵ-NSGAII with adaptive operator selection (Kollat and Reed, 2006, 2007). This means that the algorithm keeps track of the performance of the operators and adapts to the most appropriate operator. Operators of BORG are binary crossover, differential evolution, parent-centric recombination, unimodal normal distribution crossover, simplex crossover, polynomial mutation and uniform mutation, thus in total seven operators. In contrast, most other MOEA employ a single search operator (Singh et al., 2015). Second, BORG has been demonstrated to outperform other MOEAs when the population size, number of decision variables, and the complexity of the problem increases (Ward et al., 2015; Salazar et al., 2016). Given to the complex and interdependent nature of the schedule design problem as considered here, BORG is a suitable choice.

Table 2
Generational distance of Day 1 and Day 2 of five seeds.

	Day 1	*Day 2*
Seed 1	0.009	0.033
Seed 2	0.040	0.010
Seed 3	0.004	0.017
Seed 4	0.007	0.048
Seed 5	0.024	0.009

To evaluate the performance of the optimization, we consider (1) convergence, i.e., convergence of the solution set to the Pareto optimal solutions, and (2) diversity, i.e., a diverse set of solutions in the objective space (Reed et al., 2013). Convergence is measured by ε-progress. ε-progress measures whether the optimization models found a substantially better solution based on the user-defined search precision ε (Kwakkel et al., 2016).

Since the algorithm does not converge with the fully random initial population to feasible solutions, we initialize the population with 20% of solutions that are close to the base case. To guarantee diversity in the solutions, the other 80% of the initial population is sampled randomly. Due to the randomness in the optimization algorithm (including this random sample), the algorithm is analyzed for five random seeds. We merge the Pareto-optimal solutions resulting from the random seeds into a single Pareto optimal set using a non-dominated ε-sort. The combined Pareto front is used as reference set for calculating the generational distance of each of the solution sets to the combined Pareto optimal set. The smaller the generational distance, the closer the given set of solutions is to the reference set (Lwin et al., 2014).

In this research, a deterministic simulation model is used. A detailed description of the configuration of the simulation–optimization model can be found in van Schilt (2020). The simulation–optimization model is written in Python. The library Exploratory Modeling and Analysis in Python is used for implementing the optimization algorithm BORG (Kwakkel, 2017).

6. Results

In this section, the performance of the optimization model using the solution approach is discussed first. Afterward, the results, i.e., optimal schedules, of the two separate summer days (Day 1 and Day 2) are presented.

6.1. Performance of optimization model

The performance of the optimization model in combination with the solution algorithm is examined to ensure that the outcomes of the model converge to Pareto optimal solutions. In order to determine the number of function evaluations needed for this problem, the convergence of the optimization model is evaluated by means of the ε-progress. The approach converges fully when the number of improvements stabilizes.

Fig. 3 shows the results of ε-progress for Day 1 and Day 2 of five seeds. A line is drawn at the number of function evaluations where the number of improvements stabilizes. For both days, this is after around 18.000 function evaluations, meaning that the optimization has converged according to the ε-progress metrics. Therefore, we use 20.000 function evaluations in this research.

Next, to assess whether the different seeds converge to essentially the same approximate Pareto front, we calculate the generational distance of the five seeds to the combined Pareto front across the five seeds (i.e., the reference set). The smaller the generational distance, the closer the given set of solutions is to the reference set (Lwin et al., 2014).

Table 2 presents the results of the generational distance for Day 1 and Day 2 of the five seeds. For the schedules of both days, the generational distance of the Pareto fronts from the various seeds to the reference set are very small, i.e., between 0.01 and 0.05.

The outcome of ε-progress and the generational distance for this problem show that the algorithm converges to essentially the same Pareto front, i.e., the optimal combined Pareto front across five seeds.

6.2. Results per day

This section presents the results of the optimization model per day. For each day, the Pareto-optimal front is discussed first. Next, the results of the KPIs are examined for all the Pareto optimal solutions. Subsequently, we present the results of the two most divergent solutions in terms of ΔOTP in more detail. The graphs of the results of objectives and the additional indicators for all solutions of Day 1 and Day 2 can be found in Appendix B.

6.2.1. Day 1

Fig. 4 shows the Pareto-optimal front of Day 1. The difference between the base case and the Pareto-optimal solutions are displayed on the axes, measured in the percentage of improvement. The base case is the starting point and therefore, set to 0% improvement for both connectivity and on-time performance. This means that the difference (Δ) in percentage shows the increase or decrease of each objective compared to the base case. The arrows on the axes point out the direction of desirability. Both ΔOTP and ΔConnectivity should be maximized by the optimization model. This means that the point closest to the upper-right corner is the most optimal. The figure shows that the multi-objective optimization model improves the base schedule on either or both OTP and connectivity.

The results for ΔOTP in Fig. 4 present the percentage of average arrival delay minutes that the schedule increases or decreases compared to the base case of Day 1. There are two solutions that are expected to lead to a negative percentage of on-time performance, meaning more average arrival delay minutes than the base case. However, there is an increase in connectivity value for these particular solutions and therefore, these solutions are still interesting to include in the analysis.

The results for connectivity in Fig. 4 show that all Pareto solutions have a connectivity value that is at least 1.2% higher than the base case (i.e., the current actual schedule). This means that the airline would get more connectivity revenues with each of these Pareto optimal schedules compared to the actual schedule. Fig. 4 also confirms the trade-off between OTP and connectivity, namely higher OTP (less average arrival delay minutes) leads to less connectivity.

It is up to the decision maker to determine the importance of OTP versus connectivity. For example, schedule planners rather choose for a higher OTP (thus lower average arrival delay minutes) where connectivity is lower, as it does not differ significantly from a connectivity value with a lower OTP. However, in all cases, it would be better to choose a schedule resulting from the simulation–optimization model, rather than the base schedule that the airline actually uses.

Other indicators to evaluate the solutions are the total arrival delay minutes, the percentage of missed connections by passengers, and the average transfer time. Total arrival delay minutes is proportional to OTP. This means that most Pareto-optimal solutions perform better on total arrival delay minutes than the base case. Results of Day 1 show that the Pareto-optimal solution that has the same OTP as the base case has over 10% more arrival delay minutes than the base case. However, similar to the logic followed for OTP, the schedules that have more total arrival delay minutes than the base case are still interesting to include since they have a higher connectivity than the base case.

The percentage of missed connections by passengers shows whether the schedules would decrease the number of missed connections. The results of Day 1 show that a Pareto optimal solution with a higher OTP than the base case (thus higher than 0%) decreases the percentage of missed connections by passengers. This means that a better OTP

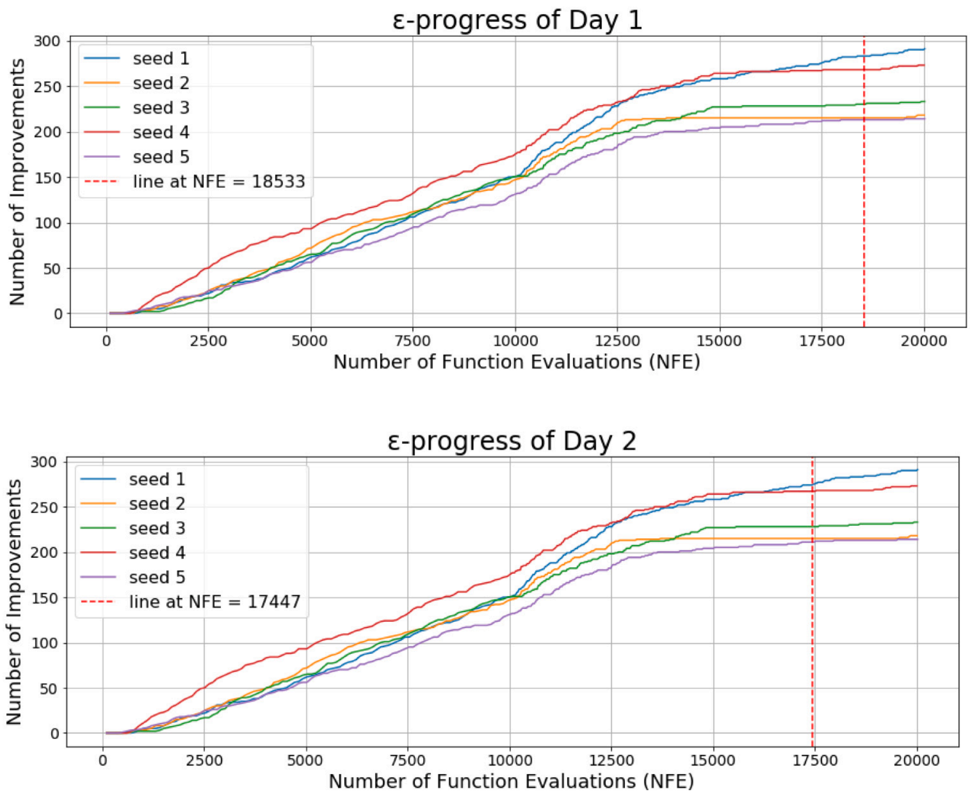

Fig. 3. ε-progress of Day 1 and Day 2 with the number of improvements per number of function evaluations of five seeds.

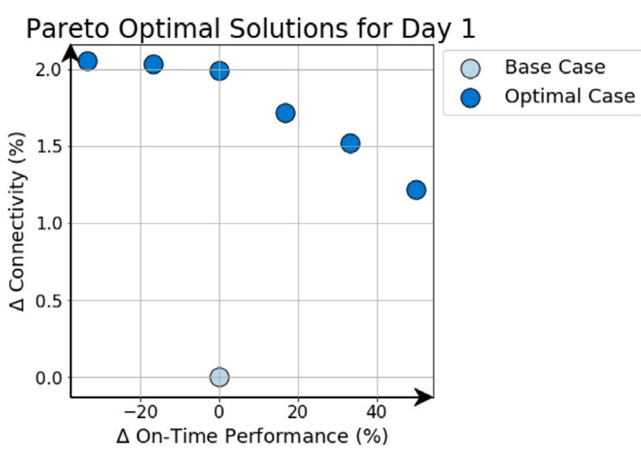

Fig. 4. Pareto front of Day 1 in percentage of difference (Δ) between the base case (0%, 0%) and the Pareto-optimal solutions for connectivity and on-time performance.

compared to the base case schedule would lead to fewer passengers missing their connections.

The average transfer time for a passenger in the Pareto schedules is significantly longer than the base case of Day 1 (between 10% and 12%).

Given the Pareto optimal front, we discuss the two most divergent solutions in more detail: (1) the solution with −33.3% ΔOTP, and (2) the solution with 50% ΔOTP. Table 3 presents the used buffer time for each fleetline, divided by the total buffer time for each fleetline. It also presents the used buffer time in percentage for the entire flight schedule. These percentages are given for both solutions. The base case is the reference value for the buffer time and therefore, is set to 0% of buffer time used. More specifically, the buffer time ($X_{r_{i,j}}$) for all rotations R of the base case is 0. The table shows that both solutions use at least 16% more buffer time than the base case. Also Solution 1.1 uses 9% less buffer time than Solution 1.2 in the total flight schedule. In the case of Solution 1.1, the optimization allocates more negative buffer time to the rotations than Solution 1.2, hence the lower amount of buffer time used. This means that rotations are shifted "forward", i.e., to the beginning of the flight schedule compared to the base case. Moreover, Solution 1.2 (with a higher ΔOTP) uses at least 17% more buffer time in each fleet line compared to the base case.

Table 3
Buffer time divided by the total buffer time available (in percentage) compared to the base case (0% of buffer time used) used by the optimal solutions of −33.33% ΔOTP (Solution 1.1) and 50% ΔOTP (Solution 1.2), per fleetline l_v and over the flight schedule ξ of Day 1.

	Buffer time used by Solution 1.1 (−33.33% ΔOTP)	Buffer time used by Solution 1.2 (50% ΔOTP)
Fleetline 1 (l_1)	14.3%	17.9%
Fleetline 2 (l_2)	17.6%	19.6%
Fleetline 3 (l_3)	16.7%	29.6%
Fleetline 4 (l_4)	14.7%	35.3%
Fleetline 5 (l_5)	17.9%	25.0%
Flight schedule (ξ)	*16.6%*	*25.6%*

Table 4
Results of the key performance indicators as compared with the base case of the optimal solutions of −33.33% ΔOTP (Solution 1.1) and 50% ΔOTP (Solution 1.2) of Day 1.

	Solution 1.1 (−33.33% ΔOTP)	Solution 1.2 (50% ΔOTP)
ΔOTP	−33.3%	50.0%
ΔConnectivity	2.1%	1.2%
ΔTotal arrival delay minutes	37.6%	−35.3%
ΔPercentage of missed connections by passengers	2.7%	−2.0%
ΔAverage transfer time	11.2%	10.4%

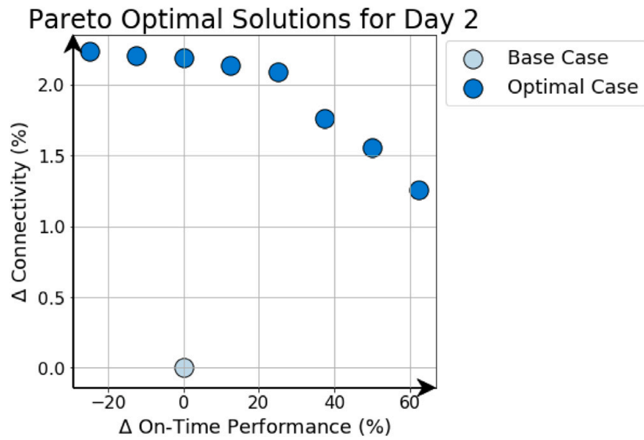

Fig. 5. Pareto front of Day 2 in percentage of difference (*Δ*) between the base case (0%, 0%) and the Pareto solutions for connectivity and on-time performance.

The KPIs for both solutions are presented in Table 4. The base case is the reference value, i.e., 0%, for all KPIs. The results confirm the trade-off between OTP and connectivity, e.g., Solution 1.1 has the highest *Δ*Connectivity and vice versa for Solution 1.2. The difference in *Δ*OTP is more than 83% and the difference in *Δ*Connectivity is around 1%. Regarding the total arrival delay minutes, the solutions differ with more than 70%. Moreover, the results also show that the schedule with the highest OTP (Solution 1.2) compared to the actual schedule leads to 2.0% less passengers missing their connection whereas the schedule with the lowest OTP (Solution 1.1) results in 2.7% more passengers missing their connections. However, the *Δ*Average Transfer Time of Solution 1.1 is slightly higher than that of Solution 1.2. In both solutions, there is more transfer time available for the passengers than in the base case.

6.2.2. Day 2

Fig. 5 shows the Pareto optimal front of Day 2. The figure shows that the multi-objective optimization model also improves the actual schedule of Day 2 on either or both OTP and connectivity. For each solution, the model improves connectivity with at least 1.3% compared to the base case.

Similar to Day 1, there are two solutions that are expected to lead to a negative percentage of on-time performance, meaning more average arrival delay minutes than the base case. However, there is an increase in connectivity value for these particular solutions and therefore, these solutions are still interesting to include in the analysis.

Regarding connectivity, Fig. 5 shows that all Pareto-optimal solutions have a higher connectivity of at least 1.3%. So, it would always be better for connectivity to choose a schedule resulting from the simulation–optimization model, rather than the base schedule that is currently in use by the airline.

Another indicator is the total arrival delay minutes. For Day 2, this is proportional to the OTP. Next, the outcomes of percentage of missed connections by passengers show that a solution with a lower OTP than the base case (*Δ*-12.5%) results in a slight decrease in percentage of missed connections by passengers. The solution with 12.5%*Δ*OTP shows a smaller decrease in the percentage of missed connections (less than -1%) compared to a solution with more average arrival delay minutes (close to −2%). So although OTP improves, it does not necessarily contribute to a lower percentage of missed connections. The average transfer time of the schedules of Day 2 is 4% to 8% higher than in the base case.

Given the Pareto optimal front, we discuss the two most divergent optimal solutions in more detail: (1) the solution with -25% *Δ*OTP, and (2) the solution with 62.5% *Δ*OTP. Table 5 presents the used buffer time

Table 5
Buffer time divided by the total buffer time available (in percentage) compared to the base case (0% of buffer time used) used by the optimal solutions of −25% *Δ*OTP (Solution 2.1) and 62.5% *Δ*OTP (Solution 2.2), per fleetline l_v and over the flight schedule ξ of Day 2.

	Buffer time used by Solution 2.1 (−25% *Δ*OTP)	Buffer time used by Solution 2.2 (62.5% *Δ*OTP)
Fleetline 1 (l_1)	43.1%	43.3%
Fleetline 2 (l_2)	5.8%	41.2%
Fleetline 3 (l_3)	23.2%	41.1%
Fleetline 4 (l_4)	32.1%	46.4%
Fleetline 5 (l_5)	11.1%	26.0%
Flight schedule (ξ)	*23.4%*	*38.6%*

Table 6
Results of the key performance indicators as compared with the base case of the optimal solutions of −25% *Δ*OTP (Solution 2.1) and 62.5% *Δ*OTP (Solution 2.2) of Day 2.

	Solution 2.1 (−25% *Δ*OTP)	Solution 2.2 (62.5% *Δ*OTP)
*Δ*OTP	−25%	62.5%
*Δ*Connectivity	2.2%	1.3%
*Δ*Total arrival delay minutes	27.9%	−54.9%
*Δ*Percentage of missed connections by passengers	2.4%	−4.9%
*Δ*Average transfer time	5.1%	7.7%

for each fleetline, divided by the total buffer time for each fleetline. It also presents the used buffer time as a percentage for the entire flight schedule. These percentages are given for both solutions. Similar to Day 1, the base case is the reference value for the buffer time and therefore, is set to 0% of buffer time used. This means that the buffer time ($X_{r_{i,j}}$) for all rotations R of the base case is 0. The table shows that Solution 2.1 and 2.2 uses 23.4% and 38.6% more buffer time than the base case, respectively. Also Solution 2.2 uses more than 15% more buffer time than Solution 2.1 in the total flight schedule. In the case of Solution 2.2, the optimization model allocates more positive buffer time to the rotations. This means that rotations are shifted "backwards", i.e., to the end of the flight schedule compared to the actual schedule, to create more room for being on-time. Moreover, the two most divergent solutions both use a similar amount of buffer time in fleetline 1. In all the other fleetlines, Solution 2.2 (with a higher *Δ*OTP) uses at least 14% more buffer time than Solution 2.1. For example, in fleetline 2, Solution 2.1 only uses 5.8% more buffer time than the base case whereas Solution 2.2 uses 41.2% more buffer time.

The KPIs for both solutions are presented in Table 6. The base case is the reference value, i.e., 0%, for all KPIs. The results again confirm the trade-off between OTP and connectivity, e.g., Solution 2.1 has the highest *Δ*Connectivity and vice versa for Solution 2.2. The difference in *Δ*OTP is more than 87% and the difference in *Δ*Connectivity is around 1%. It is up to the decision makers to determine the importance of these differences. For the KPI total arrival delay minutes, the solutions differ with more than 80%. Moreover, the results also show that Solution 2.2 (with the higher OTP) leads to around 4.9% less passengers missing their connection whereas Solution 2.1 (with the lowest OTP) results in 2.4% more passengers missing their connections. In line with this, *Δ*Average Transfer Time of Solution 2.2 is 2.6% higher than Solution 2.1. Thus, passengers have more time to transfer between flights, leading to fewer passengers missing their connections in the case of the solution with the highest *Δ*OTP (Solution 2.2). In both solutions, there is more transfer time available for the passengers than in the base case schedule.

7. Discussion

This section discusses the results of the simulation–optimization model. It presents the interpretation of the results in Section 7.1 and the limitation of this model in Section 7.2.

7.1. Interpretation of results

The multi-objective optimization model solved through MOEA BORG successfully improves the base schedule on either or both OTP and connectivity. This means that buffer scheduling resulting from the optimization model performs better than buffer allocation in the base case. It would be better to choose a schedule resulting from the simulation–optimization model, rather than the actual schedule. Due to the multi-objective nature of the problem, presenting the optimization results in a Pareto optimal front helps schedule planners to make a trade-off between OTP and connectivity. It is the task of the decision maker to determine the importance of OTP versus connectivity and choose an optimal schedule accordingly.

The percentage of buffer time used in the optimal flight schedules resulting from the multi-objective optimization model is relatively small compared to the total available buffer time. Even in the schedules with a ΔOTP of more than 50%, the percentage of extra buffer time used is smaller than 40% of the total buffer time available. The reason is that the model is not only optimizing on OTP but also on connectivity. Also, the results show that in the schedules with the lowest ΔOTP and the highest ΔOTP of both days use at least more than 16% more buffer time than the base case schedule. Thus, the optimization model uses more buffer time compared to the base case to either increase OTP or connectivity. Furthermore, the results of the two most divergent schedules of Day 1 and Day 2 imply that using less buffer time leads to a lower OTP but a higher connectivity and vice versa. Less buffer time often means that the flight schedule is tightly planned, resulting in more passengers missing their connections. Therefore, the schedule with the lowest ΔOTP has a higher percentage of missed connections by passengers than the base case or the schedule with the highest ΔOTP. One notable result is that the average transfer time of the schedule for Day 1 with the lowest ΔOTP is slightly higher (0.8%) than the schedule with the highest ΔOTP, whereas more passengers seem to miss their connections (2.7%). This implies that a higher average transfer time results in more passengers missing their connections for this solution. The main explanation is that, in this case, the optimization model creates many connections with a short transfer time and keeps a few connections with a long transfer time which on average leads to a relatively high transfer time over the flight schedule. Due to the many connections with a short transfer time, many passengers still miss their flight and therefore, this percentage stays high.

More specifically on the KPI percentage of missed connection by passengers, departure delay plays an important role. The result of the schedule of Day 2 shows that the improvement of OTP compared to the base schedule does not necessarily contribute to a lower percentage of missed connections. The reason could be that arrival delay has been improved by buffer allocation (thus a higher OTP) but it did not influence departure delay which also impacts the percentage of missed connections. Another reason is that lower OTP has a higher connectivity and lower transfer times, which leads to more connections being missed by passengers. Since the average transfer times do not vary much between the optimal schedules (a difference of around 2%), the time to catch a flight for a passenger is on average similar for schedules with a lower OTP and a higher OTP. This could also explain why a higher OTP does not necessarily lead to a lower percentage of missed connections. In contrast, the result of the schedule of Day 1 implies that the higher OTP, thus the less average arrival delay minutes, the fewer passengers miss their connection. The main difference with Day 2 is that the buffer allocation immediately impacts both arrival and departure delay on Day 1.

For the KPI average transfer time for a passenger, the boundary cases of connections play an important role. The simulation–optimization model tries to maximize connectivity, meaning that the schedules could offer more long connections (i.e., close to MACT) or cut-off short connections (i.e., close to MCT). The result of the schedule of Day 1 shows that connectivity increases while the average transfer time increases as well, since it includes more connections with a longer transfer time. The result of the schedule of Day 2 shows that the average transfer time can both be longer and shorter than the base case but are relatively close to the base case. This means that the average transfer time variability depends on the relative amount of connections that lie close to MCT versus close to MACT.

7.2. Limitations

There are three main limitations of the simulation–optimization model namely (i) a simplistic formulation of the constraints for the real-life problem, (ii) difficulty to find a feasible solution for this highly complex problem, and (iii) no convergence with stochastic variance.

The presented multi-objective optimization formulation is simplistic with respect to the constraints in a real-life schedule. Three simple base line constraints are included in this model namely (1) no extra time could be added in a day, (2) transfer time of passengers is within the minimum and maximum accepted transfer time, and (3) turnaround time between two rotations is larger than or equal to the minimum turnaround time. There are numerous additional constraints that apply in practice when designing a flight schedule, such as slot assignment, crew legalization, or resource assignment. However, finding the optimal candidate schedules with only these four relatively simple constraints is already challenging. This shows the highly complex nature of the problem. Therefore, when adding more constraints to create an optimization formulation closer to real-life, it is likely that it becomes very difficult to find optimal candidate schedules. Helpful tools to be able to add more constraints would be to relax the constraints or to add penalty costs. This gives infeasible but good performing schedules, instead of very few feasible optimal schedules, which actually might work better than the fully optimized schedules.

The multi-objective optimization formulation is solved through the MOEA BORG. As the problem is highly complex, even with only four constraints, it is difficult to find feasible schedules that do not violate the strict constraints within a reasonable time. The optimization only converged after many functions evaluation and a large run time, meaning a high computational investment is needed to find feasible optimal solutions. This challenge related to convergence shows the high complexity of solving the problem with only four constraints. This complexity could affect the usability of the optimization formulation when adding constraints.

Another challenge related to convergence through the MOEA BORG is the stochastic variance of the objectives in the optimization model. Currently, this research uses a deterministic optimization model with a simulation–optimization framework. Real airline data including many uncertainties is used for the mathematical description of the objectives in the optimization model. However, the uncertainties themselves are not explicitly modeled for the purpose of our study. This could have an impact on the performance of BORG as it is only robust with a certain level of noise, i.e., a certain level of stochastic variance. If the optimization model has a high stochastic variance on the two objectives with an extremely wide confidence interval, the simulation–optimization model is not consistent in determining whether a candidate solution performs better or worse than others. For example, if a candidate solution is evaluated twice by the optimization model with a high stochastic variance, it could give two completely different values for OTP and/or connectivity. Although this could be a challenge for the optimization model, in reality OTP of a flight cannot be predicted using only the historical data due to the high level of uncertainty. Therefore, it would be interesting to include the stochastic variance by explicitly modeling the uncertainties in the multi-objective optimization model.

8. Conclusion

The trade-off between OTP and connectivity plays a large role in allocating buffers in the schedule design of transportation services. For the aviation sector specifically, it is challenging to schedule buffers in an already tight schedule such that there is a high OTP and high connectivity, given the strict constraints on minimum turnaround times of aircraft and transfer times of passengers. Considering these constraints, it is complex to optimize a daily flight schedule for an airline on both OTP and connectivity, which are the main objectives of an airline. Although much research is performed on schedule design, the trade-off between OTP and connectivity had not yet been investigated. Therefore, we presented a novel multi-objective optimization formulation of a daily Europe flight schedule for a large international airline carrier in this paper. Solving this problem formulation through the MOEA BORG, the optimization results in candidate schedules with both a better OTP and a higher connectivity than the base schedule. Having candidate schedules instead of one optimal schedule leaves room for decision makers to explicitly decide on the importance of OTP versus that of connectivity, rather than the model already deciding on the trade-off. This means that the final optimal daily schedule is decided by the decision makers based on their expertise.

Thus, the presented multi-objective optimization formulation, in combination with the simulation–optimization framework, successfully provides support to decision makers when trading off OTP and connectivity in a daily European flight schedule. It results in candidate schedules which perform better on connectivity and OTP than the base schedule. The technique is a suitable way for solving this complex integral problem in the transportation and logistics sector.

The multi-objective optimization formulation of this paper lays the first basis for trading off OTP and connectivity in a schedule design by buffer scheduling. Only a small set of the constraints for designing a schedule are incorporated in this paper. This small set of constraints, namely the minimum turnaround time for aircrafts and the minimum and maximum acceptable transfer times of passengers, are the foundation of designing a flight schedule. However, constraints such as slots, crew legalization, and resource assignment are also necessary for decision making. To get a fully integrated optimization model for schedule design, further research could focus on incorporating these elements to support decision makers even better. Moreover, further research could also focus on including stochastic variance by explicitly modeling the uncertainties in the multi-objective optimization formulation.

CRediT authorship contribution statement

Isabelle M. van Schilt: Conceptualization, Methodology, Software, Formal analysis, Writing – original draft, Project administration. **Jonna van Kalker:** Conceptualization, Resources, Validation, Data curation, Writing – original draft, Supervision. **Iulia Lefter:** Conceptualization, Writing – review & editing, Resources, Visualization, Supervision. **Jan H. Kwakkel:** Conceptualization, Methodology, Software, Writing – original draft, Supervision. **Alexander Verbraeck:** Conceptualization, Methodology, Validation, Writing – original draft, Writing – review & editing, Supervision.

Data availability

The data that has been used is confidential.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jairtraman.2024.102547.

References

Achenbach, A., Spinler, S., 2018. Prescriptive analytics in airline operations: Arrival time prediction and cost index optimization for short-haul flights. Oper. Res. Perspect. 5, 265–279. http://dx.doi.org/10.1016/j.orp.2018.08.004.

Ageeva, Y., 2000. Approaches to Incorporating Robustness into Airline Scheduling (Ph.D. thesis). Massachusetts Institute of Technology, URL: https://hdl.handle.net/1721.1/8930.

Ahmadbeygi, S., Cohn, A., Lapp, M., 2010. Decreasing airline delay propagation by re-allocating scheduled slack. IIE Trans. 42 (7), 478–489. http://dx.doi.org/10.1080/07408170903468605.

Andradóttir, S., 1998. Simulation optimization. In: Banks, J. (Ed.), Handbook of Simulation. John Wiley & Sons, New York, NY, USA, pp. 307–334.

Barnhart, C., Cohn, A., 2004. Airline schedule planning: Accomplishments and opportunities. Manuf. Serv. Oper. Manag. 6 (1), 3–22. http://dx.doi.org/10.1287/msom.1030.0018.

Baumgarten, P., Malina, R., Lange, A., 2014. The impact of hubbing concentration on flight delays within airline networks: An empirical analysis of the US domestic market. Transp. Res. E 66, 103–114. http://dx.doi.org/10.1287/msom.1030.0018.

Burghouwt, G., de Wit, J., 2005. Temporal configurations of European airline networks. J. Air Transp. Manag. 11 (3), 185–198. http://dx.doi.org/10.1016/j.jairtraman.2004.08.003.

Burke, E.K., De Causmaecker, P., De Maere, G., Mulder, J., Paelinck, M., Berghe, G.V., 2010. A multi-objective approach for robust airline scheduling. Comput. Oper. Res. 37 (5), 822–832. http://dx.doi.org/10.1016/j.cor.2009.03.026.

Carlier, S., De Lépinay, I., Hustache, J.-C., Jelinek, F., 2007. Environmental impact of air traffic flow management delays. In: Proceedings of 7th USA/Europe Air Traffic Management Research and Development Seminar. EUROCONTROL, Barcelona, Spain, p. 16, URL: https://www.atmseminar.org/seminarContent/seminar7/papers/p_101_EC.pdf.

Chen, C., Skabardonis, A., Varaiya, P., 2003. Travel-time reliability as a measure of service. Transp. Res. Rec. 1855 (1), 74–79. http://dx.doi.org/10.3141/1855-09.

Churchill, A.M., Lovell, D.J., Ball, M.O., 2010. Flight delay propagation impact on strategic air traffic flow management. Transp. Res. Rec. 2177 (1), 105–113. http://dx.doi.org/10.3141/2177-13.

Clausen, J., Larsen, A., Larsen, J., Rezanova, N.J., 2010. Disruption management in the airline industry—Concepts, models and methods. Comput. Oper. Res. 37 (5), 809–821. http://dx.doi.org/10.1016/j.cor.2009.03.027.

Danesi, A., 2006. Measuring airline hub timetable co-ordination and connectivity: definition of a new index and application to a sample of European hubs. Eur. Transp. 34, 54–74, URL: https://hdl.handle.net/10077/5929.

Deb, K., 2014. Multi-objective optimization. In: Search Methodologies. Springer, New York, USA, pp. 403–449.

Dunbar, M., Froyland, G., Wu, C.-L., 2012. Robust airline schedule planning: Minimizing propagated delay in an integrated routing and crewing framework. Transp. Sci. 46 (2), 204–216. http://dx.doi.org/10.1287/trsc.1110.0395.

Emmerich, M.T., Deutz, A.H., 2018. A tutorial on multiobjective optimization: fundamentals and evolutionary methods. Nat. Comput. 17 (3), 585–609. http://dx.doi.org/10.1007/s11047-018-9685-y.

Fricke, H., Schultz, M., 2009. Delay impacts onto turnaround performance. Optimal time buffering for minimizing delay propagation. In: Proceedings of 8th USA/Europe Air Traffic Management Research and Development Seminar. EUROCONTROL, Napa, USA, URL: https://www.atmseminar.org/seminarContent/seminar8/papers/p_153_AO.pdf.

Fu, M., 2015. Handbook of Simulation Optimization, first ed. Springer, New York, NY, USA, http://dx.doi.org/10.1007/978-1-4939-1384-8.

Hadka, D., Reed, P., 2013. Borg: An auto-adaptive many-objective evolutionary computing framework. Evol. Comput. 21 (2), 231–259. http://dx.doi.org/10.1162/EVCO_a_00075.

Hansen, M.M., Gillen, D., Djafarian-Tehrani, R., 2001. Aviation infrastructure performance and airline cost: a statistical cost estimation approach. Transp. Res. E 37 (1), 1–23. http://dx.doi.org/10.1016/S1366-5545(00)00008-9.

Hernandez-Diaz, A.G., Santana-Quintero, L.V., Coello Coello, C.A., Molina, J., 2007. Pareto-adaptive ϵ-dominance. Evol. Comput. 15 (4), 493–517. http://dx.doi.org/10.1162/evco.2007.15.4.493.

Jacquillat, A., Vaze, V., 2017. Balancing reliability, efficiency and equity in airport scheduling interventions. In: Proceedings of the 12th USA/Europe Air Traffic Management Research and Development Seminar. EUROCONTROL, Seattle, USA, URL: https://www.atmseminar.org/seminarContent/seminar12/papers/12th_ATM_RD_Seminar_paper_89.pdf.

Jarrah, A.I., Yu, G., Krishnamurthy, N., Rakshit, A., 1993. A decision support framework for airline flight cancellations and delays. Transp. Sci. 27 (3), 266–280. http://dx.doi.org/10.1287/trsc.27.3.266.

Kafle, N., Zou, B., 2016. Modeling flight delay propagation: A new analytical-econometric approach. Transp. Res. B 93, 520–542. http://dx.doi.org/10.1016/j.trb.2016.08.012.

Kasprzyk, J.R., Reed, P.M., Hadka, D.M., 2016. Battling Arrow's paradox to discover robust water management alternatives. J. Water Resour. Plan. Manag. 142 (2), 04015053. http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000572.

Katsigiannis, F.A., Zografos, K.G., Fairbrother, J., 2021. Modelling and solving the airport slot-scheduling problem with multi-objective, multi-level considerations. Transp. Res. C 124, 102914. http://dx.doi.org/10.1016/j.trc.2020.102914.

Kim, J.Y., Park, Y., 2012. Connectivity analysis of transshipments at a cargo hub airport. J. Air Transp. Manag. 18 (1), 12–15. http://dx.doi.org/10.1016/j.jairtraman.2011.05.001.

Kollat, J.B., Reed, P.M., 2006. Comparing state-of-the-art evolutionary multi-objective algorithms for long-term groundwater monitoring design. Adv. Water Resour. 29 (6), 792–807. http://dx.doi.org/10.1016/j.advwatres.2005.07.010.

Kollat, J.B., Reed, P.M., 2007. A computational scaling analysis of multiobjective evolutionary algorithms in long-term groundwater monitoring applications. Adv. Water Resour. 30 (3), 408–419. http://dx.doi.org/10.1016/j.advwatres.2006.05.009.

Kondo, A., 2011. Impacts of delay propagation on airline operations: Network vs. point-to-point carriers. In: 2011 Integrated Communications, Navigation, and Surveillance Conference Proceedings. IEEE, Herndon, USA, pp. L4–1–L4–7. http://dx.doi.org/10.1109/ICNSURV.2011.5935340.

Kwakkel, J.H., 2017. The Exploratory Modeling Workbench: An open source toolkit for exploratory modeling, scenario discovery, and (multi-objective) robust decision making. Environ. Model. Softw. 96, 239–250. http://dx.doi.org/10.1016/j.envsoft.2017.06.054.

Kwakkel, J.H., Haasnoot, M., Walker, W.E., 2016. Comparing robust decision-making and dynamic adaptive policy pathways for model-based decision support under deep uncertainty. Environ. Model. Softw. 86, 168–183. http://dx.doi.org/10.1016/j.envsoft.2016.09.017.

Lee, L.H., Lee, C.U., Tan, Y.P., 2007. A multi-objective genetic algorithm for robust flight scheduling using simulation. European J. Oper. Res. 177 (3), 1948–1968. http://dx.doi.org/10.1016/j.ejor.2005.12.014.

Lee, S.Y., Yoo, K.E., Park, Y., 2014. A continuous connectivity model for evaluation of hub-and-spoke operations. Transportmetr. A: Transp. Sci. 10 (10), 894–916. http://dx.doi.org/10.1080/23249935.2014.906515.

L'upták, V., Droździel, P., Stopka, O., Stopková, M., Rybicka, I., 2019. Approach methodology for comprehensive assessing the public passenger transport timetable performances at a regional scale. Sustainability 11 (13), 3532. http://dx.doi.org/10.3390/su11133532.

Lwin, K., Qu, R., Kendall, G., 2014. A learning-guided multi-objective evolutionary algorithm for constrained portfolio optimization. Appl. Soft Comput. 24, 757–772. http://dx.doi.org/10.1016/j.asoc.2014.08.026.

Mitchell, M., 1996. An Introduction to Genetic Algorithms. MIT Press, Cambridge, Massachusetts, USA.

Montlaur, A., Delgado, L., 2017. Flight and passenger delay assignment optimization strategies. Transp. Res. C 81, 99–117. http://dx.doi.org/10.1016/j.trc.2017.05.011.

Olsson, N.O., Haugland, H., 2004. Influencing factors on train punctuality—results from some Norwegian studies. Transp. Policy 11 (4), 387–397. http://dx.doi.org/10.1016/j.tranpol.2004.07.001.

Peterson, E.B., Neels, K., Barczi, N., Graham, T., 2013. The economic cost of airline flight delay. J. Transp. Econ. Policy 47 (1), 107–121, URL: https://www.jstor.org/stable/24396355.

Qin, S., Mou, J., Chen, S., Lu, X., 2019. Modeling and optimizing the delay propagation in Chinese aviation networks. Chaos 29 (8), 081101. http://dx.doi.org/10.1063/1.5111995.

Reed, P.M., Hadka, D., Herman, J.D., Kasprzyk, J.R., Kollat, J.B., 2013. Evolutionary multiobjective optimization in water resources: The past, present, and future. Adv. Water Resour. 51, 438–456. http://dx.doi.org/10.1016/j.advwatres.2012.01.005.

Riley, L.A., 2013. Discrete-event simulation optimization: a review of past approaches and propositions for future direction. In: Bruzzone, A.G., Kropf, P., Riley, L.A., Davoudpour, M., Solis, A.O. (Eds.), Proceedings of the 2013 Summer Simulation Conference. Society for Computer Simulation International, Toronto, Canada, p. 47, URL: http://dl.acm.org/citation.cfm?id=2557747.

Şafak, Ö., Gürel, S., Aktürk, M.S., 2017. Integrated aircraft-path assignment and robust schedule design with cruise speed control. Comput. Oper. Res. 84, 127–145. http://dx.doi.org/10.1016/j.cor.2017.03.005.

Salazar, J.Z., Reed, P.M., Herman, J.D., Giuliani, M., Castelletti, A., 2016. A diagnostic assessment of evolutionary algorithms for multi-objective surface water reservoir control. Adv. Water Resour. 92, 172–185. http://dx.doi.org/10.1016/j.advwatres.2016.04.006.

Santos, B.F., Wormer, M.M., Achola, T.A., Curran, R., 2017. Airline delay management problem with airport capacity constraints and priority decisions. J. Air Transp. Manag. 63, 34–44. http://dx.doi.org/10.1016/j.jairtraman.2017.05.003.

Singh, R., Reed, P.M., Keller, K., 2015. Many-objective robust decision making for managing an ecosystem with a deeply uncertain threshold response. Ecol. Soc. 20 (3), URL: https://www.jstor.org/stable/26270230.

Sohoni, M., Lee, Y.-C., Klabjan, D., 2011. Robust airline scheduling under block-time uncertainty. Transp. Sci. 45 (4), 451–464. http://dx.doi.org/10.1287/trsc.1100.0361.

Sternberg, A., Soares, J., Carvalho, D., Ogasawara, E., 2017. A review on flight delay prediction. Computing Research Repository, arXiv URL: https://arxiv.org/abs/1703.06118.

Thengvall, B.G., Bard, J.F., Yu, G., 2000. Balancing user preferences for aircraft schedule recovery during irregular operations. IIE Trans. 32 (3), 181–193. http://dx.doi.org/10.1023/A:1007618928820.

Tsai, S.C., Chen, S.T., 2017. A simulation-based multi-objective optimization framework: A case study on inventory management. Omega 70, 148–159. http://dx.doi.org/10.1016/j.omega.2016.09.007.

van Schilt, I.M., 2020. Balancing Connectivity and Reliability in the Schedule Design of Airlines by Buffer Scheduling: a Proof of Concept (Master's thesis). Delft University of Technology, URL: http://resolver.tudelft.nl/uuid:66051bfc-2ba8-4aee-89f8-c464c8075ab9.

Veiseth, M., Olsson, N., Saetermo, I., 2007. Infrastructure's influence on rail punctuality. WIT Trans. Built Environ. 96, 481–490. http://dx.doi.org/10.2495/UT070451.

Vikhar, P.A., 2016. Evolutionary algorithms: A critical review and its future prospects. In: Lobiyal, D., Mohapatra, D., Nagar, A., Sahoo, M. (Eds.), 2016 International Conference on Global Trends in Signal Processing, Information Computing and Communication. IEEE, Jalgaon, India, pp. 261–265. http://dx.doi.org/10.1109/ICGTSPICC.2016.7955308.

Ward, V.L., Singh, R., Reed, P.M., Keller, K., 2015. Confronting tipping points: Can multi-objective evolutionary algorithms discover pollution control tradeoffs given environmental thresholds? Environ. Model. Softw. 73, 27–43. http://dx.doi.org/10.1016/j.envsoft.2015.07.020.

Wong, J.-T., Tsai, S.-C., 2012. A survival model for flight delay propagation. J. Air Transp. Manag. 23, 5–11. http://dx.doi.org/10.1016/j.jairtraman.2012.01.016.

Wu, C.-L., 2005. Inherent delays and operational reliability of airline schedules. J. Air Transp. Manag. 11 (4), 273–282. http://dx.doi.org/10.1016/j.jairtraman.2005.01.005.

Wu, C.-L., 2006. Improving airline network robustness and operational reliability by sequential optimisation algorithms. Netw. Spat. Econ. 6, 235–251. http://dx.doi.org/10.1007/s11067-006-9282-y.

Wu, C.-L., 2008. Monitoring aircraft turnaround operations–framework development, application and implications for airline operations. Transp. Plan. Technol. 31 (2), 215–228. http://dx.doi.org/10.1080/03081060801948233.