KU LEUVEN

FACULTY OF ECONOMICS AND BUSINESS

A model enhancement heuristic for building robust aircraft maintenance personnel rosters with stochastic constraints

De Bruecker P, Van den Bergh J, Beliën J, Demeulemeester E.

KBI_1410

A model enhancement heuristic for building robust aircraft maintenance personnel rosters with stochastic constraints

Philippe De Bruecker*

KU Leuven, Research Center for Operations Management, Leuven (Belgium)

Jorne Van den Bergh

KU Leuven, Center for Information Management, Modeling and Simulation (CIMS), Brussels (Belgium)

Jeroen Beliën

KU Leuven, Center for Information Management, Modeling and Simulation (CIMS), Brussels (Belgium)

Erik Demeulemeester

KU Leuven, Research Center for Operations Management, Leuven (Belgium)

Abstract

This paper presents a heuristic approach to optimize staffing and scheduling at an aircraft maintenance company. The goal is to build robust aircraft maintenance personnel rosters that can achieve a certain service level while minimizing the total labour costs. Robust personnel rosters are rosters that can handle delays associated with stochastic flight arrival times. To deal with this stochasticity, a model enhancement algorithm is proposed that iteratively adjusts a mixed integer linear programming (MILP) model to a stochastic environment based on simulation results. We illustrate the performance of the algorithm with a computational experiment based on real life data of a large aircraft maintenance company located at Brussels Airport in Belgium. The obtained results are compared to deterministic optimization and straightforward optimization. Experiments demonstrate that our model can ensure a certain desired service level with an acceptable increase in labour costs when stochasticity is introduced in the aircraft arrival times.

Keywords: Model enhancement, aircraft maintenance, stochastic optimization

*Corresponding author

Email addresses: philippe.debruecker@kuleuven.be (Philippe De Bruecker), jorne.vandenbergh@kuleuven.be (Jorne Van den Bergh), jeroen.belien@kuleuven.be (Jeroen Beliën), erik.demeulemeester@kuleuven.be (Erik Demeulemeester)

1. Introduction

To ensure safety in aviation, aircraft should be maintained regularly and carefully. Constructing a good workforce schedule is therefore essential to make sure that all aircraft will be maintained thoroughly in time. In the aviation industry, different forms of maintenance exist. We distinguish A-, B-, C- and D-checks, line maintenance, hangar maintenance, scheduled and unscheduled maintenance, etc. (Van den Bergh et al., 2013b). In this paper we focus on the line maintenance which includes pre-flight inspections, transit checks, daily checks (visual inspection of the aircraft, fluid level checks, general security checks, emergency equipment checks and cleanliness of the flight deck checks), weekly checks and on-call assistance (Beliën et al., 2012).

Scheduling aircraft maintenance personnel at an aircraft maintenance company entails some special problems. First, the workforce scheduling problem is heavily constrained by labour union agreements. Second, the management of the company must decide itself when exactly the maintenance should take place between the arrival and departure of the aircraft. Hence, the timing of the workload is an extra decision in the scheduling problem. Third, aircraft do not always fly on schedule and sometimes arrive with a delay. When the workforce scheduling does not anticipate delays in arrival time, the scheduled capacity may be insufficient to maintain all aircraft in time.

This paper focuses on this latter problem and presents a technique to obtain robust aircraft maintenance personnel rosters that minimize the total labour costs. Because aircraft arrive with stochastic delays, we define the stochastic robustness of these rosters as their ability to ensure a certain service level; i.e., to ensure that on average at least a certain percentage of the flights can be maintained before their scheduled time of departure (STD). To obtain this stochastic robustness, we propose a model enhancement (ME) heuristic that iteratively enhances a mixed integer linear programming (MILP) model by adding constraints based on information resulting from simulation experiments.

We illustrate our model using real life data from Sabena Technics, a large aircraft maintenance company located at Brussels Airport in Belgium.

2. Literature review

The first part of this literature review situates this paper within the existing literature on aircraft maintenance scheduling, while the second part reviews the literature on stochastic optimization.

The problem under study concerns a workforce scheduling problem for line maintenance personnel. For an overview of the literature on general workforce scheduling, we refer to Van den Bergh et al. (2013a), while a literature overview of aircraft maintenance operations can be found in Van den Bergh et al. (2013b). Many of the studied aircraft maintenance optimization problems concern maintenance routing optimization (see, e.g., (Talluri, 1998; Sarac et al., 2006; Liang et al., 2010; Sriram and Haghani, 2003)), which addresses the problem of finding optimal sequences of flights for a particular aircraft such that it can be parked in a maintenance station after a certain number of days of flying without maintenance. This problem is often solved for one airline company.

We assume, however, that the maintenance routing problem is already solved and the routes are given for several airline companies. Given a set of flights that arrive and depart at a certain airport, we try to optimize the maintenance workforce schedule at the associated maintenance station. The approach is tested on a line maintenance environment, excluding other types of maintenance, such as light and heavy maintenance. The latter include more demanding maintenance checks and are typically subject to a higher variability in workload, whereas the former is limited to routine maintenance of which the required workload can quite well be predicted beforehand. Other papers that study line maintenance problems in comparable settings to this paper include (Beliën et al., 2012), (Beliën et al., 2013), (Van den Bergh et al., 2013c), (Yan et al., 2004) and (Papakostas et al., 2010).

In contrast with (Quan et al., 2007) who present a multi-objective model that simultaneously minimizes the maintenance capacity and the makespan of completing all maintenance jobs, our model involves a single objective (minimizing maintenance personnel costs) while the other objective (minimizing the number of late flight departs due to late maintenance) is modeled through a service level constraint. The problem studied in (Safaei et al., 2011) and (Safaei et al., 2010) also requires a certain number of aircraft to be available (on time) for the next fly programs, and, hence studies a similar service level constraint. An important difference is, however, the context of a military aircraft fleet, which involves a much smaller number of aircraft (< 10) compared to our study incorporating the commercial aircraft fleet of different airline companies.

The model presented in this paper aims at developing a robust maintenance capacity schedule taking into account uncertainty in the flight arrivals and, hence, in the timing of the workload. This is a major distinction with the problem studied in (Yan et al., 2004), which is closely related to our work, except for the assumption of deterministic flight arrivals. Other papers that study aircraft maintenance scheduling assuming uncertainty in the timing of the workload include (Duffuaa and Andijani, 1999), (Mattila et al., 2008), (Petersen et al., 2012) and (Muchiri, 2009). The papers of (Safaei et al., 2011) and (Safaei et al., 2010) assume deterministic flight arrivals, but stochastic workloads (and thus stochastic maintenance durations or repair times).

Robust optimization methods are used when the uncertain parameters are assumed to take values from a certain range (Bertsimas & Sim, 2004). These methods differ from stochastic optimization methods where the uncertain parameters follow a certain probability distribution. In this paper, we are concerned with stochastic optimization as we assume that the delays in aircraft arrival times follow a certain distribution.

Several techniques have already been proposed in the literature to cope with the problems of stochastic optimization. However, no single best solution method exists. When the problem becomes very complex, simulation is often used in combination with optimization (Fu, 2002). When it comes to handling stochastic elements, regular optimization methods have some major shortcomings. Simulation on the other hand fails to implement the decision element and only evaluates a solution in a stochastic environment. Hence, combining both methods may lead to a powerful approach (Fagerholt et al., 2010). Complex simulation-optimization couplings have been described in the literature in order to make a solution more robust and applicable in a real life (stochastic) environment.

In simulation-optimization couplings, different solutions are evaluated with simulation and the search is usually guided through different solutions to obtain a good (or the best) feasible result. Chen et al. (2011) distinguish between three different subcategories of simulation optimization. Total enumeration always leads to the global optimum by evaluating all possible solutions with simulation. The second subcategory is the gradient approach. This approach imitates gradient methods in deterministic optimization methods to carry out a guided local search. Those approaches are also called model based approaches. The third subcategory are gradient-free approaches. Chen et al. (2011) refer to these as the metaheuristic approaches. The second and third subcategory can be described as guided search methods. They do not enumerate all possible solutions, and hence, do not necessarily lead to the global optimum.

In this paper, we use a technique called "model enhancement". The term *model enhancement* (ME) was used by Bachelet et al. (2007) to indicate a different way of combining simulation and optimization. It can be seen as a decomposition method like Benders' decomposition as it enhances an optimization model based on simulation results by adding constraints (Benders, 1962). While most optimization-simulation couplings focus on improving the objective function evaluated from simulation (like the simulation optimization approach), ME still focuses on optimizing the theoretical objective function. It tries to improve the solution provided by a mathematical model by the use of simulation (Bachelet et al., 2007). In their paper, Bachelet et al. (2007) only consider a deterministic problem. They assume that in practice, several modeling simplifications are needed to construct the mathematical optimization model to solve a real life problem. The resulting model therefore fails to give a correct representation of reality. Simulation is used to enhance the mathematical model and to improve the realism and applicability of the solution. In our research, we follow a similar approach, but we consider a stochastic problem. We start with a model that neglects stochasticity and use simulation to enhance the deterministic model to produce robust solutions.

3. Problem definition

At many aircraft maintenance companies, the demand for maintenance is seasonal and changes only twice a year. Furthermore, each week, the same set of flights have to be maintained. Table 1 gives a typical example of the demand faced by an aircraft maintenance company. Since the timing of the workload is a decision variable in most cases, only the scheduled time of arrival (STA) and the scheduled time of departure (STD) of each aircraft are given. Next to the time interval during which the maintenance should take place, the estimated workload in man-hours is given. Given such a recurring workload for each flight during one week, the maintenance company has to find the optimal workforce configuration in terms of the lowest total labour costs that can assure the desired service level. Such a workforce configuration defines when and how many employees should be scheduled on each day during a whole week.

Table 1: Example of the demand for aircraft maintenance

Flight	Company	STA	STD	Workload (man-hours)
1	BA	Monday 22:05	Tuesday 07:15	4
...	...	...	...	...
111	AA	Thursday 07:30	Thursday 10:40	4.25
...	...	...	...	...

Finding a feasible workforce configuration is not only constrained by satisfying the demand for maintenance services while minimizing the costs. Real life staffing and scheduling problems are also constrained by strict union regulations such as average working times, working in the weekend, break times, shift succession constraints, etc. Furthermore, aircraft do not always fly on schedule and sometimes arrive with a delay. Unexpected delays can disarrange the maintenance planning which can lead to a delay in the flight's departure time when capacity buffers are absent or too small. Of course, an aircraft maintenance company wants to ensure that 100% of the flights can be maintained in time when all flights would arrive in time (i.e., according to the STA of each aircraft). However, because flights arrive with stochastic delays, the company only allows for at most a certain average percentage of the flights (depending on the desired service level) to be maintained after their STD. We will refer to this stochastic constraint as the **service level constraint** in the optimization problem.

4. Methodology

In this section, we discuss the model enhancement (ME) technique to solve the optimization problem with the stochastic service level constraint. First, an overview of the ME algorithm is presented. Next, the three phases of the ME algorithm are discussed.

4.1. Model Enhancement

To account for the stochastic **service level constraint** in the optimization problem, model enhancement is used. With this technique, we combine simulation and optimization in order to incorporate stochastic constraints by iteratively enhancing a deterministic optimization model. During this process, the focus remains on optimizing the objective function provided by a mathematical model. Figure 1 shows the three consecutive phases that are repeated in the enhancement algorithm. One loop through phases I to III is referred to as one enhancement iteration. The enhancement iterations are indexed by index δ ranging from 0 to Δ, with Δ the total number of enhancement iterations performed.

In phase III, a mixed integer linear programming (MILP) model is solved to find the optimal workforce configuration that minimizes the total labour costs. The obtained solution provides the optimal capacity that must be available at each time period in order to maintain all aircraft in time with minimum labour costs. These MILP models of phase III are indexed as $MILP^{\delta}$, with δ ranging from 0 to Δ.

Figure 1: Model enhancement algorithm

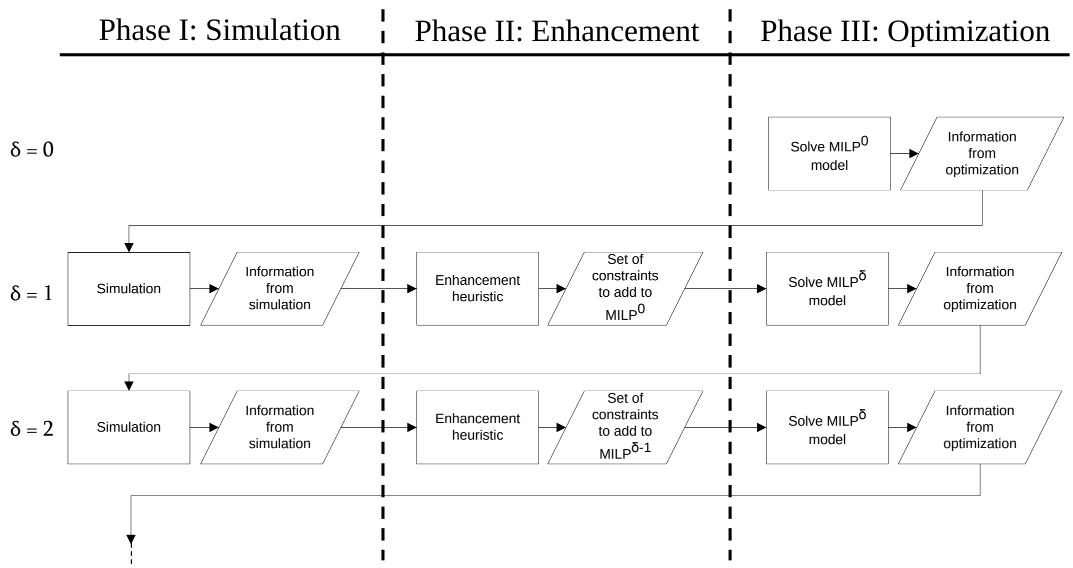

In phase I, the available capacity determined by the $MILP^{\delta-1}$ model is used to initialize a simulation model. This simulation model will simulate the assignment of workers to aircraft to perform the required maintenance during several weeks. During the simulation, aircraft arrive with stochastic delays according to a certain probability distribution based on real data. Because of possible delays, not all flights will be maintained in time. Therefore, the simulation model will result in a set of flights that could not be (completely) serviced before their STD. It will also return the average number of flights per week that cannot be maintained before the flight's STD. This is also called the average number of flights with late maintenance.

At the start of phase II, it is checked whether the service level constraint is satisfied. This constraint ensures that the average number of flights with late maintenance does not exceed the value defined in the service level agreement. When this constraint is not satisfied, the algorithm increases the capacity requirements with a certain amount during certain time intervals. To decide when and by how much the capacity requirements should be increased, a heuristic is used. This enhancement heuristic uses the information from the simulation experiment to add new constraints to the MILP model or to change the parameters of the current constraints in the MILP model. This way, solving the MILP model in the next phase will increase the available capacity around certain time periods resulting in a capacity buffer. This allows more flights to be maintained in time and hopefully to satisfy the service level constraint. If the service level constraint is satisfied at the beginning of phase II, the enhancement heuristic tries to improve the current solution by diversification. At the end of phase II, the constraints in the MILP model are adjusted and the enhancement procedure moves to phase III.

During phase III, the MILP model is solved with modified constraints. Solving this adjusted MILP model will result in a new workforce configuration with the lowest total labour costs that is more capable of handling stochastic delays with respect to the desired service level. Again, the obtained solution provides the optimal capacity at each time period which will be the input for the simulation model in phase I during the next enhancement iteration.

These three phases are successively executed until Δ enhancement iterations have passed. The next sections describe these three phases one by one.

4.2. Phase III: Optimization

The core of the ME algorithm is the MILP model that finds the optimal workforce configuration. In this section we describe the MILP formulation of the general maintenance scheduling model. In Section 5.1, the general objective function (1) and the general constraints (2) and (5) are elaborated to represent a real life problem.
We first list the parameters and sets, along with their associated indices:

$d \in \{1, 2, ..., 7\}$:	days in the week
$p \in P$:	time periods in one week
$f \in F$:	flights to be serviced
$p \in P_f$:	time periods during which flight f can be serviced. I.e.; $P_f = \{p \mid STA_f \leq p \leq STD_f\}$ with STA_f (STD_f) the scheduled time of arrival (departure) of flight f
L_f :	the workload (in man-hours) of flight f

The decision variables are:

$q_p^\delta \geq 0$:	the available capacity on period p when solving the MILP model in enhancement iteration δ
$g_{fp}^\delta \geq 0$:	the number of workers assigned to maintain flight f during time period p when solving the MILP model in enhancement iteration δ
$\mathbb{X}$:	set of different workforce scheduling variables. These can be for example decision variables that define the timing and scheduling of shifts, team sizes, etc.

The optimization model can be formulated as follows:

$$
\begin{aligned}
& \textit{Minimize}: \\
& \qquad \text{Cost function } k(\mathbb{X}) && && (1) \\
& \textit{Subject to}: \\
& q_p^\delta = \text{Capacity function } \kappa(\mathbb{X}), && \forall p \in P && (2) \\
& q_p^\delta \geq \sum_{f \in F} g_{fp}^\delta, && \forall p \in P && (3) \\
& \sum_{p \in P_f} g_{fp}^\delta = L_f * \frac{|P|}{24 * 7}, && \forall f \in F && (4) \\
& \text{Other, specific workforce scheduling constraints} && && (5) \\
& g_{fp}^\delta \geq 0, \qquad \forall f \in F && \forall p \in P && (6) \\
& q_p^\delta \geq 0, && \forall p \in P && (7)
\end{aligned}
$$

The objective function (1) describes the total labour costs faced by the company resulting from the decisions for the workforce variables in $\mathbb{X}$. These can be decisions regarding the scheduling of different types of shifts, e.g., night and day shift, the timing of those shifts, the team sizes in each shift, etc. Based on those decisions, a capacity function is used to calculate the available capacity q_p^δ for each time period p in constraint (2). Constraints (3) and (4) ensure that the assigned work for each flight (and therefore the available capacity) is sufficient to at least cover the demand for maintenance. This is referred to as the coverage constraint. In Constraint (4), $L_f * \frac{|P|}{24*7}$ is a conversion factor to match the units of L_f and $\sum_{p \in P_f} g_{fp}^\delta$. Real life staffing and scheduling problems are however not only constrained by satisfying the demand for maintenance services while minimizing the costs. Therefore, constraint (5) adds other, specific workforce scheduling constraints, such as constraints on average working times, working in the weekend, break times, shift succession constraints, etc.

During the first enhancement iteration ($\delta = 0$), we define the first MILP model as $MILP^0$. $MILP^0$ contains the original objective function and all original constraints. It can be seen as the deterministic MILP model as it assumes that all aircraft exactly arrive at their scheduled time of arrival.

Definition of $MILP^0$ (when $\delta = 0$):

$$
\begin{aligned}
& \textit{Optimize}: \\
& \qquad \text{Objective function (1)} \\
& \textit{Subject to}: \\
& \qquad \text{Constraints (2) to (7)}
\end{aligned}
$$

To account for the stochastic service level constraint that limits the number of flights that can leave with a delay, extra constraints will be added to the MILP model described above. During the next enhancement iterations ($\delta = 1$ to Δ), the $MILP^{\delta}$ model is also defined by the original objective function (1), but the model has different constraints:

Definition of $MILP^{\delta}$ (with $\delta = 1$ to Δ):

$$\begin{aligned}
&\mathit{Optimize}: && \\
& && \text{Objective function (1)} \\
&\mathit{Subject\ to}: && \\
& && \text{Constraints (2) and (3)}
\end{aligned}$$

$$\sum_{p \in P_f} g_{fp}^{\delta} \geq L_f * \frac{|P|}{24 * 7}, \qquad \forall f \in F \tag{8}$$

$$\sum_{p \in P_f^{\delta}} g_{fp}^{\delta} \geq L_f^{\delta} * \frac{|P|}{24 * 7} + \sum_{p \in P_f^{\delta}} g_{fp}^{\delta-1}, \qquad \forall f \in F^{\delta} \tag{9}$$

Constraints (5) to (7)

with parameters:

- $g_{fp}^{\delta-1} \geq 0$: the number of workers assigned to maintain flight f during time period p according to the results of the MILP model of the previous enhancement iteration.
- L_f^{δ}: This is the parameter used to increase the capacity requirements in the MILP model during enhancement iteration δ. L_f^{δ} results from the simulation model and is defined by the enhancement function (see Section 4.4.2). L_f^{δ} is the extra needed maintenance work (in man-hours) for flight f to ensure that this flight leaves without a delay.

and sets:

- $f \in F^{\delta}$: set of flights for which the extra required maintenance work L_f^{δ} must be added during enhancement iteration δ. Like L_f^{δ}, F^{δ} also results from the simulation model and is defined by the enhancement function (see Section 4.4.2).
- $p \in P_f^{\delta}$: set of all time periods p during which the extra required maintenance work L_f^{δ} can be scheduled for flight f during enhancement iteration δ. Like L_f^{δ} and F^{δ}, P_f^{δ} also results from the simulation model and is defined by the enhancement function (see Section 4.4.2).

In $MILP^{\delta}$, constraint (4) is replaced by constraints (8) and (9). Instead of the *equal* sign in constraint (4), constraint (8) uses the *greater or equal* sign. This means that the assigned work $\sum_{p\in P_f} g^{\delta}_{fp}$ for a certain flight can exceed the standard required maintenance work for that flight. Hence, this allows to assign extra work to a flight to build a capacity buffer to avoid delays in departure time.

Constraint (9) tries to avoid delays in departure time for all flights in the set F^{δ}. $\sum_{p\in P^{\delta}_f} g^{\delta}_{fp}$ captures the assigned work to flight f in the time interval P^{δ}_f during enhancement iteration δ. $\sum_{p\in P^{\delta}_f} g^{\delta-1}_{fp}$ represents the assigned work to flight f in the time interval P^{δ}_f during the previous enhancement iteration. L^{δ}_f results from the simulation model in phase I and is defined by the enhancement function (see Section 4.4.2). It is the extra needed maintenance work (in man-hours) for flight f to ensure that this flight leaves without a delay. Hence, when according to the simulation, flight f could not be maintained in time, the assigned work $\sum_{p\in P^{\delta}_f} g^{\delta-1}_{fp}$ to flight f during the previous enhancement iteration was apparently insufficient to maintain this flight in time. Therefore, the extra work L^{δ}_f must be added during this enhancement iteration to make this flight leave in time.

The information about P^{δ}_f, F^{δ} and L^{δ}_f will be made available after applying the enhancement function in phase II.

4.3. Phase I: Simulation

The first step in each enhancement iteration is to run a simulation model to evaluate the MILP solution in a stochastic environment. The simulation model simulates the assignment of workers to aircraft to perform the required maintenance during several weeks. Hence, a simulation replication simulates the maintenance during one week. The simulation model is initialized with the available capacity $q^{\delta-1}_p$ resulting from solving the MILP model during the previous enhancement iteration. During the simulation, aircraft arrive with a delay according to certain probability distributions based on real life data. The simulation model then allocates on each time period a certain number of maintenance workers to grounded aircraft (i.e., to aircraft that can be maintained during that time period). This is done according to a specific allocation rule.

Because of stochastic delays in arrival times, it is not always possible to maintain all flights in time (i.e., before the aircraft's STD). When this is the case, the remaining required maintenance is finished as soon as possible. Because of delays, a flight can require maintenance after its STD leading to more delays in successive flights. At the end of each simulated week ω, all flights that left with a delay are saved in the subset Φ_{ω}. The maintenance work that has not yet been done at the STD of flight f is saved as $\lambda_{f,\omega}$. The time window in which that flight could be maintained according to the simulation model is saved as $\Pi_{f,\omega}$.

We define the results from the simulation as follows:

Ω:	the set of all weeks ω included in the simulation experiment.
Φ_{ω} $\quad \forall\omega\in\Omega$:	the subset of all flights f that left with a delay during simulated week ω.
Φ:	the set of all different flights f that left with a delay over all simulated weeks.

$\varphi = \frac{\sum_{\omega \in \Omega} |\Phi_\omega|}{|\Omega|}$: the average number of flights per week that leave with a delay. This is also called the average number of flights with late maintenance.

$\lambda_{f,\omega}$: the work that has not yet been done at the STD of flight f with late maintenance in simulated week ω.

Ω_f: the set of weeks ω in which flight f left with a delay.

$\lambda_f = \frac{\sum_{\omega \in \Omega_f} \lambda_{f,\omega}}{|\Omega_f|}$: the average work that has not yet been done at the STD of flight f. This is the extra needed work in order to maintain flight f in time.

$\Pi_{f,\omega}$: the set of all time periods p between the simulated arrival time and the STD of flight f with late maintenance in week ω.

Π_f: the set of all time periods p between the average simulated arrival time (over all weeks in Ω_f) and the STD of flight f.

4.4. Phase II: Enhancement

Based on the simulation results, the MILP model will be enhanced to produce a more robust workforce configuration. Recall that this is a workforce configuration that can achieve a desired service level under stochastic delays in flight arrival times. Constraint (10) formally describes the service level constraint.

$$\begin{aligned} \varphi \leq\ & \text{Allowed average number of delayed flights} \\ & \text{according to the desired service level} \end{aligned} \tag{10}$$

This constraint implies that the average number of flights per week that leave with a delay (also called the average number of flights with late maintenance) must be smaller than or equal to the number that is allowed by the company. When the service level is for example 97%, at most 3% of all weekly flights (on average) can leave with a delay. Figure 2 shows the different steps of phase II.

Figure 2: Phase II: The enhancement procedure

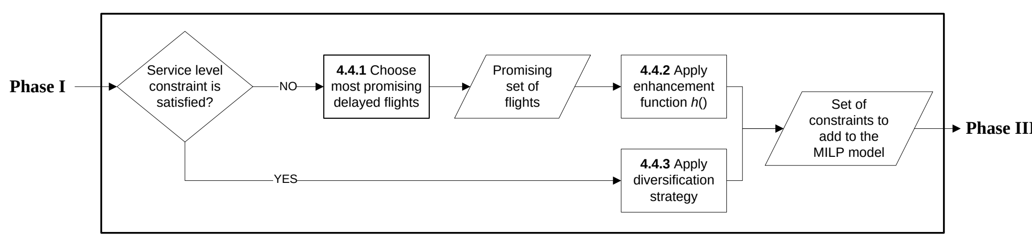

4.4.1. Choose most promising delayed flights

At the start of phase II, it is checked whether the service level constraint (10) is satisfied. When this constraint is not satisfied, we force the MILP model to increase the available capacity with a certain amount in a certain time interval for certain flights. Hence, the first step is to choose a promising set of flights for which the required maintenance work will be increased. This should be a set of flights that regularly receive late maintenance, i.e., that regularly leave with a delay. We first sort the flights in the set Φ based on the number of times that the flight leaves with a delay. Hence, the flights in Φ are sorted by decreasing value of $|\Omega_f|$. To create this set of promising flights, the first η flights are selected from the sorted set Φ according to Equation (11). These selected flights are saved in the set Φ' such that $|\Phi'| = \eta$. This set contains "promising" flights because allocating more work to a flight of set Φ' offers the best chance to decrease the value of φ.

$$\eta = \lceil \varphi - \text{Allowed average number of delayed flights} \rceil \tag{11}$$

4.4.2. Apply enhancement function

After the set Φ' has been composed, we have to decide how much extra work each flight in Φ' should receive in what time interval in order to maintain that flight completely before its STD. These decisions are made by the enhancement function h which represents the adjustments to the MILP model at each enhancement iteration. In this section, we propose a simple and intuitive formulation of function h which yielded good test results. Bachelet et al. (2007) also showed that good results can be obtained with a very simple enhancement function.

Recall that we defined the results of phase II as follows:

- F^δ: the set of flights for which the required maintenance work will be increased in the MILP model of enhancement iteration δ.
- L_f^δ: the extra required maintenance work that will be added to flight f in the MILP model of enhancement iteration δ.
- P_f^δ: the set of all time periods p during which the extra required maintenance work can be scheduled for flight f during enhancement iteration δ.

Definition of h for $\delta = 1$:

During the first enhancement iteration (i.e., the first loop through the three phases in Figure 1), phase II is executed for the first time. Therefore, there were no previous enhancements to the MILP model and the enhancement function h maps the simulation results of phase I as follows:

$$h : \Phi' \mapsto F^\delta \quad \text{by } F^\delta = \Phi' \tag{12}$$

$$h : \lambda_f \mapsto L_f^\delta \quad \text{by } L_f^\delta = \lambda_f \tag{13}$$

$$h : \Pi_f \mapsto P_f^\delta \quad \text{by } P_f^\delta = \Pi_f \tag{14}$$

Expressions (12) to (14) show that during the first enhancement iteration, the enhancement function h maps the simulation information Φ', λ_f and Π_f directly to F^δ, L_f^δ and P_f^δ respectively. F^δ, L_f^δ and P_f^δ will be used to build the extra constraints that will be added to the MILP model in phase III.

Definition of h for $\delta = 2$ to Δ:
During each enhancement iteration, the information resulting from the current iteration is saved such that it can be used in the next iteration.

We define:

$F^{\delta-1}$: the set of flights for which the required maintenance work was increased during the previous enhancement iteration.

$L_f^{\delta-1}$: the extra required maintenance work that was added to flight f during the previous enhancement iteration.

$P_f^{\delta-1}$: the set of all time periods p during which the extra required maintenance work could be scheduled for flight f during the previous enhancement iteration.

During each new enhancement iteration, the enhancement function h makes use of this previous information to map the current simulation results as follows:

$$h : \Phi', F^{\delta-1} \mapsto F^{\delta} \text{ by } F^{\delta} = \Phi' \cup F^{\delta-1} \tag{15}$$

$$h : \lambda_f, L_f^{\delta-1} \mapsto L_f^{\delta} \text{ by } L_f^{\delta} = \begin{cases} \lambda_f + L_f^{\delta-1} & \forall f \in (\Phi' \cap F^{\delta-1}) \\ L_f^{\delta-1} & \forall f \in (F^{\delta-1} \backslash \Phi') \\ \lambda_f & \forall f \in (\Phi' \backslash F^{\delta-1}) \end{cases} \tag{16}$$

$$h : \Pi_f, P_f^{\delta-1} \mapsto P_f^{\delta} \text{ by } P_f^{\delta} = \begin{cases} P_f^{\delta^{Avg}} & \forall f \in (\Phi' \cap F^{\delta-1}) \\ P_f^{\delta-1} & \forall f \in (F^{\delta-1} \backslash \Phi') \\ \Pi_f & \forall f \in (\Phi' \backslash F^{\delta-1}) \end{cases} \tag{17}$$

with:

α_f : first element in $P_f^{\delta-1}$
= previous average arrival time of flight f

β_f : first element in Π_f
= current average arrival time of flight f

$P_f^{\delta^{Avg}}$ $= [\frac{\alpha_f + \beta_f}{2}, ..., STD_f]$

Expressions (15) to (17) formally define the idea that previous and current information about the same flights can be merged. Expression (15) shows that the set $F^{\delta-1}$ is updated with the current promising delayed flights Φ' to constitute the set F^{δ}.

Expression (16) shows how function h maps λ_f and $L_f^{\delta-1}$ to L_f^{δ}. We distinguish three cases. In the first case (i.e., $f \in (\Phi' \cap F^{\delta-1})$), a promising delayed flight f (in Φ') already exists in $P_f^{\delta-1}$. This means that the extra capacity assigned to this flight during previous enhancement iterations was apparently not sufficient to maintain this flight in time. Therefore, the extra required maintenance work (λ_f) for that flight is added to $L_f^{\delta-1}$. In case two (i.e., $f \in (F^{\delta-1} \backslash \Phi')$), flight f is in the set $F^{\delta-1}$ but not in the set Φ'. In this case, the extra capacity assigned to this flight during previous enhancement iterations was sufficient to maintain this flight in time and we don't need to change the required maintenance for this flight. Therefore, function h sets L_f^δ equal to $L_f^{\delta-1}$. In the third case (i.e., $f \in (\Phi' \backslash F^{\delta-1})$), flight f is in the set Φ' but not in the set $F^{\delta-1}$. In this case, it is the first time that this flight enters the set of promising delayed flights. Because there is no extra capacity previously assigned to this flight, L_f^δ is set equal to λ_f.

To define the set P_f^δ, Expression (17) is used. Again, the same three cases as in Expression (16) are distinguished. When flight f (in Φ') already exists in $P_f^{\delta-1}$, the arrival time of flight f in both sets is averaged. $P_f^{\delta Avg}$ is then defined as the set of all time periods between the average simulated arrival time and the scheduled time of departure (cfr. supra). The reasoning behind case two and three of Expression (17) is similar to the one behind case two and three of Expression (16).

4.4.3. Apply diversification strategy

Figure 1 shows that the ME procedure starts by finding a shift configuration that is optimal for the deterministic $MILP^0$ model. After collecting the simulation results in phase I, the enhancement procedure is applied in phase II as shown in Figure 2. When the service level constraint is satisfied during the first enhancement iteration (i.e., $\delta = 1$), the whole enhancement procedure stops and the solution to the problem equals the solution to the deterministic $MILP^0$ model. When the service level constraint is not satisfied during the first enhancement iteration, the enhancement functions (12) to (14) are applied and the procedure continues. When the service level constraint is not satisfied during the next enhancement iterations (i.e., $\delta > 1$), the enhancement functions (15) to (17) are applied. When the service level constraint is satisfied during the next enhancement iterations, a local optimum is reached. To escape this local optimum, a diversification strategy is applied. To diversify the search, a different enhancement function is applied as long as the service level constraint remains satisfied:

$$h : F^{\delta-1} \mapsto F^\delta \quad \text{by} \quad F^\delta = F^{\delta-1} \backslash F^{\delta-1, L_f^{\delta-1} \leq \Psi \cdot L_f} \tag{18}$$

$$h : L_f^{\delta-1} \mapsto L_f^\delta \quad \text{by} \quad L_f^\delta = (1-\Xi) \cdot L_f^{\delta-1} \qquad \forall f \in F^\delta \tag{19}$$

$$h : P_f^{\delta-1} \mapsto P_f^\delta \quad \text{by} \quad P_f^\delta = P_f^{\delta-1} \qquad \forall f \in F^\delta \tag{20}$$

with:

$F^{\delta-1, L_f^{\delta-1} \leq \Psi \cdot L_f}$:	the set of the flights in $F^{\delta-1}$ that did not need more than Ψ% extra required maintenance work ($L_f^{\delta-1}$) compared to their scheduled workload (L_f)
Ξ:	percentage reduction in the extra required maintenance work

Expressions (18) to (20) are used to diversify the search and to guide the search through different possibilities. To achieve this, Function (19) reduces the value of $L_f^{\delta-1}$ by Ξ% for all flights in F^δ and, hence, reduces the impact of constraint (9) on the optimization model. Function (18) even eliminates a flight from $F^{\delta-1}$ when it needs less than Ψ% extra required maintenance work compared to its scheduled workload. This way, a flight will leave the set $F^{\delta-1}$ after several iterations and different flights can enter the set. This allows different possibilities to be explored, possibly leading to a better solution. The value of the diversification parameters $\Psi\%$ and Ξ is determined based on empirical tests.

5. Results and discussion

5.1. Application to a real life problem

To test the performance of the algorithm, we applied the model enhancement procedure to a real life problem at Sabena Technics, a large aircraft maintenance company located at Brussels Airport in Belgium. The main goal is to build the cheapest workforce configuration while satisfying the service level constraint.

To understand what is meant by a workforce configuration, the workforce system of Sabena Technics is explained first. At Sabena Technics, the maintenance personnel works in teams of a certain team size. Each team can only work one shift per day or have a day off. There are four possible shift types with overlapping working hours: a morning shift (M), a day shift (D), an evening shift (E) and a night shift (N). This results in a certain sequence of shift types and days off assigned to a certain team for a certain week.

Sabena Technics also organizes its workforce into two cycles as shown in Figure 3. Each cycle has its own team size and its own start and end times of each shift type. Because one cycle can contain multiple weeks (represented by the rows in each table), a team is not always assigned to the same shift sequence as the week before. Instead, a cyclic pattern is used in which the number of teams assigned to a cycle equals the number of weeks (i.e., the number of rows) in the cycle.

Figure 3: Example of a workforce configuration consisting of two cycles

Team size: 2
Morning: 05:00 - 14:00 Night: 21:30 - 07:00
Evening: 12:45 - 21:45 Day: 07:00 - 17:00

mo	tu	we	th	fr	sa	su
	N	N	N	N		

Team size: 3
Morning: 05:30 - 15:00 Night: 22:30 - 08:00
Evening: 12:45 - 22:15 Day: 07:00 - 16:30

mo	tu	we	th	fr	sa	su
N			D	E		
M	M		D	N		E
	M	E		E		
	M	E		N	N	N

In Figure 3, there is only one week in the first cycle. Therefore, there is only one team that is scheduled every week to work a night shift from Tuesday till Friday. In the second cycle there are four teams scheduled. In the first week, the first team works a night shift on Monday, a day shift on Thursday and an evening shift on Friday. The second team works a morning shift on Monday and Tuesday and so on. In the second week, the first team will work the shifts in the second row of the cycle, the second team the shifts of the third row and so on. To complete the cycle, the fourth team will then work the shifts of the first row.

Hence, constructing a workforce configuration implies that we have to decide on the number of weeks (= the number of teams = the number of rows) in each cycle, the assignment of shift types and days off to each day in each week, the start and end times of every shift type in each cycle and the team size in each cycle.

5.1.1. Optimization model (Phase III)

To solve the problem at Sabena Technics, we use a MILP model based on the model of Beliën et al. (2013). Interested readers are referred to latter work for a more profound elaboration of this MILP model and the solution technique. This MILP model has the exact same structure as the model described in Section 4.2, but we now give a specific description of the general cost function, capacity function and other workforce scheduling constraints.

Decisions about the optimal team size are not made within the model, because this would render the model non-linear. Instead, a heuristic enumeration algorithm is employed to find the best team size for each cycle (see Beliën et al. (2013)). In this enumeration heuristic, the MILP model is solved for a limited amount of time for each interesting team size combination and the best result is saved.

As was mentioned in Section 5.1, we also need to decide on the assignment of shift types and days off to each day in each week. However, incorporating the week decision in the MILP model implies too many decision variables to solve the problem efficiently. Therefore, we omit the assignment of shifts to different weeks and our scheduling decision variable only captures the number of shifts of a certain type scheduled on each day of the week. After solving the MILP model, the individual weeks can be reconstructed by hand by distributing the shifts over the different weeks for each day. The shift succession guaranteeing constraints in the MILP model make sure that a feasible workforce schedule can be constructed at the end.

We first list the sets, along with their associated indices:

$i, i' \in I$:	feasible shifts
$t \in T$:	shift types $\{M, D, E, N\}$
$i \in I_t$:	feasible shifts of shift type t (Different start and end times are possible for each shift type.)
$c \in C$:	cycles in the schedule
$d \in \{1, 2, ..., 7\}$:	days in the week
$p \in P$:	time periods in one week
$f \in F$:	flights to be serviced
$p \in P_f$:	time periods during which flight f can be serviced. I.e.; $P_f = \{p \mid STA_f \leq p \leq STD_f\}$ with STA_f (STD_f) the scheduled time of arrival (departure) of flight f

The coefficients and right hand side constants are presented below.

k_{id} :	the total cost (for one worker) of shift i on day d
h_i :	the duration of shift i (in hours)
a_{idp} :	$= 1$ if period p is included in shift i on day d; $= 0$ otherwise
b_{idp} :	the fraction of workers available to work in shift period p when assigned shift i on day d (< 1 during lunch hour)

L_f : the workload (in man-hours) of flight f
W^l : the minimum number of weeks (= teams) in a cycle
W^u : the maximum number of weeks (= teams) in a cycle
S : the minimum average number of working hours per week
U : the maximum average number of working hours per week
R : the maximum fraction of working weekends; i.e., weekends during which at least one shift is scheduled. R = (number of working weekends)/(number of weeks in the cycle)
M_c: the team size in cycle c
$|P|$: the number of time periods in a week

The decision variables are:

$q_p^\delta \geq 0$: the available capacity on period p when solving the MILP model in enhancement iteration δ
$n_c \in \{W^l, W^l+1, ..., W^u\}$: the number of weeks (= teams) in cycle c
$x_{idc} \in \{0, 1, ..., W^u\}$: the number of shifts i that is scheduled during day d in cycle c
$z_{ic} \in \{0, 1\}$: = 1, if shift i is used in cycle c and 0 otherwise
$g_{fp}^\delta \geq 0$: the number of workers assigned to maintain flight f during time period p when solving the MILP model in enhancement iteration δ
$e_{tdc}^+ \in \{0, 1, ..., W^u\}$: the number of extra weeks needed in cycle c for day d caused by shifts of type t (with $t \in \{E, N\}$) on the preceding day (in order to satisfy the shift succession guaranteeing constraints)
$e_{Ndc}^- \in \{0, 1, ..., W^u\}$: the number of extra weeks needed in cycle c for day d caused by E shifts that can be compensated by an excess in N shifts on the preceding day (in order to satisfy the shift succession guaranteeing constraints)

The optimization model can be formulated as follows:

$$Minimize:$$
$$\sum_{c \in C} \sum_{i \in I} \sum_{d=1}^{7} k_{id} M_c x_{idc} \tag{21}$$
$$Subject\ to:$$
$$\sum_{i \in I} \sum_{d=1}^{7} \sum_{c \in C} b_{idp} M_c x_{idc} = q_p^\delta, \qquad \forall p \in P \tag{22}$$
$$q_p^\delta \geq \sum_{f \in F} g_{fp}^\delta, \qquad \forall p \in P \tag{23}$$
$$\sum_{p \in P_f} g_{fp}^\delta = L_f * \frac{|P|}{24 * 7}, \qquad \forall f \in F \tag{24}$$

$$n_c \geq \lceil 1/R \rceil * \sum_{i \in I} x_{i6c}, \qquad \forall c \in C \tag{25}$$

$$n_c \geq \lceil 1/R \rceil * (\sum_{i \in I} x_{i7c} + e^+_{N7c} + e^+_{E7c}), \qquad \forall c \in C \tag{26}$$

$$\sum_{i \in I} \sum_{d=1}^{7} h_i x_{idc} \geq S n_c, \qquad \forall c \in C \tag{27}$$

$$\sum_{i \in I} \sum_{d=1}^{7} h_i x_{idc} \leq U n_c, \qquad \forall c \in C \tag{28}$$

$$x_{idc} \leq W^u z_{ic}, \qquad \forall i \in I \qquad \forall d = 1, ..., 7 \qquad \forall c \in C \tag{29}$$

$$\sum_{i \in I_t} z_{ic} \leq 1, \qquad \forall t \in T \qquad \forall c \in C \tag{30}$$

$$\sum_{i \in I} \sum_{d=1}^{7} \sum_{c \in C} a_{idp} x_{idc} \geq 1, \qquad \forall p \in P \tag{31}$$

$$n_c \geq \sum_{i \in I} x_{idc} + e^+_{Ndc} + e^+_{Edc}, \qquad \forall d \in \{1, ..., 7\} \qquad \forall c \in C \tag{32}$$

$$e^+_{N(d+1)c} \geq \sum_{i \in I_N} x_{idc} - \sum_{i \in I_N} x_{i(d+1)c}, \qquad \forall d \in \{1, ..., 6\} \qquad \forall c \in C \tag{33}$$

$$e^+_{E(d+1)c} \geq \sum_{i \in I_E} x_{idc} - \sum_{i \in I_E} x_{i(d+1)c} - e^-_{N(d+1)c}, \qquad \forall d \in \{1, ..., 6\} \qquad \forall c \in C \tag{34}$$

$$e^-_{N(d+1)c} \leq \sum_{i \in I_N} x_{i(d+1)c} - \sum_{i \in I_N} x_{idc} + e^+_{N(d+1)c}, \qquad \forall d \in \{1, ..., 6\} \qquad \forall c \in C \tag{35}$$

$$e^+_{N1c} \geq \sum_{i \in I_N} x_{i7c} - \sum_{i \in I_N} x_{i1c}, \qquad \forall c \in C \tag{36}$$

$$e^+_{E1c} \geq \sum_{i \in I_E} x_{i7c} - \sum_{i \in I_E} x_{i1c} - e^-_{N1c}, \qquad \forall c \in C \tag{37}$$

$$e^-_{N1c} \leq \sum_{i \in I_N} x_{i1c} - \sum_{i \in I_N} x_{i7c} + e^+_{N1c}, \qquad \forall c \in C \tag{38}$$

$$x_{idc} \in \{0, 1, ..., W^u\}, \qquad \forall i \in I \qquad \forall d = 1, ..., 7 \qquad \forall c \in C \tag{39}$$

$$g^\delta_{fp} \geq 0, \qquad \forall f \in F \qquad \forall p \in P \tag{40}$$

$$z_{ic} \in \{0, 1\}, \qquad \forall i \in I \qquad \forall c \in C \tag{41}$$

$$n_c \in \{W^l, W^l + 1, ..., W^u\}, \qquad \forall c \in C \tag{42}$$

$$q^\delta_p \geq 0, \qquad \forall p \in P \tag{43}$$

$$e^+_{tdc} \in \{0, 1, ..., W^u\}, \qquad \forall t \in \{E, N\} \qquad \forall d = 1, ..., 7 \qquad \forall c \in C \tag{44}$$

$$e^-_{Ndc} \in \{0, 1, ..., W^u\}, \qquad \forall d = 1, ..., 7 \qquad \forall c \in C \tag{45}$$

The objective function (21) describes the total labour costs faced by the company. The available capacity q^δ_p in constraint (22) is determined by the scheduled shifts in the left hand side of the constraint. Constraints (23) and (24) ensure that the assigned work for each flight (and therefore the available capacity) is sufficient to at least cover the scheduled demand for maintenance. This is referred to as the coverage constraint. These first three constraints are general constraints in workforce scheduling problems for aircraft maintenance and correspond to the first three constraints of the model described in Section 4.2. The rest of the constraints in the MILP model described above, are other, specific workforce scheduling constraints and correspond to constraint (5) of the MILP model in Section 4.2.

Constraints (25) and (26) are the weekend constraints. They ensure that at most a certain fraction R of the weeks in a worker's schedule contain a working weekend. To comply with the constraint on the average working hours imposed by the labour union, constraints (27) and (28) are inserted. They restrict the average number of working hours per week per worker between 36 and 38 hours. Constraints (29) and (30) ensure that the start and end times for a certain shift type are the same in each cycle. Constraint (31) is inserted because Sabena Technics wants at least one team to be present at all times as a basic capacity buffer. The shift succession constraints entail that a night shift can only be followed by another night shift (or no shift) and an evening shift can only be followed by another evening shift or a night shift (or no shift). There is no limitation to shifts succeeding a morning or a day shift. Because a team can only work one shift a day, these succession constraints ensure that there is at least 12 hours between two shifts for each team. Recall that we omit the assignment of shifts to different weeks because of efficiency issues. Therefore, shift succession constraints are not incorporated directly, but through so-called shift succession guaranteeing constraints (32 to 38). These constraints make sure that the succession constraints can be satisfied during the manual assignment of shifts to weeks. For more information on this, we refer the reader to Beliën et al. (2013). Finally, constraints (39) to (45) define the range of the decision variables.

5.1.2. Simulation model (Phase I)

The goal of the simulation model is to simulate the performance of the workforce configuration provided by the MILP model in a stochastic environment and to obtain the information described in Section 4.3. To simulate a real (stochastic) environment, stochastic delays are allocated to arriving aircraft. By analyzing real life data from Sabena Technics, we discovered that the delays in arrival times could be represented by four different probability distributions, which are then applied according to their respective probabilities. Table 2 shows how we simulated the stochastic delays in aircraft arrival times during the simulation experiment in phase I.

Table 2: Delay probability distributions

Rule	Interval (in minutes)	Probability	Distribution (in minutes)
Early arrival	$[-60, 0]$	45.81%	-60.5 + 61 Beta(6.07,1.22)
Small delay	$[1, 15]$	26.85%	0.5 + 15 Beta(0.89,1.22)
Medium delay	$[16, 270]$	27.09%	16 + 254 Beta(0.53,3.49)
Large delay	$[271, +\infty[$	0.25%	271 + 835 Beta(0.41,1.02)

Based on the simulated aircraft arrival times and the available capacity $q_p^{\delta-1}$ resulting from solving the MILP model during the previous enhancement iteration, the simulation model allocates for each time period a certain number of maintenance workers to aircraft that can be maintained during that time period. This is done according to a specific allocation rule.

While the coverage constraints (23) and (24) assign capacity to aircraft in the best possible way to maintain all flights in time with minimum capacity, no effort is made to minimize the movements of the maintenance workers between different flights. In fact, random preemption is allowed by the MILP formulation. However, attempting to model a better assignment rule in the model is not particularly helpful. First, modeling a better assignment rule in the MILP model will increase the complexity. Second, its also rather useless because the exact assignment decisions (g_{fp}^{δ}) cannot be used in a real (stochastic) environment because flights will have different arrival times each week and planners have to find a different allocation of the available capacity to arriving aircraft. Because of these two reasons, it is better to implement the assignment rule directly in the simulation model.

Finding a good assignment rule is the next challenge. Different possibilities exist with each a different impact on the performance of the workforce configuration (Van den Bergh et al., 2013c). Hence, the average number of flights with late maintenance returned by the simulation experiment can be different for different assignment rules. We choose to apply the first-leave-first-serve assignment rule in our ME procedure because of two reasons. First, this rule avoids random preemption and results in a more realistic capacity assignment. Second, assuming that all flights arrive on schedule, this assignment rule will make sure that all flights can be maintained in time with the capacity determined by the MILP model (i.e., $q_p^{\delta-1}$). Therefore, the simulation evaluation of the workforce configuration with this assignment rule will also lead to a relatively small average number of flights with late maintenance compared to other rules. Algorithm 1 shows how the simulation model allocates maintenance workers to aircraft according to this first-leave-first-serve allocation rule.

Algorithm 1 Allocation rule algorithm: First-leave-first-serve

Sort all flights $f \in F$ by increasing STD;
for all $p \in P$ **do**
  $RemainingCapacity_p = q_p^{\delta-1}$;
end for
for all $p \in P$ **do**
  for all $f \in F$ **do**
    if (Flight f is arrived and still needs maintenance) **and** ($RemainingCapacity_p > 0$) **then**
      if (Capacity still needed for flight $f \geq RemainingCapacity_p$) **then**
        Capacity assigned to flight $f = RemainingCapacity_p$;
      else
        Capacity assigned to flight $f =$ Capacity still needed for flight f;
      end if
      $RemainingCapacity_p = RemainingCapacity_p$ - Capacity assigned to flight f;
    end if
  end for
end for

Table 3: Comparison of the results of different approaches

	$MILP^0$ (1')		$MILP^0$ (15')		SFO (+ 1)		SFO (+ 2)		SFO (+ 3)		ME ($\varphi \leq 5$)		ME ($\varphi \leq 3$)		ME ($\varphi \leq 1$)	
Test Set	Costs	φ	Costs	φ	Costs	φ	Costs	φ	Costs	φ	Costs	φ	Costs	φ	Costs	φ
1_1_1_1	16066.5	5.42	15823.1	6.49	15980.0	4.56	17642.8	2.18	20530.1	1.48	15861.8	4.89	17493.4	2.80	24779.6	0.79
1_1_1_2	19354.8	6.04	19128.1	5.67	21901.7	3.04	24461.9	2.02	27397.2	1.44	19315.8	4.64	22388.9	2.85	28535.6	0.99
1_1_1_3	15545.1	6.18	15407.0	6.30	16878.1	4.39	18742.9	2.57	24887.7	1.09	16253.8	4.95	17822.5	2.52	23012.9	1.00
1_1_1_4	18024.0	6.56	17790.7	5.86	18294.0	2.84	19671.7	2.03	23800.5	1.01	17797.9	4.75	18187.1	2.79	26730.7	0.93
1_1_1_5	18038.4	6.75	17126.4	5.36	19589.9	3.16	24709.5	1.47	F	F	17332.3	4.55	23563.5	2.70	27568.4	0.99
1_1_2_1	14425.4	4.63	13218.7	4.93	14072.2	3.89	16845.9	1.63	23136.4	0.89	-	-	14493.7	2.84	18483.6	1.00
1_1_2_2	14213.5	3.87	12639.8	5.05	12692.2	3.56	15236.6	2.40	16890.8	1.77	-	-	14443.9	2.73	19176.6	0.96
1_1_2_3	18578.5	5.91	18578.5	5.91	21168.0	3.67	24051.8	1.83	F	F	18818.1	4.67	20535.8	2.64	26878.7	0.94
1_1_2_4	16421.8	4.71	16421.8	4.71	17757.9	2.73	22407.7	1.65	F	F	-	-	17686.7	2.59	23676.5	0.95
1_1_2_5	15834.5	6.31	15834.5	6.31	16218.1	3.68	18380.2	1.78	20304.2	0.94	15899.6	4.83	17293.6	2.92	19882.4	0.97
1_2_1_1	13623.7	4.06	13578.4	4.14	14029.9	1.70	15162.1	1.08	16377.8	0.95	-	-	13614.4	2.48	16348.1	0.97
1_2_1_2	16382.6	3.50	16329.2	3.72	18735.6	2.23	21532.5	1.29	F	F	-	-	16912.4	2.78	21858.1	0.84
1_2_1_3	11029.2	3.49	11004.0	4.17	13337.7	2.02	13387.6	1.47	15711.6	0.91	-	-	11112.5	2.87	15275.1	0.97
1_2_1_4	11624.8	3.75	11624.8	3.75	12961.8	2.13	14060.0	1.78	17941.3	0.68	-	-	12383.3	2.82	16197.2	0.93
1_2_1_5	13855.3	3.40	13806.5	3.43	15295.2	2.09	18401.4	1.17	F	F	-	-	14325.3	2.95	21871.0	0.76
1_2_2_1	10915.1	2.83	10057.4	3.10	10838.2	2.21	11962.4	1.95	15813.8	0.77	-	-	-	-	14258.8	0.90
1_2_2_2	12117.4	1.53	11831.6	2.78	12753.6	1.26	13849.7	1.06	16511.4	0.40	-	-	-	-	13820.4	0.77
1_2_2_3	13255.0	2.93	13179.1	2.92	15415.1	2.29	18798.6	1.27	F	F	-	-	-	-	19273.3	0.85
1_2_2_4	10290.9	4.96	10203.9	5.79	11175.9	2.27	12273.2	1.51	13604.2	1.01	-	-	11158.4	2.70	13248.5	0.92
1_2_2_5	10719.3	2.88	10719.3	2.88	11113.0	1.96	12269.7	1.13	13451.2	0.47	-	-	-	-	13307.7	0.89
Average	14515.8	4.49	14215.1	4.66	15510.4	2.78	17692.4	1.66	19025.6	0.99	14507.7*	3.99*	15521.1*	2.71*	20209.2	0.92

$MILP^0$ (X') : $MILP^0$ model solved for X minute(s)
SFO (+ Y) : Straightforward optimization model solved for 15 minutes based on a delay of Y quarter(s) for each flight
ME ($\varphi \leq Z$) : Model enhancement algorithm ran for 15 minutes with the stochastic service level constraint: $\varphi \leq Z$
φ : The average number of delayed flights according to the simulation
* : The average includes the corresponding (costs or φ) result of the $MILP^0$ (1') model for each instance indicated by " - ". Hence, when there are no ME results for a certain test set, the results of the $MILP^0$ (1') model are used to calculate the average.

5.2. *Test settings*

5.2.1. *Test sets*

The ME algorithm has been applied for several settings to a test set containing 20 instances created randomly based on real-life data dimensions from Sabena Technics. The entire set can be divided into 4 groups of 5 instances. Each group has its own specific characteristics. The results are shown in Table 3. All 20 instances (called 1_*_*_*) contain 100 flights but are created using different workload distributions. The workload is the amount of work that each flight requires. Half of the groups has flights with workloads drawn from a uniform distribution between 0 and 10 hours (1_1_*_*), while the other half has flights with workloads drawn from an exponential distribution with an average of 3.5 hours (1_2_*_*). The last difference between the 4 groups is the way how the flight arrivals are chosen. Half of the groups has flights with peak arrivals (1_*_1_*), while the other half has flights with uniform arrivals (1_*_2_*). In the case of peak arrivals, more flights require maintenance either in the morning or the afternoon, while the demand for maintenance is uniformly distributed over the entire day in the case of uniform arrivals. Finally, 5 random instances are generated (1_*_*_1, 1_*_*_2, 1_*_*_3, 1_*_*_4 and 1_*_*_5) for each of the above cases resulting in 2x2x5 = 20 instances.

Based on the flight information in each test set, a workforce configuration needs to be constructed resulting in a weekly shifts schedule. The resolution of our scheduling decision is fifteen minutes. Hence, $|P|$ is set to (7 x 24 x 4 =) 672. This means that decisions have to be made every fifteen minutes and the planning horizon (equal to one week) is divided in 672 quarters of an hour.

We now discuss Table 3 which shows the results of the computational experiments.

5.2.2. *$MILP^0$ results*

Columns 2 to 5 of Table 3 present the results of the $MILP^0$ model. This is the deterministic MILP model since it only looks at the Scheduled Time of Arrival (STA) of each flight to construct the workforce configuration. No delays are taken into account or anticipated. Solving this $MILP^0$ model is also the first step in the enhancement algorithm. Columns 2 and 4 show the resulting total labour costs from the $MILP^0$ model that was solved with a time limit of one minute (1') and fifteen minutes (15') respectively. All MILP models were solved using the IBM CPLEX optimizer. The third and fifth column show the average number of delayed flights ($= \varphi$) that results from simulating the respective workforce configuration found by the $MILP^0$ model. These two columns show how well the results from the deterministic model perform in a stochastic environment without any enhancements to the MILP model.

5.2.3. *SFO results*

Columns 6 to 11 present the results of the Straightforward Optimization (SFO) technique. SFO can be seen as a mathematical optimization technique that uses approximations or even relaxations instead of stochastic variables. SFO uses the average delay time of a flight to approximate stochastic delays. SFO is a very simple technique to anticipate stochastic delays and is therefore used as an alternative method to ME in our performance analysis. In fact, the same $MILP^0$ (15') model is used to obtain the results in columns 6, 8 and 10. However, prior to solving the MILP model, a delay of respective one, two and three quarters is added to the STA of each flight to anticipate respectively average, large and very large stochastic delays. From Table 2 we calculated that the average delay time is 12 minutes. But since we are working with time periods of 15 minutes, the average delay time is rounded to 15 minutes, i.e., one quarter. Also note that we solve the SFO model for 15 minutes because the ME algorithm will also be given 15 minutes of computation time. Columns 7, 9 and 11 of Table 3 show the average number of delayed flights ($= \varphi$) that results from simulating the respective workforce configuration found by the SFO ($+Y$) model. " F "indicates that the respective model failed to find a single feasible solution in the allowed time span. Hence, three quarters appears to be the maximum delay to which the SFO MILP model is able to find a feasible solution in most cases.

5.2.4. ME results

The last six columns of Table 3 show the results of the ME algorithm for three different service level constraints. In the first case (see columns 12 and 13), the service level constraint allows for at most five flights on average to leave with a delay. The next two cases allow for at most three flights and one flight respectively. For each of the test sets, fifteen ME iterations were performed ($\Delta = 15$). During each optimization phase, the optimization of the MILP model in the ME procedure is limited to one minute, which, according to preliminary test, results in acceptable optimality gaps. During the simulation phase, the results of each MILP model are used to set the capacity at each period $p \in P$ in the simulation model. In our case, we perform a simulation with a length of 500 weeks, which takes just about a second to run. Simulation tests indicated that a simulation length of 500 weeks already results in robust estimates for φ. Recall from Section 4.3 that φ is the average of $|\Phi_\omega|$ over all simulated weeks ω. Statistical analysis of the results indicated that $|\Phi_\omega|$ follows a Poisson distribution with average φ. Hence, the variance of $|\Phi_\omega|$ equals φ. Based on the simulation results, the MILP model is enhanced during the enhancement phase. When the service level constraint is satisfied at the beginning of phase II, the diversification procedure is applied. During the diversification, we set Ξ to 0.4 (= 40%) and Ψ to 0.25 (= 25%) to diversify the search. The whole ME procedure is programmed in C++ using the IBM ILOG CPLEX API for the optimization phase.

Because each MILP optimization is limited to one minute and the simulation time is negligible, it takes about 15 minutes to apply the ME procedure (of fifteen iterations) on each test set. In Table 3, " - " is used to indicate that the respective test set is not solved with ME because the simulation evaluation of the $MILP^0$ (1') solution already satisfies the service level constraint that is specified on top of the column. Note that the simulation evaluation of the $MILP^0$ (1') result is used for this purpose instead of the $MILP^0$ (15') result. This is because solving MILP models in the ME algorithm is also limited to only one minute. Hence, when the simulation evaluation of the $MILP^0$ (1') model already satisfies the stochastic constraint, there is no need to enhance the model, and the ME algorithm will not be executed.

5.3. Performance analysis

5.3.1. Ability to satisfy the service level constraint

The first performance measure that we consider is the ability of the ME algorithm to satisfy a predefined service level constraint. To test this, three scenarios with a different service level constraint are investigated. As can be seen in columns 13, 15 and 17 of Table 3, ME ensures the satisfaction of the service level constraint for each test instance for each scenario, even if the allowed average number of flights with late maintenance is reduced from 5 to 3 and even to 1. While ME ensures the satisfaction of the service level constraint for each test instance, the simulation evaluation of the SFO solutions cannot be accurately predicted in advance. Theoretically, φ should be zero for "SFO (+ 1)" because an average delay of one quarter is expected for each flight (see Section 5.2.3) which is also accounted for in the SFO (+ 1) MILP model. However, the results in column 7 of Table 3, prove otherwise. As opposed to ME, SFO is clearly unable to provide extra capacity at the right moments.

5.3.2. Ability to achieve lower costs compared to SFO

Until now, we showed that our ME approach can provide a robust solution to our problem (i.e., a solution that satisfies the service level constraint). In this section we investigate the cost of robustness as we expect the labour costs to rise when the service level is increased. A higher service level requires a higher capacity buffer which entails higher costs. Since our focus remains on minimizing the labour costs, this is our second performance measure.

Since solving each MILP model in the ME approach is limited to one minute, the result of the $MILP^0$ (1') model is the result of the initial ME iteration (i.e., for $\delta = 0$). Two situations can occur. Either, the $MILP^0$ (1') model does not satisfy the service level constraint and the enhancement procedure starts, or, the $MILP^0$ (1') model does satisfy the service level constraint. When the latter is the case, and $MILP^0$ is solved to optimality after one minute, stopping the procedure as proposed in Section 4.4.3 would be optimal. However, the MILP represents an extremely complex problem that cannot be solved to optimality in a reasonable time limit. Therefore, $MILP^0$ (1') and even $MILP^0$ (15') will never result in an optimal solution to the problem. Improving a stochastic feasible solution is left out of the focus of this paper and is a possible direction for future research. Hence, the cases indicated by " - " in Table 3 are left out of our performance analysis and we only consider the cases where $MILP^0$ (1') does not satisfy the service level constraint.

The ME approach sometimes results in lower costs compared to the $MILP^0$ (1'), even when the average number of delayed flights is reduced to satisfy the service level. This can be observed by comparing the costs of ME ($\varphi \leq 5$) with $MILP^0$ (1') in columns 2 and

12 and the costs of ME ($\varphi \leq 3$) with $MILP^0$ (1') in columns 2 and 14. The first reason for this counter-intuitive observation is that adding the extra constraints during the ME approach facilitates the optimization procedure resulting in lower costs within the same time limit. Second, the ME algorithm features a diversification strategy, which is in fact a local search method. However, most of the instances follow our intuition and show that the labour costs rise with increasing service level.

We now compare the ME results to the SFO results with respect to the increase in labour costs. SFO can be used as a benchmark because it is the most simple approach, and it should therefore be outperformed by the ME approach for the addition of the extra complexity not to be in vain. We solved three SFO models based on delays ranging from one quarter to three quarters. Since different delays result in different values for φ, the SFO results can be used as a benchmark for different ME settings. Hence, when the result of an SFO model satisfies the service level constraint (i.e.,"$\varphi \leq 5$", "$\varphi \leq 3$" or "$\varphi \leq 1$" resp.), the respective ME algorithm should result in lower labour costs. Table 3 indeed reveals that for all instances where the simulation evaluation of the SFO approach satisfies the "$\varphi \leq 5$", the "$\varphi \leq 3$" or the "$\varphi \leq 1$" constraint, the respective ME approach always outperforms the respective SFO approach in terms of labour costs.

5.3.3. Diversification to escape local optima

While solving the problem, our focus remains on minimizing the total labour costs. Therefore, the ME algorithm will not terminate once a solution is found that satisfies the service level constraint. Instead, a diversification procedure is used to escape from local optima as elaborated in Section 4.4.3. Table 4 analyzes the performance of the diversification technique. For each of the three ME scenarios ("ME ($\varphi \leq 5$)", "ME ($\varphi \leq 3$)" and "ME ($\varphi \leq 1$)" resp.), the minimum, maximum and average percentage improvement resulting from diversification is reported. This is the percentage difference between the costs of the best ME solution (as reported in Table 3) and the costs of the first feasible solution found during the ME procedure that satisfies the service level constraint.

Table 4: Percentage improvements obtained by diversification

	% improvement by diversification over all test sets for three scenarios:		
	ME ($\varphi \leq 5$)	ME ($\varphi \leq 3$)	ME ($\varphi \leq 1$)
Minimum	0.00%	0.00%	0.00%
Maximum	5.22%	23.47%	6.91%
Average	1.57%	8.22%	0.99%

In the first scenario (ME ($\varphi \leq 5$)), we allow on average for at most five flights with late maintenance instead of three or one. Therefore, it is relatively easy to satisfy the service level constraint in this first scenario. Minor changes are required to the available capacity resulting from the $MILP^0$ (1') model which means that the first solution that satisfies the service level constraint will already be very good in terms of labour costs. This renders the diversification strategy rather useless in this scenario, which results in a relatively small average percentage improvement of 1.57% (see Table 4).

In the second scenario (ME ($\varphi \leq 3$)), it becomes more difficult to satisfy the service level constraint without a drastic increase in labour costs compared to the $MILP^0$ (1') model. Nevertheless, the ME algorithm still finds the first solution that satisfies the service level constraint quite fast in this scenario. However, the resulting labour costs increase significantly. Hence, the diversification strategy has more potential in this case and we expect the average percentage improvement to increase. Table 4 shows that the average percentage improvement has indeed increased from 1.57% to 8.22%. In the best case, the diversification strategy even lowers the costs of the first feasible solution by 23.47%.

Figure 4: Demonstration of the power of ME to improve a feasible solution

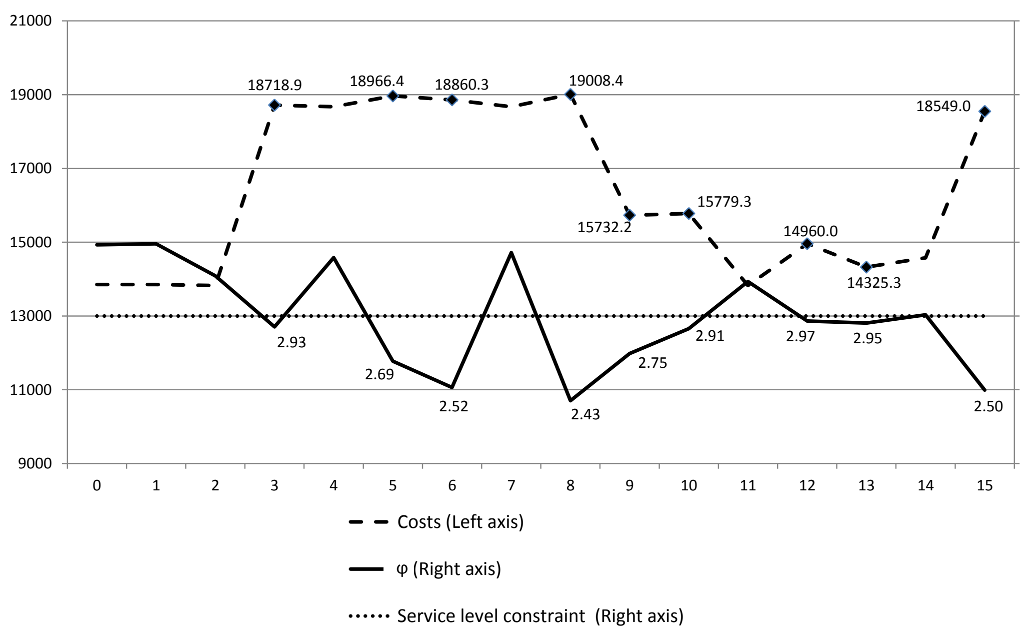

Note: This example shows the results of the ME procedure applied to test instance 1_2_1_5 under the stochastic constraint: $\varphi \leq 3$. After the optimization of the $MILP^0$ model, 15 ME iterations (shown on the horizontal axis) are performed. Stochastic feasible solutions are indicated with diamonds on the graph.

Satisfying the service level constraint becomes even more difficult in the third scenario (ME ($\varphi \leq 1$)). As expected, the labour costs of the first solution that satisfies the service level constraint increase drastically in this case. Hence, the diversification strategy has even more potential in this case and we again expect the average percentage improvement to increase. However, the results in Table 4 prove otherwise. Because of the very strict service level constraint ($\varphi \leq 1$), the ME procedure spends much more time on finding a first feasible solution. Therefore, less time remains for the diversification strategy to improve the initial feasible solution. This results in a very low average percentage improvement of only 0.99%. But in the best case, the diversification strategy lowers the costs of the first feasible solution by 6.91%.

Figure 4 demonstrates the ME procedure applied on one specific problem instance. During the first three ME iterations in Figure 4, Equations (15) to (17) are used to find the first feasible result (18718.9). Then, Equations (18) to (20) are applied to find a better result that still satisfies the stochastic service level constraint. Figure 4 shows that the ME approach is capable of finding three times a better feasible result with respect to the labour costs (15732.2, 14960.0 and 14325.3 resp.) during the next iterations. Note that for each of the feasible solutions the service level constraint ($\varphi \leq 3$) is indeed satisfied because $\varphi \leq 3$.

6. Conclusion and future research

In this paper, we present a heuristic approach for building robust aircraft maintenance personnel rosters. A model enhancement (ME) heuristic is constructed to optimize a mixed integer linear programing (MILP) model with a stochastic service level constraint. This constraint implies that only a certain average percentage of all weekly flights cannot be maintained in time when aircraft arrive with stochastic delays. The three phases of the ME algorithm are presented: a simulation model, an enhancement procedure and the mixed integer linear programming model.

We successfully applied our model to a real life problem setting at Sabena Technics, a large aircraft maintenance company located at Brussels Airport in Belgium. We illustrate the performance of the ME algorithm with a computational experiment and compare the results to deterministic optimization and straightforward optimization (SFO). We report the results in terms of labour costs and the average number of flights with late maintenance according to the simulation experiment. We tested the ME algorithm on 20 instances based on real data and allowed the ME algorithm to run for 15 iterations on each instance.

Experiments first demonstrate that the ME approach always succeeds in finding a feasible solution that satisfies a predefined stochastic service level constraint. Second, the cost of robustness appears to be lower for the ME approach than for the SFO approach. The cost of robustness is the increase in labour costs required to construct a capacity buffer in order to satisfy the service level constraint. Third, experiments also prove the power of the diversification strategy in the ME algorithm to guide the search to find better solutions once the stochastic service level constraint has been satisfied.

Finally, we propose some interesting extensions to this paper for future research. First, the stochastic model could be extended by uncertainty in capacity (i.e., absenteeism) and uncertainty in the workload. Second, different allocation rules can be investigated to improve the model. Certain flights can be made more important than other flights for example. And finally, the diversification strategy can be improved. For example, a tabu mechanism (such as in Tabu Search) or some other metaheuristic mechanism could be implemented to guide the search to find better feasible solutions.

Bachelet, B., Yon, L. 2007. Model enhancement: Improving theoretical optimization with simulation. *Simulation Modelling Practice and Theory* **15** (6) 703–715.

Beliën, J., Demeulemeester, E., De Bruecker, P., Cardoen, B., Van den Bergh, J. (2013). Integrated staffing and scheduling for an aircraft line maintenance problem. *Computers & Operations Research*

Beliën, J., Demeulemeester, E., Cardoen, B. (2012). Improving workforce scheduling of aircraft line maintenance at Sabena Technics. *Interfaces* **42** (4) 352–364.

Benders, J.F. (1962). Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik* **4** (1) 238–252.

Bertsimas, D. and Sim, M. (2004). The Price of Robustness. *Operations Research* **52** (1) 35–53.

Chen, C.-H., Lee, L. H. (2011). *Stochastic Simulation optimization.* Singapore: World Scientific.

Duffuaa, S. O., Andijani, A. A. (1999). An integrated simulation model for effective planning of maintenance operations for Saudi Arabian Airlines (SAUDIA). *Production Planning and Control* **10** 579–584.

Fagerholt, K., Christiansen, M., Hvattum, L. M. , Johnsen, T. A. V., Vabø, T. J. (2010). A decision support methodology for strategic planning in maritime transportation. *Omega* **38** (6) 465–474.

Fu, M.C. (2002). Optimization for simulation: theory vs. practice. *INFORMS Journal on Computing* **14** (3) 192–215.

Liang, Z., Chaovalitwongse, W. A., Huang, H. C., Johnson, E. L. (2010). On a New Rotation Tour Network Model for Aircraft Maintenance Routing Problem. *Transportation Science* **45** 109–120.

Mattila, V., Virtanen, K., Raivio, T. (2008). Maintenance Decision Making in the Finnish Air Force Through Simulation. *Interfaces* **38** 187–201.

Muchiri, A. K. (2009). Application of Maintenance Interval De-Escalation in Base Maintenance Planning Optimization. *Entreprise Risk Management* **1** 63–75.

Papakostas, N., Papachatzakis, P., Xanthakis, V., Mourtzis, D., Chryssolouris, G. (2010). An approach to operational aircraft maintenance planning. *Decision Support Systems* **48** 604–612.

Petersen, J. D., Solveling, G., Clarke, J. P., Johnson, E. L., Shebalov, S. (2012). An Optimization Approach to Airline Integrated Recovery. *Transportation Science* **46** 482–500.

Quan, G., Greenwood, G. W., Liu, D., Hu, S. (2007). Searching for multiobjective preventive maintenance schedules: Combining preferences with evolutionary algorithms. *European Journal of Operational Research* **177** 1969–1984.

Safaei, N., Banjevic, D., Jardine, A. K. S. (2010). Bi-objective workforce-constrained maintenance scheduling: a case study. *Journal of the Operational Research Society* **62** 1005–1018.

Safaei, N., Banjevic, D., Jardine, A. K. S. (2011). Workforce-constrained maintenance scheduling for military aircraft fleet: a case study. *Annals of Operations Research* **186** 295–316.

Sarac, A., Batta, R., Rump, C. M. (2006). A branch-and-price approach for operational aircraft maintenance routing. *European Journal of Operational Research* **175** 1850–1869.

Sriram, C., Haghani, A. (2006). An optimization model for aircraft maintenance scheduling and re-assignment. *Transportation Research Part A* **37** 29–48.

Talluri, K. T. (1998). The four-day aircraft maintenance routing problem. *Transportation Science* **32** 43–53.

Van den Bergh, J., Beliën, J., De Bruecker, P., Demeulemeester, E., De Boeck, L. (2013). Personnel scheduling: A literature review. *European Journal of Operational Research* **226** (3) 367–385.

Van den Bergh, J., De Bruecker, P., Beliën, J., Peeters, J. (2013). Aircraft maintenance operations: state of the art. *FEB@Brussel research paper.*

Van den Bergh, J., De Bruecker, P., Beliën, J., De Boeck, L., Demeulemeester, E. (2013). A three-stage approach for aircraft line maintenance personnel rostering using MIP, discrete event simulation and DEA. *Expert Systems with Applications* **40** (7) 2659–2668.

Yan, S., Yang, T.-H., Chen, H.-H. (2004). Airline short-term maintenance manpower supply planning. *Transportation Research Part A: Policy and Practice* **38** 615–642.

FACULTY OF ECONOMICS AND BUSINESS
Naamsestraat 69 bus 3500
3000 LEUVEN, BELGIË
tel. + 32 16 32 66 12
fax + 32 16 32 67 91
info@econ.kuleuven.be
www.econ.kuleuven.be