

Recoverable robust single day aircraft maintenance routing problem

Stephen J Maher¹, Guy Desaulniers², and François Soumis²

¹School of Mathematics and Statistics, University of New South Wales, Sydney NSW 2052, Australia.

²École Polytechnique de Montréal and GERAD, Department of Mathematics and Industrial Engineering, Montréal Québec H3C 3A7, Canada.

Abstract

Aircraft maintenance planning is of critical importance to the safe and efficient operations of an airline. It is common to solve the aircraft routing and maintenance planning problems many months in advance, with the solution spanning multiple days. An unfortunate consequence of this approach is the possible infeasibility of the maintenance plan due to frequent perturbations occurring in operations. There is an emerging concept that focuses on the generation of aircraft routes for a single day to ensure maintenance coverage that night, alleviating the effects of schedule perturbations from preceding days. In this paper, we present a novel approach to ensure that a sufficient number of aircraft routes are provided each day so maintenance critical aircraft receive maintenance that night. By penalising the under supply of routes terminating at maintenance stations from each overnight airport, we construct a single day routing to provide the best possible maintenance plan. This single day aircraft maintenance routing problem (SDAMRP) is further protected from disruptions by employing the recoverable robustness technique. To efficiently solve the recoverable robust SDAMRP acceleration techniques, such as identifying Pareto-optimal cuts and a trust region approach, have been applied. The SDAMRP is evaluated against a set of flight schedules and the results demonstrate a significantly improved aircraft maintenance plan. Further, the results demonstrate the magnitude of recoverability improvement that is achieved by employing recoverable robustness to the SDAMRP.

Key words: Airline planning, maintenance, robustness, recoverability.

1 Introduction

The aircraft routing problem is typically solved following the fleet assignment, many months in advance of the day of operations. At this stage, maintenance planning is performed to ensure that each aircraft in the fleet is checked at regular intervals. The resulting maintenance plans can span multiple days, making them susceptible to schedule perturbations. While robust planning is becoming more popular for airline

problems, perturbations occurring in operations can still significantly alter the planned solution. As such, any schedule disruptions may restrict maintenance critical aircraft from entering a maintenance base, unless costly intervention is provided by the operations control centre.

There have been numerous approaches presented for solving the aircraft maintenance routing problem, with and without considering the effects of schedule perturbations. Lacasse-Guay *et al.* [17] provides a good review of current business practices that shape the formulation of the aircraft routing problem (ARP). In [17], the authors explain that the various approaches fall into three broad categories, *Strings*, *Big-cycle* and *One-day routes*. The *string* model involves the construction of flight routes spanning between visits to maintenance stations. The flight routes, or strings, are constructed to be maintenance feasible and can be performed by all aircraft of the one fleet. The *big cycle* model has the objective of identifying one single route that covers multiple days and includes every scheduled flight. This modelling approach is most common for schedules that are identical every day. In the resulting solution, each overnight stop in the cycle represents a possible starting point for an aircraft. Also, the cycle is constructed to be maintenance feasible by scheduling maintenance visits at appropriate intervals throughout.

Finally, the *one-day routes* approach is designed under the assumption that disruptions from preceding days render the maintenance planning for an airline infeasible most of the time. Given this assumption, aircraft routing can be seen as a two-stage decision process. In the first stage (the planning stage), one-day routes are planned to ensure aircraft maintenance feasibility by providing a sufficient number of routes out of each overnight airport that terminate at maintenance stations. In the second stage, which is performed every night before the operations, specific aircraft are assigned to the one-day routes. Specifically, this process ensures that each maintenance critical aircraft at the start of the day terminate at a maintenance station. The one-day routes model used in the first stage differs from the string and big-cycle models discussed above by constructing flight routes that only span a single day. Routes originating from each overnight airport are constructed for each day of the planning horizon. The one-day routes planning model will be developed further throughout this paper.

Traditionally, the aircraft maintenance routing problem is presented using either a big-cycle or string formulation. Examples of maintenance routing problems solved to identify a big-cycle are presented by Feo and Bard [10], Clarke *et al.* [6] and Gopalan and Talluri [14]. Alternatively, examples of the string formulation are presented by Barnhart *et al.* [3] and Sriram and Haghani [31]. The aircraft routing problems developed using these two approaches schedule aircraft maintenance over multiple days under the expectation that few disruptions will occur. Given the prevalence of disruptions to airline operations, there is a high probability that planned maintenance schedules will be affected each day. While the maintenance schedule can be corrected by the operations control centre with the use of aircraft swaps during the day, this significantly increases the operational costs of the airline. Ideally any intervention by the operations control centre is performed at night to correct any changes to maintenance schedules.

The one-day routes approach to the ARP is an attempt to address the down-the-line impacts that disruptions have on maintenance scheduling. In practice, an optimisation problem using the one-day routes approach is solved each night, focusing on the aircraft requiring maintenance the following night. This approach attempts to generate a sufficient number of routes that terminate at maintenance stations, which will be allocated to maintenance critical aircraft. Heinhold [16] presents an approach employed by Southwest Airlines that calculates the expected number of aircraft at each overnight airport requiring maintenance the following day. This expectation is used to form a model to ensure that an adequate number of aircraft routes departing from each overnight airport terminate at maintenance stations. By solving this problem for close-in planning, only small modifications are required in response to disruptions from preceding days to recover the planned maintenance schedule.

The concept of one-day routes is investigated further by Lapp and Cohn [19] proposing the use of a multi-stage optimisation problem to improve maintenance reachability. Lapp and Cohn [19] develop a model using the concept of a line-of-flight (LOF), which is a sequence of flights performed by an aircraft during a single day. The optimisation problem modifies LOFs to provide maintenance visits for maintenance critical aircraft stationed at each overnight airport. This is similar to the approach presented by Heinhold [16], however the multi-stage approach improves the tractability of the problem. The input for this problem is the set of originally planned LOFs departing from each overnight airport. A subset of LOFs terminate at maintenance stations and are described as a maintenance line-of-flight (MLOF). The proposed optimisation model identifies overnight airports that have an inadequate number of MLOFs. By selecting a LOF and MLOF from airports with an under and over supply of MLOFs respectively, a single swap is performed to allow the LOF to terminate at a maintenance station. This process is performed to increase the maintenance reachability from all airports in the network. The authors demonstrate an increase in the maintenance reachability compared to the originally planned LOFs, reporting reasonable runtimes. Since only one swapping opportunity is identified along each LOF, there is a possibility to further improve the maintenance routing by allowing more rerouting options for each aircraft. The single day aircraft maintenance routing problem (SDAMRP) developed in this paper employs the one-day routes approach to provide an adequate number of MLOFs in the aircraft routing stage. A major contribution of this paper is the novel approach developed for considering aircraft maintenance for a single day schedule.

The development of the one-day routes approach for the ARP is motivated by the impact of schedule disruptions on medium and long term maintenance plans. While the one-day routes approach aims to provide a feasible maintenance schedule at the start of each day, disruptions throughout the day can still affect this maintenance plan. Robust planning is a common approach employed to reduce the susceptibility of an aircraft routing solution to schedule disruptions. There are two main types of robust planning, *proxy* and *feedback* robustness, each displaying individual strengths and weaknesses.

Proxy robust approaches attempt to improve the operational performance of an aircraft routing

solution by enhancing favourable planning characteristics. Examples of proxy robust approaches are presented by Lan *et al.* [18] and Borndorfer *et al.* [5], which attempt to minimise the expected propagated delay. An alternative form of proxy robustness identifies features from the planning solution that are expected to aid the recovery process. This form of robust planning focuses on the concept of recoverability, which is defined as the effort required by operational controllers during recovery as a result of planning stage decisions. Proxy robust approaches that explicitly attempt to improve the recoverability of aircraft routing solutions are presented by Ageeva [1], Eggenberg [9] and Rosenberger *et al.* [28].

By comparison, feedback robust approaches explicitly evaluate of the robustness, or recoverability, of the planning stage solution during the optimisation process. The evaluation stage of a feedback robust approach helps to develop a less conservative robust planning solution compared to equivalent proxy robust models. Recoverable robustness, introduced by Liebchen *et al.* [20], is an example of a feedback robust approach that attempts to improve the planning stage recoverability. In particular, the objective of recoverable robustness is to identify a planned solution that is recoverable with *limited effort*. This concept is applied to the tail assignment problem in Froyland *et al.* [12], which is solved to identify a single day tail assignment with a minimum expected recovery cost. This paper presents an extension upon the recoverable robust tail assignment problem presented in [12] by applying this approach to real-world, large-scale airline schedules. Additionally, we present an aircraft routing problem that is solved with direct consideration of maintenance requirements for a single fleet type. This same level of maintenance planning is not provided in the tail assignment problem presented by Froyland *et al.* [12]. The aircraft routing problem with maintenance considerations is a novel approach modelled using one-day routes and the recoverable robustness technique further strengthens this planning method. A further contribution of this paper is the analysis of connection cost functions regarding their impact on the recoverability improvements and solution runtimes on the recoverable robust problem.

This paper is presented in two stages, the formulation of the single day aircraft maintenance routing problem and the application of recoverable robustness. The purpose of this structure is to highlight the individual strengths of each of these techniques. Section 2 develops the model for the single day aircraft maintenance routing problem, providing a full description of the techniques used in the problem formulation. A description of recoverable robustness for the single day aircraft maintenance routing problem is presented in Section 3. Using the recoverable robustness technique, the planning aircraft routing solution is evaluated by solving a recovery subproblem. The description of the recovery problem used in the evaluation stage is presented in Section 3. To solve the problems presented in this paper, we apply Benders' decomposition and column generation. These solution methods are discussed in Section 4. The results of our computational experiments are reported in Section 5, which is followed by our conclusions in Section 6.

2 Single day aircraft maintenance routing problem

The single day aircraft maintenance routing problem (SDAMRP) is solved at the planning stage to identify a set of aircraft routes for a single fleet that ensure maintenance feasibility at the end of the day. There are two key aspects of the SDAMRP i) finding a feasible aircraft routing solution, and ii) ensuring maintenance feasibility for all aircraft. First, a feasible aircraft routing solution requires exactly one aircraft to be assigned to each flight in the schedule. It is also required that only the aircraft available to the airline are assigned to the flight schedule. Second, aircraft maintenance feasibility is given by ensuring each aircraft requiring maintenance at the end of the day terminates at a maintenance station. This requirement is achieved by providing a sufficient number of aircraft routes from each origination airport that terminate at maintenance stations.

There are a number of maintenance checks that must be performed on aircraft to satisfy aviation regulatory requirements. The maintenance check performed most frequently is a type *A* check, which is modelled in the SDAMRP to be required once every six days. This assumption ensures that a sufficient amount of time is given to achieve high utilisation of aircraft while not exceeding any maintenance requirements. While performing a type *A* check once every six days may be appropriate for most airlines, it is trivial to alter this to match individual business practices.

Maintenance routing in the SDAMRP is based upon the concept of maintenance misalignments. A *maintenance misalignment* occurs when an aircraft requiring maintenance at the end of the day is assigned a one-day route that does not terminate at a maintenance station. If there exists any maintenance misalignments in an aircraft routing solution, costly aircraft swaps must be performed during the day to satisfy maintenance requirements. As a contribution of the SDAMRP, the number of maintenance misalignments is penalised in the objective in an attempt to provide a sufficient number of routes departing from each overnight airport that terminate at maintenance stations.

2.1 Aircraft flight routes

The SDAMRP is modelled using the one-day routes formulation that is presented in Lacasse-Guay *et al.* [17]. A one-day route p is defined as a sequence of flights that is performed by an aircraft during a single day, originating and terminating at overnight airports. The set of flights that are used to construct the one-day routes is given by N , which is described as the original schedule. Each pair of sequential flights in the one-day route is described as a feasible connection. A feasible connection (i, j) has the properties i) the destination airport of flight i is identical to the origin airport of flight j , and ii) the time between the arrival of i and the departure of j is less than a minimum turn time. All feasible connections between the flights in N are contained in C .

The set of one-day routes feasible for the SDAMRP are partitioned by the origination airport into disjoint sets. All routes p that may be assigned to aircraft originating from overnight airport b are

contained in the set P_b . The decision variables y_p equal 1 if one-day route p is operated by an aircraft and 0 otherwise. Each one-day route is described by the parameters a_{jp} that equal 1 if flight j is contained in route p and 0 otherwise. The cost of one-day route p , c_p , is dependent on the length of the connections contained in that route, which can be weighted as a form of proxy robust approach. An example of a proxy robust approach based upon connection lengths is presented by Grönkvist [15].

Aircraft are finite resources for an airline that are strategically positioned across airports throughout the network. The number of aircraft that are positioned at overnight airport b at the beginning of each day is given by R_b . It is possible to deduce the value of R_b by analysing the flow balance for each airport in the network. This parameter provides an upper bound on the number of one-day routes that can originate from b in a feasible solution. However, in practice the number routes originating from b is equal to R_b . Since this model is designed for a cyclic daily schedule, the flow balance of the schedule ensures that R_b aircraft terminate at overnight airport b to begin the following days flying.

2.2 Maintenance misalignment

The SDAMRP attempts to satisfy the maintenance requirements for each aircraft by penalising the expected number of maintenance misalignments at each overnight base. This is achieved by introducing the variables M_b and parameters o_p . The variables M_b are forced by the model constraints to equal the number of one-day routes departing from overnight airport b that terminate at a maintenance station. The parameters o_p equal 1 if route p terminates at a maintenance station and 0 otherwise.

The maintenance requirements of each aircraft at the start of the day can be described by one of two states, either requiring maintenance that evening or not. Therefore, the expected number of aircraft that require maintenance departing from each overnight airport is conveniently modelled by a binomial distribution. Using the assumption that maintenance is required once every six days, the probability of an aircraft requiring maintenance is $1/6$. To minimise the number of maintenance misalignments the penalty function $F_b(M_b)$ is introduced. This function penalises the difference between the expected required maintenance routes, given by the binomial distribution, and the actual number of maintenance routes, given by the variables M_b . While the penalty function is conveniently described by $F_b(M_b)$, it is necessary to define this function in terms of integer variables in order to formulate the SDAMRP as a mixed-integer program. The specific details regarding this penalty function is provided in Section 2.4.

2.3 Mathematical model

The parameters and variables described above are used to formulate an aircraft maintenance routing problem that is solved for a single day flight schedule. The complete formulation of the SDAMRP is presented below,

$$\min \sum_{b \in B} \sum_{p \in P_b} c_p y_p + \sum_{b \in B} F_b(M_b), \quad (1)$$

$$\text{s.t.} \quad \sum_{b \in B} \sum_{p \in P_b} a_{jp} y_p = 1 \quad \forall j \in N, \quad (2)$$

$$\text{(SDAMRP)} \quad \sum_{p \in P_b} y_p = R_b \quad \forall b \in B, \quad (3)$$

$$M_b = \sum_{p \in P_b} o_p y_p \quad \forall b \in B, \quad (4)$$

$$y_p \in \{0, 1\} \quad \forall b \in B, \forall p \in P_b, \quad M_b \geq 0 \quad \forall b \in B. \quad (5)$$

The objective of the SDAMRP minimises the cost associated with aircraft routing and any penalties resulting from maintenance misalignments. Constraints (2) ensure that every flight j contained in the original schedule N is included on a single one-day route. An upper bound on the number of aircraft departing from each overnight airport is given by constraints (3). Maintenance planning is provided by the inclusion of constraints (4), which count the number of routes departing from overnight airport b that terminate at a maintenance station. By counting the number of aircraft routes terminating at a maintenance station, the maintenance misalignment at each overnight airport b can be penalised by the function $F_b(M_b)$.

2.4 Penalty function for maintenance misalignments

As explained in Section 2.2, the function $F_b(M_b)$ penalises the difference between the expected required and actual number of maintenance routes departing from airport b . The actual number of maintenance routes is given by M_b and the expected number is obtained from a binomial distribution, $B(R_b, \phi)$ where $\phi = 1/6$. Given this distribution, the variables p_{bi} are defined as the probability that exactly i aircraft at overnight airport b require maintenance at the end of the day. The following expression is used to calculate the value of p_{bi} ,

$$p_{bi} = \binom{R_b}{i} \phi^i (1 - \phi)^{R_b - i}. \quad (6)$$

Using the values for these variables and the actual number of maintenance routes given by M_b , it is possible to quantify the expected number of maintenance misalignments at each overnight airport b . This expectation, $E_b(M_b)$, is given by,

$$E_b(M_b) = \begin{cases} \sum_{i=M_b+1}^{R_b} (i - M_b) p_{bi} & \text{if } 0 \leq M_b < R_b, \\ 0 & \text{if } M_b = R_b. \end{cases} \quad (7)$$

The penalty function included in the objective of the SDAMRP is described by $F_b(M_b) = \tau E_b(M_b)$, where τ is an arbitrary positive weight. In the case that $c_p > 0, \forall b \in B, \forall p \in P_b$ in objective (1) it may

be necessary to use a relatively large value for τ in the penalty function to increase the efficacy of this approach.

It is clear from equation (7) that the function $F_b(M_b)$ is not linear, hence a reformulation is required to construct the SDAMRP as a mixed-integer program. The reformulation of $F_b(M_b)$ and the SDAMRP involves modelling the penalty function using a set of integer variables. To describe this, we first define Δ_b^i as a sequence of step sizes,

$$\Delta_b^i = F_b(i-1) - F_b(i) \quad \forall i = 1, \dots, R_b, \quad (8)$$

which is a decreasing sequence in i . Thus, the penalty function can be defined as the sum of this sequence,

$$F_b(M_b) = \begin{cases} \sum_{i=M_b+1}^{R_b} \Delta_b^i & \text{if } 0 \leq M_b < R_b, \\ 0 & \text{if } M_b = R_b. \end{cases} \quad (9)$$

Now, in the current form it is not possible to use equation (9) in the mixed-integer programming formulation of the SDAMRP. The difficulty with equation (9) is the lower bound in the summation, $M_b + 1$. By introducing the set of binary variables $\xi_b^i, i = 1, \dots, R_b$, it is possible to define $M_b = \sum_{i=1}^{R_b} \xi_b^i$. The reformulation of the SDAMRP is performed by replacing M_b by $\sum_{i=1}^{R_b} \xi_b^i$ in constraints (4) and $F_b(M_b)$ by $F_b(0) - \sum_{i=1}^{R_b} \Delta_b^i \xi_b^i$. This results in an linear objective function and a mixed-integer programming formulation of the SDAMRP. Furthermore, the non-negativity requirements on M_b in (5) are no longer necessary. These constraints are replaced by binary requirements on the ξ_b^i variables.

3 Recoverable robustness

The recoverable robust single day aircraft maintenance routing problem (SDAMRP-RR) is a feedback robust approach that attempts to achieve maintenance feasibility while simultaneously improving the recoverability of the aircraft routing solution. A key feature of recoverable robustness is identifying a planning solution that is recoverable in *limited effort*. In this paper *limited effort* is defined as the minimum expected recovery cost and reduced number of routing changes, which also describes improved recoverability. Formulating the aircraft routing problem as the SDAMRP, the recoverability of the planning stage solution is evaluated by solving an aircraft recovery problem for a set of disruption scenarios. This evaluation details the expected recovery costs and the level of intervention required as a result of the planning stage decisions. Minimising the weighted sum of the maintenance planning costs and the evaluation results is expected to provide improved recoverability for the SDAMRP. In this section a description of the recovery problem is provided along with the details regarding the integration of the planning and evaluation stages.

3.1 Evaluation stage

The recoverable robustness technique is solved with two distinct stages, the planning and evaluation stages. The planning stage involves solving the SDAMRP and the evaluation stage solves a recovery problem to explicitly evaluate the recoverability of the planning stage solution. A recovery problem is solved for each disruption scenario s contained in S , minimising the recovery costs and the number of changes made from the planned solution. The results from this evaluation is used to improve the recoverability of the planning stage solution.

3.2 Aircraft recovery one-day routes

The one-day routes formulation is used to define the aircraft routes in the recovery problem. However, in comparison to the planning problem, there are significant differences affecting their construction. Similar to the SDAMRP, the decision variables y_p^s equal 1 if route p is operated by an aircraft in scenario s and 0 otherwise. Each of the one-day routes originate from overnight bases, however the connections used prior to the disruption in scenario s are fixed. The fixing of these connections permits the omission of the related flights from the schedule and connection network. As such, an alternative set of flights, N^s , is used to define the recovery schedule for each scenario s and all feasible connections between the flights in N^s are contained in C^s .

Since the recovery flight schedule defined by the set N^s omits all flights preceding the disruption described by scenario s , the location of each aircraft at the start of the disruption must be known. The possible locations of an aircraft at the start of a disruption are contained in L , which describes all airports l within the flight network. At the time of a disruption, the location l for grounded aircraft is given by the current airport and if an aircraft is operating a flight, l is the destination airport of that flight. The parameters v_{lp}^s are introduced and equal 1 if in scenario s the aircraft operating one-day route p in the planning stage is positioned at airport l at the start of the disruption.

3.3 Recovery policies

The recovery problem is solved with a full set of recovery options, including aircraft rerouting and flight delays and cancellations. In the event of a disruption, the aircraft routes in the planning stage solution may not be feasible for the disrupted flight schedule N^s , hence the construction of new routes is required. The new aircraft routes are generated in the recovery process to respect the origination locations and maintenance requirements. A necessary feature of aircraft recovery problems is that maintenance critical aircraft are assigned one-day routes that still terminate at maintenance stations. This is enforced in the generation of recovery one-day routes in the SDAMRP-RR.

Flight delays are an important recovery action that involves retiming flights to increase the number of potential flight connections. Flight retiming options are implemented in the SDAMRP-RR using the

technique of flight copies, which is the duplication of each flight with progressively later departure times. Thengvall *et al.* [32] presents an example of this approach implemented for an aircraft recovery problem and Mercier and Soumis [24] describe its use for flight retiming in an integrated airline planning problem.

All flights in the recovery problem are assigned to an aircraft or, if that is not possible, they are cancelled. Flight cancellations in the SDAMRP-RR are modelled with the variables z_j^s , which equal 1 if flight j is cancelled in scenario s at a cost of c_j^s . The cost of a flight cancellation is modelled in an attempt to capture both the realised and unrealised costs of lost revenue and passenger dissatisfaction respectively.

The indirect costs of flight cancellations associated with passenger dissatisfaction are modelled in the SDAMRP-RR with the introduction of a *loss rate* parameter g^{LR} . The loss rate indicates the willingness of passengers to rebook with the airline, for example i) $g^{LR} < 1$ indicates that passengers are not discouraged to travel with the airline in the future and some of the revenue lost from the flight cancellation is recaptured by providing alternative travel arrangements, ii) $g^{LR} = 1$ does not account for any recapture or loss of future bookings, and iii) $g^{LR} > 1$ describes the situation where passengers are less likely to rebook with the airline in the future. The results will demonstrate the effect different values of g^{LR} has on the efficacy of the recoverable robustness technique.

Employing flight delays and cancellations in the recovery problem unfortunately destroys the flow balance of the original schedule. As a result, it is not guaranteed that a sufficient number of aircraft will terminate at each overnight airport as required for the following days flying. To address this aspect of recovery problems for the cyclic schedules used in this paper, the parameter R_b specifies the minimum number of aircraft required to terminate at each overnight base b . In the model constraints for the aircraft recovery problem, the parameters t_{bp}^s are introduced which equal 1 if route p in scenario s terminates at overnight airport b and 0 otherwise.

An important feature of the recoverable robustness technique is the objective of minimal deviation from the planned solution. This is implemented by attempting to construct aircraft routes in the recovery problem with the same connections used by aircraft in the solution to the planning stage. Now, each recovery problem in the evaluation stage is solved with a more restricted set of flights compared to the planning stage and the use of flight delays also affects departure times. Hence, the set \bar{C}^s is defined to implement the objective of minimal deviation by containing all connections common between the flight schedules of the planning stage and evaluation stage of scenario s . As stated previously, a one-day route p is defined by a sequence of connected flights, so the parameters e_{ijp} and e_{ijp}^s are defined to equal 1 if the connection $(i, j) \in \bar{C}^s$ is included in route p for the planning stage and recovery scenario s respectively. In addition, the objective of minimal deviation also requires the additional variables ϵ_{ij}^{s+} and ϵ_{ij}^{s-} to count the changes between the planning and evaluation stages.

The mathematical model for the SDAMRP-RR is given by,

$$\begin{aligned} \min \quad & \sum_{p \in P} c_p y_p + \sum_{b \in B} \left\{ F_b(0) - \sum_{i=1}^{R_b} \Delta_b^i \xi_b^i \right\} \\ & + \sum_{s \in S} w^s \left\{ \sum_{p \in P^s} c_p^s y_p^s + \sum_{j \in N^s} g^{LR} c_j^s z_j^s + \sum_{(i,j) \in C^s} g^{SW} \epsilon_{ij}^{s-} \right\}, \end{aligned} \quad (10)$$

$$\text{s.t.} \quad \sum_{b \in B} \sum_{p \in P_b} a_{jp} y_p = 1 \quad \forall j \in N, \quad (11)$$

$$\sum_{p \in P_b} y_p \leq R_b \quad \forall b \in B, \quad (12)$$

$$\sum_{i=1}^{R_b} \xi_b^i - \sum_{p \in P_b} o_p y_p = 0 \quad \forall b \in B, \quad (13)$$

$$\sum_{b \in B} \sum_{p \in P_b^s} a_{jp}^s y_p^s + z_j^s = 1 \quad \forall s \in S, \forall j \in N^s, \quad (14)$$

$$\sum_{p \in P_b^s} y_p^s \leq R_b \quad \forall s \in S, \forall b \in B, \quad (15)$$

$$\sum_{b_1 \in B} \sum_{p \in P_{b_1}^s} t_{b_2 p}^s y_p^s \geq R_{b_2} \quad \forall s \in S, b_2 \in B, \quad (16)$$

$$\sum_{b \in B} \sum_{p \in P_b^s} v_{lp}^s y_p^s - \sum_{b \in B} \sum_{p \in P_b} v_{lp}^s y_p = 0 \quad \forall s \in S, \forall l \in L, \quad (17)$$

$$\sum_{b \in B} \sum_{p \in P_b^s} o_p^s v_{lp}^s y_p^s - \sum_{b \in B} \sum_{p \in P_b} o_p v_{lp}^s y_p \geq 0 \quad \forall s \in S, \forall l \in L, \quad (18)$$

$$\sum_{b \in B} \sum_{p \in P_b} e_{ijp} y_p - \sum_{b \in B} \sum_{p \in P_b^s} e_{ijp}^s y_p^s = \epsilon_{ij}^{s+} - \epsilon_{ij}^{s-} \quad \forall s \in S, \forall (i, j) \in \bar{C}^s, \quad (19)$$

$$y_p \in \{0, 1\} \quad \forall b \in B, \forall p \in P_b, \quad y_p^s \in \{0, 1\} \quad \forall s \in S, \forall b \in B, \forall p \in P_b^s, \quad (20)$$

$$z_j^s \in \{0, 1\} \quad \forall s \in S, \forall j \in N^s, \quad \epsilon_{ij}^{s+}, \epsilon_{ij}^{s-} \geq 0 \quad \forall s \in S, \forall (i, j) \in \bar{C}^s, \quad (21)$$

$$\xi_b^i \in \{0, 1\} \quad \forall b \in B, \forall i \in \{1, \dots, R_b\}. \quad (22)$$

The SDAMRP-RR combines the planning stage of the SDAMRP with an evaluation stage provided by aircraft recovery subproblems. The objective of the SDAMRP-RR is the sum of the objective from the SDAMRP and the weighted cost of recovery and penalties resulting from connection changes in each scenario. The modifications described in Section 2.4 to formulate the SDAMRP as a mixed-integer program have been applied in the formulation of the SDAMRP-RR, given by objective (10) and constraint (13). Constraints (11)-(13) describe the SDAMRP as the planning stage for the recoverable robust problem.

The constraints related to the evaluation stage of the SDAMRP-RR are given by (14)-(19). Constraints (14) are the flight coverage constraints that include the decision variables z_j to describe flight cancellations. An upper bound on the number of aircraft one-day routes departing from each overnight base in a feasible recovery solution is given by constraints (15). Since flow balance is not guaranteed in

the recovery schedule, the number of aircraft required to terminate at each overnight base is provided by constraints (16).

The link between the planning and recovery stages is provided by constraints (17)-(19). Constraints (17) are included to provide the number of aircraft located at each airport at the start of a disruption for each scenario s . As a key feature of the SDAMRP-RR, the aircraft expected to receive maintenance at the end of the day are still assigned a one-day route that terminates at a maintenance station in the recovered solution. Constraints (18) ensure that the recovered solution provides at least the same number of routes originating from airport l that terminate at maintenance stations as given in the solution to the planning stage. Finally, constraints (19) are included in the SDAMRP-RR to count the number of connections used in the planned solution that are not operated as a result of recovery, which is a measure of *limited effort*.

4 Solution methodology

The SDAMRP and SDAMRP-RR are large-scale optimisation problems that require the use of decomposition techniques to develop efficient solution approaches. The decomposition techniques employed in this paper are column generation and Benders' decomposition. Both of these solution techniques have been widely demonstrated to improve the tractability and reduce solution runtimes of large-scale optimisation problems.

The aim of solving the SDAMRP and SDAMRP-RR is to identify an implementable aircraft routing solution with improved maintenance planning. For the SDAMRP, this is achieved by imposing integrality requirements on the variables y_p . In the SDAMRP-RR, the solution process involves two different stages, both of which are presented in (10)-(22) with integer one-day route variables. While integrality is important in the evaluation stage, it is not required to identify an implementable aircraft routing solution. Through experiments it is observed that the evaluation stage LP solution is integer optimal for approximately 75% of the scenario set. Thus, solving the evaluation stage as a continuous problem does not greatly affect the solution quality. This modification to the SDAMRP-RR aids the application of Benders' decomposition.

4.1 Benders' decomposition

Benders' decomposition is a solution approach originally proposed to solve mixed-integer programming problems [4]. This solution approach has received much attention in regards to stochastic programming problems, through the implementation of the popular L-shaped method. The structure of the SDAMRP-RR is similar to a stochastic program, where the first-stage (deterministic) variables are related to the planning aircraft routing problem and the second-stage (probabilistic) variables are related to the recovery problems for each scenario. Given the similarity between the SDAMRP-RR and stochastic

programs, Benders' decomposition is applied in the form of the L-shaped method to achieve an efficient solution approach.

In the SDAMRP-RR, there is a clear separation of variables between the planning and recovery problems to define the master and subproblems respectively. The subproblems given by this decomposition, which we label as the primal Benders' subproblem (PBSP- s), are defined for each scenario s , described with the variables y_p^s , z_j^s , ϵ_{ij}^{s+} and ϵ_{ij}^{s-} and the related constraints. Each subproblem is solved for a given disruption scenario s to evaluate the effect of planning decisions on the cost and ease of recovery.

The Benders' master problem (BMP) for the SDAMRP-RR is defined as a planning aircraft routing problem given by the SDAMRP. The BMP is formulated to contain the variables y_p^s and the constraints (11)-(13). In addition, a set of cuts are included to reflect the decisions made in each PBSP- s . The decision variables φ^s are introduced in the added cuts to provide a lower bound on the objective value of the PBSP- s for each scenario s . The BMP formulated without any added cuts is identical to the SDAMRP. Since the added cuts restrict the feasible region, all feasible solutions to the BMP are feasible for the SDAMRP.

Evaluating the recoverability of the planning solution is performed by solving the PBSP- s with fixed solution values from the BMP. The fixed solution to the BMP, given by $\bar{y}_n = \{\bar{y}_{np}, \forall b \in B, \forall p \in P_b | \bar{y}_{np} = 1\}$, represents the best possible planning solution in iteration n achieved with the previously realised evaluation information. Fixing the solution to the planning stage significantly reduces the complexity of the original problem by permitting the evaluation stage to be solved purely as a series of recovery problems.

The PBSP- s is defined as an aircraft recovery problem given by,

$$\mu_s(\bar{\mathbf{y}}_n) = \min \sum_{p \in P^s} c_p^s y_p^s + \sum_{j \in N^s} g^{LR} c_j^s z_j^s + \sum_{(i,j) \in C^s} g^{SW} \epsilon_{ij}^{s-}, \quad (23)$$

$$\text{s.t.} \quad \sum_{b \in B} \sum_{p \in P_b^s} a_{jp}^s y_p^s + z_j^s = 1 \quad \forall j \in N^s, \quad (24)$$

$$\sum_{p \in P_b^s} y_p^s \leq R_b \quad \forall b \in B, \quad (25)$$

$$\sum_{b_1 \in B} \sum_{p \in P_{b_1}^s} t_{b_2 p}^s y_p^s \geq R_{b_2} \quad b_2 \in B, \quad (26)$$

(PBSP- s)

$$\sum_{b \in B} \sum_{p \in P_b^s} v_{lp}^s y_p^s = \sum_{b \in B} \sum_{p \in P_b} v_{lp}^s \bar{y}_{np} \quad \forall l \in L, \quad (27)$$

$$\sum_{b \in B} \sum_{p \in P_b^s} o_p^s v_{lp}^s y_p^s \geq \sum_{b \in B} \sum_{p \in P_b} o_p v_{lp}^s \bar{y}_{np} \quad \forall l \in L, \quad (28)$$

$$\sum_{b \in B} \sum_{p \in P_b^s} e_{ijp}^s y_p^s + \epsilon_{ij}^{s+} - \epsilon_{ij}^{s-} = \sum_{b \in B} \sum_{p \in P_b} e_{ijp} \bar{y}_{np} \quad \forall (i,j) \in \bar{C}^s, \quad (29)$$

$$y_p^s \geq 0 \quad \forall p \in P^s, \quad z_j^s \geq 0 \quad \forall j \in N^s, \quad (30)$$

$$\epsilon_{ij}^{s+}, \epsilon_{ij}^{s-} \geq 0 \quad \forall (i,j) \in \bar{C}^s. \quad (31)$$

The objective of the PBSP- s minimises the cost of recovery, in particular flight delay and cancellation costs, and penalises any changes from the connections set in the planning stage. The dual variables for this problem are defined as $\alpha^s = \{\alpha_j^s, \forall j \in N^s\}$, $\beta^s = \{\beta_b^s, \forall b \in B\}$, $\delta^s = \{\delta_b^s, \forall b \in B\}$, $\gamma^s = \{\gamma_l^s, \forall l \in L\}$, $\lambda^s = \{\lambda_l^s, \forall l \in L\}$ and $\rho^s = \{\rho_{ij}^s, \forall (i,j) \in \bar{C}^s\}$ for the constraints (24)-(29) respectively.

The general application of Benders' decomposition solves the PBSP- s to identify either a feasibility or optimality cut in each iteration. A feasibility cut is added to the BMP in the event that the PBSP- s is proved to be infeasible, conversely an optimality cut is added if a feasible solution exists. Because of the similar structure of the PBSP- s and BMP, an initial feasible solution can be constructed for the PBSP- s using $\bar{\mathbf{y}}_n$. Consequently, only optimality cuts are generated from the PBSP- s and added to the BMP. The Benders' optimality cut is given by,

$$\varphi^s \geq \sum_{j \in N^s} \alpha_j^s + \sum_{b \in B} \left\{ R_b (\beta_b^s + \delta_b^s) + \sum_{l \in L} \sum_{p \in P_b} v_{lp}^s \bar{y}_{np} (\gamma_l^s + o_p \lambda_l^s) + \sum_{(i,j) \in \bar{C}^s} \sum_{p \in P_b} e_{ijp} y_p \rho_{ij}^s \right\}. \quad (32)$$

Given an optimal solution to the PBSP- s , the dual solution used to construct the cut given by inequality (32) represents an extreme point of the dual feasible region. Since the PBSP- s is a highly degenerate problem, there are potentially many extreme points in the dual problem that provide the same objective value. The strength of the cuts constructed from each of these extreme points when added to the BMP is not identical, as such there are many approaches that may be applied to select the strongest.

The rate of convergence of the Benders' decomposition algorithm is highly dependent on the strength of the cuts added to the master problem. A common approach used to identify the extreme point in

the dual problem providing the strongest optimality cut, called the Pareto-optimal cut, is presented by Magnanti and Wong [22]. The use of the Magnanti-Wong method for identifying Pareto-optimal cuts is demonstrated in the solution process of various airline optimisation problems [12, 23, 26]. These papers report significant improvements in solution runtimes as a result of implementing this method. The Magnanti-Wong method has been applied to find Pareto-optimal cuts from the solution of the PBSP- s in this paper in an attempt to achieve similar runtime improvements.

The BMP is solved to find the best possible planning solution relative to the cuts added from the PBSP- s . All Benders' cuts ω added to the BMP from the PBSP- s for scenario s are contained in the set Ω^s . The BMP is given by,

$$\Phi_n = \min \sum_{p \in P} c_p y_p + \sum_{b \in B} \left\{ F_b(0) - \sum_{i=1}^{R_b} \Delta_b^i \xi_b^i \right\} + \sum_{s \in S} w^s \varphi^s, \quad (33)$$

$$\text{s.t.} \quad \sum_{b \in B} \sum_{p \in P_b} a_{jp} y_p = 1 \quad \forall j \in N, \quad (34)$$

$$\sum_{p \in P_b} y_p \leq R_b \quad \forall b \in B, \quad (35)$$

$$\text{(BMP)} \quad \sum_{i=1}^{R_b} \xi_b^i - \sum_{p \in P_b} o_p y_p = 0 \quad \forall b \in B, \quad (36)$$

$$\begin{aligned} \varphi^s - \sum_{b \in B} \left\{ \sum_{l \in L} \sum_{p \in P_b} v_{lp}^s \bar{y}_p (\gamma_l^s + o_p \lambda_l^s) + \sum_{(i,j) \in \bar{C}^s} \sum_{p \in P_b} e_{ijp} y_p \rho_{ij}^{s\omega} \right\} \\ \geq \sum_{j \in N^s} \alpha_j^s + \sum_{b \in B} R_b (\beta_b^s + \delta_b^s) \quad \forall s \in S, \forall \omega \in \Omega^s, \end{aligned} \quad (37)$$

$$y_p \in \mathbb{Z}^+ \quad \forall b \in B, \forall p \in P_b, \quad \xi_b^i \in \{0, 1\} \quad \forall b \in B, \forall i \in \{1, \dots, R_b\}, \quad (38)$$

$$\varphi^s \geq 0 \quad \forall s \in S. \quad (39)$$

The objective of the BMP minimises the cost of aircraft routing, the penalties arising from maintenance misalignments and the weighted sum of the lower bounds on the PBSP- s objective values. The constraints (34)-(36) are identical to constraints (2)-(4) from the SDAMRP after applying the modifications described in Section 2.4. The additional constraints representing Benders' optimality cuts constructed from the dual solutions to the PBSP- s are given by (37).

The lower and upper bounds on the objective value for the PBSP- s in iteration n are given by the solution value of φ_n^s and the objective value $\mu_s(\bar{\mathbf{y}}_n)$ respectively. The iterative solution process between the BMP and the PBSP- s progressively restricts the feasible region with the addition of cuts to improve upon the lower bounds given by φ^s . This improvement in the bounds for each scenario results in a strictly non-decreasing sequence for the lower bound on the objective value of the original problem, given by Φ_n . Simultaneously, the addition of cuts also has the effect of reducing the upper bound on the objective value of the original problem. However, the sequence of upper bounds from each iteration

is not strictly non-increasing. The upper bound on the original problem in iteration n is given by,

$$\Phi_n^{UB} = \sum_{p \in P} c_p \bar{y}_{np} + \sum_{b \in B} \left\{ F_b(0) - \sum_{i=1}^{R_b} \Delta_b^i \bar{\xi}_{nb}^i \right\} + \sum_{s \in S} \mu_s(\bar{\mathbf{y}}_n), \quad (40)$$

where $\bar{\xi}_n = \{\bar{\xi}_{nb}^i, \forall b \in B, \forall i \in \{1, \dots, R_b\} | \bar{\xi}_{nb}^i = 1\}$ are the fixed solution values for the maintenance count variables. The best upper bound identified during the solution process is labelled as $\hat{\Phi}^{UB} = \max_n \{\Phi_n^{UB}\}$. Using the lower and upper bounds on the original problem, the optimality gap in iteration n is calculated by,

$$Gap_n = \frac{\hat{\Phi}^{UB} - \Phi_n}{\Phi_n}. \quad (41)$$

The addition of cuts to the BMP is expected to improve the lower bound on the objective function value. This expectation is given by the relative difference between the lower and upper bounds, given by φ_n^s and $\mu_s(\bar{\mathbf{y}}_n)$ respectively, for each scenario subproblem. Papadakos [26] describes a termination condition for the Benders' decomposition solution process based upon this expectation, which is given by,

$$\frac{\mu_s(\bar{\mathbf{y}}_n) - \varphi_n^s}{\Phi_n} \leq \varepsilon. \quad (42)$$

The tolerance level, ε , that is used for this problem is 10^{-4} . The violation of this condition for scenario s indicates that an additional cut is required to improve the lower bound given by the objective value of the BMP.

The optimality gap and cut addition condition given by relations (41) and (42) define the termination criteria for the Benders' decomposition solution process for the SDAMRP-RR. The solution process will terminate when either Gap_n falls below some tolerance level, which is set at 1% for the SDAMRP-RR, or inequality (42) is satisfied by all scenarios in a single iteration. While it is expected that the optimality gap stopping condition will be reached first, the latter condition is implemented primarily to identify when cuts must be added from each subproblem. The benefit of the optimality gap condition is that it takes a global view of the problem, terminating the solution process when a desired gap is reached.

4.2 Column generation

Column generation is applied to optimisation problems formulated with a large number of variables displaying a special combinatorial structure. The three problems SDAMRP, BMP and PBSP- s display a problem structure permitting the use of column generation. Since one-day route variables are used in the formulation for each of these problems, the implementations of the column generation solution approach are very similar. A description of the implementation of column generation will be provided with respect to the PBSP- s in this section. For completeness, the differences in the application of the column generation solution approach for each of the problems will also be highlighted and explained.

The column generation solution process initially defines a restricted master problem (RMP) as a restriction on the original problem by containing only a subset of all possible variables. In this paper,

the restriction is defined by the sets $\bar{P}_b \subseteq P_b$ and $\bar{P}_b^s \subseteq P_b^s$ which replace P_b and P_b^s respectively in the SDAMRP, BMP and PBSP- s . The solution to the RMP provides an upper bound on the objective value of the original problem, which is improved with the addition of variables to the sets \bar{P}_b and \bar{P}_b^s .

4.2.1 Column generation for the PBSP- s

The one-day route variables in the PBSP- s define paths through a connection network that originate and terminate at overnight bases. Given this variable structure, the set of all possible variables P_b^s can be described by the feasible region of a network flow problem. Using the dual variables defined in Section 4.1, the reduced cost of a one-day route p originating from base b is given by,

$$\hat{c}_p^{sb} = c_p^{sb} - \sum_{j \in N^s} a_{jp}^s \alpha_j^s - \beta_b^s - \sum_{b_2 \in B} t_{b_2 p}^s \delta_{b_2}^s - \sum_{l \in L} \gamma_{lb}^s - \sum_{(i,j) \in \bar{C}^s} e_{ijp}^s \rho_{ij}. \quad (43)$$

In the optimal solution to the RMP, all variables contained in \bar{P}_b^s have a reduced cost of $\hat{c}_p^{sb} \geq 0$ and $\hat{c}_p^{sb} = 0$ for all variables contained in the basis. This indicates that any improvement in the objective value of the RMP is only achieved through the addition of variables from the sets $P_b^s \setminus \bar{P}_b^s, \forall b \in B$. By setting the objective function of a shortest path problem to (43), the solution will identify variables with a reduced cost $\hat{c}_p^{sb} \leq 0$ that are expected to improve the objective function value of the RMP. The column generation subproblem for the PBSP- s describes a classical shortest path problem, permitting the use of a wide variety of dedicated solution algorithms. Since the connection network used for the PBSP- s forms an acyclic directed graph, the shortest path problem can be efficiently solved using an algorithm for acyclic networks presented in Ahuja *et al.* [2].

4.2.2 Column generation for the BMP and SDAMRP

Since the variables in the BMP describe the flow of aircraft through a connection network similar to the PBSP- s , a shortest path problem can also be solved to identify the minimum reduced cost variables. The major difference between the shortest path problems solved to identify negative reduced cost columns for the BMP and the PBSP- s is the form of the objective function. Since the RMP formulated for the BMP has a different structure to the PBSP- s , a different set of dual variables must be defined. The dual variables for the BMP are given by $\mathbf{u} = \{u_j, \forall j \in N\}$, $\mathbf{v} = \{v_b, \forall b \in B\}$, $\mathbf{w} = \{w_b, \forall b \in B\}$ and $\boldsymbol{\kappa} = \{\kappa^{s\omega}, \forall s \in S, \forall \omega \in \Omega^s\}$ for constraints (34)-(37) respectively. Therefore, the reduced cost function of a variable in the BMP originating from overnight airport b is given by,

$$\hat{c}_p^b = c_p^b - \sum_{j \in N^s} a_{jp} u_j^s - v_b + o_p^b w_b - \sum_{s \in S} \sum_{\omega \in \Omega^s} \sum_{(i,j) \in \bar{C}^s} e_{ijp} \kappa^{s\omega}. \quad (44)$$

Solving a shortest path problem with objective function (44) identifies negative reduced cost variables for the BMP. Since the connection network defined for the BMP can be sorted in a topological order, this shortest path problem can be efficiently solved by the algorithm described for the PBSP- s .

The connection networks defined for the SDAMRP and BMP are identical, as such the same set of feasible columns is defined for both problems. This implies that the shortest path problem employed to identify negative reduced cost columns for the BMP can also be used for the SDAMRP. The only modification required is the elimination of the final term in equation (44) to define the objective function of the shortest path problem for the SDAMRP. It follows that the same solution algorithm used for the PBSP-*s* and the BMP is applicable for the SDAMRP.

4.3 Acceleration techniques

The standard implementation of the Benders' decomposition solution process is commonly affected by slow convergence to the optimal solution. One aspect affecting the convergence of the algorithm is the efficacy of the cuts generated by the subproblems. A number of approaches have been proposed to address this difficulty by improving the strength of each added cut to tighten the feasible region of the master problem quickly and efficiently [11, 22, 25]. As stated in Section 4.1, the Magnanti-Wong method [22] has been implemented in the solution process for the SDAMRP-RR to identify Pareto-optimal cuts.

To achieve fast convergence of the Benders' decomposition solution approach for large-scale optimisation problems, additional enhancement techniques are required. The enhancements of the two-phase method, the use of parallel computing and a trust region method have been implemented for the SDAMRP-RR. The following sections discuss each of these acceleration techniques, describing their implementation in the solution process for the SDAMRP-RR.

4.3.1 Two-phase method

The two-phase method is implemented by Froyland *et al.* [12] and is demonstrated to be a successful approach to reduce solution runtimes. This method is an adaptation of the three-phase method developed for solving integrated airline planning problems [7, 23, 26]. The additional phase of the three-phase method solves the Benders' decomposition subproblems to integral optimality. Given the objective of the SDAMRP-RR is to identify an implementable solution with improved recoverability, it is only necessary to execute two phases of the three-phase algorithm.

The first stage of the two-phase method solves the linear programming (LP) relaxations of the BMP and the PBSP-*s*. The runtimes to solve the LP relaxations of the BMP and the PBSP-*s* is much less than the respective integer programs, allowing for the addition of more valid cuts early in this solution process. The first stage terminates when one of the stopping conditions presented in Section 4.1 is satisfied, solving the LP relaxation of the SDAMRP-RR to optimality. Since the convex hull of the feasible region for the SDAMRP-RR is contained within the LP relaxation, all added cuts in the first stage are valid for the integer problem.

The integrality requirements for the variables in the BMP are reintroduced in the second stage of the two-phase method. This stage solves the BMP as an integer program, applying branch-and-price to

identify the integer optimal solutions. The solution to the BMP in each iteration of the Benders' decomposition solution approach is a feasible solution to the original problem, with the best implementable solution identified when a stopping condition from Section 4.1 is satisfied.

4.3.2 Parallel computing

The Benders' decomposition framework lends itself to parallel computing methods as demonstrated by Linderoth and Wright [21]. Since each of the recovery subproblems in the SDAMRP-RR are completely separable and solved independently, it is possible to solve each of these concurrently. The concurrent approach for solving the Benders' decomposition subproblems does not affect the solution methods for each individual problem, as such the generation and addition of optimality cuts remains the same. The only difference between the sequential and concurrent approach is the implementation of the Benders' decomposition process to include the distribution of subproblems onto a number of compute threads.

The distribution of the subproblems for the SDAMRP-RR attempts to achieve the minimum amount of idle time for all threads. In the first iteration, the subproblems are assigned to threads in numerical order, where the first n subproblems are each solved on one of the n threads. Once a thread finishes executing the solution process for the current subproblem, the next subproblem in the list is then solved on that thread. The evaluation stage for the current iteration concludes when all subproblems have been assigned to a thread and the execution of each thread has completed.

For the subsequent iterations, the solution runtime for each subproblem in the previous iteration is used in the assignment process of subproblems to threads. The subproblems are sorted by solution runtime in descending order and the n subproblems with the longest runtime in the previous iteration are each assigned to one of the n threads. As in the previous case, once a thread completes the execution of the solution process, the next subproblem in the runtime sorted list is assigned to that thread.

Significant time reduction of the Benders' decomposition solution process is observed by the implementation in a parallel computing environment. The results will detail the difference in the optimality gaps achieved with the sequential and concurrent solution approaches.

4.3.3 Trust region

An undesired effect of the Benders' decomposition solution approach is the movement of the master problem solution to different parts of the feasible region between consecutive iterations. The large distances between solutions may render many of the added cuts ineffective and increase the number of iterations executed in the Benders' decomposition solution process. A trust region addresses this issue by forcing the solution in the current iteration to be found within a neighbourhood around the master problem solution from the previous iteration. This is implemented with either a regularisation term in the master problem objective function [29] or as an additional set of constraints [21, 30].

The trust region in the SDAMRP-RR introduces a set of constraints to *measure* the distance between

solutions of consecutive iterations. As explained in Section 2, a one-day route describes a set of connected flights to be performed by a single aircraft. Therefore, the optimal solution to the BMP is described by the set of connections that are used in the aircraft one-day routes. Thus, a measure of the distance between solutions is the number of different connections that are used.

To implement a trust region in the SDAMRP-RR the following set of constraints are added to the BMP,

$$\sum_{b \in B} \sum_{p \in P_b} e_{ijp} y_p + \epsilon_{ij}^+ - \epsilon_{ij}^- = \sum_{b \in B} \sum_{p \in P_b} e_{ijp} \bar{y}_{(n-1)p} \quad \forall (i, j) \in C. \quad (45)$$

This set of constraints is similar to the linking constraints (29) in the PBSP- s , as such the same definition for the counting variables ϵ_{ij}^+ and ϵ_{ij}^- is used. To minimise the number of connection changes between solutions of consecutive iterations, the term $\sum_{(i,j) \in C} g^{TR} \epsilon_{ij}^-$ is added to the objective function of the BMP.

While the trust region is used to restrict the feasible region of the master problem, it is important to ensure that the resulting objective function value is less than the best found upper bound. This is achieved by including a constraint to impose an upper bound on the objective function value. Such a constraint for the BMP is given by,

$$\sum_{p \in P} c_p y_p + \sum_{b \in B} \left\{ F_b(0) - \sum_{i=1}^{R_b} \Delta_b^i \zeta_b^i \right\} + \sum_{s \in S} w^s \varphi^s \leq \hat{\Phi}^{UB}. \quad (46)$$

The addition of the trust region and the upper bound on the objective value introduces a two-stage solution process for the BMP. The first stage solves the BMP as described in Section 4.1 to identify the current lower bound on the objective of the original problem. The second stage adds the constraints (45) and (46) to the BMP, which is then solved to identify a solution within a neighbourhood of the solution from the previous iteration. While it is only necessary to solve the second-stage of this process to generate optimality cuts, the lower bound identified in the first stage is used in the stopping condition described by (41). In addition, it is possible to replace the right-hand side of (46) with $\min\{\zeta \hat{\Phi}, \hat{\Phi}^{UB}\}$, where $\zeta \geq 1$. This ensures the solution to the second stage will remain close to the current lower bound and the solution from the previous iteration. Once a terminating condition is satisfied, the optimal solution to the SDAMRP-RR is found by continuing the Benders' decomposition process by only executing the first stage of this process.

Solving the BMP using a trust region identifies a suboptimal solution at each iteration of the Benders' decomposition solution process. While the optimality of the BMP affects the cuts generated in the evaluation stage, it is demonstrated that the finite convergence of the solution process does not require an optimal BMP solution [8, 13, 27]. For example, Geoffrion and Graves [13] describe a solution process that only identifies feasible suboptimal solutions to the BMP at each iteration. Optimality is achieved when no further feasible solutions exist with an objective value less than the best found upper bound. While our implementation is similar to Geoffrion and Graves [13], the stopping conditions presented in

Section 4.1 significantly alters the solution approach. Terminating the Benders' decomposition solution process prior to all feasible solutions being identified greatly improves the solution runtimes.

5 Computational results

The improved maintenance planning approach presented by the SDAMRP is evaluated by comparing the resulting solution with that of an aircraft routing problem solved without any maintenance considerations. To provide this set of benchmark results, an aircraft routing problem is formulated from the SDAMRP with the maintenance counting constraints (4) eliminated and the term $\sum_{b \in B} F_b(M_b)$ removed from the objective. Both the SDAMRP and the benchmark problem is solved for four different flight schedules and the resulting number of maintenance misalignments is compared.

The recoverability improvement achieved by solving the SDAMRP-RR is evaluated by simulating the recovery of disruption scenarios for the SDAMRP and SDAMRP-RR solutions and comparing the resulting recovery costs. A selection of the flight schedules used to evaluate the SDAMRP are employed in the analysis of the SDAMRP-RR.

5.1 Description of schedule data and parameters

The characteristics of the flight schedules used for our computational experiments are given in Table 1. The key features of each flight schedule is the set of airports where one-day routes terminate, called overnight bases, and the subset of these where maintenance can be performed.

	F267_A49	F578_A153	F1165_A289	F3370_A526
Flights	267	578	1165	3370
Aircraft	49	153	289	526
Airports	20	50	97	73
Overnight(Maintenance) bases	12(1)	41(2)	67(5)	73(10)

Table 1: Flight schedule details.

Traditionally the aircraft routing problem is solved following the fleet assignment, which partitions the schedule into sets of flights, one for each fleet type. Since the aircraft assigned to a single partition of the schedule are all of the same type, the traditional formulation of the aircraft routing problem describes a feasibility problem. Solving the SDAMRP, and similarly the SDAMRP-RR, as a feasibility problem is accomplished by setting the cost of each one-day route to zero, i.e. $c_p = 0, \forall b \in B, \forall p \in P_b$. This problem formulation is referred to as solving the SDAMRP with a zero connection cost function (ZCF).

5.1.1 Parameters for the SDAMRP-RR

To thoroughly review the SDAMRP-RR, experiments are performed by also solving the planning stage using a proxy robust connection cost function introduced by Grönkvist [15] (GCF). The GCF is implemented by Froyland *et al.* [12], and hence a direct comparison to [12] is made by applying the identical parameter settings in this paper. Solving the SDAMRP-RR using the ZCF and GCF will demonstrate the impact that connection cost functions have on the solution runtimes and recoverability improvement.

The definition of *limited effort* is given in Section 3 in terms of expected recovery costs and the number of changes made during recovery. The latter part of this definition is modelled using a swap cost g^{SW} , which can be viewed as a measure of the effort required to perform any aircraft swaps. This parameter directly affects the flexibility of the recovery solution by minimising the number of connection changes that are made.

The largest costs in the evaluation stage of the SDAMRP-RR are related to flight delays and cancellations. Flight delay decisions are made during the construction of recovered one-day routes. As such, the one-day route cost is directly related to the cost of any delay decisions that are made. Alternatively, flight cancellation decisions are given by a set of variables in the SDAMRP-RR, hence the cancellation costs are explicitly included in the objective function. The cost of a flight cancellation includes realised and unrealised costs, with a *loss rate* g^{LR} presented in Section 3.3 as a quantitative measure of passenger dissatisfaction. To thoroughly analyse the effect of different loss rates on the SDAMRP-RR, the results are presented with $g^{LR} \in \{0.5, 1.0, 1.5, 2.0\}$.

A selection of the flight schedules presented in Table 1 are used to evaluate the SDAMRP-RR, namely the F267_A49 and F578_A153 schedules. These schedules are selected to provide an example of the potential recoverability improvements that can be achieved by the application of recoverable robustness. In [12], the recoverable robust tail assignment problem is evaluated on a schedule that contains 53 flights and 10 aircraft. Hence, the evaluation of the SDAMRP-RR extends upon the analysis presented in [12] by assessing the applicability of the recoverable robust framework to larger flight schedules.

5.1.2 Evaluation scenarios for the SDAMRP-RR

The evaluation stage of the SDAMRP-RR involves assessing the recoverability of the current planning solution by solving recovery problems for a set of disruption scenarios. The disruption type selected for each scenario in the evaluation stage is an airport closure. The motivation for selecting this disruption type is that airport closures generally require significant intervention by the airline during recovery. Further, it is believed that airport closure scenarios will provide a good example of the benefits that are achievable with the recoverable robustness framework. In practice, it is trivial to implement different disruption types in the evaluation stage to match the desired outcomes of each airline.

The airport closure scenarios are generated from the original flight schedule to represent closures at different airports with a range of starting times and durations. The airports where at least 5% of all

arrivals or departures occur are selected as affected airports for the closure scenarios. For the F267_A49 data set there are 6 airports that satisfy this criteria and for the F578_A153 there are 4. To assess the impact of disruptions on maintenance planning, the airports where aircraft maintenance can be performed are also included in the set of affected airports for both data sets.

To provide a broad range of scenarios at each airport the earliest and latest start times for the closures are 6:30am and 5:00pm local time respectively, and the durations range between 1 and 5 hours. The set of scenarios is constructed so there are equal numbers of closure start times and durations. Using the F267_A49 data set as an example, if the evaluation stage is solved with 24 scenarios, 4 are generated for each airport by using 2 different starting times and durations. While each of the scenarios are not equally likely, we wish to improve the recoverability across the complete set of scenarios, as such a uniform distribution is used.

5.2 Analysis of maintenance misalignment

The solution approach used as a benchmark for this analysis assumes that the aircraft routing and maintenance planning are performed separately. Given the aircraft routing solution, modifications are made *a posteriori* to ensure all aircraft receive maintenance as required. This is the airline planning approach considered by Feo and Bard [10], Gopalan and Talluri [14], Sriram and Haghani [31], and Lapp and Cohn [19]. The solution to the SDAMRP attempts to identify an aircraft routing solution that requires few manual modifications to satisfy all maintenance requirements. As such, we believe that a comparison against the standard aircraft routing solution is acceptable for this analysis.

The improved maintenance planning by the SDAMRP is assessed by calculating the total number of maintenance misalignments across all airports. The number of maintenance misalignments at overnight airport b is calculated by

$$MM_b = \max \left\{ 0, \frac{R_b}{6} - \sum_{i=1}^{R_b} \xi_b^i \right\}, \quad (47)$$

and the total number of maintenance misalignments is given by $MM_{Tot} = \sum_{b \in B} MM_b$. This value is calculated for both the standard aircraft routing problem and the SDAMRP. The maintenance misalignment results are reported in Table 2.

The results in Table 2 demonstrate that the SDAMRP significantly reduces the number of main-

<i>Misalignment</i>	F267_A49	F578_A153	F1165_A289	F3370_A526
Aircraft Routing	2.67	3.5	6.67	3.0
SDAMRP	1.17	0.5	0.0	0.0
Improvement (%)	56.25	85.71	100.0	100.0
ARP - Runtime (sec)	1.61	6.4	155.82	67675.0
SDAMRP - Runtime (sec)	7.27	6.97	168.84	65087.0

Table 2: Maintenance misalignments and solution runtimes for different flight schedules.

tenance misalignments compared to the standard aircraft routing solution. The improvement in the number of maintenance misalignments across all data sets is at least 56.25%, and for two of the flight schedules the number of misalignments is reduced to zero. This result is very important since the SDAMRP is solved at the equivalent stage as the standard aircraft routing problem, so any maintenance misalignments that exist requires an additional maintenance planning problem to be solved. The results demonstrate that it is possible to completely eliminate the need for any further modification of the aircraft routing solution by solving the SDAMRP.

5.3 Enhancement approaches for the SDAMRP-RR

Benders' decomposition is applied to the SDAMRP-RR to improve the tractability of the problem. However, the standard implementation of this solution approach exhibits poor convergence to the optimal solution. A number of techniques are introduced in Section 4.3 to improve the convergence of the Benders' decomposition solution approach and reduce solution runtimes. In this section the performance of the parallel computing and trust region enhancements are assessed, demonstrating the improved convergence achieved by their implementation. The experiments solve the SDAMRP-RR for the F267_A49 and F578_A153 data sets using the GCF and 50 scenarios in the evaluation stage.

The evaluation of the different enhancement techniques for the F267_A49 and F578_A153 data sets is presented in Table 3. It is clear that the use of all enhancements is necessary to achieve the best convergence for the Benders' decomposition solution process for both data sets. This is not surprising since the two enhancement techniques focus on two separate parts of the solution process. However, it is interesting to note that the trust region method achieves a better convergence compared to using a parallel computing environment for the F267_A49 data set. This demonstrates the strength of the trust region method to focus the solutions of the Benders' master problem between consecutive iterations,

Swap Penalty	All Enhancements	Parallel Only	Trust Region Only	No Enhancements
<i>F267_A49</i>				
2000	1.26%	1.64%	1.34%	3.75%
5000	1.29%	6.88%	2.44%	11.92%
7500	2.82%	11.68%	2.84%	19.06%
<i>F578_A153</i>				
2000	0.33%	0.65%	0.57%	4.19%
5000	0.41%	1.91%	9.54%	9.37%
7500	0.63%	3.44%	16.88%	12.8%

Table 3: Optimality gap, given by (41), of the SDAMRP-RR solved for the F267_A49 and F578_A153 data sets using different enhancement techniques. The SDAMRP-RR is solved with the GCF and 50 evaluation scenarios. Bold entries are the experiments where $Gap_n < 0.01$ is achieved within the maximum runtime.

improving the efficacy of the Benders' optimality cuts.

A different conclusion can be drawn from the results for the F578_A153 data set, providing a richer analysis of the enhancement techniques. It is observed that using parallel computing achieves a greater improvement in the convergence of the Benders' decomposition solution approach compared to the trust region method. Since the F578_A153 data set is much larger than the F267_A49 data set, a greater proportion of the solution runtime per iteration is required for the evaluation stage. Therefore, the benefits from distributing the recovery problems to a number of parallel compute threads will be more pronounced for the F578_A153 data set. It is expected as the problem size and the number of scenarios in the evaluation stage increases, the use of parallel computing will have a greater impact on the convergence of the Benders' decomposition solution process.

5.4 Analysis of planning recoverability

To assess the recoverability improvement given by the SDAMRP-RR, the solutions to both the SDAMRP and SDAMRP-RR are evaluated against a set of disruption scenarios. In the optimisation process of the SDAMRP-RR, the planning stage solution is evaluated against a set of 150 disruption scenarios for both data sets. To provide a fair comparison between the solutions of the SDAMRP and SDAMRP-RR, an alternative set of 300 disruptions scenarios are generated for the recoverability evaluation. Since this alternative set of scenarios is different to the 150 used in the evaluation stage of the SDAMRP-RR, any advantage to the recoverable robust solution in this analysis is removed.

The recoverability of the solutions to the SDAMRP and SDAMRP-RR is given by the average cost of recovery over the set of 300 evaluation scenarios. The solution to each model, given by $\hat{\mathbf{y}}$, is used as an input to a recovery problem defined by the PBSP- s , and the cost of recovery for scenario s is given by the objective function value $\mu_s(\hat{\mathbf{y}})$. Therefore, we define the recoverability of the aircraft routing solution from the SDAMRP as $Rec = \sum_{s \in S} w^s \mu_s(\hat{\mathbf{y}})$. We define the recoverability achieved using the solution to the SDAMRP-RR in the same manner, which is labelled as Rec_{RR} .

There are many features and parameters of the SDAMRP-RR that affect the efficacy of the recoverable robustness technique and solution runtimes. The key features affecting the possible recoverability improvement are i) the value of the swap cost g^{SW} , ii) the value of the loss rate g^{LR} , and iii) the number of scenarios used in the evaluation stage. Both the swap cost g^{SW} and the loss rate g^{LR} are closely related with a trade-off between the number of aircraft swaps and flight cancellations observed in the recovery solutions. As such, the results for the SDAMRP-RR using different values for g^{SW} and g^{LR} will be presented in Section 5.4.1. Solving the recovery problem is a time consuming feature of the recoverable robustness approach, consequently the number of scenarios that are used in the evaluation stage of the SDAMRP-RR has a great effect on the solution runtimes. The result of increasing the number of scenarios will be presented in Section 5.4.2. The following sections analyse the effects that each of these features has on the number of maintenance misalignments, recoverability improvement and

solution runtimes.

5.4.1 Varying the swap cost and loss rate parameters

The swap cost and loss rate parameters are airline specific values that greatly affect recoverability improvement and solution runtimes of the SDAMRP-RR. Since the SDAMRP-RR is solved as a planning problem the runtime required to find the optimal solution is not of critical importance. However it is possible to reduce the solution runtimes for the SDAMRP-RR through the intelligent selection of these parameters. The values of the swap cost and loss rate parameters also affect the flexibility of the recovery solution, which directly impacts on the potential recoverability improvement.

The improvement in the recoverability of the SDAMRP from applying recoverable robustness is given by the relative difference between Rec and Rec_{RR} . This is calculated by $Improve = |Rec - Rec_{RR}| / \min\{Rec, Rec_{RR}\}$ and is given as a percentage in Figure 1. The results presented in this figure demonstrate the recoverability improvement achieved using different swap costs, loss rates and connection cost functions. It is clear from this figure that the greatest improvement in recoverability is achieved with a swap cost $g^{SW} = 7500$ regardless of the value of g^{LR} with only one exception. This is a feature of the recoverable robustness technique since the solution to the SDAMRP-RR attempts to minimise the number of changes from the planned solution. As the value of g^{SW} increases, the effect of this parameter forces the solution to the SDAMRP-RR to require fewer changes during recovery. Therefore, the solution to the SDAMRP will require comparatively more aircraft swaps, resulting in a higher recovery cost.

The flexibility in the recovery problem has a significant effect on the efficacy of the recoverable

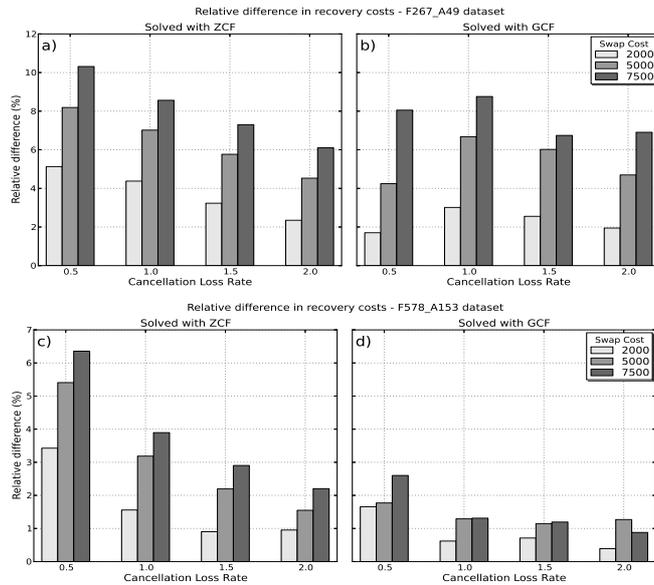


Figure 1: Relative improvement in the recoverability of the SDAMRP by applying recoverable robustness. 150 scenarios are used in the evaluation stage of the SDAMRP-RR.

robustness technique. The values of g^{LR} have a direct correlation with the flexibility of the recovery problem, whereby larger values limit the number of flight cancellations that can be made. Figure 1 demonstrates that as the value of g^{LR} increases, the reduced flexibility in the recovery problem negatively affects the improvement in recoverability achieved by the SDAMRP-RR. This effect of reduced flexibility is observed for both of the data sets and connection cost functions.

Flexibility in the planning stage is also observed to affect the potential recoverability improvement of the SDAMRP-RR. In Figure 1, a greater recoverability improvement is achieved for both data sets using the ZCF. Since the ZCF does not impose a cost on connections in the network, the construction of one-day routes is only affected by the maintenance misalignment constraints and the added Benders' cuts. By contrast, the GCF favours the use of specific connections in the construction of one-day routes, reducing the efficacy of the Benders' cuts. Therefore, the Benders' cuts have a greater impact on the one-day route construction when the ZCF is used compared to the GCF, hence permitting a greater recoverability improvement for the SDAMRP-RR.

The runtimes for each of the cases presented in Figure 1 is given in Figure 2. A maximum runtime for the SDAMRP-RR is set at 10 hours (36000 seconds), which is exceeded for most of the cases presented. When the maximum runtime is exceeded, the method used to find the best integer solution to the SDAMRP-RR depends on the current phase of the solution process. If the algorithm terminates during Phase 1, the solution to the SDAMRP-RR is found by solving the BMP to integral optimality with the current set of generated cuts. Otherwise, if the algorithm terminates during Phase 2, the solution to the SDAMRP-RR is given by the best solution to the BMP found during the solution process.

The vast majority of experiments presented in Figure 2 exceed the maximum allowable runtime. This result alludes to the complexity of the SDAMRP-RR and the obvious necessity to implement

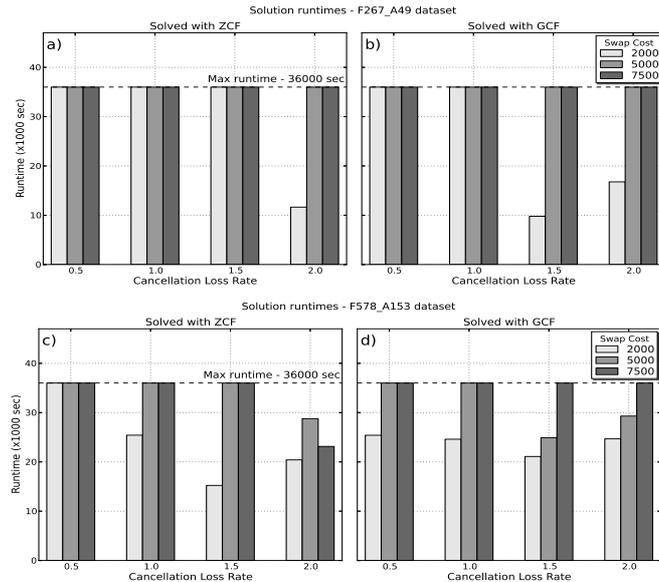


Figure 2: Runtime of the SDAMRP-RR. 150 scenarios in the evaluation stage of the SDAMRP-RR.

enhancement techniques to improve the convergence of the Benders' decomposition solution process. Reviewing the results in Section 5.3 it is clear that enhancement techniques are required to identify the optimal solution even when only 50 scenarios are used in the evaluation stage. Thus, the results presented in Figure 2 are not surprising. In addition, the results from Figures 1 and 2 demonstrate that even suboptimal solutions to the SDAMRP-RR achieves a significant improvement in recoverability.

As a final note, the runtimes for the SDAMRP-RR solved with the ZCF, Figures 2a and 2c, are generally greater than when the GCF is used, Figures 2b and 2d. The use of the ZCF in the planning stage of the SDAMRP-RR increases the degeneracy of the problem which greatly affects the convergence of the Benders' decomposition solution process. One cause of degeneracy is symmetry in the BMP, which requires a greater number of cuts from the PBSP-*s* to eliminate identical solutions. The GCF helps to eliminate the symmetry in the BMP through the cost of one-day routes, also reducing the number of optimal solutions. While the runtimes required for the SDAMRP-RR using the ZCF are much greater than when the GCF is used, Figure 1 demonstrates that the former achieves a better recoverability improvement. This presents a trade-off between solution runtime and quality, arising from the flexibility of the BMP.

5.4.2 Number of evaluation scenarios

This section examines the effect that increasing the number of scenarios has on the improvement in recoverability. In Section 5.4.1, the solution to the SDAMRP-RR is solved with 150 scenarios in the evaluation stage. To further analyse the recoverable robustness technique, experiments are performed by solving the SDAMRP-RR with 250, 350 and 500 scenarios in the evaluation stage. The set of scenarios in each of these experiments are generated using the process described in Section 5.1.2. Similar to the experiments in Section 5.4.1, the solutions to the SDAMRP and SDAMRP-RR are evaluated against a set of 300 disruption scenarios.

Two different results are observed from the different data sets as the number of evaluation scenarios increases. Firstly, Figure 3a demonstrates that the number of scenarios used in the evaluation stage

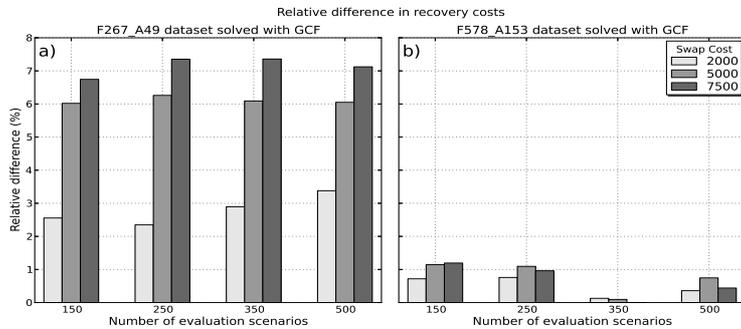


Figure 3: Relative improvement in the recoverability of the SDAMRP by applying recoverable robustness with different sets of scenarios in the evaluation stage of the SDAMRP-RR . Loss rate is set to $g^{LR} = 1.5$.

	Master Iterations				Optimality Gap			
Swap Penalty	Evaluation Scenarios				Evaluation Scenarios			
	150	250	350	500	150	250	350	500
2000	7.0	6.0	2.0	5.0	0.96%	0.47%	26.92%	1.71%
5000	9.0	9.0	2.0	4.0	0.92%	0.81%	30.3%	6.64%
7500	9.0	7.0	2.0	2.0	1.25%	1.42%	25.32%	35.27%

Table 4: Number of master iterations and optimality gap of the SDAMRP-RR solved for the F578_A153 data set with different sets of evaluation scenarios.

of the SDAMRP-RR does not greatly affect the recoverability improvement for the F267_A49 data set. In particular, the variation in the recoverability improvement caused by an increase in the number of scenarios is much less than that achieved using different swap penalty values. This indicates that it is more effective to intelligently select the values for the parameters g^{SW} and g^{LR} than to increase the number of scenarios.

For the F578_A153 data set, a much greater variation in the recoverability improvement as the number of scenarios increases is observed in Figure 3b. In particular, the increase in the number of evaluation scenarios causes the recoverability improvement using the F578_A153 data set to decrease to very low levels. For example, using 500 scenarios in the evaluation stage with $g^{SW} = 2000$, the SDAMRP-RR is solved to find an aircraft routing solution that has a 0.36% improvement in recoverability. In addition, using 350 evaluation scenarios the recoverability improvement is almost negligible. Given that the runtimes required to achieve this improvement have been capped at 10 hours, this result is not encouraging for practical application. It is clear for the F578_A153 data set that the selection of the g^{SW} and g^{LR} parameter values is much more effective in improving the solution quality and runtimes compared to increasing the number of evaluation scenarios.

The poor recoverability improvement for the F578_A153 data set is the result of the large runtimes required in the evaluation stage of the SDAMRP-RR. The maximum runtime is exceeded for all but four of the experiments presented in Table 4, which is indicated by optimality gaps greater than 1%. In Figure 3b, the worst recoverability improvement is observed when 350 scenarios are used in the evaluation stage. By comparing Figure 3b with Table 4, it is clear that these experiments also display the worst runtime performance. Only 2 master problem iterations are performed for the experiments using each of the swap penalties examined. This demonstrates the runtime spent solving the evaluation scenarios have a significant impact on the recoverability improvement achieved by the SDAMRP-RR. Also, large optimality gaps are observed for all experiments using 350 scenarios in the evaluation stage, explaining the poor recoverability improvement presented in Figure 3b.

5.5 Maintenance misalignment in the SDAMRP-RR

The results in Section 5.2 demonstrate that the SDAMRP significantly reduces the number of maintenance misalignments compared to a standard aircraft routing solution. While the application of recoverable robustness to the SDAMRP attempts to improve the recoverability of the planned solution, this should not come at the expense of maintenance planning. In stochastic programming problems, the deterministic solution provides a lower bound on the optimal stochastic solution, indicating that there is a cost associated with considering uncertain information. This is observed when comparing the solutions to the SDAMRP and SDAMRP-RR, where a lower bound on the number of maintenance misalignments is given by the solution to the SDAMRP. Figure 4 presents a comparison of the number of maintenance misalignments for the SDAMRP, SDAMRP-RR and the standard aircraft routing problem. The *dashed* line represents the SDAMRP solution, which is the best possible solution, and the *dot-dashed* line is the solution to the standard aircraft routing problem.

The results in Figure 4 for the SDAMRP-RR are all within the gap between the SDAMRP and aircraft routing solution, with most solutions close to the lower bound. While the SDAMRP-RR is a conservative planning approach, the results demonstrate that this solution does not greatly affect the reduction in maintenance misalignments achieved by the SDAMRP. It is interesting to note that the number of maintenance misalignments for the SDAMRP-RR when solved with the ZCF is equal to the solution of the SDAMRP for a large number of cases. This can be explained by the increased flexibility that is provided by not artificially imposing a cost on the connection lengths. Since the aircraft routing problem is highly degenerate, there are many solutions that can achieve the same maintenance

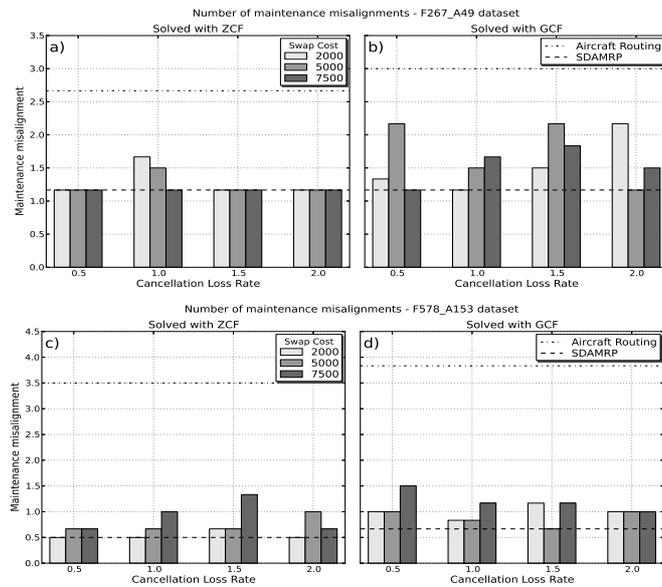


Figure 4: Assessing the impact of recoverable robustness on maintenance misalignment. 150 scenarios in the evaluation stage of the SDAMRP-RR.

misalignment result. In this case, the application of recoverable robustness to the SDAMRP helps to select the solution from this set that provides the best recoverability.

It can be seen in Figure 4 that the smallest deterioration in the maintenance misalignments is achieved by solving the SDAMRP-RR with the F578_A153 data set. The most interesting observation is that solving the SDAMRP-RR for the F578_A153 data set displays the least variation from the SDAMRP solution across all parameter values. This is an important result since it demonstrates that the application of recoverable robustness to larger data sets does not greatly affect the maintenance planning of the SDAMRP. It can be concluded that the solution to the SDAMRP-RR provides an improvement in recoverability while still achieving a significant reduction in the number of maintenance misalignments.

6 Conclusions

This paper introduces two models that reduce the susceptibility of maintenance planning solution to schedule perturbations. The SDAMRP is formulated using a novel modelling approach to count the number of maintenance misalignments in the aircraft routing solution. This model is extended with the SDAMRP-RR by applying the recoverable robustness technique, improving the recoverability of the maintenance planning solution. Both modelling approaches are contributions to the airline planning literature.

The two models are evaluated using a variety of flight schedules to demonstrate the versatility of the proposed approaches. The SDAMRP is demonstrated to significantly reduce the number of maintenance misalignments, with a minimum improvement of 56.25%. Additionally, the application of recoverable robustness achieves a recoverability improvement without any great impact to the maintenance planning solution. The SDAMRP-RR presents a large-scale optimisation problem that is solved by integrating Benders' decomposition and column generation. A number of enhancement techniques are also implemented with the results demonstrating significantly improved convergence of the solution process by employing parallel computing techniques.

Potential future research for this problem is the extension of the single day maintenance routing problem to construct maintenance plans for longer time frames or consider other maintenance check types. Additionally, the recoverable robust framework investigated in this paper is greatly affected by the problem size and number of evaluation scenarios. As such, a direction for future research is to identify techniques to reduce the runtimes required in the evaluation stage of the recoverable robustness framework. Fast solution runtimes for the evaluation stage will directly affect applicability of the recoverable robustness technique to a wider range of large-scale optimisation problems.

References

- [1] Y. Ageeva. Approaches to incorporating robustness into airline scheduling. Master's thesis, Massachusetts Institute of Technology, August 2000.
- [2] R. Ahuja, T. Magnanti, and J. Orlin. *Network flows: theory, algorithms, and applications*. Prentice Hall, 1993.
- [3] C. Barnhart, N. L. Boland, L. W. Clarke, E. L. Johnson, G. L. Nemhauser, and R. G. Shenoi. Flight string models for aircraft fleetling and routing. *Transportation Science*, 32(3):208–220, 1998.
- [4] J. F. Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4(1):238–252, 1962.
- [5] R. Borndörfer, I. Dovica, I. Nowak, and T. Schickinger. Robust tail assignment. In *Proceedings of the fiftieth annual symposium of AGIFORS*, 2010.
- [6] L. Clarke, E. Johnson, G. Nemhauser, and Z. Zhu. The aircraft rotation problem. *Annals of Operations Research*, 69:33–46, 1997.
- [7] J.-F. Cordeau, G. Stojković, F. Soumis, and J. Desrosiers. Benders' decomposition for simultaneous aircraft routing and crew scheduling. *Transportation Science*, 35(4):375–388, 2001.
- [8] G. Côté and M. A. Laughton. Large-scale mixed integer programming: Benders'-type heuristics. *European Journal of Operational Research*, 16(3):327–333, 1984.
- [9] N. Eggenberg. *Combining Robustness and Recovery for Airline Schedules*. PhD thesis, École Polytechnique Fédérale De Lausanne, December 2009.
- [10] T. Feo and J. Bard. Flight scheduling and maintenance base planning. *Management Science*, 35(12):1415–1432, 1989.
- [11] M. Fischetti, D. Salvagnin, and A. Zanette. A note on the selection of Benders' cuts. *Mathematical Programming*, 124(1-2):175–182, 2010.
- [12] G. Froyland, S. J. Maher, and C.-L. Wu. The recoverable robust tail assignment problem. *Transportation Science*. To appear.
- [13] A. M. Geoffrion and G. W. Graves. Multicommodity distribution system design by Benders' decomposition. *Management Science*, 20(5):822–844, 1974.
- [14] R. Gopalan and K. T. Talluri. The aircraft maintenance routing problem. *Operations Research*, 46(2):260–271, 1998.

- [15] M. Grönkvist. *The Tail Assignment Problem*. PhD thesis, Chalmers University of Technology and Göteborg University, August 2005.
- [16] A. Heinhold. Maintenance reachability at Southwest Airlines. In *Proceedings of the AGIFORS Scheduling and Strategic Planning Study Group Meeting, Dallas*, 2008.
- [17] E. Lacasse-Guay, G. Desaulniers, and F. Soumis. Aircraft routing under different business processes. *Journal of Air Transport Management*, 16(5):258–263, 2010.
- [18] S. Lan, J.-P. Clarke, and C. Barnhart. Planning for robust airline operations: optimizing aircraft routings and flight departure times to minimize passenger disruptions. *Transportation Science*, 40(1):15–28, 2006.
- [19] M. Lapp and A. Cohn. Modifying lines-of-flight in the planning process for improved maintenance robustness. *Computers & Operations Research*, 39(9):2051–2062, 2012.
- [20] C. Liebchen, M. Lübbecke, R. Möhring, and S. Stiller. The concept of recoverable robustness, linear programming recovery, and railway applications. In R. Ahuja, R. Möhring, and C. Zaroliagis, editors, *Robust and Online Large-Scale Optimization*, volume 5868 of *Lecture Notes in Computer Science*, pages 1–27. Springer Berlin Heidelberg, 2009.
- [21] J. Linderoth and S. Wright. Decomposition algorithms for stochastic programming on a computational grid. *Computational Optimization and Applications*, 24(2-3):207–250, 2003.
- [22] T. Magnanti and R. Wong. Accelerating Benders’ decomposition: algorithmic enhancement and model selection criteria. *Operations Research*, 29(3):464–484, 1981.
- [23] A. Mercier, J. Cordeau, and F. Soumis. A computational study of Benders’ decomposition for the integrated aircraft routing and crew scheduling problem. *Computers & Operations Research*, 32(6):1451–1476, 2005.
- [24] A. Mercier and F. Soumis. An integrated aircraft routing, crew scheduling and flight retiming model. *Computers & Operations Research*, 34(8):2251–2265, 2007.
- [25] N. Papadakos. Practical enhancements to the Magnanti-Wong method. *Operations Research Letters*, 36(4):444–449, 2008.
- [26] N. Papadakos. Integrated airline scheduling. *Computers & Operations Research*, 36(1):176–195, 2009.
- [27] W. Rei, J.-F. Cordeau, M. Gendreau, and P. Soriano. Accelerating Benders’ decomposition by local branching. *INFORMS Journal on Computing*, 21(2):333–345, 2009.

- [28] J. M. Rosenberger, E. L. Johnson, and G. L. Nemhauser. A robust fleet-assignment model with hub isolation and short cycles. *Transportation Science*, 38(3):357–368, 2004.
- [29] A. Ruszczyński. A regularized decomposition method for minimizing a sum of polyhedral functions. *Mathematical Programming*, 35(3):309–333, 1986.
- [30] T. Santoso, S. Ahmed, M. Goetschalckx, and A. Shapiro. A stochastic programming approach for supply chain network design under uncertainty. *European Journal of Operational Research*, 167(1):96–115, 2005.
- [31] C. Sriram and A. Haghani. An optimization model for aircraft maintenance scheduling and re-assignment. *Transportation Research Part A: Policy and Practice*, 37(1):29–48, 2003.
- [32] B. G. Thengvall, J. F. Bard, and G. Yu. Balancing user preferences for aircraft schedule recovery during irregular operations. *IIE Transactions*, 32(3):181–193, 2000.