Scheduling a Dynamic Aircraft Repair Shop

Maliheh Aramon Bajestani and J. Christopher Beck

Department of Mechanical & Industrial Engineering University of Toronto Toronto, Ontario M5S 3G8, Canada {maramon, jcb}@mie.utoronto.ca

Abstract

We study a dynamic repair shop scheduling problem in the context of military aircraft fleet management. A number of flights, each with a requirement for a specific number and type of aircraft, are already scheduled over a long horizon. The overall scheduling horizon is considered to be composed of multiple time periods. The goal is to maintain a full complement of aircraft. We need to assign aircraft to flights and schedule repair activities while considering the repair capacity limit and aircraft failures. The number of aircraft awaiting repair dynamically changes over time because of aircraft failures. We designed three re-scheduling policies using different optimization and heuristic techniques to solve the dynamic problem over successive time periods. Experimental results demonstrate that the wave coverage is higher if the optimization technique, logic-based Benders decomposition, is used to solve the problem over longer time periods more frequently.

Introduction

In industries using expensive machinery, it is common to repair rather than replace a machine when it breaks down. For example, it is far too expensive for a railroad or airline company to keep stock on-hand to replace failed machines. The need for repair, however, generates a set of new decisions: "How many repair resources (e.g., repairpersons) should be allocated?", "Where should repairs take place?", and "When should they be done and using which resources?". In this paper, we study an aircraft repair shop. When aircraft fail, the management process must dynamically react to failures by scheduling and re-scheduling repair activities to maximize aircraft availability. A high-quality schedule capable of dealing with uncertainty and adjusting to unexpected events leads to an efficient repair operation.

Motivated by the case study in Safaei et al. (Safaei, Banjevic, and Jardine 2010), we address the problem of aircraft fleet management where a number of flights are planned over a long horizon consisting of several time periods. Every flight, called "wave", has a requirement for a specific number of aircraft of different types. Aircraft flow over a long horizon is illustrated in Figure 1.

The goal is to construct a repair schedule that will maximize wave coverage while allowing for aircraft failures in systematic pre- and post-flight checks. Each aircraft failure requires a set of repair activities with known processing



Figure 1: A flow chart representing aircraft flow among waves, checks, and the repair shop over a long horizon.

times and resource requirements to be scheduled using resources with a limited capacity.

Simply stated, we view the dynamic problem as successive scheduling problems over time periods. We use three scheduling techniques including logic-based Benders decomposition (LBBD), mixed-integer programming (MIP), and a dispatching heuristic developed in our previous work (Aramon Bajestani and Beck 2011) to solve the scheduling problem in each time period. To revise the schedule, we design three different policies based on when and how the re-scheduling is done.

Empirical studies indicate that solving the scheduling problems more frequently over longer time periods using LBBD results in the best performance and that solving the scheduling problems more frequently is more important than solving them over longer time periods.

The main contributions of our paper are:

- generalizing an offline scheduling problem studied in (Aramon Bajestani and Beck 2011) to its dynamic counterpart;
- demonstrating how to adapt existing solution techniques to a dynamic problem;
- empirically analyzing the impact of applying different rescheduling strategies to determine when and how to respond to real-time events.

The following section defines our problem, reviews the scheduling algorithms used for each time period and discusses the literature on scheduling a repair system. We go on to describe the proposed re-scheduling strategies, present our experiments and results, suggest possible directions for future work, and provide a conclusion.

Background

In this section, we present the formal definition of the problem, review the solution approaches for scheduling the repair shop, and discuss the literature on scheduling a repair system.

Problem Definition

The problem at time 0 is shown schematically in Figure 2. The circles represent aircraft. It is assumed that the total number of aircraft is constant over a long horizon. In our example, at the beginning, three aircraft are ready for the pre-flight check; others are in the repair shop awaiting repair before they can proceed to the pre-flight check. A number of waves (five are shown) and their corresponding pre- and post-flight checks are already scheduled over a long horizon.



Figure 2: Snapshot of the problem at time 0 over a long horizon.

At the beginning, we schedule the repair activities over a time period, for example the interval containing the first three waves. The schedule may not be executed as is. We dynamically react to actual aircraft failures by re-scheduling the repair activities after, for example, each wave.

The goal in each time period is to assign aircraft to waves to maximize coverage while respecting constraints on maintenance capacity. The coverage is the extent to which the aircraft requirements of the waves are met. The scheduling problem is under the constraints that the repair shop has limited capacity and the aircraft are subject to breakdown which can be detected in pre- or post-flight checks. We assume that once an aircraft fails, it goes to the repair shop and waits until its repair operations are performed.

We use the following notation to represent the problem.

- N is the set of aircraft. λ_n is the failure rate of the aircraft $n \in N$.
- K is the set of aircraft types. For each aircraft type k ∈ K, there are A_k aircraft ready (i.e., not in the repair shop) at the beginning of the first time period. λ

 k is the mean failure rate over all aircraft of type k.
- *R* is the set of repair resources (called "trades"). The maximum capacity of trade *r* ∈ *R* is *C_r*.
- W is the set of waves and D is an ordered set of due dates in the time period. D consists of the wave start-times plus a big value, B sorted in ascending order. Each wave, $w \in$ W has a start-time, $st_w \in D$, and an end-time et_w . Each wave requires a_{kw} aircraft of type k.
- J is the set of jobs in the time period. Each job is associated with a specific aircraft, and I_k denotes the set of repair jobs for aircraft type k. M_r is the set of jobs requiring trade r. Each job might require more than one trade to be completed. The processing time of job j on trade r is p_{jr} and c_{jr} is the capacity of trade r required by job j.

To understand how the dynamic problem over the long horizon can be viewed as scheduling problems over successive time periods, assume that we start repairing the failed aircraft and assigning them to the waves based on the computed schedule at time 0. A wave might start while a repair is under way in the repair shop. If some aircraft fails the pre-flight check, it goes to the repair shop. Each failed aircraft requires a set of repair activities with known processing times and resource requirements. At the repair shop, some of the previously failed aircraft might be already repaired, some might be under repair, and others might be awaiting repair. Once the failed aircraft enter the repair shop, we have a new repair scheduling problem with a new set of jobs, including the recently failed aircraft and the previously failed aircraft whose repairs are still under way or are not yet started. The new problem has an added constraint, namely that the repairs currently under way cannot be disrupted.

Scheduling for a Time Period

This section reviews the details of mixed-integer programming, logic-based Benders decomposition, and the dispatching heuristic presented in (Aramon Bajestani and Beck 2011).

Mixed Integer Programming The variables are defined in Table 1 and the model is shown in Figure 3.

Var.	Definition
Z_{kw}	The number of aircraft of type k assigned to fly in wave w
x_{ij}	$x_{ij} = 1$ if the <i>i</i> th due date is assigned to job <i>j</i>
st_{jr}	The start-time of job j on trade r
U_{kw}	The number of aircraft of type k whose repair due date is st_w
E_{kw}	The expected number of available aircraft type k for wave w
et_{jr}	The end-time of job j on trade r

Table 1: The decision variables (top) and inferred variables (bottom) for the MIP model.

The objective function (1) maximizes the number of aircraft assigned to a wave subject to a limit on the number of aircraft required and the expected number available (Constraint (5)). Equation (2) calculates the number of aircraft of type k whose repair due date is st_w . Equation (3) calculates the expected number of available aircraft for the first wave. Equation (4) calculates the expectation of availability for the other waves. The first term includes those aircraft available but not used for the previous wave and those newly arrived from the repair shop. The second term includes all aircraft that are available because they have completed waves since the previous wave started. ξ_k^{pre} and ξ_k^{post} denote the probability of failure associated with aircraft type k in pre- and post-flight checks, respectively: $\xi_k^{pre} = (1 - e^{-\alpha \bar{\lambda}_k})$ and $\xi_k^{post} = (1 - e^{-\beta \bar{\lambda}_k})$, where $\alpha < \beta$ to reflect deterioration of the aircraft type k area $\lambda < \beta$. of the aircraft through use. As tracking the history of the aircraft is prohibitive to find the actual probability of failure for each, they are distinguished based on their type and the failure rate of each aircraft type $\bar{\lambda}_k$ is used to estimate the probability of failure. Constraint (6) ensures that exactly one due date is assigned to each job. Equation (7) calculates

$$\max.\sum_{w=1}^{W}\sum_{k=1}^{K}Z_{kw} \tag{1}$$

s.t.
$$U_{kw} = \sum_{j \in I_k} x_{ij}$$
, if $d_i = st_w$ (2)

$$E_{k1} = (A_k + U_{k1})(1 - \xi_k^{pre}), \ \forall k$$

$$E_{kw} = (E_{k(w-1)} - Z_{k(w-1)} + U_{kw})(1 - \xi_k^{pre}) +$$
(3)

$$\sum_{v=1}^{w-1} Z_{kv} (1 - \xi_k^{post}) (1 - \xi_k^{pre}),$$

if $st_{w-1} < et_v \le st_w, \forall w \ne 1, k$ (4)
 $Z_{kw} \le \min(E_{kw}, a_{kw}), \forall k, w$ (5)

$$\sum_{i=1}^{|D|} x_{ij} = 1, \ \forall j \tag{6}$$

$$st_{jr} + p_{jr} = et_{jr}, \forall j, r$$

$$et_{jr} \le \sum_{i=1}^{|D|} x_{ij} d_i, \ \forall j, r \tag{8}$$

(7)

$$\sum_{j \in M_r} c_{jr} \le C_r, \text{if } st_{jr} \le t < et_{jr}, \ \forall t, r$$
(9)

$$x_{ij} \in \{0, 1\}, \ \forall i, j$$
 (10)

$$0 \le E_{kw} \le |N|, \ \forall k, w \tag{11}$$

$$st_{jr}, et_{jr} \in \mathbb{Z}^+ \cup \{0\}, \ \forall j, r$$

$$(12)$$

$$Z_{kw} \in \mathbb{Z}^+ \cup \{0\}, Z_{kw} \le |N|, \ \forall k, w \tag{13}$$

Figure 3: The global MIP model for one time period.

the end-time of the jobs. The end-time of each job is guaranteed to be less than or equal to the assigned due date by constraint (8). Constraint (9) enforces the capacity limit of each trade, where t denotes the discrete time during which job j is under way.

Logic-based Benders Decomposition A logic-based Benders decomposition (LBBD) method can be formulated where the master problem assigns aircraft to waves to maximize wave coverage over the current time period and the sub-problems create the maintenance schedules given the due dates assigned by the master problem solution. The master problem is solved using MIP, while constraint programming (CP) is used for the scheduling sub-problems.

The Due-Date Assignment Master Problem (DAMP): MIP Model To formulate the master problem as a MIP model, we use a binary variable x_{ij} for each $j \in J$ and $i \in D$ with the same meaning as in the global MIP model. A MIP formulation of DAMP is as follows: max. Objective (1)

s.t. Constraints (2) to (6), (10), (11), (13)

$$\sum_{j \in M_r, \sum_{i=1}^{|D|} x_{ij} d_i \le st_w} c_{jr} p_{jr} \le st_w C_r, \quad \forall r, w$$
(14)

MIP cuts (15)

The master problem incorporates a number of the constraints in the global MIP model. It does not represent the start-times of jobs nor does it fully represent the capacity of the trades. As is common in Benders decomposition, the master problem includes a relaxation of the sub-problems (Constraints (14)) and Benders cuts (Constraints (15)).

The Sub-problem Relaxation Constraint (14) is the relaxation of the capacity of a trade, expressing a limit on the area of jobs that can be executed. The limit is defined using the area bounded by the capacity of the trade and the time intervals $[0, st_w]$ for each wave w, plus [0, B] where Bis the maximum due date assigned to the jobs on the trade. The area of each interval must be greater than or equal to the sum of the areas of the jobs that finish by the end of the interval.

The Benders Cuts We demonstrate the intuition with an example before defining the cut. Assume that for a given trade with five jobs and a due date set, $D = \{14, 17, 20, 100\}$, the current master solution is: $x_{21} = 1, x_{12} = 1, x_{43} = 1, x_{14} = 1$, and $x_{15} = 1$. Job 1 is assigned to the second due date, 17, Job 2 has the first due date, 14, and so on. If the current solution is infeasible due to the resource capacity of the trade, then we know that at least one of the jobs must have a later due date. We can, therefore, constrain the sum of the consecutive x_{ij} up to and including the ones assigned to 1 to be one less than the number of jobs. In our example, the cut would be:

$$(x_{11} + x_{21}) + (x_{12}) + (x_{13} + x_{23} + x_{33} + x_{43}) + (x_{14}) + (x_{15}) \le 5 - 1$$

Formally, assume that in iteration h, the solution of the DAMP assigns a set, Q, of due dates to the jobs on trade r. Assume further that there is no feasible solution on trade r with the assignments in Q.

The cut after iteration h is:

$$\sum_{j \in M_r} \sum_{i \in I_{jh}^r} x_{ij} \le |M_r| - 1, \quad \forall r$$
(16)

where $I_{jh}^r = \{i' | i' \leq i, \text{ and } x_{ij}^h = 1\}$ is the set of due dates indices less than or equal to the due date index assigned to job j and $|M_r|$ is the number of jobs on trade r.

Job Scheduling Sub-problem Given a set of due dates assigned to the jobs on a trade, the goal of the job scheduling sub-problem (JSSP) is to assign start-times to the jobs to satisfy the due dates and the trade capacity. The JSSP for each trade can be modeled using cumulative constraints (Hooker 2005). We use a CP formulation:

Cumulative(
$$[t_j | d_j^h], [p_{jr} | d_j^h], [c_{jr} | d_j^h], C_r), \forall r$$

 $0 \le t_j \le d_j^h - p_{jr}, \forall j, r$ (17)

where t is an array of variables such that t_j is the start-time of job j, d is an array of values such that d_j^h is the due date assigned to job j in master problem in iteration h. The variables p_{jr}, c_{jr}, C_r are as defined above. Constraint (17) enforces the time windows: the job cannot be started later than $d_i^h - p_{jr}$.

A Dispatching Heuristic The dispatching heuristic, inspired by the Apparent Tardiness Cost (ATC) heuristic (Pinedo 2005), is a list-scheduling heuristic. It prioritizes repair activities based on how early the corresponding aircraft type is needed, the processing time of each job, and the relative type demand. The ranking index we use is as follows:

$$I_j = ST(k_j) \exp(-\frac{FN_j}{FC_j}), \quad \forall j$$

If we let k_j denote the type of aircraft j, then $ST(k_j)$ is the start-time of the first wave that requires an aircraft of type k_j . FN_j is the fraction of the total number of aircraft of type k_j required by the first wave that requires k_j , and FC_j is the maximum proportion of the capacity needed by job j over all its required trades, as follows.

$$FC_j = \max_r \left(\frac{p_{jr}c_{jr}}{ST(k_j)C_r}\right), \quad \forall r$$

The heuristic sorts the jobs in ascending order of the index and then iterates through the jobs, scheduling each job at its earliest available time.

Literature Review

Queuing theory is often used to model repair systems [(Iravani, Krishnamurthy, and Chao 2007) and the references therein]. Queuing theory has a long-term definition of optimality resulting in some sort of repair policy commonly assuming that the repair resources are unary capacity (i.e., one repair is carried out at a time). A repair policy determines the order under which the repair activities should be carried out.

Queuing theory does not model the combinatorics of the scheduling problem. In our problem, the optimization of scheduling performance at discrete time points (i.e., before each flight) is of interest, and the repair resources have a discrete capacity. Therefore, we believe that better performance can be achieved by dealing directly with the combinatorics and explicitly scheduling the repair shop to meet the waves.

Dynamic scheduling is well-suited to handle the uncertain and combinatorial structure of the scheduling problem. Dynamic scheduling concerns the allocation of resources to activities over time when the real-time events occur during the execution of previously determined schedule (Aytug et al. 2005).

The real-time event studied in this paper is a job-related event (Vieira, Hermann, and Lin 2003) because of the uncertainty involved in the systematic pre- and post-flight checks. Some of the repaired aircraft cannot accomplish their assigned flight: they are diagnosed as failed and must return to the repair shop.

When and how to respond to the real-time events are two independent variables in dynamic scheduling problems addressed in (Vieira, Hermann, and Lin 2003; Aytug et al. 2005; Bidot et al. 2009). In our problem, the length of each time period, and the frequency of re-scheduling determine how and when we react to the aircraft failures, respectively.

Re-scheduling Strategies

Our strategies have three main parts: scheduling the repair activities, observing the aircraft failures while executing the computed schedule, and responding to dynamic events by re-scheduling the repair activities. They start by scheduling the repair activities over one time period at time 0. The length of time period defines the scheduling horizon over which the repair activities are scheduled. The re-scheduling strategies start executing the repair schedule while observing the aircraft failures. The frequency of re-scheduling determines when our strategies dynamically respond to the aircraft failures.

We use the techniques reviewed in the Background section to schedule the repair activities over one time period. Three different policies denoted as P_{ij} are designed in which *i* and *j* define the length of scheduling horizon and the frequency of re-scheduling in number of waves, respectively.

The three policies discussed here are:

- P_{11} : This policy has a scheduling horizon with a length of one wave and re-schedules after every wave. In Figure 4, we show that P_{11} schedules one wave at a time (i = 1) and re-schedules after each wave (j = 1).
- P₃₁: This policy has a scheduling horizon with a length of three waves and re-schedules after every wave. In contrast to P₁₁, for P₃₁ (Figure 5), the scheduling horizon is three waves but re-scheduling is still done after each wave.
- P₃₃: This policy has a scheduling horizon with a length of three waves and re-schedules after every third wave (Figure 6).



To model the dynamic events, we simulate the aircraft failures in pre- and post-flight checks. Every aircraft either passes or fails each check. If the aircraft fails, a new



set of repair activities with known processing times and resource requirements is added to the repair shop. If the aircraft passes, it flies the wave. To model the aircraft deterioration, we increase the failure rate of the aircraft by γ percent each time it flies a wave. For example, consider that λ_n is the initial failure rate of the aircraft $n \in N$. Its failure rate after flying t waves will be equal to $\lambda_n (1 + \gamma)^t$.

The observed wave coverage for each wave is the number of aircraft flying the wave divided by the number required.

Experimental Results

The next sub-section describes the problem instances and the experimental details. We then compare the impact of using different scheduling techniques and re-scheduling policies on the observed wave coverage.

Experimental Setup

For our problem instances, the number of aircraft, the number of trades, and the total number of waves are set to $\{10, 15, 20, 25, 30\}, \{4\}$, and $\{30\}$ respectively. Each combination has 5 instances for a total of 25 instances. Each instance is simulated 10 times.

Aircraft The number of aircraft types is equal to $\frac{|N|}{5}$, where |N| is the number of aircraft. The aircraft are randomly assigned to different types, and the initial failure rate for each aircraft is randomly chosen from the uniform distribution of [0, 0.5]. The failure rate of an aircraft is increased by $\gamma = 5\%$ each time it is used. The values of α and β are 1 and 3, respectively (see the Mixed Integer Programming section).

Trades The capacity limit for each trade is $C_r = 10$.

Repair Jobs The repair jobs at time 0 each require half the trades, on average. Subsequent repair jobs require all trades. This difference was done to have enough repair jobs for the successive scheduling problems. The capacity of trade r used by job j, c_{jr} , is drawn from [1, 10] while the processing time, p_{jr} , is drawn from [r, 10r]. At time 0, having 80% of the aircraft in the repair shop results in |J| = 0.8|N| repair jobs.

Waves The plane requirement for each wave is randomly generated from the integer uniform distribution $[1, a_k]$ where a_k denotes the number of aircraft of type k. The length of each wave is drawn with uniform probability from [3, 5]. To find an appropriate start time for the first wave (not too early or too late) and subsequent waves, $T = 1.2 \times LB$ is defined where $LB = \max_r(S_r)$. The sum of the processing areas of the jobs in each trade, r, divided by the trade capacity is denoted by S_r . The processing areas in each trade are summed over the jobs in the repair shop at time zero. The start time of each wave is generated as $st_1 = rand[\frac{T}{3}, \frac{T}{2}]$ for the first

wave, and $st_w = et_{w-1} + rand(0, 40)$ for $1 < w \le 30$. As mentioned earlier the total number of waves is 30.

Dynamic events To simulate an aircraft failure, we generate a random value from the uniform distribution [0, 1] for each aircraft at each check. If the random value is less than the aircraft's probability of failure, the aircraft fails; otherwise, it passes. The aircraft's probability of failure in pre- and post-flight checks are calculated using $(1 - e^{-\alpha\lambda_n})$ and $(1 - e^{-\beta\lambda_n})$, respectively. As mentioned earlier, λ_n is the failure rate of aircraft $n \in N$ which increases by $\gamma = 5\%$ each time the aircraft flies a wave. Note that, passing the preflight check of a wave does not necessarily mean that the aircraft flies the wave. If the number of available aircraft are more than the requirements, the aircraft that fly are randomly selected to meet the requirements.

The time-limit to schedule the repair activities in each scheduling horizon is 600 seconds. We execute the best feasible schedule found before the time-limit if MIP times out. In the case that LBBD times out, the schedule created by the dispatching heuristic is executed as LBBD cannot create a feasible schedule when it times-out.

The scheduling uses IBM CPLEX 12.1 and IBM ILOG Solver/Scheduler 6.7, and the simulation is coded in C++.

Computational Results

In this section, we discuss our results to answer a number of different questions.

Question 1 What is the impact of using a complete technique vs. the dispatching heuristic on the mean observed wave coverage?

We expect a complete technique to achieve higher wave coverage because it incorporates known information on uncertainty into scheduling the repair activities, while the dispatching heuristic does not have this property.

Figure 7 shows the mean observed coverage up to wave $w \in \{1, 2, ..., 25\}$ for different scheduling techniques over all three policies. The mean observed coverage up to wave w is $O_w = \frac{\sum_{i=1}^w \nu_i}{w}$, where ν_i denotes the coverage of wave i. As illustrated, LBBD achieves about a 40% higher mean coverage over all waves than either MIP or the dispatching heuristic. The MIP algorithm also takes the probabilistic information into account creating a repair schedule but it times out on 72% of scheduling problems without finding a feasible solution. The dispatching heuristic then is used to create the repair schedule. Therefore, using MIP and the dispatching heuristic results in waves with almost the same coverage as the dispatching heuristic alone.

Table 2 presents further data for all scheduling techniques: the mean observed coverage, the percentage of waves with less than or equal to 0.3 coverage, and the percentage of waves with more than or equal to 0.7 coverage. The data indicate the clear superiority of LBBD over the dispatching heuristic.

Question 2 Does P_{31} policy provide the waves with higher coverage than P_{33} and P_{11} policies using the optimization



Figure 7: Mean observed coverage for different scheduling techniques.

Scheduling	Mean Observed	% Waves with a	% Waves with a
Technique	Coverage	Coverage ≤ 0.3	Coverage ≥ 0.7
LBBD	0.71	6.07	55.80
MIP	0.49	27.31	21.40
Heuristic	0.51	26.39	26.43

Table 2: The mean observed coverage, the percentage of waves with less than or equal to 0.3 coverage, and the percentage of waves with more than or equal to 0.7 coverage up to wave 25 over all the policies.

technique LBBD?

We expect that P_{31} with LBBD will produce better coverage because it schedules over a longer horizon and adjusts the schedule as soon as aircraft failures occur. Although P_{31} with the dispatching heuristic also responds quickly to the aircraft failures, it does not incorporate the length of the scheduling horizon into the ranking index for repair activities and always repair the aircrafts for the earliest future.

Figure 8 shows the mean observed coverage for different policies using LBBD up to wave $w \in \{1, 2, ..., 25\}$. The P_{31} policy leads to consistently higher coverage.

Figure 9 displays the cumulative percentage of the waves with a coverage less than or equal to ω for LBBD and the dispatching heuristic, where ω denotes the values on the xaxis. The best performing approach will have a fewer waves with a low coverage and more waves with a high coverage. Therefore, its curve will be closer to the lower right-hand corner. As illustrated, in LBBD, P_{31} performs better than the two other policies. In contrast, in the dispatching heuristic, P_{31} results in waves with the same coverage as P_{11} .

Question 3 Does P_{11} have more waves with very low coverage than P_{33} using LBBD?

We expect that P_{33} will result in fewer waves with very low coverage beacuse it takes the possibility of a more distant future into account when creating the repair schedule, while P_{11} policy repairs the aircraft at the earliest possible time.

Figure 10 demonstrates that our intuitions are correct that the P_{33} policy has fewer waves with very low coverage than



Figure 8: Mean observed coverage for different policies using LBBD.



Figure 9: The percentage of waves with a coverage less than or equal to ω , where ω denotes the values on the x-axis.

the P_{11} policy using LBBD.

Question 4 Is a quicker reaction to the dynamic events more important than scheduling over a longer horizon?

The P_{31} policy changes the repair schedule after each wave and trades-off the coverage among three consecutive waves by scheduling over a longer horizon. In contrast, the P_{11} policy schedules for one wave and reacts after each wave while the P_{33} policy reasons over a longer term without a quick response to the dynamic events.

As already shown in Figure 8, the P_{31} policy results in a higher mean coverage. The superiority of policy P_{31} indicates that both features of quick response to the dynamic events and long-term reasoning contribute to the overall performance, but the question is which one contributes more.

To answer the question, the waves are partitioned into buckets of size 3. We expected that P_{33} would achieve a higher mean coverage over each bucket than P_{11} because it reasons about the trade-off among the three waves. However, Figure 11 demonstrates that both policies achieve equal performance until wave 15 and the policy P_{11} then does better. This observation indicates that quick reaction to the air-



Figure 10: Percentage of waves with a coverage less than or equal to 0.3 using LBBD.



Figure 11: Mean observed coverage over each wave bucket using LBBD.

craft failures has more significant impact on the observed coverage than long-term reasoning.

To obtain more insight, the mean observed coverage for the first, second, and third waves in each bucket is shown in Figures 12, 13 and 14, respectively. As illustrated, P_{11} results in a higher coverage than P_{33} for the waves scheduled later in each bucket. The mean difference between two policies ($P_{11} - P_{33}$) is equal to -5, 2, and 10 percent for the first, second, and third waves in buckets, respectively. Intuitively, we expect that P_{11} provides the earlier waves in buckets with a higher coverage than P_{33} because the later policy trades off between three waves when assigning the aircraft to the waves. However, the observation contradicts our expectation and is an evidence that the quick reaction is more important than scheduling over longer horizon.

Summary The following conclusions are supported by empirical observations:

- LBBD provides the waves with a higher coverage than the dispatching heuristic.
- The P₃₁ policy results in waves with a higher coverage than P₁₁ policy using LBBD and with the same coverage using the dispatching heuristic.

- The P_{33} policy is shown to result in fewer waves with very low coverage than P_{11} policy using LBBD.
- Scheduling over a longer horizon and quickly adjusting the schedule based on the real events are the features contributing to the increase in the observed coverage. Furthermore, it is shown that the quick reaction to the dynamic events is more important than the long scheduling horizon.

Future Work

Although the experiments show that scheduling over a longer horizon and responding quickly to disruptions using optimization techniques in a dynamic and uncertain environment yield better performance, the generality of this observation remains in question. One promising direction in future work would be to establish a formal framework to determine how long the scheduling horizon should be and how quickly we should respond to real events. Bidot (Bidot 2005) presented the first steps toward a generic theoretical framework in his PhD thesis.

In this paper, to measure system performance, we have focused on the value of mean observed coverage, not taking into account the computational cost of applying different policies. Although P_{33} is shown to be the dominant policy, it has a greater computational cost and potentially increases the unnecessary changes in the schedule (Aytug et al. 2005). Optimizing over a longer scheduling horizon and rescheduling repair activities once failed aircraft enter the repair shop explain the high computational cost and the increased unnecessary changes, respectively. How to quantify costs in order to evaluate different policies is another interesting direction for future work.

Conclusion

In this paper, we address a dynamic aircraft scheduling problem in a repair system. The goal is to meet the aircraft requirements for each wave by assigning the failed aircraft to the flights considering the maintenance capacity and the aircraft failures. The number of failed aircraft dynamically changes because of aircraft breakdowns. Our proposed solution approaches solves the dynamic problem as successive scheduling problems over multiple time periods. We use three different scheduling techniques developed in our previous work and three re-scheduling policies to schedule the repair activities on-line with dynamic reaction to the aircraft failures. The length of the scheduling horizon and the frequency of re-scheduling are the features defining our three policies.

The computational results show that an optimization approach using logic-based Benders decomposition, scheduling over a longer horizon, incorporating the known information on aircraft failures, and adjusting the repair schedule as soon as new jobs enter the repair shop yield higher mean coverage. The results also provide evidence that quick reaction to the aircraft failures is more important than scheduling over a longer horizon as the policy with a higher frequency of re-scheduling does better than the policy with



Figure 12: Mean observed coverage for the first waves in buckets.



Figure 13: Mean observed coverage for the second waves in buckets.



Figure 14: Mean observed coverage for the third waves in buckets.

longer scheduling horizon on the waves scheduled later in each time period.

References

Aramon Bajestani, M., and Beck, J. C. 2011. Scheduling an aircraft repair shop. In *Proceedings of the Twenty-First International Conference on Automated Planning and Scheduling (ICAPS'11).*

Aytug, H.; Lawley, M.; McKay, K.; Mohan, S.; and Uzsoy, R. 2005. Executing production schedules in the face of uncertainties: A review and future directions. *European Journal of Operational Research* 161:86–110.

Bidot, J.; Vidal, T.; Laborie, P.; and Beck, J. 2009. A theoretical and practical framework for scheduling in a stochastic environment. *Journal of Scheduling* 12:315–344.

Bidot, J. 2005. *A general framework integrating techniques for scheduling under uncertainty*. Ph.D. Dissertation, Ecole Nationale d'Ingenieurs de Tarbes, France.

Hooker, J. 2005. A hybrid method for planning and scheduling. *Constraints* 10:385–401.

Iravani, S. M. R.; Krishnamurthy, V.; and Chao, G. 2007. Optimal server scheduling in nonpreemptive finite-population queueing systems. *Queueing System* 55:95–105.

Pinedo, M. L. 2005. *Planning and Scheduling in Manufacturing and Services*. Springer Series in Operations Research. Springer.

Safaei, N.; Banjevic, D.; and Jardine, A. K. S. 2010. Workforce constrained maintenance scheduling for aircraft fleet: A case study. In *Proceedings of Sixteenth ISSAT International Conference on Reliability and Quality in Design*, 291–297.

Vieira, G.; Hermann, J.; and Lin, E. 2003. Rescheduling manufacturing systems: a framework of strategies, policies and methods. *Journal of Scheduling* 6:36–92.