OPTIMAL PREDICTIVE MAINTENANCE SCHEDULING OF AIRCRAFT

Saakaar Bhatnagar, Eleonore Jacquemet

Stanford University

ABSTRACT

For an airline, aging aircraft fleets present a challenge in that they need more maintenance, and are less efficient than newly purchased aircraft. Ideally, when a carrier has a mixture of older and newer aircraft at their disposal, they would always want to fly the newer, more efficient versions. But continuously flying the newer fleet is impossible, because regular maintenance of these aircraft is required, which make them unavailable for that time period. Combined with yearly trends in passenger flying and stochastic fuel prices,optimal predictive maintenance of more efficient aircraft is beneficial to the airline, which wants to maximise revenue.

In this paper, we frame the given problem as an MDP and solve it using several algorithms, exploring the comparison of average earned profits between an optimal maintenance policy and other policies like a greedy policy, which is more commonly used.

1. INTRODUCTION

With the introduction to the market once every few years of new aircraft with improved aerodynamics, fuel efficiency and better avionics, flying is getting cheaper, more fuel efficient and environment friendly, while also enabling airlines to become more and more profitable. A purchased aircraft has a life span dependent on pressurization cycles, and short haul aircraft last anywhere between 15-20 years in service, while some larger aircraft (like the B-747) from the 1980s are still in service today.

This means that any airline will always have, at a given time, a mixture of older and newer aircraft in their fleet. With the introduction of new technology, flying the older models on commercial routes becomes less desirable, and in the ideal world a carrier would always want to fly the newest, most efficient aircraft they have on a route, because using that aircraft will generate more profits.

But that isn't always possible, since aircraft must undergo unscheduled and scheduled maintenance from time to time, in regulation with aviation law and safety procedures. There will hence be a portion of an airline fleet, at any given time that would be unavailable to operate. This would consist of a mixture of old, almost retired aircraft along with almost new, technologically superior aircraft.

So why should the timing of maintenance schedules for aircraft be important? It turns out that passenger volumes (represented in aviation statistics by a very important fraction called load factor) are seasonal; i.e there are certain times of the year in which there is a lot more travel, compared to others. For example, summer months are the holiday season in which passenger volume surges. In these months, ticket prices surge as well, and in these months the profit per passenger increases. Infact, this seasonality is so pronounced that several airlines globally that have not managed do well in these holiday months have, in recent years declared bankruptcy in the following fall months due to a dip in demand and not enough capital accumulation from the holiday months [1].

Hence, in order for carriers to fully capitalise on the seasonal demand for flying, it would be ideal for them to fly aircraft such that their profit per passenger flown is maximum as per season. This is a very complicated problem due to sizes of aircraft fleet, unscheduled maintenance requirements and further exacerbated by the fluctuating price of Jet-A fuel, the number one driver of costs in the aviation industry.

This paper is an attempt to come up with a simplified, 2 aircraft (versus a diverse fleet) model of the problem and solve it as a Markov Decision Process. We attempt to demonstrate the benefit of prior planning of scheduling maintenance for older and newer models of aircraft, in order to maximise expected revenue. We show comparisons with a greedy policy, which in the present case is to simply fly the more efficient aircraft all the time until it is not available due to required maintenance. We investigate the effect of the search horizon on the expected profits, and conclude with results and how this work can be expanded.

2. LITERATURE REVIEW

Due to the increasing costs and competitiveness of flying, there has been active research to try and reduce the overhead costs of maintenance for airlines. One of the biggest areas of work is to enable predictive maintenance of airline fleets using big data, to improve component failure prediction, increase availability, and optimize profits.

Feo and Bard [9] presented a model that both located maintenance stations and developed flight schedules, in order to meet the cyclic maintenance requirement for aircraft. Using a closed loop network as a model, with nodes representative of geographical locations where maintenance is possible, and training their network using data from a fleet of 727s, they attempted to eliminate some maintenance bases(nodes) in the network to bring down costs for airlines.

The University of Maryland [7] developed a heuristic model to optimize aircraft maintenance scheduling and reassignment. Sriram and Haghani constructed in their paper an horizon cyclic schedule with maintenance constraints for heterogeneous fleet of aircraft. The maintenance-scheduling problem is modeled as well as a closed loop model which takes as input an Origin Destination pair for each aircraft. Maintenance costs minimization is based on maintenance scheduling and aircraft re-assignments (which induce a penalty taking into account in the model).

Higle and Johnson in their paper 'Flight Schedule Planning with Maintenance Considerations' [7] used a model of maintenance opportunities. Indeed, they incorporate the different maintenance requirements by counting maintenance opportunities in the time window where the maintenance should occurred. They chose the best time to execute the maintenance while maintaining the schedule.

In light of the above approaches, we wish to modify the traditional closed-loop network approach taken by modelling a similar problem as a Markov Decision Process. Instead of reducing costs as most of the above methods do, we aim to directly maximize expected profits by making the profit our utility function.

3. METHODOLOGY

3.1. The Bellman Equation

The fundamental approach in this work is to obtain an optimal action for the current time step, simulate the availability of aircraft in the next time step, and then obtain the optimal action for the next time step, and so on. The fundamental equation solved for this is the Bellman Equation:

$$U_t^*(s,a) = max_a(R(s,a) + \gamma \sum_{s'} T(s'|s,a)U_{t+1}^*(s'))$$

An explanation of every component of the above equation in this problem, is as follows:

- Reward Function: Estimated using real world data, it is of the following form:

  $$((T.P * L.F - O.C) * Capacity * NFlights)$$

  where

  - T.P: Ticket Price: dependent on demand(load factor).
  - L.F: Load Factor: Obtained at every time step using real data [2], depends on the month to capture seasonality of travellers.

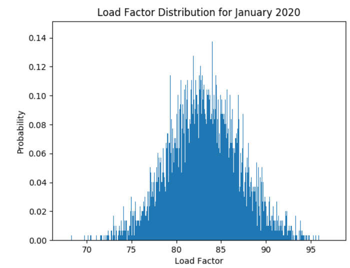

(a) Load Factor Distribution, January

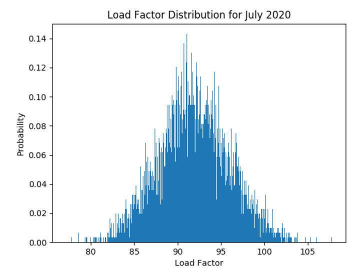

(b) Load Factor Distribution, July

Fig. 1: Load Factor distribution for different months. Note the difference in mean values.

  - O.C: Operating cost : Operating cost per seat per flight of aircraft, obtained from real world data. [4] [3]. Note that it is modelled as independent of load factor. It can be further broken down into:

    $$O.C = F.C + FuelCost + M.C$$

    where

    * F.C= Fixed costs: Approximated at 40 $,per seat per flight
    * Fuel Cost=Cost per seat per flight. Sampled from a distribution based on real world data prediction. US Energy Information administration published short-term energy outlook on fuel price. They detailed futures price and futures price 95 % confidence interval. [5] **Note: We assume that the total cost of fuel is independent of the number of passengers flying, since the weight of passengers flying is a fraction of the total airborne weight.**
    * M.C= Maintenance Costs per seat per flight
  - Capacity: Seating capacity of aircraft

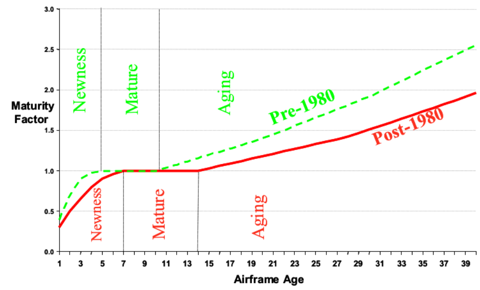

Fig. 2: Boeing Maturity Curve[4] depicting factor of maintenance cost due to aging. This means older aircraft have higher M.C in the formula above, reducing profits

  – NFlights: No. of flights being simulated

- T(s'|s,a): Estimated based on the recent history of flying of the aircraft, by keeping track of counts. A function updates the "counts" based on the action taken in the timestep, which modifies T(s'|s,a) for the next timestep.
- Gamma: For this problem we select it to be 1, because we have finite horizons and future rewards are as important as current rewards.
- States: In this MDP states represent availability of aircraft to fly at the current time step. For example, a state of [1,0] represents that aircraft 1 is available to use and aircraft 2 is not (is in maintenance).
- Action: Represents which aircraft(1 or 2) we use a the current timestep. The action space at a timestep is dependent on the state we are in. For example, the action space at state [1,1] is [1,2] and at [1,0] is [1].

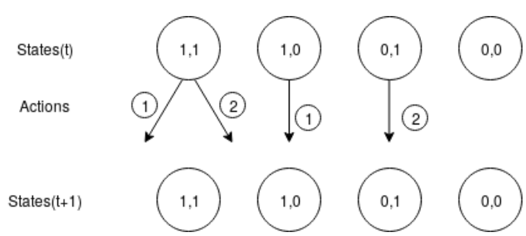

Fig. 3: Depiction of MDP

It is important to note that the game does not end if state [0,0] is reached. Achieving that state simply means no flight was possible at that timestep. The reward associated with such a state is obviously the most negative.

In the current problem, we use data on a monthly basis as available on the internet at the Bureau of Transport Statistics [2]. Hence, one time step represents a month.

We used two algorithms, namely forward search and policy iteration.

3.2. Algorithms used

3.2.1. Forward Search

The forward search algorithm uses recursion to search the tree to some horizon d to find the optimal action to take at the current timestep. It solves the given MDP exactly upto the given horizon.

Algorithm 4.6 Forward search

```
function SelectAction(s, d)
    if d = 0
        return (NIL, 0)
    (a*, v*) ← (NIL, −∞)
    for a ∈ A(s)
        v ← R(s, a)
        for s′ ∈ S(s, a)
            (a′, v′) ← SelectAction(s′, d − 1)
            v ← v + γT(s′ | s, a)v′
        if v > v*
            (a*, v*) ← (a, v)
    return (a*, v*)
```

Fig. 4: The Forward Search Algorithm [6]

3.2.2. Policy Iteration

The policy iteration algorithm computes the optimal policy by following two steps:

- Policy evaluation: The algorithm estimates the reward of a given policy.
- Policy improvement: The algorithm finds the best policy among the possible policies.

Algorithm 4.2 Policy iteration

```
function PolicyIteration(π_0)
    k ← 0
    repeat
        Compute U^{π_k}
        π_{k+1}(s) = arg max_a (R(s, a) + γ Σ_{s′} T(s′ | s, a) U^{π_k}(s′)) for all states s
        k ← k + 1
    until π_k = π_{k−1}
    return π_k
```

Fig. 5: Policy Iteration Algorithm [6]

3.3. Execution of Policy and Time Stepping

With the theoretical approach declared above, we can solve the MDP at every timestep. With the absence of real-world data on the results of execution of various types of policies, we instead use a simulation of timestepping. The simulation proceeds according to the same probabilities defined in T(s'|s,a). We perform as many time steps as nMonths, the number of months we wish to simulate.

Because the results are stochastic, we carry out many simulations, with each simulation executing timesteps nMonths from the same initial time. The results demonstrate the convergence of simulated rewards.

4. RESULTS

4.1. Results from Forward Search

4.1.1. Convergence of Simulations

In this section, we aim to demonstrate convergence of simulations. Due to the inherently stochastic nature of the problem, we demonstrate convergence of simulated rewards by tracking the running average of the same against number of simulations. Every simulation is independent of the previous one.

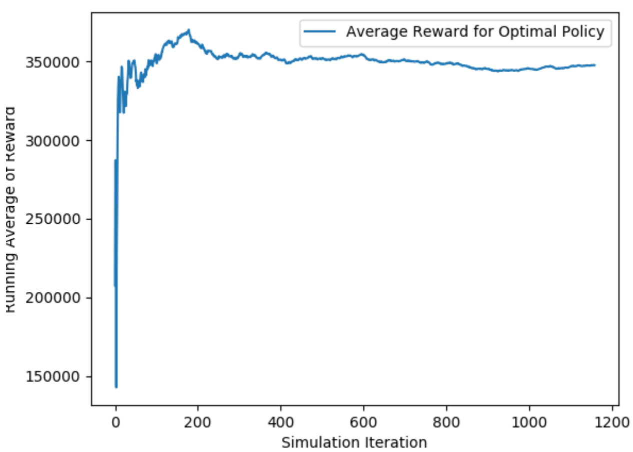

Fig. 6: Simulated reward for 5 months using a search horizon of 3. The running average converges to a value, which is the expected reward of following an optimal policy.

4.1.2. Effect of Horizon

In this section, we show the effect of search depth on both the expected simulated reward and the running time of the simulation. We calculate an average 10-month reward using various search horizons and compare the results.
With an increase of search depth, we do not see much improvement between converged optimal policy results and greedy policy results, however the time to compute rises very quickly. This suggests that while using some lookahead(depth) is a good idea, looking to far ahead into the future does not help.

Depth	1	3	5	7
Time to Converge(s)	7.36	50.5	976.5	22664
Optimal Reward(x10^5)	7.7	8.14	8.23	8.31

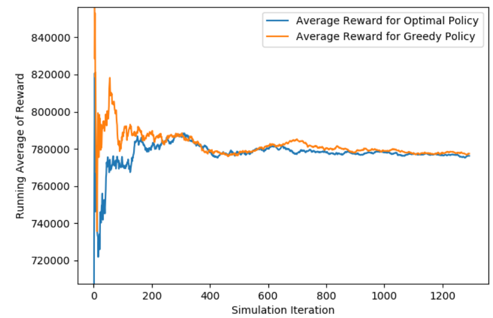

Fig. 7: Effect of horizon experiment: Depth=1, which implies greedy search.

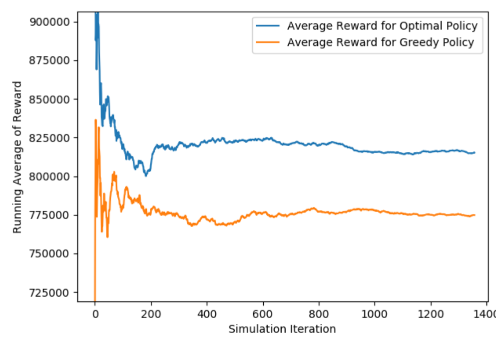

Fig. 8: Effect of horizon experiment: Depth=3. Does better than greedy search.

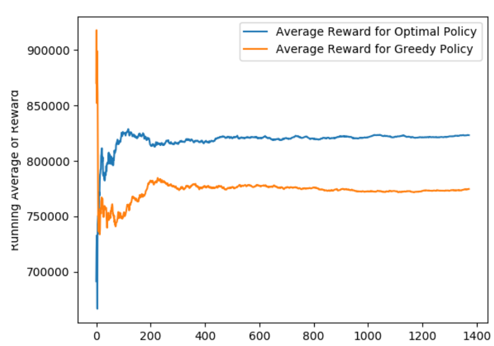

Fig. 9: Effect of horizon experiment: Depth=5. Approximately the same convergent reward as d=3.

4.2. Policy Iteration

So far we employed an online method. The drawback of employing an online method for this problem is that while obtaining the optimum action for the last months we end up looking at months further ahead which aren't factored into our reward. Hence, now we use an offline method to predict an optimal policy which gives us the best policy for our frame of search (10 months), without considering the future after that.

The policy iteration algorithm aims to find the best policy among the possible policies given the number of future months to look at. It obtains the optimal reward by iterating through all possible policies for that number of future months. Due to stochasticity in the problem dynamics, we ran policy iteration multiple times until convergence.

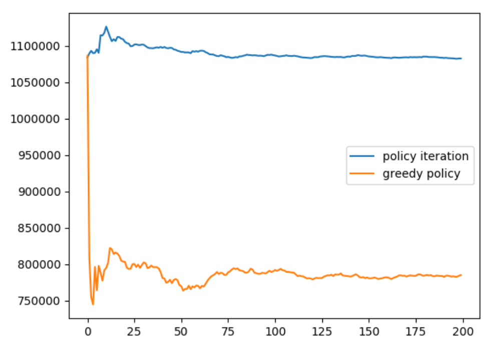

Fig. 10: Policy iteration: converged reward for a 10 month plan.

From the plot it is immediately visible that policy iteration gives us a much better result than forward search does. This was expected since since policy iteration gives us the best plan for exactly 10 months. However, it does not take into account time after that.

From looking at the resultant optimal policies, we can see that in the most crowed months (June, July, August), these policies use the new aircraft(i.e the action is always 1). On the contrary, the greedy policies may not be able to use the new aircraft as it has been used in previous months and should go to maintenance.

5. CONCLUSIONS AND FUTURE WORK

The aim of the work was to demonstrate how planning into the future can help optimize policy to maximise a reward function over a specified time period. Specifically, how this can be applied to the problem of aviation scheduling in order to make optimum use of improved technology, seasonal load factor cycles, while keeping taking into account stochastic variations in the main drivers of industry cost such as fuel price.

We were able to demonstrate the usefulness of such planning, and also showed comparative studies of problem hyperparameters such as search depth.

We demonstrated two different algorithms and showed the expected rewards from them.

This work can be easily and rapidly expanded to an industrial scale with the availability of more accurate real world data, on various time scales.

- Although the present study has been based on monthly data, in industry the practice we have tried to optimise occurs on a time scale closer to a weekly basis.
- In industry, many more factors would have to be taken into account such as the size and demographics of the fleet, further routing of aircraft(here we assume the aircraft flies between just 2 ports for the entire time), availability of maintenance facilities at given ports, to name a few.
- The results shown are heavily dependent on the problem dynamics. Modeling the probabilistic problem dynamics in this case is a rather subjective matter. A better way to do it would be using real world data on maintenance requirements with continuous usage of an airframe, but such information is not public and is highly region/law dependent.

6. CONTRIBUTIONS

6.1. Saakaar Bhatnagar

Part of the literature review, mathematical formulation of problem, partial preparation of real world data, forward search algorithm implementation and it's results.

6.2. Eleonore Jacquemet

Part of literature review, partial preparation of real world data, policy iteration algorithm implementation and it's results.

7. REFERENCES

[1] Wendover Productions: Why So Many Airlines are Going Bankrupt, Youtube

[2] Bureau of Transport Statistics: Passenger Travel Facts and Figures(United States Department of Transportation)

[3] Wendover Productions: Why Flying is So Expensive, Youtube

[4] The Maintenance Costs of Aging Aircraft: Insights from Commercial Aviation, Matthew Dixon: Page 14

[5] US Energy Information Administration: Short-term Energy Outlook

[6] Decision Making Under Uncertainity, Mykel Kochenderfer, Page 81,99

[7] Flight Schedule Planning with Maintenance Considerations, Julia L. Higle Anne E. C. Johnson

[8] An optimization model for aircraft maintenance scheduling and re-assignment, Chellappan Sriram, Ali Haghani

[9] Thomas Feo, Jonathan Bard: Flight Scheduling and Maintenance Base Planning, Management Science, Vol 35, No 12, December 1989