A Markov Decision Process Model for Airline Meal Provisioning

by

Jason Goto, University of British Columbia

<This article is based on the award-winning paper in the 1999 Student Paper Prize Competition>

ABSTRACT

An airline caterer seeks to provide a meal quantity for each flight that closely matches final on-board passenger load. Faced with preparation lead-time, the caterer must estimate required meal quantities well in advance of departure. Passenger load may vary considerably during this lead-time. Thus, adjustments are often required as more information becomes available.

We describe the application of a Markov decision process to the meal ordering process in an attempt to identify ordering policies that outperform actual practice. Marginal improvements in overage and shortage are achieved with the model: overage costs are reduced by approximately 17% and the proportion of flights short-catered are reduced by 33%. In addition, we evaluate the model over a range of scenarios to determine the costs associated with achieving a low proportion of short-catered flights.

INTRODUCTION

The passenger airline industry operates on low profit margins with many competitors. Airline carriers sustain profitability through operational efficiency improvements and by maintaining or increasing market share.

Classic applications of Operations Research in the airline industry include yield management, route selection, and optimization of flight and maintenance schedules. Other applications include crew scheduling, analysis of customer demand, and fuel consumption control. The majority of these applications focus on operations central to the airline.

In the current competitive environment, some carriers are attempting to generate savings through efficiency improvements in their periphery operations. Savings generated from increased efficiency may be directed toward improving customer service. Inflight meal provisioning is one such area worthy of pursuit as it involves high volumes and significant costs, and has direct impact on customer service.

The results and methodology described in this document are the result of an applied industry research project involving the Centre for Operations Excellence (COE) and Canadian Airlines International Ltd. The COE provides business solutions to industry problems that require specialized research capabilities, through the Faculty of Commerce at the University of British Columbia.

INFLIGHT CATERING

Inflight meal catering refers to the provisioning of a meal service for each passenger during a flight. This service is typical of long duration flights. The complexity of the meal service varies by class of passenger service and flight destination.

At several key decision points, prior to the departure of a flight, the final passenger load is estimated and the meal order is adjusted accordingly. Typically, the caterer may access information concerning tickets sold, passengers checked in, and the number of stand-by passengers. The final passenger load estimate is based on this information and the personal judgement of the catering staff. The estimation, monitoring, and adjustment of the required meal quantity is referred to as *meal ordering*.

Excess costs are incurred when the meal quantity on the aircraft exceeds the passenger load. Furthermore, customer service costs are incurred when the meal quantity is lower than the passenger load. Excess provisioning is referred to as *overage* and short-provisioning is referred to as *shortage*. The caterer seeks to provide a meal quantity that closely matches the passenger load at departure.

THE PROBLEM

The solution to the problem appears to be trivial: simply observe the number of tickets sold for a given flight and provision a meal quantity accordingly. We invested considerable effort early in the project to develop an understanding of the processes, stakeholders, data, and costs associated with meal provisioning. As a result, we developed a good appreciation of the problem, and found that the process of meal ordering was challenging due to two basic reasons:

1. Significant lead time is required to produce a meal order. Meal provisioning involves preparation, cooking, assembling, chilling, and transporting the meal order, and in some airports, large flights depart within minutes of each other.
2. The passenger load may vary considerably within the lead time. Last minute ticket purchases, missed flights, stand-by passengers, and upgrade coupons all contribute to variability in passenger load.

In this project we attempt to address the following problem statement: *Given the information available and the current processes, can the variability in meal provisioning be reduced?* Reducing variability in meal provisioning would simultaneously reduce overage and shortage.
Our approach was to model the meal ordering process as a finite-horizon discrete-time Markov decision process (MDP) on a flight by flight basis. A Markov decision process is a sequential decision problem where the set of actions, rewards, and transition probabilities depend only on the current state of the system and the current action selected. The history of the problem has no effect on the current decisions. Markov decision processes are commonly applied in inventory control settings where demand for a product follows a known (or estimated) probability distribution.

METHODOLOGY

For each flight that requires a meal order, the caterer is faced with a series of decisions. The caterer must estimate the initial quantity of meals to build, and later decide whether to adjust the order as more information becomes available. Figure 1 depicts a generalization of the timeline preceding the departure of a flight.

Note that there are five opportunities to set or adjust the meal order quantity before the flight departs. Costs and penalties are imposed depending on the time of the adjustment. For example, a fixed van delivery charge is incurred on any adjustments that occur after the meal order has been delivered to the aircraft.

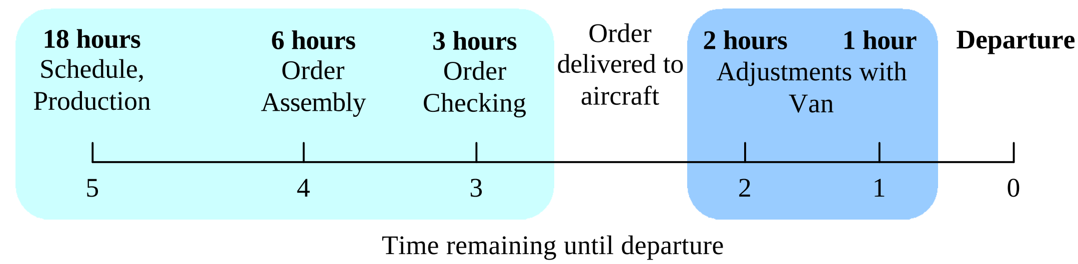

Figure 1. Timeline of events / Chronologie des événements

In formulating the meal ordering process as an MDP, we represent the states of the model as all possible combinations of booked passenger load and meal quantities. We can choose actions to move between states. For example, if the current state of the system is 95 meals and 100 booked passengers, we may choose to adjust the meal order quantity to 100 meals. This adjustment of +5 meals is an action. However, we may expect the booked passenger load to increase by 10 passengers with a probability *p*, or stay the same with a probability of *1-p*. Such a system is depicted in Figure 2.

Figure 2. A simplified one-stage decision problem / Un problème de décision à une étape

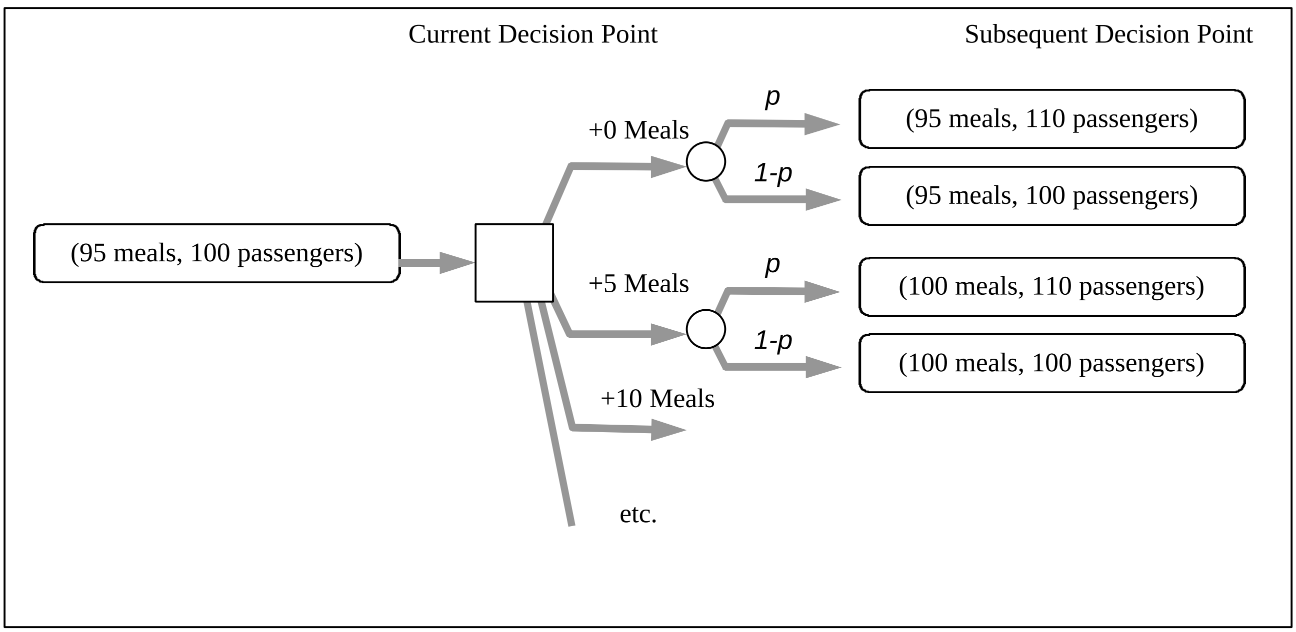

We can see that depending on the transition probabilities *(p, 1-p)* and the chosen action, the system will arrive in a certain state for the following decision point. Each action has an associated cost, where a reward is represented by a negative cost. The decision maker chooses the action that results in the minimum expected cost, based on the transition probabilities. This example describes a simplified instance of a single stage problem. To better represent the meal ordering process, we consider all possible states, actions and transition probabilities. A Markov decision process evaluates the problem over the entire decision-making horizon (five stages), and identifies the actions that minimize expected cost, given the state of the system, and the decision stage. The optimal order adjustments tables are referred to as *decision rules*, and the collection of all decision rules is referred to as a *policy*. A decision rule is created for each decision point in the timeline.

Using the pre-departure booked passenger load obtained from Canadian Airlines historical tables, we generated the transition probabilities on a flight-by-flight basis. Significant effort was spent in identifying appropriate parameters for modelling the probability distributions of change in passenger load between decision points. We then applied meal costs and penalties as described in discussions with Inflight Service staff and based on data collected from historically billed information.

To test the validity of the policy, we holdback a portion of the pre-departure data. Thus, a portion of the data is used for developing the transition probabilities, thereby producing an optimal policy. We apply this optimal policy to the held back data, and compare the performance of the optimal policies to the performance observed in actual practice. We measure performance in terms of *provisioning error*, which is final meal quantity minus passenger load at departure, calculated on a flight instance basis (i.e. daily for each flight number). Positive provisioning error represents overage and negative provisioning error represents shortage.

All analysis was conducted on data spanning from December 1998 to January 1999 inclusive. Ten months of the data were used for developing the transition probabilities, and the remaining two months of data were used as a validation dataset.

RESULTS

The model was evaluated over a sample of 55 flight routes and the performance of the ordering policies were observed. During initial testing and evaluation, it became evident that the terminal cost associated with short-catering played a large role in the behavior of the policy. A low terminal cost resulted in an ordering policy biased towards shortage, whereas a high terminal cost resulted in an ordering policy biased towards overage. We evaluated the model over a range of terminal costs and chose the terminal cost that most closely matched current practice.

The optimal policies were applied to the test dataset and the distribution of provisioning error was compared to the distribution of provisioning error observed in actual practice. Surprisingly, in many cases current meal ordering performance was close to optimal. This may indicate that the meal ordering staff exercise good judgement, and/or the meal ordering staff have access to information unavailable to the model.

The optimal policies provide some encouraging results. In the sample group, it was estimated that application of the optimal policies would result in a savings of 17% in overage costs. In addition, the proportion of flights short-catered would simultaneously drop by 33%. Thus, savings are generated while improving service level of the operations.

In addition, the model allows us to evaluate the expected cost of various levels of performance. Using the model, we were able to estimate the cost associated with zero short-catering on flights departing from the Vancouver station. This result or other scenarios may be useful in determining how much to invest in improving the processes to achieve a desired level of service.

CONCLUSION

Given the current processes and information available, it is possible to marginally reduce the variability in meal provisioning. While the Markov decision process model does not drastically improve meal ordering performance, we may obtain modest gains in cost savings and service level. The model provides Canadian Airlines with a reference point for their current meal provisioning performance, as well as enables them to estimate the costs associated with improving their service level. Furthermore, the model may be used for scenario analysis, as it is easily adaptable to changes in business practices or cost structures.