

Airline Network Inventory Control Policy: An Application of Markov Decision Processes and Simulation

by

Kyle Adrian Biswanger

B.A.Sc. Mechanical Engineering, University of British Columbia, 1999

A Thesis submitted in partial fulfillment of the requirements for
the degree of Master of Science (Business Administration)

in

The Faculty of Graduate Studies

(Department of Commerce and Business Administration)

We accept this thesis as conforming
to the required standards

The University of British Columbia

November 2000

©Kyle Adrian Biswanger, 2000

In presenting this thesis in partial fulfilment of the requirements for an advanced degree at the University of British Columbia, I agree that the Library shall make it freely available for reference and study. I further agree that permission for extensive copying of this thesis for scholarly purposes may be granted by the head of my department or by his or her representatives. It is understood that copying or publication of this thesis for financial gain shall not be allowed without my written permission.

Department of Management Science, faculty of commerce

The University of British Columbia
Vancouver, Canada

Date Dec 10th 2000

Abstract

In this thesis an airline network inventory control policy is developed. The inventory problem being faced differs from traditional hierarchical inventory problems. In this problem the inventory system is composed of a network of warehouses where demand is observed between warehouses and inventory "consumed" from one warehouse is "provided" to another.

The control policy is developed using two modeling techniques: Dynamic Programming and Simulation. Dynamic programming, or Markov Decision Processes, is used to help define a control policy. The limitations imposed by the use of Dynamic Programming are overcome by implementing the results in a Simulation model.

The policy resulting from this analysis demonstrates the potential to save \$390,000 per year with minimal operational impact to Canadian Airlines International. This represents a 60% reduction in non-service related equipment movements.

Table of Contents

Abstract.....	ii
Table of Contents.....	iii
List of Tables.....	iv
List of Figures.....	v
Acknowledgements.....	vii
1.0 Background.....	1
1.1 Problem Description.....	2
1.2 An Illustrative Example.....	2
1.3 General System Description.....	4
1.4 Definitions.....	4
1.5 Canadian Airlines System Description.....	4
1.6 Data Availability.....	8
2.0 Literature Review.....	9
3.0 Methodology.....	12
3.1 Gather Information.....	12
3.2 Develop Markov Decision Processes.....	13
3.3 Develop Simulation.....	14
4.0 Markov Decision Process: Model Formulations.....	15
4.1 States and Actions.....	15
4.2 Model Assumptions.....	15
4.3 Model Cost Function.....	17
4.4 Model 1: 1-Node Model Formulation.....	20
Model 2: 2-Node Model Formulation.....	23
4.5 Model 3: 3-Node Model Formulation.....	26
5.0 Markov Decision Process: Model Results.....	30
5.1 Model 1 Results:.....	32
5.2 Model 2 Results:.....	42
5.3 Model 3 Results.....	51
6.0 Simulation: Model Formulation.....	57
6.1 Simulation Model Inputs.....	57
6.2 Simulation Model Logic.....	59
6.3 Simulation Model Metrics.....	59
6.4 Simulation: Alternative Equipment Provisioning Policy.....	60
7.0 Simulation: Model Results.....	64
7.1 USSR Results.....	65
7.2 Reducing the Number of Cities.....	68
7.3 Inventory Imbalance Stability.....	70
7.4 Optimizing λ	72
7.5 The effect of S^L , $\gamma = 1$	75
7.6 Selecting a Warm up Period.....	77
7.7 Final Simulation Models.....	78
7.8 Cost Savings Analysis.....	79
8.0 Conclusion.....	80
9.0 Future Recommendations.....	81
Glossary.....	82
Bibliography.....	83
Appendix A: Solving the 3 Node Model.....	84

List of Tables

TABLE 1 – MDP MODEL COST STRUCTURES.....	19
TABLE 2 – MDP POLICIES	31
TABLE 3 – 1 NODE MODEL PARAMETERS.....	33
TABLE 4 – 1 NODE OPTIMAL VALUE OF X^H	39
TABLE 5 – 2 NODE MODEL PARAMETERS.....	42
TABLE 6 – 2 NODE OPTMAL VALUE OF X^H	48
TABLE 7 – SUMMARY OF PROVISIONING POLICIES.....	64

List of Figures

FIGURE 1 – GRAPHICAL REPRESENTATION OF PROBLEM	3
FIGURE 2 – AN EXAMPLE OF EQUIPMENT FLOW IMBALANCE	4
FIGURE 3 – AIRLINE TRAFFIC	5
FIGURE 4 – GETTING THE EQUIPMENT MANIFEST	7
FIGURE 5 – TRADITIONAL HIERARCHAL INVENTORY MODEL	10
FIGURE 6 – NETWORK INVENTORY MODEL	11
FIGURE 7 – PICTORIAL REPRESENTATION OF PROJECT PLAN	12
FIGURE 8 – 1 NODE MDP MODEL	20
FIGURE 9 – 2 NODE MDP MODEL	23
FIGURE 10 – 3 NODE MDP MODEL	26
FIGURE 11 – 1 NODE OPT SL 100, CONTROL SURFACE: COST STRUCTURE 1,1,2	34
FIGURE 12 – 1 NODE OPT SL 20, CONTROL SURFACE: COST STRUCTURE 1,1,2	35
FIGURE 13 – EXPECTED COSTS OF THE 1 NODE POLICY.....	36
FIGURE 14 – 1 NODE SR MODEL SUMMARY	37
FIGURE 15 – 1 NODE DH POLICY SUMMARY.....	37
FIGURE 16 – 1 NODE COMPARISON OF OPT AND DH POLICIES.....	38
FIGURE 17 – 1 NODE POLICY COST COMPARISONS.....	39
FIGURE 18 – 1 NODE SR POLICY PERFORMANCE, EXPECTED COST	40
FIGURE 19 – 1 NODE SR POLICY, RELATIVE MAGNITUDE OF EXPECTED COSTS.....	40
FIGURE 20 – 1 NODE SR POLICY PERFORMANCE, % PERFORMANCE	41
FIGURE 21 – 2 NODE OPT, CONTROL SURFACE: COST STRUCTURE 3,2,3	43
FIGURE 22 – 2 NODE OPT POLICY SUMMARY	44
FIGURE 23 – 2 NODE SR POLICY EXECTED COSTS	45
FIGURE 24 – 2 NODE DH POLICY SUMMARY.....	46
FIGURE 25 – 2 NODE POLICY COST COMPARISONS.....	47
FIGURE 26 – 2 NODE SR POLICY PERFORMANCE, EXPECTED COST	49
FIGURE 27 – 2 NODE SR POLICY, RELATIVE MAGNITUDE OF EXPECTED COSTS.....	49
FIGURE 28 – 2 NODE SR POLICY PERFORMANCE, % PERFORMANCE	50
FIGURE 29 – 3 NODE OPT POLICY SUMMARY	51
FIGURE 30 – 3 NODE SR POLICY SUMMARY	52
FIGURE 31 – 3 NODE DH POLICY SUMMARY.....	53
FIGURE 32 – 3 NODE POLICY COST COMPARISONS.....	54
FIGURE 33 – 3 NODE SR POLICY PERFORMANCE, EXPECTED COSTS.....	55
FIGURE 34 – 3 NODE POLICY PERFORMANCE, RELATIVE MAGNITUDE OF EXPECTED COSTS	55
FIGURE 35 – 3 NODE SR POLICY PERFORMANCE, % PERFORMANCE	56
FIGURE 36 – HISTOGRAM OF NORMALIZED CITY SIZE.....	62
FIGURE 37 – PERCENT SYSTEM VOLUME BY STATION GROUP.....	62
FIGURE 38 – USSR, MINIMUM INVENTORY IMBALANCE, % OF S^L , $\gamma = 1.3$	66
FIGURE 39 – USSR, # OF TIMES INVENTORY IMBALANCE DECREASES BELOW $-S^L$, $\gamma = 1.3$	67
FIGURE 40 – USSR, AVERAGE EXCESS EQUIPMENT MOVED, $\gamma = 1.3$	67
FIGURE 41 – RSSR, MINIMUM INVENTORY IMBALANCE DEVIATION, %, $\gamma = 1.3$	68
FIGURE 42 – RSSR, # OF TIMES INVENTORY IMBALANCE DECREASES BELOW $-S^L$, $\gamma = 1.3$	69
FIGURE 43 – RSSR, AVERAGE EXCESS EQUIPMENT MOVED, $\gamma = 1.3$	69
FIGURE 44 – STATION STABILITY, LARGE STATION, $\lambda = 0.1$	70
FIGURE 45 – STATION STABILITY, MEDIUM STATION, $\lambda = 0.1$	71
FIGURE 46 – STATION STABILITY, LARGE STATION, $\lambda = 0.2$	71
FIGURE 47 – STATION STABILITY, MEDIUM STATION, $\lambda = 0.2$	71

FIGURE 48 – STATION STABILITY, LARGE STATION, $\lambda = 0.125$	72
FIGURE 49 – STATION STABILITY, MEDIUM STATION, $\lambda = 0.125$	73
FIGURE 50 – STATION STABILITY, LARGE STATION, $\lambda = 0.15$	73
FIGURE 51 – STATION STABILITY, MEDIUM STATION, $\lambda = 0.15$	73
FIGURE 52 – STATION STABILITY, LARGE STATION, $\lambda = 0.175$	74
FIGURE 53 – STATION STABILITY, MEDIUM STATION, $\lambda = 0.175$	74
FIGURE 54 – RSSR, % AVAILABLE COST REDUCTION CAPTURED, $\gamma = 1.3$, Π	75
FIGURE 55 – RSSR, MINIMUM INVENTORY IMBALANCE, % OF S^L , $\gamma = 1.0$	76
FIGURE 56 – RSSR, # OF TIMES INVENTORY IMBALANCE DECREASES BELOW $-S^L$, $\gamma = 1.0$	76
FIGURE 57 – RSSR, SYSTEM TOTAL EXCESS EQUIPMENT MOVEMENTS, $\lambda = 0.125$, $\gamma = 1.3$	78

Acknowledgements

I would like to thank the following individuals and organizations for their help in providing the opportunity to work on this thesis and their help in actually completing this thesis: Canadian Airlines International, and the Centre for Operations Excellence for providing me with an exciting project to work on, Dr. Marty Puterman and Dr. Hong Chen for helping me complete this thesis, Stephen Jones for making my decision to come the COE and easy one, Lauren Gray for her substantial help and creative input developing materials for the conferences at which I presented this project, my family and friends for there support, and fellow student Steven Kabanuk who offered insight and made long nights at the COE a little more bearable.

1.0 Background

This project was developed for the Inflight Service department of our¹ partnership company, Canadian Airlines International (CAI). CAI is one of Canada's two major airlines², with service to many domestic and international cities. The objective of the Inflight Service department is to enhance the primary product of the airline, passenger transportation. To do this, the Inflight Service department provides services such as meals, reading materials and video entertainment on many of CAI's flights.

To support inflight operations, CAI must monitor and maintain adequate levels of stock at the many cities it operates flights to. One of the difficulties of maintaining stock levels is managing the daily inflows and outflows of equipment at each station³. Equipment is loaded onto each flight, decreasing the departing city's inventory. Upon the completion of a flight, the equipment is unloaded and added to the arrival city's inventory. The problem of maintaining adequate inventory levels is amplified by the existence of many levels of service, each requiring a unique set of equipment, and the volatility in passenger demand.

The airline industry is cost competitive. This requires CAI to examine its costs and seek new ways of doing business that either enhance service or reduce costs. Two areas, related to inflight service equipment, where costs are incurred include 1) the amount of equipment or inventory in the CAI system, and 2) the provisioning of that equipment to each flight that CAI operates. The cost incurred from the amount of inventory in the CAI system consists of the capital required to purchase the equipment and the opportunity cost of that capital. The cost incurred from the provisioning of equipment to each flight is composed of several components, the largest of which is the handling charges for preparing the equipment for loading, and cleaning the equipment after its use.

The current inventory control policy that CAI uses to manage its inflight service equipment is called the deadhead policy. Under this policy each and every flight is loaded to capacity so that inflows and outflows of equipment are balanced. This results in reduced levels of inventory, at the cost of loading equipment onto aircraft that may never be used; most flights do not fly to capacity with passengers. Therefore, one potential area to reduce costs is the provisioning of inflight service equipment (trays, cutlery, carriers, ovens, etc) to aircraft, and this is the topic of this thesis.

The remainder of this section is organized as follows. First the general problem being addressed will be described with the aid of an example. Then the general characteristics of the problem will be identified. The section concludes with a description of the Canadian Airlines system.

¹ The partnership is with the COE- Centre for Operations Excellence UBC

² During the course of this project, Air Canada acquired CAI. Both airlines are currently undergoing merger activities

³ A station is a city that CAI operates flights to and from

1.1 Problem Description

Aircraft provisioning, or determining the amount of equipment to load onto an aircraft, involves the management of thousands of commissary and equipment items. CAI operates approximately 220 flights per day⁴, from over 30 flight kitchens (stations), and each flight may require up to 20,000 pieces of equipment. The equipment may range from coffee pots and high temperature ovens to china and cutlery.

The objective of this study is to devise an equipment provisioning policy that will **reduce the amount of excess equipment boarded onto aircraft and simultaneously ensure that significant equipment flow imbalances are mitigated.**

An equipment provisioning policy is the set of rules that defines the quantities of equipment to be loaded onto an aircraft. When determining the amount of equipment to load onto an aircraft, two things must be considered, the amount of equipment required to provide service to the passengers, and the impact of loading excess equipment on operational costs.

Two general approaches to an equipment provisioning policy are described below. The first is the equipment provisioning policy that CAI currently uses, and the second is an alternative equipment provisioning policy.

1.1.1 CAI's Current Equipment Provisioning Policy

The current equipment provisioning policy dictates that equipment be loaded onto aircraft to capacity⁵, regardless of customer load. This policy has the advantage of balancing the quantities of equipment into and out of a station. However, given that the average occupancy of CAI aircraft is 70%, a significant portion of the equipment loaded onto aircraft is never used to provide inflight service.

1.1.2 Alternative Equipment Provisioning Policy

An alternative equipment provisioning policy would load equipment onto aircraft in relation to the anticipated customer load. This policy would reduce the amount of excess equipment boarded onto flights. However, the inbound and outbound equipment quantities wouldn't necessarily be equal, resulting in a need to balance inventory between stations.


1.2 An Illustrative Example

The following example highlights the difference between the current equipment provisioning policy and an alternative equipment provisioning policy.

⁴ Based on 1999 operating statistics

⁵ capacity is considered to be the amount of equipment required to provide service to the passenger capacity of the aircraft; strictly, this definition of capacity is not true. Additional equipment could be bulk loaded as belly cargo, however, this possibility is ignored for the purposes of this study.

Figure 1 – Graphical Representation of Problem



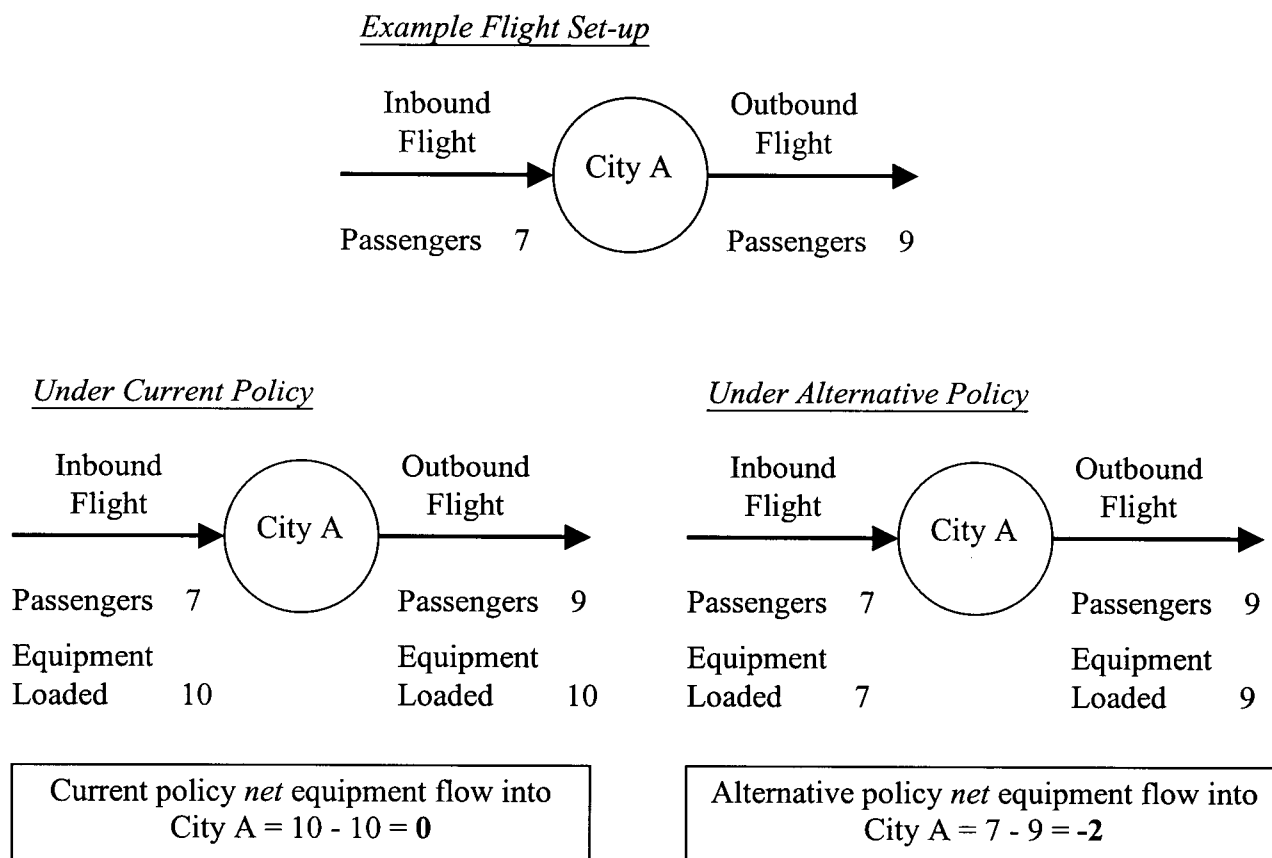
Let us start with a ten seat aircraft and let us assume that there are only seven passengers on the ten seat aircraft, as shown in Figure 1. Under the current provisioning policy ten units⁶ of equipment will be loaded onto the aircraft. During the flight only seven units are actually used to provide service, the remaining three units are not. Upon completion of the flight, all equipment is removed from the aircraft and transported to the flight kitchen. At the flight kitchen the equipment is cleaned and stored for later use. Although the three units of equipment incur costs because of handling charges, breakage and fuel burn, they do not provide any tangible service benefit.

Under an ideal alternative provisioning policy enough equipment is loaded to provide service to the seven passengers. Equipment not required for service is not loaded. In this example, 3 units of excess equipment are not loaded onto the aircraft, resulting in reduced costs. However, when equipment is loaded in relation to passenger demand, equipment flow imbalances can occur. An example of how an imbalance can occur is shown in Figure 2. As a consequence, an alternative equipment policy requires extra movements of equipment to balance inventories between stations. This introduces new logistical concerns and additional costs.

In summary, we are faced with the decision of whether to load equipment onto aircraft to capacity or to customer load. If we choose to load equipment onto aircraft to capacity, we eliminate the additional logistical difficulties but we ignore potential cost savings. However, if we choose to load equipment onto aircraft in relation to customer load, we reduce operational costs, but we introduce new logistical difficulties and possibly additional costs.

⁶ one unit of equipment is the amount of equipment required to provide service to one passenger

Figure 2 – An Example of Equipment Flow Imbalance



1.3 General System Description

The general problem being addressed can be described as an equipment flow on a network of stations (cities). Each station has inventories of several types of equipment. There is a flight schedule that defines all flights to and from the stations of the network. Each scheduled flight has an associated equipment manifest to be loaded prior to its departure. The equipment that is loaded onto an aircraft will be subtracted from the departing station's inventory and added to the arriving station's inventory. Finally, there is a passenger demand for each flight, which determines the required amount of equipment.

1.4 Definitions

Several terms, which may appear to be synonymous, can have a different meaning. In order to avoid confusion, and provide clarity for the remainder of the thesis, several terms and definitions are provided in the glossary at the end of this thesis.

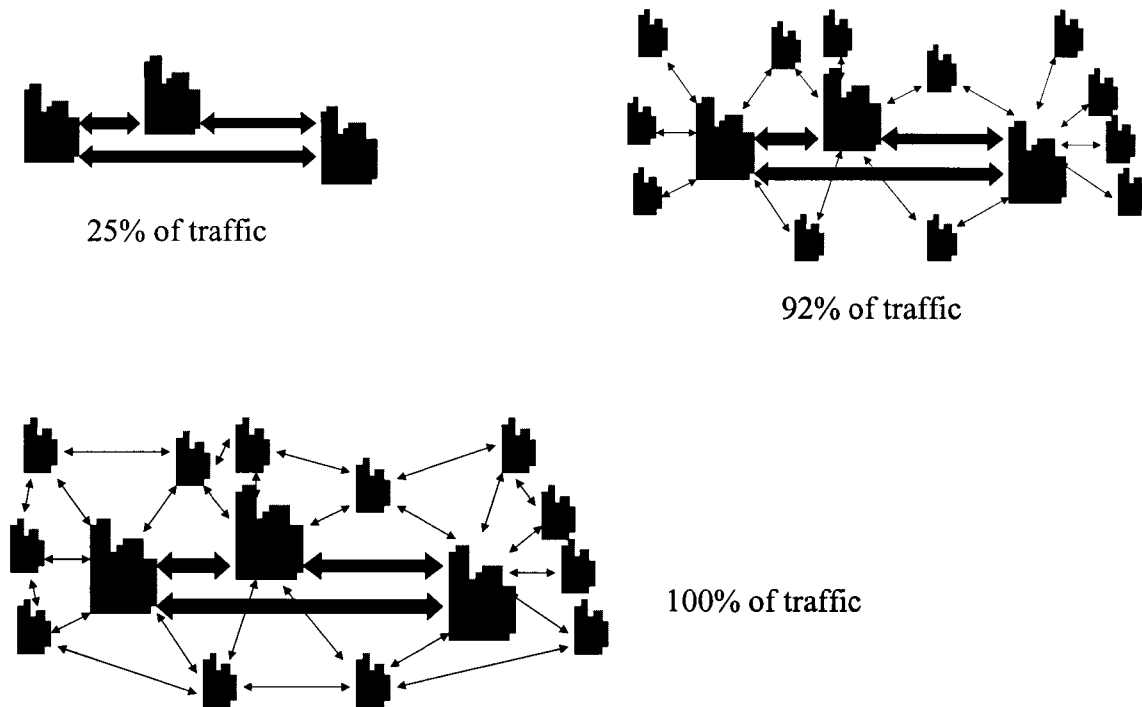
1.5 Canadian Airlines System Description

Defining the CAI system characteristics allows us to translate the general case to the specific case. This requires a description of the network characteristics, the flight schedule, the method for determining a flight's equipment manifest, the equipment inventories and the nature of the passenger demand.

1.5.1 Network Description

The Canadian Airlines network of cities is characterized by three high volume cities. The air traffic between those three cities accounts for 25% of the total Canadian Airlines traffic⁷. The flights between the three main cities and flights directly connected to them represent 92% of the total air traffic. Directly connected flights are flights that fly to and from at least one of the three main cities. The remaining 8% of the traffic neither arrives nor departs from the three main cities.

Figure 3 – Airline Traffic



1.5.2 Flight Schedule

A flight schedule is a listing of flights between city pairs, and the corresponding arrival/departure times, effective date range of operation, days of the week the flight operates, and the *aircraft type* serving it.

The final schedule can be difficult to understand. Flights that routinely appear every day can unexpectedly disappear for a day, a month or indefinitely and then reappear. In addition, the flight routings ensure that a flight offering from one city to another isn't always matched by a flight offering in the opposite direction. In this sense the schedule is not predictable or "balanced". A final addition to the schedule complexity is mechanical uncertainty. Aircraft maintenance requirements cannot always be predicted. An aircraft that is scheduled for a specific route can be recalled for repairs. In order to maintain the schedule, the aircraft routings are adjusted. As a result, the specific aircraft scheduled to provide a specific flight is only tentatively known 48hrs in advance.

⁷ based on 1999 operating statistics

1.5.3 Equipment Manifest

The entire equipment manifest is determined by a combination of the meal service, galley configuration and the aircraft providing the flight. The meal service variation on Canadian Airlines flights is impressive. There are over 140 different meal service types in the 2000 meal schedule. The meal service types represent different combinations of meal type⁸, tray type, china type, and type of service accompanying the meal. This variation is accompanied by over 70 unique galley configurations. The galley configurations are a function of aircraft, seating configuration of the aircraft⁹, origin/destination¹⁰, and assigned galley code¹¹. Finally, not every aircraft of the same type is exactly the same. There are minor variations from aircraft to aircraft that affect a small quantity of components loaded onto aircraft¹². The entire process of determining the equipment manifest is summarized in Figure 4.

1.5.4 Equipment Inventories

The inflight service equipment is stored at the caterer's flight kitchen while not in use. It is the responsibility of the caterer to report inventory levels when requested, and these numbers are confirmed via periodic flight kitchen audits. The typical amount of inventory that CAI plans to have available at a given flight kitchen is approximately 36hrs worth of operating equipment. In other words, each flight kitchen is allocated enough inventory of equipment so that it can continue its operations with no incoming equipment for approximately 36hrs.

It is important to remember for every flight into a station, one must leave. Therefore, each station is constantly "shipping" and "receiving" equipment.

1.5.5 Passenger Demand

Passenger demand on any flight is a difficult quantity to predict. The volatility in passenger demand is observed in two separate spheres. First, the passenger demand varies significantly from flight to flight. Second, the passenger demand on a particular flight varies considerably in the hours leading up to departure.

Comparing passenger demand across flights can be accomplished by examining the population of observed passenger demands, and the passenger demand on city pair flights. The previous three years of CAI flight data demonstrate that passenger demand exhibits an hour of the day effect, a day of week effect and a month of year effect. The amplitude of variation is largest in the hour of day effect, followed by month of year and finally day of week. An examination of the city pair relationships, reveals that the demand difference is non zero and seasonal. Finally, the observed patterns in the difference of city pair demand are difficult to characterize.

The anticipated passenger demand varies considerably during the final hours prior to departure. The characterization of this change was one component of a thesis conducted by Jason Goto[6]. The results of his work demonstrate a passenger demand that is uncertain and volatile in the hours leading up to departure.

⁸this is a general grouping such as hot dinner, cold lunch, cold breakfast

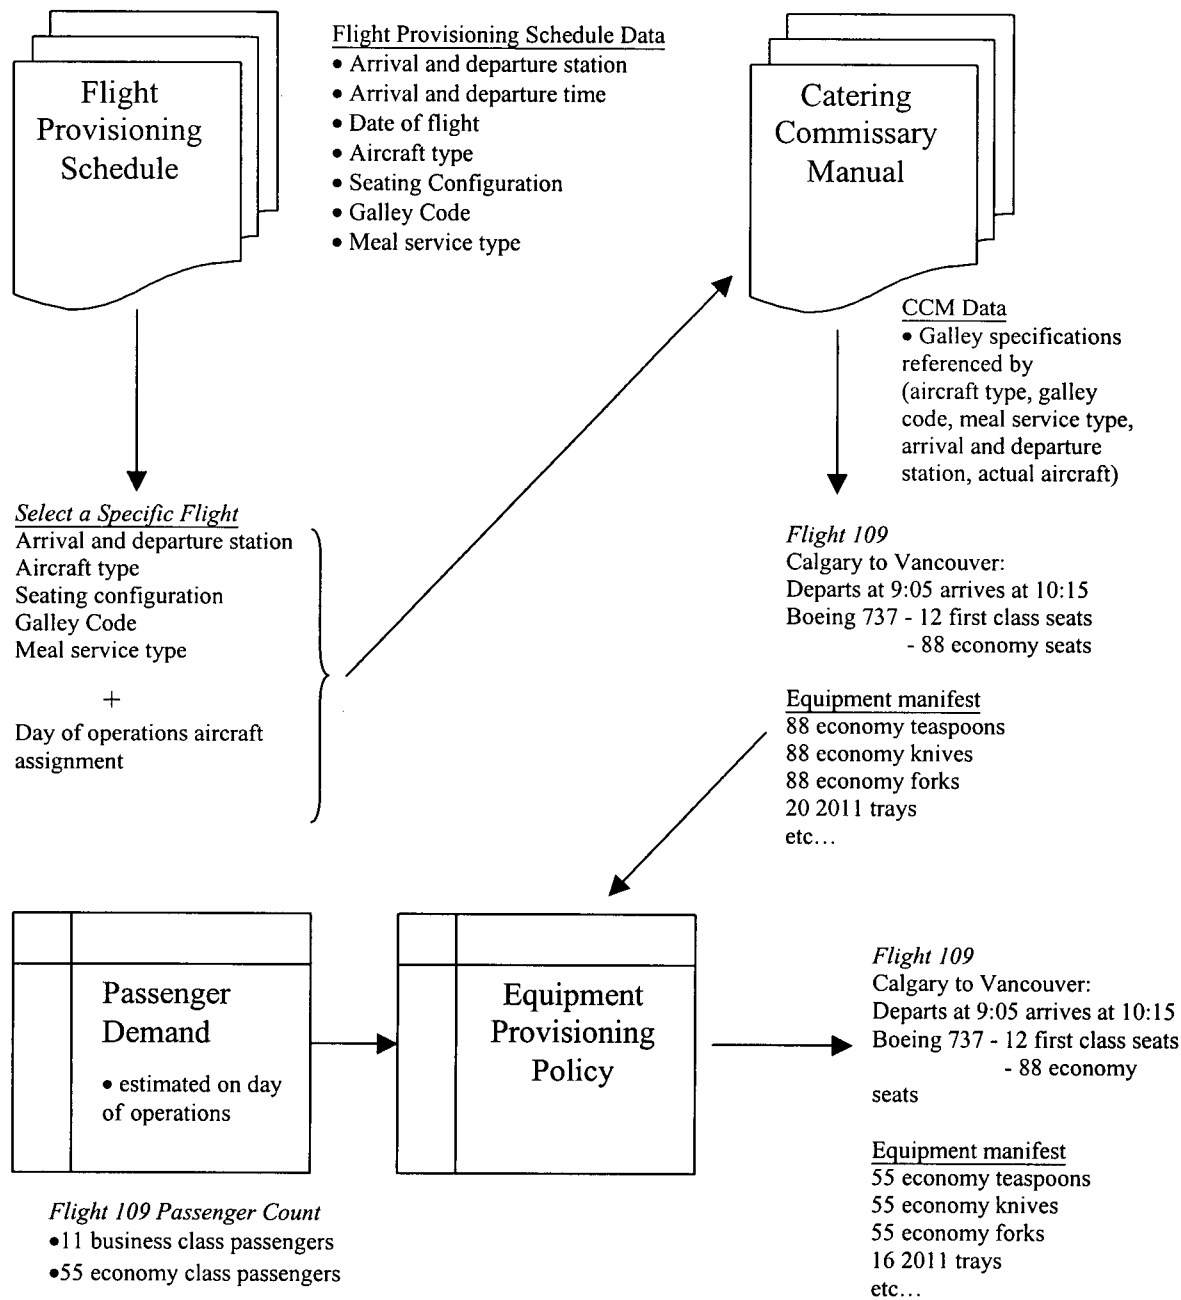
⁹ The business class section of some aircraft can be removed and replaced with additional economy class seating

¹⁰ The equipment manifest is a function of destination, for example when travelling to the Pacific Rim, an Asian tea trolley is substituted for the regular beverage trolley. Another example is the increased provisioning of alcoholic beverages on flights to certain destinations

¹¹ each flight has a particular galley code assigned to it, this is a high level code and gives an indication to the kinds of equipment to be loaded, but does not allow the identification of a specific equipment manifest

¹² Specifically, the aircraft type affects oven types and contour carriers

Figure 4 – Getting the Equipment Manifest



1.6 Data Availability

The two main data sources used in the Inflight Service department prior to this project were the Meal Pages and the Catering Commissary Manual. Both documents existed in an “electronic paper” format, meaning the data were stored in desktop publishing software, as it would look in the final printed reports.

1.6.1 Meal Pages

The Meal Pages is a meal service and equipment provisioning schedule. It is the result of the application of meal service selection rules to the flight schedule. It contains all the original flight schedule information along with additional information such as the type of meal service, high-level galley code, miscellaneous service items and additional comments for the caterers.

1.6.2 Catering Commissary Manual

The Catering Commissary Manual is a document that contains all of the galley configurations. This includes itemized lists of the types and quantities of equipment that comprise a galley. In addition, it provides assembly and loading instructions, complete with diagrams of the galleys and aircraft galley locations. This document is used in conjunction with the meal pages by over thirty flight kitchens that service CAI flights

2.0 Literature Review

There has been substantial investigation into the development of inventory management systems. The basic methods for managing inventory include the Fixed Order Quantity (FOQ) models, Economic Order Quantity (EOQ) models, and the ABC classification system. These models, with the exception of the ABC system, address inventory control when there is a clear hierarchy of inventory flow; where there is a clearly identified production center or distributor that fulfills the orders of retailers. Inventory levels are observed at all levels in the hierarchy, and external demand is observed at the retailer level of the hierarchy. The ABC model dictates which components of the inventory stock to actively manage. The model addressed by this thesis is unique in that there is no hierarchy, demand is only observed from other warehouses, and the inventory is not consumed rather it is reused.


The simplest inventory management policy is a stationary policy where a fixed quantity of inventory arrives at a warehouse at fixed intervals. The inventory levels at the warehouse are depleted by some external demand. In practice this type of policy would be readjusted on a regular basis, either adjusting the magnitude of shipments or adjusting the period between shipments.

An improvement to the stationary policy is the FOQ or fixed order quantity model. In this model a fixed amount of inventory is ordered whenever the inventory level falls below a re-order level.

An improvement to the FOQ model, would allow order quantities to vary. EOQ or economic order quantity can be used to determine the optimal amount of inventory to order. Standard assumptions for an EOQ model include: demand certainty, constant demand, no stock outs, and a constant lead time. The EOQ model can be extended to account for variable lead times and variable customer demand.

Finally, the ABC model categorizes the types of inventory by the costs incurred. Type A inventory is high cost impact items, meaning high unit cost and low volume, high unit cost high volume or low unit cost and high volume. Type B and C contain lower cost inventory items. The results of the ABC analysis indicate which inventory items to manage aggressively, and which inventory items to put less emphasis on. In practice, this method would be combined with an inventory control model, focusing mainly on type A inventory items.

Figure 5 – Traditional Hierarchal Inventory Model



Typical research topics in this area include examining the optimality of inventory policies applied to inventory models of this form.

D. Atkins and P. Iyogun [1] investigated lower bounds on performance in coordinated multi-product inventory systems. In this type of inventory system, orders are placed for multiple inventory items. The question to be answered is how much of each type of inventory should be requested at each order point. Given that there are many possible solutions to an inventory problem of this type, it becomes necessary to define the optimal solution so that the proposed solutions can be compared to the performance of the optimal solution. However, the lower bound is dependent upon the model formulation. For example, standard assumptions include deterministic demand, no backlogging, instantaneous delivery and defined systems costs. If any of these assumptions are changed the new formulation must be assessed.

Many papers have been written on the effectiveness of lot-sizing policies. One of the papers by J. Mitchell [2] extended previous results from R. Roundy. The parameters of the model solved include an N-retailer inventory system with no backlogging, constant demand, with setup and storage costs at each facility. This work proved the 98% effectiveness of optimal power-of-two policy.

M. Fu [3] developed sample path derivative estimates of performance with respect to s and S for the (s, S) policy. The result of this work was sample path derivative estimators for average inventory, back order and cost per period for an (s, S) inventory policy. The estimators can be used for sensitivity analysis and gradient-based optimization techniques.

F. Chen [4] examined the effectiveness of stationary policies. The advantage of a stationary policy is the predictability of the impact to operations. The disadvantage is that stationary policies can be costly. In this paper, the effectiveness of a stationary policy used in a multi

echelon system with deterministic demand and backlogging was determined. The effectiveness of the stationary policy is 70% for both the multistage serial system and the one warehouse multi retailer system. It should be noted that in a randomly generated set of 1000 stationary policies the average effectiveness was 99%. In addition, it was demonstrated that integer ratio policies are only 70% effective.

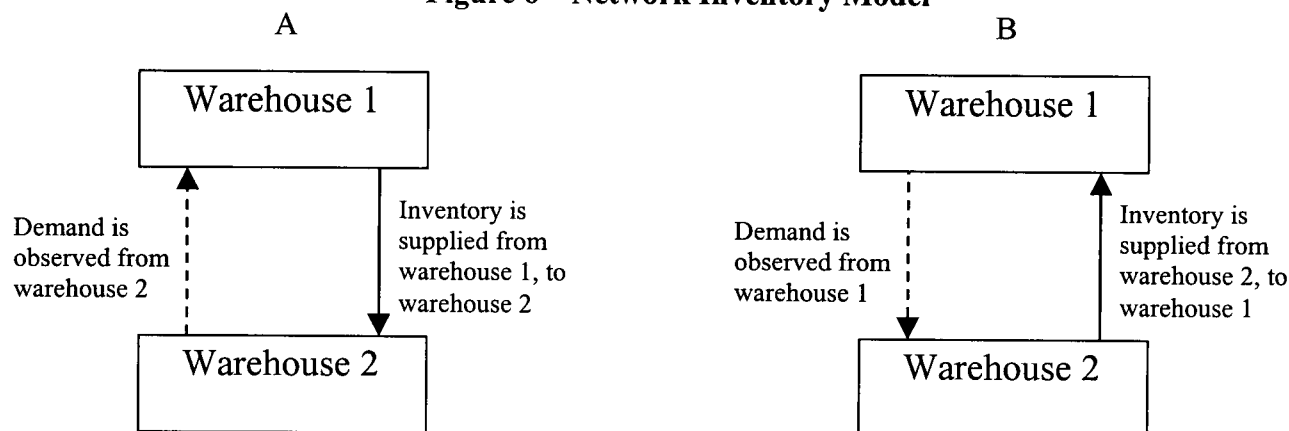
R. Anupindi and S.Tayur [5] developed a simulation model to manage a single stage that produces multiple items. The problem solved was that of production scheduling to achieve service level targets. The model assumes a cyclic schedule that determines production sequence of items, and the number of times an item is produced in a cycle. Given the schedule, the production of each item is determined by an (s,S) inventory policy. The demand for each type of product is random and the inter arrival times between demands are also random. Finally, a simulation was developed to determine “good” values of S and s.

As mentioned above the types of problems solved are all very similar. There is a warehouse that supplies external demand and places orders for new inventory. The demand is either deterministic or stochastic. The reorder quantities are either fixed or variable. The reorder times are either fixed or variable. Some models include backlogging, others do not. The model can be extended to include a production center, other warehouses and downstream warehouses. However, the basic model is inventory is provided by some source and consumed by some sink.

The problem being addressed in this thesis has many similarities to the traditional inventory models. But, there are three key distinctions:

1. The inventory system is composed of a network of warehouses
2. Demand is only observed between warehouses
3. Inventory that is “consumed” from one warehouse is “provided” to another.

Figure 6 – Network Inventory Model

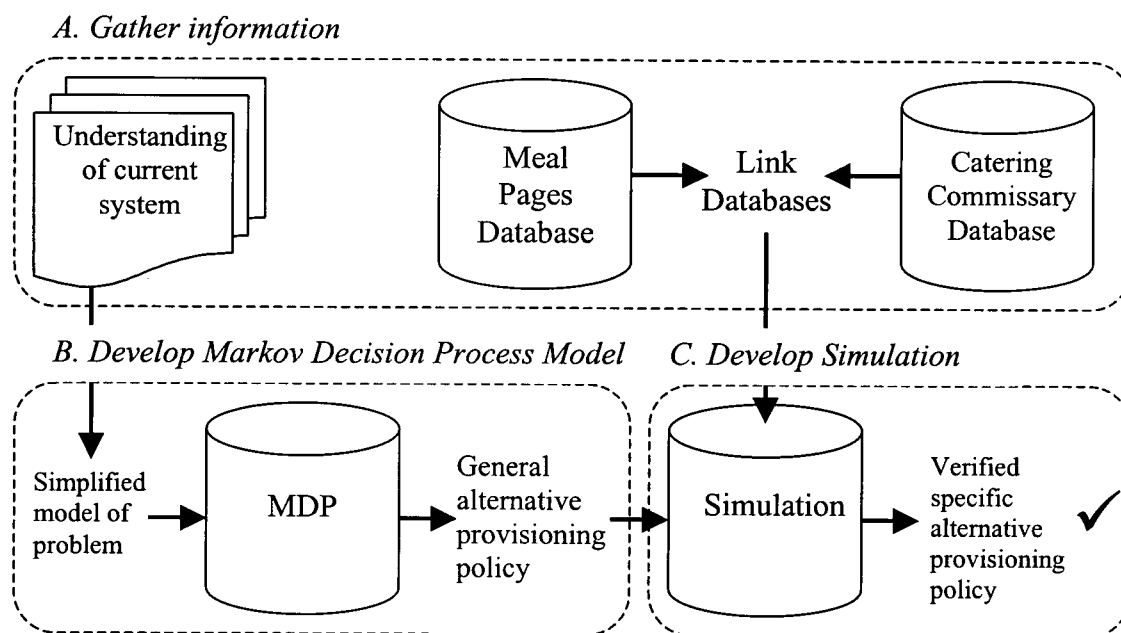


Let us consider two warehouses in this type of inventory model. In Figure 6 under heading A we see that warehouse 1 is supplying inventory to warehouse 2 according to some demand observed between warehouse 1 and 2. Under heading B we see that the opposite occurs. Therefore we see that there is no inventory hierarchy – inventory flows from warehouse to warehouse in circles, rather than down through the inventory system.

3.0 Methodology

Before appropriate modeling techniques could be selected, the operations of inflight service had to be understood. After the information gathering phase was completed it was apparent that a Markov Decision Process (MDP) model was a natural fit to the problem. An MDP analysis would generate an optimal equipment provisioning policy and permit an analysis of a generalized problem. From the optimal policy it would be possible to devise a suitable alternative equipment provisioning policy. Unfortunately, due to the heavy computational burden inherent in MDP analysis, the complete CAI network could not be modeled. Therefore, the results of the MDP analysis were used in the development of a simulation model. The brute force approach of simulation techniques permitted an in-depth analysis of the complete CAI system. The simulation was based upon operating data from CAI. Consequently it accurately approximated the actual system under study. The result of the simulation was an alternative equipment provisioning policy that was formulated using the most information available. The relationships between the three phases are represented in Figure 7.

Figure 7 – Pictorial Representation of Project Plan



3.1 Gather Information

The purpose of this phase of the project was to develop a conceptual framework of Inflight Service operations. This included understanding the entire process of provisioning equipment, from scheduling to fulfillment. In addition, we identified the data requirements of the project and took appropriate measures to ensure that the data requirements were fulfilled. The results of this phase are presented in the previous section, 1.0 Background.

3.1.1 Develop a Conceptual Framework

In order to understand the operations in Inflight service, we spent a significant amount of time working on project related activities at the client site. This allowed a daily interaction with Inflight service personnel, access to department meetings, tours of catering operations and an understanding of computer systems in use. Many areas of CAI were consulted including

inventory management, catering operations, and inflight service. In addition we had the opportunity to meet with SABRE, the developers of the inventory management software available at CAI, as well as to discuss inventory management with CAI's partner airline AMR. This resulted in a rigorous understanding of the operations.

3.1.2 Access to Required Data

Generally, data management at companies is directly linked to operational activities. As a result, the data are in a format that is suited to operations. This can make accessing the data in a manner other than intended, extremely difficult. At CAI, the data in the meal pages and the catering commissary manual presented access issues.

In order to construct a complete equipment flow model of the CAI system, data from both the Meal Pages and from the Catering Commissary Manual were required. These information sources are described in section 1.6. With the help of CAI personnel and COE associate Sean Baird, the design and implementation of the **Meal Pages database** and the **Catering Commissary database** was completed. These two databases have been integrated into CAI's operations. The benefit to CAI was the ability to streamline the main operations associated with these two databases as well as automate some ancillary activities. The benefit to this project was the direct access to required operational information that, as a result of being implemented, had the most current and up to date data.

3.2 Develop Markov Decision Processes

A standard problem to be addressed in inventory theory is the existence and performance of an optimal policy. Without an understanding of the structure of an optimal policy, it is impossible to gauge the performance of other policies precisely. MDP models can be used to identify and evaluate the performance of an optimal policy. MDP models, and dynamic programming in general, do have one drawback - the computational burden required to solve the model. Given that the network to be solved contains 31 nodes, it is unlikely that a 31 node MDP model can be formulated and solved in reasonable amount of time (or even lifetime). Consequently, the MDP models usefulness is limited to the analysis of simplified versions of the CAI network.

Three separate models of the Canadian Airlines operational system were formulated. Each of the three models gave a different perspective on what the alternative policy should be¹³. The results of the three separate models were combined to form **an alternative equipment provisioning policy** for use in the simulation. In addition the use of MDP allowed for a comparison of the alternative policy to the optimal policy for each model.

3.2.1 What is MDP?

A decision can be thought of as a choice of action based upon available data, to affect a future outcome. Generally, the action and resulting outcome are associated with a reward or cost. In special cases there are a sequence of decisions to be made. In addition, the outcome of a specific action may be deterministic or stochastic. A Markov Decision Process is a model of a stochastic sequential decision process. The defining characteristic of a Markov Decision Process is the fact that "the history of the problem has no effect on future decisions." (Puterman 1994) All that is required to make a future decision is the current state; the path to that state is irrelevant.

¹³ Each model will have an optimal policy *for that specific model*, and the policy is not necessarily optimal for the actual Canadian Airlines system. The results of three separate models combined demonstrate the form of an alternative provisioning policy.

3.2.2 MDP and Equipment Provisioning

Equipment provisioning involves the decision of how much *excess* equipment to load onto aircraft. The amount of excess equipment loaded onto aircraft will have two effects. First, the inventory levels at the departure station will decrease while the inventory levels at the arrival station will increase. Second, the loading of excess equipment will incur costs. We are not just concerned with the provisioning of equipment to a single flight, but all flights. As such we are faced with a sequence of similar decisions. The final elements of the problem arise from the calculation of the future inventory of equipment at a station. The future inventory level is equal to the current inventory level plus the amount of equipment loaded onto arriving flights, minus the amount of equipment loaded onto departing flights. If the equipment is loaded onto the aircraft in relation to passenger demand, then the equipment flow is stochastic. Finally, when making the decision, we are not concerned with how the current inventory level came to be. We are only concerned with current value itself.

3.3 Develop Simulation

The purpose of this phase of the project is to validate the effectiveness of an alternative provisioning policy suggested by the MDP analysis. Due to the difficulty involved in modeling the complete CAI system as an MDP, the decision was made to use another operations research tool, simulation. Simulation is used when the real world system under study is too complex to model using other tools.

3.3.1 What is Simulation?

A simulation can be of many types, stochastic or deterministic, discrete event based or continuous, however the fundamentals of simulation are generally the same. Generally, we are interested in simulating the operations of real world facilities and processes. "The facility or process of interest is usually called a system, and in order to study it scientifically we often have to make assumptions about how it works. These assumptions, which usually take the form of mathematical or logical relationships, constitute a model that is used to try and gain some understanding of how the corresponding system behaves." [7]

The main distinction between Simulation Models and MDP models is that simulation models can only test policies; they cannot be used to determine policies. MDP analysis will define a policy.

3.3.2 Simulation and Equipment Provisioning

For this project, simulation was used to model the complete 31 node CAI network. The alternative provisioning policy generated from the MDP analysis was modeled using simulation. This allowed the incorporation of system specific details, such as passenger demand by market segment, multiple aircraft types, and a complicated flight schedule. The simulation model was constructed in MS EXCEL, incorporating relevant operational data from CAI.

4.0 Markov Decision Process: Model Formulations

A Markov Decision Process model is composed of states, actions, assumptions, rewards, and state transitions. States represent the current condition of the model. For example, in these models one of the factors that is of interest is inventory level, and so the current level of inventory is one portion of the state. Actions are the means by which a decision is enacted in the model. In this case we would like to decide how much equipment to load onto an aircraft, and as such the action is how much equipment is loaded onto an aircraft. The assumptions simplify the real world system under study. The rewards or costs are given as a result of achieving a particular state. Finally, the state transitions define how to determine a future state given the current state and the current action.

This section begins with a description of the states and actions of the models. Then the assumptions used in the three different model formulations are presented. Next, the basis for the cost function is developed. Finally, each model is presented in sequence, starting with the 1-node and ending with the 3-node.

4.1 States and Actions

The state at time t is defined to be the quantity of the vector of inventory imbalances at each station at time t , and the current value of the passenger demand for all flights at time t . For example, in the one node model, there is one inventory imbalance for each time t , and one passenger demand for each time t . Therefore, the state is defined by the value of the inventory imbalance at time t and the demand on the flight at time t .

Inventory imbalance is defined as the difference between the inventory level and the inventory allotment¹⁴. An inventory imbalance can be positive or negative, and it measures the deviation of the inventory level from the specified allotment. The use of inventory imbalance allows us to ignore the cumbersome nature of dealing with inventory allotments and inventory levels.

The action set at time t is defined to be the vector of excess equipment movements for each flight at time t . Each action at time t is equal to the amount of excess equipment to be loaded onto each flight at time t . The value of an action is restricted to be between 0 and the aircraft capacity minus observed passenger demand. This means that the total equipment loaded must at least be equal to passenger demand and must at most be equal to the aircraft capacity.

4.2 Model Assumptions

The models, as formulated, are based on several assumptions. Each assumption is listed below, and then described in detail in the following sections:

- 1) One single type of aircraft
- 2) One class of seat
- 3) C aggregate seats on aircraft
- 4) Simplified flight schedule
- 5) Sufficient inventory to satisfy demand out
- 6) Single type of equipment
- 7) Approximate passenger demand distributions

¹⁴ The allotment is determined outside of the scope of this model

4.2.1 One Single Type of Aircraft

For simplicity, we have assumed that the aircraft in the model are all of the same type. In the CAI system there are over 9 different types of aircraft.

4.2.2 Aircraft Have One Class of Seats

Most aircraft in the CAI system have two classes of seats, business and economy. In the MDP models no distinction is made between the classes of seats.

4.2.3 C aggregate seats on aircraft

All aircraft at CAI have a combined seating capacity over 100 seats. These seats are treated as C aggregate seats. This is required to ensure that the formulated models are solvable.

4.2.4 Simplified Flight Schedule

The actual flight schedule is a complicated web of flights between cities. Several flights can arrive in a city before a flight departs. In the MDP model we assume that for each epoch, one flight arrives and departs at each station for every other station in the model. In the 1-node (station) model, one flight arrives and one flight departs from the one station. In the 2-node model, one flight arrives and one flight departs from each of the two stations. In the 3-node model, two flights arrive and two flights depart from each of the three stations. This leads to a flight schedule that does not allow aircraft to build up at a station. Also, the simplified schedule has balanced flights between city pairs- for every flight going from A to B there is a flight going from B to A.

4.2.5 Sufficient Inventory to Satisfy Demand Out

As previously discussed, each station has an inventory allotment equal to approximately 36hrs of operations. During a period of a week it is possible that sufficient differences between in inbound and outbound passenger demand could exist to exhaust a stations inventory. However, such an imbalance would be observable and rectifiable before it occurred. For example, with 18hrs of inventory remaining a decision could be made to rebalance stock from other cities, and this inventory would arrive within the eighteen remaining hours. Therefore, it is assumed that sufficient inventory is available to satisfy passenger demand on any aircraft. The formulation of the model incorporates bounds on deviation from inventory allotments and enforces those bounds with automatic reshipments.

4.2.6 Single Type of Equipment

An average equipment manifest will contain many types of equipment in varying quantities. For the MDP model we will represent the maximum possible equipment to be loaded as C units of equipment. This representation is a consolidation of equipment into a single type of equipment, as well as an aggregation of the quantity of equipment into C groups.

4.2.7 Simplified Passenger Demand

As described in the previous section, the passenger demand is difficult to characterize and predict. It varies depending on several factors, including market segment (origin and destination), time of year and time of day. In the MDP model, we have assumed all passenger distributions to be triangular (0.4C, 0.7C, 1.0C). Varying demand patterns by region will be addressed in the simulation model.

4.3 Model Cost Function

The cost function is what motivates the model to choose actions. When evaluating a series of available choices, the model will compute the corresponding costs and choose the action associated with the lowest cost. It is perhaps one of the most important components of the model. To define the cost function we must first revisit the project objective:

To reduce the amount of excess equipment boarded onto aircraft and simultaneously ensure that significant equipment flow imbalances are mitigated

The objective statement describes two measures of success, equipment flow imbalance and excess equipment loaded. The cost function incorporates both of these, and one more- the rebalance action. The complete cost function will be presented below, followed by a description of each component. Finally the section concludes with an interpretation of cost multipliers

4.3.1 The Complete Cost Function

The complete cost function is given by:

$$r_t(s_t, a_t) = K_1 * |S_t| + K_2 * a_t + K_3 * R_t$$

The cost function will penalize inventory imbalances, movements of excess equipment and

r_t = reward at time t

K_1 = inventory imbalance penalty multiplier

K_2 = excess equipment movement penalty multiplier

S_t = inventory imbalance at time t

a_t = amount of excess equipment loaded in period t

R_t = rebalance action in period t (enforces inventory imbalance limits, discussed later)

K_3 = equipment rebalance penalty multiplier

rebalance actions. When comparing costs generated from the cost function, it is the lower cost that is preferred.

4.3.2 Inventory Imbalance Penalty

The inventory imbalance can be measured by subtracting the current inventory level from the inventory allotment. Alternatively, if the system is assumed to start with a zero inventory imbalance, then the cumulative sum of equipment flows to and from a station will determine the inventory imbalance.

$$\sum_{i=1}^t X_t^I = I_t - I^A = S_t$$

X_t^I = equipment flow imbalance in period t

I_t = inventory level in period t

I^A = inventory allotment

S_t = inventory imbalance in period t

We can change the relative importance of the deviation from inventory allotment by multiplying it by a constant. Finally, we can represent the contribution of the inventory imbalance penalty to the cost function as:

$$K_1 * |S_t|$$

K_1 = inventory imbalance penalty multiplier

S_t = inventory imbalance in period t

4.3.3 Movement of Excess Equipment Penalty

The excess equipment boarded onto an aircraft can be measured by subtracting the demand for equipment on an aircraft (passenger demand) from the total amount of equipment loaded.

$$a_t = X_t^L - d_t$$

a_t = amount of excess equipment loaded in period t

X_t^L = total equipment loaded in period t

d_t = demand for equipment in period t

The demand for equipment is equal to the passenger demand on the flight, d_t . The excess quantity of equipment is defined by the amount of additional equipment we choose to load, a_t . We wish to minimize the amount of excess equipment movement in every period. We can represent the contribution of excess equipment movement penalty to the cost function as:

$$K_2 * a_t$$

a_t = amount of excess equipment loaded in period t

K_2 = excess equipment movement penalty multiplier

Again, K_2 allows us to change the relative importance of this penalty.

4.3.4 Equipment Rebalance Penalty

As discussed in the assumptions section, an automatic equipment rebalance is incorporated in the model. This rebalance is activated if the inventory imbalance is permitted to grow too large and is not corrected. Operationally, this would correspond to a bulk shipment of equipment either by train, truck or aircraft. The rebalance action is directly measurable; R_t represents its value. We

wish to minimize the automatic rebalance actions. The contribution of the equipment rebalance penalty to the total cost function can be represented as

$$K_3 * R_t$$

R_t = automatic rebalance action in period t

K_3 = equipment rebalance penalty multiplier

4.3.5 K Term Interpretations

The value of the K terms in the cost function relative to each other will affect the model results. Therefore, a complete understanding of their individual significance is warranted. The K values will be discussed in the order K_2 , K_3 and K_1 .

The value of K_2 has a direct interpretation in the real world system under study. If the objective function is measured in dollars, then the value of K_2 is the cost of supporting the logistics to provide a unit of equipment to the aircraft. Specifically, this encompasses such costs as handling charges, cleaning charges, fuel burn and expected breakage.

The value of K_3 also has a direct interpretation in the real world system under study. The value of K_3 is the cost of all activities required to rebalance equipment. This includes packing, transportation, unpacking, cleaning, and additional breakage costs.

The value of K_1 does not have a direct interpretation in the real world system under study. Its model representation is a penalty for deviation from allotment. However, its real world interpretation does not represent a completely identifiable cost. There are costs that could be associated with inventory imbalances, for example square footage at a caterer's flight kitchen. But assuming that there are sufficient units of equipment to satisfy the demand requirements, what is the value of additional units of equipment? Operationally, the only "cost" is the probability of stock out, and this will be discussed in greater detail in the simulation section. For the purposes of the MDP model, the cost function encourages lower inventory imbalances - that is what we need it to do.

4.3.6 Cost structure Representations

Choosing specific ratios for the relative values of the K terms allows us to identify and characterize the resulting cost structure. A list of cost structures used in the MDP analysis is given in Table 1. This is known as a two way factorial design for K_1 and K_2 .

Table 1 – MDP Model Cost Structures

<i>Cost Structure Name</i>	K_1	K_2	K_3
113	1	1	3
123	1	2	3
133	1	3	3
213	2	1	3
223	2	2	3
233	2	3	3
313	3	1	3
323	3	2	3
333	3	3	3


These specific cost structures were chosen to allow the identification of trends with the change of specific K terms. For example, examining the results of the 113, 123 and 133 cost structures will reveal the influence of increasing and decreasing K_2 . The cost structures provided allow for easy analysis of various relative magnitudes of K_1 and K_2 .

K_3 was kept constant for two reasons. The first is the fact that moving a unit of equipment in a rebalance action is at least as expensive as moving a unit of equipment as excess equipment. If this were not the case then CAI would not bother to ship excess equipment at all. The second is that this choice reduces the number of analyzed models by a factor of three, allowing us to focus on the more industrially relevant questions of this investigation.

4.4 Model 1: 1-Node Model Formulation

The simplest representation of the CAI network is a collection of cities that act independently of one another. This framework leads to a 1-Node model of the problem. In this case we are modeling the network in the presence of demand and inventory information asymmetry, specifically, we only have access to information from station A when making the decisions. This model will yield a solution that optimizes only the station being analyzed without any consideration of the network around it. A graphical representation of the problem is given in Figure 8.

Figure 8 – 1 Node MDP model

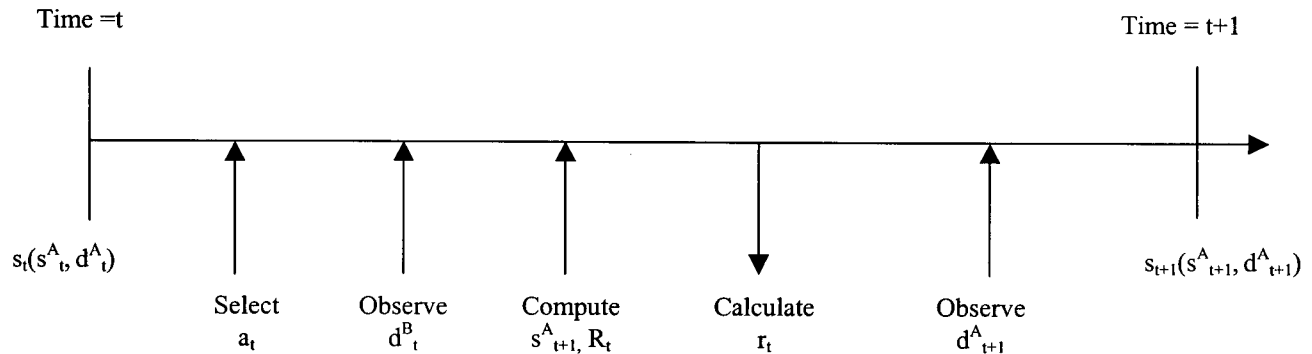


4.4.1 Definitions

- N = number of epochs
- D_t^A = random variable representing passenger demand out of city A
- D_t^B = random variable representing passenger demand in to city A
- d_t^A = observed passenger demand on a flight in period t out of city A
- d_t^B = observed passenger demand on a flight in period t into city A
- S^A = inventory imbalance at city A
- S^L = upper limit on absolute inventory imbalance
- a_t = amount of excess inventory to load onto aircraft from city A in period t
- K_1 = inventory imbalance penalty multiplier
- K_2 = excess equipment movement penalty multiplier
- K_3 = equipment balancing penalty multiplier
- C = aircraft passenger capacity
- R_t = equipment rebalance action in period t
- r_t = cost in period t

4.4.2 Timeline of the model

We start at time t in state s_t . The state of one node model is defined as the current value of passenger demand out of the station (d_t^A), and the current inventory imbalance at the station s_t^A . An action (a_t) is selected by using information contained in the state. The passenger demand into the station (d_t^B) is then observed. The future inventory imbalance at time $t = t+1$ can be calculated by using the current state, action and demand in. The rebalance action is determined



by comparing the calculated inventory imbalance to the automatic rebalance limit. Once the rebalance action is known, the actual future inventory imbalance can be determined. Finally a new passenger demand out is observed. The combination of the new passenger demand out and the actual inventory imbalance at time $t=t+1$ defines the future state, s_{t+1} .

4.4.3 State Space

The state space for this model is the cross product of the passenger demand out space and the inventory imbalance space. A typical state is written as $s_t=(d_t^A, s_t^A)$. The presence of demand out in the state space is peculiar. However, when we realize that we are trying to choose the amount of excess equipment to send, the three variables that we would like in a one-node problem are demand in, demand out and deviation from inventory allotment. The model is based on asymmetric demand information; demand in is not observed prior to the choice of action. Consequently, demand in is not included in the state space, but the observed quantities, demand out and inventory imbalance, are.

4.4.4 Action Space

The action space consists of all possible quantities of excess equipment that can be boarded onto a given aircraft. The action space starts at zero and stops at the capacity of the aircraft minus the demand on the aircraft. An action equal to 1 means 1 unit of excess equipment has been boarded onto the aircraft.

4.4.5 State Transitions

The transition of s^A can be broken down into two regions, (1) normal inventory accounting and (2) automatic rebalance. In the normal inventory accounting region, a future realization of inventory imbalance (s_{t+1}^A) is determined by adding s_t^A and demand in, and subtracting demand out and the current action. In the automatic rebalance region s_{t+1}^A is determined by comparing s_{t+1}^A to S^{L15} . If the absolute value of s_{t+1}^A is greater than S^L , a rebalance action is initiated and the s_{t+1}^A is set equal to zero. If the absolute value of s_{t+1}^A is less than S^L , the s_{t+1}^A is equal to the result from the normal inventory accounting region.

¹⁵ note S^L is not a decision variable, it is a model parameter

4.4.6 The Complete Formulation

$$T = \{1, \dots, N\}$$

State Space: $S = D^A \times S^A$

$$D^A = \{0, 1, \dots, C\}$$

$$D^B = \{0, 1, \dots, C\}$$

$$S^A = \{s : -S^L \leq s \leq S^L, \text{int}\}$$

Actions: $A_{d_t^A} = \{a_t\} = \{0, 1, 2, \dots, C - d_t^A\}$

Rewards: $r_t(s_{t+1}, a_t) = K_1 * |s_{t+1}| + K_2 * a_t + K_3 * |R_t|$

$$r_N(s_N) = 0$$

$$R_t = \begin{cases} 0 & \text{if } |s_t^A - d_t^A - a_t + d_t^B| \leq S^L \\ s_t^A - d_t^A - a_t + d_t^B & \text{if } |s_t^A - d_t^A - a_t + d_t^B| > S^L \end{cases}$$

$$P(D_t^A = d) = p_d$$

$$P(D_t^B = d) = q_d$$

State
Transition:

$$s_{t+1}^A = \begin{cases} s_t^A - d_t^A - a_t + d_t^B & \text{if } |s_t^A - d_t^A - a_t + d_t^B| \leq S^L \\ 0 & \text{if } |s_t^A - d_t^A - a_t + d_t^B| > S^L \end{cases}$$


Bellman Eqn:
$$V_t(d_t^A, s_t^A) = \max_{a_t} \left\{ \sum_{d_t^B} q_{d_t^B} \left[(K_1 * |s_{t+1}| + K_2 * a_t + K_3 * |R_t|) + \sum_{d_{t+1}^A} p_{d_{t+1}^A} * V_{t+1}(d_{t+1}^A, s_{t+1}) \right] \right\}$$

$$V_t(d_N^A, s_N) = 0$$

4.5 Model 2: 2-Node Model Formulation

The motivation for extending the first model is the fact that the first model has only one node. As a result its ability to capture the essential elements of a network is limited. The next logical extension is a 2-Node problem. This model is analogous to a single city pair. The 2-Node model is represented in Figure 9.

Figure 9 – 2 Node MDP Model



4.5.1 Definitions

All the definitions are included for completeness. In the 2-node model we have added two variables, S^B and a^B .

N = number of epochs

D_t^A = random variable representing passenger demand out of city A

D_t^B = random variable representing passenger demand in to city A

d_t^A = observed passenger demand on a flight in period t out of city A

d_t^B = observed passenger demand on a flight in period t into city A

S^A = inventory imbalance at city A

S^B = inventory imbalance at city B

S^L = upper limit on absolute inventory imbalance

a_t^A = amount of excess inventory to load onto aircraft from city A to city B in period t

a_t^B = amount of excess inventory to load onto aircraft from city B to city A in period t

K_1 = inventory imbalance penalty multiplier

K_2 = excess equipment movement penalty multiplier

K_3 = equipment balancing penalty multiplier


C = aircraft passenger capacity

R_t = equipment rebalance action in period t

r_t = cost in period t

4.5.2 Timeline

As in the previous model, the timeline starts with the model in state s_t at time t . The actions are selected, the rebalance quantity is determined, and then a cost is incurred. Finally, the appropriate state transition occurs moving the state to s_{t+1} . The main difference between the 2N model and the 1N model is the fact that both demands are observed prior to making a decision.



4.5.3 State Space

The state space for this model is the union of the passenger demand out space, the passenger demand in space and the inventory imbalance at station A space. Capitalizing on the fact that the model represents a closed system reduces the size of the state space. In any epoch, if we are given the inventory imbalance at one station, we know what the inventory imbalance at the other station must be.

4.5.4 Action Space

The action space for the 2-Node model incorporates an additional action. Each component of the action space is exactly the same as in the 1-node model.

4.5.5 State Transition

A future state, s_{t+1} , is dependent on two things, 1) the probability of observing particular values for d^A_{t+1} and d^B_{t+1} and 2) the total result of equipment movements in period t . The inventory accounting procedures for this model are almost identical to the one node model, except that action b is now incorporated.

4.5.6 The Complete Formulation

$$T = \{1, \dots, N\}$$

State Space:

$$S = D^A \times D^B \times S^A$$

$$D^A = \{0, 1, \dots, C\}$$

$$D^B = \{0, 1, \dots, C\}$$

$$S^A = \{s^A : -S^L \leq s^A \leq S^L, \text{int}\}$$

$$s_t^B = -s_t^A$$

Actions:

$$A_{d_t^A, d_t^B} = \{a_t^A, a_t^B\}$$

$$a_t^A = \{0, 1, 2, \dots, C - d_t^A\}$$

$$a_t^B = \{0, 1, 2, \dots, C - d_t^B\}$$

Rewards:

$$r_t(s_{t+1}, a_t) = 2 * K_1 * |s_{t+1}^A| + K_2 * (a_t^A + a_t^B) + K_3 * |R_t|$$

$$r_N(s_N) = 0$$

$$R_t = \begin{cases} 0 & \text{if } |s_t^A + d_t^B + a_t^B - d_t^A - a_t^A| \leq S^L \\ s_t^A + d_t^B + a_t^B - d_t^A - a_t^A & \text{if } |s_t^A + d_t^B + a_t^B - d_t^A - a_t^A| > S^L \end{cases}$$

$$P(D_t^A = d) = p_d$$

$$P(D_t^B = d) = q_d$$

State
Transition:

$$s_{t+1}^A = \begin{cases} s_t^A + d_t^B + a_t^B - d_t^A - a_t^A & \text{i } |s_t^A + d_t^B + a_t^B - d_t^A - a_t^A| \leq S^L \\ 0 & \text{i } |s_t^A + d_t^B + a_t^B - d_t^A - a_t^A| > S^L \end{cases}$$

Bellman Eqn:


$$V_t(d_t^A, d_t^B, s_t^A) = \max_{a_t} \left\{ 2 * K_1 * |s_{t+1}^A| + K_2 * (a_t^A + a_t^B) + K_3 * |R_t| + \sum_{d_{t+1}^A} \sum_{d_{t+1}^B} p_{d_{t+1}^A} q_{d_{t+1}^B} * V_{t+1}(d_{t+1}^A, d_{t+1}^B, s_{t+1}^A) \right\}$$

$$V_N(d_N^A, d_N^B, s_N^A) = 0$$

4.6 Model 3: 3-Node Model Formulation

The 3-node model is the smallest "network" that we can create. It consists of three city pairs. It is directly analogous to the three high volume cities of the CAI network. The 3-node model is represented in Figure 10.

Figure 10 – 3 Node MDP model



As can be seen, the movement from a 2-Node model to the 3-Node model involves a significant increase in the size and complexity of the model.

4.6.1 Definitions

N = number of epochs

D_t^1 = random variable representing passenger demand from station A to station B

D_t^2 = random variable representing passenger demand from station A to station C

D_t^3 = random variable representing passenger demand from station B to station A

D_t^4 = random variable representing passenger demand from station B to station C

D_t^5 = random variable representing passenger demand from station C to station A


D_t^6 = random variable representing passenger demand from station C to station B

d_t^1 = demand from station A to station B in period t

d_t^2 = demand from station A to station C in period t

d_t^3 = demand from station B to station A in period t
 d_t^4 = demand from station B to station C in period t
 d_t^5 = demand from station C to station A in period t
 d_t^6 = demand from station C to station B in period t
 a_t^1 = amount of excess inventory to load from station A to station B in period t
 a_t^2 = amount of excess inventory to load from station A to station C in period t
 a_t^3 = amount of excess inventory to load from station B to station A in period t
 a_t^4 = amount of excess inventory to load from station B to station C in period t
 a_t^5 = amount of excess inventory to load from station C to station A in period t
 a_t^6 = amount of excess inventory to load from station C to station B in period t
 S^A = inventory imbalance at station A
 S^B = inventory imbalance at station B
 S^C = inventory imbalance at station C
 S^L = inventory imbalance limit
 R_t = equipment rebalance action in period t
 C = aircraft capacity
 K_1 = deviation from allotment penalty multiplier
 K_2 = excess equipment movement penalty multiplier
 K_3 = rebalance inventory multiplier

4.6.2 Timeline



As in the two previous models, the timeline starts with the model in state s_t at time t . The actions are selected, the rebalance quantity is determined, and then a cost is calculated. Finally, the appropriate state transition occurs moving the state to s_{t+1} .

4.6.3 State Space

The state space for this model is the union of the six passenger demand spaces and two inventory imbalance spaces (the third inventory imbalance is defined by the first two inventory imbalances).

4.6.4 Action Space

The action space for this model is the set of six actions. Each action can take a value between 0 and the aircraft capacity minus the corresponding passenger demand.

4.6.5 State Transition

Passenger demand is a random variable; as such a future realization of a state is dependent on the future inventory imbalance and the probability of obtaining a particular set of demands $\{d_{t+1}^1, d_{t+1}^2, d_{t+1}^3, d_{t+1}^4, d_{t+1}^5, d_{t+1}^6\}$.

The process of summing equipment flows into and out of a station is the same, but the process for determining the rebalance action is significantly different. In the 3-Node model, we must know which station has an excessive inventory imbalance and which of the two stations (perhaps both) to rebalance the inventory with. The method for determining the rebalance action and the stations participating in the rebalance action is provided in the complete formulation section.

4.6.6 Complete Formulation

The complete formulation is provided on the following page

$$T = \{1, \dots, N\}$$

State Space:

$$\begin{aligned} S &= D^1 \times D^2 \times D^3 \times D^4 \times D^5 \times D^6 \times S^{AB} \\ S^{AB} &= \{(s^A, s^B) : -S^L \leq s^A + s^B \leq S^L\} \\ -S^L &\leq s^A \leq S^L \\ -S^L &\leq s^B \leq S^L \\ s_t^A + s_t^B + s_t^C &= 0 \\ s^A, s^B &\text{ integer} \end{aligned}$$

Actions:

$$A_s = \{a_t^1, a_t^2, a_t^3, a_t^4, a_t^5, a_t^6\}$$

$$a_t^i = \{0, 1, 2, \dots, d_t^i\}$$

Rewards:

$$r_t(s_{t+1}, a_t) = K_1 * (|s_{t+1}^A| + |s_{t+1}^B| + |s_{t+1}^A + s_{t+1}^B|) + K_2 * (\sum_{i=1}^6 a_t^i) + K_3 * |R_t|$$

$$r_N(s_N) = 0$$

$$P(D_t^i = d) = p_d^i$$

$$c_{t+1}^A = s_t^A - d_t^1 - d_t^3 + d_t^2 + d_t^4 - a_t^1 - a_t^3 + a_t^2 + a_t^4$$

$$c_{t+1}^H \geq c_{t+1}^M \geq c_{t+1}^S$$

$$c_{t+1}^B = s_t^B - d_t^2 - d_t^5 + d_t^1 + d_t^6 - a_t^2 - a_t^5 + a_t^1 + a_t^6$$

s = small
m = medium
h = high

$$c_{t+1}^C = -s_t^A - s_t^B - d_t^4 - d_t^6 + d_t^3 + d_t^5 - a_t^4 - a_t^6 + a_t^3 + a_t^5$$

$$R_t = 0 \quad \text{if} \quad c_{t+1}^H \leq S^L \quad \text{and} \quad c_{t+1}^S \leq -S^L$$

$$R_t = |c_{t+1}^H| \quad \text{if} \quad |c_{t+1}^M| \geq S^L \quad \text{and} \quad |c_{t+1}^H| \geq |c_{t+1}^S|$$

$$R_t = |c_{t+1}^S| \quad \text{if} \quad |c_{t+1}^M| \geq S^L \quad \text{and} \quad |c_{t+1}^H| < |c_{t+1}^S|$$

$$R_t = |c_{t+1}^S| \quad \text{if} \quad |c_{t+1}^M| < S^L \quad \text{and} \quad |c_{t+1}^H| \geq |c_{t+1}^S|$$

$$R_t = |c_{t+1}^H| \quad \text{if} \quad |c_{t+1}^M| < S^L \quad \text{and} \quad |c_{t+1}^H| < |c_{t+1}^S|$$

State

Transition:

$$s_{t+1}^H = c_{t+1}^H; s_{t+1}^M = c_{t+1}^M; s_{t+1}^S = c_{t+1}^S \quad \text{if} \quad c_{t+1}^H \leq S^L \quad \text{and} \quad c_{t+1}^S \leq -S^L$$

$$s_{t+1}^H = s_{t+1}^M = s_{t+1}^S = 0 \quad \text{if} \quad |c_{t+1}^M| \geq S^L \quad \text{and} \quad |c_{t+1}^H| \geq |c_{t+1}^S|$$

$$s_{t+1}^H = s_{t+1}^M = s_{t+1}^S = 0 \quad \text{if} \quad |c_{t+1}^M| \geq S^L \quad \text{and} \quad |c_{t+1}^H| < |c_{t+1}^S|$$

$$s_{t+1}^H = -c_{t+1}^M; s_{t+1}^M = c_{t+1}^M; s_{t+1}^S = 0 \quad \text{if} \quad |c_{t+1}^M| < S^L \quad \text{and} \quad |c_{t+1}^H| \geq |c_{t+1}^S|$$

$$s_{t+1}^H = 0; s_{t+1}^M = c_{t+1}^M; s_{t+1}^S = -c_{t+1}^M \quad \text{if} \quad |c_{t+1}^M| < S^L \quad \text{and} \quad |c_{t+1}^H| < |c_{t+1}^S|$$

Bellman Eqn:

$$V_t(d_t^i, s_t) = \max_{a_t} \left\{ K_1 * (|s_{t+1}^A| + |s_{t+1}^B| + |s_{t+1}^A + s_{t+1}^B|) + K_2 * (\sum_{i=1}^6 a_t^i) + K_3 * |R_t| \right. \\ \left. + \sum_{d_{t+1}^1} \sum_{d_{t+1}^2} \sum_{d_{t+1}^3} \sum_{d_{t+1}^4} \sum_{d_{t+1}^5} \sum_{d_{t+1}^6} p_d^1 p_d^2 p_d^3 p_d^4 p_d^5 p_d^6 * V_{t+1}(d_{t+1}^i, s_{t+1}) \right\}$$

$$V_N(d_N^i, s_N) = 0$$

5.0 Markov Decision Process: Model Results

The goal of the MDP analysis is to generate an alternative equipment provisioning policy, and to determine the performance of that policy. To achieve the goal we must first develop an optimal policy. Then an alternative policy can be developed using the optimal policy as a basis. Next, the current policy, the deadhead policy must be developed. Finally, we can use the results of the optimal policy and the deadhead policy to quantify the performance of the alternative policy.

The results of the MDP analysis can be summarized as:

1. The alternative provisioning policy suggested by the MDP analysis is the Single Reference (SR) policy.
2. As the size of the network increases, the SR policy performance also increases
3. The SR policy information requirements are flight demand and inventory imbalance of the station, the DH policy only requires flight demand

The remainder of this section is devoted to presenting the results in greater detail. The results of each model analysis are presented in order of increasing complexity, starting with the 1 Node and ending with the 3 Node. In this section, many policies are presented that have similar names. For clarification, these policies are listed on the following page in Table 2. Lastly, throughout the analysis many references will be made to cost structures. As discussed in section 4, the cost structure is defined as the 3 values for K_1 , K_2 and K_3 . Typically a cost structure will be written as K_1, K_2, K_3 or as $K_1 K_2 K_3$. For example, a cost structure of 1,2,3 or 123 means $K_1 = 1$, $K_2 = 2$ and $K_3 = 3$.

Table 2 – MDP Policies

Policy	Policy Name	Description
DH	Deadhead	General policy, all flights are loaded to capacity
SR	Single Reference	General policy, compare inventory imbalance and reference level to determine excess equipment to load
OPT	Optimal	General policy, the policy resulting from the application of the Bellman equation
1N DH	1 Node Deadhead	Flights into and out of the only station are always loaded to capacity
1N SR	1 Node Single Reference	Flights out of the only station follow the SR policy
1N OPT	1 Node Optimal	The policy resulting from an application of the Bellman equation for the 1 Node model
SL 20	1 Node Optimal, SL = 20	The policy resulting from an application of the Bellman equation for the 1 Node model with SL = 20
SL 100	1 Node Optimal, SL = 100	The policy resulting from an application of the Bellman equation for the 1 Node model with SL = 100
2N DH	2 Node Deadhead	Flights into and out of both stations are always loaded to capacity
2N SR	2 Node Single Reference	Flights out of both stations follow the SR policy
2N OPT	2 Node Optimal	The policy resulting from an application of the Bellman equation for the 2 Node model
3N DH	3 Node Deadhead	Flights into and out of all three stations are always loaded to capacity
3N SR	3 Node Single Reference	Flights out of all three stations follow the SR policy
3N OPT	3 Node Optimal	The policy resulting from an application of the Bellman equation for the 3 Node model

The single reference provisioning policy will base the decision of how much excess equipment to load on two pieces of information, the current inventory imbalance, and the reference level. If the current inventory imbalance is less than the reference level then no excess equipment is to be loaded. If the current inventory imbalance is greater than the reference level then the aircraft is to be loaded to capacity.

5.1 Model 1 Results:

The 1N models were coded using VBA and MS access. The total computation time for 18 parameterized versions of the optimal policy, 108 parameterized versions of the SR policy and the 9 parameterized versions of the DH policy was approximately 1 hour.

The results of the 1-Node (1N) analysis can be summarized as follows:

- The 1N alternative policy suggested by analysis is the 1-Node Single Reference(1N SR) policy
- The difference in expected costs between the 1N SR policy and the 1N OPT policy is small, the maximum difference is less than 5%.
- The 1N SR policy outperforms the 1N DH policy when the 1N OPT policy outperforms the 1N DH policy
- Acceptable 1N SR policies exist in cost structures 123 , 133 and 233, for these three cost structures the cost of the 1N SR is on average 63% that of the 1N DH. Over all cost structures, the cost of the 1N SR is on average 142% that of the 1N DH, and 104% that of the 1N OPT.

We will begin the discussion of the 1N results by examining the 1N OPT policy. The formulation and results of the 1N SR policy, will be presented. This will be followed by a discussion of the results from the 1N DH policy. In each of these sections there is a discussion of the effect of K1 and K2 on policy performance. Finally, the effectiveness of the 1N SR policy will be presented.

5.1.1 1 Node Optimal Policy

The 1N MDP model was parameterized and solved. This process was repeated for eighteen parameter configurations. Table 3 summarizes the parameters used for identifying the characteristics of the 1N optimal policy.

Table 3 – 1 Node Model Parameters

<i>Model #</i>	K_1	K_2	K_3	N	S^L	
1	1	1	3	10	11	100
2	1	1	3	10	11	100
3	1	1	3	10	11	100
4	2	1	3	10	11	100
5	2	1	3	10	11	100
6	2	1	3	10	11	100
7	3	1	3	10	11	100
8	3	2	3	10	11	100
9	3	3	3	10	11	100
10	1	1	3	10	11	20
11	1	2	3	10	11	20
12	1	3	3	10	11	20
13	2	1	3	10	11	20
14	2	2	3	10	11	20
15	2	3	3	10	11	20
16	3	1	3	10	11	20
17	3	2	3	10	11	20
18	3	3	3	10	11	20

S^L - inventory imbalance limit
 N - number of epochs
 C - aircraft capacity
 K_1 - inventory imbalance penalty multiplier
 K_2 - excess equipment movement penalty multiplier
 K_3 - rebalance inventory multiplier


The results of the eighteen parameterized 1N OPT models for each epoch are presented in graphical form in appendix A - model 1 results. The selection of various K 's allows for the identification of trends and patterns that result from varying the magnitude of K_1 , K_2 and K_3 relative to each other.

The analysis of the 1N OPT models is broken into to groups, the models with S^L equal to 100, and the models with S^L equal to 20. We will refer to these models as the SL 100 and SL 20 respectively.

5.1.1.1 SL 100 models

The **SL 100 models** are a simplified version of the 1N model. Given the model parameters, when S^L is set at 100, the rebalance behavior of the model can never be triggered. This yields a simplified version of the complete model results. The control surface of one epoch in one of the models is presented in Figure 11. The control surface allows the identification of an action given a particular demand and inventory imbalance. For example, if the demand in epoch 7 is equal to 7 and the inventory imbalance is equal to 9, then the number of units of excess equipment to load (action) is 3.

Figure 11 – 1 Node OPT SL 100, Control Surface: Cost Structure 1,1,2



Comparing across the 9 SL 100 models, for all epochs several trends are identified. These can be summarized as:

Each epoch has three regions:

1. Ship no extra equipment
2. Ship as much excess as possible
3. Decision region connecting regions 1 and 2

Each epoch has a Decision Region:

The decision region is linear with increasing inventory

The decision region is linear with increasing demand

The bounds of the decision region move as K_1 and K_2 vary

The Cost Multipliers have several effects:

As K_1 increases the decision region is pushed to the left

As K_2 increases the decision region is pushed to the right

$K_1 > K_2$ decision region is pushed to the right

$K_2 > K_1$ decision region is pushed to the left


If $K_1 = 0$ all three regions are zero, K_1 is driving term

5.1.1.2 SL 20 models

The **SL 20 models** incorporate the full functionality of the one node model. Rebalance actions can be triggered. The behavior of the SL 20 models is similar to, but more complex than, the SL 100 models. The control surface of one epoch in one of the models is presented in Figure 12. The interpretation of the control surface is the same as for the SL 100 models.

Figure 12 – 1 Node OPT SL 20, Control Surface: Cost Structure 1,1,2

Control Surface, Epoch 8



The behavior of the SL 20 models can be described as follows:

Each epoch has at most seven regions

1. Trigger rebalance (inventory negative) - move all possible excess equipment
2. Increase inventory level - move no extra equipment
3. Decrease inventory level - move all possible excess equipment
4. Trigger rebalance (inventory positive) - move no extra equipment
5. Decision region connecting 1 and 2
6. Decision region connecting 2 and 3
7. Decision region connecting 3 and 4

Each epoch has at most three Decision Region:

Decision region connecting 1 and 2 is not smooth

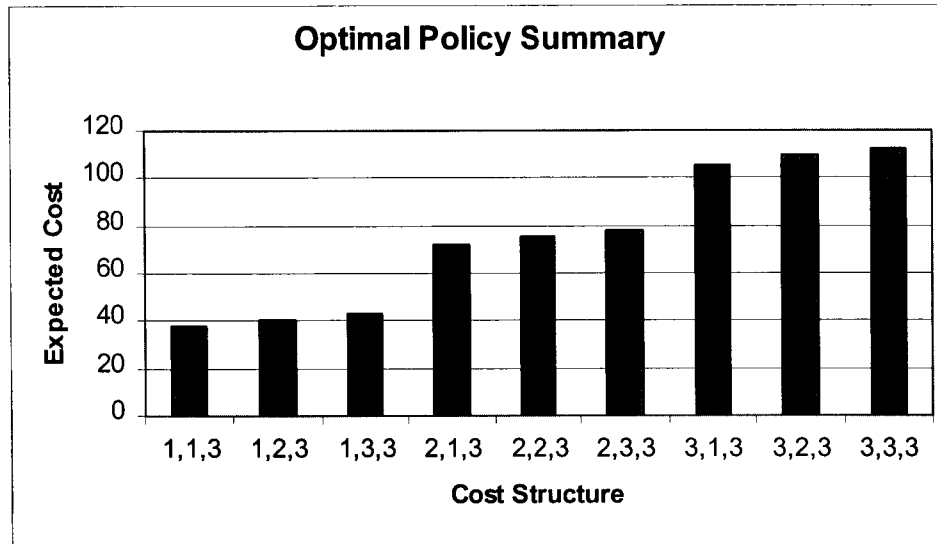
Decision region connecting 2 and 3 is smooth

Decision region connecting 3 and 4 is not smooth

Basic Characteristics

These models are the same as the SL 100 models, with the additional effect of the two trigger-rebalance regions. The results of the 1N optimal policy are presented in Figure 13 below. The expected costs are presented for each cost structure, and denote the expected cost of starting in epoch zero with an inventory imbalance of zero.

Figure 13 – Expected Costs of the 1 Node Policy




5.1.2 1 Node Alternative Provisioning Policy

The initial analysis of the 1N OPT policy suggests that a good approximation of the optimal policy exists. An alternative equipment provisioning policy would approximate 3 of the regions identified, 1-move no extra equipment, 2-move all available excess equipment and the 3-region connecting these two (middle region). The trigger rebalance regions observed outside these three are mathematical curiosities, they have little practical significance. If a decision was made to trigger a rebalance, why not instead conduct a rebalance that turn- that would avoid paying to move the same equipment twice. Approximating the decision region 3 as a step function generates an alternative policy that can be described as:

- 1) If current inventory imbalance is less than a reference inventory level (X^H), load no extra equipment
- 2) If current inventory imbalance is more than a reference inventory level (X^H), load all available extra equipment

This policy has its decisions based on a single reference inventory level (X^H) and as such is referred to as the Single Reference policy (SR). This model was coded, parameterized with the SL 20 parameters, and solved for values of X^H from -2 to 10. The results of the SR model for the optimal value of X^H , $X^H = 3$, are presented in Figure 14.

Figure 14 – 1 Node SR Model Summary




The expected costs for the 1N SR policy are very similar to the expected costs of the optimal policy. The only significant difference is that the expected cost for the 1N SR policy is at most 5% higher than the expected cost for the optimal policy for every cost structure. This suggests that the 1N SR policy is a very good approximation of the 1N OPT policy.

5.1.3 1-Node Deadhead Policy

The results of the 1N DH policy are obtained by adjusting the action selection rule to ensure that for any given demand, the amount of total equipment moved is always equal to the capacity of the aircraft. Therefore, given a passenger demand on an aircraft, the amount of excess equipment loaded is equal to the aircraft capacity minus the passenger demand. The results from the 1N DH policy are presented in Figure 15.

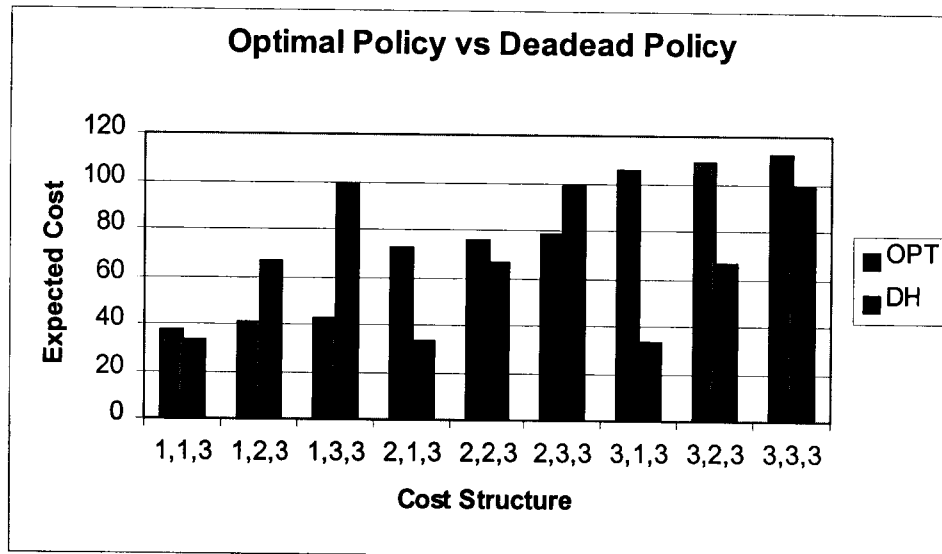
Figure 15 – 1 Node DH Policy Summary



5.1.4 Optimality and the 1 Node Optimal Policy

Generally when we refer to an optimal policy we mean that it is the best policy available. In the case of MDP analysis, and this thesis, the optimal policy is defined as the policy resulting from the application of the bellman equation. However, referring to Figure 16 we see that the 1N DH policy outperforms the 1N OPT policy for some cost structures, which is counter to the definition of the optimal policy. The short answer to this is that the Deadhead policy is not in the set of policies that we maximize over for the 1N model. The Deadhead policy is applied to a different 1N model.

Figure 16 – 1 Node Comparison of OPT and DH Policies



The fact that the 1N DH policy formulation results in a different model being solved was not discovered until after the formulation of the first MDP model. Rather than discarding the results, we should instead focus on them. This model has illuminated a very important fact- the strength of the 1N DH policy lies in convention. The 1N DH policy has the benefit of perfect information (knowing all equipment flows to and from all stations) because of the fact that equipment flows do not change. The cost that the 1N DH policy pays for the perfect information is the excess equipment that must be moved in order to maintain constant inbound and outbound equipment flows.


5.1.5 Effectiveness of Alternative Equipment Provisioning Policy

Recall that the objective of the MDP analysis was to develop an alternative equipment provisioning policy. The alternative provisioning policy has been identified as the SR policy. The objective of this section is to qualify the performance of the 1N SR policy. An acceptable alternative policy would outperform the 1N DH policy. A good policy would not deviate to far from the 1N OPT policy. It then follows that when qualifying the performance of an alternative policy, it must be compared to the optimal policy and to the deadhead policy.

5.1.5.1 Policy Cost Comparisons

In this section we will compare the expected costs generated by each model. The expected costs for all 1N policies, DH, OPT and SR with X^H ranging from -2 to 10, are presented in Figure 17.

Figure 17 – 1 Node Policy Cost Comparisons



For example, looking at policy type = OPT (1N Optimal policy) and K Values = 1,1,3 ($K_1 = 1$, $K_2 = 1$, $K_3 = 3$) reveals an expected cost of approximately 37. Although only three K values appear on the Y-axis, 113, 223 and 333 the values for all 9 cost structures are provided in the order 113, 123, 133, 213, 223, 233, 313, 323, 333.

Three conclusions can be drawn from this graph. The 1N DH policy outperforms the 1N OPT policy for some cost structures. The 1N SR policy outperforms the 1N DH policy in at most three cost structures. The optimal value of X^H (the value of X^H that yields the lowest cost) is between 2 and 4. The optimal values of X^H for each cost structure are given in Table 4.

Table 4 – 1 Node Optimal Value of X^H

Cost Structure	Optimal X^H
1,1,3	3
1,2,3	3
1,3,3	4
2,1,3	3
2,2,3	3
2,3,3	3
3,1,3	2
3,2,3	3
3,3,3	3

For all but two cost structures, $X^H = 3$ is the optimal value. In the two cost structures where the optimal value of X^H is not 3, the expected costs for $X^H = \text{"the optimal value"}$ and $X^H = 3$ differ by less than 1%.

5.1.5.2 SR Performance

Two performance metrics are used to evaluate the effectiveness of the 1N SR policy. The first performance metric is a comparison of the magnitude of the costs for each policy. This is summarized in Figures 18 and 19. From these two figures it can be seen that for all cost structures the 1N SR policy does not deviate too far from the optimal policy. However, only for those cost structures where the optimal policy outperforms the 1N DH policy does the 1N SR policy outperform the 1N DH policy. In the worst case, the 1N SR policy achieves a cost 300% greater than the 1N DH policy. In the best case, the 1N SR policy achieves a cost 55% lower than the 1N DH policy.

Figure 18 – 1 Node SR Policy Performance, Expected Cost


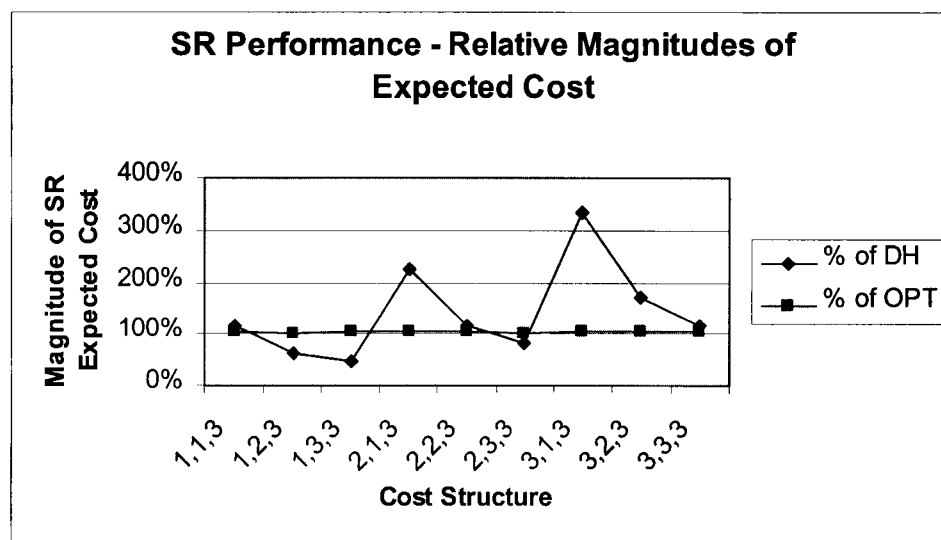


Figure 19 – 1 Node SR Policy, Relative Magnitude of Expected Costs



The second performance metric is the proportion of the expected cost reduction available that the 1N SR policy captures (%performance). The expected cost reduction available is defined as the difference in the expected costs of the 1N OPT policy and the 1N DH policy. In the case where the expected cost of the 1N DH policy is lower than the expected cost of the 1N OPT or 1N SR policy the % performance is zero. This metric is summarized in Figure 20.

$$\%performance = \frac{SR - Optimal}{DH - Optimal}$$

Figure 20 – 1 Node SR Policy Performance, % Performance




Figure 20, demonstrates that the 1N SR policy can capture a substantial portion of the cost reduction available. However, it also demonstrates that the % performance is highly sensitive to the selection of X^H and to the cost structure

5.1.5.3 SR policy Conclusions

The 1N SR policy is a good approximation of the optimal policy. However, attention must be made in the determination of the reference inventory level as this can significantly affect performance. More importantly, great consideration must be given to the identification of the prevailing cost structure, as this can make the difference between an alternative policy that reduces costs and one that significantly increase costs.

5.2 Model 2 Results:

The results of the 2-Node analysis can be summarized as follows:

1. The 2N alternative policy is the 2N SR policy
2. The 2N SR policy outperforms the 2N DH policy in all but 3 cost structures, the 1N SR policy was outperformed by the 1N DH policy in 6 cost structures implying that adding more nodes to the problem will generate a further increase
3. In the 6 cost structures where the 2N SR policy outperforms the 2N DH policy, the expected cost of the 2N SR policy (using optimal X^H) is on average 62% that of the 2N DH, over all cost structures this value is 92%
4. The 2N OPT policy has a lower expected cost than the 1N OPT policy, the 2N OPT policy has similar decision regions to the 1N OPT policy, and the effect of K_1 and K_2 on the 2N OPT policy are opposite from the 1N OPT policy
5. The effect of K_1 and K_2 are the same for the 2N SR policy and the 1N SR policy
6. C++ should be used to develop the 3N model if results are to be achieved in my lifetime

The 2N models were originally coded and solved using VBA and MS Access. The total computation time to solve nine parameterized versions of the optimal policy was approximately 168 hours of computation time. The nine parameterized versions of the DH policy were solved in approximately 2 and 1/2 hours. Finally, the 108 parameterized versions of the SR policy were solved in 33 hours. However, the SR policy was re-coded in C++ in order to verify the original results. The 108 parameterized versions of the SR model coded in C++ were solved in less than one minute.

5.2.1 Characteristics of 2 Node Optimal Policy

The model formulation presented in section 3.7 was coded and solved. Table 5 presents the different parameterizations of the model that were solved.

Table 5 – 2 Node Model Parameters

Model #	K_1	K_2	K_3	N	S^L
1	1		3	10	11 20
2	1		3	10	11 20
3	1		3		11 20
4	2		3		11 20
5	2		3		11 20
6	2		3		11 20
7	3		3		11 20
8	3	2	3		11 20
9	3		3		11 20

S^L - limit of deviation from inventory allotment

N - number of epochs

C - aircraft capacity

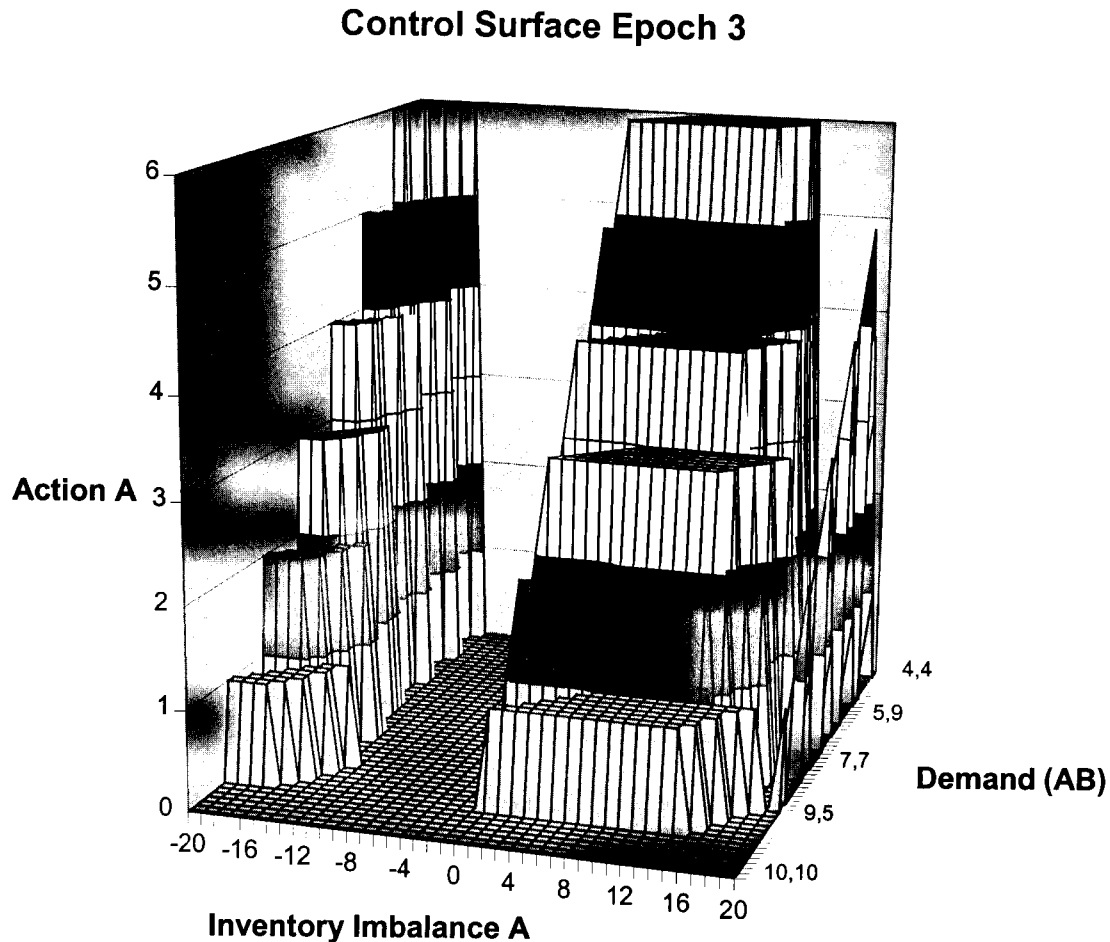
K_1 - deviation from allotment penalty factor

K_2 - excess equipment movement penalty factor

K_3 - rebalance inventory factor

The optimal solution to the 2N model is similar to the optimal solution to the one node model. The solution to the 2-node 323 optimal policy is presented as Figure 21.

Figure 21 – 2 Node OPT, Control Surface: Cost Structure 3,2,3



The above graph is the action space for station A only. Given an inventory imbalance at station A, and the demand on flights A and B, the action for flight A can be determined. For example, given that the current epoch is 3, the inventory imbalance at station A is -4, the demand on flight A of 9, the demand on flight B of 5 we can determine the number of units of excess equipment units to load on flight A. Referring to inventory imbalance = -4 and demand (AB) = 9,5 we see that the action on flight A is 0. To determine Action B, we would need to consult the station B control surface.

If the results are analyzed on a per station¹⁶ basis, then each component of the action space can be described as follows:

¹⁶ There are 2 stations in the two node problem. The action space consists of two actions for each state; one action for each station. The results, meaning the action space corresponding to a particular state are difficult to present and interpret using a graphical 3 dimensional representation. For simplicity, each component of the action space is graphed separately.

Each epoch has at most seven regions

1. Trigger rebalance (inventory negative) - move all possible excess equipment
2. Increase inventory level - move no extra equipment
3. Decrease inventory level - move all possible excess equipment
4. Trigger rebalance (inventory positive) - move no extra equipment
5. Decision region connecting 1 and 2
6. Decision region connecting 2 and 3
7. Decision region connecting 3 and 4

Decision regions

Decision region connecting 1 and 2 is not smooth


Decision region connecting 2 and 3 is smooth

Decision region connecting 3 and 4 is not smooth

Given that the 2N OPT is similar to the 1N OPT, it would follow that the SR policy would be applicable to the 2N model as well. This will be addressed in more detail in section 5.2.2.

The expected cost for the 9 parameterized models are presented in Figure 22. Surprisingly, the expected costs for the 2N OPT policy are less than the expected costs for the 1N OPT policy. Upon closer examination it can be seen that this is due to the differences in model formulation, namely, the 2N model contains demand information for both station. More specifically, the 2N OPT policy has the benefit of knowing the demand on both flights, the inventory imbalance at both stations when making the decisions on what to load on both flights. Clearly, the additional information will result in a superior policy, and this is borne out in the 2N OPT policy expected costs.

Figure 22 – 2 Node OPT Policy Summary



5.2.2 2 Node Alternative Provisioning Policy

The 2N alternative equipment provisioning policy is based upon the decision regions identified from the 2N OPT policy. Given the similarity between the 1N OPT policy decision regions and the 2N OPT policy decision regions, it is appropriate to apply the SR policy as the alternative policy for the 2N model. The 2N SR policy can be described as:

- 1) If current inventory is less than a reference inventory level (X^H), load no extra equipment
- 2) If current inventory is more than a reference inventory level (X^H), load all available extra equipment

The 2N SR policy is similar as the 1N SR policy. A key distinction between the 2N SR policy and the 2N OPT policy is the fact that the decision made for the 2N SR policy does not use all the information available in the state. In particular, when deciding what amount of excess equipment to ship from one of the stations, for arguments sake station A, the policy will consider two quantities; the observed demand level on the departing flight from station A and the current level of inventory at station A. The policy will ignore the observed demand level on a flight departing from station B that is going to arrive at station A. The policy will also ignore the current level of inventory at station B. The reverse is true when deciding how much equipment to be loaded on the flight from station B.

The 2N SR policy was solved for various values of the reference inventory level X^H . For each reference level X^H , the model was solved using the same parameters presented in Table 5. The expected costs for the case where $X^H = 3$ are presented in Figure 23.




Figure 23 – 2 Node SR Policy Expected Costs


The expected costs for the 2N SR policy are higher than the expected costs for the 1N SR policy, however, there is also another station contributing to those costs. In all cases the 2N SR expected costs are lower than two times the 1N SR expected costs.

5.2.3 2 Node Deadhead policy

The 2N DH policy is similar to the 1N DH policy. Given any passenger demand on a flight, the amount of excess equipment will be equal to the aircraft capacity minus the observed passenger demand. The key difference between the 2N DH policy and the 1N DH policy is that the 2N DH policy takes into account the cost implications of a second station.

The 2N DH policy was solved using the same parameters presented in Table 5. The expected costs for the 9 parameterized models are presented in Figure 24.

Figure 24 – 2 Node DH Policy Summary



The expected costs for the 2N DH policy are similar to the expected costs for the 1N DH policy. The only difference is that in all cases the expected cost for the 2N DH policy is exactly two times the expected cost for the 1N DH policy.

5.2.4 Optimality and the 2 Node Optimal Policy

The results of the 2N policies (OPT, SR and DH) demonstrate that the 2N OPT policy is always optimal. This result is expected because all three policies are applied to the same model formulation, which means that by definition the policy resulting from the application of the Bellman equation is the optimal policy.

5.2.5 Effectiveness of the 2 Node Single Reference Policy

In order to determine the effectiveness of the 2N SR policy, its results must be compared to the 2N DH policy and the 2N OPT policy.

5.2.5.1 Policy Cost Comparisons


The expected costs for all parameterized 2N models are presented in Figure 25.

From Figure 25 it can be seen that:

- 1) The 2N OPT policy outperforms all other policies by a considerable margin
- 2) For some cost structures the deadhead policy outperforms the single reference policy
- 3) Negative values of X^H in the 2N SR policy result in expected costs identical to the 2N DH policy
- 4) The optimal value of X^H (inventory reference level) varies depending on cost structure

Each of the above observations will be discussed in greater detail.

Figure 25 – 2 Node Policy Cost Comparisons



In all cases the 2N OPT policy costs are at least 70% lower than the costs of the 2N DH policy and 22-70% lower than the cost of the 2N SR policy. This can be thought of as the value of using perfect information. In the DH policy and the SR policy, information in the state is disregarded when making the decision. The result is a considerably lower cost for the 2N OPT policy.

The value of X^H resulting in the lowest expected cost for the 2N SR policy changes given the cost structure parameterization. The optimal value of X^H for each cost structure is presented in Table 6.

Table 6 – 2 Node Optmal Value of X^H

Cost Structure (K's)	Optimal Value of X^H
1,1,3	2
1,2,3	2
1,3,3	3
2,1,3	-1
2,2,3	2
2,3,3	2
3,1,3	-1
3,2,3	-1
3,3,3	2

As in the 1N model, the 2N DH policy outperforms the 2N SR policy, for some cost structures. This occurs in cost structures 2,1,3 , 3,1,3 and 3,2,3. However, there are fewer cost structures in which the 2N DH policy outperforms the 2N SR policy than in which the 1N DH policy outperforms the 1N SR policy.

When X^H is negative, the 2N SR policy behaves exactly like the 2N DH policy. This should not come as a surprise when we realize that the starting inventory imbalance is zero. Therefore, when X^H is less than zero, the action selected by the 2N SR policy is always move all available excess equipment, which is exactly the 2N DH policy.

5.2.5.2 2 Node SR Performance

Two performance metrics are used to gauge the performance of the 2N SR policy. These are the same two metrics used for the 1N model analysis.

The first performance metric is magnitude of expected costs. This metric is summarized in Figures 26 and 27. There are two important observations from Figure 26. First, as K_1 increases, the difference in expected cost between the 2N SR policy and the 2N DH policy decreases, in three cases the 2N DH policy performs better than the 2N SR policy. In addition, as K_1 increases the difference between the 2N SR policy and the 2N OPT policy increases. Second, there are several cost structures for which the 2N SR policy is below the 2N DH policy and close to the 2N OPT policy. As can be seen in Figure 27, the 2N SR policy expected cost ranges from 130%-700% of the 2N OPT policy, and 35%-200% of the 2N DH policy.

Figure 26 – 2 Node SR Policy Performance, Expected Cost





Figure 27 – 2 Node SR Policy, Relative Magnitude of Expected Costs




The second performance metric is the proportion of expected cost reduction available that is captured by the 2N SR policy (%performance). The expected cost reduction available is defined as the difference in the expected costs of the 2N OPT policy and the 2N DH policy. The proportion captured is therefore equal to the expected cost of the 2N SR policy minus the expected cost of the 2N OPT policy, all over the expected cost of the 2N DH policy minus the expected cost of the 2N OPT policy.

$$\%performance = \frac{SR - Optimal}{DH - Optimal}$$

This metric is summarized in Figure 28. As can be seen, the most favorable cost structure is 113 followed by 123 and then by 233. It is visually apparent by the presence of many "0% cost reduction captured" that the 2N DH policy does indeed outperform the 2N SR policy for some

cost structures. It can also be seen that $X^H=2$ performs very well in the most favorable cost regimes. Finally, comparing Figure 28 to Figure 20 we see that the 2N SR policy captures available cost reduction for more cost structures than does the 1N SR policy, 6 for the 2N and 3 for the 1N. However, due to the superior performance of the 2N OPT policy in comparison to the 1N OPT policy, the average available cost reduction captured by the 2N SR policy, is less than the 1N SR policy, 23% for 2N SR versus 27% for 1N SR. Considering only non-zero performance regions for 1N SR, this value is 60% for 2N SR versus 77% for 1N SR.

Figure 28 – 2 Node SR Policy Performance, % Performance



5.3 Model 3 Results

The 3N models were coded in C++. The total computation time required to solve all the parameterized versions of the model was approximately 350 hours. Interested readers are directed to Appendix A for an in depth discussion of computation problems with solving the 3 Node Model.

The 3N results can be summarized as follows:

1. The 3N SR policy outperforms the 3N DH policy in all but 1 cost structure, implying that adding more nodes to the problem will generate a further increase
2. In the 8 cost structures where the 3N SR policy outperforms the 3N DH policy, the expected costs of the 3N SR policy (using optimal X^H) is on average 50% that of the 3N DH, over all cost structures this value is 58%
3. The 3N OPT policy is always optimal, and the expected cost of the OPT policies may be linearly related to the number of stations
4. The computational burden presented by the 3N model indicates that the development of a 4N model is impractical


The model formulation presented in section 3.8 was coded and solved using C++. In order to formulate the problem, a rebalance algorithm was created to determine which stations would be involved in an automatic rebalance action, if required.

5.3.1 Characteristics of the 3 Node Optimal Policy

An analysis of a scaled down version of the model (aircraft capacity =2, 2 bins of demand), resulted in an understanding of 3N OPT policy. The policy will select the minimum amount of excess equipment to load onto departing aircraft in order to maintain a zero inventory imbalance.

The full version of the 3N OPT policy was solved for the 9 different cost structures used in the previous two models. The expected costs for the 9 parameterized versions of the 3N OPT model are presented in Figure 29.

Figure 29 – 3 Node OPT Policy Summary




Comparing the 3N OPT policy to the 2N OPT policy, we see that the 3N expected costs are approximately 1.5 times the cost of the 2N costs. This indicates that the expected costs for the optimal policies with complete information (ie the 2N and the 3N) may be linearly related to the number of nodes in the model.

5.3.2 3 Node SR Policy

The results for the 3N SR policy are summarized in Figure 30

Figure 30 – 3 Node SR Policy Summary




The expected costs for the 3N SR policy are larger than the expected costs for the 3N OPT policy. In addition, they are 1.6-2.0 times the expected costs for the 2N SR model. Although the stations only increase by a factor of 1.5, the number of flights increased by a factor of 2. Therefore it would appear that the 3N SR policy is benefiting from additional flight volumes in a larger network.

5.3.3 3 Node DH policy

The 3N DH policy is similar to the 2N and 1N DH policy. The only difference is that the policy is applied to all six flights per epoch. The results from the 3N DH policy are summarized in Figure 31. The expected costs for the 3N DH policy are noticeably higher than the 3N SR policy and the 3N OPT policy. The expected costs for the 3N DH policy are exactly three times higher than the expected costs for the 2N DH policy, and exactly six times higher than the expected costs for the 1N DH policy. This indicates that the expected costs for all DH policies are linearly related to the number of flights per epoch. The effect of K1 and K2 on 3N DH expected costs are exactly the same as the effect of K1 and K2 on 2N DH and 1N DH expected costs. K1 has no effect, and K2 has a correlation of 1 with expected cost.

Figure 31 – 3 Node DH Policy Summary




5.3.4 Effectiveness of the 3 Node Alternative Provisioning Policy

The expected costs for all 3N policies are presented in Figure 32. It can be seen that:

1. The 3N OPT policy has the lowest expected cost for all cost structures
2. The 3N SR policy has lower expected costs than the 3N DH policy
3. The 3N SR policy has higher expected costs than the 3N OPT policy
4. The optimal value for reference level (X^H) is 3
5. When X^H is less than 0, the 3N SR policy degenerates into the 3N DH policy

Figure 32 – 3 Node Policy Cost Comparisons



The metrics used for the 2N and 1N models will be used to quantify the performance of the 3N SR policy.

5.3.4.1 3 Node SR Policy Performance

Figures 33 and 34 present the expected costs and relative magnitudes of expected costs for the 3N models (The 3N SR policy is presented with the optimal value of X^H).

Referring to Figure 33, it can be seen that the 3N SR policy has a lower expected cost for 8 of the 9 cost structures. Recall that the 2N SR policy was 6 of 9, and the 1N SR policy was 3 of 9. In addition, in the cost structure where the 3N SR policy has higher expected cost than the DH policy, the difference between the 3N SR policy and 3N DH policy is small.

Figure 34 demonstrates that the 3N SR policy deviates substantially from the 3N OPT policy. This indicates that there is significant cost reduction still available that is not captured.

Figure 33 – 3 Node SR Policy Performance, Expected Costs






Figure 34 – 3 Node Policy Performance, Relative Magnitude of Expected Costs



The proportion of available cost reduction captured by the 3N SR policy is presented in Figure 35. The 3N SR policy has a reasonable performance ($>30\%$) for 8 of the 9 cost structures. Comparing Figure 35 to Figure 28 and Figure 20 it can be seen that an increase in the number of nodes has increased the number of cost structures in which the 3N SR policy has a non zero performance. In addition the magnitude of the percent performance has also increased from the 2N to the 3N SR policy. The average performance, across all cost structures, for the 2N SR policy with $X^H = 3$ is 36%. The average performance, across all cost structures, of the 3N SR policy with $X^H = 3$ is 53%.

Figure 35 – 3 Node SR Policy Performance, % Performance



In summary, the 3N SR policy is superior policy than the 3N DH policy for 8 of 9 cost structures. However, the 3N SR policy does not capture all of the available cost reduction; defined as the difference between the 3N OPT policy and 3N DH policy. The 3N SR policy results indicate that an increase in the number of nodes, and more importantly, the number of flights per epoch results in an increase in the performance of the SR policy.

6.0 Simulation: Model Formulation

A simulation model is a representation of the real world system under study. Entities are defined; in the case of this simulation the entities are aircraft. The entities follow some logical flow; in this case an entity enters the simulation, acquires equipment from its departing station, flies to the arriving station and drops off its equipment there. In many simulations there are metrics for measuring the performance of the simulation. There can be deterministic events and stochastic events in a simulation. In this simulation an example of a deterministic event is a flight leaving its departure station and arriving at its arrival station. We know with certainty (at least in the simulation) that the flight will arrive at the arrival station. An example of a stochastic event would be the number of passengers that are actually on the aircraft. Finally, the simulation comes to life by executing the simulation logic many times over, a process known as replication, with each replication randomly choosing a different set of stochastic events.

In this section we will present the simulation model formulation. There are four main components in model formulation; model inputs, model logic, model metrics, and provisioning policy. Each component will be discussed in more detail.

6.1 Simulation Model Inputs

The simulation model has three inputs, namely:

1. Flight schedule
2. Passenger demand
3. Equipment manifest

The simulation allows for an equipment provisioning policy to be implemented in a test environment. The simulation flight schedule defines the number and time of flights between the stations. Each flight in the simulation flight schedule has a simulated passenger demand. Finally, each flight has an equipment manifest.

6.1.1 Flight Schedule

The flight schedule was obtained from the Meal Pages database¹⁷. The flight schedule defines the following for every flight CAI offers:

- Date of flight
- Arrival station, departure station
- Arrival time, departure time
- Aircraft type and seating configuration
- Galley code
- Meal service type
- Miscellaneous service descriptions

A representative 2-week sample from the schedule was selected. The sample data set contained 2986 flights, over 14 days, operating between 31 flight stations, with nine different aircraft types.

¹⁷ For more information please refer to section 1.6.1

6.1.2 Modeling Passenger Demand

The purpose of modeling passenger demand is to provide a method for determining the number of passengers on any given flight in the flight schedule. There are several available methods for modeling passenger demand for use in the simulation. One method is to model the passenger demand for each city pair for each direction. However, this would require 1860 different distributions. An alternative approach is to cluster the stations in groups that share common characteristics in order to reduce the total required number of distributions and also minimize the loss of information.

For this simulation 9-passenger demand distributions were generated. The first six passenger demand distributions are for general geographic areas that have similar passenger demand characteristics. For example, if CAI operated flights to many destinations in the Middle East region, and the characteristics of passenger demand on those flights were similar, one demand distribution would be created for all flights to the many destinations in the Middle East region. The remaining three passenger distributions are for the three main cities (three main city pairs) of the CAI network. Specifically, each of the distributions is based only on observed passenger demand between one city pair.

From the discussion of the CAI network in section 1.4.5, it can be seen that this choice of passenger distribution parameterization will capture most of the essential elements of the system. The use of a passenger distribution for each of the three main city pairs permits an accurate representation of a large number of flights. In addition, the majority of the flights to geographic regions is directly linked to one of the three main stations, and do in fact have similar demand profiles.

6.1.3 Equipment Manifest

In order to complete the equipment flow model, the quantity and type of equipment flown on each aircraft must be determined.

From the development of the Catering Commissary Database¹⁸, it became apparent that the variation in equipment manifest on the Tray set up (TSU) level for economy class passengers was minimal. In fact, the only major difference in equipment manifests was the number of TSUs loaded per passenger. On long over seas flights two meals are provided for each Y class passenger, requiring two TSUs to be loaded per passenger. On short domestic and trans border flights, one meal is provided per passenger. Therefore the economy class passenger equipment manifest can be approximated as a tray set up unit with several standard components.

However, the variation in equipment manifest on the TSU level for business class passengers is much more extreme. Fortunately, there are several reasons to exclude business class passengers from this analysis. The comparative size of business class is small, on average 12.8% aircraft seating. The risk involved in adjusting business class service is very high, business class passenger's account for a disproportionately large percentage of an Airline's revenue. The variation in the equipment manifest for business class passengers is significant, meaning there are many more types of equipment to be balanced. Finally the occupancy percentage of business class is higher than economy, meaning there is less excess equipment loaded in the first place. From the above, business class passengers will be excluded from the analysis.

¹⁸ For more information please refer to section 1.6.2

In summary, the equipment manifest used in the simulation is approximated as a tray set up unit with several standard components. Accordingly, the analysis is restricted to economy equipment only.

6.2 Simulation Model Logic

The model logic is as follows,

1. Initialize replication with all stations starting with an inventory imbalance of 0, and time equal to starting date
2. Determine passenger demand for each flight
3. Apply the provisioning policy to determine excess equipment loaded onto all flights
4. Calculate total equipment loaded for each flight (passenger demand + excess equipment loaded)
5. Calculate total outbound equipment for each station
6. Calculate total inbound equipment for each station
7. Calculate next days inventory imbalance for each station
8. If the current date is less than the ending date, increment the date by one day and return to step 2

The simulation is evaluated on a flight-by-flight basis; however, statistics are gathered on a daily basis.

6.3 Simulation Model Metrics

In order to analyze the output it is necessary to develop simulation performance metrics. These include:

1. Quantity of excess equipment moved
2. Magnitude of inventory imbalances

6.3.1 Quantity of Excess Equipment Moved

This metric is a means of measuring the cost reduction potential of an alternative policy. In order to measure the cost reduction we would first determine the reduction in excess equipment movements. This is calculated by determining the excess equipment movements resulting from the deadhead policy X_{DH}^E . Then determining the excess equipment movements resulting from an alternative policy X_{SR}^E . The reduction in excess equipment movements is found by subtracting X_{SR}^E from X_{DH}^E . Finally, the reduction in excess equipment movements is multiplied by the incremental cost associated with moving a unit of excess equipment, yielding the cost reduction potential of a policy.

This metric is easy to measure. The equipment provisioning policy prescribes the number of excess equipment movements for each flight. Therefore, the quantity of excess equipment moved can be measured by summing over the appropriate time interval.

6.3.2 Magnitude of Inventory Imbalances

This metric is a measure of the risk of stock out for an alternative policy. The risk of stock out can be used to determine required inventory allotment levels, because with any feasible solution, the risk of a stock out should be very close to zero. Therefore if an alternative policy has a high risk of stock out at a station, then additional equipment would have to be purchased for that station in order to reduce the risk of a stock out to acceptable levels. The level of inventory increase / decrease required to maintain a probability of stock out close to zero will have cost

ramifications with respect to purchasing and supplying additional safety stock. Ideally the required additional inventory will be zero.

The inventory imbalances are also directly measurable for each simulated day of operations. The important metric is the number of times the inventory imbalance decreases past the safety stock level. Measuring the inventory imbalance with respect to the safety stock level will allow us to assess the risk of stock out. Whenever the policy requires the use of safety stock, it is running the risk of stock out. The relative risk of the policy can be approximated by counting the number of times it required the use of safety stock, at each station.

6.4 Simulation: Alternative Equipment Provisioning Policy

In this section, we transform the single reference policy (SR) resulting from the Markov Decision Process (MDP) analysis¹⁹ into the universal simulation single reference policy (USSR). The objective of this section is to identify the differences between the simulation and MDP models, and modify the SR policy accordingly.

6.4.1 Developing the Universal Simulation Single Reference Policy (USSR)

The simulation model has several differences from the MDP model that should be taken into consideration, and used to modify the SR policy. There are four main differences:

- 1) The simulation model has multiple sizes of aircraft
- 2) Each station has a varying number of flights
- 3) There are multiple passenger demand distributions
- 4) Any given station does not have flights to all other stations

All four points highlight a key distinction between the MDP models and the simulation model; the MDP models are symmetric, in flights, passenger demand and aircraft types. The simulation model is not symmetric.

The SR policy resulting from the MDP analysis is repeated for convenience:

- 1) If current inventory imbalance is less than a reference level move no extra equipment
- 2) If current inventory imbalance is more than a reference level move all available extra equipment

Recall that the value of the reference level defined the provisioning policy. In the MDP model, the reference level for all stations was equal.

¹⁹ Refer to Section 5.0

Applying the SR policy to a simulation with 31 stations, each with a different number of flights per day, we get the USSR policy:

$$X_{ijk}^L = \begin{cases} 0 & \text{if } S_{ij} \leq X_i^H \\ C_{ijk} - D_{ijk} & \text{if } S_{ij} > X_i^H \end{cases}$$

i = station index, a value between 1 and 31
j = day index, a value between 1 and 14
k = flight index, each station i will have n_{ij} flights on day j
 C_{ijk} = aircraft capacity
 D_{ijk} = passenger demand
 X_{ijk}^L = excess equipment loaded
 X_i^H = reference inventory imbalance level
 S_{ij} = inventory imbalance for station i on day j

Remembering the fact that the simulation model is not symmetric, it is not appropriate for the value of reference level X_i^H to be equal for all stations. Therefore, the reference level X_i^H should be unique for each station, meaning that the values of the set of reference levels define the policy.

In order to quantify the non symmetric nature of the simulation model, we can calculate the average sum total of the number of seats flown out of each station in one day. We will call this quantity Daily Volume (V).

$$V_i = \frac{\sum_{j=1}^{14} \sum_{k=1}^{n_{ij}} C_{ijk}}{14}$$

i = station index, a value between 1 and 31
j = day index, a value between 1 and 14
 k_j = flight index, each station i will have n_{ij} flights on day j
 V_i = Daily volume out of station i
PC = Aircraft capacity

V_i is calculated for each station. This is then divided by the sum total of V_i for all stations and multiplied by 100. This is the % of system V, and is shown for all stations in Figure 36. Figure 36 shows that the station size (% of system V) varies significantly across all stations. There are 19 small stations (% of system V < 2%), 9 medium size stations (% of system V < 10%), and three very large stations (% of system V > 10%). The 19 small stations account for 16.4% of the system V. The 9 medium stations account for 24.2% of the system V. The 3 large stations account for 59.4% of the system V.

Figure 36 – Histogram of Normalized City Size





Figure 37 – Percent System Volume by Station Group



Given that we wish to modify X_i^H to allow for simulation asymmetry, it follows that the value of X_i^H should be related to V_i .

$$X_i^H = \lambda * V_i$$

λ = V multiplier

This approximation will allow the USSR policy to vary depending on station size. It is important to note that it is the choice of λ combined with the set of V_i that now define X_i^H , and therefore the policy. It should be noted that as λ becomes increasingly negative, the USSR policy approaches the DH policy. In the following section we will determine lambda by trying various values for lambda and selecting the best one.

6.4.2 Inventory imbalance limit (S^L), safety stock level

Another important difference between the Simulation model and the MDP model is the value of the inventory imbalance limit (S^L). In the MDP models, this value was constant across all stations, owing to the homogeneous nature of the stations. However, for reasons cited in the previous section the value of S^L should also vary by station.

In the CAI system, each station has an inventory allotment, I_i^A . The value of I_i^A is equal to 1.5 times the maximum daily use of equipment. We can approximate this as $1.5V_i$ ²⁰. The safety stock level I_i^S , will then be defined as I_i^A minus the value of S_i^L .

The inventory imbalance limit S_i^L must be some value between 1 and $1.5 V_i$. For any given day, the amount of equipment used is approximately equal to V_i , so the inventory imbalance limit must be greater than this. In addition, we do not want to add additional equipment, therefore the maximum absolute inventory imbalance must be less than the allotment, or $1.5 V_i$.

Defining the inventory imbalance limit is important for two reasons. First it is required in order to determine if inventory imbalances are acceptable. Secondly, the choice of S_i^L defines the safety stock level. The safety stock level will be a value between $0.5V_i$ and 0, corresponding to a S_i^L between $1 V_i$ and $1.5 V_i$. For the simulation S_i^L will be defined as:

$$S_i^L = \gamma V_i$$

γ = a constant between 1 and 1.5

The choice of γ then becomes very important- if it is too high, very low safety stock is available, if it is too low the station is forced to carry extra inventory. The value of gamma is arbitrarily set at 1.3 for the simulation model. This value provides a "middle" of the road approach, with a safety stock level of $0.2 V_i$. The effect of varying the value of γ will be investigated before final simulation results are obtained.

²⁰ Strictly, the inventory allotment is equal to the maximum possible equipment use for the entire schedule. The simulation only incorporates 2 weeks of the schedule, and therefore DV_i is a conservative approximation

7.0 Simulation: Model Results

The goal of developing the simulation model is to test the single reference policy resulting from the MDP analysis in a model that more closely approximates the actual CAI system. The simulation model formulation section has described how the CAI system was translated into the simulation model. This section will present the results of the single reference policy applied in the simulation.

The simulation model presented in section 6.0 was formulated and solved using MS Excel and VBA. Each parameterized version of the model was solved using 100 replications, with each replication using a distinct set of random numbers (for generating passenger demand).

The provisioning policy used in the simulation was modified as a result of the analysis, generating two variants of the USSR policy; the RSSR policy and the 3SSR policy. These policies are summarized in table 8.

Table 7 – Summary of Provisioning Policies

Policy	Policy Name	Policy Description
SR	Single Reference	If inventory imbalance is greater than reference level, ship all available excess equipment, if inventory imbalance is less than reference level, ship no excess equipment
DH	Deadhead	Always ship all available excess equipment
USSR	Universal Simulation Single Reference	All 31 simulated stations follow a customized SR policy, the reference level varies by station
RSSR	Reduced Simulation Single Reference	19 small stations follow the DH policy, the 12 other stations follow a customized SR policy
3SSR	3 Station Simulation Single Reference	28 stations follow the DH policy, the three largest stations follow a customized SR policy

The main results of the simulation analysis can be summarized as:

1. The USSR policy generates unacceptable inventory imbalances
2. The RSSR policy results in a reduction in excess equipment movements of 3200/day, reducing 60% of all possible excess equipment movements
3. The 3SSR policy results in a reduction in excess equipment movements of 1000/day, reducing 19% of all possible excess equipment movements
4. The resulting cost savings for the RSSR policy are estimated to be \$390,000 annually
5. The resulting cost savings for the USSR policy are estimated to be \$120,000 annually
6. The DH policy has a total of 5300 excess movements per day, out of a total of 24,000 total equipment movements per day

The following sections present the results for the simulation model in greater detail. The method for analysis is as follows:

- Step 1 – determine if policy reduces excess equipment movements and maintains acceptable inventory imbalances
- Step 2 – determine range of optimal value of λ
- Step 3 – assess stability of model resulting from range of values of λ ; select one value of λ
- Step 4 – determine suitable warm-up period
- Step 5 – determine reduction in excess equipment movements for policy with optimal stable value of λ , adjusting for warm-up period

7.1 USSR Results

The simulation was parameterized and solved with γ equal to 1.3²¹ for several values of λ (recall that the value of λ defines the SR policy). The results demonstrated that several stations exhibited unacceptably high inventory imbalances, and that these stations are the 19 small stations. In addition, the cost reduction potential of the USSR policy is realized for values of λ between -0.1 and 0.3 . The value captured increases exponentially with increasing λ . The next step is to reformulate the provisioning policy; 19 small stations will use the DH policy, and the 9 medium stations and 3 large stations will use the SR policy.

Figure 38 shows the minimum inventory imbalance as a percent of S^L_i (%MII).²² Each line in Figure 38 represents one station. For each value of λ , the simulation is replicated 100 times. The minimum value of inventory imbalance²³ (MII) across the 100 replications is recorded for each station. The MII is then divided by S^L_i (inventory imbalance limit for station i) and multiplied by negative one. The %MII is presented for each station for values of λ from -1.1 to $+1.1$ in 0.1 increments. Finally the results of the simulation using a DH policy are presented for comparison on the left hand side of the x axis.


As λ increases, the %MII also increases. However, it can be seen that the %MII doesn't stabilize for the large stations (the three increasing series below 50%). In addition the %MII reaches its maximum value after $\lambda = 0.3$. From an initial inspection it can be seen that quite a few stations (19) are experiencing inventory imbalances that are significantly higher than their S^L_i .

²¹ the value of γ will determine the value of S^L , and therefore influence the simulation metrics as defined in 6.4.2

²² S^L_i is a positive number, to make %MII a positive number it must be multiplied by -1

²³ remember that inventory imbalance can be negative

Figure 38 – USSR, Minimum Inventory Imbalance, % of S^L , $\gamma = 1.3$



An interesting question is why do some stations exhibit a non-zero minimum inventory deviation for the DH policy? The DH policy is supposed to ensure that inventory imbalances are zero, and there are clearly non-zero values for inventory imbalances. The answer comes from an investigation of the flight schedule. Over some period of time, for example a week, the number of aircraft into a station should equal the number of aircraft out of a station. An investigation of the schedule demonstrates that this is the case. However, over a day, this is not necessarily true. For example, consider a station following the deadhead policy, a station, which has only one flight that departs on Monday, and returns on Tuesday. On Monday, the station will show a negative inventory imbalance. On Tuesday the inventory imbalance will return to zero. Therefore over short periods it is possible for stations following the Deadhead policy to experience non zero inventory imbalances.

Figure 39, presents the total number of times the inventory imbalance decreased past the inventory imbalance limit for all 100 replications. Again each line represents one station. As can be seen the inventory imbalances generated by the USSR policy are too large to be acceptable. Interestingly, it is only the small stations that exhibit this behavior.

Figure 39 – USSR, # of Times Inventory Imbalance Decreases Below $-S^L$, $\gamma = 1.3$




Figure 40 – USSR, Average Excess Equipment Moved, $\gamma = 1.3$




Figure 40 shows the average excess equipment moved from each station. As can be seen the number of excess equipment movements of the USSR policy are minimized for values of λ greater than 0.3, meaning the cost reduction is maximized²⁴. Cost reduction from the USSR policy is not realized until λ exceeds -0.1 , and from there the cost reduction increases exponentially with increasing λ .

²⁴ Recall the discussion from section 6.3.1

7.2 Reducing the Number of Cities

The simulation was re-formulated and solved with γ equal to 1.3 for several values of λ . For this policy the “small” stations with unacceptable inventory imbalances are set to use the DH policy. The remaining 12 stations continued to operate on the SR policy. This combined policy of DH for 19 stations and SR for 12 stations is referred to as the reduced simulation single reference (RSSR) policy.

The results from RSSR policy demonstrate that the inventory imbalance for all stations is acceptable for $\lambda \leq 0.2$. For $\lambda \geq 0.3$ some additional inventory would have to be purchased for three of the stations. The cost reduction for the RSSR policy increases exponentially with increasing λ , reaching a maximum of 67% at $\lambda = 0.3$. Consequently, it can be seen that the 19 stations removed (61% of stations) account for only 33% of the potential cost reduction.

Figure 41 presents the %MII for the reduced simulation. Figure 41 demonstrates the impact of removing the small stations. %MII increases with increasing λ , and the large stations' %MII still do not stabilize, but the maximum %MII is approximately 100%. This indicates that the RSSR policy generates acceptable levels of inventory imbalance. This finding is further supported by Figure 42, which shows that there are only three stations that exhibit inventory imbalances below S^L_i . Further, this only occurs 3 times over 100 replications of 14 days when λ is greater than 0.2.

Figure 41 – RSSR, Minimum Inventory Imbalance Deviation, %, $\gamma = 1.3$




Figure 42 – RSSR, # of Times Inventory Imbalance Decreases Below $-S^L$, $\gamma = 1.3$





Figure 43 shows the average excess equipment moved for the reduced simulation model by city. Similar to what was observed for the USSR policy, the cost reduction for the RSSR policy is zero until $\lambda = -0.1$. From there, cost reduction increases exponentially until $\lambda = 0.3$. When $\lambda = 0.3$ the RSSR policy still generates excess equipment movements as a result of the 19 stations following the DH policy. The maximum value captured as 67% for $\lambda = 0.3$ onward. The % value captured is defined as 100 multiplied by one minus the excess equipment movements generated by the RSSR policy divided by the excess equipment movements generated by the DH policy.

Figure 43 – RSSR, Average Excess Equipment Moved, $\gamma = 1.3$




7.3 Inventory Imbalance Stability

Up to this point, the stability of the inventory imbalances generated by the provisioning policy has been ignored. No consideration has been given to the overall trends displayed by the inventory imbalances over time and across replications.

An investigation into inventory imbalances by station over the 14 days indicates that for the RSSR policy, $\lambda \geq 0.2$ the inventory imbalance variance is increasing. For $\lambda = 0.1$ the inventory imbalance variance is non-increasing. Therefore the next step is to determine an optimal stable value of λ between 0.1 and 0.2.

Figures 44 and 45 are box plots of inventory imbalances for the 100 replications, by day, for two of the 12 stations not using the DH policy. From previous analysis in section 7.2, we know that the magnitude of the inventory imbalances is satisfactory. However, the issue of long-term sustainability has not been addressed. One way to determine if the system is operating at steady state after the 14 days of simulated operations is to examine the variances of inventory imbalances across all 14 days for each station. An indicator of stability would be a non-increasing variance. From Figures 44 and 45, it can be seen that the variance is constant if we discount the first few days²⁵. Referring to the days after day 4, we see that the variance does change but it isn't strictly increasing. Rather, it seems to be bounded. Finally, referring to the mean inventory imbalance, we see that it is oscillating around zero. The conclusion is that when $\lambda = 0.1$, the RSSR policy generates stable long-term inventory imbalance variances.

Figure 44 – Station Stability, Large Station, $\lambda = 0.1$



²⁵ The first few days will have a lower variance simply because they are closer to the starting initial condition of 0 inventory imbalance. This issue will be addressed in greater detail in section 7.4.

Figure 45 – Station Stability, Medium Station, $\lambda = 0.1$




Figure 46 – Station Stability, Large Station, $\lambda = 0.2$




Figure 47 – Station Stability, Medium Station, $\lambda = 0.2$




Figure 46 and 47 are box plots of inventory imbalances for the 100 replications, by day when $\lambda = 0.2$, for the same two stations in Figure 44 and 45. The notable difference between $\lambda = 0.2$ and $\lambda = 0.1$, is that the variances appear to be increasing for $\lambda = 0.2$. Although the mean still oscillates around 0, the increasing variance indicates that the simulation with $\lambda = 0.2$ has not yet reached steady state. The importance of this finding is that the previous analysis for $\lambda = 0.2$ does not consider the possibility that the inventory imbalances could increase in the days following the end of the simulation. Therefore the conclusions drawn about the RSSR policy performance when $\lambda = 0.2$ is based upon data that is not representative of the actual performance of the RSSR policy. In order to fully investigate this, the simulation period would have to be extended until the results for $\lambda = 0.2$ stabilized. This possibility is left for future investigation

7.4 Optimizing λ

From section 7.2 we know that the RSSR policy has an acceptable level of inventory imbalance for all stations. We also know from the previous section that the RSSR policy is stable for $\lambda = 0.1$ and unstable for $\lambda = 0.2$. Finally, we know that the value captured by the RSSR policy increases exponentially with increasing λ .

Up to this point the RSSR policy has been evaluated for values of λ ranging from -1.1 to 1.1 in 0.1 increments. The goal of this section is to narrow the analysis to the range of $0.1 \leq \lambda \leq 0.2$, in 0.05 increments, in order to find the optimal value of λ . This is defined as the largest value of λ for which a stable solution exists. From the following analysis it is determined that this value of λ is 0.125 .

Figures 48 to 53 present box plots of inventory imbalance for various values of λ . Some key observations from the graphs are:

1. The patterns in the graphs are apparent across different values of λ
2. Variances increase with increasing λ
3. For a single value of λ , there appears to be a weekly pattern for both the large and medium station
4. For the large station, the inventory imbalances are stable for all values of λ
5. For the medium station, the inventory imbalances are stable for $\lambda = 0.125$, questionably stable for $\lambda = 0.15$ and not stable for $\lambda = 0.175$

Figure 48 – Station Stability, Large Station, $\lambda = 0.125$




Figure 49 – Station Stability, Medium Station, $\lambda = 0.125$




Figure 50 – Station Stability, Large Station, $\lambda = 0.15$




Figure 51 – Station Stability, Medium Station, $\lambda = 0.15$




Figure 52– Station Stability, Large Station, $\lambda=0.175$





Figure 53– Station Stability, Medium Station, $\lambda = 0.175$



There are patterns that are observed across the different values of λ . Referring to the large station box plots (Figures 48, 50 and 52), we can see that day one is negative, day two is lower than day one, day three is higher and positive and so on. This observation is expected as the flight schedule does not change, and the 100 different random number streams (one for each replication) are the same across each value of λ .

Variances for a particular day for both stations do increase with increasing λ . Again this is expected, as a higher value of lambda would result in a lower number of excess equipment movements. This in turn would increase the possible inventory imbalances, which would directly increase the variance.


For a single value of λ , for either station, there are weekly patterns in the inventory imbalances. Referring to Figure 49, we see that day one is high, days two through five are lower and about the same, day six is high, and day seven is lower than days 2 through five. Continuing on we see that day eight is high, days 9-12 are lower and about the same, day 13 is high and day 14 is lower than days 9-12. Again this is expected as the flight schedule from the Meal Pages database is based on a seven-day pattern. Although there are other patterns and random events, the seven-day pattern usually dominates.

Referring to Figures 48, 50 and 52 we can see that inventory variances are not increasing, neglecting the increase for the first few days. This is evident in Figures 48 and 50, and is a little questionable for Figure 52. It should be noted that to increase confidence in the conclusion that variances are non-increasing, the simulation schedule should be increased from 14 days to a minimum of 28 days.

Referring to Figures 49, 51 and 53 we can see that the inventory variances are not increasing for $\lambda = 0.125$. For $\lambda = 0.15$ it is inconclusive, and for $\lambda = 0.175$ the variance appears to be increasing. Again to increase confidence in these conclusions, the simulation schedule should be increased.

As a result of the above analysis, it is possible to select either $\lambda = 0.15$ or $\lambda = 0.125$. However, in order to be conservative, the choice is $\lambda = 0.125$. The RSSR policy with $\lambda = 0.125$ will capture 61.6% of the available cost reduction. This result must be modified to take into account the required warm up period.

Figure 54 – RSSR, % Available Cost Reduction Captured, $\gamma = 1.3$, II



7.5 The effect of S^L , $\gamma = 1$

The inventory imbalance limit, S^L , has been set to $1.3V_i$ for all previous analyses. Recall from section 6, that the safety stock resulting from this choice is $0.2*V_i$. In the event that this amount is deemed too low, the results for the RSSR policy with $S^L = 1.0V_i$ are presented below.

Changing the value of S^L will only affect the threshold that determines whether the inventory imbalances are acceptable for a given policy. Figures 55 and 56 present the inventory imbalances for the RSSR policy for values of λ between 0.1 and 0.2. As can be seen the decrease in γ has resulted in several values of λ which demonstrate unacceptable inventory imbalances. However, for $\lambda=0.125$, there is only one station which has an inventory imbalance greater than S^L , and this only occurs twice in 100 replications of 14 days. If the users of this policy still felt that this was unacceptable, the offending station could be removed from the SR policy, and placed on the DH policy or a small amount of equipment could be added to its safety stock. The case where $\gamma = 1.5$ is uninteresting as it will have a better inventory imbalance performance than when $\gamma=1.3$; and $\gamma=1.3$ already has acceptable inventory imbalances for all stations at $\lambda=0.125$.

Figure 55 – RSSR, Minimum Inventory Imbalance, % of S^L , $\gamma = 1.0$





Figure 56 – RSSR, # of Times Inventory Imbalance Decreases Below $-S^L$, $\gamma = 1.0$



7.6 Selecting a Warm up Period

The next step in quantifying the RSSR policy is to determine its performance relative to the DH policy. In order to do this the RSSR policy results must not include the simulated warm-up period, because this will generate results that are not entirely characteristic of the RSSR policy. The policy will not operate in an environment where every two weeks the starting inventory imbalance is zero.


The results for the simulation model presented thus far have been inclusive of the warm-up period. Up to this point, the inclusion of the warm-up period did not significantly affect results. Specifically, the only concern has been the performance of the policies relative to one another and the maximum/minimum inventory imbalance. The performance of policies relative to one another is not influenced significantly by the warm-up period – all policies would receive about the same benefit / detriment. The maximum / minimum inventory imbalance would always occur outside of the warm-up period. At this point we are interested in determining the appropriate reduction in excess equipment movements. As such it is important to ensure only steady state results, those excluding the warm-up period, are analyzed.

The purpose of the warm-up period is to remove results from the simulation run that are not representative of the normal expected results. In this case, we are interested in removing the period where inventory imbalances are abnormally low; i.e. the period where the inventory imbalances are highly correlated to the initial condition. There are several approaches to determining this period. Three possible approaches are listed below:

1. Arbitrarily truncate the first seven days
2. Inspect the box plots of inventory imbalance, and select only those points that occur after the average variance in inventory imbalance has been reached
3. Inspect the excess equipment movements, and select only those points that occur after the moving average is lower than the total average

The first option is simple, but would need to be supported with additional evidence. The second option seems reasonable, and given the data analysis from section 7.3, this would occur sometime around day 4. The third option is also reasonable. However when the actual data is examined, there is significant volatility in the excess equipment movements by station. Encouragingly, the system total excess equipment is relatively constant. Referring to Figure 57, the warm-up period would extend to day four. As can be seen, the warm-up period will have a minimal impact on average excess equipment movements.

Figure 57 – RSSR, System Total Excess Equipment Movements, $\lambda=0.125$, $\gamma=1.3$



As a result of the excess equipment analysis and the inventory imbalance analysis coming to the same conclusion, the warm-up period will be set to 4 days. It should be noted that Figure 57 suggests that the warm-up period selection will not affect the results significantly. Still, in all remaining simulation analysis, the first four days of simulation output are removed.

7.7 Final Simulation Models

The final simulation models are based upon the following parameters:

$$\lambda = 0.125$$

$$\gamma = 1.3$$

Warm up period = 4 days

Replications 100

The policies are evaluated based upon the number of times that inventory level decreases past S^L , reduction in excess equipment movements and the total possible excess equipment movements.

7.7.1 RSSR Policy

The results of this policy indicate that there is never an inventory imbalance below S^L . The number of excess equipment movements is reduced by approximately 3200. The number of excess equipment movements for the simulation deadhead policy SDH policy is approximately 5300 per day.

7.7.2 3SSR Policy

The 3SSR policy is an application of the SSR policy to the three largest stations only. Under this policy, the three largest stations will load equipment according to a reference inventory level, and the other 28 stations follow the deadhead policy. The results of this policy indicate that there is never an inventory imbalance below S^L . In addition, the number of excess equipment movements is reduced by approximately 1000 per day. The number of excess equipment movements for the SDH policy is approximately 5300 per day.

7.8 Cost Savings Analysis

The main benefit of implementing an alternative provisioning policy is a potential reduction in costs. There are four main areas where cost reductions could arise:

- *reduced inventory levels*: safety stock levels could decrease as a result of lower equipment use
- *reduced shrinkage*: each time a piece of equipment is used, there is a chance that it may be broken; if we reduce usage, we should reduce breakage
- *reduced fuel consumption*: the amount of fuel consumed is proportional to the total weight of the aircraft; reduction in excess equipment reduces total weight of aircraft
- *reduced handling costs*: all equipment is cleaned and assembled for boarding, and is cleaned again at the arrival station, whether it is used or not

It is conceivable that safety stock levels could decrease. The safety stock levels are defined as a multiple of daily consumption. Therefore, if daily consumption decreases, the resulting safety stock required could also decrease. This is supported by the results of the simulation analysis, which indicate not one of the 31 stations experienced inventory deviations below S^L . However, when the actual costs associated with inventory are tallied up, we see that the total system wide inventory costs for the equipment under analysis is less than \$500,000. Recognizing that we would only be able to reduce a fraction of that, the one-time savings become something around \$50,000. Clearly, this is not something worthy of further investigation.

CAI records historical breakage statistics for many of its items in inventory. The breakage is generally attributed to normal wear and tear, and to this end the breakage percentage is combined with the cost per item to generate a breakage cost per use. For the items in this investigation, the breakage cost per use is approximately \$0.09. Multiplying this amount by the reduction in excess equipment movements calculated in the previous section we see that the RSSR policy saves \$288 / day, (100K annually) and the 3SSR policy saves \$90/day (33K annually).

The fuel burn on aircraft is proportional to the take off weight of the aircraft. If the amount of excess equipment loaded onto an aircraft is reduced, then the weight of the aircraft is reduced. A rough estimate of potential cost savings yielded an optimistic estimate of \$5000 annually. This amount is insignificant and does not warrant further investigation.

Equipment is cleaned and assembled every time it is used. The exact cost of this is determined by the caterer contract, but is generally in the range of \$0.25-\$0.75. This corresponds to a system wide savings of \$800 to \$2400 / day (290K-875K annually) under the RSSR policy, and a savings of \$250 to \$750 / day (90K-275K annually) under the 3SSR policy. This potential cost savings combined with the breakage costs is significant and is worthy of future investigation.

8.0 Conclusion

This thesis presents a network inventory control policy for an airline, the single reference policy. This policy was developed using Markov Decision Process models in conjunction with a Simulation model.

The three MDP models provided an objective basis for determining an alternative equipment provisioning policy and quantifying the performance of the policy. The one-node MDP model demonstrated that the deadhead policy was a useful local optimization policy. The two-node MDP model demonstrated that the single reference policy could outperform the Deadhead policy, however the cost structure of the system was important. The three-node MDP model demonstrated that the single reference policy's performance would improve as the number of nodes increased. The three node single reference policy, reduced costs related to the movement of excess equipment by 56% on average.

The Simulation model was used to test the single reference policy using more detailed operational data. Three variants of the single reference policy were developed and tested: the universal simulation single reference policy, the reduced simulation single reference policy and the three-city single reference policy. The results of the universal simulation single reference (USSR) indicated that reduction in excess equipment movements were exponentially related to the reference level chosen in the USSR policy. The second was that the USSR policy generated unacceptably high inventory imbalances for many of the "small" stations. The reduced simulation single reference policy (RSSR) indicated that the inventory imbalance levels were acceptable, and stable, for certain values of the reference level. The RSSR policy reduced the excess equipment movements by 60%²⁶. The final model was the 3SSR model. This model reduced the excess equipment movements by 19%.

The results of the simulation were used in a cost savings analysis that demonstrated the annualized potential cost savings for the RSSR model were a minimum of \$390,000 per year, and for the 3SSR model were a minimum of \$120,000 per year. The operational impact of implementing either the RSSR or the 3SSR equipment provisioning policy is minimal. The only additional data requirement is a daily estimation of inventory levels.

In order to complete this thesis, we required operational data. A result of this project was the development and implementation of two databases at Canadian Airlines International, the Meal Pages Database and the Catering Commissary Database. Both databases have replaced the previous processes for developing a meal provisioning schedule and for maintaining the aircraft equipment configurations / loading procedures. These by themselves, represent a significant contribution to operations.

²⁶ for $\lambda=0.125$

9.0 Future Recommendations

The results of this investigation are based on several models, each with its own assumptions. In order to increase the confidence in the model results, the following steps should be taken:

1. Incorporate a longer flight schedule for simulation
2. Include more passenger demand distributions
3. Investigate airline-caterer relationship

The simulation model as formulated relies on 2 representative weeks of flight data. This corresponds to almost 3000 flights. However, with the continued implementation of the Meal Pages database, additional flight information will be available. This would permit a broader sample period.

The passenger demand data available is very extensive, and would permit further investigation. Additional areas of analysis include increasing the number of distributions used to accommodate more regions, developing passenger distributions for each season, and developing passenger distributions for each direction.

Finally, the cost savings analysis of the simulation results is based upon limited information about the nature of the Caterer-Airline relationship. The cost reductions cited are based upon previous caterer service pricing. However, significant variation occurs from caterer to caterer and from station to station. This would suggest that the prices as quoted do not accurately reflect the true cost of the service. Therefore to improve the accuracy of the estimated cost savings, further investigation into the airline-caterer relationship is required.

Glossary

Aircraft –	an air vehicle
Aircraft type –	the model of aircraft, (B737, B747 etc.)
Aircraft route –	the listing of flight routes, in order, that a particular aircraft will follow
City pair –	two of the cities from the network of cities
DH –	deadhead policy, always move all available excess equipment
DV _i –	daily Volume for station I, equal to the average daily total of seats flown out of station I
Flight –	an aircraft transporting passengers between two cities
Flight leg –	one origin and one destination city from a flight route
Flight route –	the origin and destination cities of a flight, a flight route can have more than two cities, for example, consider a flight from Vancouver to London that has a stopover in Calgary
Inventory allotment –	the working equipment assigned to each station
Inventory imbalance –	the difference between the current inventory level and the inventory allotment
S ^L –	inventory imbalance limit, for MDP analysis if S ^L is exceeded and automatic rebalance action is initiated, for simulation S ^L is used to determine acceptable models
SR –	single reference policy, if inventory imbalance is greater than X ^H , move all available excess equipment, if less than X ^H move no excess equipment
SSR –	simulation single reference policy, all simulated stations follow the SSR policy
Station –	a station is a city that flights fly to and from
RSSR –	reduced simulation single reference policy, 19 simulated stations follow the DH policy, 12 simulated stations follow the SSR policy
3SSR –	3 city simulation single reference policy, only the three largest stations follow the SSR policy, all other stations follow the DH policy
X ^H –	reference inventory imbalance level, for use with SR policy and derivative policies

Bibliography

- [1] D. Atkins and P. Iyogun 1987 "A Lower Bound on a Class of Coordinated Inventory / Production Problems", *Operations Research Letters* **6**, 63-67
- [2] J. Mitchell 1987 "98%-Effective Lot-Sizing for One-Warehouse Multi-Retailer Inventory Systems with Backlogging", *Operations Research* **35**, 399-404
- [3] M. Fu 1992 "Sample Path Derivatives for (s,S) Inventory Systems", *Operations Research*, **42**, 2, 351-364
- [4] F. Chen 1995 "Stationary Policies in Multiechelon Inventory Systems with Deterministic Demand and Backlogging", *Operations Research*, **46**, 3, S26-S34
- [5] R. Anupindi and S. Tayur 1997 "Managing Stochastic Multiproduct Systems: Model, Measures and Analysis", *Operations Research*, **46**, 3, S98-S111
- [6] Goto, Jason 1999 "A Markov Decision Process Model for Airline Meal Provisioning", Graduate Thesis, UBC Management Science
- [7] A. Law and D. Kelton "Simulation Modelling and Analysis 2nd Edition" 1991 McGraw Hill Inc.
- [8] R. Sargent 1992 "Validation and Verification of Simulation Models", Proceedings of the 1992 Winter Simulation Conference, pp 37-47
- [9] T. Naylor, J. Finger 1967 "Verification of Computer Simulation Models", *Management Science*, **14**, 2, B92-B101
- [10] L. Schruben 1980 "Establishing the Credibility of Simulations", *Simulation*, March 1980, 101-105
- [11] G. Kleindorfer, L. O'Neil, R. Ganesan 1998 "Validation in Simulation: Various Positions in the Philosophy of Science", **44**, 8, 1087-1099

Appendix A: Solving the 3 Node Model

The final execution time required to solve the 3N problem was approximately 38hrs. The total computation time required to solve all the parameterized versions of the models (9) was 336 hours. In order to solve the 3N problem in 38hrs, several concessions had to be made. These concessions are summarized below:

1. Sum of inventory = 0
2. No record of action space for each state
3. Storing expected values for each combination of inventory levels
4. Smaller passenger demand distribution

Without these changes to the model, the computation time would have been excessively long. The unreduced model has approximately 8.1×10^9 individual states for each epoch. Given 11 epochs, that total number of states for the model is 8.9×10^{10} . There are 15 relevant statistics for each state, 14 integer data type and 1 float data type. This yielded a required storage space of approximately 2.9×10^{12} bytes. In addition, for each state, the algorithm would have to solve a number of iterations equal to the size of the action space. Depending on the size of the demand space the number of iterations is in the range of 1 to 117,649. Table 7 details the reduction in storage space required for each successive model reduction. Table 8 details the number of iterations required for each successive model reduction. Assuming that the model with no reduction could have been solved, and that the space requirements would not have restricted computation speed, it would have required 1.9×10^6 years or execution time to solve all 9 parameterized versions of the model. The following sections describe each model reduction.

Table A1 – 3 Node Storage Space

Model Reduction	Inventory Configs	Demand Configs	# States / Epoch	Total # States	Bytes / State	Total # Bytes
No Model Reduction	68921	117649	8.11E+09	8.92E+10	32	2.85E+12
Sum of Inventory = 0	1261	117649	1.48E+08	1.63E+09	32	5.22E+10
		117649	1.48E+08	1.63E+09		
Expected values	1261	N/A	1261	1.39E+04	8	1.11E+05
Passenger demand	1261	N/A	1261	1.39E+04	8	1.11E+05

Table A2 – 3 Node Number of Iterations

Model Reduction	Total # States	Max Iterations / State	Total Iterations
No Model Reduction	8.92E+10	1.18E+05	1.15E+17
Sum of Inventory = 0	1.63E+09	1.18E+05	2.11E+15
No action space	1.63E+09	1.18E+05	
Expected values	1.39E+04	1.18E+05	1.80E+10
Passenger demand	1.39E+04	1.56E+04	2.38E+09

Sum of Inventory Levels = 0

Every state is defined by three inventory levels and six passenger demands. Given that the starting inventory level for all stations is zero, it follows that the sum of inventory levels at all other times must also be zero. Therefore, although there are 68,921 ($41 \times 41 \times 41$ ²⁷) different combinations of inventory levels, there are only 1261 combinations of inventory levels that add up to zero. It is important to note that this model reduction does not alter the model solved; rather it capitalizes on the fact that the model is a closed system.

No Record of Action Space for each state

In the previous two models, for every state (inventory levels at all stations, and all passenger demands to and from each station) the actions selected by the model were stored for later analysis. Removing the actions from the state reduced the storage space required for each state by 12bytes (6X2). Unfortunately this reduction removed the possibility of examining the decision regions as was previously done in the 1 and 2 node models.

Storing expected values for each configuration of inventory levels

The 3N models, like the 1N and 2N models require the demand information for two things. The first is to calculate an immediate cost, the second is to determine the expected value of all future costs given a current state. Fortunately, the stochastic nature of the problem only occurs from epoch to epoch. The behavior of the model in the current epoch given the current state is deterministic. This fact removes the necessity of storing demand information for all states in epochs later than the epoch currently being examined. Therefore, all states in a given epoch that share a common set of inventory imbalances can be equivalently represented as the set of inventory imbalances and the expected cost associated with it. Given that the choice has been made not to store the actions for every state, there is no purpose for storing the demand configurations. As such, storing only expected values of cost for a set of inventory imbalances does not limit further limit the analysis.

Smaller passenger demand distribution

The first three model reductions solved the space requirement problem presented by the unreduced model. Unfortunately, there were still a daunting number of iterations to be performed. The only method for reducing the number of iterations was to reduce the number of states. This implied either reducing the number of epochs, reducing the allowable inventory imbalances, or reducing the passenger demand space. Reducing the number of epochs would only have a linear effect on the number of states, and therefore was not considered a viable option. Reducing the allowable inventory levels would enforce even smaller bounds on the auto rebalance limit. Given that the rebalance limit should in some way be proportional to the amount of equipment that could possibly enter or exit a station in any given epoch, the progression to a 3-node model from a 2-node model should actually see the rebalance limit increase. The only option remaining was to reduce the size of the passenger demand space. Although this was not a desirable option, it did have one positive- a small reduction in demand space would have a very large impact on the total number of states. The final choice was to reduce the demand space from a discrete triangular distribution (4,7,10) with 7 points to a discrete triangular distribution (5,7,9) with 5 points. This decision will improve the results of the alternative provisioning policy, and the optimal policy, because it reduces the number of situations where balancing inventory requires moving 6 units of excess equipment.

²⁷ Remember that the auto rebalance action will not permit inventory levels greater than 20 or less than -20. This gives 41 possible values for inventory at each station.